Improving Efficiency of Commercial Vehicle Operations Using Real-Time Information: Potential Uses and Assignment Strategies

AMELIA C. REGAN, HANI S. MAHMASSANI, AND PATRICK JAILLET

Advances in communication, automatic vehicle location, and geographic information system technologies have made available several types of real-time information with benefits for commercial vehicle operations. Continuous updates on vehicle locations and demands create considerable potential for developing automated, real-time dispatching systems. The potential benefits of a diversion strategy in response to real-time information are explored under idealized conditions, and the technologies that are available for use in commercial vehicle operations and selected results derived from simulation are described. The results illustrate potential savings from simple diversion strategies under real-time information and highlight the need for methodological development to support improved truckload carrier operations decisions.

Telecommunications and information technologies provide unprecedented opportunities for using real-time information to enhance the productivity, performance, and energy efficiency of the commercial transportation sector. Achieving the benefits of real-time information requires development of fleet operating strategies, including vehicle assignment and dispatching rule with increased flexibility, along with suitable decision support methodologies. There appears to be virtually no methodology in the literature intended specifically for truckload or other surface carrier operations under the kind of real-time information possible with emerging technologies. The lack of methodological development applies to both the analysis of carrier operations to evaluate the effectiveness of real-time information and to actual tools that could be used by carriers to take advantage of such information. The area of vehicle routing and scheduling, including dynamic vehicle allocation and load assignment models, has evolved rapidly in the past few years, both in terms of underlying mathematical basis and actual commercial software tools (*1–3*). Although these approaches may well be adaptable to operations under real-time information availability, they are currently unable to take full advantage of such information because their underlying formulations do not recognize possible decisions that are meaningful only under real-time information. One such decision is the possibility to divert in response to customer demands.

After briefly describing some of the technologies available, this work identifies and explores potential uses of real-time information for the efficient management of truckload carrier operations. In particular an en route "diversion" strategy in response to unfolding customer demands is proposed and analyzed. A simulation model to explore the profitability of such diversion strategies under various operational conditions and demand arrival patterns is described and conditions under which such strategies might be profitable are derived. Findings suggest that meaningful potential exists for improving truckload carrier operations. These findings and related operational issues are discussed.

In truckload operations, carriers typically know only a portion of the loads that must be moved before the beginning of the day. Typically 60 percent of a given day's loads may be accepted on the same day that they are moved (*1*). The assignment of an available driver to a load therefore takes place almost in real time or at least shortly after the request is received. In addition, the load acceptance decision made by a carrier must be executed in real time and may have a significant impact on the carrier's ability to accept other loads later in the day or in the days that follow. This research explores ways to make "good" assignment decisions, and ultimately load acceptance decisions, that lead to overall cost-effective operations but rely on local (current) rather than on long-term or forecasted information. Although various forecasting methods may be used to estimate future demands, this information is not reliable in practice because of the large number of possible origin and destination combinations and the inherent randomness of the process (*1*).

INTRODUCTION TO TECHNOLOGIES

Automatic vehicle location (AVL) systems are finding increasing application in a variety of contexts, including truckload and less than truckload trucking companies; local delivery and courier services; fire, rescue, and police departments; utility companies; security companies; public transportation companies; high-value and hazardous materials shippers; and taxi and limousine services. Not all applications require the same degree of accuracy. The dispatcher in a long-distance trucking company will most likely derive the same benefit from knowing the locations of the company vehicles to within 50, 1000, or even 10 000 m, whereas a police dispatcher may need to determine, with certainty, on which streets tracked vehicles are located.

Although global positioning satellite (GPS) technology is often perceived as the leading AVL technology, the companies with a major market share in the long-haul trucking AVL market do not employ GPS technology in their standard products. One of these systems uses a group of nationwide specialized mobile radio towers with optional Loran-C location tracking, and the other uses a

A. C. Regan and H. S. Mahmassani, Department of Civil Engineering, University of Texas, Austin, Tex. 78712. P. Jaillet, Department of Management Science and Information Systems, University of Texas, Austin, Tex. 78712.

network of two geosynchronous satellites to perform tracking and communication, with Loran-C an optional addition. The 500- to 1000-m accuracy these systems provide is adequate for dispatchers to estimate which highways the trucks are on. Commercial applications that require street-level location information could benefit from increased accuracy. Applications that include navigation, either on board or at a central location, require more accurate position information than 500- to 1000-m estimates, as these do not ensure street-level accuracy.

In most AVL applications, the position location obtained must be transmitted to the dispatch center over an available communication link. Although position estimates and even point-to-point routing could, with an appropriate microcomputer, be calculated on board the vehicle, the vehicle's position must be transmitted back to the dispatch center for display. Communication links available for this purpose differ in cost and sophistication. VHF, cellular, or subtitle link may all be used, with digital cellular becoming more and more widely available. The link used typically depends on the frequency of communication and the distance between the dispatch center and the vehicles. The communication cost of such a system may be high, with messages costing as much as ten cents per brief packet for satellite communication systems, about five cents per packet for transmission over 800 or 900 MHz trunked radio lines, and more for standard cellular in which rates are determined by the minute rather than the data packet. If the vehicle locations are 'polled' often by the central dispatcher, these costs add up quickly (*4*).

This study is most interested in irregular route common carrier operations. Discussions with operators of trucking companies of various sizes have made clear that whatever the particular technologies chosen, AVL and two-way communication systems will be necessary for many trucking companies to compete in a market where the location and magnitude of demands for service are highly dynamic (*5*). A 1992 survey performed at the University of Texas of just under 300 carrier companies pointed out the fact that carriers agree that AVL and two-way communications technologies will lead to improvements in many aspects of their operations. Figures 1 and 2 share some of their responses about what they perceive as the potential benefits of these technologies. Although it is clear that these technologies are beginning to see widespread use, it is equally clear that the full potential of these technologies will not be realized until responsive real-time dispatching tools become available.

