THE CONTRIBUTION OF ADVANCED PACKAGE ARRIVAL INFORMATION TO EFFICIENT GROUND DELIVERIES BY INTERNATIONAL COURIERS

by

Michael A. Haughton
School of Business and Economics
Wilfrid Laurier University

Abstract

This paper quantifies the value of real-time access to information on the customs clearance status of packages consigned from overseas for subsequent truck transportation in enhancing efficiency in international couriers' ground delivery operations. The value is quantified by determining the unnecessary ground transportation costs of second-best tactics that become necessary when truck dispatch/routing personnel lack real-time he access to customs clearance information. Second-best tactics studied include forecasting each package's clearance status and adjusting ground delivery routes when actual arrival information becomes known. This study's findings indicate that the value can be substantial, amounting to hundreds of dollars per day for the key transportation resources: vehicles and drivers. The study also finds that use of forecasts is not always the most cost-effective second-best tactic, even when the forecasts might be considered very accurate.

1. Introduction

Initiatives such as the suite of package flow technologies being developed by UPS[1] epitomize ongoing efforts by major international courier companies refine the quality and reliability of real-time information on parcels from point of origin (where the consignor transfers custody to the courier) through to possession by the overseas consignee. A part of the economic justification for these initiatives is the expectation of lower cost of ground (truck) delivery operations. This paper models the potential cost reductions (in terms of costs associated with distance traveled and drivers/vehicles used to deliver packages to consignees). The research task amounts to addressing a

[1]Press release dated September 13, 2005 at the UPS website (www.pressroomom.ups.com/pressrelease/current/0,1088,4600,00.html)

retrospective "what if" question for international couriers that have undertaken or are pondering the aforementioned initiatives: *without the real-time package status information made accessible through those initiatives, how much costlier would it be to perform ground delivery operations*? Answers to this question clarify the significance of ground delivery costs in decisions to embark on these initiatives

The question is addressed by comparing the transportation cost of the optimal vehicle dispatch/routing tactic that is possible in situations of access to real-time package arrival information with the transport cost of second-best (suboptimal) tactics that are worth considering when such access is lacking. In dealing with this question, the remainder of the paper is organized into six sections. The first of these outlines the difficulties of optimizing ground delivery operations when the courier company's dispatch/routing personnel are forced to operate with imperfect and/or tardy package arrival information. Section 3 details next-best tactics that partially compensate for these deficiencies in the package arrival information. Section 4 reviews the relevant literature on information uncertainty in vehicle routing in order to clarify this paper's contribution to that literature. Section 5 then describes the key features of the research methodology. Section 6 presents the findings on the transportation cost comparisons, and Section 7 summarizes the study's main insights.

2. Problem Context

At each node and link in the multi-stage trans-border flow of parcels from consignor to consignee, the potential for unpredictable delay creates uncertainty about when a package will reach the consignee. For example, in light of the heightened post-9/11 worry about security vulnerabilities that might exist when merchandise traverse national borders, speedy customs clearance is less certain; e.g., more delays due to the more than doubling of the inspection rate of many shipments to the US, even after the initial panic of 9/11 had subsided (Douglas, 2002). The continued importance of this source of delays is reflected in the development of Flex® Global View by UPS as a way to transmit timely notification to relevant personnel and to customers about items held up in Customs for detailed inspection.

With all the activities to be performed at the distribution centre (DC) between arrival of parcels and the start of their dispatch by truck, 11th hour receipt of information on these delays makes it difficult for the courier's dispatch/routing personnel to optimally adjust the day's truck delivery plans. That is, for air arrivals, activities such as aircraft unloading, package sorting, ground vehicle loading, etc., leave little time to rerun vehicle routing programs to optimally reconfigure the routes. And even if there is time to finish that computer processing task, one still faces the additional and potentially bigger hassles of reissuing new instructions to drivers to follow the reconfigured routes and rearranging packages in the vehicles. UPS portrays an unwelcome consequence of deficient real-time information flows as follows: "*a driver often had to spend valuable time 'touching cardboard' – checking and rearranging packages in the back of the package car in the order he/she felt would be the most optimum for delivery*"[2]. This is especially problematic if the reconfigured routes are markedly different from the originally planned ones.

