Competitive Performance Assessment of Dynamic Vehicle Routing Technologies Using Sequential Auctions

Miguel Andres Figliozzi, Hani S. Mahmassani, and Patrick Jaillet

Technologies for a dynamic truckload pickup-and-delivery problem in a competitive environment by use of sequential auctions are compared. In this environment, demands arrive randomly over time and are described by pickup-and-delivery locations and hard time windows. On demand arrival, carriers compete for the loads in a second-price auction. Four fleet assignment technologies with different degrees of sophistication are tested with simulations. The technologies differ in how they deal with the combinatorial and stochastic elements of the online problem. A one-step look-ahead dynamic vehicle routing technology that tries to estimate the impacts of current decisions on serving future loads (not yet arrived) is introduced. The performance of each technology is analyzed in relation to different demand characteristics.

The principal focus of the research described here is to compare and evaluate the performance of dynamic vehicle routing (DVR) technologies in a competitive environment. This investigation covers the truckload pickup-and-delivery problem, which entails the dynamic operation of trucking fleets that provide service to a general pattern of stochastic time-sensitive customer loads. The research uses an innovative approach in which DVR technologies are compared in a marketplace environment; each arriving load triggers an auction in which carriers compete with each other to win the right to service the load.

On the demand side, the motivation for this work is twofold. The growing demand for customer-responsive, made-to-order manufacturing is increasing the time sensitivity of customer demands. This trend is shifting the logistics and transportation process from one that relies on long planned lead times to one that relies on real-time operations. The operation of Dell computers in Texas exemplifies this trend.

The explosive growth in business-to-business electronic commerce, which is changing the way in which business is conducted, motivates the usage of marketplaces. A specific example of these changes is the increasing use of private exchanges, in which a company invites selected suppliers to interact in a real-time marketplace, compete, and provide the required services. Private exchanges are growing in number. A report published in mid-2002 estimated that as of June 2003, 15% of all Fortune 2000 companies would have set up private exchanges (*1*). Furthermore, the same source indicated that an additional 28% of all Fortune 2000 companies planned to implement a private exchange by the end of 2003. These market changes have produced a recent shift away from fixed pricing and toward flexible pricing. The typical dynamic pricing mechanisms are auctions, in which prices and allocations are based on bids.

On the supply side, the use of different DVR technologies reflects the asymmetric nature of competition. Even though carriers may compete in the same market, they are endowed with inherently different resources, ranging from physical assets such as fleets and facilities to communication and decision support systems. Furthermore, the levels of adoption of communication technology and expertise by carriers may vary greatly (*2*). This research presents the results of simulation experiments performed to test different DVR technologies. These technologies can be reduced to algorithms when simulated; however, they are referred to herein as "technologies," as their real-life implementation requires much more than implementation of an algorithm. Implementation may require upgrades in communication and decision support systems, software, computational power, trained personnel, as well as understanding of the nature and complexity of the DVR problem.

The paper is organized as follows: the next section describes the marketplace framework and operation. Previous work related to the DVR problem is then discussed. This is followed by descriptions of the different DVR technologies simulated in this paper and the simulation framework developed to illustrate and evaluate technologies through numerical experiments. The experimental results are analyzed and discussed, followed by concluding comments.

MARKET DESCRIPTION

This paper focuses on the performance of different DVR technologies in sequential auction transportation marketplaces. The marketplace enables the sale of cargo capacity mainly on the basis of price, yet it still satisfies customer level-of-service demands. The specific focus of the study is the reverse auction format, in which shippers post loads and carriers compete over them (bidding). The auctions operate in real time, and transaction volumes and prices reflect the status of demand and supply. A framework for the study of transportation marketplaces has been presented by Figliozzi et al. (*3*).

The market comprises shippers, which independently call for shipment procurement auctions, and carriers, which participate in them (it is assumed that the likelihood of two auctions being called at the same time is zero). Auctions are performed one at a time as shipments arrive at the auction market. Shippers generate a stream of shipments with corresponding attributes, according to predetermined probabil-

M. A. Figliozzi, University of Sydney, Institute of Transport Studies, Faculty of Economics and Business, Sydney NSW 2006, Australia. H. S. Mahmassani, Department of Civil and Environmental Engineering, Martin Hall, University of Maryland, College Park, MD 20742. P. Jaillet, Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139-4307.

Transportation Research Record: Journal of the Transportation Research Board, No. 1882, TRB, National Research Council, Washington, D.C., 2004, pp. 10–18.

ity distribution functions. A shipment attribute is its reservation price, or the maximum amount that the shipper is willing to pay for the transportation service. It is assumed that the auction announcement, bidding, and resolution take place in real time, thereby precluding the option of bidding on two auctions simultaneously.

The auction type used is the second-price or Vickrey auction. The second-price auction for one load operates as follows:

1. Each carrier submits a single bid.
2. The winner is the carrier with the lowest bid (which must be below the reservation price; otherwise, the auction is declared void).
3. The item (shipment) is awarded to the winner.
4. The winner is paid either the value of the second-lowest bid or the reservation price, whichever is the lowest.
5. The other carriers (not winners) do not win, pay, or receive anything.

A powerful characteristic of (one-item) Vickrey auctions is that the optimal strategy is simply to bid the true cost of serving the shipment. In economic terminology, this property is called a "truth-revealing mechanism." Furthermore, it is assumed here that the mechanism implemented in the marketplace for the sequence of bids is truth revealing at each bidding instance. This powerful assumption is necessary because the objective of the paper is to compare different DVR technologies. Bidding strategies are therefore reduced to marginal-cost bidding, removing any strategic or speculative element from the results obtained. Readers interested in the latter might refer to the work of Krishna (*4*) for a general treatment of auctions and to the work of Figliozzi et al. (*5*) for an analysis of carrier bidding and behavior in transportation auctions.

FORMULATIONS AND SOLUTIONS OF DVR PROBLEM

The DVR problem is a relaxation of the static vehicle routing problem, in which information about the demand or shipments to be served unfolds over time. Stochastic arrival times and shipment characteristics differentiate the DVR problem from the vehicle routing problem. Stochasticity transforms an NP (nonpolynomial) hard combinatorial optimization problem (with complete information) into a decision-making problem under uncertainty (partial information), while it preserves all the intricacies associated with the original NP hard problem. Powell et al. presented an extensive discussion of dynamic network modeling problems that arise in logistics and distribution systems, including a priori optimization and on-line decision policies for stochastic routing problems (*6*).

