

Planning Strategies for Home Health Care Delivery

Mike Hewitt, Nisha Nataraj

Department of Industrial and Systems Engineering, Rochester Institute of Technology, Rochester, NY 14534,
mrheie@rit.edu, nxn4881@rit.edu,

Maciek Nowak

Information Systems and Operations Management, Graduate School of Business, Loyola University, Chicago, IL,
60611, mnowak4@luc.edu

In home health care, continuity of care, wherein a patient is always visited by the same nurse, can be just as important as cost, as it is closely correlated to quality of care. While a patient typically receives care for two to three months, such that assigning a nurse to a patient impacts operations for lengthy periods of time, previous research focusing on continuity of care uses planning horizons that are often a week or shorter. This paper computationally demonstrates that considering a long planning horizon in this setting has significant potential for savings. Initially, a deterministic setting is considered, with all patient requests during the planning horizon known *a priori*, and the routing cost of planning for two to three months is compared with the cost when planning is done on a weekly basis. With inherent uncertainty in planning for such a long time horizon, a methodology is presented that anticipates future patient requests that are unknown at the time of planning. Computational evidence shows that its use is superior to planning on a weekly basis under uncertainty.

1. Introduction

Home health care refers to the care provided by skilled and non-skilled caregivers for patients, typically homebound, within their own homes. Services available for in-home care range from hospice care to wound and pain management to routine self-care activities such as feeding, bathing and cleaning. Patients are typically recommended for home health care by their primary physician or after discharge from a hospital and the episode of care often

lasts 60-90 days, during which time most individuals need to be visited multiple times each week by their caregivers.

In-home care is an attractive alternative to nursing homes and hospitals because it typically lowers costs and improves a patient's quality of life. For example, Balinsky (1999) reports that the cost of caring for "high-tech, ventilator dependent children ... at home is 87 percent less than a hospital setting." With all the associated benefits, it is not surprising that the use of home health care has become widespread. A study conducted by the National Association for Home Care and Hospice (NAHC, 2010) found that in 2008 approximately 33,000 care-givers made visits to approximately 12 million patients across the country. However, providing care in the home on that scale does involve significant transportation resources; the same study found that home health care workers drove close to 5 billion miles in the United States in 2006.

Given the numerous budget cuts that healthcare providers have faced (NYAHS and HCA, 2011), many are focused on reducing costs, particularly those related to transportation. However, cost reductions often negatively impact the primary goal of most providers: maintaining a high level of quality care. One important quality metric used by providers is continuity of care, wherein the goal is to see patients with the same caregiver on each visit. Continuity of care has many benefits, including reducing the time and effort a caregiver spends reviewing a patient's history and care plan, and fostering close relationships between patients and their caregivers. However, maintaining a high level of continuity of care also reduces planning flexibility and opportunities for savings. When patients are assigned to the same caregiver for two to three months, there are few opportunities for re-optimization when new patients require care.

Models and solution methods have been presented for the planning problem faced by home health care agencies. However, the length of the planning horizon to use has not been studied. While the combination of long periods of care and continuity of care suggests that long planning horizons should be considered when applying these solution methods, the existing literature rarely looks ahead more than one week. Table 1 summarizes the length of the planning horizon used in the representative literature that considers continuity of care in some capacity. Note that all papers use a week or less, with some extending the

horizon on a rolling basis. Groër et al. (2009) run a condensed test extending the one week horizon of their algorithm using four weeks of aggregated data to plan for a fifth week, finding a degradation in their results when compared to week-by-week planning. In this paper, the planning horizon is extended to 2-3 months and the advantages of considering a longer time frame when planning are quantified.

Paper	Author	Year	Continuity of Care	Length of Horizon
An integrated spatial DSS for scheduling and routing Home-Health-Care nurses	Begur et al.	1997	No	Week
A home health care routing and scheduling problem	Cheng and Rich	1998	No	Single Day
A hybrid setup for a hybrid scenario: Combining heuristics for the home health care problem	Bertels and Fahle	2006	No	Single Day
LAPS CARE - An operational system for staff planning of home care	Eveborn et al.	2006	Weighted Objective	Advanced planning followed by daily planning
A hybrid approach to solve the periodic home health care problem	Steeg and Schröder	2008	Weighted Objective	Week
The consistent vehicle routing problem*	Groer et al.	2009	Constraint	Week/Five weeks
Home health care problem: An extended multiple traveling salesman problem	Kergosien et al.	2009	Constraint modeling preference	Single Day
Dynamic periodic fixed appointment scheduling for home health	Bennett and Erera	2011	Maintained through single nurse scheduling	Rolling horizon over a year
Workforce management in periodic delivery operations*	Smilowitz et al.	2011	Weighted Objective	Week

Table 1 Planning horizons considered in literature – starred papers are in a package delivery setting

To do so, an existing planning model and method are adapted to compare the transportation cost and staffing requirements of plans with differing planning horizons. It is initially shown that longer planning horizons enable significant reductions in transportation cost and staffing requirements in a deterministic setting where all current and future patient service requests are known at the time of planning. The planning problem is modeled as a variant of the Consistent Vehicle Routing Problem (ConVRP) presented in Groër

et al. (2009) and solved with a modified version of their ConVRP Record to Record travel algorithm (ConRTR). The modification to the ConRTR results in improved performance with problem instances that vary in the number of visits required by each patient over the problem horizon. The modified ConRTR is used to implement two planning strategies: a rolling, weekly strategy that is similar to what is often proposed in the literature, and a long term planning strategy where planning is done once for the entire horizon.

