7-19-2012

A Method for Designing Efficient Routes for Home Healthcare Agencies Considering Uncertain Future Demand

Nisha Nataraj

Rochester Institute of Technology

A METHOD FOR DESIGNING EFFICIENT ROUTES FOR HOME HEALTHCARE AGENCIES CONSIDERING UNCERTAIN FUTURE DEMAND

A Thesis

Submitted in partial fulfillment of the

requirements for the degree of

Master of Science in Industrial Engineering

in the

Department of Industrial & Systems Engineering

Kate Gleason College of Engineering

by

Nisha Nataraj

July 19, 2012

DEPARTMENT OF INDUSTRIAL AND SYSTEMS ENGINEERING

KATE GLEASON COLLEGE OF ENGINEERING

ROCHESTER INSTITUTE OF TECHNOLOGY

ROCHESTER, NEW YORK

CERTIFICATE OF APPROVAL

M.S. DEGREE THESIS

The M.S. Degree Thesis of Nisha Nataraj

has been examined and approved by the

thesis committee as satisfactory for the

thesis requirement for the

Master of Science degree

Approved by:

____

Dr. Mike R. Hewitt, Thesis Advisor

____

Dr. Michael E. Kuhl

DEDICATION

To my parents, for their ever-patient support and for believing in me when I didn’t. To my grandparents, for showing me that hard work always has its rewards, even when the fruit is not immediate. And to my grandaunt, for giving me an education.

ACKNOWLEDGMENT

This thesis, and all the lessons learned along the way, would not have been possible without my advisor, Dr. Mike Hewitt. I am especially grateful for his constant guidance, encouragement and faith. Understanding and interpreting data is a difficult skill, but his ever-present curiosity and keen interest has been a valuable lesson, something I will treasure for the rest of my graduate life. I know that all the hours spent reviewing my work could not have been easy, but his patient attention to detail has taught me that we can always do better – the end result is always so much more satisfying. Thursday afternoons will forever remind me that research can be fun.

I would like to thank Dr. Michael Kuhl for taking the time and effort to be on my thesis committee. His feedback, always prompt, has taught me the value of perspective and to never lose sight of the bigger picture. I am indebted to both Dr. Hewitt and Dr. Kuhl for encouraging me to continue on with my graduate studies and giving me the advice needed to do so.

I owe much to the Industrial and Systems Engineering department. I am obliged to both the people and the resources needed for me to work on my research and successfully complete my thesis.

Every graduate student has that task in front of them that seems at one point insurmountable; their Everest. I will be forever grateful to my sister-in-law, Namitha, and close friend, Suraj, for spending all those hours patiently answering my programming questions, but more so, for instilling in me the belief that I could learn to code. The CS tutoring lab was a great environment to work in, and I am thankful to the tutors for their help.

I can never completely thank my family, my parents and my brother, Nakul, enough for their unconditional love and support. Last, but far from the least, I would like to express my deepest gratitude to Tejas. His patient love and understanding has given me the strength to face every little failure through every single day.

ABSTRACT

Faced with increasing demand for home health care that is rising faster than the supply of resources needed to deliver it and drastic budget cuts, many home care agencies are struggling to remain operational. There is a need for efficient routing that doesn't compromise on the quality of care achieved when a patient is visited by the same nurse over the entire period of care, also known as care-giver continuity or continuity of care. Because care periods often last more than 60 days, care-giver continuity causes scheduling decisions to have a long-term impact by potentially restricting the agency from making alternative assignments that could reduce routing costs. Our research aims to understand and quantify the benefit of utilizing time horizons of 60-90 days when making routing decisions under the constraint of continuity of care. We do so by defining the Home Health Care Routing Problem (HHCRP) as a variant of the Vehicle Routing Problem (VRP), known as the Consistent VRP, that includes the continuity of care requirement. Unlike related literature on this problem which considers planning horizons of at most a week, computational experiments in a variety of settings suggest the importance of considering planning horizons of 2-3 months when developing schedules for care-givers. Given that it is almost impossible to have complete information about future patients that far ahead, we also present a method that enables planners to design schedules for care-givers in the face of such uncertainty and demonstrate its effectiveness computationally.

TABLE OF CONTENTS

LIST OF TABLES AND FIGURES

1. INTRODUCTION

A study conducted by the National Association for Home Care and Hospice (NAHC, 2006) showed that in 2006, home health care workers drove close to 5 billion miles in the United States. As of 2008, there were approximately 958,000 home health employees across the country (NAHC, 2010). Home Care refers to the care provided by skilled and non-skilled care-givers for patients, typically homebound, within their own homes. There are a wide range of services available for in-home care, ranging from hospice care to wound and pain management. Patients are typically recommended home health care by their primary physician or after discharge from a hospital. The episode of care often lasts 60-90 days, during which time most individuals need to be visited multiple times each week by their care-givers.

In-home care is an attractive alternative to nursing homes and hospitals because of lower costs and improved quality of life benefits. For example, the average in-home management cost for ventilator-dependent adults on a daily basis was only $235 in the year 1992, whereas the associated daily cost when hospitalized was $719 (Bach, Intintola, Alba, & Holland, 1992). Lower costs can also enable patients to be seen for longer periods of time in turn resulting in the better management of conditions. Some quality of life benefits are that patients can live in the comfort of their own homes and have the option of assistance with routine self-care activities such as feeding, bathing and cleaning. Home care has also been found to reduce the likelihood of re-admission, not only improving quality of life, but saving money as well. One such example is a study on psychiatric patients in Connecticut that determined that 46% of those that were hospitalized without first participating in the in-home crisis intervention program were subsequently re-admitted after their discharge, compared to a readmission rate of only 12% for those who participated in the program (Pigott & Trott, 1993).

Home Health Care (HHC) is a growing industry. In 2008, around 33,000 care-givers made visits to approximately 12 million patients across the country (NAHC, 2010). In spite of the large demand for home care, agencies have been struggling to stay afloat. According to a report released by the Home Care Association (HCA) of New York and the New York Association for Homes and Services for the Aging (NYAHSA), since 2008 , Medicaid home services has lost over $430 million in budget cuts, (NYAHSA & HCA, 2011). To make matters worse, more than 70% of agencies were reported to be operating at a loss in 2008. These large scale operations suggest that gasoline costs can have a significant impact on the economics of home health care. With worsening operating losses, high fuel costs and bleak financial scenarios, it is critical for agencies to reduce transportation costs associated with making visits to patients' homes.

However, the need for efficient routing goes beyond agencies' current financial state of affairs. Efficient routes not only reduce costs but also allow care givers to spend more time at patients' homes and reduce non-value-added travel time. Furthermore, agencies can accordingly plan for anticipated staffing requirements ahead of time. A study on the benefits of the implementation of a Decision Support System for scheduling in home care organizations in Sweden found reduced short-term sick leave amongst nurses and a reduction in the number of missed visits (Eveborn et al., 2009). In addition, organization managers reported better staffing and budgetary control.

Other challenges that agencies currently face is that schedulers typically use relatively unsophisticated planning systems to design routes for nurses. As a result, planning is done for a single patient at a time, on-demand, and as we discovered during interviews at a large local home health agency, schedulers often end up relying on manual processes to schedule nurses.

Therefore, at the end of the day, it is clear that efficient routes have a far reaching impact that transcends monetary benefits.

An important quality of care metric in home health is continuity of care. With continuity of care, patients are always seen by the same care-giver, resulting in saved time and effort associated with familiarizing new care-givers with patients' histories and care plans, as well as fostering close relationships and trust between patients and their care givers. Continuity also has other benefits. In the meta-analysis by Cabana & Jee (2004) in a general health-care setting, the authors found that continuity of care over time resulted in reduced hospitalizations, improved patient satisfaction and improved receipt of preventive services such as mammograms and MMR vaccinations. One challenge that maintaining high levels of continuity of care presents to agencies is that once a care-giver (henceforth referred to as a nurse) is assigned to a patient, that assignment should remain fixed for the remainder of the care period. With periods of care typically lasting 2-3 months, this means that the scheduling decisions made by an agency have a long-term impact on their daily costs. However, agencies rarely consider this long-term impact when making scheduling decisions and research related to the scheduling problem faced by home health care agencies rarely considers planning horizons longer than one week.

Based on interviews with schedulers at a large local home health care agency, Figure 1 shows the typical process of scheduling patients at a large local home health care agency. Referrals from hospitals and other sources are built into the database of the scheduling system. Each patient is then assessed during an initial visit to determine the kind and frequency of care that is required. Once the evaluation is complete, the internal staff then filter the visit requests according to discipline, depending on whether the care-provider is a nurse, home health aide,

therapist or some other skilled/unskilled care-provider. Schedules and assignments are created regularly (ensuring that the number of visits for each patient is at least as good as the compliance standards) and updated throughout the week as changes are made. As mentioned earlier, a patient often needs to be seen multiple times a week. The frequency and days on which a patient needs to be seen are usually pre-determined, and stay fixed for remainder of the care period (for example, a patient may need to be seen three days a week, on Mondays, Wednesdays and Fridays for 60 days). In spite of the availability of scheduling software, we found that the planners often choose to design schedules manually. While this is a cumbersome process in itself, the added difficulty of keeping track of continuity of care requirements makes it even harder. The nurse then makes the visit, and the planners schedule the next visit. The cycle repeats until the end of the episode of care. Highlighted in Figure 1, is the stage of the scheduling process that can be automated with an algorithm that produces daily routes for nurses. On occasion, changes may need to be made after a schedule has already been created. For example, a nurse might call in sick or a patient might request a change in schedule. However, we consider any post-processing required to the schedule to be outside the scope of this research.

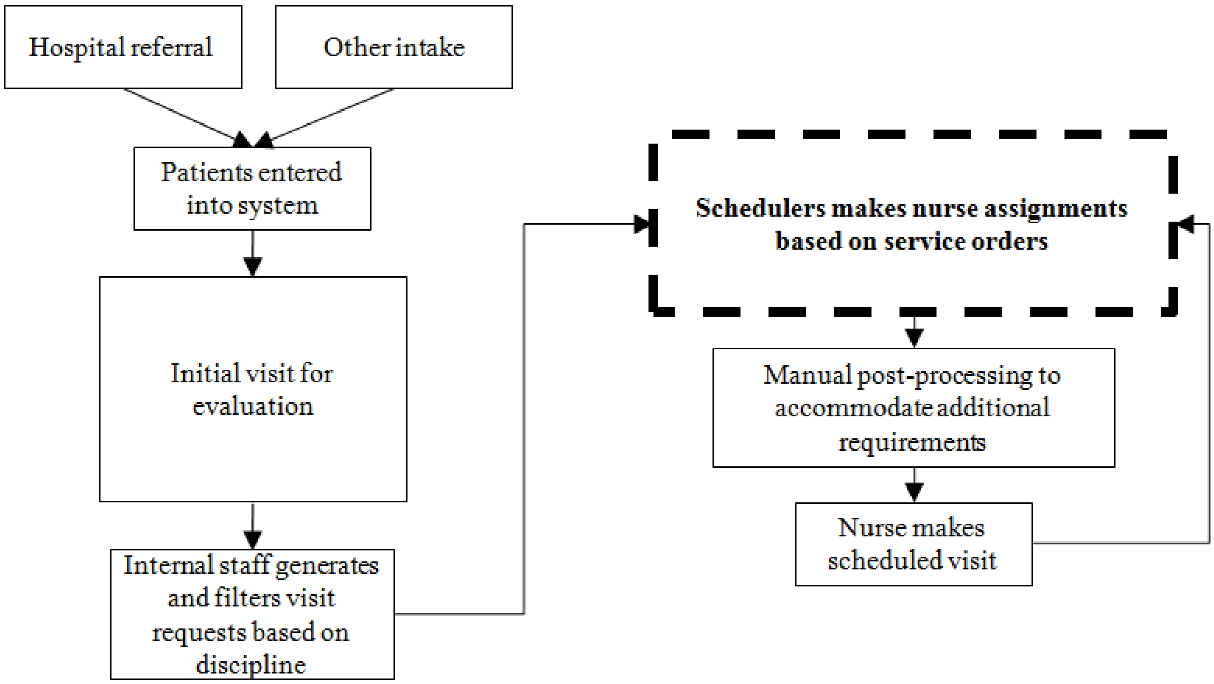

Figure 1 - Process of scheduling a visit for a new patient at a typical agency

In this research we computationally demonstrate the benefits of considering long planning horizons (2-3 months) when maintaining high levels of continuity of care. We do so by modeling the scheduling problem faced by home health care agencies as a type of Vehicle Routing Problem (VRP), called the Home Health Care Routing Problem (HHCRP). Our first experiments focus on a deterministic setting, where all patients to be seen during the planning horizon are known. However, with planning horizons of this length, such an assumption is clearly unrealistic. Thus, we next present a method that, when executed in a rolling-horizon manner, anticipates future patients requiring care while assigning known patients to nurses. Because many agencies are small and have unsophisticated information systems, we designed this method to have low data requirements to capture uncertainty.

The rest of the thesis is organized as follows: We first formally define the problem and provide a mathematical description of the same. Next, we review the relevant literature

discussing the different VRP variants relevant to the HHCRP and previous work in home health care routing. A background of the Consistent Record to Record algorithm designed by Groer, Golden, & Wasil (2008) is provided, followed by a description of its adaptation to the home health care setting. We then describe our experimental design and describe two planning strategies that are central to establishing the impact of long-term continuity. Following this, we present an enhancement of the Consistent Record to Record algorithm for when long planning horizons are considered and a variant of the method that recognizes uncertainty with respect to future patients needing care. Finally, we discuss our conclusions and ideas for future work.

2. PROBLEM STATEMENT

In this section, we discuss the objectives of this research and then formally define the HHCRP followed by a Mixed Integer Program (MIP) formulation of the problem to provide a precise mathematical definition of the objective and constraints of the problem.

2.1. Objectives

The goal of this research is to determine the importance of considering long planning horizons when developing schedules for nurses in a home health care setting where continuity of care is a critical quality metric. At present, most agencies plan "on-demand", typically considering each patient as they come, on a case-by-case basis. Furthermore, previous literature in the home health care and vehicle routing settings rarely consider planning horizons that are longer than a week. This research emphasizes the benefits achieved from planning horizons that are 2-3 months long.

The specific objectives for this research are:

1. To understand the routing costs and staffing requirements of planning over weekly basis when compared with a long-term planning horizon over the duration of 2-3 months using a routing method that is adapted to the home health care setting;
2. To develop a method that can handle fluctuation in the number of patients over a planning horizon of 2-3 months and demonstrate its capabilities in scenarios with growing demand and steady demand;
3. To develop a method for long-term planning which anticipates future patients and has low data requirements to model uncertainty; and
4. To computationally study the capabilities and limitations of the above methods.

The results of this research will enable home health care agencies to design more efficient routes and anticipate staffing needs ahead of time. Efficient routes will also allow nurses to have more time to spend with their patients. The anticipated impact of this research can also apply to other routing environments, such as package delivery, where consistency is a desired metric. By quantifying the benefits achieved from long-term planning, we emphasize the underlying need for routing methods that are designed for longer planning horizons.

2.2. Problem Description

The HHCRP is defined on a graph on which the agency and patient homes are represented as nodes, and routes connecting these nodes are represented as arcs. Nurses set out from either their homes or the agency at the beginning of the day and visit those patients who need to be seen for that day. Patients require multiple visits over a planning horizon of 60-90 days. The days that a patient requires visits and the frequency of these visits are assumed to be fixed beforehand. A realistic assumption based on our conversation with the agency is that the frequency and days of visit are determined in advance based on the type of care needed, so that it is mutually acceptable to both the patient and the agency. Once a visit schedule and frequency are decided, it stays fixed for the remainder of the patient's episode of care. For example, a patient may need to be seen twice a week, every Monday and Thursday for the next 8 weeks.

We make some assumptions about the problem. First, we do not differentiate patients based on condition, or nurses based on skill/qualifications. Many agencies work with home health aides who assist with routine self-care and hospice care, and agencies typically do not differentiate aides based on skill/qualifications. Second, the number of nurses is not a fixed input to the problem. Additional nurses are introduced according to necessity, typically because an

economical routing decision cannot be made without exceeding the maximum number of hours that existing nurses may work on a given day. In a realistic setting, agencies often have the option of outsourcing to third-party providers during periods of high demand.

2.3. Mathematical Model

The sets used in the formulation of the MIP and the subsequent algorithm are defined as follows – *N* is a set of nurses that need to visit the set of all patients, *P* over a horizon, *D*. *P`* is a subset of the set *P*, and includes all patients that need visits, but excludes the agency/depot. In order for a route to be feasible, the total time that a nurse works on a given day, *d*, must be at most *T*.