Regan et al. analyzed the opportunities and challenges of using real-time information for fleet management (*7–9*). They also formulated and evaluated (using simulations) various heuristics for the dynamic assignment of vehicles to loads under real-time information. Subsequent work by Yang et al. (*10, 11*) introduced a static optimization-based approach and tested it against the previously developed heuristic rules. Their approach solves static snapshots of the DVR problem with time windows by use of an exact mathematical programming formulation (which is the basis for two of the technologies studied in this paper). As new input occurs, static snapshot problems are solved repeatedly, which allows a complete reassignment of trucks to loads at each arrival instance. Mahmassani and colleagues studied DVR strategies for fleet size operations, in which computational and response times are important constraints (*12, 13*). They also studied strategies for DVR under high arrival rates and priority loads.

A growing body of work focuses on the solution of the stochastic DVR problem. Powell (*14–16*) and Powell et al. (*17*) proposed a formulation based on a Markov decision process and several formulations using stochastic programming. Gendreau et al. (*18*) and Ichoua et al. (*19*) used tabu search to solve a DVR problem with soft time windows. Gendreau et al. suggested the use of information about future requests to solve the DVR problem (*18*). This paper delves further into this idea, presenting a methodology that uses information about future requests to estimate the cost of servicing a new load. More recently, Larsen et al. studied the DVR problem with different degrees of dynamism (defined as the percentage of demands that carriers typically do not know in advance) (*20*).

DVR TECHNOLOGIES

Carriers must keep in mind the marginal cost for each transaction, especially in a sequential auction that implements a truth-revealing mechanism. The loaded distance is not included in the marginal cost because it is assumed that all carriers have the same cost per mile; therefore, the addition or subtraction of a constant to or from all the bids (e.g., the loaded distance of an arriving shipment) does not alter the ranking of bids. Besides, if all carriers include the loaded distance in their bids, that term cancels out when profits are computed (the payment, which in this case is the second bid, and the winner's cost include the same constant: the shipment loaded distance). Shippers' reservation prices also do not include the loaded distance.

In real-time situations, this is an increasingly difficult task when optimal decision making involves the solution of larger NP hard problems and the necessity of taking into account the stochastic nature of future demands. This paper presents four levels of DVR technologies. These technologies are presented in an order that shows increasing levels of sophistication.

1. Base or naïve technology. A carrier that uses this type of technology simply serves shipments in the order in which they arrive. If the carrier has just one truck, it estimates the marginal cost of an arriving shipment, s, simply as the additional empty distance incurred when s is appended to the end of the current route. If the carrier has more than one truck, the marginal cost is the cost of the truck with the lowest appending cost. This technology does not take into account the stochastic or combinatorial aspect of the cost estimation problem and is considered one of the simplest possible. Nonetheless, it provides a useful benchmark against which to compare the performance of more complex and computationally demanding technologies.
2. Static truck optimal (STO). Carriers that use this technology optimize the static vehicle routing problem at the truck level. If the carrier has just one truck, it estimates s as the additional empty distance incurred when s is inserted or appended to the current route. If the carrier has more than one truck, the marginal cost is the cost of the truck with the lowest inserting or appending cost. Like the previous technology, it does not take into account the stochastic nature of the problem. This technology roughly stands for the best that a myopic truck driver (i.e., one who ignores the future but who has real-time information) can achieve.
3. Static fleet optimal (SFO). The carrier that uses this technology optimizes the static vehicle routing problem at the fleet level. If the carrier has just one truck, the technology is equivalent to the previous case. If the carrier has more than one truck, the marginal cost is

the increment in empty distance that results from the addition of *s* to the total pool of trucks and loads yet to be serviced. If the problem were static, this technology would provide the optimal cost. Again, as for the two previous technologies, it does not take into account the stochastic nature of the problem. This technology roughly stands for the best that a myopic fleet dispatcher (i.e., one who ignores the future but who has real-time information) can achieve. A detailed mathematical statement of the mixed integer program (MIP) formulation used by SFO is given elsewhere (*10, 11*). STO is a special case of the general SFO formulation.

4. One-step look-ahead fleet optimal (1LFO). As was the case for the previous carrier, this carrier optimizes the static vehicle routing problem at the fleet level. This provides the static marginal cost (*smc*) for the addition of *s*. However, this carrier also knows the distribution of load arrivals over time and their spatial distributions (this paper does not discuss how the carrier has acquired this information). Hence, the carrier can simulate whether and how much winning *s* affects the marginal cost of serving the next arriving load; this is the dynamic marginal cost (*dmc*) of serving *s*. Unlike the previous types of technologies, this carrier takes into account the stochasticity of the problem to estimate the static cost of serving *s* and the effect on the marginal cost of serving the next arriving shipment (a one-step look ahead). This technology roughly stands for what a fleet dispatcher with real-time information and knowledge of the future (yet unrealized probabilistic demands) can do. However, 1LFO is not an optimal technology; rather, it is a heuristic that tries to estimate how serving *s* affects the cost of serving the next shipment.

Define L as the set of loads (won bids) yet to be serviced by the carrier when s arrives; define $c{:}L \rightarrow \Re$ as a function that, given a set of loads and current fleet deployment return the minimal empty distance required to serve all the loads included in L; and define $S = \{s_1, \ldots, s_n\}$ as a simulated set of n loads, in which each of these loads represents a realization of the next unknown arriving load (immediately after s) by using the probability distribution functions that the shippers use to generate loads. Then

$$smc(s) = c(L + s) - c(L)$$

$$dmc(s) = \sum_i [c(L + s + s_i) - c(L + s_i)]/n \qquad s_i \in S$$

By use of the 1LFO technology, the marginal cost used to bid is the average between *smc* (s) and *dmc* (s). The average is taken to correct the static estimation of *smc* by taking into account the impact of serving s on the next bid.

Three cases are possible:

1. *smc* (s) < *dmc* (s), in which serving the arriving demand leads to a fleet deployment that causes a higher *dmc* (relative to the current *smc*); therefore, the bid is increased;
2. *smc* (s) > *dmc* (s), in which serving the arriving demand leads to a fleet deployment that causes a lower *dmc* (relative to the current *smc*); therefore, the bid is decreased; and
3. *smc* (s) = *dmc* (s), in which serving the arriving demand leads to a fleet deployment that does not affect the value of *dmc* (relative to the current *smc*); therefore, the bid is equal to *smc* (s).

The response or solution time is a key consideration in real-time applications. However, given that the objective of this paper is to analyze how much can be gained by using different technologies, it is assumed that carriers have enough computational power that they can always bid before another request comes in.