The potential for savings is next demonstrated in a stochastic setting where there is uncertainty with respect to when and where future patients will request care. To do so, another adaptation of the ConRTR is presented, based on an easily calculated point estimate of future patient requests. This revised algorithm anticipates these future requests when assigning known patients to nurses. Given that many agencies have rudimentary information systems and rely on manual planning processes, requirements for the method include that it rely on data that agencies already have and that it is easily explained, in hopes that doing so will increase the adoption rate of the planning method. Some automated decision support systems have already been developed for home health care agencies. The system developed by Begur et al. (1997) integrates a nurse routing component into a Spatial Decision Support System that uses a Geographical Information System to aid schedule-making. LAPS CARE, a system that does not constrain continuity but includes it as an objective, is developed and shown to be beneficial for Swedish home care organizations by Everbom et al. (2006). However, neither of these systems considers a long term planning horizon.

This paper makes three contributions to the literature on vehicle routing and home health care logistics. First, in the context of vehicle routing, the ConRTR is enhanced, leading to significant improvement in its performance when there is considerable variation in the number of visits each customer requires. Second, in the context of home health care logistics, significant savings are found when assigning nurses to patients with planning horizons that are longer than what is typically seen in literature and in industry. Third, in the context of both home health care logistics and vehicle routing, the first solution approach for the ConVRP with stochastic customers is presented. While the methodology developed here is motivated by home health care, the ConVRP has been applied to other

problems such as small package delivery and grocery delivery and this methodology is pertinent for those dynamic environments as well. Ultimately, the results in this paper show that planning for a longer period of time can lead to considerable gains in operational efficiency.

The rest of the paper is organized as follows. Section 2 formally defines the problem and related literature, the application of the ConRTR to the home health care setting and the enhancement of the ConRTR. Section 3 describes two planning strategies that, in a deterministic setting, are central to establishing the relationships between continuity of care, planning horizon, and cost. Section 4 presents a planning method for a stochastic version of the problem where not all patient requests for visits are known. Finally, the conclusions and scope for future work are discussed in Section 5.

2. Problem Description and ConRTR Enhancement

The planning problem faced by a home health care agency is often modeled as a variant of the Vehicle Routing Problem (VRP) (Begur et al. 1997, Cheng and Rich 1998, Bertels and Fahle 2006, Steeg and Schröder 2008, Kergosien et al. 2009). While the majority of existing work on the planning problem faced by home health care agencies view the problem as a static planning problem, Bennett and Erera (2011) study a dynamic home health care routing and scheduling problem for a single-nurse variant where some patient requests are not known in advance.

This paper presents a model of the planning problem faced by a home health care agency using the ConVRP (Groër et al. 2009). The ConVRP aims to design routes that ensure a customer is visited by the same driver/vehicle each time he or she requires a visit. The ConVRP is similar to the Periodic VRP (Beltrami and Bodin 1974, Russell and Igo 1979), an extension of the VRP where vehicle routes are constructed to serve customers over a multi-period planning horizon and visit schedules are chosen from a set of possible visit frequencies. However, the ConVRP differs in that visit schedules are assumed to be determined *a priori*, as is the case in our problem. Other papers (Francis et al. 2006, 2007, Zhong et al. 2007, Smilowitz et al. 2012) have studied consistency or driver familiarity (which is a variation of consistency) in the small package industry, but consistency is primarily modeled through an objective rather than a constraint.

This problem is modeled on a graph where agency and patient locations are represented by nodes and routes connecting those nodes are represented by arcs. The nurses begin from a central agency before visiting patients in this problem, which is modeled using a single depot. Patients require a certain number of visits per week (typically two or three) over a period of care that is generally 60-90 days. The days that a patient requires a visit are fixed *a priori* as agencies determine this in advance based on the type of care prescribed. Once a visit schedule is established, it typically stays fixed for the remainder of the patient's period of care. An example schedule is for a patient to be seen twice a week, every Monday and Thursday, for eight weeks. For the purpose of this analysis, patients are not differentiated based on condition, or nurses based on skill/qualifications. Also, the number of nurses is flexible as agencies often have the option of outsourcing patient visits to third-party providers during periods of high demand.