REAL-TIME ASSIGNMENT STRATEGY: DIVERSION

Because of the length of some empty moves made to pick up loads, it is possible that new information on demands to be serviced may arrive while a driver is en route to a pickup. Assuming that time windows for movements are flexible, this new demand information may be used to order demands in such a way as to reduce empty distances driven. Quasi-continuous dispatcher-to-driver communication makes it possible to divert a driver en route to a pickup location to an alternative load, thereby inducing a resequencing or reassignment of the original load. Such diversion strategies are not generally feasible under current operations because dispatcher-driver communication takes place at discrete instances only, typically at a load pickup or delivery point (*5*).

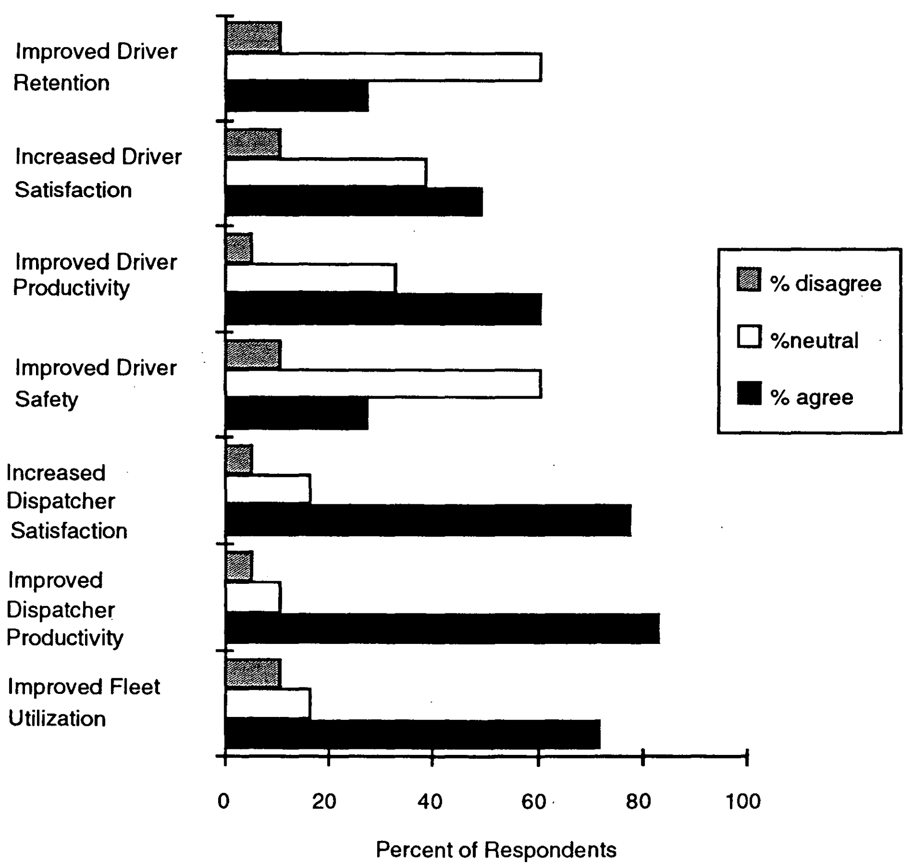

FIGURE 1 User assessment of two-way communication and AVL system benefits: Part 1.

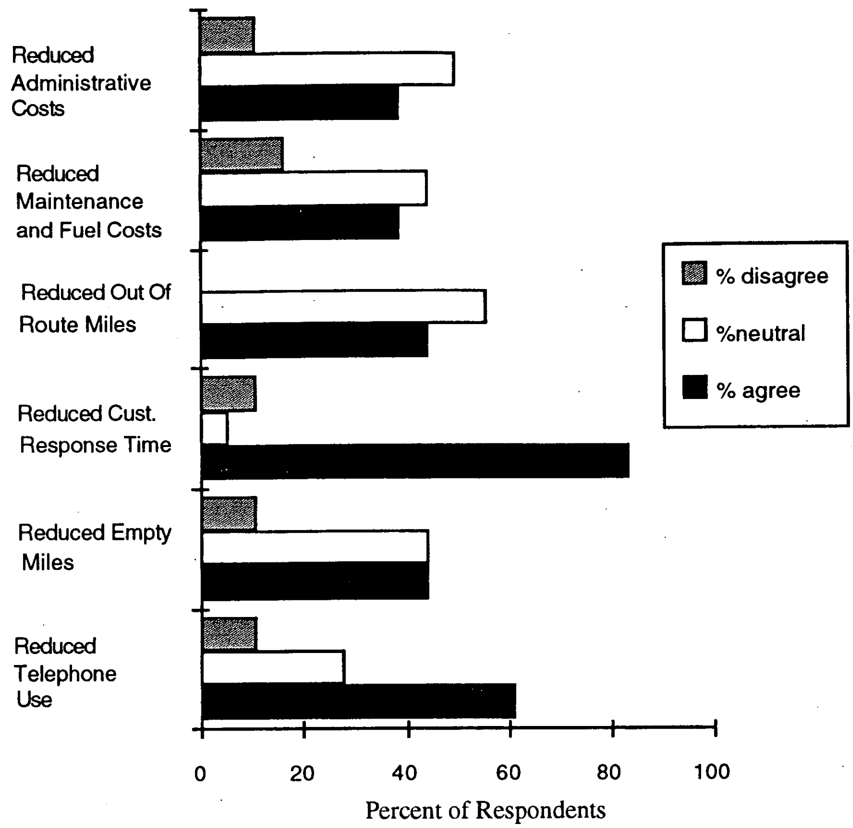

FIGURE 2 User assessment of two-way communication and AVL system benefits: Part 2.

The relative improvement possible under this strategy depends on the relative locations of the alternative pickup and delivery points. Under some distributional assumptions about the locations of these points, the interest is in the probability that diverting the driver to a new demand while en route to a previously assigned pickup will be beneficial. In even the simplest case, it is difficult to derive this probability analytically because the various cost components are not independent. For this reason, these probabilities and various other performance measures are evaluated through simulation of such diversion strategies over service horizons of varying lengths, under different arrival stream distributions, and under load acceptance rules that either require all loads to be serviced or allow less profitable loads to be rejected. The scenarios examined up to this point are not intended to exactly replicate actual operating conditions, but to provide a simplified representation that allows derivation of basic insights into the potential benefits of real-time information and the factors that affect these benefits, as well as the identification and design of strategies that merit examination under more realistic operating conditions.

