Decision support for consumer direct grocery initiatives

Martin Savelsbergh, Georgia Institute of Technology, Atlanta, USA

Business-to-Consumer e-commerce has not only increased the demand for package delivery services, such as those offered by Federal Express and UPS, but has also led to the development of new home delivery service models and activities such as grocery delivery services. By bringing goods "the last mile" to the customer's front door, home delivery is seen by many as the ultimate value-added service for consumers. Many of the companies that have started to provide such services have found, however, that consumer direct delivery poses an enormous logistical challenge and are struggling to create direct delivery strategies that are profitable and can satisfy customer expectations. For example, the difficult challenge of matching grocery store prices while incurring high distribution costs led to closing of the San Francisco-based operations for Peapod and the declaration of bankruptcy by Webvan.

The fulfillment process for most consumer direct businesses can be divided into three phases: (1) order capture and promise, (2) order sourcing and assembly, and (3) order delivery. Our research effort focuses primarily on the interactions between order promise (deciding on a delivery time) and order delivery (devising efficient delivery schedules). Better integration of order promise and order delivery decisions has the potential to substantially improve profitability, especially for those consumer direct businesses offering "attended" deliveries. Attended deliveries are those where the consumers must be present and may be necessary for security reasons (e.g. expensive computer equipment) or because goods are perishable (e.g milk, flowers), and are a vital feature of many consumer direct service models. To avoid delivery failures as much as possible, it is customary in attended home delivery services for the company and customer to mutually agree on a narrow delivery window, or time slot. Thus, it is important to be able to quickly determine in which time slots each delivery request can be serviced given the already committed deliveries.

We are not only interested in how to effectively and efficiently evaluate the *feasibility* of each request, but also how to dynamically evaluate the *profitability* of each request in an effort to create improved decision technology to support home delivery services. Home delivery is a fairly new phenomenon, at least at this scale, and thus few models and algorithms have been proposed and studied that help create an understanding of the complexities and intricacies of these distribution problems. We propose and analyze new routing and scheduling problems that capture key features of home delivery, including the Home Delivery Problem (HDP), which allows the vendor to select which deliveries are accepted and the time slots for the accepted deliveries, and the Home Delivery Problem with Incentives (HDPI) which allows the vendor to dynamically compute the prices presented to a customer for the different time slots. Though the issue of whether to accept an order in hand or wait for more profitable delivery later is very similar to the core issues addressed in revenue management, the solution approach must be very different due to the cost structures involved. The costs associated with an airline seat do not depend on who occupies the seat nor on who purchases the rest of the tickets in the same fare class. In home delivery, though, the costs associated with accepting an order in a certain time slot depend on the delivery address and time slot of that order as well as on the delivery locations and time slots of the other accepted orders. We propose algorithms for both problems based on insertion techniques to help vendors make these decisions.

In the Home Delivery Problem (HDP) we have to construct a set of delivery routes for a specific day in the not too distant future. Requests from a known set of customers for a delivery on that particular day arrive in real-time and are considered up to a certain cut-off time, which precedes the actual execution of the planned delivery routes. We have to accept or reject each delivery request as it arrives. We assume that each request, if accepted, consumes d_i of the vehicle capacity and results in a revenue of r_i. There is a homogeneous set of m vehicles with capacity Q to serve the accepted orders. At each point in time t, we know for each customer i the probability $p_i(t)$ that customer i will place a request for delivery between t and the cut-off time. The requests are equally likely to occur at any point in time, so we have no information on the order in which requests may be realized. The objective is to maximize the total profit resulting from executing the set of delivery routes, i.e., total revenues minus total costs, where we assume that total cost depends linearly on the total distance traveled. To increase the level of service, we guarantee that the actual delivery will take place during a 1-hour time slot on the day of delivery. The 1-hour delivery time slots are non-overlapping and cover the entire day, e.g., 8.00 - 9.00, 9.00 - 10.00, ..., 19.00 - 20.00. We assume that for each delivery request we have a time slot profile that identifies which time slots will be acceptable to the associated customer. If we accept a request for delivery, we have to commit, at that time, to an acceptable time slot from the profile.

If we only determine whether a request for delivery can be feasibly handled, we ignore the fact that additional more profitable requests may still arrive, i.e., requests that have a higher revenue or, that can be incorporated more cost effectively because of their location. When a delivery request materializes, we have to decide whether to accept it or not, and, if we accept it, we have to decide which of the time slots in the customer's time slot profile to assign to it. Conceptually this can be done by computing for every possible realization of future delivery requests:

- the probability that this realization of future delivery requests occurs
- the optimal solution value for the resulting vehicle routing with time windows (VRPTW) solution *without* the request under consideration (but with all already accepted requests and all future realized requests)
- for each time slot in the time slot profile associated with the request under consideration, the value of the optimal VRPTW solution *with* the request served during the specified time slot (along with all of the already accepted requests and all future realized requests)

With this information we can compute the expected total profit if the request is rejected and the expected total profit if the request is accepted in each of the time slots in the customer's time slot profile. If the expected total profit without the request is larger than the expected total profit with the request for all time slots, then we reject the request. Otherwise, we accept the request and assign the time slot which resulted in the largest expected total profit.