To solve the ConVRP, Groër et al. (2009) propose the ConVRP Record to Record travel algorithm (ConRTR) and show that it is capable of producing high-quality solutions for instances of the ConVRP with up to several thousand customers very quickly while enforcing perfect consistency of service. It ensures this consistency by designing a set of template routes which visit all customers with multiple requests over the planning horizon. The template routes are not implemented in practice. Instead, they constitute a framework from which daily routes are derived. Figure 1 provides an example of one such template route and the daily routes developed from it. The ConRTR is an iterative algorithm that alternates between deriving and improving the template routes based on artificial bounds on vehicle capacity and route length, then deriving and improving daily routes based on the set of template routes. The daily routes are used to improve the artificial bounds on the template routes in an iterative loop.

While the ConRTR can produce high quality solutions for instances that consider a planning horizon of multiple months, when deriving and improving template routes it does not differentiate customers by the number of visits they require. Whether a patient is to be visited twice or five times over the course of the time period, the template route visits each once, with the only distinction based on the service time required for each patient. With a two or three month planning horizon, the difference in the number of visits can become

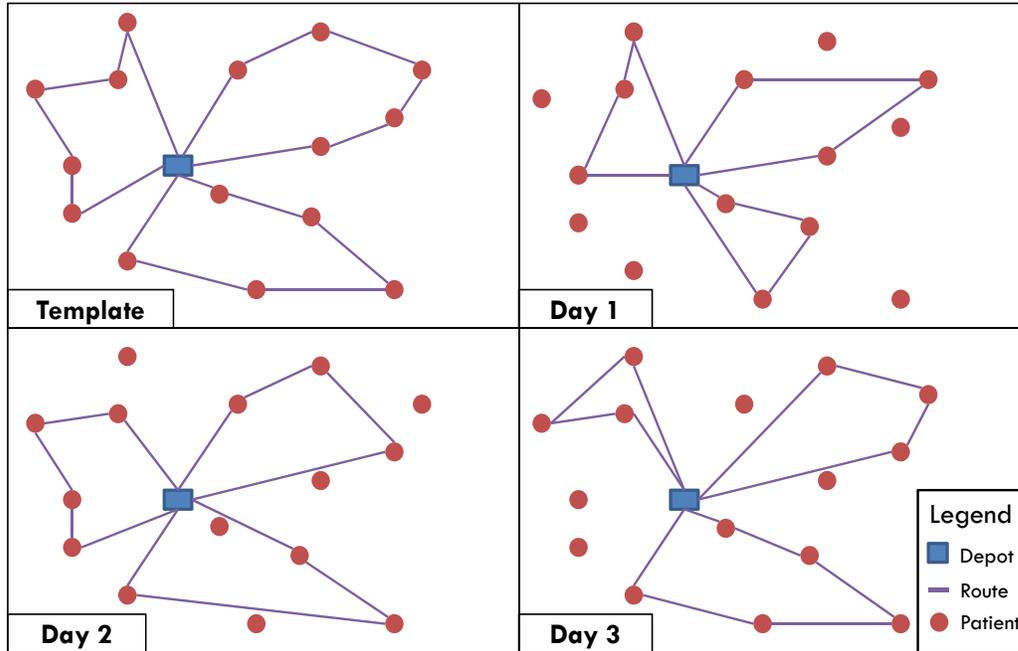


Figure 1 Deriving daily routes from a template route

quite pronounced (e.g., one patient has two visits per week while another has three, or one patient requires service for three months while another three weeks). Without considering this information, some efficiencies may be lost. For example, if two patients living in the same vicinity each require one month of service over a two month horizon, one for the first month and the other for the second, it may be most efficient to service them with the same nurse. However, the ConRTR may not place them on the same route as it cannot interpret that each patient only requires service for half of the problem horizon.

In order to distinguish patients based on the number of visits required, the service time per visit for each patient is discounted based on the number of days the patient must be visited over the problem horizon. Given a required service time, s_i , for patient i , with n_i visits over the length of the problem horizon H , the discounted service time, s_i^* , is expressed as $s_i^* = \frac{n_i}{H}(s_i)$. For example, if a patient requires three visits per week, each lasting 60 minutes, for six weeks of an eight week problem horizon (with five days per week), the discounted visit time used for designing the template routes would be $\frac{3*6}{5*8} * 60 = 27$ minutes.

In the extreme case where a patient requires service every day over the course of the problem horizon, the visit time would not be discounted. Therefore, when the template routes are constructed, those patients requesting service most frequently will have the least discounted service time. Returning to the example with two patients each requesting one month of service over a two month horizon, each patient's service time is reduced by 50% and it is more likely that the ConRTR may be able to fit both on the same template route.

The effectiveness of this adjustment is evaluated by testing the ConRTR with and without the discounted service times. The problem instances used for this evaluation are those described in Section 3.1. It is found that the use of discounted service times enables the ConRTR to produce solutions that, on average, require 4.34 fewer nurses and 70.82 fewer hours of travel. It is also noted that the use of discounted service times reduces the average run time of the ConRTR by more than 50%, from 126.2 seconds to 56.7. Discounted service times are used for testing throughout the remainder of the paper.