Diversion Probabilities Under Simple Assumptions

To begin with the most basic case, while a driver is en route to a load origin, information about another load (and in this initial case, only one other load) to be moved becomes available. Answers to the following questions are desired: What is the probability, given various diversion decision rules, that the driver will be diverted to serve the new load first? What is the probability that following such diversion decision rules will result in a reduction of overall distance traveled? And, what is the associated expected reduction in travel?

To clarify, consider in Figure 3, a vehicle that begins at the center, *c*, of a circle and moves toward the origin of a loaded movement between Points X_1 and X_2, where these points are uniformly and randomly generated over the area of the circle. Given a diversion point (the point at which another load to be moved becomes available) some fraction of the distance from the center of the circle and origin X_1, the probability is derived that the distance between the diversion point to a new origin X_3 will be less than the distance from the diversion point to origin X_1. Let α, $0 \leq \alpha \leq 1$ denote the fraction of the distance from the center to X_1 traveled to reach the diversion point. The probability that the distance from the

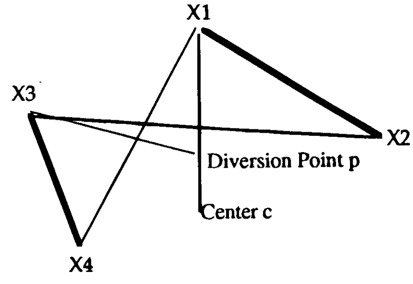

FIGURE 3 Diversion example.

diversion point to the new origin is less than that to the old origin is given by $(1 - \alpha)^2/2)$, as shown hereafter.

Let $B(c, r)$ denote the circle of center c and radius r, and $d(x, y)$ the Euclidean distance between points x and y. Consider two random points in $B(c, r)$, say, X_1 and X_3. For $0 \leq \alpha \leq 1$, let $W_1(\alpha)$ be the point on the segment (c, X_1) such that $d[c, W_1(\alpha)] = \alpha[d(c, X_1)]$. Define the following two random variables $Y_1 = d[W_1(\alpha), X_1]$ and $Y_2 = d[W_1(\alpha), X_3]$, where Y_1 and Y_2 represent the distances from the potential diversion point to the current and potential load origins.

Let Z be the radial distance of $W_1(\alpha)$ so $Z = d[c, W_1(\alpha)]$, and $f_z(\cdot)$ be its probability density function:

$$P(Y_2 > Y_1) = \int^{P} (Y_2 < Y_1 \mid Z = z) f_z(z)\, dz$$
$$= \int^{P} \{X_3 \in B[W_1(\alpha), z/\alpha - z]\} f_z(z)\, dz \qquad (1)$$

Because $W_1(\alpha)$ is a random point in $B(c, \alpha)$,

$$P(Y_2 < Y_1) = \int_0^{\alpha} [(1 - \alpha) z/\alpha]^2 (2z/\alpha^2)\, dz = (1 - \alpha)^2/2 \qquad (2)$$

If a myopic strategy of diverting to the new demand origin, X_3, is followed if it is closer to the diversion point than origin X_1, then $(1 - \alpha)^2/2$ represents the fraction of loads for which one actually diverts. This probability $[P(Y_2 < Y_1)]$ is shown graphically as a function of the diversion point location parameter, α, in Figure 4. However, under this strategy, even if the diversion decision at point $\alpha = 0$ is evaluated, the resulting average savings in terms of reduced distance traveled while serving the two loads is less than 1 percent, and diverting at points further downstream actually results in a slight increase in traveled distance, on average.

A more plausible diversion strategy would also consider the relative distances between the destination point of the first movement and the origin point of the next load. In Figure 3 these are given by $d(X_2, X_3)$ and $d(X_4, X_1)$. In this case diversion is chosen if

$$d(p, X_3 + d(X_4, X_1) < d(p, X_1) + d(X_2, X_3) \qquad (3)$$

Analytic derivation of the corresponding diversion probability under this strategy is no longer straightforward because the respective distances are not independent. The diversion likelihood and associated expected benefit are evaluated using a simulation program under the following underlying assumptions:

- A circular work area with a radius of 1 unit of travel as in Figure 3,
- Uniformly and independently generated demand locations,
- Euclidean (straight-line) travel distances, and,
- Diversion results compared with a 'base' case where demands are serviced in order of their arrival.

Simulation was used to evaluate the case shown in Figure 3, in which the vehicle begins at the center of the circle, c, and only two demands are served. A total of 75,000 independent trials were executed for each of 10 values of α, the diversion point fraction, that varied between 0 and 1. When $\alpha = 0$, that is, resequencing occurs before departing for the first demand, as should be expected, resequencing occurs in half the cases. The fact that resequencing occurs in more than 10 percent of the cases when the diversion decision is evaluated at the origin point of the first load, that is, $\alpha = 1$, is somewhat counterintuitive and results from the cases in which the loaded movement of the candidate load takes the vehicle close to the origin of the original load. Average savings resulting from such a diversion strategy, in which the demand horizon (the number of demands served in a single simulation instance) is only two loads, vary from 7 percent of the total distance traveled if demands are taken in order with $\alpha = 0$, down to 1 percent with $\alpha = 1$. These results are shown in Figure 4.