In all cases it is assumed that a carrier bids only if a feasible solution has been found. If serving *s* unavoidably violates the time window of a previously won shipment, the carrier simply abstains from bidding or submits a high bid that exceeds the reservation price of *s*.

SIMULATION FRAMEWORK

This paper studies truckload carriers that compete over a square geographic region. It is assumed that the length of each square side is 1 unit of distance. For convenience, trucks travel at a constant speed equal to one unit of distance per unit of time. Demands for truckload pickup and delivery arise over this area and over time. Origins and destinations of demands are uniformly distributed over the square area, so the average loaded distance for a request is 0.52 unit of distance. All the arrivals are random; the arrival process follows a time Poisson process. The expected interarrival time ($E[T]$) is $1/(K\lambda)$, where λ is the demand request rate per vehicle and K is the total market fleet size. Roughly, the average service time for a shipment is 0.77 unit of time (λ is equal to approximately 1.3 unit of time). The service time is broken down into 0.52 unit of time, which corresponds to the average loaded distance, plus 0.25 unit of time, which approximates the average empty distance (average empty distances vary with the arrival rates and the time windows considered). Three different Poisson arrival rates per truck per unit of time are simulated:

- $\lambda = 0.5$ (uncongested),
- $\lambda = 1.0$ (congested), and
- $\lambda = 2.0$ (extremely congested).

The shipments have hard time windows. Three different time windows are simulated. In all cases it is assumed that the earliest pickup time is the arrival time of the demand to the marketplace. The latest delivery times (LDTs), in an order that reflects increasing slackness, are as follows:

LDT1 = arrival time + 1 × (shipment loaded distance + 0.25) + 1 × uniform (0.0, 1.0)
LDT2 = arrival time + 2 × (shipment loaded distance + 0.25) + 2 × uniform (0.0, 1.0)
LDT3 = arrival time + 4 × (shipment loaded distance + 0.25) + 4 × uniform (0.0, 1.0)

The time windows are called TW1, TW2, and TW3, respectively. The first type of time window (TW1) provides hardly enough scheduling flexibility. The opposite can be said about the last type (TW3). The reservation price distribution is uniform (1.42, 1.52) for all shipments. In all cases, reservation prices exceed the maximum marginal cost possible (therefore, the highest bid) in the simulated area (≈1.41 units of distance). It is also assumed that all vehicles and loads are compatible; no special equipment is required for specific loads. In all simulations, four carriers are competing for the demands. Of the four carriers, three compete by using the base or naïve technology; the other carrier uses one of the more sophisticated technologies.

Multiple performance measures are used to evaluate the various technologies. The first is total profits, which equal the sum of all payments received by won auctions minus the empty distance incurred to serve all won shipments (as mentioned earlier, shipment loaded distances are not included in the bids and loaded distances cancel out when profits are computed). The second performance measure is the number of auctions won or the number of shipments served, which is an indicator of market share. The third is the carrier's average empty distance, or the average distance from the destination of one load to the origin of the next load served. Average empty distance is a measure of the scheduling efficiency of the DVR technology. The fourth performance measure is average shipment (served) loaded distance.

ANALYSIS OF EXPERIMENTAL RESULTS

Figure 1 compares the absolute performances of the STO and SFO technologies. The graphs in the two columns have comparable scales to facilitate visual comparisons. The results correspond to a marketplace with four carriers, three of which use the naïve technology and one of which uses a more sophisticated technology (either STO or SFO). Each carrier has a fleet of two trucks. Figure 1 shows how SFO outperforms STO in the number of won bids and the average empty distance, for the most part with wider time windows and slightly congested arrival rates. With wide time windows (TW4) and congestion (arrival rate = 8 arrivals per unit of time), SFO outperforms STO in all four performance measures.

From a competitive point of view, it is essential to see the relative performance of carriers that use these technologies with respect to