3. Home health care planning strategies in a deterministic setting

In this section two planning strategies are presented, week-by-week and long-term, under a deterministic setting wherein all patient visit requests are known *a priori*. The ConRTR algorithm is implemented for both strategies to evaluate the impact that planning horizon lengths have on routing costs with consistent continuity of care. The computational analysis is initially conducted under a deterministic setting to motivate the development of a method that can be applied to a stochastic setting and to generate bounds that can be used to evaluate the performance of that method.

Week-by-week planning strategies have been used in previous research (Begur et al. 1997, Steeg and Schröder 2008). In this strategy, only those patients that require service for the coming week are considered when assigning patients to nurses, with assignments and routes made in previous weeks carried forward. When a new patient requires service, he or she is added to a route from a previous week. In order to implement this strategy, the ConRTR is initially applied to the first week of service. For each of the following weeks, the template routes from the previous week are loaded and all new patients for the current

week are inserted in a greedy manner that also yields feasible daily routes. This is done without changing the assignments from previous weeks such that consistency is maintained for all patients.

For the long-term planning strategy, the ConRTR is applied once at the beginning of the horizon, taking into consideration all patients who are to be seen over the problem horizon. Figure 2 illustrates the difference between the two strategies. In this small example, consider a two week planning horizon with two nurses, A and B, that may service at most four patients per week. Six patients require service, with patients 2, 3, 4 and 6 to be serviced both weeks, while patients 1 and 5 require care in week 2. With week-by-week planning, in the first week nurse A may be used to service all patients in an effort to limit fixed costs associated with using an additional nurse, while minimizing travel costs. However, when this assignment is implemented in the second week, nurse B is required to handle the new demand and both nurses must make the longer trip to the 3, 4, 5, and 6 cluster of patients. With long-term planning, nurse B would start serving patient 2 in the first week to maintain consistency and save this nurse from making the longer trip to the more distant patient cluster. While this results in both nurses used for both weeks, it shows that long-term planning can quickly lead to savings in travel cost.

3.1 Experimental setup

Problems with long-term horizons for the ConVRP have not been studied, so the instances used for testing in this paper were randomly generated. The procedure for generating these instances was similar to that used for the ConVRP instances. As planned visits rarely occur on the weekend, a five-day week and a four-week month are used. Each instance models an “urban” setting, wherein patients are randomly distributed in a circle with a radius of five miles, or a “rural” setting, wherein patients are randomly distributed in a circle with a radius of 15 miles. A nurse travels 30 miles per hour in an urban setting and 40 miles per hour in a rural setting. Each nurse’s workday is 10 hours. Based on conversations with a home health care agency, each patient visit requires one hour. A set of 200 initial patients is created with locations and length of care randomly generated. While the number of patients requesting service in each week is not constant, to model

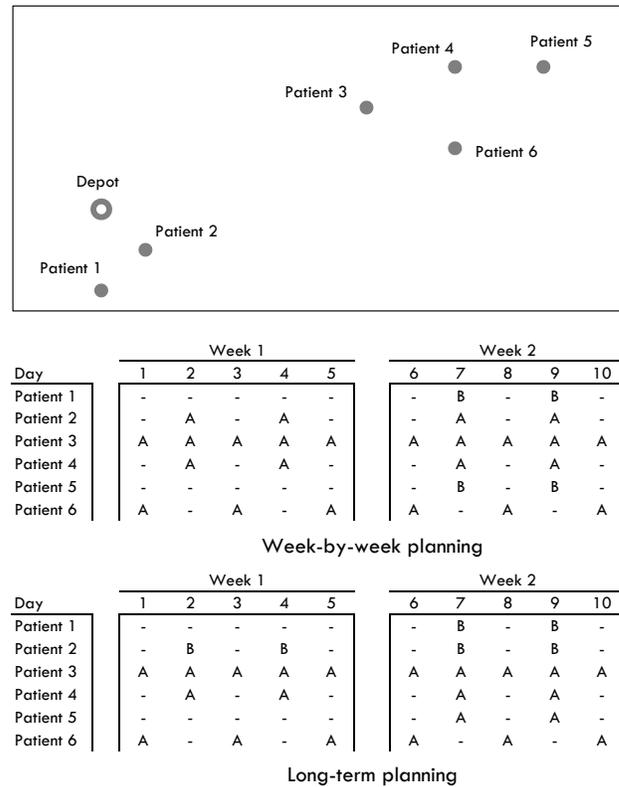


Figure 2 Nurse assignments for week-by-week and long-term planning

an agency in a “steady-state” of demand, the number of patients whose period of care begins after week one is set equal to the number of patients whose period of care ends over the entire planning horizon. The length of care for each patient and the set of patients beginning service after the first week are randomly generated such that this steady-state is maintained. The weekly demand schedule is generated assuming a 70% chance that a patient will require a visit on each day.