Unfortunately, this approach is computationally intractable because the number of possible future realizations is too large and because the VRPTW is an NP-hard problem. To develop a more tractable approach, we can approximate the above conceptual solution through an insertion heuristic that solves a single instance of a modified vehicle routing problem with time windows (VRPTW) each time a request materializes. The objective is to maximize profit where we accept the fact that we may not be able to satisfy all requests due to limited capacity or time. If the request under consideration is part of the constructed set of delivery routes, it is more valuable to include this request rather than wait for future requests so it is accepted; if the request under consideration is not part of the constructed set of delivery routes, it is rejected.

The insertion heuristic consists of two phases. In the first phase, all accepted delivery requests are inserted using a simple insertion criterion (assuming we are inserting customer j between two consecutive customers $i-1$ and i in an existing route)

$$r_j - (c_{i-1,j} + c_{j,i} - c_{i-1,i}), \tag{1}$$

where, because we are trying to maximize profit, we always select the customer with the largest criterion value as the one to be inserted next. In the second phase, the remaining customers are inserted until there are no more feasible insertions due to limited capacity. The size of each delivery request in the second phase is adjusted downward by its probability of being received before the cut-off time, i.e., the size is set to $p_i(t)d_i$ for request i if we are currently at time t where $p_i(t) = 1$ for the request under consideration.

We have experimented with several different insertion criteria to evaluate the potential profitability associated with an insertion in the second phase, including:

- In computing the profitability associated with inserting a visit to customer j between $i-1$ and i, the estimate must reflect that the revenue from j will only be realized with a certain probability and the neighboring customers $i-1$ and i may or may not be present in the final set of delivery routes. It may be better to compute the cost of an insertion relative to two already accepted requests, say u and v. The expected length of the path between u and v, assuming route $(1, ..., u, ..., v, ..., n+1)$, can be computed as follows

$$\sum_{j=u}^{v-1} \sum_{k=u+1}^{v} d_{jk} p_j(t) p_k(t) \prod_{l=j+1}^{k-1} (1 - p_l(t)). \tag{2}$$

We now calculate the expected length with and without request j as part of the path between u and v, and we take the difference between these two values as the cost for inserting j. The expected revenue $p_j(t)r_j$ minus this cost yields the value of the insertion. This calculation is similar to what has been used to estimate the cost of a section of an *a priori* route in the stochastic vehicle routing problem and can be rather expensive computationally when vehicle capacity is large or individual probabilities are quite low.

- This next criterion can be viewed as an estimate and/or simplification of the preceding criterion. We consider both an optimistic and a pessimistic cost of the insertion, represented by values v_1 and v_2 respectively. The first value is based on the current set of delivery routes

$$v_1 = c_{i-1,j} + c_{j,i} - c_{i-1,i} \tag{3}$$

where we evaluate inserting customer j between two consecutive customers $i-1$ and i in a route. The second value is based on the set of delivery routes for accepted customers only, i.e., the set of routes at the end of Phase 1

$$v_2 = c_{k-1,j} + c_{j,k} - c_{k-1,k} \tag{4}$$

where we evaluate inserting customer j between two consecutive customers $k-1$ and k in a route. The value v_1 represents an optimistic view on the insertion cost, because it assumes that at least some of the as of yet unrealized delivery requests on the route will materialize, and v_2 represents a pessimistic view on the insertion cost, because it assumes that none of the as of yet unrealized delivery requests on the route will materialize. The value of the insertion for request j is set to

$$p_j(t)r_j - (v_2 - p_{i-1}(t)p_i(t)(v_2 - v_1)). \tag{5}$$

In the HDP, we do not consider the fact that customers might be willing to pay a premium to have their order delivered in their preferred time slot or that customers might be willing to accept a delivery in a different time slot if offered a discount. In the Home Delivery Problem with Incentives (HDPI), we assume that customer time slot preferences are modeled with a time slot price profile. This profile specifies the price the customer is willing to pay for a delivery in each particular time slot. Consequently, by setting the prices of the time slots appropriately, we can influence a customer to select a desired time slot.

It is natural to link incentives to the expected profit calculations from the previous section. We can evaluate the potential profitability associated with inserting the delivery in each time slot, as well as the profitability associated with not serving the customer. The difference between the expected profit with and without the request currently under consideration then serves as the marginal contribution to profit. A simple rule for setting incentives is to charge a premium equal to the absolute value of the marginal contribution to profit for all time slots where the marginal contribution to profit is negative, and to provide a discount equal to 10%, for example, of the marginal contribution to profit for all time slots where the marginal contribution to profit is positive. Even though our focus is on dynamic marginal contribution based incentive schemes, we also investigate incentive schemes based on historical customer behavior or on currently accepted deliveries only as they offer many advantages including computational efficiency.

We have conducted several experiments, via simulation, to determine the impact on total profit of using more sophisticated decision support techniques. We compare our algorithms for the HDP and HDPI with methods used in practice, as well as analyze the impact of instance characteristics on the performance of our proposed technology. We demonstrate, for example, the change in performance of different approaches over dense and sparse distributions of customers, different time window widths, and vehicle capacities. Results reveal substantial increases in profitability from improved technology, especially in certain classes of instances.

References

[1] A. Campbell and M.W.P. Savelsbergh (2002). Decision Support for Consumer Direct Grocery Initiatives. Report TLI-02-09, Georgia Institute of Technology.