Let c_{ij} represent the cost associated with traveling from node *i* to node *j*. The parameter w_{id} is considered to be 1 if a patient *i* requires a visit on day *d*. Let a_{id} be the arrival time at patient *i* on day *d*. The parameter s_{id} is the service or visit time for each patient *i*. We define the decision variable x_{ijnd} to be 1 if a nurse *n* travels from node *i* to node *j* on day *d* and 0 otherwise. The decision variable y_{ind} is considered to be 1 if a nurse *n* visits a patient *i* on day *d* and 0 otherwise. Finally, the notation *o* represents the depot or agency from which a nurse may start his/her day.

$$Minimize\ \sum_{d \epsilon D} \sum_{n \epsilon N} \sum_{i \epsilon P} \sum_{j \epsilon P} c_{ij} x_{ijnd} \qquad (1)$$

s.t.

$$\sum_{j \epsilon P} x_{ojd} = 1 \qquad \forall d \in D, \forall n \in N \qquad (2)$$

$$\sum_{i \epsilon P} x_{iod} = 1 \qquad \forall d \in D, \forall n \in N \qquad (3)$$

$$\sum_{n \epsilon N} y_{ind} = w_{id} \qquad \forall i \in P', \forall d \in D \qquad (4)$$

$$\sum_{i \epsilon P} x_{ijnd} = y_{ind} \qquad \forall j \in P, \forall n \in N, \forall d \in D \qquad (5a)$$

$$\sum_{i \epsilon P} x_{jind} = y_{ind} \qquad \forall j \in P, \forall n \in N, \forall d \in D \qquad (5b)$$

$$w_{id\alpha} + w_{id\beta} - 2 \le y_{ind\alpha} - y_{ind\beta} \le -(w_{id\alpha} + w_{id\beta} - 2) \quad (6)$$

$$\forall\ days\ d_\alpha and\ d_\beta, \alpha \ne \beta, \forall i \in P, \forall n \in N$$

$$a_{id} + x_{ijnd}(s_{id} + t_{ij}) - (1 - x_{ijnd})T \le a_{jd} \quad (7)$$

$$\forall i \in P, \forall j \in P', \forall n \in N, \forall d \in D$$

$$a_{id} + x_{ijnd}(s_{id} + t_{ij}) + (1 - x_{ijnd})T \ge a_{jd} \quad (8)$$

$$\forall i \in P, \forall j \in P', \forall n \in N, \forall d \in D$$

$$\sum_{i\epsilon P, i\epsilon P} x_{ijnd} \le |S| - 1 \quad (9)$$

$$for\ S\ \subseteq P, with\ 2 \le |S| \le P\ , \forall n \in N, \forall d \in D$$

$$0 \le\ a_{id} + w_{id}(s_{id} + t_{io}) \le Tw_{id} \qquad \forall i \in P' \quad (10)$$

$$x_{ijnd} \in \{0,1\};\ y_{ijnd} \in \{0,1\}; a_{ind} \ge 0; \quad (11)$$

$$\forall i \in P, \forall j \in P, \forall n \in N, \forall i \in D$$

The objective (1) of the HHCRP is to minimize the total cost (or equivalently, the total distance) of all routes for all nurses over all days in the horizon. Constraints (2) and (3) make sure that a nurse starts and ends his/her day from the depot. Constraint (4) ensures that nurses visit patients on each day that they require a visit and do so at most once. Each patient has just a single successor and predecessor as ensured by (5a, 5b). Constraint (6) ensures that the same nurse visits a patient when multiple visits are required. Constraints (7) and (8) determine the arrival time at a patient's home. Sub-tours are eliminated with constraint (9). Finally, (10) imposes restrictions on the maximum time that a nurse may work on a given day.

While we formulate the problem as an MIP, we cannot solve it as such. This is because typically MIPs are used for smaller problems with fewer nodes, say, around 100 patients. However, we're looking at a mid-sized agency with around 200-400 patients and also

significantly increasing the size of the problem by considering a planning horizon of 2-3 months. Because current solvers cannot solve these problems in realistic run-times, we instead choose to use heuristics for the HHCRP.

3. RELATED WORK

In this section we first present an overview of the Vehicle Routing Problem and its variants that are relevant to the HHCRP, followed by a review of some work that has been carried out in the home care routing setting.

3.1. The Vehicle Routing Problem

The Vehicle Routing Problem was first introduced in 1959 by Dantzig and Ramser as the Truck Dispatching Problem (Dantzig & Ramser, 1959). The objective was to develop the shortest possible set of routes for a fleet of gasoline trucks to deliver gasoline from a major terminal to several smaller stations. It has since spawned a class of problems encompassing different applications ranging from emergency services to home delivery, while dealing with a variety of constraints such as capacity, time windows, dynamic requests, etc.

Many practical instances of the VRP involve problem sizes that consider up to several thousand customers (Kytojoki, Nuortio, Braysy, & Gendreau, 2007; Li, Golden, & Wasil, 2004) and cannot be solved to optimality in realistic runtimes. Baldacci, Mingozzi, & Roberti (2011) provide a review of recent exact algorithms that have been developed for the Capacitated VRP (CVRP) and VRP with Time Windows (VRPTW) variants. Currently, exact algorithms are capable of solving instances with at most a few hundred customers in a single day planning horizon whereas home care agencies often see a few thousand patients a week. For larger problems, the focus has been on developing classical heuristics and meta-heuristics that solve the problem to near-optimality in shorter times. Local search heuristics tend to have a comparatively restricted exploration of the search space; Meta-heuristics rectify this, and are able to explore the

search space better by allowing worsening solutions. A review of these is available in Laporte, Gendreau, Potvin, & Semet (2000).

3.2. Variants of the VRP

The HHCRP is often modeled as a variant of the VRP (Cheng & Rich, 1998; Steeg & Schröder, 2008). In the HHCRP, each nurse's schedule for the day is identified by a route and patients by their locations. Both problems also include a service or visit time associated with each visit. Additionally, there is a restriction on maximum length of day (that includes both travel and service times) for drivers and nurses. Several variants of the VRP focus on issues that are important in the context of home health such as having multiple depots (for nurses to have the option of starting their day from their own homes), periodicity (because patients often require multiple visits), and consistency (which translates to continuity of care in the home health setting).

In the HHCRP, patients need to be visited periodically throughout their episode of care. For example, a patient might need physical therapy three times a week for two months. The Periodic VRP (PVRP) considers a planning horizon of several days, such as a week. Customers often need to be visited more than once during the planning horizon. The PVRP has been previously applied to grocery distribution, waste collection, etc., and a review of these applications is available in Francis, Smilowitz, & Tzur (2008). A Variable Neighborhood Search (VNS) is used in Hemmelmayr, Doerner, & Hartl (2009) and Pirkwieser & Raidl (2008) to solve the PVRP. VNS, first presented in Mladenovic & Hansen (1997) is a Meta-heuristic that explores increasingly distant neighborhoods of the current solution using search operators and moves to a new solution only if an improvement is made.

The Consistent VRP (ConVRP), first described by Gröer et al. (2008) aims to design routes that ensure that a customer is visited by the same driver/vehicle each time he or she requires a visit. The ConVRP is inherently periodic because it considers a time frame during which a customer has to be visited multiple times by the same driver. While the majority of their analysis is based on a horizon of only a single week, they also conducted some tests on a real-world instance spanning 5-weeks, using 4 weeks of template routes to derive routes for the fifth week. From a VRP modeling standpoint, driver consistency is equivalent to continuity of care within a home health care setting. The authors were one of the first to focus exclusively on consistency and did so in the small package shipping industry. Other papers (Smilowitz, Nowak, & Jiang, 2011; Zhong, Hall, & Dessouky, 2007) have previously looked at driver familiarity, which is a variation of consistency. With driver familiarity, if a driver repeatedly visits a particular region, he/she will know the routes better and as a result, delivery performance improves. This work can be applied in the small package delivery industry where it is important for drivers to know their way around a region well in order to find addresses quickly.

Real-life applications of the VRP often have several sources of uncertainty. Taking them into consideration makes methods more likely to produce solutions that are robust with respect to different realizations of uncertain parameters. When considering extremely long planning horizons in the VRP or home health care setting, it is unlikely that every patient who needs a visit in future weeks is known ahead of time. Incorporating this uncertainty into the planning process may yield better routes.

There are typically two methods for incorporating stochastic information into a VRP – one option is to model it as a Chance Constrained Program (CCP), and the other is to model it as a Stochastic Program with Recourse (SPR). In the first case, routes are designed such that there

is a minimum expected level of success. However, this method does not consider the cost of recourse action (such as returning to the depot to restock in case of exceeded capacity or skipping absent customers) if the routes fail. SPRs, on the other hand, try to minimize the expected cost of the routes as well as the cost of recourse actions. The review by Gendreau, Laporte, & Seguin (1996a) provides a good overview of the stochastic VRP and its variants.

The different sources of uncertainty typically considered in an SVRP include demand, customers, and travel times. In VRPs with stochastic demands, customer locations are deterministic; however, demands are unknown and are realized only once the vehicle arrives at the location. On the other hand, with stochastic customers, demands are deterministic, but customer locations and arrival times are not known before hand. Flatberg, Hasle, Kloster, Nilssen, & Riise (2005) review the existing literature in the stochastic VRP, including the SVRP with travel times. Some papers also consider the combined effect of stochastic demand and customers (Gendreau, Laporte, & Seguin, 1995; Gendreau, Laporte, & Seguin, 1996b)**.** In the first paper, the authors use an integer-L shaped algorithm to solve the VRP with Stochastic Demands and Customers (VRPSCD) exactly, observing that stochastic customers make the problem harder. In the second paper, the authors develop the TABUSTOCH algorithm to solve the VRPSCD and compare it against the optimal solution obtained from the integer-L shaped algorithm.

From the point of view of the HHCRP, because patients are not associated with any sort of "demand", the largest source of uncertainty is customers. With the VRPSC, typically, the set of customers who can potentially make a request is known ahead of time. However, some of these customers may not need to be visited, because of which the recourse action involves just skipping absent customers. Some of the earliest work on the VRPSC is available in the review by

Gendreau, Laporte, Seguin (1996). In Bent & Van Hentenryck (2004), the authors propose a Multiple-Scenario Approach to solve the partially dynamic VRPTW with Stochastic Customers. The paper exploits stochastic information about dynamic customers by creating a projected plan with predicted customers in order to improve the quality of the solution. A consensus function then chooses the most similar plan from multiple plans that contain both static (known) and dynamic requests.

Campbell & Savelsbergh (2005) define the Home-Delivery Problem (HDP) for grocery delivery with stochastic customers. The goal is to be able to visit as many customers as possible while maximizing the total profit from these visits. The authors compare different profitability-based insertion heuristics to create a set of routes that will help evaluate whether to accept or reject incoming customer requests. There has not been much work done as yet in stochastic and dynamic home care routing. However, Bennett & Erera (2011) study a dynamic home health care routing and scheduling problem for a single-nurse variant where all patient requests are not known in advance. The authors consider a rolling horizon of a year and handle requests for new patients as they arrive. They achieve this by developing distance and capacity-based insertion heuristics with an objective of maximizing the number of patients seen. This is one of the only few papers available in the literature to consider a planning horizon that far out in a home health setting.

3.3. Home health care routing

In comparison to the literature available on the VRP, there has been relatively little work done in the HHCRP. Kergosien, Lenté, & Billaut (2009) model the home health care problem as a multiple-Traveling Salesman Problem (m-TSP) with a slight variant of continuity, in that some services may need to be performed by a specific nurse and such preferences are modeled as a

constraint. The problem is formulated as an Integer Linear Program (ILP) and solves instances with up to 40 locations to optimality. Cheng & Rich (1998) also formulate the HHCRP as an MIP. A nurse routing component is integrated into a Spatial Decision Support System in Begur, Miller, & Weaver (1997) along with a Geographical Information System (GIS) to aid schedule-making. The authors use the sequential Clark & Wright savings algorithm to assign routes to nurses and a Nearest Neighbor heuristic to improve these individual routes. Two papers, (Bertels & Fahle, 2006; Steeg & Schröder, 2008) use a hybrid of Constraint Programming (CP) and a meta-heuristic to solve the HHCRP. The latter considers the idea of continuity of care by trying to minimize the number of different nurses that visit a single patient. In the paper by Eveborn, Flisberg, & Rönnqvist (2006) the authors build a decision support system, LAPS CARE, for Swedish home care organizations. The problem is formulated as a set partitioning model and solved using a repeated matching heuristic. They incorporate the importance of continuous care by using a weighted objective function. The LAPS CARE system is designed to create plans a few days in advance and handle any last minute changes as necessary. The impact and numerous benefits of LAPS CARE on the Swedish home care industry is well documented in Eveborn et al. (2009).

While models (and solution methods) have been presented for the HHCRP, it is not yet known how they should be used. In particular, the combination of long periods of care and continuity of care suggests that long planning horizons should be considered. We summarize in Table 1 the length of the planning horizon used in previous VRP and HHCRP literature that considers continuity of care in some capacity. We note that few (if any) approaches consider planning horizons longer than a week.

Paper	Author	Year	Continuity of Care	Length of Horizon
An Integrated Spatial DSS for Scheduling and Routing Home-Health-Care Nurses	Begur, Miller and Weaver	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: Combing 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, Flisberg and Rönnqvist	2006	Weighted Objective	Few days; Advance planning is followed by last minute changes
A Hybrid Approach to Solve the Periodic Home Health Care Problem	Steeg and Schröder	2008	Weighted Objective	Week
Home Health Care Problem: An Extended Multiple Traveling Salesman Problem	Kergosien, Lenté and Billaut	2009	Constraint	Single Day
The Consistent Vehicle Routing Problem (Package delivery setting)	Groer, Golden and Wasil	2009	Constraint	Week
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 (Package delivery setting)	Smilowitz, Nowak and Jiang	2011	Weighted Objective	Week

Table 1 - Length of planning horizon considered in the literature

4. BUILDING CONSISTENT ROUTES: THE CONRTR IN A HOME HEALTH SETTING

Our analysis of routing costs over long planning horizons uses the Consistent Record-to-Record (ConRTR) Algorithm, which is capable of solving large VRP instances, with up to several thousand customers (Gröer et al., 2008). The ConRTR can solve these large problems very quickly while inherently maintaining continuity of care. The algorithm is designed to minimize the total distance traveled over all the routes while treating consistency (continuity of care in our context) as a constraint.

4.1. Background

To ensure that customers are always visited by the same driver, the ConRTR creates a set of *template routes* for all customers with multiple requests over the planning horizon. The template routes are not actually traversed by vehicles; rather, they work as a framework from which each daily route may then be derived. By deriving each day's routes from the templates, it is ensured that each customer is always visited by the same driver. As mentioned earlier, in the VRP and the home health care setting, a common restriction is the maximum allowable duration of a driver or nurse's day. Because the templates consider customers over the entire period and are never implemented themselves, it is hard to determine what the bounds on the duration of a driver's day should be to produce feasible daily routes. The problem arises when the derived daily routes are found to exceed the maximum allowable length of the day, because all customers do not need to be seen every day. For example, the algorithm may design a template route for a vehicle considering all customers over, say, a week. Because all customers do not need to be visited every day, the duration of the template route will be greater than what is allowed on any given day. To fix this, an "expansion factor" is used to estimate the maximum allowable time

for a template route and prevent it from being exceeded. These bounds on the travel time are frequently recalculated by deriving daily routes, and checking for feasibility.

The ConRTR algorithm (see Algorithm 1) starts with an initial solution of template routes generated by the modified Clark and Wright algorithm and iteratively makes changes to the solution using one point, two point and two opt move operators, as shown in Figure 2. One-point moves change the position of a single node in the solution, and two-point moves swap the positions of two nodes in the solution. Two-opt moves, however, remove and replace two edges from the solution. A two-opt move may be either inter- or intra-route. The second stage of the ConRTR is based on Record to Record algorithm developed by Li, Golden, & Wasil (2004) which considers both “uphill” (non-improving) and “downhill” moves (improving) in the hope of finding better template solutions.