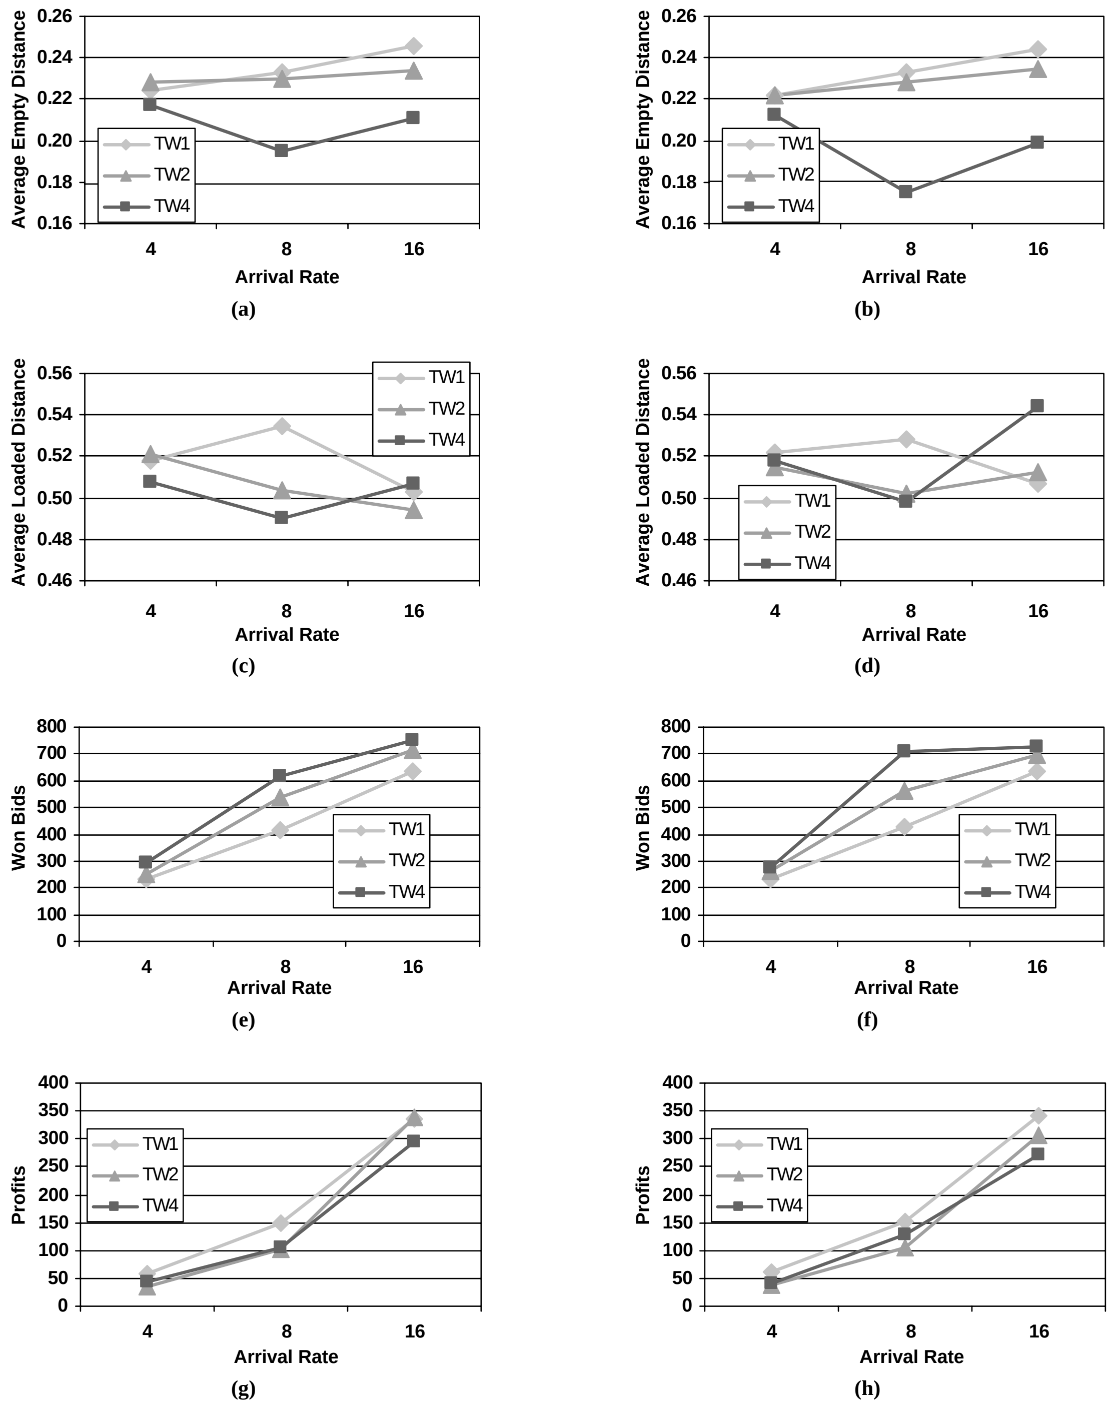

FIGURE 1 Comparison of performances of carriers using STO (*a*, *c*, *e*, *g*) and SFO (*b*, *d*, *f*, *h*) with that of carriers using naïve technology.

that of carriers that use naïve technologies. Figure 2 compares the relative performances of the carriers that use STO and SFO with those of the carriers that use the naïve technology. As expected, the sophisticated technologies in general outperform the naïve one. However, relative performance critically depends on the arrival rate and time windows that characterize the demand. Shorter time windows clearly have a negative effect on profit performance. The first type of time window (TW1) is so short that it provides few opportunities to improve schedules. The plots of won bids and profits look like flat lines. As time windows get larger, there are more opportunities to swap load service orders and to swap loads among trucks. For wider time windows, plots of won bids versus arrival rates have a concave curve.

Under low arrival rates, the carriers with optimal formulations have few shipments to work with. Under very high arrival rates, all the carriers' fleets are fully used, irrespective of their intrinsic tech-