Customs clearance processes factor into the parcel arrival uncertainty that causes this downstream problem because customs authorities can select any given package for detailed scrutiny. A package may be selected either as routine execution of an inspection policy requiring examination of randomly selected packages or because of suspicion that the package exceeds some threshold risk of contravening customs regulations and/or posing a security threat. Therefore, a courier company's investments in telecommunication links with customs authorities and in package tracking systems enhance reliability in the transmission of sufficiently advanced notice of which packages have been selected for detailed customs scrutiny; i.e., packages for which such time-consuming scrutiny will cause them to miss their originally scheduled flight. Without these investments to facilitate the route reoptimization efforts of dispatch/routing personnel, only second-best tactics to address parcel arrival uncertainty are possible. The question of interest here is that, compared to the best case scenario of being able to perform daily route reoptimization, how much more costly are ground delivery operations in the case of (forced) use of second-best (next-best) tactics.

3. Next-Best Tactics

The next-best alternatives considered here involve activities before knowing package arrival, between knowing package arrivals and dispatching the pre-planned routes, and post-delivery recourse to handle packages that could not be accommodated on pre-planned routes. One pre-arrival activity involves forecasting which packages will arrive as scheduled (and hence, which will be non-arrivals because they might be selected for detailed customs scrutiny). In the period between having package arrival information and dispatching the pre-planned routes, the activities, inclusive of what UPS characterizes as drivers "touching cardboard", involve opportunistic ad hoc route adjustments to maximize deliveries while controlling transportation costs. Adjustments focused on cost control involve

[2]Pressroom fact sheet at the UPS website (www.pressroom.ups.com/mediakits/factsheet/0,2305,1134,00.html)

instructing drivers to bypass consignee addresses for the surprise non-arrivals (packages predicted to arrive as originally scheduled but did not) and adjustments to increase the number of delivered packages by making detours to visit consignee addresses for the surprise arrivals (predicted to be non-arrivals but did arrive). Since the forecast-based (pre-planned) routes are not always able to accommodate all surprise arrivals, a recourse action may be necessary to deliver the remaining surprise arrivals on that day (the assumption here is that all of a day's arriving packages will be delivered on that day). The recourse is likely to involve running a set of late-morning or early-afternoon routes for those packages. This paper examines three possible tactics. The key difference among the tactics can be viewed in terms of a trade-off between having excess capacity on the pre-planned routes and having to make adjustments to handle surprise arrivals. Description of each tactic follows.

Tactic 1: Optimistically project that all of the day's sent packages will arrive as originally scheduled. Plan the routes on the basis of the total of the day's optimistic projection plus yesterday's surprise non-arrivals (assumed to arrive with certainty today). Since this satisfies the largest possible delivery requirement (all arriving packages) the only possible cost-saving route change is to bypass consignee addresses for packages that do not arrive as scheduled. That move is not possible for today's surprise non-arrivals that were among yesterday's non-arrivals; i.e., since yesterday's non-arrivals will surely arrive today, they will be slated for a delivery and therefore cannot be bypassed.

Tactic 2: Forecast the arrival outcome for each package that is sent with the goal of arriving today and, as in Tactic 1, plan the routes on the basis of the total number of packages comprising today's forecasted arrivals and yesterday's non-arrivals. Treatment of surprise non-arrivals is also similar. Regarding insertions of surprise arrivals into pre-planned routes, this tactic avoids the task of determining the best insertions and conveying potentially complicated detour instructions to drivers by considering only the simplest of insertions. Specifically, a surprise arrival is inserted only when it is consigned to an address that is already on a pre-planned route; i.e., the driver is already slated to visit that address anyway so one need only place the consignee's parcel on the driver's truck (subject to truck capacity constraints). Unless these simple insertions cover all surprise arrivals, the aforementioned recourse of dispatching additional drivers to serve the remaining (un-inserted) surprise arrivals will be employed.