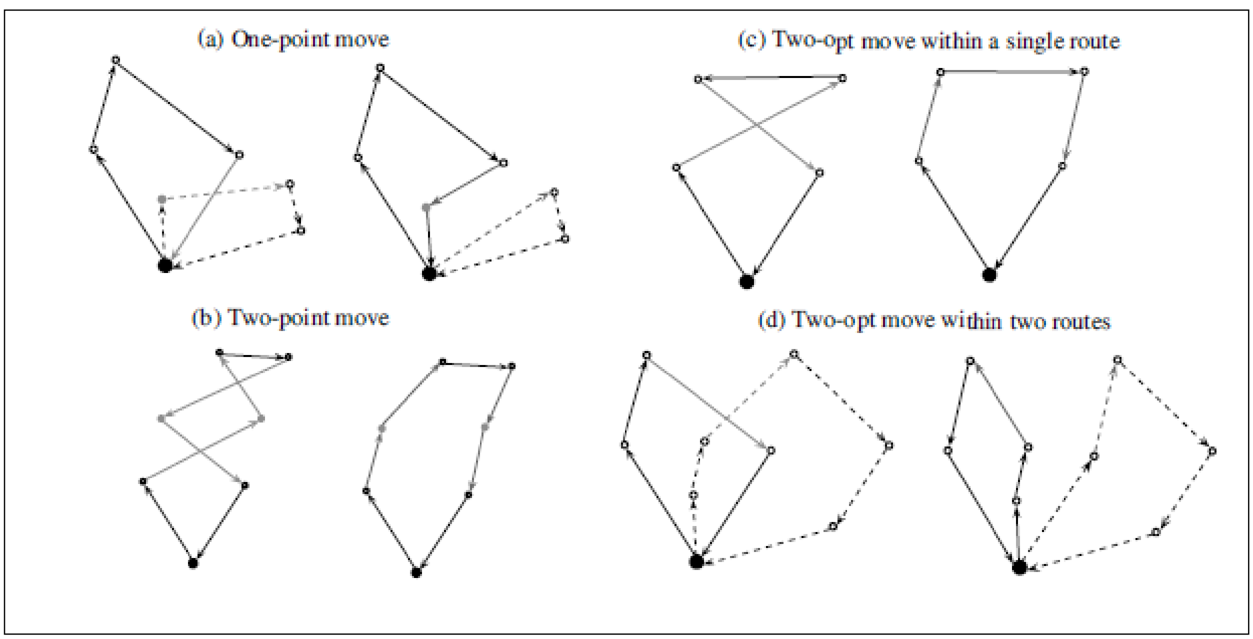

Figure 2 - Search operators used in the ConRTR algorithm (Gröer et al., 2008)

Algorithm 1 – The Consistent Record to Record Algorithm (Gröer et al., 2008)

Input:

- A set, *P,* of patients, with their visit frequency and geographic location known,
- A set of nurses each with maximum total travel time *T*, and *D* days in a planning horizon.
- *C* denotes the current set of template routes being considered
- *F* denotes the most recently generated set of template routes that is known to lead to feasible routes for each day
- *F** represents the set of template routes that leads to the lowest total travel time for the *D* days.
- *I* is the number of iterations in the diversification phase,
- *J* is the maximum number of non-improving iterations allowed before returning a solution,
- α represents a tolerance for the amount of deterioration allowed in the local search
- λ is a parameter used in the Clarke-and-Wright algorithm to quickly generate multiple initial solutions.
- Given a set of template routes *S*, let *f(S)* represent the total travel time of all *D* routes if they are feasible.

Ouput: Set of routes for all customers across all days in the horizon

1. Initialization

a) Set $I = 30$, $J = 5$, $\alpha = .01$, $l = 1$, and $\lambda = \{0.6, 1.0, 1.4\}$

b) Set $C = F = F^* = \emptyset$.

c) Partition the set of N customers into two groups, G_1 containing all customers requiring service on two or more days and G_2 containing all customers requiring service on only one day.

d) Compute an expansion factor $E = |G_1|/\mu_{daily}$ where μ_{daily} is the mean number of stops required on each day and $|N|$ is the number of customers in the template. Make an initial estimate for the maximum capacity of the template routes by setting $Q_{template} = Q \times E = Q_0$ and estimate the length of day of the template routes by setting $T_{template} = T/\sqrt{E} = T_0$.

e) For all customers in $G1$, set the demand amount and service time to be the mean values of these quantities taken across all days that the customer needs a visit.

2. **Create an initial set of template routes**

a) Generate an initial set of template routes C for the customers in G1 using the modified Clarke-and-Wright algorithm with parameter $\lambda\ [l]$ and maximum travel time $T_{\text{template}} = T_0$.

b) For each day d, create routes by removing customers from C not requiring service on day d and then inserting customers from G_2 requiring service only on day d..

c) If the routes for all D days are feasible, set $F = C$, $Q_{old} = Q_{template}$, $T_{old} = T_{template}$. Go to Step 3.

d) If at least one route on the D days is not feasible, then calculate the mean capacity violation (V_Q) and the mean travel time violation (V_T) across all routes, and tighten the template constraints by setting $Q_{template} = Q_{template} - V_Q/2$ and $T_{template} = T_{template} - V_T/2$.

Return to Step 2.a and try to generate a set of feasible template routes.

3. **Diversification Phase:** Modify the current feasible template routes $C = F$. If $f(C) < f(F^*)$, set $F^* = C$.

a) Set *Record* equal to the total travel time of all routes in the current template C. Set $Deviation = \alpha \times Record$.

b) For i = 1 to I

i) Apply the one-point move, two-point move, and two-opt move with record-to-record travel to the current template routes C. Accept any improving move and only those deteriorating moves where the total travel time of all template routes is less than $Record+Deviation$ and all routes satisfy the template constraints.

4. Improvement Phase: Improve the current solution C. Set $k = 0$.

a) Apply the one-point move, two-point move, and two-opt move, accepting only improving moves until no further improvements can be found.

b) Construct routes for each of the D days by applying the customer removal and insertion procedures.

c) If the routes for all D days are feasible, set $F = C$. Compute the minimum slack amount across all daily routes in terms of capacity (S_Q) and travel time (S_T). Relax the template constraints by setting $Q_{template} = Q_{template} + S_Q/2$ and $T_{template} = T_{template} + S_T/2$. Go to Step 5.

d) If at least one route on the D days is not feasible, then compute the mean violation (V_Q and V_T) as in Step 2f and tighten the template constraints by setting $Q_{template} = Q_{template} - V_Q/2$ and $T_{template} = T_{template} - V_T/2$. Set $k = k + 1$. If $k < 5$, continue to find feasible improvements by returning to Step 4.a. Otherwise, return to the last known feasible template, setting $C = F$ and go to Step 5.

5. If $f(C) < f(F^*)$, set $F^* = C$. If the objective function value of the current template routes C is less than Record, set $j = 0$. Set $j = j+1$. If $j < J$, return to Step 3 and continue modifying the

template. Otherwise, we have been unable to improve the current solution for J iterations, so stop modifying the current template and go to Step 6.

6. Set $l = l + 1$ and return to Step 2 to generate a new initial solution if $l \leq 3$. Otherwise, use the best set of template routes (F^*) found during the search to generate the routes for each of the D days and return.

4.2. ConRTR in a home care setting

In order to adapt the ConRTR to a home health care setting (see Algorithm 2), we make some minor changes. The ConRTR partitions the set of all customers to be visited across a given period into two groups – one set of customers need multiple visits in the period, and the other need only a single visit. The first adaptation of the ConRTR eliminates the need for the latter, because we make the fair assumption that patients requiring home care always need more than one visit over a horizon. Because of this, deriving daily routes is a simple matter of removing patients who do not need to be seen on a given day as shown in Figure 3. However, in the original ConRTR (Algorithm 1) there is an additional step that involves also adding patients who need to be seen only once, since the issue of visit consistency does not arise for such patients. The second adaptation concerns the demand and capacity associated with each customer and vehicle in a typical package delivery setting. In the modified ConRTR, we remove any capacity restrictions from the algorithm, because patients aren't typically associated with "demands".

The algorithm was coded in C++ using the classes and functions available in the open source VRPH library (See Appendix A.1.)

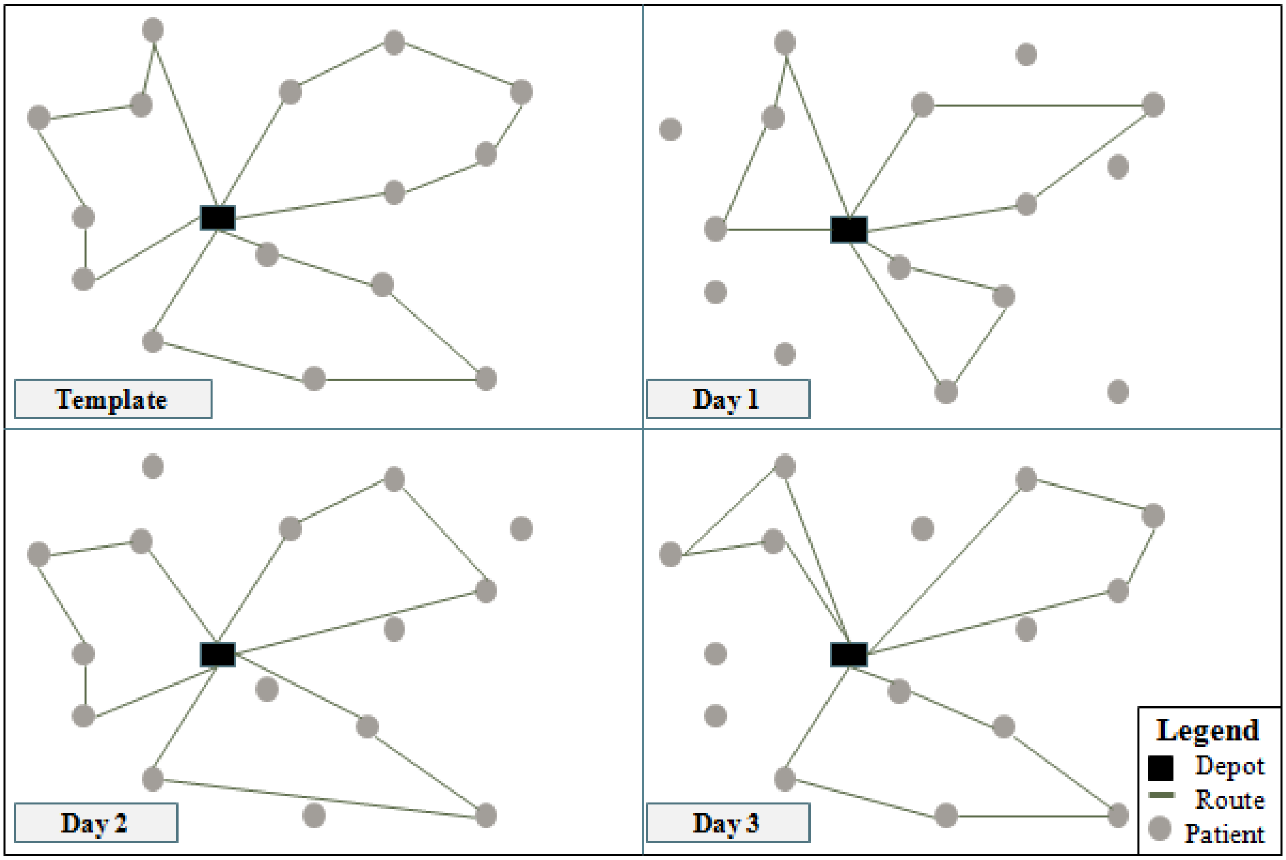

Figure 3 – Deriving daily routes from template routes in the ConRTR. In this example we assume a 3 day planning period. The routes for each day are derived by removing patients who do not need to be seen on a given day from the template routes, depicted in the top left corner.

Algorithm 2 – The Modified Consistent Record to Record Algorithm

Input:

- A set, *P*, of patients, with their visit frequency and geographic location known,
- A set of nurses each with maximum total travel time *T*, and *D* days in a planning horizon.
- *C* denotes the current set of template routes being considered
- *F* denotes the most recently generated set of template routes that is known to lead to feasible routes for each day

- F^* represents the set of template routes that leads to the lowest total travel time for the D days.
- I is the number of iterations in the diversification phase,
- J is the maximum number of non-improving iterations allowed before returning a solution,
- α represents a tolerance for the amount of deterioration allowed in the local search
- λ is a parameter used in the Clarke-and-Wright algorithm to quickly generate multiple initial solutions.
- Given a set of template routes S, let $f(S)$ represent the total travel time of all D routes if they are feasible.

Ouput: Set of routes for all patients across all days in the horizon

1. Initialization

a) Set $I = 30$, $J = 5$, $\alpha = .01$, $l = 1$, and $\lambda = \{0.6, 1.0, 1.4\}$ as suggested in Gröer et al., (2008)

b) Set $C = F = F^* = \emptyset$;

c) Compute an expansion factor $E = |P|/\mu_{daily}$ where μ_{daily} is the mean number of stops required on each day and $|P|$ is the number of patients in the template. Make an initial estimate for the maximum length of day of the template routes by setting $T_{template} = T/\sqrt{E} = T_0$.

d) For all customers in P, set the service time to be the mean values of these quantities taken across all days that the patient needs a visit.

2. Create an initial set of template routes

a) Generate an initial set of template routes C for the patients in P using the modified Clarke-and-Wright algorithm with parameter λ *[l]* and maximum travel time $T_{\text{template}} = T_0$.

b) For each day d, create routes by removing patients from C not requiring service on day d.

c) If the routes for all D days are feasible, set $F = C$, $T_{old} = T_{template}$. Go to Step 3.

d) If at least one route on the D days is not feasible, then calculate the mean travel time violation (V_T) across all routes, and tighten the template constraints by setting $T_{template} = T_{template} - V_T/2$. Return to Step 2.a and try to generate a set of feasible template routes.

3. **Diversification Phase:** Modify the current feasible template routes $C = F$. If $f(C) < f(F^*)$, set $F^* = C$.

a) Set *Record* equal to the total travel time of all routes in the current template C. Set $Deviation = \alpha \times Record$.

b) For $i = 1$ to I

i) Apply the one-point move, two-point move, and two-opt move with record-to-record travel to the current template routes C. Accept any improving move and only those deteriorating moves where the total travel time of all template routes is less than *Record*+*Deviation* and all routes satisfy the template constraints.

4. Improvement Phase: Improve the current solution C. Set $k = 0$.

a) Apply the one-point move, two-point move, and two-opt move, accepting only improving moves until no further improvements can be found.

b) Construct routes for each of the D days by applying the patient removal procedure.

c) If the routes for all D days are feasible, set $F = C$. Compute the minimum slack amount across all daily routes in terms of travel time (S_T). Relax the template constraints by setting $T_{template} = T_{template} + S_T/2$. Go to Step 5.

d) If at least one route on the D days is not feasible, then compute the mean violation (V_T) as in Step 2.f and tighten the template constraints by setting $T_{template} = T_{template} - V_T/2$. Set $k = k + 1$. If $k < 5$, continue to find feasible improvements by returning to Step 4.a. Otherwise, return to the last known feasible template, setting $C = F$ and go to Step 5.

5. If $f(C) < f(F^*)$, set $F^* = C$. If the objective function value of the current template routes C is less than Record, set $j = 0$. Set $j = j+1$. If $j < J$, return to Step 3 and continue modifying the template. Otherwise, we have been unable to improve the current solution for J iterations, so stop modifying the current template and go to Step 6.

6. Set $l = l + 1$ and return to Step 2 to generate a new initial solution if $l \leq 3$. Otherwise, use the best set of template routes (F^*) found during the search to generate the routes for each of the D days and return.

5. COMPUTATIONAL ANALYSIS: IMPACT OF CONTINUITY ON LONG-TERM PLANNING

In this section, we first present the planning strategies used to establish the impact of continuity and then describe the tests performed on them in a variety of settings. Next, we present an enhancement to the ConRTR that is designed for when the method is executed on problems with long planning horizons and show its efficacy computationally.

5.1. Planning strategies

We consider two planning strategies to understand the long-term impact of continuity. Both these strategies are implemented using the ConRTR. The first is a Week By Week planning strategy (WBW) typically suggested by many papers in the literature (Begur et al., 1997; Steeg & Schröder, 2008). In this strategy (see Algorithm 3), only those patients that need to be seen for a given week are considered, with decisions and routing assignments made in previous weeks carried over. New patients may need to be visited over the course of the horizon, and these patients are added on to the previous week's routes during the first week they need to be seen. At the beginning of each week, we assume that we have complete information about every patient that needs to be seen for that week.

Algorithm 3 – Week by Week Planning

Input: Weekly set of known patient locations and visit frequencies

Ouput: Set of routes for each week in the horizon

1. Set $i = 1$
2. While ($i \leq horizon$)
    a. If template routes from week (i-1) = Φ

i. Run modified ConRTR for all patients in week *i*. Derive daily routes from template routes to get daily route lengths

ii. Total horizon route length = Σ (Daily route lengths for week *i*)

iii. Save template routes

iv. Set $i = i + 1$

b. Else, read template routes from week (*i*-1)

i. Insert all new patients for week i into template routes without changing the assignments from previous weeks

ii. Derive daily routes for week *i* from new templates

iii. Total horizon route length = Total horizon route length + Σ (Daily route lengths for week *i*)

iv. Set $i = i + 1$

The second is a Long-term Template Planning (LTP) strategy that takes into account that some patients may not need to be visited right away at the start of the horizon, (see Algorithm 4). We consider this strategy to understand the benefits of planning 2-3 months ahead. We anticipate that by incorporating future patients early on in the plan, it is possible to design more strategic routes with lower costs.