The instances are summarized in Table 2, where column “Begin/End” provides the total number of patients whose period of care begins or ends after week one (i.e., for a value of 35, 35 new patients begin service over the problem horizon and 35 end service). The algorithms are coded in C++ and executed on a machine with a 1.80 GHz AMD P820 triple-core processor with 4 GB of RAM.

Instance	Begin/End	Geographic setting	Horizon
35U8	35	Urban	8 weeks
70U8	70	Urban	8 weeks
35R8	35	Rural	8 weeks
70R8	70	Rural	8 weeks
55U12	55	Urban	12 weeks
110U12	110	Urban	12 weeks
55R12	55	Rural	12 weeks
110R12	110	Rural	12 weeks

Table 2 Problem instance characteristics

3.2 Deterministic computational results

Table 3 presents a comparison of the two planning strategies with respect to the total time nurses spend in transit during the planning horizon and the average number of patients a nurse visits on each day. Long-term planning leads to a significant decrease in travel time (equivalent to approximately six 50-hour work weeks) and enables each nurse to visit almost one more patient per day. Note that the savings in travel time are larger for the instances that model a rural setting. This is to be expected given the greater distances between patients.

Instance	Travel Time (Hours)			Average # Patients Served per Day per Nurse	
	WbW	LT	Savings	WbW	LT
35U8	668.6	556.6	112	5.3	6.6
70U8	848.2	693.4	154.8	4.8	6.1
35R8	1,367.10	1,091.85	275.25	4.4	5.1
70R8	1,880.55	1,451.85	428.7	4.1	4.7
55U12	1,115.40	953.6	161.8	4.9	5.7
110U12	1,693.40	1,442.40	251	4.4	4.9
55R12	2,384.40	2,055.15	329.25	4.4	5
110R12	2,955.45	2,262.60	692.85	4	4.7
Average	1614.14	1313.43	300.71	4.54	5.35

Table 3 Comparison of travel time and patients served under week-by-week (WbW) and long-term (LT) planning

Figure 3 examines the utilization of the nurses, measured as the amount of time spent with patients relative to the total time working (time with patients plus travel time). This figure shows the utilization for one instance; all other instances exhibit similar behavior. While the fraction of time spent with patients is high for both strategies, it steadily decreases over the problem horizon when planning is done on a week-by-week basis. With this strategy, as more patients end their period of care, it becomes difficult to find new patients to efficiently fill the empty slots. Considering all patients simultaneously allows for the development of routes that account for these fluctuating requests.

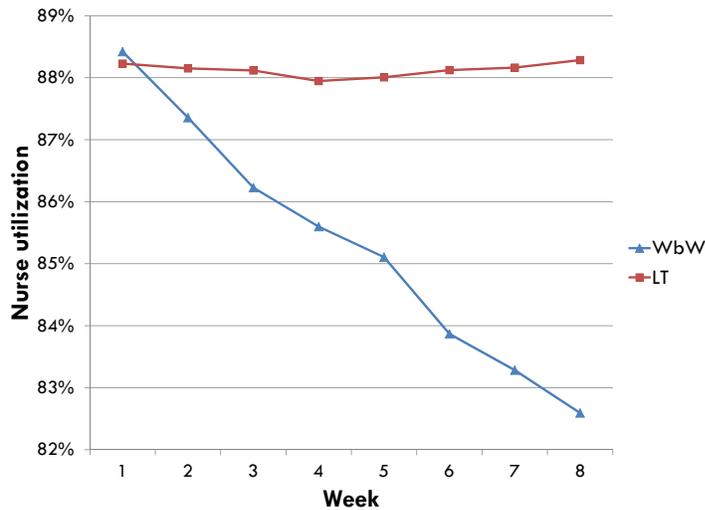


Figure 3 Utilization of nurses under week-by-week (WbW) and long-term (LT) planning

It is evident that planning over a longer period can reduce the variable costs associated with transportation. Determining the impact on the fixed costs associated with the number of nurses staffed is also of interest. Table 7 reports the average and standard deviation of the number of nurses used per week over the entire horizon. By employing long-term planning, seven fewer nurses are needed on a weekly basis. Also, the variation in the number of nurses needed each week is significantly lower, which should help with staffing issues.

A month-by-month planning strategy was also tested, as a hybrid between the week-by-week and long-term strategies. Long-term planning still lead to a savings of over 180 hours in travel time on average, and the use of five fewer nurses on an average weekly basis.