Tactic 3: The only significant distinction between this tactic and Tactic 2 is that detour-causing route insertions of surprise arrivals are allowed in order to reduce use of the recourse action. Tactic 3 reflects a scenario of belief that that tasks of determining the optimal insertions and ensuring that drivers efficiently follow the detour instructions can be effectively done. True, given the aforementioned time constraints caused by the frenetic pace of pre-dispatch activities at the DC, Tactic 3 might not always be easy to deploy. However, based on the previously noted use of ad hoc route adjustments at UPS prior to implementation of its suite of package flow technologies, Tactic 3 is judged to be, at least, *occasionally* feasible. Bolstering this judgment is the fact that the adjustments for Tactic 3 involves the ostensibly lighter burden of adding delivery stops to a driver's pre-planned route instead of a more significant overhaul of the route.

4. Literature Review and Extensions

The literature on problems caused by the presence of stochastic/uncertain elements in vehicle routing can be considered as containing many branches. One of these deals with the analysis of next-best routing alternatives that deserve consideration when access to relevant customer/demand information is too late for deployment of the response tactic that makes optimal use of transportation resources: daily route reoptimization. The present study falls within the ambit of that branch, which itself has several inter-related branches: see, for example, reviews by Gendreau, Laporte and Séguin (1996) and Baita, Ukovich, Pesenti, and Favaretto (1998). Papers on that issue can be traced back to Golden and Yee (1979), and includes, for example, Waters (1989), Benton and Rosetti (1992), and Haughton (1998, 2000). Together covering a variety of routing contexts, these papers confirm the superiority of route reoptimization over alternative tactics that involve modifying routes that are pre-planned using *a priori* optimization (i.e., based on *probabilistic estimation* of customer demand as opposed to *known* customer demand). Still these papers, as well as, for example, Bertsimas, Chevri, and Peterson (1995), also acknowledge the validity and practical necessity of second-best tactics. In particular, they acknowledge that the time constraint on performing tasks beyond mere (re)running of route optimization algorithms contribute to this necessity.

In gauging the value of advanced information in eliminating costly dependence on next-best tactics, one of this study's extensions of the extant literature is that it illustrates this value in realistic and useful monetary terms instead of just physical units (e.g., travel distance). More importantly, while anchored to the existing literature, the tactics studied here augment those covered in that literature on three key issues. The first extension is that this study explicitly accounts for forecasts of customer requirements (Tactics 2 and 3). Secondly, this study observes practical limitations on the insertion of customers into pre-planned routes. In particular, Tactic 2 captures the idea that if a driver has to make a detour, there is the risk that his unfamiliarity with the area within which he has to make the detour might cause him to take a circuitous path and increase travel distance (transportation cost). Thirdly, this study extends the range of structures of recourse routes. The literature's standard model is that, once depleted before completing its route, a vehicle will immediately return to the depot to be refilled, then be re-dispatched to resume its route by visiting the customers whose requests are yet to be fully satisfied. This can cause low vehicle utilization because each affected route gets a separate vehicle re-dispatch to visit only a few customers. Therefore, this paper considers an alternative approach: serve the affected customers en bloc later in the day; i.e., the latter two second-best tactics.

5. Research Methodology

The study involved simulation experiments. The hypothetical context for the experiments is a single distribution centre comprising 1000 consignees, each with a 50% chance of being sent a package from overseas through that DC each day. The DC is located at the centre of the service area, and the spatial dispersion of consignees conforms to a uniform distribution. The service area comprised the actual road network and residential and commercial consignee addresses within a 100x100 kilometre area in southwestern Ontario in Canada. Daily parcel size for each consignee follows the same gamma distribution of $\alpha = 4$, $\beta = 25$ (so mean = 100 units, standard deviation = 50 units) and is independent over time and independent of the package sizes sent to other consignees. The fleet of delivery trucks is homogeneous with a capacity of 2500 units per vehicle; i.e., each vehicle can serve an average of about 25 consignees (deliver 25 packages) on a single run.