Algorithm 4 – Long-term Template Planning

Input: Set of known patient locations and visit frequencies for all weeks in the horizon

Ouput: Set of routes for all weeks in the horizon

1. Run modified ConRTR for all patients across horizon

a. Derive daily routes from template routes to get daily route lengths

b. Total horizon route length = Σ (Daily route lengths over horizon)

Figures 4 and 5 below illustrate the difference between the week by week and the long-term template planning strategies. In this very simple example, we consider a 3 week long planning horizon and assume that each nurse can visit at most 3 patients in a day. In both approaches, Patients 1, 4 and 6 need to be seen from weeks 1 through 3, but patients 2, 5 and 3 only enter the home care program in weeks 2 and 3. When planning on a weekly basis, the nurse-patient assignments from week 1 are carried over to week 2 and only those patients for a given week are considered. When building long-term plans however, all patients, including those who don't start their home care program until weeks 2 or 3 are considered. While both planning strategies use the ConRTR templates, they differ in how much information about future patients they incorporate into the early planning stages. As a result, in week 1, with the WBW strategy, we need only one nurse to visit all patients. In week 2, however, two new patients make requests and a second nurse is needed in order to be able to complete all visits. In the long-term plan however, the nurse-patient assignments are slightly different – Nurse A visits patients 1 and 4, and nurse B visits patient 6. Because we planned for the patients who need to only be visited in weeks 2 and 3 ahead of time, we are able to create a schedule with a more balanced workforce over the length of the 3 week planning horizon.

	Week 1					Week 2					Week 3				
Days→	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Patient 1	✓	✓	✓	-	-	✓	✓	✓	-	-	✓	✓	✓	-	-
Patient 2	-	-	-	-	-	✓	-	✓	✓	-	✓	-	✓	✓	-
Patient 3	-	-	-	-	-	-	-	-	-	-	-	✓	-	✓	-
Patient 4	✓	-	✓	✓	✓	✓	-	✓	✓	✓	✓	-	✓	✓	✓
Patient 5	-	-	-	-	-	✓	✓	-	-	✓	✓	✓	-	-	✓
Patient 6	-	✓	-	✓	✓	-	✓	-	✓	✓	-	✓	-	✓	✓

Week 1: Nurse A 1-4-6

Decisions Carry Forward

Week 2: Nurse A 1-4-6, Nurse B 2-5

Decisions Carry Forward

Week 3: Nurse A 1-4-6, Nurse B 2-5-3

Figure 4 - Week by Week Planning

	Week 1					Week 2					Week 3				
Days→	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Patient 1	✓	✓	✓	-	-	✓	✓	✓	-	-	✓	✓	✓	-	-
Patient 2	-	-	-	-	-	✓	-	✓	✓	-	✓	-	✓	✓	-
Patient 3	-	-	-	-	-	-	-	-	-	-	-	✓	-	✓	-
Patient 4	✓	-	✓	✓	✓	✓	-	✓	✓	✓	✓	-	✓	✓	✓
Patient 5	-	-	-	-	-	✓	✓	-	-	✓	✓	✓	-	-	✓
Patient 6	-	✓	-	✓	✓	-	✓	-	✓	✓	-	✓	-	✓	✓

Week 1: Nurse A 1-4, Nurse B 6

Week 2: Nurse A 1-4-5, Nurse B 2-6

Week 3: Nurse A 1-4-5, Nurse B 2-3-6

Figure 5 - Long-term Template Planning

5.2. Experimental design

Our experimental design consists of 16 randomly generated instances spanning four factors with two levels each. The four factors are the number of new patients added per week, the density of population (i.e. rural or urban), the length of the planning horizon and the number of initial patients. We used a full factorial design; the levels for each factor are given in Figure 6.

		Levels	
		Low	**High**
Factors	No. of new patients/week (%I)	2.5%	5%
	Spread (S)	78.5 sq. miles (Urban)	706.5 sq. miles (Rural)
	Length of horizon (L)	2 months	3 months
	No. of initial patients (I)	200	400

Figure 6 - Factors and Levels for Experimental Design

We next describe some assumptions regarding the instances. First, we assume that each period is a five-day week. Therefore, a two-month horizon consists of 40 days, or 8 five-day weeks, and a three-month horizon consists of 60 days or 12 five-day weeks. Second, each patient has a 70% chance of being seen on each day of the week. In the rare event that no days are selected for a patient, we assume that the patient doesn't need to be visited. Once a visit frequency and schedule is determined, it does not change for the rest of the patient's episode of care. Third, we consider a circular region with a radius of 5 miles for urban settings and 15 miles for rural settings. We do so because we assume that the most a nurse will travel in an urban setting is 10 minutes; at a speed of 60 miles per hour, that translates to a distance of 10 miles or 1 mile per minute. Similarly, in a rural setting, we assume that a nurse will travel at most 30 minutes to visit a patient, or equivalently, at a speed of 60 miles per hour, 30 miles. Therefore, urban instances are spread over a 5 mile radius with an area of 78.5 sq. miles and rural settings are spread over a 15 mile radius and an area of 706.5 square miles. Fourth, each nurse may work at most ten hours a day. Fifth, travel time between two patients is determined by the geographical distance, *d*, calculated by the Haversine formula as given below, where *lat* is the latitude in radians, *lon* is the longitude in radians and *R* is the radius of the earth

$$a=\sin^2(\Delta lat/2)+\cos(lat_1)*\cos(lat_2)*\sin^2(\Delta long/2)$$

$$c=2*atan2(a^{1/2},(1-a)^{1/2})$$

$$d = R*c$$

Sixth, for our experiments we consider only a single depot where nurses start and end their day at the agency. Seventh, for our initial analysis, based on our interviews with the home care agency, we assume that each patient needs to be visited for an hour. However, we also discuss our subsequent experimentation studying instances with 30-minute visit durations. Finally, while our initial experimental design considers a growing demand scenario where new patients only enter the home care program, later experiments consider the possibility of patients also leaving the program as their care period ends as in a period of steady demand.

We use the following performance measures in our analysis:

1. **Total Travel Time:** Travel time is the total time spent by each nurse, *N*, in traveling to each patient in the set *P,* as required, and measured across all days, *D* in the planning horizon.

$$\textit{Total travel time} = \sum_N \sum_D \textit{Total daily travel time per nurse}$$

2. **Average Value Added Utilization:** Value-added (VA) utilization represents the amount of time nurses spend providing care. It is determined by the average amount of time nurses spend with their patients over the average total time they work by week.

$$\textit{Avg. VA Utilization per week} = \frac{\sum_N \textit{Total weekly visit time per nurse}}{\sum_N \textit{Total weekly route length per nurse}}$$

3. **Dispersion in Number of Nurses:** The dispersion is measured by the standard deviation in the number of nurses working on a weekly basis. In addition, we also report the average number of nurses scheduled across the horizon, and the maximum number of nurses scheduled at any week during the planning horizon. These metrics will help establish the staffing requirements for a given planning horizon.

4. **Average number of patients:** This metric is used to determine the number of patients, on average, seen by each nurse on a daily basis over the horizon. A higher value is desirable because it indicates that within a given restriction on the maximum number of hours a nurse may work, he or she is able to visit more patients.

$$Avg.\ No.\ of\ Patients = \frac{\sum_D Total\ number\ of\ patients\ visited\ per\ day}{\sum_D Total\ number\ of\ nurses\ scheduled\ per\ day}$$

5.3. Long-term Impact of Continuity

Table 2 below describes each instance used in our experiments. Columns #, I, %I, and N detail the instance number, the number of initial patients for each instance, percentage of new patients added on per week, and the number of new patients added on per week respectively. Column S represents the spread of the instance as described in Figure 6. Column H represents the length of the planning horizon in weeks. We use a 4-part notation to describe each instance - the number of initial patients is represented by I; the number of new patients per week is represented by N; urban spreads are represented by U and rural spreads by R; the horizon is represented by H. For example, 400I-10N-U-8H refers to an instance of 400 initial patients with 10 new patients each week in an urban setting over an 8-week long horizon.

#	I	% I	N	S	H
200I-5N-U-8H	200	2.5%	5	Urban	8 weeks
200I-10N-U-8H	200	5.0%	10	Urban	8 weeks
400I-10N-U-8H	400	2.5%	10	Urban	8 weeks
400I-20N-U-8H	400	5.0%	20	Urban	8 weeks
200I-5N-R-8H	200	2.5%	5	Rural	8 weeks
200I-10N-R-8H	200	5.0%	10	Rural	8 weeks
400I-10N-R-8H	400	2.5%	10	Rural	8 weeks
400I-20N-R-8H	400	5.0%	20	Rural	8 weeks
200I-5N-U-12H	200	2.5%	5	Urban	12 weeks
200I-10N-U-12H	200	5.0%	10	Urban	12 weeks
400I-10N-U-12H	400	2.5%	10	Urban	12 weeks
400I-20N-U-12H	400	5.0%	20	Urban	12 weeks
200I-5N-R-12H	200	2.5%	5	Rural	12 weeks
200I-10N-R-12H	200	5.0%	10	Rural	12 weeks
400I-10N-R-12H	400	2.5%	10	Rural	12 weeks
400I-20N-R-12H	400	5.0%	20	Rural	12 weeks

Table 2 - Instance Details

Tables 3(a)-(d) below show the savings achieved from comparing the total travel time required by the routes produced by the weekly planning strategy in comparison with the long-term template planning strategy. We categorize the Tables 3(a)-(d) by the spread (S) and the length of the horizon (L). As expected, the long-term solutions are almost always better than those from the weekly planning save for two instances ("400I-20N-U-8H" and "400I-10N-U-12H") in which the travel times for both LTP and WBW are almost identical. These results are also intuitive in that we see greater savings in rural settings, because nurses have to travel greater distances to see their patients. As a result, a poor routing decision in earlier weeks can have a greater impact simply because travel times are so much greater in rural areas. In the column "Travel Time Savings, WBW-LTP", we see that the savings in travel time alone can be as high as 50 hours depending on the size of the problem. This means that nurses will have an additional 50 hours that they can spend with patients rather than on travel. A comparison in the savings in travel time between urban and rural settings shows that agencies in rural areas have a lot more to gain by considering long-term planning, while the benefits for agencies in urban areas are not as high.

#	Travel Time (Hrs)		Travel Time Savings WBW-LTP (Hrs)
	WBW	LTP	
200I-5N-U-8H	273.5	267.9	**5.6**
200I-10N-U-8H	350.9	340.9	**10.0**
400I-10N-U-8H	525.6	515.4	**10.1**
400I-20N-U-8H	647.7	647.4	**0.3**

(a) Instances in an urban setting over an 8 week horizon

#	Travel Time (Hrs)		Travel Time Savings WBW-LTP (Hrs)
	WBW	LTP	
200I-5N-R-8H	607.7	584.1	**23.6**
200I-10N-R-8H	896.6	850.7	**45.9**
400I-10N-R-8H	1570.3	1542.1	**28.2**
400I-20N-R-8H	1574.4	1557.9	**16.5**

(b) Instances in a rural setting over an 8 week horizon

#	Travel Time (Hrs)		Travel Time Savings WBW-LTP (Hrs)
	WBW	LTP	
200I-5N-U-12H	423.7	400.0	**23.7**
200I-10N-U-12H	485.7	462.2	**23.5**
400I-10N-U-12H	965.5	965.5	**0.0**
400I-20N-U-12H	1127.2	1120.7	**6.6**

(c) Instances in an urban setting over an 12 week horizon

#	Travel Time (Hrs)		Travel Time Savings WBW-LTP (Hrs)
	WBW	LTP	
200I-5N-R-12H	1136.5	1070.6	**65.9**
200I-10N-R-12H	1239.7	1185.2	**54.5**
400I-10N-R-12H	2671.1	2667.3	**3.8**
400I-20N-R-12H	2620.4	2598.0	**22.4**

(d) Instances in a rural setting over a 12 week horizon

Table 3 – Savings in travel time from comparing WBW and LTP planning strategies

Table 4 compares the average number of nurses needed to visit all patients over the planning horizon in the WBW and LTP strategies. While the average number of nurses is always higher for LTP, the standard deviation in the number of nurses is significantly lower for this approach. Unlike in the WBW approach, a lower standard deviation in the LTP approach allows agencies to have better plan for their staffing requirements ahead of time. The reason for the larger number of nurses in LTP is clear from Table 5; the average number of patients seen by each nurse in LTP is lower than WBW. We can conclude that in spite of a larger number of nurses in the LTP strategy, we still see considerable savings in travel time because routes that aren't myopic are planned much more efficiently. This observation correlates to the average value-added utilization plotted over time in Figure 7. In the first few weeks, LTP sees lower utilization than WBW. However, we start to see an improvement in the utilization a few weeks in, corresponding to a decrease in the efficiency of the routes as more new patients make requests. Because we cannot re-optimize routes every week with the WBW strategy, new patients can only be incorporated in existing routes by inserting them into old routes, as determined by the least increase in route length. This limits the relevancy of the routing decisions

in the WBW strategy to only a given week under consideration. The continual insertion without re-optimization over a period of time leads to worsening solution quality as compared to the LTP approach where we design routes for every patient, including those in the future weeks, all at once.

#	WBW			LTP		
	Average	Std dev	Max.	Average	Std dev	Max.
200I-5N-U-8H	25.0	1.8	28	26.6	0.8	28
200I-10N-U-8H	28.8	3.4	34	31.5	0.5	33
400I-10N-U-8H	50.9	3.5	58	53.1	1.2	56
400I-20N-U-8H	55.2	6.1	64	62.7	0.8	65
200I-5N-R-8H	28.1	1.0	29	30.8	1.1	33
200I-10N-R-8H	32.7	3.1	38	37.1	1.6	39
400I-10N-R-8H	59.7	3.2	65	60.8	1.4	68
400I-20N-R-8H	64.3	6.9	76	70.0	2.3	77
200I-5N-U-12H	26.6	2.0	30	29.4	0.4	30
200I-10N-U-12H	29.9	4.5	38	35.4	0.8	36
400I-10N-U-12H	53.4	4.1	61	57.0	0.9	60
400I-20N-U-12H	61.8	9.4	78	71.1	1.7	75
200I-5N-R-12H	30.7	2.5	36	32.2	1.4	36
200I-10N-R-12H	33.7	4.1	41	40.6	2.2	44
400I-10N-R-12H	63.9	5.4	73	68.7	1.3	74
400I-20N-R-12H	69.7	8.9	85	80.5	3.1	89

Table 4 - Number of nurses: WBW v/s. LTP

#	WBW	LTP
200I-5N-U-8H	6.1	5.7
200I-10N-U-8H	5.7	5.2
400I-10N-U-8H	6.1	5.5
400I-20N-U-8H	6.0	5.1
200I-5N-R-8H	6.0	5.8
200I-10N-R-8H	6.0	5.3
400I-10N-R-8H	6.0	5.6
400I-20N-R-8H	5.7	4.9
200I-5N-U-12H	5.4	4.9
200I-10N-U-12H	5.2	4.5
400I-10N-U-12H	5.3	5.0
400I-20N-U-12H	5.4	4.5
200I-5N-R-12H	5.0	4.9
200I-10N-R-12H	5.1	4.7
400I-10N-R-12H	4.9	4.6
400I-20N-R-12H	5.1	4.4

Table 5 - Average number of patients seen per nurse per day: WBW v/s. LTP

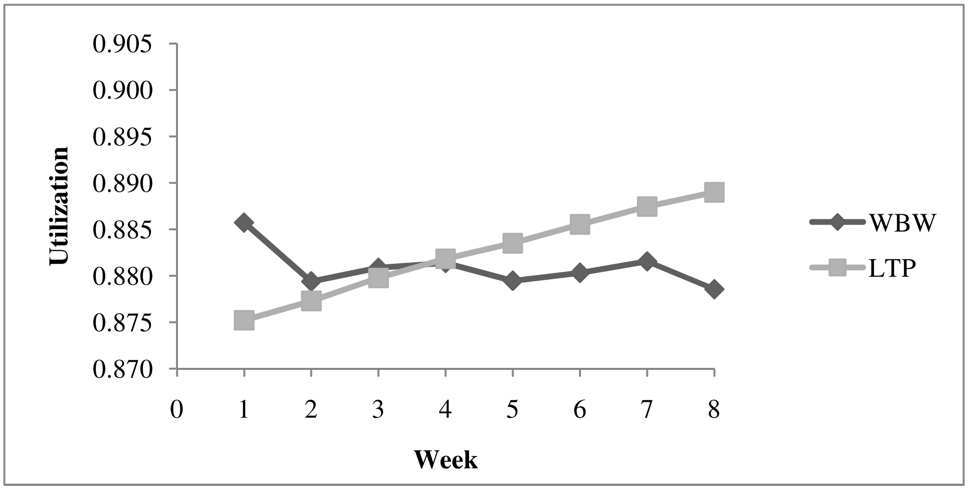

Figure 7 - Average weekly value-added utilization for a sample rural setting with 400 initial patients over an 8 week horizon.

5.4. Discounted Long-term Template Planning

The above experiment makes it clear that with continuity of care, it is indeed beneficial to consider long planning horizons. However, as seen in Table 4, the LTP approach has one significant drawback – the average number of nurses required to visit a given set of patients is always higher than the WBW approach. We attribute the increase in the number of routes to the fact that not all patients need to be seen as frequently, and that some patients need to be seen more often than others. This difference in frequency becomes more pronounced with patients that do not require care until late in the planning. The result is that the long-term template plan has routes with relatively poor value-added utilization in the first few weeks (Figure 7) compared to the later weeks. While such results are to be expected given the nature of the data and the way the templates are generated, it was worth investigating to see if the utilization of these routes could be improved by taking into account the frequency of visits.