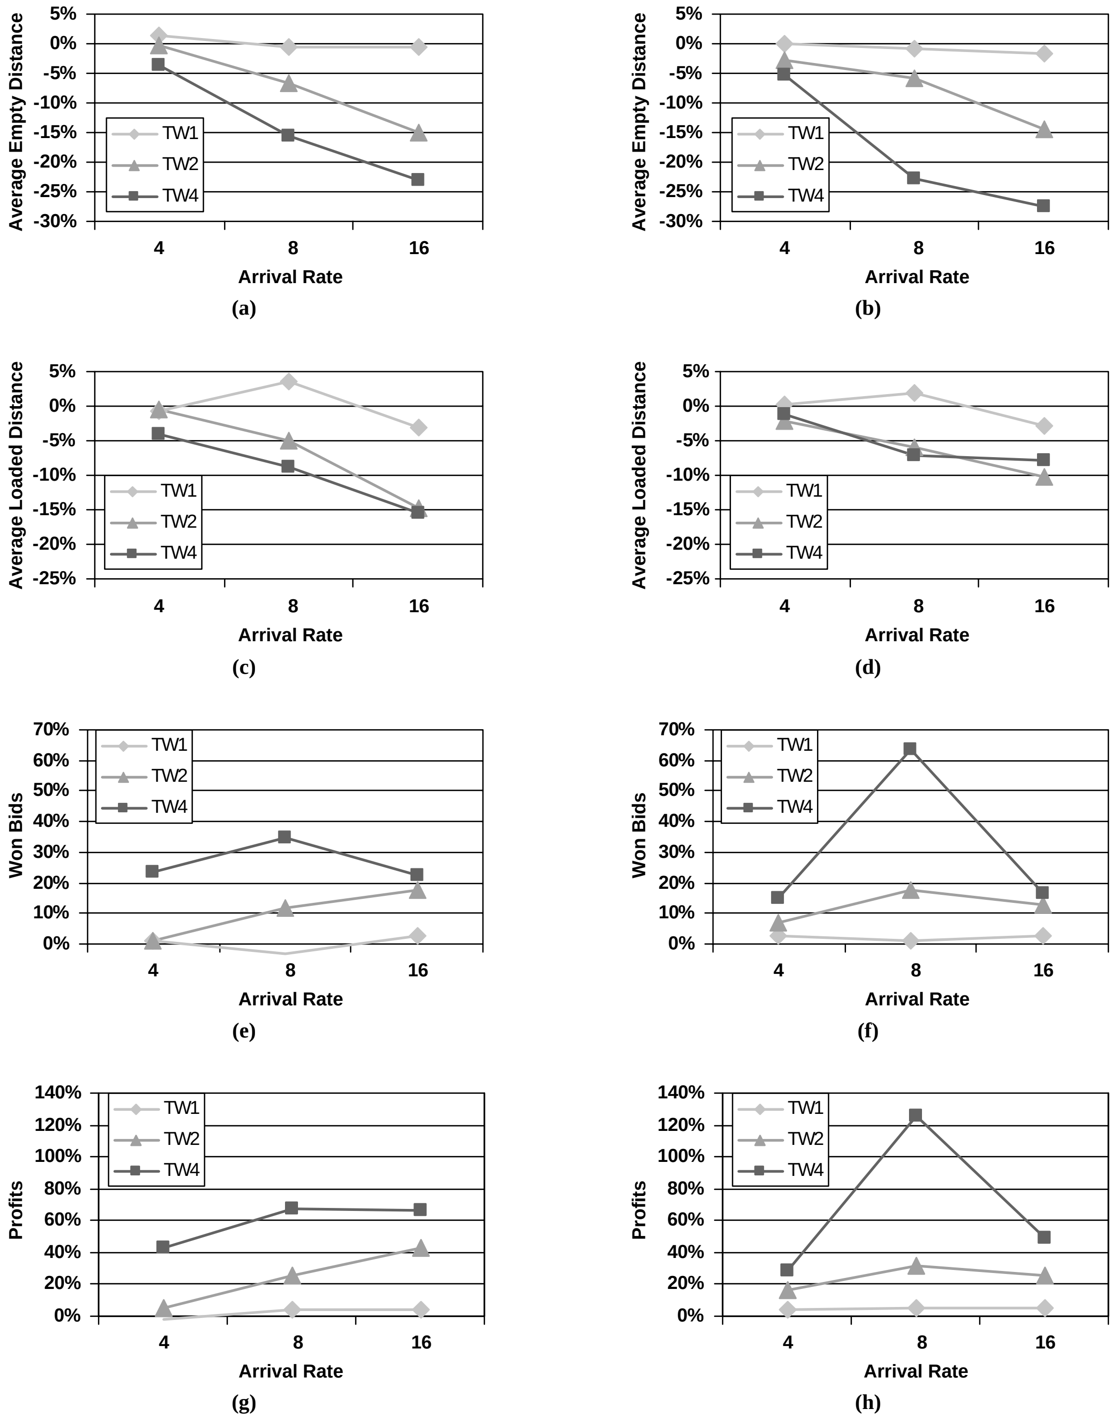

FIGURE 2 Percentage change for carriers using STO (*a*, *c*, *e*, *g*) and SFO (*b*, *d*, *f*, *h*) compared with the values for carriers using naïve technology.

nology or efficiency. Therefore, there is an optimal arrival rate that maximizes the competitiveness of static optimization techniques.

Average loaded distance is included as a performance measure because it reveals how the optimal static technologies are working. The MIP formulations manage to join together a larger number of short loaded distance shipments. As shipments arrive one at a time, the likelihood that a small shipment will fit into the existing routes increases as its loaded distance decreases. In this paper, reservation prices in the marketplace setting are independent of the shipments' loaded distances. In a marketplace in which shipment reservation prices strongly depend on loaded distances, however, carriers that use the sophisticated technologies would surely exhibit a different profit performance, although won bids, loaded distances, and empty distances could, by and large, remain unchanged.

Figures 3 and 4 depict graphics similar to those described above, but they illustrate the absolute and relative competitiveness of