The methodology consisted of simulating how the cost comparisons of interest are affected by two factors: (i) forecasting accuracy (number of correctly predicted arrival/non-arrival outcomes as a proportion of the number of sent parcels) and (ii) the customs clearance (arrival) rate of packages (average percentage of sent packages that arrive to the DC as scheduled). Two levels were used for each factor (0.90, 0.975 for forecast accuracy and 0.85, 0.95 for arrival rate) to yield a total of 4 factor combinations.

To convert each tactic's travel distances and requirements for ground delivery resources (drivers and vehicles) into monetary terms, the inputs for the experiment were cost estimates from the most recent published edition in the series of quadrennial studies by Transport Canada to determine the cost of commercial trucking operations. In that study, the province of Ontario's per kilometre cost for a typical package delivery truck (a two-axle van) was estimated at $2.88 when trucks operate at the highest utilization level. Though noted in that study as rarely attained, $2.88 was chosen over higher unit costs (of up to $3.52 for lower utilization levels) in order to present conservative estimates of the efficiency benefits of advanced package arrival information. Also, the fixed cost component (estimated at $60.32 per vehicle per day in Ontario) was used to determine the sunk cost of simply having excess (standby) vehicles to cope with uncertainty about package arrivals; i.e., on days when the vehicle is not dispatched a cost of $60.32 is still incurred. A sunk cost also applies to the other delivery resource held on standby: drivers. The related cost (driver wages) of $154/day was also taken from the said Transport Canada study.

The experiments, which were run over 250 simulated days, started with randomly generating package sending outcomes for each day using a Bernoulli process (p = 0.5 for the sending probability) to determine which of the 1000 consignee addresses packages are being sent to. For each simulated day, five vehicle routing solutions, corresponding to the four second-best tactics and the optimal tactic (route reoptimization covering today's actual arrivals and yesterday's surprise non-arrivals) were run. A trip distance of 200 kilometres was taken as a reasonable upper limit on a single driver's workload for a day. So in the simulation that constraint was applied to the four tactics as previously described; i.e., no single reoptimized route, pre-planned route, adjusted route (through address insertions and/or deletions) or recourse route was allowed to exceed 200 kilometres. So while a vehicle can be "double shifted" (dispatched on a pre-planned route and then later on a recourse route) a driver cannot be; i.e., a driver/vehicle dispatch for a recourse route requires an extra driver but not an extra vehicle. Roadshow®, an established vehicle routing program in commercial trucking operations, was used to construct each day's reoptimized, pre-planned, and recourse routes. Procedures already specified in describing the different second-best tactics were used for route adjustments.

At the end of the 250-day experiments, the daily differences of interest between the ideal of daily route reoptimization and the other tactics were collected and averaged, both in terms of physical units (travel distances driver/vehicle dispatches) and their monetary equivalents. The 250-day simulation was replicated 5 times, each replicate using a unique set of random number streams for the three stochastic factors: packages sent, sent packages that reach the DC without delay, and, based on the parameter value for forecast accuracy, the prediction (arrival or non-arrival) for each sent package. The combination of 5 replicates and a simulation run-length of 250 days was adequate to provide an average error margin of ±4.72% (with a 95% confidence level) in estimating the mean daily cost differences. That combination was derived from the established procedure of incrementally increasing the number of replicates and run length to determine the point at which the error margin stabilized. The variance reduction principle of common random numbers (CRN) –see, e.g., Law and Kelton (1991)– was used in order to minimize confounding the comparisons and the analysis of factor effects. Thus, for any given replicate, the same set of random numbers was used for all 4 factor combinations.