The Discounted Long-term Template Planning method (DLTP) modifies the way the ConRTR templates are built by differentiating patients based on how frequently they need to be visited. With the LTP approach, those patients that don't need to be seen until later weeks are

also included in the templates. Because of this, daily routes in the weeks preceding a request made by such a patient are largely under-utilized. Factoring in the frequency of visits when building the templates therefore enables the algorithm to better allocate capacity.

We discount the service time S_i, for all nodes i by the number of visits, n_i, i requires over the horizon, H. The discounted service time, S_i^*, is expressed as,

$$S_i^* = \frac{n_i}{H}(S_i)$$

For example, if a patient needs three visits a week, each lasting an hour, for 6 weeks out of an 8 week horizon, the discounted visit time used for designing the template routes would be (3*6/5*8)*60 = 27 minutes.

Algorithm 5 – Discounted Long-term Template Planning

Input: Set of known patient locations and visit frequencies for all weeks in the horizon

Ouput: Set of routes for all weeks in the horizon

1. For all patients across horizon, calculate discounted visit times
2. Run modified ConRTR for all patients across horizon to get a set of template routes using discounted service times
    a. Derive daily routes from template routes by re-setting original visit times to get daily route lengths
    b. Total horizon route length = Σ (Daily route lengths over horizon)

We ran the DLTP approach over the same instances studied in the previous experiment (see Table 2) and found that this method did exceptionally well as clear from the results presented in Tables 6(a) - (d). The discounted template planning approach significantly outperformed the long-term template planning strategy in every instance as indicated by columns 6 and 7 ("Travel Time Savings"). The average percentage of savings achieved through the LTP approach was 3% compared to 12% with the DLTP approach.

#	Travel Time (Hrs)			Travel Time Savings (Hrs)	
	WBW	**LTP**	**DLTP**	**WBW - LTP**	**WBW – DLTP**
200I-5N-U-8H	273.5	267.9	242.8	5.6	**30.7**
200I-10N-U-8H	350.9	340.9	320.3	10.0	**30.6**
400I-10N-U-8H	525.6	515.4	494.1	10.2	**31.5**
400I-20N-U-8H	647.7	647.4	634.1	0.3	**13.6**

(a) Instances in an urban setting over an 8 week horizon

#	Travel Time (Hrs)			Travel Time Savings (Hrs)	
	WBW	**LTP**	**DLTP**	**WBW - LTP**	**WBW – DLTP**
200I-5N-R-8H	607.7	584.1	549.6	23.6	**58.1**
200I-10N-R-8H	896.6	850.7	789.6	45.9	**106.9**
400I-10N-R-8H	1570.3	1542.1	1388.5	28.2	**181.8**
400I-20N-R-8H	1574.4	1557.9	1392.6	16.5	**181.9**

(b) Instances in an urban setting over a 12 week horizon

#	Travel Time (Hrs)			Travel Time Savings (Hrs)	
	WBW	**LTP**	**DLTP**	**WBW - LTP**	**WBW – DLTP**
200I-5N-U-12H	423.7	400.0	384.4	23.7	**39.3**
200I-10N-U-12H	485.7	462.2	448.3	23.5	**37.4**
400I-10N-U-12H	965.5	965.5	946.7	0.0	**18.8**
400I-20N-U-12H	1127.2	1120.7	1103.3	6.6	**23.9**

(c) Instances in a rural setting over an 8 week horizon

#	Travel Time (Hrs)			Travel Time Savings (Hrs)	
	WBW	**LTP**	**DLTP**	**WBW - LTP**	**WBW – DLTP**
200I-5N-R-12H	1136.5	1070.6	1004.5	65.9	**132.1**
200I-10N-R-12H	1239.7	1185.2	1091.5	54.5	**148.2**
400I-10N-R-12H	2671.1	2667.3	2346.5	3.8	**324.6**
400I-20N-R-12H	2620.4	2598.0	2391.7	22.4	**228.7**

(d) Instances in a rural setting over a 12 week horizon

Table 6 – Savings in travel time from comparing WBW, LTP and DLTP strategies

Figure 8 below compares the average value-added weekly utilization for the same sample rural setting represented in Figure 7. There is a definite improvement in the average utilization with the DLTP strategy which also translates into a reduction in the number of nurses needed to make all visits. This behavior can be explained by Table 8 where we see that the average number of patients seen per nurse is higher for DLTP than WBW or LTP. The standard deviation in the average number of nurses in DLTP also reduces in as seen in Table 7. With the DLTP approach we always need a considerably fewer number of nurses than the WBW or the LTP approaches. Therefore, the DLTP approach eliminates the one disadvantage that LTP had compared to WBW.

#	WBW			LTP			DLTP		
	Average	**Std dev**	**Max.**	**Average**	**Std dev**	**Max.**	**Average**	**Std dev**	**Max.**
200I-5N-U-8H	25.0	1.8	28	26.6	0.8	28	23.2	0.8	24
200I-10N-U-8H	28.8	3.4	34	31.5	0.5	33	26.1	0.7	27
400I-10N-U-8H	50.9	3.5	58	53.1	1.2	56	46.3	0.8	48
400I-20N-U-8H	55.2	6.1	64	62.7	0.8	65	54.0	0.0	54
200I-5N-R-8H	28.1	1.0	29	30.8	1.1	33	25.0	0.0	25
200I-10N-R-8H	32.7	3.1	38	37.1	1.6	39	32.4	1.4	35
400I-10N-R-8H	59.7	3.2	65	60.8	1.4	68	49.1	0.8	51
400I-20N-R-8H	64.3	6.9	76	70.0	2.3	77	60.2	2.6	68
200I-5N-U-12H	26.6	2.0	30	29.4	0.4	30	26.2	0.8	27
200I-10N-U-12H	29.9	4.5	38	35.4	0.8	36	32.0	0.7	33
400I-10N-U-12H	53.4	4.1	61	57.0	0.9	60	52.3	0.8	54
400I-20N-U-12H	61.8	9.4	78	71.1	1.7	75	65.0	0.9	66
200I-5N-R-12H	30.7	2.5	36	32.2	1.4	36	29.0	1.6	33
200I-10N-R-12H	33.7	4.1	41	40.6	2.2	44	35.0	1.7	39
400I-10N-R-12H	63.9	5.4	73	68.7	1.3	74	59.7	1.3	66
400I-20N-R-12H	69.7	8.9	85	80.5	3.1	89	64.4	1.5	66

Table 7 - Number of nurses, comparing WBW, LTP and DLTP strategies

#	WBW	LTP	DLTP
200I-5N-U-8H	6.1	5.7	6.5
200I-10N-U-8H	5.7	5.2	6.3
400I-10N-U-8H	6.1	5.5	6.1
400I-20N-U-8H	6.0	5.1	5.6
200I-5N-R-8H	6.0	5.8	6.6
200I-10N-R-8H	6.0	5.3	6.1
400I-10N-R-8H	6.0	5.6	6.1
400I-20N-R-8H	5.7	4.9	5.4
200I-5N-U-12H	5.4	4.9	6.0
200I-10N-U-12H	5.2	4.5	5.2
400I-10N-U-12H	5.3	5.0	5.6
400I-20N-U-12H	5.4	4.5	5.2
200I-5N-R-12H	5.0	4.9	6.0
200I-10N-R-12H	5.1	4.7	5.4
400I-10N-R-12H	4.9	4.6	5.3
400I-20N-R-12H	5.1	4.4	5.5

Table 8 - Average number of patients seen per nurse per day, comparing WBW, LTP and DLTP strategies

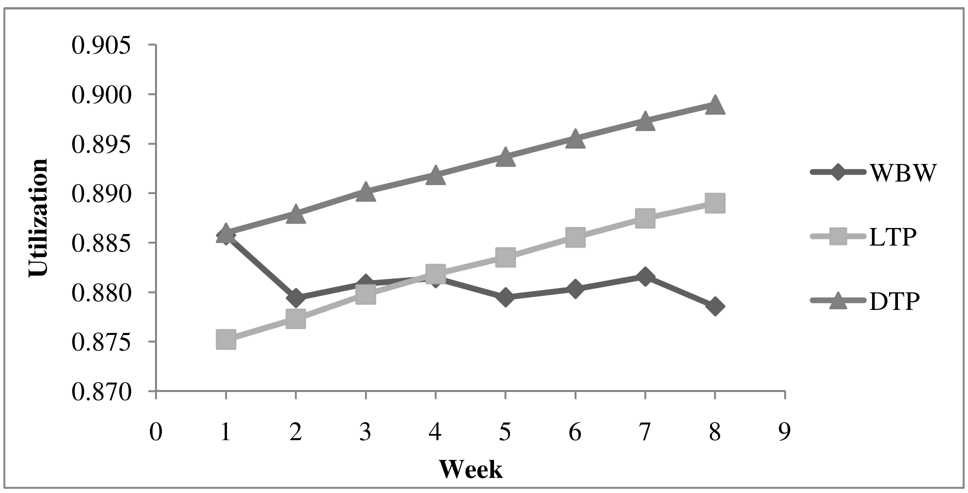

Figure 8- Average weekly utilization for a sample rural setting with 400 patients and an 8 week horizon

It is interesting to compare the weekly total route times associated with the WBW, LTP and DLTP strategies. By doing so, we can determine the magnitude of weekly savings in travel time in the LTP and DLTP approaches and identify the ("break-even") point in the horizon where savings begin to be seen. It is easy to see that with the WBW strategy, the first few weeks always have more efficient routes because we make decisions most relevant to that point of time. This myopic planning has its benefits in the first few weeks. However, as the weeks go on, long-term planning becomes more prudent. With the DLTP strategy, we start to almost always see

savings quicker than with the LTP strategy as seen in Table 9. The magnitude of weekly savings over time is presented for each instance in Appendix A.2.

#	Week	
	LTP	**DLTP**
200I-5N-U-8H	5	1
200I-10N-U-8H	5	3
400I-10N-U-8H	4	2
400I-20N-U-8H	5	5
200I-5N-R-8H	4	2
200I-10N-R-8H	4	3
400I-10N-R-8H	4	1
400I-20N-R-8H	5	2
200I-5N-U-12H	5	3
200I-10N-U-12H	6	6
400I-10N-U-12H	7	6
400I-20N-U-12H	7	7
200I-5N-R-12H	4	2
200I-10N-R-12H	6	5
400I-10N-R-12H	8	2
400I-20N-R-12H	7	5

Table 9 - Week we start to see savings

5.5. Implementation

Our algorithms were coded in C++ on a machine with a 1.80 GHz AMD P820 triple-core processor (4 GB Ram). It should be noted that all of our experiments ran fairly quickly. The average computation times in seconds across instances described previously in Table 2 are given in Table 10 below. The DLTP approach was faster for all instances than the LTP approach. This is presumably because on average the DLTP approach finds the best solution much faster than the LTP approach as shown in Table 11 with the exception of a few instances.

Instance Size	Avg. Computational Times (seconds)	
	LTP	**DLTP**
200 Initial Patients	126.2	56.7
400 Initial Patients	399.0	258.0

Table 10- Average Computational Times

#	Avg. time taken to find best solution (seconds)	
	LTP	DLTP
200I-5N-U-8H	45.5	18.1
200I-10N-U-8H	75.6	32.7
400I-10N-U-8H	160.0	54.8
400I-20N-U-8H	215.4	362.1
200I-5N-R-8H	73.5	53.7
200I-10N-R-8H	64.3	72.2
400I-10N-R-8H	163.8	561.4
400I-20N-R-8H	335.2	1498.0
200I-5N-U-12H	68.4	27.7
200I-10N-U-12H	100.0	44.7
400I-10N-U-12H	195.6	102.0
400I-20N-U-12H	791.3	118.7
200I-5N-R-12H	137.1	38.9
200I-10N-R-12H	234.1	59.5
400I-10N-R-12H	281.0	248.9
400I-20N-R-12H	658.7	279.3

Table 11 - Time taken to find best solution

5.6. Experiments with reduced visit times

To determine the impact of the duration of visit times on the magnitude of savings, we conducted some tests on reducing the length of the each visit from 60 minutes to 30 minutes. The details instances used in this experiment are described in Table 12.

#	I	% N	N	S	H
400I-5N-U-8H	400	2.5%	10	Urban	8 weeks
400I-10N-U-8H	400	5.0%	20	Urban	8 weeks
400I-5N-R-8H	400	2.5%	10	Rural	8 weeks
400I-10N-R-8H	400	5.0%	20	Rural	8 weeks
400I-5N-U-12H	400	2.5%	10	Urban	12 weeks
400I-10N-U-12H	400	5.0%	20	Urban	12 weeks
400I-5N-R-12H	400	2.5%	10	Rural	12 weeks
400I-10N-R-12H	400	5.0%	20	Rural	12 weeks

Table 12 - Instances for experiments with reduced visit times

As indicated in Tables 13 (a) and (b), with 30 minute visit times, the LTP strategy performs much better than in the experiments with 60 minute visit times. As a result, the difference in performance between LTP and DLTP approaches is reduced. This behavior can be

explained by the fact that the strength of the DLTP strategy lies in the discounting of visit times when building the template routes. However, with reduced visit times, the impact that this discounting can make reduces.

#	30 Minute Visit Times: Savings in Travel Time		60 Minute Visit Times: Savings in Travel Time		30 Minute Visit Times : Max. No. of Nurses			60 Minute Visit Times : Max. No. of Nurses		
	WBW-LTP (Hrs)	WBW-DLTP (Hrs)	WBW-LTP (Hrs)	WBW-DLTP (Hrs)	WBW	LTP	DLTP	WBW	LTP	DLTP
400I-5N-U-8H	20.9	33.6	10.2	31.5	29	28	27	58	56	48
400I-10N-U-8H	35.9	40.6	0.3	13.6	36	33	32	64	65	54
400I-5N-R-8H	82.9	130.5	28.2	181.8	37	40	32	65	68	51
400I-10N-R-8H	45.5	97.1	16.6	181.9	42	46	36	76	77	68
400I-5N-U-12H	34.6	50.3	0.0	18.9	33	33	30	61	60	54
400I-10N-U-12H	60.5	84.9	6.6	23.9	41	40	36	78	75	66
400I-5N-R-12H	102.0	276.5	3.8	324.6	47	46	44	73	74	66
400I-10N-R-12H	108.8	292.4	22.4	228.7	49	54	53	85	89	66

(a) Comparison of savings in travel time when visit time is varied

#	30 Minute Visit Times: No. of nurses									60 Minute Visit Times: No. of nurses								
	WBW			LTP			DLTP			WBW			LTP			DLTP		
	Avg	Std dev	Max	Avg	Std dev	Max	Avg	Std dev	Max	Avg	Std dev	Max	Avg	Std dev	Max	Avg	Std dev	Max
400I-5N-U-8H	26.0	2.1	29	25.9	0.5	25.9	24.3	0.8	24.3	50.9	12.0	58	53.1	1.4	56	46.3	0.7	48
400I-10N-U-8H	31.1	3.6	36	31.7	0.5	31.7	30.2	0.5	30.2	55.2	36.9	64	62.7	0.6	65	54.0	0.0	54
400I-5N-R-8H	33.6	2.4	37	33.5	1.3	33.5	28.1	1.7	28.1	59.7	10.1	65	60.8	2.0	68	49.1	0.6	51
400I-10N-R-8H	36.3	3.6	42	39.3	2.2	39.3	32.3	1.8	32.3	64.3	47.6	76	70.0	5.3	77	60.2	7.0	68
400I-5N-U-12H	27.6	2.8	33	30.1	0.9	30.1	27.7	0.6	27.7	53.4	16.9	61	57.0	0.9	60	52.3	0.6	54
400I-10N-U-12H	34.6	5.1	41	36.5	1.2	36.5	34.1	1.2	34.1	61.8	88.2	78	71.1	2.8	75	65.0	0.9	66
400I-5N-R-12H	38.3	3.6	47	39.1	1.6	39.1	36.3	1.7	36.3	63.9	29.0	73	68.7	1.6	74	59.7	1.6	66
400I-10N-R-12H	40.2	5.3	49	44.9	3.4	44.9	41.5	4.1	41.5	69.7	79.8	85	80.5	9.8	89	64.4	2.4	66

(b) Number of nurses when visit time is varied

Table 13 – Comparison of 30 minute versus 60 minute visit times

We also notice that with reduced visit times, nurses can see more patients in a day (Table 14), leading to fewer routes as seen in Table 13b. Previous experiments also seem to suggest that fewer routes lead to increased savings. However, there is another, more significant reason that we see greater savings in the experiments with reduced visit times. By visiting more patients

each day on average than in the 60 minute visit time experiments, nurses travel greater distances. On one hand, this increased travel can magnify the impact of a poor routing decision in the WBW approach and on the other, leave scope for larger improvements in the LTP and DLTP approaches. On average, we see approximately 7% more savings in travel time across all instances with the 30 minute visit times as compared with the 60 minute visit times.