Instance	WbW		LT	
	Average	Standard deviation	Average	Standard deviation
35U8	31	2.8	24.5	0.4
70U8	39.6	6.7	30.3	0.2
35R8	36.8	2.8	31.2	0.2
70R8	46.6	6.5	40.1	0.2
55U12	36.1	3.8	30.6	0.1
110U12	50.4	9.1	44.4	0.1
55R12	41.1	4.7	36.1	0.2
110R12	54.8	8.8	45.2	0.3
Average	42.05	5.65	35.3	0.21

Table 4 Comparison of nurses employed under week-by-week (WbW) and long-term (LT) planning

4. Long Term Planning Under Uncertainty

While the deterministic experiments show that there are advantages to considering a long planning horizon, these cost reductions may be difficult to achieve without perfect information. A stochastic model is necessary to take into account the variable behavior that is likely in any planning scenario, particularly as the planning horizon is extended. Uncertainty may be modeled using a variety of techniques, but nearly all require the determination of probability distributions. For home health agencies that do not employ sophisticated information systems, such methods may be difficult to implement. Thus, a straightforward methodology for the stochastic ConVRP is presented here, relying on basic data that agencies typically already have.

Many applications of the VRP involve uncertain parameters, such as customer demands, customer locations, and/or travel times. Most relevant to this work is the VRP with stochastic customers (VRPSC) (Gendreau et al. 1996), where demands are known, but customer locations and, if the problem is periodic, days to be seen, are not known beforehand. Other related work includes Bent and Van Hentenryck (2004), which proposes a Multiple-Scenario Approach to solve the partially dynamic Vehicle Routing Problem with Time Windows with Stochastic Customers, and Campbell and Savelsbergh (2005), which defines the Home-Delivery Problem (HDP) for grocery delivery with stochastic customers and presents different profitability-based insertion heuristics to create a set of routes that will help evaluate whether to accept or reject incoming customer requests.

Using historical data, an agency can generally derive an aggregated expectation of the average number of patients to be visited in a given week. This information may be used to anticipate requests and adjust routing accordingly. For this model, this expectation is used to determine how many new patients request service each week. By anticipating these future requests and planning for them accordingly, a more efficient set of template routes may be created, leading to improved daily routes in future weeks.

Prior to creating the first set of template routes with the ConRTR, the expected number of patients visited per week is used to create a set of “dummy” patients that serve as a placeholder for potential future requests. With these dummy patients added to the initial known set of patients, the template routes are constructed for the entire problem horizon. The template routes are used to create the daily routes for the first week. Then, on a weekly, rolling-horizon basis, daily routes are derived at the beginning of each week based on the template routes and information that is learned as new, actual patients request service. It is assumed that at the beginning of each week all patients whose care period begins that week are known.

In order to create a dummy patient, several parameters regarding every future patient must be estimated: (1) the week his or her period of care begins, (2) his or her location, and, (3) his or her weekly visit schedule. To evaluate the first parameter, for each week the number of patients still requesting service based on the initial schedule is compared with the expected number of patients for an average week. As home care is typically prescribed for a fixed period of time, it is assumed that the agency knows when each initial patient’s period of care will end. When there are fewer known patients than expected it is assumed that new patients begin their period of care in this week and these patients are added to the set of requests as dummy patients. The care period for a dummy patient continues through to the end of the problem horizon. For the second and third parameters to be estimated, the location is fixed at the depot and the patient is assigned a visit every day of the week. As these dummy patients are placeholders, any arbitrary location should suffice and the depot is often close to the geographic center of all requests.

This methodology is illustrated in Figure 4. Assume a four-week planning horizon, with 14 known patients in week 1, and an expectation of 13 patients each week. Also, it is

known that in week 2, the period of care ends for one patient, decreasing the number of known patients in week 2 to 13. Thus, the difference between the number of known patients in week 2 and the number of expected patients is zero. No dummy patients are added. Week 3 sees two additional patients end service, with two fewer real patients than expected patients. Therefore, two dummy patients are added at the depot. Week 4 follows with three more patients ending service and three dummy patients added.

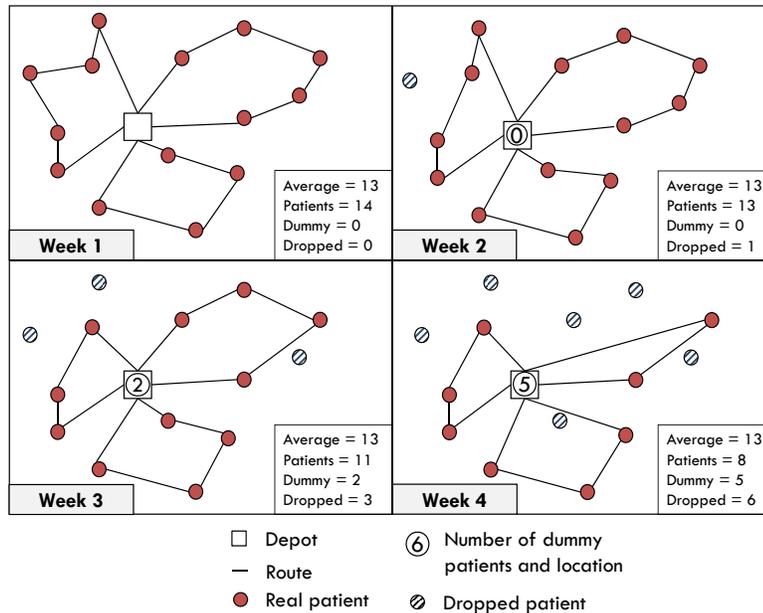


Figure 4 Modeling expectation of future patients with dummy patients

After the dummy patients are generated and added to the schedule of known patients requiring service, the ConRTR is applied over the entire planning horizon to derive template routes. These template routes are then used to create the daily routes for week 1. At the beginning of each of the following weeks, the actual requests for service are realized. If there are slots reserved for dummy patients on template routes, a dummy patient is selected for replacement in a greedy fashion such that the real patient is placed on the geographically closest route of a dummy patient and inserted into the route at the point where travel cost is minimized. If there are more new real patients than dummy patients