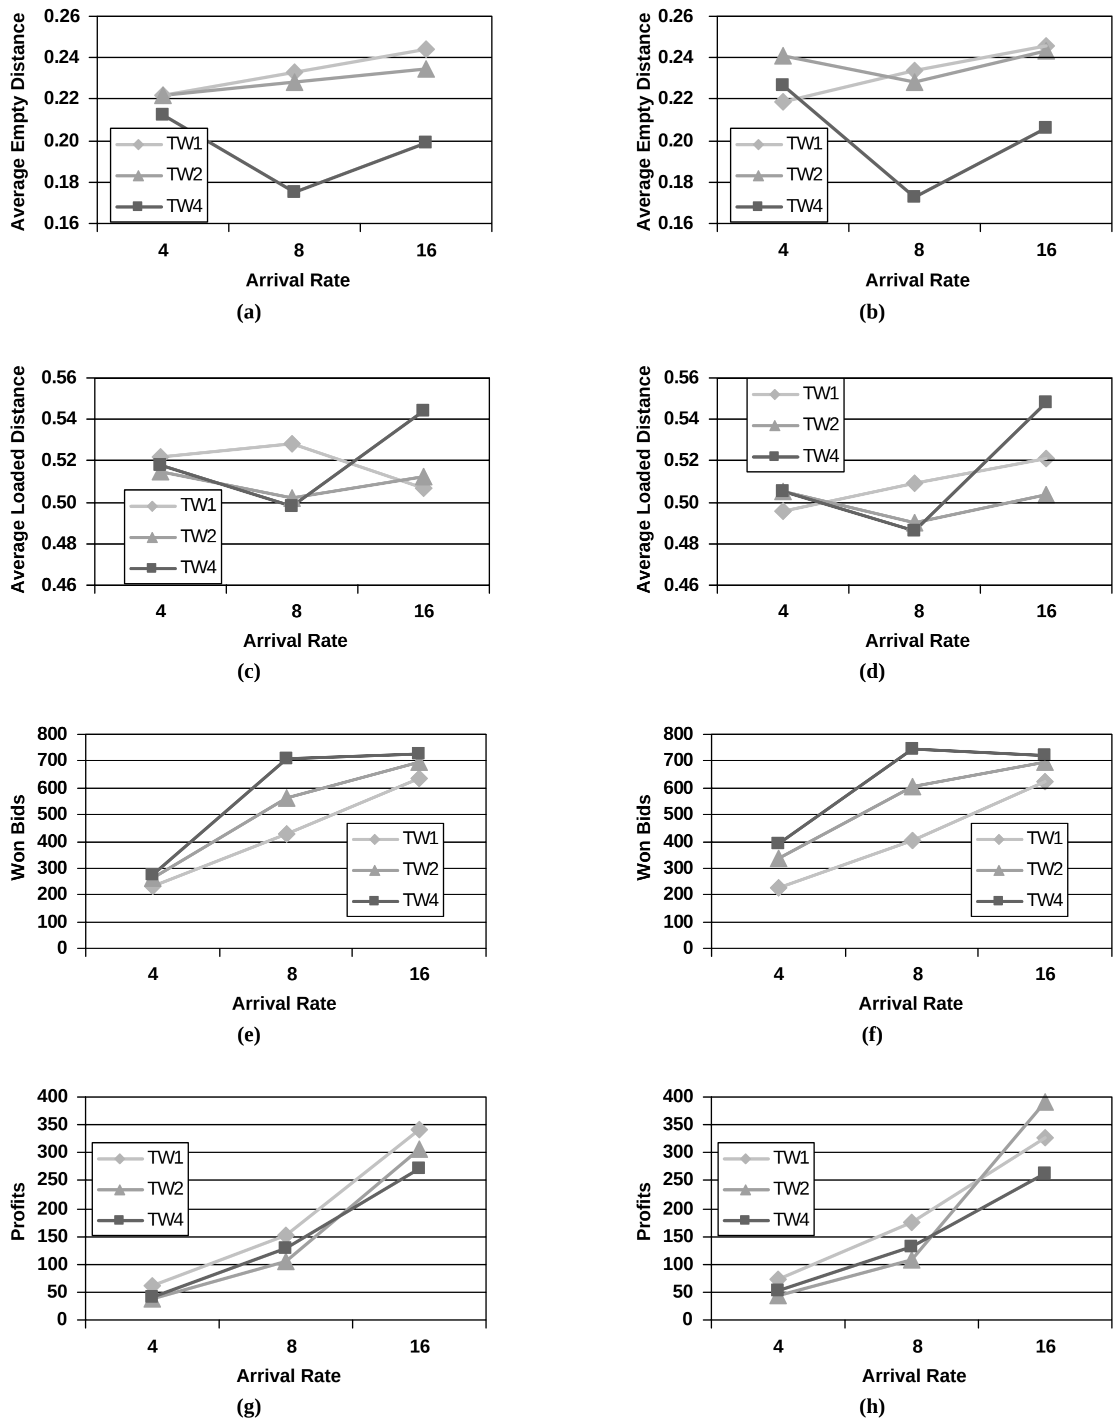

FIGURE 3 Comparison of performances of carriers using SFO (*a*, *c*, *e*, *g*) and 1LFO (*b*, *d*, *f*, *h*) with that of carriers using naïve technology.

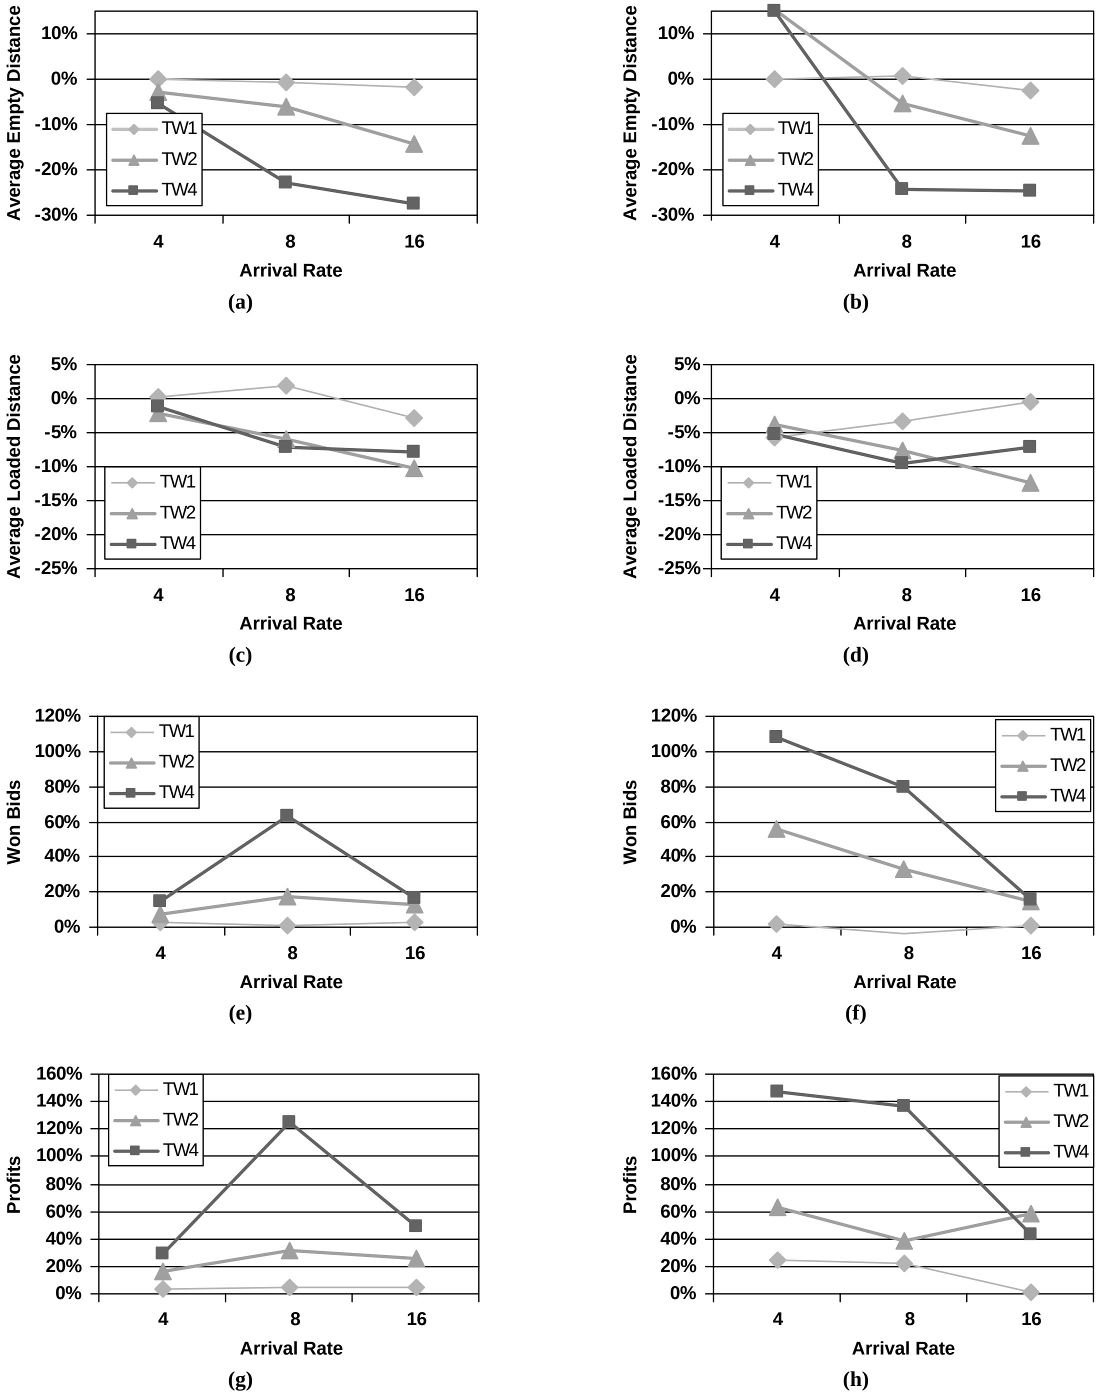

FIGURE 4 Percentage change for carriers using SFO (*a*, *c*, *e*, *g*) and 1LFO (*b*, *d*, *f*, *h*) compared with the values for carriers using naïve technology.