#	30 Min Visit Times: No. of patients seen per nurse			60 Min Visit Times: No. of patients seen per nurse		
	WBW	LTP	DLTP	WBW	LTP	DLTP
400I-5N-U-8H	11.8	11.8	12.6	6.1	5.5	6.1
400I-10N-U-8H	10.7	10.4	11.0	6.0	5.1	5.6
400I-5N-R-8H	8.8	8.9	10.6	6.0	5.6	6.1
400I-10N-R-8H	9.0	8.3	10.1	5.7	4.9	5.4
400I-5N-U-12H	11.6	10.6	11.5	5.3	5.0	5.6
400I-10N-U-12H	10.1	9.6	10.2	5.4	4.5	5.2
400I-5N-R-12H	8.2	8.0	8.7	4.9	4.6	5.3
400I-10N-R-12H	8.8	7.9	8.5	5.1	4.4	5.5

Table 14 - Average number of patients seen per nurse per day: 30 minute versus 60 minute visit times

5.7. Experiments in a state of steady demand

Our previous experiments did not consider patients whose care periods ended during the planning horizon. However, we also wanted to study scenarios which include both patients who are "added" to the system, as well as patients who "drop" out of the system at some point in the horizon to model a system in a state of steady demand. We assume that once a patient is added, he or she cannot be dropped from the system, and similarly, once a patient is dropped, he or she cannot be added again. We randomly generated instances with an equal number of patients who are added and dropped over the horizon. These instances are described in Table 15. The column "Add/Drop" describes the number of patients added and dropped over the entire horizon. For example, instance 1 has a planning horizon of 8 weeks. If we consider that the number of new patients added on average each week is expressed as a percentage of the number of initial

patients in week 1, we have 5 new patients (2.5% of 200 initial patients) on average each week from weeks 2 through 8. At the end of the horizon we can expect to have 35 new patients. The number of patients ending their care can also be calculated in a similar manner. While the number of added patients need not necessarily match the number of dropped patients ever week, we designed the instances so that these numbers even out at the end of the horizon.

#	I	Add/Drop	S	H
200I-5N-U-8H	200	35	Urban	8 weeks
200I-10N-U-8H	200	70	Urban	8 weeks
200I-5N-R-8H	200	35	Rural	8 weeks
200I-10N-R-8H	200	70	Rural	8 weeks
200I-5N-U-12H	200	55	Urban	12 weeks
200I-10N-U-12H	200	110	Urban	12 weeks
200I-5N-R-12H	200	55	Rural	12 weeks
200I-10N-R-12H	200	110	Rural	12 weeks

Table 15 – Steady demand experiment instances

The results, as seen in Table 16, show that long-term planning definitely has relevance in this type of scenario as well, both in terms of savings in travel time as well as the number of nurses required. Figure 9 shows the average utilization of a sample instance with 200 patients in a rural setting over an 8 week horizon. The WBW strategy shows a decrease in utilization over time, unlike the DLTP strategy which stays relatively constant. This trend, also observed in our previous experiments that only considered adding patients, is presumably because the efficiency of each route drops as patients are inserted into the template routes (without re-optimizing so as to maintain continuity) in later weeks. We decided to include the LTP approach in a portion of our study to see compare its performance with DLTP. As with growing demand experiments, DLTP continues to result in better savings in the steady demand experiments. When averaged across the instances, the DLTP approach results in a percentage savings of 18%.

#	Travel Time (Hrs)			Savings in Travel Time (Hrs)	
	WBW	LTP	DLTP	WBW-LTP	WBW-DLTP
200I-5N-U-8H	334.3	309.1	278.3	25.3	**56.0**
200I-10N-U-8H	424.1	378.9	346.7	45.3	**77.5**
200I-5N-R-8H	911.4	816.8	727.9	94.6	**183.6**
200I-10N-R-8H	1253.7	1078.0	967.9	175.8	**285.8**
200I-5N-U-12H	557.7	488.2	476.8	69.5	**80.9**
200I-10N-U-12H	846.7	686.0	721.2	160.7	**125.4**
200I-5N-R-12H	1589.6	1473.3	1370.1	116.4	**219.5**
200I-10N-R-12H	1970.3	1688.4	1508.4	281.9	**461.9**

Table 16 – Savings in travel time for steady-demand experiments: WBW v/s. DLTP

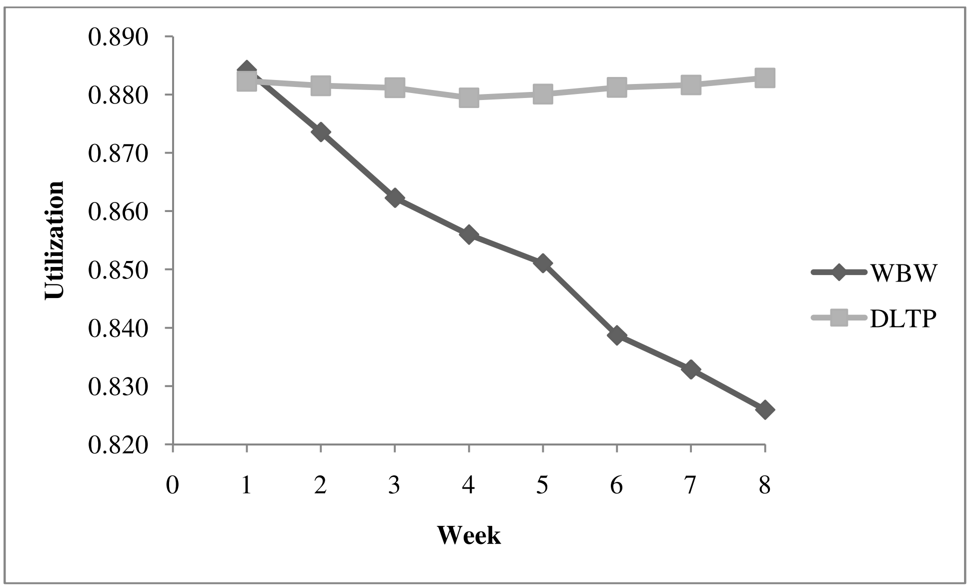

Figure 9 - Average weekly utilization for a sample rural setting with 200 initial patients over an 8 week horizon in a steady demand experiment

Table 18 compares the average number of patients seen per nurse. The DLTP approach allows each nurse to see visit more patients on a daily basis. As with our previous experiments, the standard deviation and the average number of nurses over the horizon (Table 17) are much lower in the DLTP approach than the WBW. The similarities between the results in the steady state demand and the growing demand experiments are encouraging, because they imply that long-term planning has relevance in a range of scenarios.

#	WBW	DLTP
200I-5N-U-8H	5.3	6.6
200I-10N-U-8H	4.8	6.1
200I-5N-R-8H	4.4	5.1
200I-10N-R-8H	4.1	4.7
200I-5N-U-12H	4.9	5.7
200I-10N-U-12H	4.4	4.9
200I-5N-R-12H	4.4	5.0
200I-10N-R-12H	4.0	4.7

Table 17 – Average number of patients seen per nurse per day, steady demand experiments

#	WBW			DLTP		
	Average	Std dev	Max.	Average	Std dev	Max.
200I-5N-U-8H	31.0	2.8	36	24.5	0.4	27
200I-10N-U-8H	39.6	6.7	51	30.3	0.2	32
200I-5N-R-8H	36.8	2.8	41	31.2	0.2	35
200I-10N-R-8H	46.6	6.5	58	40.1	0.2	44
200I-5N-U-12H	36.1	3.8	42	30.6	0.1	32
200I-10N-U-12H	50.4	9.1	69	44.4	0.1	45
200I-5N-R-12H	41.1	4.7	50	36.1	0.2	40
200I-10N-R-12H	54.8	8.8	69	45.2	0.3	53

Table 18 – Number of nurses in the steady demand experiments: WBW v/s. DLTP

5.8. Long-term Template Planning with Uncertain Information

While the previous experiments show that long-term planning has clear advantage, it is not possible for agencies to know fully in advance which patients need to be visited when, and just as importantly, where, over an 8 or 12-week planning horizon. There are different ways of modeling this uncertainty depending on the nature of the problem and the information available such as scenario-based planning as in Bent & Van Hentenryck (2004) or a two stage plan with recourse action (Hvattum, Lokketangen, & Laporte, 2006). The disadvantage with using these approaches is that we typically need some information about the probability of a customer making a request and his/her location ahead of time. However, home health agencies may not have access to that kind of information. We prefer a method that has low data requirements and that agencies can easily use. Given historical data, it is possible to obtain estimates of the average

number of patients to be seen in a given month. Additionally, it is reasonable to assume that agencies typically know how long their current patients need to be seen. Therefore, the most significant unknowns from a scheduling and routing standpoint are, when new patients make requests, where these patients are located, and a weekly visit schedule that is mutually acceptable to both the patient and the agency.

The first problem which is knowing when these new patients need visits can be solved by comparing the known weekly average number of patients for a given week with the estimated weekly average and filling in any "deficit" with dummy patients at the depot. If the weekly average for a given week is higher than the estimated weekly average, we do not add any dummy customers and leave the set of patients as is. By locating these dummy patients at the depot, we circumvent the second unknown, namely, where new patients are located. While we do not have an estimate for the travel time involved in visiting these dummy patients, we can artificially reduce each nurse's daily visit capacity so as to intelligently leave room in their schedules for new patients when these visits actually need to be made. Finally to account for the third unknown, the weekly visit requirements are set to be such that each dummy patient must be visited every day of the week. While doing so takes a rather pessimistic view, it allows us to avoid making any presumptions about the nature and frequency of future patient visit requirements. To summarize, we model uncertainty by comparing weekly averages and locating any missing patients at the depot.

The scenario just described is illustrated as a simple example in Figure 10. Assume we have a 4 week long planning horizon, and 15 patients in week 1 with an estimated weekly average of 14 patients. We know that in week 2, two patients drop out of the routes because their episode of care has ended and they no longer need to be visited, bringing the number of patients

to 13. The difference between the number of patients in week 2 and the estimated weekly average is 1, and therefore we add a single dummy patient, located at the depot. Similarly in weeks 3 and 4, based on the number of patients who drop out, we determine the number of dummy patients to be 3 and 6 respectively. More importantly, we can make these decisions ahead of time.

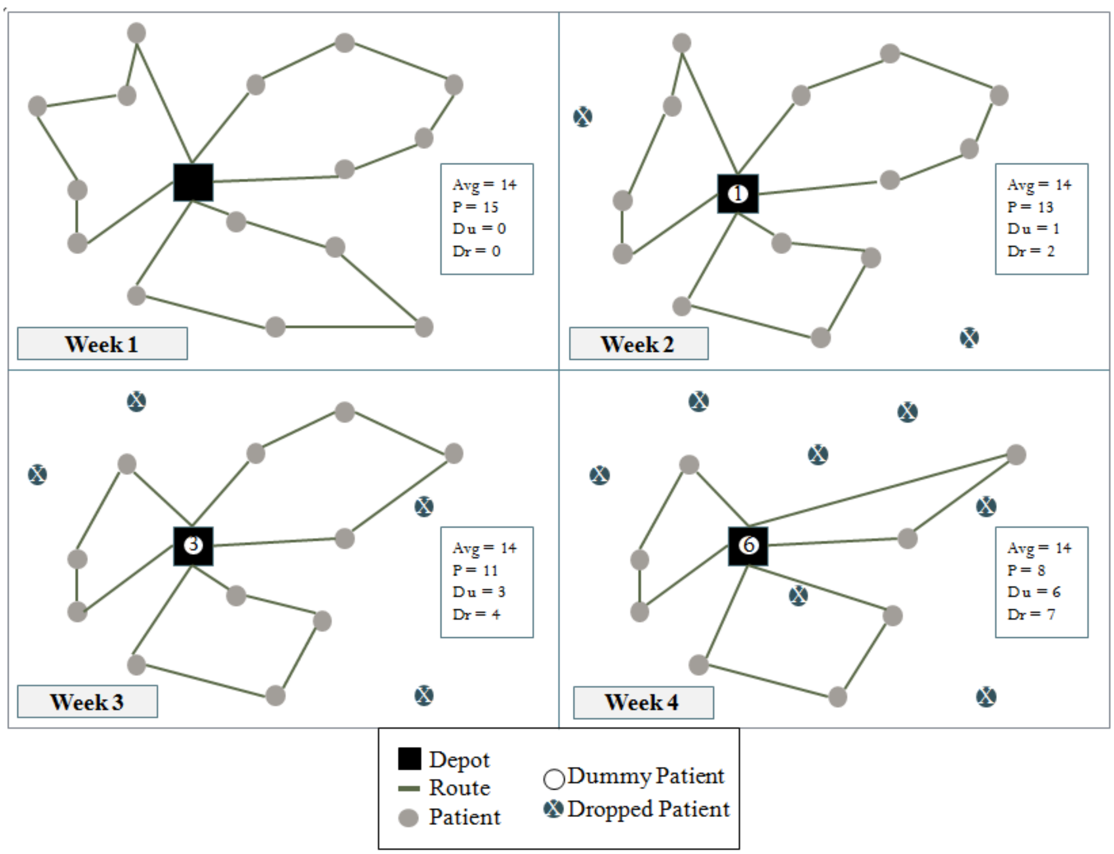

Figure 10 – Discounted Long-term Template Plan with Uncertain Information. 'Avg' represents the estimated weekly average, 'P' represents the number of known patients, 'Du' represents the number of dummy patients and 'Dr' represents the number of patients who drop out of the system because their episode of care ends.

We can then develop a long-term plan that includes the patients who continue to be visited through the length of the horizon, the patients who end their care at some point during and, the dummy patients. In the example above, we would include all nodes marked "patient", "dummy patient" and "dropped patient" over the 4 week horizon. The solution from this

template plan, the Discounted Long-term Template with Uncertain Information (DLTP (UI)), can then be used to derive routes on a week by week basis, as information about new patients becomes available. At this point, the dummy patients are taken out of the template solution and new patients are inserted into the template routes, in place of the dummy patients. Algorithm 6 describes the planning and implementation stages of the DLTP (UI). It is important to note that Step 5 is carried out only when information about patients in a given week becomes available, typically, at the beginning of the week.

Algorithm 6 – Discounted Long-term Template Planning with Uncertain Information

---

Input: Set of known patient locations (including patients who need to be visited throughout the horizon and patients who end their episode of care in the duration of the horizon) and visit frequencies for all weeks in the horizon; Estimated average number of patients per week

Ouput: Set of routes for each week in the horizon

Stage I: Planning

1. Set $i = 1$

2. While ($i \leq horizon$)

   a. If (Average number of patients for week i < *estimated weekly average*)

      i. Add (*estimated weekly average – average number of patients for week i*) dummy patients with location set at depot and weekly visit frequency = 5

      ii. Set $i = i + 1$

   b. Elseif (*Average number of patients for week i* > *estimated weekly average*)

      i. Set $i = i + 1$

3. Using discounted service times, run modified ConRTR for all known and dummy patients across horizon from step 2 to output a set of template routes as in Discounted Long-term Template Planning (Algorithm 5).

4. Set $i = 1$

Stage II: Implementation

5. While $(i \leq horizon)$

a. If (Average number of patients for week i < *estimated weekly average*)

i. Replace dummy patients with new patients for week i from template routes obtained in step 3.

b. Elseif (*Average number of patients for week i* > *estimated weekly average*)

i. Insert new patients for week i into template routes without changing the assignments from previous weeks

c. Derive daily routes for week i from new templates

d. Total horizon route length = Total horizon route length + Σ (Daily route lengths for week i)

e. Set $i = i + 1$

The instances used in the DLTP (UI) approach are the same as in the add-drop experiments, the details of which are provided in Table 15. The results from the experiments with uncertain information are shown in Table 19. Expectedly, the DLTP (UI) approach cannot compete with the results from the DLTP approach and these differences are reported in column 9 ("Savings in Travel Time DLTP (UI) – DLTP"). More importantly however, DLTP (UI) does better than the WBW strategy, as seen in column 8 ("Savings in Travel Time WBW – DLTP

(UI)"). Table 20 reports the average number of nurses and the standard deviation associated. We see that the standard deviation for the DLTP (UI) approach is still comparatively high, but is always lower than the standard deviation in the WBW approach. The comparison of the average number of patients seen per nurse per day in Table 21 reveals that DLTP (UI) is not far behind DLTP and even surpasses it in several instances.

One possible explanation for this behavior is that with the DLTP approach, all patients, including the ones that aren't visited until much later in the planning horizon are built into the template routes. In DLTP (UI), the number of patients on the long-term template routes is determined by the forecasted weekly average. Patients who only need to be seen in the last few weeks, and are in "excess" of the forecasted weekly average for planning purposes alone, are not incorporated into the first stage of the DLTP (UI) approach. As a result, the patients who need to be visited only in the last few weeks are inserted into the template routes in the implementation stage of the algorithm, as they make requests. Because of this, the number of patients in the long-term template routes is higher in DLTP compared to DLTP (UI). Therefore, the DLTP (UI) designs templates that focus on patients that are seen during a majority of the planning horizon, and makes decisions for those patients seen only toward the end that are relevant to that particular week. The DLTP-UI strategy combines the short-term benefits of WBW planning with the long-term benefits of long-term planning. As a related point of interest, an examination of the routes revealed that a majority of the new patients represented by dummy patients tend to be clustered on the same template route produced by stage I of the algorithm. On the other hand, those patients that are not on the template routes and are only inserted in stage II of the algorithm when they need to be seen are spread over different routes.