This analysis is extended beyond the first diversion decision. After serving the load selected, the vehicle begins to move toward the unsatisfied demand. Again, a new demand arises along the way, creating a new diversion opportunity. With a demand horizon of 100 loaded movements, evaluated sequentially on a pairwise basis, simulation results indicate overall benefits (of the diversion strategy) in the range of 1.5 to 12.5 percent of overall distance traveled, depending on the diversion point fraction, relative to the base case of servicing demands in the order in which they arrive. Figure 5

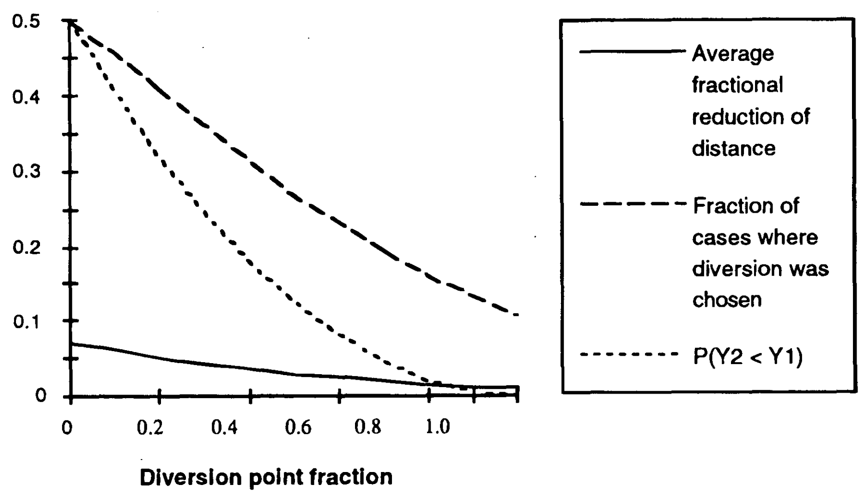

FIGURE 4 Probability of diversion when distances to load origins and between two demand points are considered.

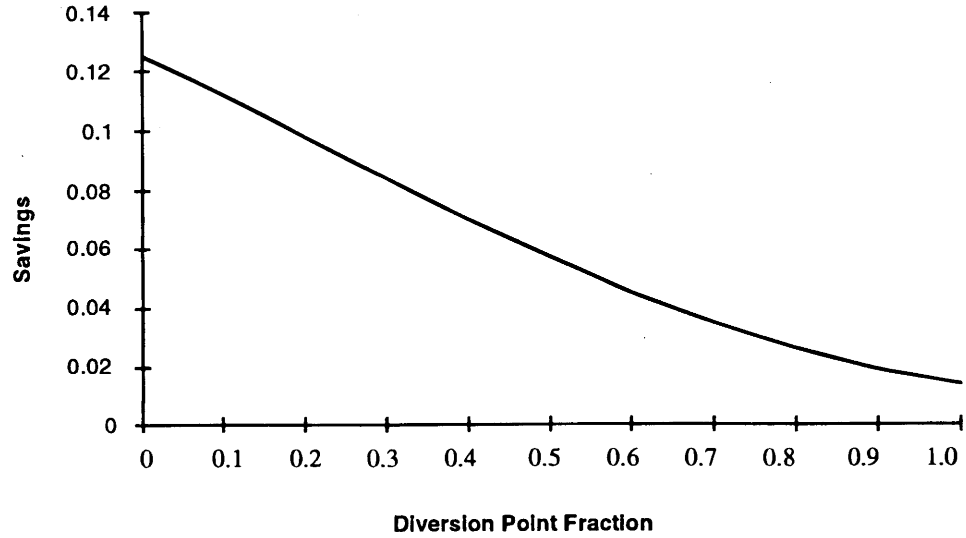

FIGURE 5 Reduction of travel with diversion strategy as fraction of overall travel with 100 demands served.

shows the overall reduction in travel cost as the diversion point fraction is varied from 0 to 1. A regression model that assumes an exponential functional form approximates this curve as

$$\text{Reduction} = (0.151)0.111^{\alpha} \quad (4)$$

Note that in the simulations the new load is assumed to be known by the time the vehicle reaches the diversion decision point. If α, the diversion point fraction, is a uniform random variable taken between 0 and 1, which would correspond to a scenario in which the new load may become known at any point along the route between the vehicle's last load destination and the next load origin, then the average reduction in travel is more than 6 percent of the total distance traveled.

In addition to these average numbers, it is important to gain insight into the worst-case performance of this strategy. If the service horizon is short, say, less than 10 demands served, it is possible to make one or more diversion decisions that result in an overall cost (distance) increase over the demand horizon. However, over a longer service horizon, diversion outperforms the base case more than 99 percent of the cases. Figure 6 gives the expected probability of overall gain and loss, respectively, along with the associated magnitudes of the gains and losses over different demand horizons. Each set of numbers is based on 10,000 simulated realizations of the corresponding sequence of random demand locations. Of course, overall gain here corresponds to the sequence, not to individual diversion decisions. The expected gains (or losses) are given in terms of fractions of the overall distance traveled under the base case (no diversions). The reported decreases (and increases) are conditional values given that the particular sequences experienced a decrease (or increase) under the diversion rule. The expected overall gain (or loss) over a sequence of calls is given by

$$\begin{aligned} \text{E[gain]} = {} & \text{E[gain} \mid \text{gain} > 0]\, p(\text{gain} < 0) \\ & - E[\text{gain} \mid \text{gain} < 0]\, p(\text{gain} < 0) \end{aligned} \quad (5)$$

The results in Figure 6 indicate that even with a horizon with as few as 10 diversion points the probability of overall loss is only 11.6 percent, with a corresponding expected loss of 2.3 percent of the overall distance under the base case. On the other hand, the 84.5 percent likelihood of gain is accompanied by an expected gain greater than three times the expected loss (7.1 percent reduction in overall cost). Note that for the 10-demand case the likelihood of 0 gain (no diversions chosen at all) is about 3.9 percent. The probability of loss rapidly decreases with the service horizon considered, to less than 1 percent with 50 demands (accompanied by an insignificant loss of under 1 percent, whereas the corresponding gain is over 6 percent with over 99 percent probability of gain. The fact that the probability of loss and the expected conditional loss are extremely small makes the diversion operating strategy appear to be somewhat of a win-win strategy. The simple diversion criterion of comparing the relative distances to serve a pair of loads sequentially appears relatively robust. If it suggests that diversion is profitable diversion is done, with a very high probability of realizing some meaningful benefit, and if it does not, the original plan is followed.

To further reduce the likelihood of loss over a finite number of decisions one can introduce a threshold in the diversion rule, whereby the local gain is required to exceed some minimum level to trigger a diversion, as follows:

If

$$\begin{aligned} & d(p, X_3) + d(X_4, X_1) < d(p, X_1) + d(X_2, X_3) \\ & \quad - T[d(p, X_1) + d(X_2, X_3)] \end{aligned} \quad (6)$$

then

divert and serve load X_3 to X_4 first,

where

p = current diversion point,
X_1, X_2 = origin and destination locations of current load,
X_3, X_4 = origin and destination of newly arrived load, and
T = threshold multiplier corresponding to the minimum relative improvement associated with a given diversion.