SFO and 1LFO technologies. The knowledge about the stochastic nature of the demand makes an obvious difference in the number of won bids (Figure 4). Averaging of *smc* and *dmc* seems to more accurately reflect the true marginal cost of serving a shipment. However, the relative performance of 1LFO decreases with increasing arrival rates. Under very high arrival rates (arrival rates = 32 arrivals per unit of time in Figures 3 and 4), the respective performances of SFO and 1LFO are comparable. This seems to suggest that at high arrival rates, *smc* is approximately equal to *dmc*.

Figure 5 illustrates the results for a carrier that uses the SFO technology and that is competing against three carriers that are using the STO technology. The graphs on the left represent the absolute performance of the carrier using SFO; the graphs on the right represent the relative performance of the carriers using SFO against the relative performance of carriers using STO. The same patterns discussed in Figures 1 and 2 hold here. Optimization at the fleet level is more competitive under moderately congested conditions and with wider time windows. Unlike the previous case, however, the average loaded

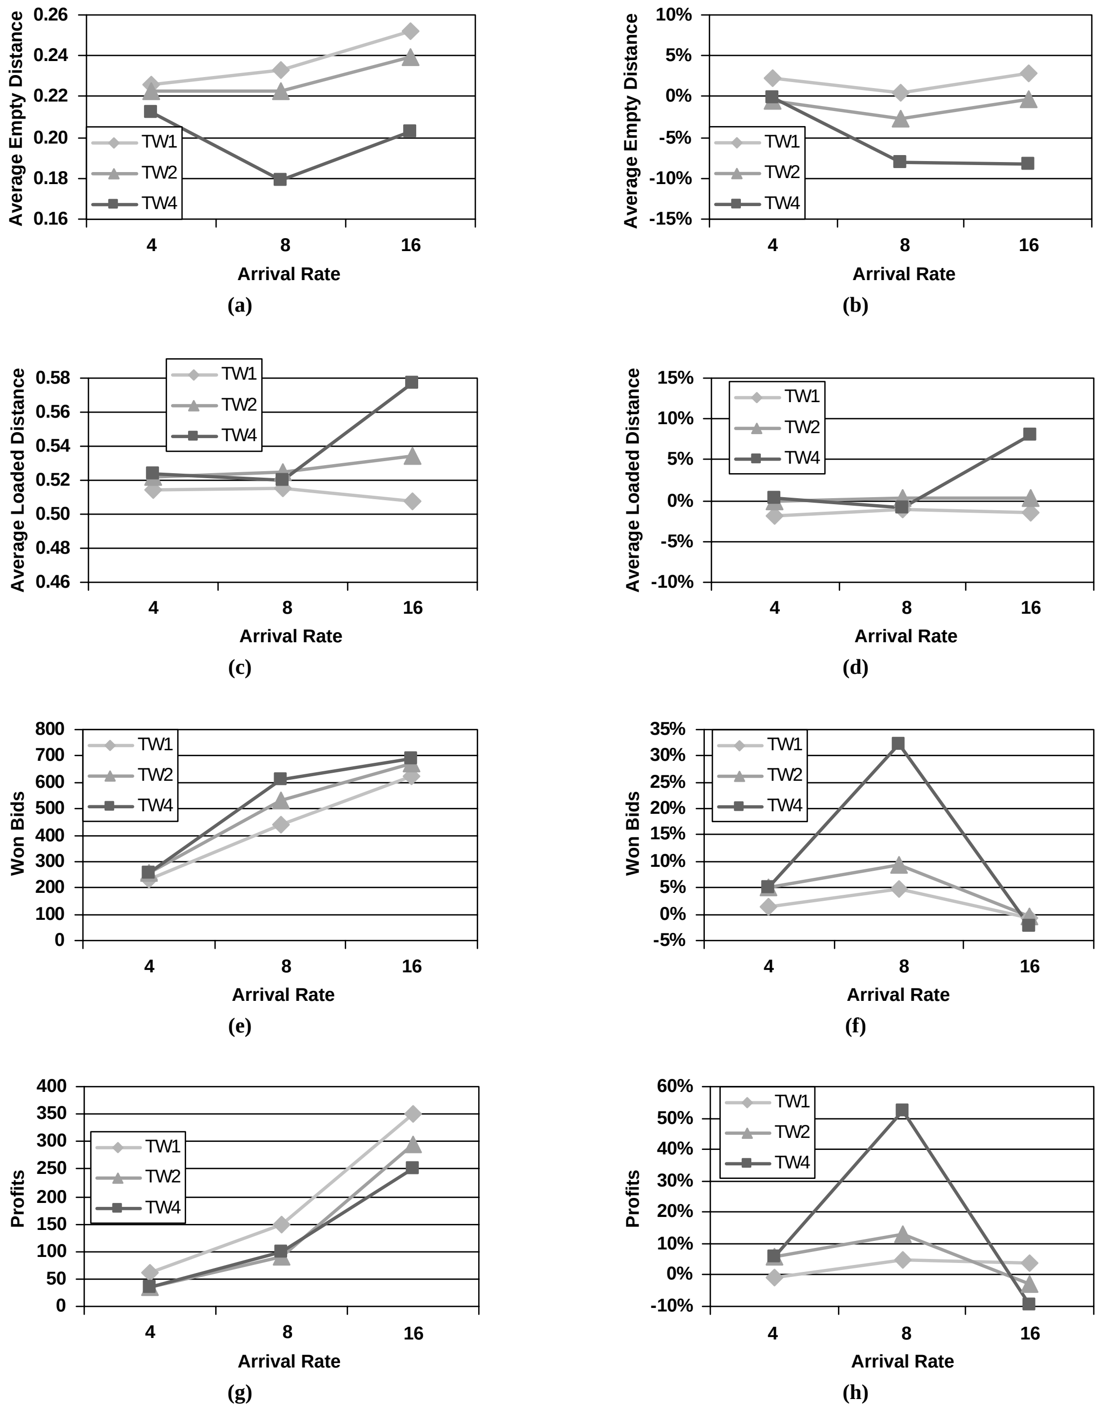

FIGURE 5 Comparison of SFO and STO: SFO absolute values (*a*, *c*, *e*, *g*) and percentage change for carriers using SFO compared with the values for carriers using STO (*b*, *d*, *f*, *h*).

distance does not vary as much, except for the case with a high arrival rate and wide time windows.