#	Travel Time (Hrs)			Savings in Travel Time (Hrs)	
	WBW	DLTP	DLTP (UI)	WBW-DLTP (UI)	DLTP (UI)-DLTP
200I-5N-U-8H	334.3	278.3	305.7	28.7	27.4
200I-10N-U-8H	424.1	346.7	399.9	24.2	53.3
200I-5N-R-8H	911.4	727.9	793.2	118.3	65.3
200I-10N-R-8H	1253.7	967.9	1067.8	185.9	99.9
200I-5N-U-12H	557.7	476.8	525.3	32.4	48.5
200I-10N-U-12H	846.7	721.2	788.7	58.0	67.5
200I-5N-R-12H	1589.6	1370.1	1443.7	145.9	73.6
200I-10N-R-12H	1970.3	1508.4	1796.4	173.9	288.0

Table 19 - Savings in travel time from comparing WBW, DLTP and DLTP (UI) strategies

#	WBW			DLTP			DLTP (UI)		
	Avg	Std dev	Max.	Avg	Std dev	Max.	Avg	Std dev	Max.
200I-5N-U-8H	31.0	2.8	36	24.5	0.4	27	25.9	1.4	28
200I-10N-U-8H	39.6	6.7	51	30.3	0.2	32	32.2	2.6	38
200I-5N-R-8H	36.8	2.8	41	31.2	0.2	35	29.3	3.1	33
200I-10N-R-8H	46.6	6.5	58	40.1	0.2	44	34.9	4.9	41
200I-5N-U-12H	36.1	3.8	42	30.6	0.1	32	28.4	1.7	29
200I-10N-U-12H	50.4	9.1	69	44.4	0.1	45	39.6	3.4	46
200I-5N-R-12H	41.1	4.7	50	36.1	0.2	40	32.3	4.2	38
200I-10N-R-12H	54.8	8.8	69	45.2	0.3	53	41.1	5.6	47

Table 20 - Number of nurses, comparing WBW, DLTP and DLTP (UI) strategies

#	WBW	DLTP	DLTP (UI)
200I-5N-U-8H	5.3	6.6	6.1
200I-10N-U-8H	4.8	6.1	5.7
200I-5N-R-8H	4.4	5.1	5.6
200I-10N-R-8H	4.1	4.7	5.6
200I-5N-U-12H	4.9	5.7	6.2
200I-10N-U-12H	4.4	4.9	5.4
200I-5N-R-12H	4.4	5.0	5.7
200I-10N-R-12H	4.0	4.7	5.4

Table 21 - Average number of patients seen per nurse per day for WBW, DLTP and DLTP (UI) strategies

We conducted some tests on the DLTP (UI) approach with 30 minute visit times for the instances described in Table 15. We found that while the instances with an 8-week planning horizon still showed improvement over the WBW approach as seen in column 5 (“30 Minute Visit Times: Savings in travel time, WBW-DLTP (UI) (Hrs)”) of Table 22, most of the instances over the 12 week planning horizon didn’t register any improvement. However, the average

number of nurses and the associated standard deviation for the 30 minute visit times, as seen in Table 23, are still better in the DLTP (UI) approach as compared to the WBW. As with 60 minute visit time experiments reported in Table 21, we notice that the average number of patients seen by each nurse per day is greater in DLTP-UI than DLTP (See Table 24).

#	30 Minute Visit Times: Savings in Travel Time		60 Minute Visit Times: Savings in Travel Time	
	WBW-DLTP (UI) (Hrs)	DLTP (UI)-DLTP(Hrs)	WBW-DLTP (UI) (Hrs)	DLTP (UI)-DLTP(Hrs)
200I-5N-U-8H	**17.2**	21.4	28.7	27.4
200I-10N-U-8H	**13.9**	49.7	24.2	53.3
200I-5N-R-8H	**25.3**	107.9	118.3	65.3
200I-10N-R-8H	**62.1**	183.1	185.9	99.9
200I-5N-U-12H	**-0.8**	78.3	32.4	48.5
200I-10N-U-12H	**-6.2**	168.4	58.0	67.5
200I-5N-R-12H	**14.4**	250.9	145.9	73.6
200I-10N-R-12H	**-2.3**	456.7	173.9	288.0

Table 22 – Comparison of savings in travel time for DLTP (UI): 30 minute v/s. 60 minute visit times

#	30 Minute Visit Times: No. of nurses									60 Minute Visit Times: No. of nurses								
	WBW			DLTP			DLTP (UI)			WBW			DLTP			DLTP (UI)		
	Avg	Std dev	Max	Avg	Std dev	Max	Avg	Std dev	Max	Avg	Std dev	Max	Avg	Std dev	Max	Avg	Std dev	Max
200I-5N-U-8H	17.0	2.6	21	13.3	0.8	15	13.5	1.6	15	31	2.8	36	24.5	0.4	27	26.5	1.4	28
200I-10N-U-8H	20.7	3.8	27	16.6	0.6	19	15.5	2.1	19	39.6	6.7	51	30.3	0.2	32	33.0	2.6	38
200I-5N-R-8H	20.8	2.1	26	19.0	0.5	24	17.4	3.0	21	36.8	2.8	41	31.2	0.2	35	28.9	3.1	33
200I-10N-R-8H	27.2	3.5	33	22.9	0.8	30	19.6	5.1	26	46.6	6.5	58	40.1	0.2	44	34.3	4.9	41
200I-5N-U-12H	19.0	2.6	23	14.8	1.0	17	15.3	1.9	17	36.1	3.8	42	30.6	0.1	32	28.2	1.7	29
200I-10N-U-12H	27.7	5.0	36	20.8	0.4	23	20.6	2.3	22	50.4	9.1	69	44.4	0.1	45	40.1	3.4	46
200I-5N-R-12H	23.7	2.6	28	19.8	0.5	27	18.8	4.2	24	41.1	4.7	50	36.1	0.2	40	32.0	4.2	38
200I-10N-R-12H	30.3	5.0	38	24.9	0.5	36	23.3	5.9	32	54.8	8.8	69	45.2	0.3	53	40.5	5.6	47

Table 23 – Number of Nurses for DLTP (UI): 30 minute v/s. 60 minute visit times

#	WBW	DLTP	DLTP (UI)
200I-5N-U-8H	9.8	12.2	12.2
200I-10N-U-8H	9.3	11.2	12.2
200I-5N-R-8H	7.8	8.5	9.6
200I-10N-R-8H	7.0	8.3	10.3
200I-5N-U-12H	9.4	12.0	11.7
200I-10N-U-12H	8.1	10.5	10.7
200I-5N-R-12H	7.7	9.1	10.1
200I-10N-R-12H	7.2	8.6	9.8

Table 24 - Average number of patients seen per day for DLTP (UI): 30 minute v/s. 60 minute visit times

We conclude that the DLTP (UI) approach is quite competitive and in the absence of perfect information, it will provide agencies with a good solution for long-term planning with an improvement over weekly planning, especially because of the added advantage of more steady staffing requirements that can be determined ahead of time.

6. CONCLUSIONS

In the healthcare setting, continuity of care is an extremely important quality of care metric. In order to maintain continuity of care and cost-effective schedules and routes, our computational results show that it is worth considering long planning horizons, particularly in rural settings and relatively short visit times. Our experiments indicate that long-term planning is relevant even in scenarios with a varying patient base, such as when patients' care periods end and they no longer need visits, and when new patients make requests to be seen. In an industry that is short staffed and overworked, the savings in travel time will afford nurses more time with their patients as opposed to on the road.

By factoring in the frequency of visits a patient requires over the length of the horizon when developing the template routes, we were able to significantly improve the efficiency of the routes as compared to the long-term planning strategy. The Discounted Long-term Template Planning approach has three major benefits. First, we see an increase in the magnitude of travel time savings, especially in those instances with long visit times. Second, nurses are able to see more patients per day on average. And finally, the average number of nurses required to visit all patients and the standard deviation associated is considerably lower than in the weekly planning approach. Agencies can therefore have a better idea of their staffing requirements and plan accordingly ahead of time.

The final contribution of this research is the simple stochastic routing methodology proposed. The advantages of this approach are that it does not need much effort or data to implement and we are still able to incorporate a long-term plan into a weekly plan, as soon more information about new customers becomes available.

7. FUTURE WORK

There is plenty of scope for future experimentation. For one, the problem can be extended to model patient and nurse preferences, as well as incorporate nurse skills when making assignments. Secondly, many agencies allow nurses to start and end their day from home, without needing them to stop by at the agency. Another possibility is that some patients may need to be seen only during certain times. From a VRP modeling standpoint, including multiple depots and time windows while maintaining continuity of care will give the problem an interesting additional dimension.

It is obviously not possible to know with complete certainty which patients enter home care programs in future. We demonstrated the potential for reducing weekly routing costs when information about future patients is unknown through a relatively simple planning strategy. However, as evidenced by the large gap between the DLTP (UI) and the DLTP solutions, there is still a lot of room for improvement. The research in this paper indicates the strong need to develop good forecasting methods in conjunction with stochastic routing algorithms to be able to accurately determine when and where future patients need to be visited.

Another avenue of interest is to examine the robustness of long-term planning to dynamic changes (such as a nurse calling in sick, or a patient requesting that a visit be cancelled) that are no doubt a daily occurrence in the home care setting. Finally, in order for the methods proposed in this research to be successful in a realistic setting, it is important that the long-term planning approaches are evaluated over a rolling horizon, of, say, a year. By carrying forward decisions made over a planning horizon of 2-3 months to the next planning horizon on a rolling-basis for a

year, it is possible to evaluate the performance of the long-term planning methods and study its strengths and limitations in a real-life setting.

REFERENCES

Bach, J., Intintola, P., Alba, a., & Holland, I. (1992). The ventilator-assisted individual. Cost analysis of institutionalization vs rehabilitation and in-home management. *Chest*, *101*(1), 26-30. doi:10.1378/chest.101.1.26

Baldacci, R., Mingozzi, A., & Roberti, R. (2011). Recent exact algorithms for solving the vehicle routing problem under capacity and time window constraints. *European Journal of Operational Research*. Elsevier B.V. doi:10.1016/j.ejor.2011.07.037

Begur, S. V., Miller, D. M., & Weaver, J. R. (1997). An Integrated Spatial DSS for Scheduling and Routing Home-Health-Care Nurses. *Interfaces*, *27*(4), 35-48. doi:10.1287/inte.27.4.35

Bennett, A. R., & Erera, A. L. (2011). Dynamic periodic fixed appointment scheduling for home health. *IIE Transactions on Healthcare Systems Engineering*, *1*(1), 6-19. doi:10.1080/19488300.2010.549818

Bent, R. W., & Van Hentenryck, P. (2004). Scenario-Based Planning for Partially Dynamic Vehicle Routing with Stochastic Customers. *Operations Research*, *52*(6), 977-987. doi:10.1287/opre.1040.0124

Bertels, S., & Fahle, T. (2006). A hybrid setup for a hybrid scenario: combining heuristics for the home health care problem. *Computers & Operations Research*, *33*(10), 2866-2890. doi:10.1016/j.cor.2005.01.015

Cabana, M., & Jee, S. (2004). Does continuity of care improve patient outcomes. *Journal of Family Practice*, *53*(December), 974-981. Retrieved from http://www.jfponline.com/pages.asp?aid=1830

Campbell, A. M., & Savelsbergh, M. W. P. (2005). Decision Support for Consumer Direct Grocery Initiatives. *Transportation Science*, *39*(3), 313-327. doi:10.1287/trsc.1040.0105

Cheng, E., & Rich, J. L. (1998). A home health care routing and scheduling problem. *Rice University, Texas, Tech. Rep. TR98-04*, 1-24. Citeseer. Retrieved from http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.51.8869&rep=rep1&type=pdf

Dantzig, G. B., & Ramser, J. H. (1959). The Truck Dispatching Problem. *Management Science*, *6*(1), 80-91. doi:10.1287/mnsc.6.1.80

Eveborn, P., Flisberg, P., & Rönnqvist, M. (2006). LAPS CARE - An operational system for staff planning of home care. *European Journal of Operational*, *171*, 962-976. Retrieved from http://www.sciencedirect.com/science/article/pii/S0377221705001207

Eveborn, P., Rönnqvist, M., Einarsdóttir, H., Eklund, M., Lidén, K., & Almroth, M. (2009). Operations Research Improves Quality and Efficiency in Home Care. *Interfaces*, *39*(1), 18-34. doi:10.1287/inte.1080.0411

Flatberg, T., Hasle, G., Kloster, O., Nilssen, E. J., & Riise, A. (2005). Dynamic and stochastic aspects in vehicle routing-a literature survey. *STF90 A05413, SINTEF Applied Mathematics*. Citeseer. Retrieved from http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.108.8134&rep=rep1&type=pdf

Francis, P., Smilowitz, K., & Tzur, M. (2008). The period vehicle routing problem and its extensions. In B. Golden & et al. (Eds.), *The Vehicle Routing Problem: Latest Advances and New Challenges* (pp. 73-102). Springer, Heidelberg. Retrieved from http://www.springerlink.com/index/m871p8h230856065.pdf

Gendreau, M., Laporte, G., & Seguin, R. (1995). An exact algorithm for the vehicle routing problem with stochastic demands and customers. *Transportation science*, *29*(2), 143-155. Institute for Operations Research and the Management Sciences. Retrieved from http://cat.inist.fr/?aModele=afficheN&cpsidt=3536014

Gendreau, M., Laporte, G., & Seguin, R. (1996). Stochastic vehicle routing. *European Journal of Operational Research*, *88*(1), 3-12. doi:10.1016/0377-2217(95)00050-X

Gendreau, M., Laporte, G., & Séguin, R. (1996). A tabu search heuristic for the vehicle routing problem with stochastic demands and customers. *Operations Research*, 469–477. JSTOR. Retrieved from http://www.jstor.org/stable/171707

Groer, C., Golden, B., & Wasil, E. (2010). A library of local search heuristics for the vehicle routing problem. *Mathematical Programming Computation*, 79-101. doi:10.1007/s12532-010-0013-5

Gröer, C., Golden, B., & Wasil, E. (2008). The Consistent Vehicle Routing Problem. *Manufacturing & Service Operations Management*, *11*(4), 630-643. doi:10.1287/msom.1080.0243

Hemmelmayr, V., Doerner, K., & Hartl, R. (2009). A variable neighborhood search heuristic for periodic routing problems. *European Journal of Operational Research*, *195*(3), 791-802. Elsevier B.V. doi:10.1016/j.ejor.2007.08.048

Hvattum, L. M., Lokketangen, A., & Laporte, G. (2006). Solving a dynamic and stochastic vehicle routing problem with a sample scenario hedging heuristic. *Transportation Science*, *40*(4), 421-438. doi:10.1287/trsc.1060.0166

Kergosien, Y., Lenté, C., & Billaut, J. C. (2009). Home health care problem: An extended multiple traveling salesman problem. *4th Multidisciplinary International Conference on Scheduling: Theory and Applications (MISTA'09), Dublin (Irlande)* (pp. 10–12). Retrieved from http://www.mistaconference.org/2009/papers/085-092-110-P.pdf

Kytojoki, J., Nuortio, T., Braysy, O., & Gendreau, M. (2007). An efficient variable neighborhood search heuristic for very large scale vehicle routing problems. *Computers & Operations Research*, *34*(9), 2743-2757. doi:10.1016/j.cor.2005.10.010

Laporte, G., Gendreau, M., Potvin, J.-Y., & Semet, F. (2000). Classical and modern heuristics for the vehicle routing problem. *International Transactions in Operational Research*, *7*(4-5), 285-300. doi:10.1111/j.1475-3995.2000.tb00200.x

Li, F., Golden, B., & Wasil, E. (2004). Very large-scale vehicle routing: new test problems, algorithms, and results. *Computers & Operations Research*, *32*, 1165 - 1179. doi:10.1016/j.cor.2003.10.002

Mladenovic, N., & Hansen, P. (1997). Variable neighborhood search. *Computers & Operations Research*, *24*(11), 1097-1100. doi:10.1016/S0305-0548(97)00031-2

NAHC. (2010). *Basic statistics about Home Care, Updated 2010*. Retrieved from http://www.nahc.org/facts/10HC_Stats.pdf

NYAHSA, & HCA. (2011). *State Budget Cuts Have Imperiled New York ' s Home Health Care Delivery System*. Retrieved from http://www.hca-nys.org/documents/VitalSignsPR.pdf

Pigott, H. E., & Trott, L. (1993). Translating Research into Practice: The Implementation of an In-Home Crisis Intervention Triage and Treatment Service in the Private Sector. *American Journal of Medical Quality*, *8*(3), 138-144. doi:10.1177/0885713X9300800306

Pirkwieser, S., & Raidl, G. R. (2008). A variable neighborhood search for the periodic vehicle routing problem with time windows. *Proceedings of the 9th EU/MEeting on Metaheuristics for Logistics and Vehicle Routing, Troyes, France* (pp. 1-7). Retrieved from http://webhost.ua.ac.be/eume/workshops/eume08/papers/paper21.pdf

Reinelt, G. (1991). TSPLIB - A Traveling Salesman Problem Library. *ORSA Journal on Computing*, *3*(4), 376-384. Retrieved from http://joc.journal.informs.org/content/3/4/376.short

Smilowitz, K., Nowak, M., & Jiang, T. (2011). Workforce management in periodic delivery operations. *Transportation Science*, 1-31.