6. Findings

Table 1 shows the extent to which tactics based on lack of access to advanced parcel arrival information (Tactics 1 through 3 respectively) are costlier than the tactic based on having that access (daily route reoptimization). The comparison criterion is the mean daily excess transportation cost over the cost for daily route reoptimization. As indicated earlier, transportation cost covers the travel distance cost ($2.88/kilometre) and the sunk cost of having underused standby vehicles and drivers. The table depicts two expected patterns. First, more accurate forecasts reduce the excess cost of the second-best tactics (except for tactic 1, where forecasting accuracy is not an issue). Case in point is that for Tactic 2, which is most dependent on accurate forecasts, the increase in accuracy from f = 0.90 to f = 0.975 almost halves the daily excess cost: from an average $1787 to an average of $1028. The second pattern is that, Tactic 3, by being the most reactive to short-notice availability of parcel arrival information, consistently has the lowest (best) values of all the second-best tactics. If, as previously discussed, one can be certain that Tactic 3 is feasible under conditions of late access to that information, then the suggestion from the mean daily costs in Table 1 is that it should be the preferred second-best tactic. Still, its daily excess costs are not always insubstantial, so being able to switch from this tactic to the more advantageous position of having access to advanced package arrival information still yields material savings in the cost of ground delivery operations. For example, at the very high level of 97.5% for forecasting accuracy, with an average of 95% of the packages arriving on-time (unimpeded by customs clearance procedures), such access can eliminate excess cost of $112/day.

TABLE 1
Daily Excess Truck Routing Cost of Second-Best Tactics over Route Reoptimization

Forecast Accuracy	**Arrival Rate**	**Excess Cost for Tactic 1**	**Excess Cost for Tactic 2**	**Excess Cost for Tactic 3**
0.9	0.85	$700	$1787	$508
0.9	0.95	$223	$1755	$559
0.975	0.85	$700	$1030	$110
0.975	0.95	$223	$1028	$112

Compared to Tactic 3, the other two second-best tactics expend less effort to deal with delayed availability of package arrival information so there is more certainty about their feasibility. Therefore, for some couriers, the excess vehicle routing costs of these tactics vis-à-vis route reoptimization are probably better indicators of the benefits of advanced package arrival information for ground delivery operations (this caveat is implied in all of this paper's references to Tactic 3's performance. This caveat provides strengthens support for beliefs that the benefits can be considerable. For example, even under the most favourable factor combinations considered in this study (97.5% for forecast accuracy and 95% for the package clearance rate), if Tactic 3 cannot be deployed then the better of the other two second-best tactics (Tactic 1) incurs an excess transportation cost of $223 per day.

This particular result is instructive from another standpoint: the fact that Tactic 1 outperforms the other forecast-based tactic (Tactic 2) shows that as a next-best response to the problem of tardy package arrival information, simply providing a capacity buffer (Tactic 1) can sometimes be more effective than putting efforts into forecasting. In fact, Table 1 shows that for one of the four factor combinations considered here, Tactic 1 outperforms both forecast-based tactics (it outperforms Tactic 2 in all four cases). Clearly, the forecast-based tactics involve more non-routing resources than Tactic 1 (e.g., a forecasting system and the effort to forecast package arrivals as well as to analyze and effect route adjustment options) so even when they have lower transport costs than Tactic 1, inclusion of their non-routing costs could make them inferior to Tactic 1. For the three cases in which at least one forecast-based tactic outperforms Tactic 1 on excess transportation cost, Table 2 shows the gap in cost performance between that tactic and Tactic 1. The gap indicates the largest non-routing cost that the forecast-based tactic can incur without being inferior to Tactic 1.

Still, while the immediately preceding discussion indicates that using slack capacity on the pre-planned routes could be atop the pecking order of feasible second-best tactics, the central argument of this paper remains. That is, even the more desirable second-best tactics incur ground transportation costs that can run into hundreds of dollars per day, costs that are eliminable with advanced package arrival information. This finding of apparently material cost savings appears to have some real world corroboration: the projected annual savings of 100 million miles per year from the UPS package flow technologies (see footnote 2) translate to approximately 5 kilometres per vehicle per day for its fleet of 88,000 delivery trucks. This value is within the range of 1.96 to 7.57 kilometres per vehicle for the fleet required by Tactic 3 (see Table 3). These simulation results for Tactic 3 provide a reasonable basis for real world validation checks because that tactic is closest to what UPS has switched from in continuing to refine its package flow technologies. That is, the switch from a situation in which delayed upstream information cause drivers to spend valuable time making the last minute adjustment of rearranging packages in the order they feel would be the most optimum for delivery.