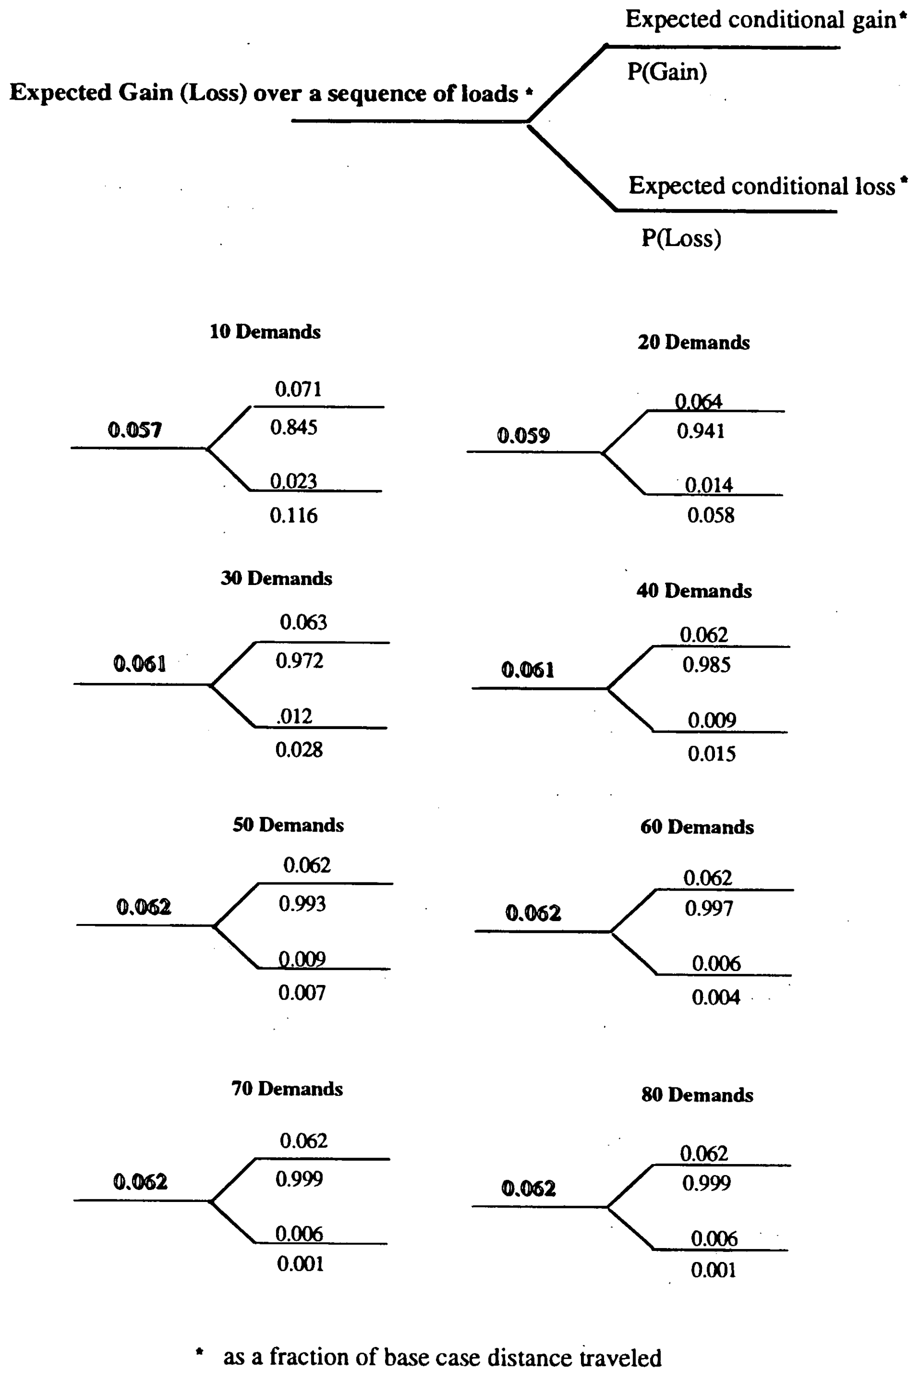

FIGURE 6 Benefits of diversion (distances to serve both loads considered).

This multiplier was varied from 0 to 0.5 (50 percent in the present analysis. Results suggest that a threshold value of about 10 percent of the base case cost yields the best performance. However, although the addition of a threshold for diversions reduces the risk of bad diversions, any threshold that precludes many positive diversions results in a reduction of expected benefits. The addition of the threshold rule for diversion increased the overall benefits by an amount between 1 and 0.3 percent, depending on the demand threshold. However, more significantly, the thresholds cut nearly in half the already low probability of loss in each case.

In these simulations, the diversion point fraction, α, varied uniformly between 0 and 1. The diversion decision was based on an entire sequence of moves associated with the current and new demand points. Alternatively, Figure 7 shows the results of a diversion strategy under a strictly myopic or greedy strategy of diverting to the closest origin point, that is, if $d(p, X_3) < d(p, X_1)$ then divert to load X_3.