Steeg, J., & Schröder, M. (2008). A hybrid approach to solve the periodic home health care problem. *Operations Research Proceedings 2007*, 297–302. Springer. Retrieved from http://www.springerlink.com/index/r82766653070t160.pdf

Vincenty, T. (1975). Direct and inverse solutions of geodesics on the ellipsoid with application of nested equations. *Survey review*, *XXII*(176). Retrieved from http://www.ingentaconnect.com/content/maney/sre/1975/00000023/00000176/art00008

Zhong, H., Hall, R. W., & Dessouky, M. (2007). Territory Planning and Vehicle Dispatching with Driver Learning. *Transportation Science*, *41*(1), 74-89. doi:10.1287/trsc.1060.0167

8. APPENDIX A

A.1 The VRPH Library

The library of Heuristics for the Vehicle Routing Problem is a free, open-source library written in C++ and available for download at http://www.coin-or.org/projects/VRPH.xml. The library was designed and developed by Chris Groer and Bruce Golden (Groer et al., 2010). It contains a number of local search heuristics including the record to record algorithm (Li et al., 2004) that is a part of the ConRTR algorithm.

The VRPH library has a number of classes and functions that are common to heuristics for variants of the VRP, such as search operators like the one-point and two-point moves as well as algorithms to generate an initial solution, including the Clarke and Wright algorithm and the Sweep algorithm. The details about the library and its capabilities are available in the paper by (Groer et al., 2010).

Figure 11 shows the VRPH functions used in the modified ConRTR. The object V is an instantiation of the class *VRP* in the VRPH. The class VRP contains details about the instance such as the size, number of patients and their locations, visit frequency, maximum length of day, etc., that is input in the TSPLIB format(Reinelt, 1991). It also "contains a number of data structures and methods that can be used to access problem information, modify the solution, and output the solution to buffers and files."(Groer et al., 2010). The object CW is an instantiation of the class *ClarkeWright* and contains functions to build an initial solution from the Clarke and Wright algorithm. Finally, the object VV is an instantiation of the class *VRPViolation* that stores information about the violation of maximum length of day constraints in order to suitably tighten the template travel time as required in algorithm 1. The functions and their input arguments that

are used in each stage of the modified ConRTR are listed in the figure. The VRPH library and its modifications for the purpose of this thesis is available in the attached CD.

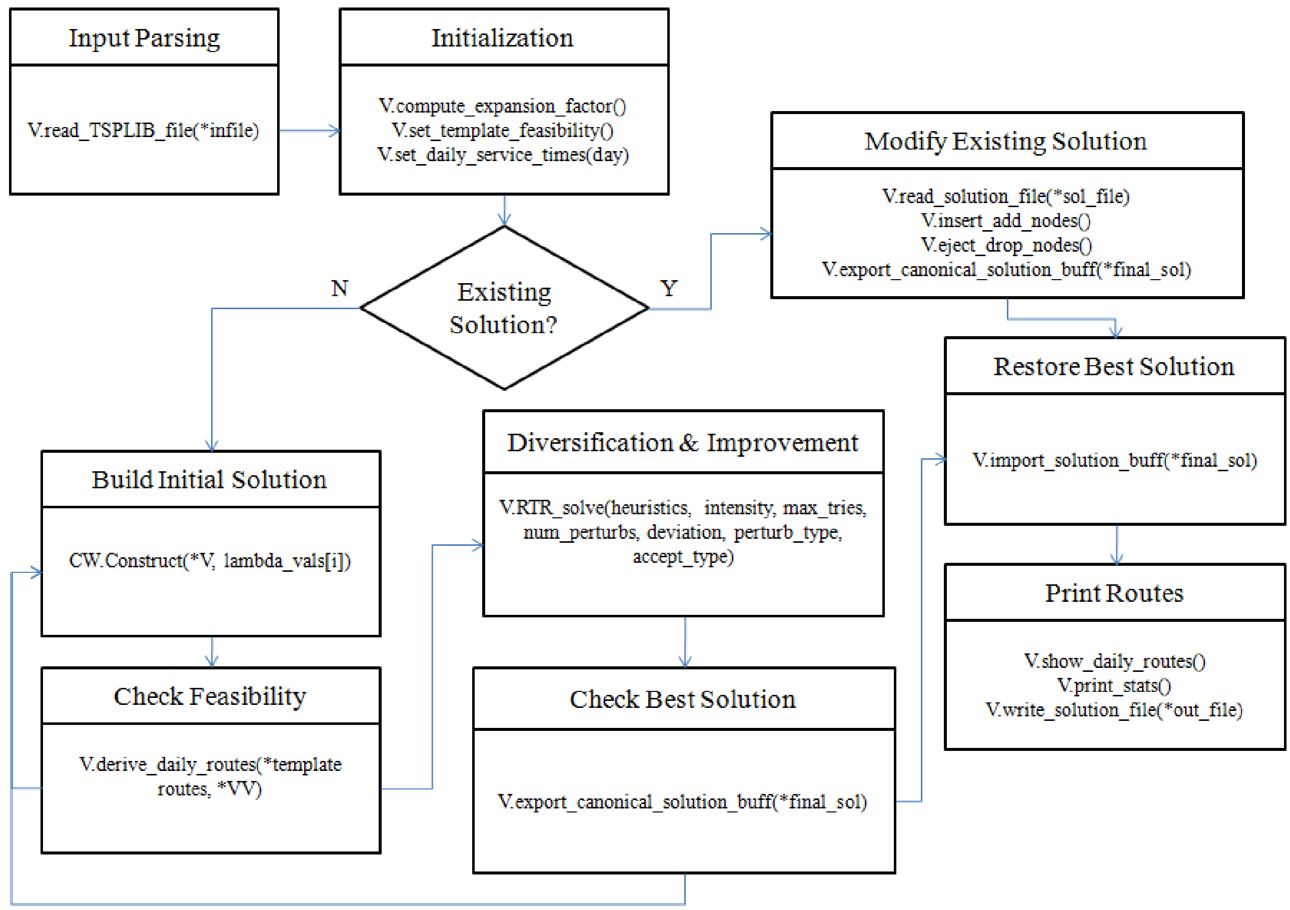

Figure 11 - VRPH Functions used in the modified ConRTR

A description of the directory structure and the files is provided here. The supplement to the VRPH library is available in *VRPH_Supplement.rar*. Upon extraction, the contents of the root directory become available. The folder, *Data*, lists all the instances used in the HHCRP as well as a set of benchmark instances for the VRP, all of which are formatted for input. The data files need to be listed in the same file as the executable. The following are the relevant input arguments that may be provided:

1. -f "file_name.vrp": Reads instance details from given file name

2. -out "file_name.out": Writes solution generated by the heuristic to the given file name (optional)

3.-sol "file_name.out": Reads initial solution from existing file (optional)

The folder *VRPH WBW* contains files needed to solve the VRPH based on a weekly planning strategy. The *VRPH_Supplement/VRPH WBW/MSVC_2008/VRPH* directory contains a *VRPH.sln* file that contains projects for building the VRPH library and the vrp_rtr application. When building the solution, it must be ensured that the *VRPH_Supplement/VRPH WBW/inc* directory is listed in the correct path under additional include directories.

1. vrp_rtr: implementation of the Con_RTR algorithm used for week by week planning

The folder *VRPH LTP* contains files needed to solve the VRPH based on a long-term planning strategy. The *VRPH_Supplement/VRPH LTP/MSVC_2008/VRPH* directory contains a *VRPH.sln* file that contains projects for building the VRPH library and the applications listed below. Again, when building the solution, it must be ensured that the *VRPH_Supplement/VRPH LTP/inc* directory is listed in the correct path under additional include directories.

1) vrp_rtr: implementation of the Con_RTR algorithm used for long-term planning

2) vrp_rtr_service_factor: implementation of the Con_RTR with discounted service times used for long-term planning

3) vrp_rtr_weekbyweek: implementation of the weekly planning strategy with uncertain information.

4) data_generator: generates instances for the different HHCRP experiments, given input parameters for number of patients, length of horizon, maximum length of day for nurses, location of patients, visit requirements for patients and duration of visit for patients.

A.2 Magnitude of Weekly Savings

In this section we discuss the magnitude of weekly savings for the growing and steady-state demand experiments. The weekly savings are calculated by comparing the weekly costs associated with WBW and LTP, DLTP and DLTP (UI) and are presented in Tables 25-28. A negative value indicates that the routes in the WBW approach have a lower total travel time for a given week compared with the LTP, DLTP or DLTP (UI) approaches. Through the tables provided below, we can also determine the week we start to see savings as presented in Table 9 in Section 5.4. Tables 27 and 28, show that the magnitude of weekly savings in DLTP (UI) approach is better than that of the DLTP approach for the first one or two weeks. As discussed in section 5.8, the DLTP (UI) designs long-term template routes that focus on the patients that are seen for the majority of the planning horizon and ignores those patients that are seen only in the last few weeks, presumably resulting in the improved initial savings. However, this changes as the magnitude of weekly savings in DLTP quickly becomes greater than DLTP-UI, as the benefits of long-term planning with perfect information become more pronounced.

#	Weeks											
	1	2	3	4	5	6	7	8	9	10	11	12
200I-5N-U-8H	-2.9	-1.6	-0.8	-0.1	0.6	2.0	4.1	4.3	-	-	-	-
200I-10N-U-8H	-6.8	-3.5	-1.7	0.0	2.2	4.0	7.3	8.6	-	-	-	-
400I-10N-U-8H	-4.9	-2.1	-0.7	0.5	2.5	3.2	5.3	6.4	-	-	-	-
400I-20N-U-8H	-13.6	-8.2	-5.1	-2.7	2.1	6.4	9.6	11.9	-	-	-	-
200I-5N-R-8H	-5.2	-1.8	-0.4	3.1	5.0	6.2	7.0	9.8	-	-	-	-
200I-10N-R-8H	-18.6	-8.8	-3.4	4.9	9.2	14.8	20.8	27.1	-	-	-	-
400I-10N-R-8H	-18.4	-3.7	-2.0	0.9	7.8	10.3	11.9	21.5	-	-	-	-
400I-20N-R-8H	-35.5	-19.2	-8.0	-4.8	9.4	14.3	24.6	35.8	-	-	-	-
200I-5N-U-12H	-3.5	-2.5	-0.6	-0.2	0.2	1.7	3.2	3.4	3.8	4.9	6.4	6.8
200I-10N-U-12H	-7.4	-5.2	-4.0	-2.2	-1.4	1.7	2.7	5.0	6.2	7.8	8.9	11.3
400I-10N-U-12H	-9.9	-8.0	-4.9	-3.0	-1.9	-0.9	0.2	1.8	4.7	5.3	7.4	9.3
400I-20N-U-12H	-21.2	-14.5	-13.4	-7.8	-4.8	-2.1	3.3	7.7	10.2	12.8	16.6	19.8
200I-5N-R-12H	-11.2	-5.4	-4.1	1.4	4.3	5.3	7.8	9.8	10.6	11.9	16.2	19.2
200I-10N-R-12H	-23.4	-16.8	-14.0	-9.2	-4.0	3.9	8.8	12.0	18.2	21.0	25.8	32.0
400I-10N-R-12H	-38.0	-25.5	-16.8	-13.6	-6.8	-4.6	-1.9	5.0	13.1	29.8	31.1	32.1
400I-20N-R-12H	-53.0	-38.0	-30.9	-18.8	-11.1	-0.9	8.0	18.4	24.5	32.4	39.5	52.2

Table 25 - Magnitude of weekly savings in hours WBW v/s. LTP – Growing Demand

#	Weeks											
	1	2	3	4	5	6	7	8	9	10	11	12
200I-5N-U-8H	0.2	1.6	2.3	2.9	3.8	5.1	7.3	7.5	-	-	-	-
200I-10N-U-8H	-4.4	-1.0	0.9	2.7	4.8	6.7	9.8	11.2	-	-	-	-
400I-10N-U-8H	-2.4	0.5	2.0	3.2	5.2	5.9	8.1	9.1	-	-	-	-
400I-20N-U-8H	-11.9	-6.5	-3.5	-1.0	3.7	8.0	11.3	13.6	-	-	-	-
200I-5N-R-8H	-0.9	2.6	3.8	7.3	9.2	10.4	11.5	14.1	-	-	-	-
200I-10N-R-8H	-10.3	-1.1	4.1	12.5	16.5	22.1	28.3	34.7	-	-	-	-
400I-10N-R-8H	0.5	15.2	17.0	19.5	27.2	29.6	31.3	41.5	-	-	-	-
400I-20N-R-8H	-15.5	1.3	12.8	16.1	30.3	35.1	45.4	56.3	-	-	-	-
200I-5N-U-12H	-2.1	-1.1	0.8	1.1	1.5	3.0	4.5	4.7	5.1	6.1	7.6	8.0
200I-10N-U-12H	-6.2	-4.0	-2.7	-0.9	-0.2	2.8	3.8	6.1	7.3	8.9	10.0	12.3
400I-10N-U-12H	-8.2	-6.4	-3.3	-1.3	-0.4	0.6	1.6	3.4	6.2	6.8	9.0	10.9
400I-20N-U-12H	-19.7	-13.2	-12.1	-6.4	-3.4	-0.6	4.8	9.3	11.7	14.3	17.9	21.1
200I-5N-R-12H	-6.0	0.0	1.2	6.9	9.8	10.9	13.3	15.3	16.2	17.5	21.9	25.0
200I-10N-R-12H	-16.2	-9.3	-6.1	-0.9	4.3	12.0	17.0	20.0	26.0	28.9	33.1	39.4
400I-10N-R-12H	-11.2	1.4	10.2	13.1	19.9	22.1	24.9	31.8	39.8	56.4	57.6	58.7
400I-20N-R-12H	-37.0	-22.0	-14.9	-2.5	6.0	16.1	25.1	35.6	42.3	50.9	58.0	71.1

Table 26 - Magnitude of weekly savings in hours WBW v/s. DLTP – Growing Demand

#	Weeks											
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**
200I-5N-U-8H	2.1	4.3	5.2	5.9	6.6	9.9	10.3	11.7	-	-	-	-
200I-10N-U-8H	-1.3	0.6	3.5	7.4	12.1	14.4	17.4	23.4	-	-	-	-
200I-5N-R-8H	4.9	12.4	18.2	19.0	23.6	31.3	33.3	40.9	-	-	-	-
200I-10N-R-8H	-2.3	9.6	23.3	28.7	35.9	54.1	62.5	74.1	-	-	-	-
200I-5N-U-12H	-2.0	0.2	1.7	3.9	5.6	7.4	7.5	8.5	10.7	11.9	12.6	12.8
200I-10N-U-12H	-9.9	-4.3	-1.1	4.2	6.2	9.5	11.0	14.4	16.0	21.8	26.8	30.9
200I-5N-R-12H	-7.0	-3.5	-2.1	1.6	8.4	11.2	18.9	28.7	31.1	38.7	43.1	50.5
200I-10N-R-12H	-8.1	0.6	8.3	16.8	25.9	36.5	40.3	51.6	60.5	70.0	76.7	83.0

Table 27- Magnitude of weekly savings in hours WBW v/s. DLTP – Steady Demand

#	Weeks											
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**
200I-5N-U-8H	3.6	3.7	3.5	2.7	3.2	4.3	3.5	4.2	-	-	-	-
200I-10N-U-8H	2.6	1.9	1.7	7.0	2.8	1.5	1.8	4.8	-	-	-	-
200I-5N-R-8H	13.9	11.8	12.0	11.9	12.5	15.6	16.9	23.5	-	-	-	-
200I-10N-R-8H	24.5	24.2	20.2	16.7	17.8	22.9	26.0	33.6	-	-	-	-
200I-5N-U-12H	4.0	3.4	3.2	2.7	1.8	1.5	1.2	1.2	2.4	3.4	3.8	3.9
200I-10N-U-12H	3.5	3.1	2.0	3.6	2.0	3.5	3.3	4.9	5.1	5.8	9.0	12.3
200I-5N-R-12H	15.7	14.4	14.1	12.1	6.9	7.9	6.8	9.0	8.8	12.9	15.4	22.0
200I-10N-R-12H	18.1	16.5	15.9	12.6	13.6	9.8	9.6	9.4	9.4	14.5	19.8	24.7

Table 28 - Magnitude of weekly savings in hours WBW v/s. DLTP (UI) – Steady Demand