in a week, the additional new patients are placed on routes following the procedure for week-by-week planning. These new template routes are then used to create the set of daily routes for the week, continuing in this fashion until the end of the problem horizon.

4.1 Stochastic computational results

The instances described in Table 2 were again used for the computational experiments. The initial set of experiments were performed with the expected number of weekly patients equal to the average number of patients that are to be seen in a week over the planning horizon. The sensitivity of the algorithm to this expected value is tested later. Table 5 compares the performance of long-term planning under uncertainty (LTUU) with the deterministic solutions found using week-by-week planning (WbW) and long-term planning (LT). By anticipating future service requests with an estimate of the expected number of patients to be seen in a week, LTUU was able to provide significant savings in travel time over WbW. However, with all patient requests known at the time of planning, LT outperforms LTUU. It is evident that improving the forecasting of future patient requests can lead to considerable cost savings.

Instance	Travel Time (Hours)			LTUU % change from	
	WbW	LT	LTUU	WbW	LT
35U8	669	557	611	8.6%	-9.8%
70U8	848	693	800	5.7%	-15.3%
35R8	1367	1092	1190	13.0%	-9.0%
70R8	1881	1452	1602	14.8%	-10.3%
55U12	1115	954	1051	5.8%	-10.2%
110U12	1693	1442	1577	6.9%	-9.4%
55R12	2384	2055	2166	9.2%	-5.4%
110R12	2955	2263	2695	8.8%	-19.1%
Average	1614	1313	1461	9.5%	-11.3%

Table 5 Percentage change with long-term planning under uncertainty (LTUU) from week-by-week (WbW) and long-term (LT) planning

In Table 6, the performance of LTUU is measured using nurse staffing requirements. LTUU uses fewer nurses on an average weekly basis than both WbW and LT. When

performing LT, the template routes visit all patients, including those not visited until much later in the planning horizon. As a result, there may be template routes that contain patients whose period of care begins early in the planning horizon and patients whose period of care begins much later. Such a template route may yield under-utilized daily routes in the early weeks of the planning horizon, and, multiple such daily routes may be combined into a single route. However, with LTUU, the initial set of template routes only visit the set of known patients and the dummy patients. Specifically, when patients begin care in one of the last few weeks of a planning period, and these patients are in excess of the expected number of patients, they are not considered when deriving template routes during the initial planning phase of LTUU. Instead, those patients are inserted into the nearest daily routes the week that their period of care begins. Similarly, note that with LT each nurse visited 5.35 patients per day (on average), while with LTUU each nurse visited 5.71 patients per day.

Instance	WbW		LT		LTUU	
	Average	Standard deviation	Average	Standard deviation	Average	Standard deviation
35U8	31.0	2.8	24.5	0.4	26.5	1.6
70U8	39.6	6.7	30.3	0.2	33.0	2.8
35R8	36.8	2.8	31.2	0.2	28.9	3.2
70R8	46.6	6.5	40.1	0.2	34.3	5.1
55U12	36.1	3.8	30.6	0.1	28.2	1.8
110U12	50.4	9.1	44.4	0.1	40.1	3.5
55R12	41.1	4.7	36.1	0.2	32.0	4.2
110R12	54.8	8.8	45.2	0.3	40.5	5.7
Average	42.1	5.7	35.3	0.2	32.9	3.5

Table 6 Number of nurses used with week-by-week (WbW), long-term (LT) and long-term planning under uncertainty (LTUU) planning

The majority of dummy patients are placed on the same template route, which is not surprising given that they are all located at the depot. Thus, LTUU prescribes plans with some nurses primarily tasked with visiting the patients that are not initially known. This suggests that a planning policy where nurses are held in reserve in anticipation of future

patients requiring care could be effective. However, such a policy may conflict with work balancing objectives.

The sensitivity of LTUU to the expected number of patients to visit each week is also tested. For each instance, LTUU is executed with the expected number of weekly patients both above and below the actual average number of weekly patients and the results for instance 200I-70-R-8H are reported in Table 7 (other instances have similar results). In the first column the value of the expected number of weekly patients is compared to the actual average number of patients seen in a week for that instance. Not surprisingly, LTUU performs best when the two values are equal, but even when the expected value is skewed from the average, LTUU is superior to WbW, suggesting that an exact forecast of the average number of weekly patients is not necessary for the method to yield savings.