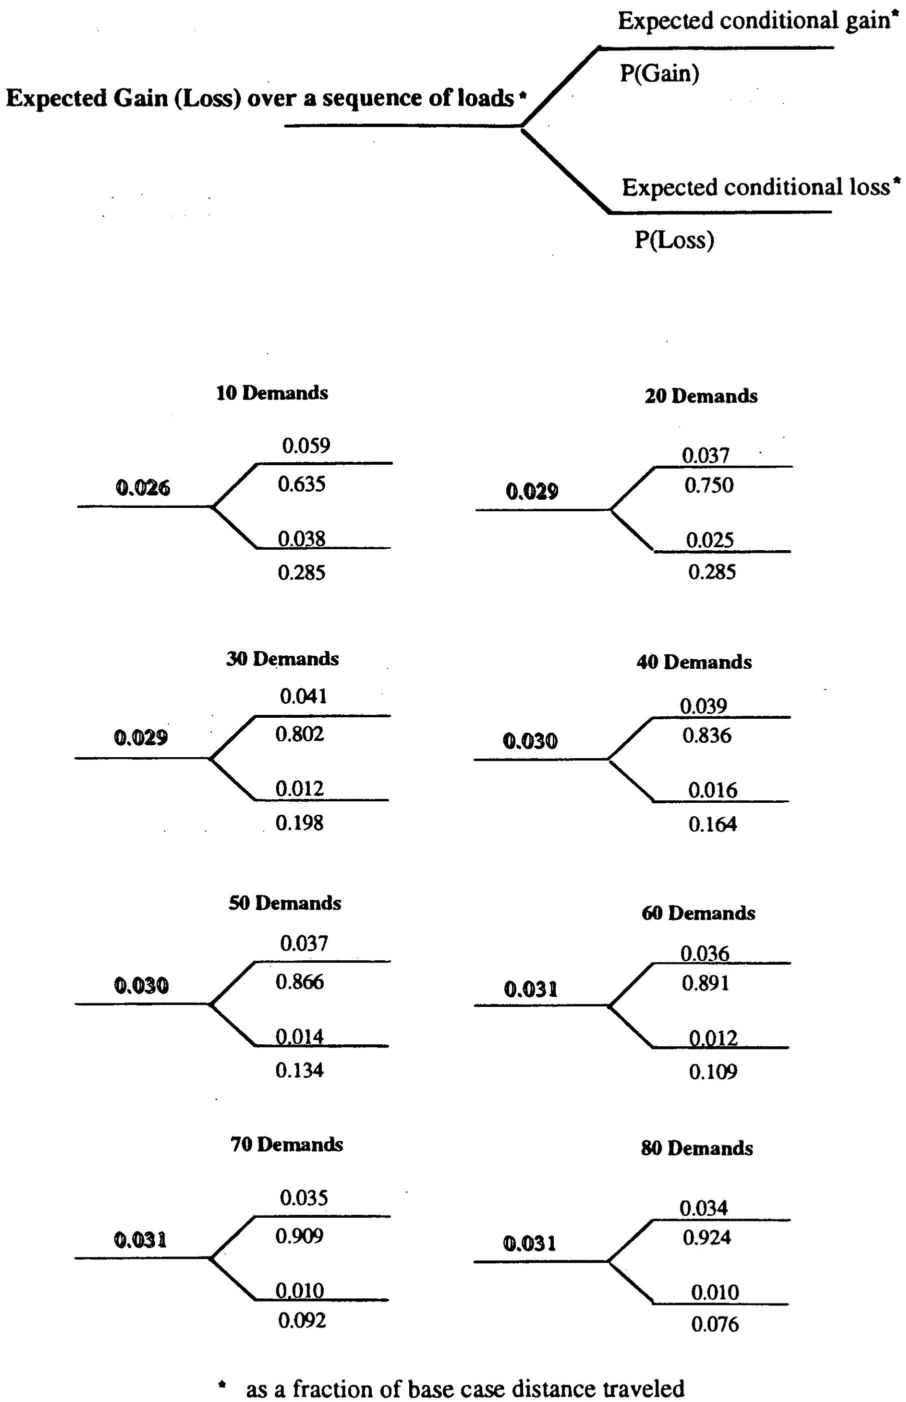

FIGURE 7 Benefits of purely greedy diversion strategy.

Two interesting points can be noted about these two sets of results in Figures 6 and 7. The first is that despite the limited information employed, that of which origin is closer, the greedy diversion strategy consistently leads to a reduction in overall expected travel over the demand horizons considered. The second is that by slightly increasing the amount of information considered, by also considering the distances that must be traveled empty between the two loads, the overall benefits double from 3.1 percent of the distance traveled to 6.2 percent over a demand horizon of 60 points. More importantly, the risk of overall loss is reduced significantly. For example, the probability that the diversion strategy will result in an overall cost increase with 60 demands is reduced from 0.109 (in the greedy case) to 0.004 when the distances to the load destinations are also considered. This is nearly 2 orders of magnitude less. Overall, these results demonstrate the potential power of reacting to even small amounts of real-time information on the state of the system.

EXTENSIONS OF SINGLE-VEHICLE, TWO-DEMAND CASE

This exploration under highly idealized conditions suggests that even a simple local diversion strategy is highly likely to result in a reduction of overall distance traveled. After considering only two demands at a time, the analysis is extended to consider several demands in a particular decision to divert and to look at demands that are uniformly generated in space but arrive according to a Poisson arrival stream as well as from a uniform distribution. In addition, operational constraints in which every demand must be served, and those in which one has the freedom to accept or reject demands according to the cost of serving them have been explored (*6*).

Naturally the performance of a given diversion strategy can be compared relative to several possible benchmarks or base cases with differing results. In addition to the base case of serving demands in the order in which they arrive, an "intelligent base" case in this analysis, is also considered next.

Poisson Arrival Stream, Optimal Resequencing, All Demands Accepted

This scenario has the following assumptions:

- Demands are generated from a Poisson arrival stream over time,
- The rate of arrival is rapid enough that more than one new demand may arrive while the vehicle is en route to a pick-up, and
- Demands diverted away from or not chosen for diversion are added to a queue and resequenced optimally with respect to overall distance traveled before being served.

This diversion strategy is compared with two different base cases. The first assumes service in the exact order of arrival, as considered previously, whereas the second, an intelligent base case, assumes that any demands waiting for service are resequenced optimally at the completion of each loaded movement. This intelligent base strategy is applied with optimal resequencing of up to five demands at a time. This itself results in solutions that are only 1 to 2 percent higher than those attained under the comparable diversion strategy in terms of overall travel distance. Under the assumption that all demands must be served, with demands generated from a Poisson arrival stream and with an arrival rate rapid enough to produce diversion opportunities, this intelligent base scenario leads to savings of more than 12 percent of the base case travel distance, and the en route diversion strategy tends to improve on the intelligent base by about 1 to 2 percent.