TABLE 2
Truck Routing Cost Comparison of Excess Capacity Tactic (Tactic 1) with Forecast-Based Tactics

Forecast Accuracy	**Arrival Rate**	**Tactic 1 Rank among next-best Tactics (Cost)**	**Tactic 3 Cost (superiority to Tactic 1)**
0.9	0.85	2nd Best ($700)	$508 (*$192)*
0.975	0.85	2nd Best ($700)	$110 (*$590*)
0.975	0.95	2nd Best ($223)	$112 (*$111*)

Note: Based on the paired *t*-tests, all comparisons were statistically significant at 0.010.

TABLE 3
Excess travel distance per vehicle in the fleet used by each tactic

Forecast Accuracy	Arrival Rate	Tactic 1	Tactic 2	Tactic 3
0.9	0.85	9.39	26.17	6.91
0.9	0.95	3.25	26.26	7.57
0.975	0.85	9.39	13.78	1.96
0.975	0.95	3.25	13.91	1.96

7. Conclusions

Using the context of ground delivery of packages that have an uncertain arrival status from overseas consignors, this paper assesses the value of advanced package arrival information as a way to eliminate unwelcome effects of the uncertainty. The value of eliminating those effects translates to the capability of switching from next-best tactics that compensate for absent advanced package arrival information to an optimal tactic: daily route reoptimization. This study finds that such a switch can yield substantial savings of hundreds of dollars per day in the costs of truck delivery of packages. For example, if the switch is from the second-best tactic of providing excess capacity on the delivery routes, then the savings range between $223 and $700 per day. These findings are consistent with expectations of international courier companies that their initiatives to assure real-time access to package arrival information (e.g., via sophisticated package tracking technologies and inter-organizational information system linkages with customs authorities) will enhance efficiency in their ground delivery operations. In addition to this conclusion, another noteworthy result from the study is that next-best options that involve forecasting package arrivals, even with highly accurate forecasts, are not always more effective than next-best options that merely provide excess capacity on pre-planned delivery routes (Tactic 1). Validation testing of the simulation findings suggests they are robust enough to be extrapolated to real-world truck delivery operations.

References

Baita, F.; Ukovich, W.; Pesenti, R.; Favaretto, D., 1998. Dynamic routing-and-inventory problems: A review. *Transportation Research-A*, 32 (8), 585-598.

Benton, William. C.; Rossetti, and Manuel D., 1992. The vehicle scheduling problem with intermittent demands. *Computers and Operations Research*, 19 (6), 521-531.

Bertsimas, D.J.; Chevri, P.; Peterson, M., 1995. Computational approaches to the stochastic vehicle routing problem. *Transportation Science*, 29 (4), 342-352.

Douglas, M., 2002. A strong sense of security. *Inbound Logistics*, retrieved online (http://www.inboundlogistics.com).

Gendreau, Michel; Laporte, , Gilbert; Séguin, René, 1996. Stochastic vehicle routing. *European Journal of Operational Research*, 88 (4), 3-12.

Golden, Bruce L.; Yee, J. R., 1979. A Framework for probabilistic routing. *AIIE Transactions*, 11 (2), 109-112.

Haughton, M., 1998. The performance of route modification and demand stabilization strategies in stochastic vehicle routing. *Transportation Research Part B: Methodological*, 32 (8), 551-566.

Haughton, M., 2000. Quantifying the benefits of route reptimization under stochastic customer demands. *Journal of the Operational Research Society*, 51 (3), 302-332.

Laporte, Gilbert; Louveaux, François V.; Van Hamme, Luc, 2002. An integer L-shaped algorithm for the capacitated vehicle routing problem with stochastic demands. *Operations Research*, 50, 415-423.

Law, Averil; Kelton, W. David: 1991. *Simulation Modeling & Analysis*. (2nd ed.). McGraw-Hill.

Transport Canada and Trimac Logistics Ltd. Consulting Services (2000), "Operating Costs of Trucks in Canada", available online at http://www.tc.gc.ca/pol/en/Report/OperatingCost2000, 2005. Last accessed December 3, 2005.

Waters, C. D. J., 1989. Vehicle scheduling problems with uncertainty and omitted customers. *Journal of the Operational Research Society*, 40 (5), 1099-1108.