Expected number of patients in a week (E)	Travel Time (Hours)		% decrease
	WBW	LTUU	
Average -10		1,625.00	13.6%
Average -5		1,633.16	13.2%
Average	1,880.60	1,601.70	14.8%
Average +5		1,660.68	11.7%
Average +10		1,631.28	13.3%

Table 7 Percentage change in cost between LTUU and WbW with adjusted expected number of patients

5. Conclusions and Future Work

Many home health care agencies can benefit from planning methods that enable them to reduce costs while still providing high levels of continuity of care, a metric often associated with quality care. In this paper we first illustrate that significant savings in travel time and nurse-staffing requirements can be realized by planning for a period of time that is much longer than what is typically seen in the vehicle routing literature. In the course of doing so we present an enhancement of an existing heuristic for the Consistent VRP that enables it to find better solutions in less time when applied to problems with variability in

the amount of service required by each patient. And, because planning for a long period of time introduces considerable uncertainty into the problem, we present a long term planning strategy that anticipates future service requests based on an easily calculated point estimate of those requests. We show computationally that this strategy is superior to planning on a rolling horizon, week by week basis.

There are multiple avenues for extending the research presented in this paper. Within the context of modeling home health care operational realities more closely, we can modify the planning methods to take into account skill and licensing requirements when assigning a patient to a nurse and support nurses departing from multiple depots. Within the context of enhancing the performance of the method that anticipates future patient requests we intend to study other ways of modeling the location of a future patient. Lastly, the problem itself could be viewed as a dynamic or online scheduling and routing problem, which is how many agencies make scheduling decisions now.

References

- Balinsky, W. 1999. Pediatric home care: reimbursement and cost benefit analysis. *Journal of Pediatric Health Care* **13**(6) 288–294.
- Begur, S.V., D.M. Miller, J.R. Weaver. 1997. An integrated spatial DSS for scheduling and routing home-health-care nurses. *Interfaces* **27**(4) 35–48.
- Beltrami, E.J., L.D. Bodin. 1974. Networks and vehicle routing for municipal waste collection. *Networks* **4**(1) 65–94.
- Bennett, A. R., A. L. Erera. 2011. Dynamic periodic fixed appointment scheduling for home health. *IIE Transactions on Healthcare Systems Engineering* **1**(1) 6–19.
- Bent, R.W., P. Van Hentenryck. 2004. Scenario-based planning for partially dynamic vehicle routing with stochastic customers. *Operations Research* 977–987.
- Bertels, S., T. Fahle. 2006. A hybrid setup for a hybrid scenario: combining heuristics for the home health care problem. *Computers & Operations Research* **33**(10) 2866 – 2890.
- Campbell, A.M., M.W.P. Savelsbergh. 2005. Decision support for consumer direct grocery initiatives. *Transportation Science* **39**(3) 313–327.

- Cheng, E., J.L. Rich. 1998. A home health care routing and scheduling problem. Rice University, Texas, Tech. Rep. TR98-04.
- Eveborn, P., P. Flisberg, M. Rönnqvist. 2006. LAPS CARE- an operational system for staff planning of home care. *European Journal of Operational Research* **171**(3) 962 – 976.
- Francis, P., K. Smilowitz, M. Tzur. 2006. The period vehicle routing problem with service choice. *Transportation Science* **40** 439–454.
- Francis, P., K. Smilowitz, M. Tzur. 2007. Flexibility and complexity in periodic distribution problems. *Naval Research Logistics* **54** 136 – 150.
- Gendreau, M., G. Laporte, R. Séguin. 1996. Stochastic vehicle routing. *European Journal of Operational Research* **88**(1) 3–12.
- Groër, C., B. Golden, E. Wasil. 2009. The consistent vehicle routing problem. *Manufacturing & Service Operations Management* **11**(4) 630–643.
- Home Care Association of New York State and New York Association of Homes & Services for the Aging. 2011. State budget cuts have imperiled New York’s home health care delivery system. Tech. rep.
- Kergosien, Y., C. Lenté, J.C. Billaut. 2009. Home health care problem: An extended multiple traveling salesman problem. *4th Multidisciplinary International Conference on Scheduling: Theory and Applications (MISTA’09), Dublin (Ireland)*. 10–12.
- National Association for Home Care and Hospice. 2010. Basic statistics about home care. Tech. rep.
- Russell, R., W. Igo. 1979. An assignment routing problem. *Networks* **9**(1) 1–17.
- Smilowitz, K., M. A. Nowak, T. Jiang. 2012. Workforce management in periodic delivery operations. Forthcoming in *Transportation Science*.
- Steeg, J., M. Schröder. 2008. A hybrid approach to solve the periodic home health care problem. *Operations Research Proceedings 2007* 297–302.
- Zhong, H., R.W. Hall, M. Dessouky. 2007. Territory planning and vehicle dispatching with driver learning. *Transportation Science* **41**(1) 74–89.