The dashed lines in Figure 8 show the various distances (costs) compared when choosing to divert, resequence, or serve the demands as they arrived in the case where two demands are in the queue and while the driver is en route to the current demand, origin X_1. The strategy chooses the minimum cost case of the $n!$ alternative orderings, where n is the number of demands in queue. It is assumed that n is a small number, say, less than 6, since to enumerate all alternatives for even slightly larger queues would take a prohibitive amount of time. This assumption of short queues makes sense in the trucking application where drivers typically have one or two jobs queued at most. Using the notation in Figure 8, when a new demand arises the minimum of the following six quantities corresponding to all possible service sequences is evaluated:

$$a\colon d(p, X_1) + d(X_2, X_3) + d(X_4, X_5)$$

$$b\colon d(p, X_1) + d(X_2, X_5) + d(X_6, X_3)$$

$$c\colon d(p, X_3) + d(X_4, X_5) + d(X_6, X_1)$$

$$d\colon d(p, X_3) + d(X_4, X_1) + d(X_2, X_5)$$

$$e\colon d(p, X_5) + d(X_6, X_1) + d(X_2, X_3)$$

$$f\colon d(p, X_5) + d(X_6, X_3) + d(X_4, X_1) \quad (7)$$

Cases e and f represent a diversion to the new load; Cases c and d represent diversion to a load already in the queue, one that was previously found unprofitable to divert to or to place first in the queue but was resequenced because of the information on the new demand. Cases a and b represent no-diversion cases, that is, the vehicle proceeds as before, but the second and third loads may be resequenced if that is beneficial.

Various extensions of these rules were explored, and it seems that under the assumption that eventually all demands must be serviced, a strategy that allows diversion but limits the number of times that one diverts before some demand is serviced is better than one that allows diversion whenever it is locally better. If diversion is allowed whenever it appears (locally) beneficial, under these assumptions costs may be low early in the service horizon and then considerably higher at the end.

Investigation of Poisson Arrival Stream, Optimal Resequencing, Loads Accepted or Rejected on Basis of Cost to Existing Route

This investigation has a different assumption from the preceding case with respect to load acceptance, namely, many demands are generated over time and loads may be accepted or rejected. This case assumes a rapid arrival rate for new demands. Whenever a new demand becomes known and space is available in the queue adding it to the current queue is considered. Rather than inserting the new demand into the existing route, resequencing the route in light of the new demand is considered. Because to optimally resequence the whole route would be computationally expensive (and possibly infeasible) the marginal cost of adding a demand to one of the first five slots in the queue is determined. If the additional empty distance needed to service the first four demands along with the new demand exceeds a given threshold value the new load is rejected. Otherwise the load takes an empty slot in the queue. These (up to) five demands are then resequenced optimally. Under the diversion strategy the load acceptance or rejection decision is made as soon as the demand becomes known if the vehicle is moving empty and as soon as it becomes empty if it is moving loaded. In the intelligent base case load acceptance decisions are made for all loads that have become known during the last period of service immediately after service is complete. These loads are evaluated for acceptance or rejection using the same logic as that in the diversion case (marginal cost to add to the first five slots in the queue) in the order in which they arrived. It was an a priori thought that the diversion strategy would perform well under these circumstances. However, it appears that excessive diversion creates a sort of "zig-zag" effect where a

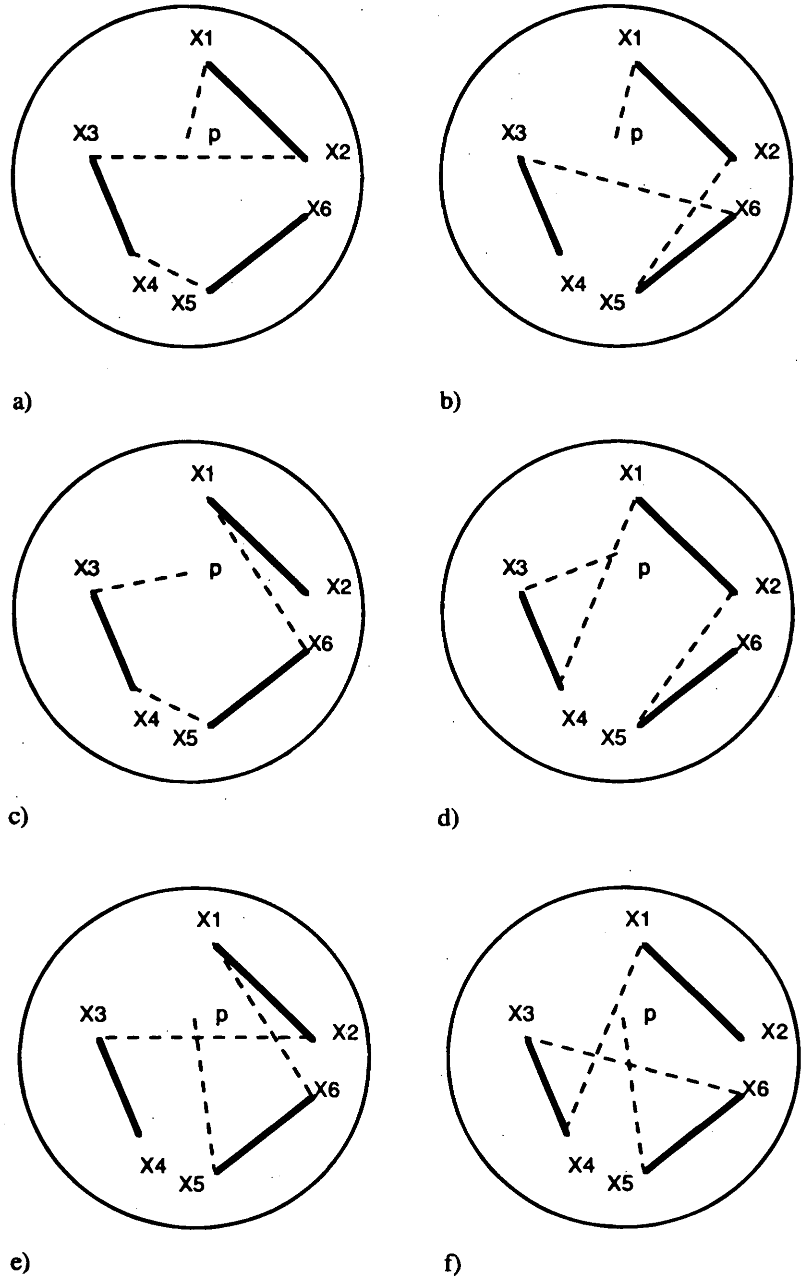

FIGURE 8 **Alternatives with two queued demands and a third arrival: *a*, no change in plan; *b*, no diversion, resequence; *c*, divert to previously considered demand; *d*, divert to previously considered demand; *e*, divert to new demand; *f*, divert to new demand. (Dashed lines represent empty movements, and solid lines, loaded movements.)**