CONCLUSIONS

A sequential auction framework was used to compare the competitiveness of different DVR technologies. The technologies were evaluated under different demand conditions. It was shown that under severely constrained DVR problems the performance gains obtained with sophisticated technologies are scarce. Major performance gains are obtained under moderately congested conditions as well as with ample time windows.

The methodology proposed for the testing of DVR technologies seems adequate for the evaluation of competitive performance, especially for logistics and transportation problems that are embedded in dynamic stochastic environments or that support e-commerce marketplaces and activities. The paper also introduced a DVR strategy that assumes knowledge about the shipment arrival and other characteristics. Simulation results show that this new strategy outperforms myopic ones, particularly in uncongested marketplaces. Further research is needed, however, to fully tap the competitive edge provided by knowledge about the stochastic nature of future demand arrivals.

ACKNOWLEDGMENT

This study was supported by a grant from the National Science Foundation.

REFERENCES

1. Hoffman, W., J. Keedy, and K. Roberts. The Unexpected Return of B2B. *The McKinsey Quarterly,* 2002, p. N3. www.mckinseyquarterly.com/article_abstract.asp?ar=1210&L2=24&L3=47. Accessed June 2003.
2. Regan, A. C., and T. F. Golob. Freight Operators' Perceptions of Congestion Problems and the Application of Advanced Technologies: Results from a 1998 Survey of 1200 Companies Operating in California. *Transportation Journal,* Vol. 38, No. 3, 1999, pp. 57–67.
3. Figliozzi, M. A., H. S. Mahmassani, and P. Jaillet. Framework for Study of Carrier Strategies in Auction-Based Transportation Marketplace. In *Transportation Research Record: Journal of the Transportation Research Board, No. 1854,* TRB, National Research Council, Washington, D.C., 2003, pp. 162–170.
4. Krishna, V. *Auction Theory.* Academic Press, New York, 2002.
5. Figliozzi, M. A., H. S. Mahmassani, and P. Jaillet. Modeling Carrier Behavior in Sequential Auction Transportation Markets. Presented at 10th International Conference on Travel Behavior Research, Lucerne, Switzerland, Aug. 2003.
6. Powell, W. B., P. Jaillet, and A. Odoni. Stochastic and Dynamic Networks and Routing. In *Handbook in Operations Research and Management Science, Vol. 8. Network Routing* (M. O. Ball, T. L. Magnanti, C. L. Monma, and G. L. Nemhauser, eds.), Elsevier, Amsterdam, Netherlands, 1995, pp. 141–296.
7. Regan, A. C., H. S. Mahmassani, and P. Jaillet. Dynamic Decision Making for Commercial Fleet Operations Using Real-Time Information. In *Transportation Research Record 1537,* TRB, National Research Council, Washington, D.C., 1996, pp. 91–97.
8. Regan, A. C., H. S. Mahmassani, and P. Jaillet. Dynamic Dispatching Strategies Under Real-Time Information for Carrier Fleet Management. In *Transportation and Traffic Theory* (J. B. Lesort, ed.), Pergamon, 1996, pp. 737–756.
9. Regan, A. C., H. S. Mahmassani, and P. Jaillet. Evaluation of Dynamic Fleet Management Systems: Simulation Framework. In *Transportation Research Record 1645,* TRB, National Research Council, Washington, D.C., 1998, pp. 176–184.
10. Yang, J., P. Jaillet, and H. S. Mahmassani. On-Line Algorithms for Truck Fleet Assignment and Scheduling Under Real-Time Information. In *Transportation Research Record: Journal of the Transportation Research Board, No. 1667,* TRB, National Research Council, Washington, D.C., 1999, pp. 107–113.
11. Yang, J., P. Jaillet, and H. S. Mahmassani. Real-Time Multi-Vehicle Truckload Pickup and Delivery Problem. *Transportation Science,* Vol. 38, 2004, pp. 135–148.
12. Mahmassani, H. S., Y. Kim, and P. Jaillet. Local Optimization Approaches to Solve Dynamic Commercial Fleet Management Problems. In *Transportation Research Record: Journal of the Transportation Research Board, No. 1733,* TRB, National Research Council, Washington, D.C., 2000, pp. 71–79.
13. Kim, Y., H. S. Mahmassani, and P. Jaillet. Dynamic Truckload Truck Routing and Scheduling in Oversaturated Demand Situations. In *Transportation Research Record: Journal of the Transportation Research Board, No. 1783,* TRB, National Research Council, Washington, D.C., 2002, pp. 66–71.
14. Powell, W. B. A Stochastic Formulation of the Dynamic Vehicle Allocation Problem. *Transportation Science,* Vol. 20, 1986, pp. 117–129.
15. Powell, W. B. An Operational Planning Model for the Dynamic Vehicle Allocation Problem with Uncertain Demands. *Transportation Research,* Vol. 21B, No. 3, 1987, pp. 217–232.
16. Powell, W. B. A Stochastic Formulation of the Dynamic Assignment Problem, with an Application to Truckload Motor Carriers. *Transportation Science,* Vol. 30, No. 3, 1996, pp. 195–218.
17. Powell, W. B., M. T. Towns, and A. Marar. On the Value of Optimal Myopic Solution for Dynamic Routing and Scheduling Problems in the Presence of User Noncompliance. *Transportation Science,* Vol. 34, No. 1, 2000, pp. 50–66.
18. Gendreau M., F. Guertin, J. Potvin, and E. Taillard. Parallel Tabu Search for Real-Time Vehicle Routing and Dispatching. *Transportation Science,* Vol. 33, 1999, pp. 381–390.
19. Ichoua, S., M. Gendreau, and J.-Y. Potvin. Diversion Issues in Real-Time Vehicle Dispatching. *Transportation Science,* Vol. 34, No. 4, 2000, pp. 426–438.
20. Larsen, A., O. Madsen, and M. Solomon. Partially Dynamic Vehicle Routing—Models and Algorithms. *Journal of Operational Research Society,* Vol. 53, 2002, pp. 637–646.