vehicle is en route and then diverts and then diverts again. It appears that without additional constraints to restrict the amount of diversion a comparable intelligent base case in which the first few queued demands are optimally resequenced performs better. Table 1 provides a summary of the results presented in the last few sections.

CONCLUSIONS

Trucking operations consume a vast quantity of economic and environmental resources. A reduction in overall travel of even a few percentage points would represent a significant savings to both suppliers and consumers of trucking services. The U.S. Department of Transportation estimated that in 1991 motor vehicle fuel purchases accounted for 7.9 percent of common carrier costs or about $8.7 billion nationally (*7*). If 10 percent of these vehicles had a 5 percent reduction in fuel consumption, $43.5 million would be saved each year. In this work we identify and explore potential uses of real-time information for the efficient management of truckload carrier operations. Findings suggest that the diversion strategy examined may result in reduced travel distances and hence improved efficiency under certain conditions. Such strategies could become one part of an overall assignment and load acceptance strategy for truckload operations. The exploration of idealized scenarios suggests that a reduction of overall travel distance of between 5 and 10 percent would not be unreasonable. Although they are not intended to exactly replicate actual operating conditions, these scenarios do provide a simplified representation that allows the extraction of basic insights into the potential benefits of real-time information, the factors that affect their benefits, and the identification and design of strategies that merit examination under more realistic operating conditions.

Continuing developments include extending this analysis to a more 'realistic' scenario with respect to geographic region studied and customer demand stream. In addition, a fleet of vehicles rather than a single truck is examined. As the availability of automatic vehicle location and two-way communication technologies improves, and as the cost of equipping vehicles with these technologies decreases, more and more fleets will incorporate these technologies into their daily operations. Recent interviews with carrier company executives and fleet managers suggest that they are eager to incorporate communications technologies and optimization tools into their operations but that much work remains to be done in terms of developing such tools in a manner that is responsive to actual operating realities. An additional benefit of such tools is that they would enable companies to find good solutions to a complicated multiobjective problem. An addition to the goal of reducing overall distance driven is that of matching a driver with the load that best meets his or her needs. Discussions with industry executives have pointed out that irregular route truck drivers may stay on the road for more than 3 weeks at a time before finding an opportunity to pull a load in the direction of their home base. Flexible assignment strategies would improve the chances of finding a load that meets the preferences of an individual driver. A data base management system that would likely be an adjunct to any computer-based dispatching system could make the preferences of individual drivers easily accessible. It is clear that there are many potential uses of new technologies in commercial vehicle operations in general, and freight carrier operations in particular. Although technologies are beginning to see widespread use, it is equally clear that their full potential will not be realized until responsive real-time dispatching tools become available.

TABLE 1 Summary of Key Results

Scenario	Results
Demands generated a fraction of the way towards the first demand, 2 demands evaluated in each simulation.	Savings of 1-7 percent of the base travel distance depending upon the diversion point fraction.
Demands generated a fraction of the way towards the current demand, 100 demands evaluated in each simulation.	Savings of 1.5-12.5 percent of the base travel distance depending upon the diversion point fraction.
Demands generated from a Poisson arrival stream, all demands served, optimal resequencing of up to first five demands, diversion to new or queued demand.	Savings of 13-14 percent base case travel distance when compared to the base base, 1-2 percent of the intelligent base case travel distance when compared to the intelligent base case.
Demands generated from a Poisson arrival stream, accepted or rejected for service based upon space in queue and the cost of providing service given the current queued demands, optimal resequencing of up to first five demands, diversion to new or queued demand.	No comparison to the base case, little or no savings when compared to intelligent base case because of zig-zag effect.

Notes: The base case refers to serving the demands in the order in which they arrive. The 'intelligent base' refers to the case in which en-route diversion is not allowed, but a short queue of demands is resequenced optimally prior to the start of new service.

REFERENCES

1. Powell, W. B., A Comparative Review of Alternative Algorithms for the Dynamic Vehicle Allocation Problem. In *Vehicle Routing: Methods and Studies* (B. L. Golden and A. A. Assad, eds.). Elsevier Science Publishers, Amsterdam, The Netherlands, 1988, pp. 249–291.
2. Golden, B. L., and A. Assad. Perspectives in Vehicle Routing: Exciting New Developments. *Operations Research,* Vol. 34, No. 5, Nov.–Dec. 1986, pp. 803–810.
3. Dejax, P. J., and T. G. Crainic. A Review of Empty Flows and Fleet Management Models in Freight Transportation. *Transportation Science,* Vol. 21, No. 4, Nov. 1987, pp. 227–249.
4. French, R. L. Introduction to Intelligent Vehicle Highway Systems (IVHS). Presented at VNIS '91 International Conference on Vehicle Navigation and Information Systems, Dearborn, Mich., Oct. 1991.
5. Regan, A. C. Emerging Technologies in Automatic Vehicle Location and Two-Way Communication Systems for Common Carrier Operations. Presented at 1993 International Motor Carrier Symposium, Williamsburg, Va., May 1993.
6. Regan, A. C., and H. S. Mahmassani. *Evaluating Real-Time Dispatching Strategies for Truckload Trucking Operations: Simulation Model Results.* Working paper. Department of Civil Engineering, University of Texas, Austin, July 1994.
7. Schmitt, R. R., and T. D. Feinberg (eds.). *Transportation Statistics Annual Report 1994.* U.S. Department of Transportation, 1994.

Publication of this paper sponsored by Committee on Transportation Supply Analysis.