5-1-2011

The Vehicle routing problem with a volunteer workforce

Leonardo Gala

The Vehicle Routing Problem with a Volunteer Workforce

By

Leonardo Mark Gala Jr.

A Thesis Submitted in Partial Fulfillment of the Requirements for the Degree of Master of Science in Industrial Engineering

Supervised by

Professor Dr. Mike Hewitt

Department of Industrial and Systems Engineering

Kate Gleason College of Engineering

Rochester Institute of Technology

Rochester, New York

May 2011

Approved by:

____

Dr. Mike Hewitt, Professor

Thesis Advisor, Department of Industrial and Systems Engineering

____

Dr. Ruben Proano, Professor

Committee Member, Department of Industrial and Systems Engineering

Thesis Release Permission Form

Rochester Institute of Technology

Kate Gleason College of Engineering

Title:

The Vehicle Routing Problem with a Volunteer Workforce

I, Leonardo Mark Gala Jr., hereby grant permission to the Wallace Memorial Library to reproduce my thesis in whole or part.

Leonardo Mark Gala Jr.

Date

Dedication

I dedicate the work presented in this research to my parents Iris Pons-Gala and Leonardo Gala Sr. Without their unwavering belief in me, this project would not have been started, let alone finished. A lesson I am glad to have learned from them is that you have to work hard for what you want. It makes reaching your goals that much more satisfying.

Acknowledgements

I have to thank my thesis advisers: Dr. Mike Hewitt and Dr. Ruben Proano. Dr. Hewitt's constant support and enthusiasm throughout this project was something not replicable. Without his guidance, I would have refrained from a graduate degree oriented towards research and optimization all together. I have learned plenty of skills in optimization and software development from him that will help me in my future endeavors, but the skills I will cherish the most are what to do when things do not go exactly to plan. Dr. Proano's wealth of operations research and optimization knowledge was extremely helpful in orienting me throughout this project. What I am most thankful for, though, is his ability to find the strengths and weaknesses of your work and how it is presented.

Abstract

The Vehicle Routing Problem with a Volunteer Workforce

Leonardo Mark Gala Jr.

Supervising Professor Dr. Mike Hewitt

Non-profit organizations like the Meals on Wheels association of America rely on a volunteer workforce to prepare and deliver meals to approximately one million homebound citizens nation-wide. At the community level, hundreds of volunteers are routed through rural, sub-urban, and urban sectors daily. These communities can benefit from optimization techniques that effectively route volunteers. Lack of volunteer availability requires Meals on Wheels to maintain a waiting list for people who require meals but cannot be incorporated into the current delivery schedule. The consistency of delivery routes is also of concern, as there are service and operational benefits gained when volunteers develop meaningful relationships with the people they serve. This research focuses on optimizing a Vehicle Routing Problem where efficient routing, meeting all demand, and consistent assignments are valuable. The three competing goals are aggregated into a single weighted function. A Tabu Search heuristic with variable neighborhood structures is then applied to solve the problem. Analysis is presented on each weight's impact on the competing objectives. The Tabu Search heuristic is bench marked against a current leading paper in consistent vehicle routing with comparable results. Finally, a large-scale instance similar in size to those serviced by Meals on Wheels is solved.

Contents

Algorithms

Tables

Figures

1 Introduction

On a daily basis, non-profit organizations like the Meals on Wheels Association of America (MOW) deliver approximately one million meals throughout communities nationwide. Within each of these communities, Americans aged 60 and older rely on government funded programs like MOW to meet their dietary needs for sustaining a healthy life style. Due to economic or health limitations, these elderly citizens do not have the means to provide or prepare meals for themselves. In addition to the aging population, MOW services individuals who are incapable of sustaining themselves due to medical limitations. The assistance MOW provides enables their clients to remain comfortable in their homes instead of requiring them to relocate to subsidized housing or nursing homes, either at personal or government expense. However, what makes this organization unique from commercial delivery or food service programs is that the MOW workforce is comprised of dedicated and caring volunteers from each community that hand deliver the meals. MOW routes between 800,000 and 1.2 million volunteers nationwide with the collective goal of ending senior hunger (Meals On Wheels Association of America 2010).

The volunteer fleet that supports MOW both strengthens the effect the program has on each community and makes operational efforts difficult. As volunteers serve the community, they become more adept at navigating their routes and create relationships with the MOW cliental on a personal level. More often than not, the volunteers transcend the role of delivery driver and develop friendships and adopt feelings of concern for their clients well being. For example, a close relationship between volunteer and elderly client lead to the aversion of a major medical emergency, as volunteer Anne Larky was able to interpret unusual behavior in one of her clients- prompting hospitalization and the resulting prevention of permanent brain damage

(Devlin 2010). However, volunteers have limited availability throughout the week to service routes due to work and personal schedules. This fluctuation in availability not only makes meeting daily demand difficult, but decreases the likelihood of consistent service for all clients.

Unlike private delivery organizations, MOW does not have the resources to hire a fleet of dedicated delivery drivers. With the national demand for meals increasing by 290% from 1980 to 2002, funding and resource management for food preparation have been of significant importance for each county home to MOW (O'Shaughnessy 2004). Thanks to contributions from private organizations and cost reduction efforts, dropping overhead to 16% of the total budget, in 2009 alone, the program was able to serve thousands more meals nationally than the previous year (Meals On Wheels Association of America 2010). Finances are not the only scarce resource to the organization though, as without their volunteer support, meals could not be delivered. Although financial management has lead to the continued success of MOW, there can be greater benefits realized with operational improvements focused on workforce management and volunteer routing.

Like major parcel delivery companies such as FedEx and UPS, MOW shares the common problem of making home deliveries in large sized communities, containing demand at varied locations, and with delivery routes limited by time and vehicle capacity. Although industry leaders in parcel delivery have the engineering and financial resources to invest in technology that handles these routing scenarios, many community based MOW centers do not. However, this research will show that the adaptation of some readily available optimization algorithms can improve MOW current routing schemes while balancing workforce goals at minor cost of implementation.

Logistics issues, regarding efficient routing and scheduling of delivery vehicles pertaining to home delivery, at the commercial level can be handled with mathematical optimization and Operations Research techniques. In general, optimization seeks to find the best alternative from a group of many choices. This best alternative maximizes or minimizes the value of a objective function- in the case of vehicle routing that function may pertain to travel cost (Rardin 1998). The possible choices that can be made are generally unique to each scenario, as a set of constraints such as operational hours in a day, or vehicle capacity, limit the decisions that can be made. How the best decisions are arrived at depends on the algorithm being applied to solve the problem.

The problem of local delivery can be formulated as the Vehicle Routing Problem (VRP). The VRP is a generalization of the Traveling salesman problem, where a single vehicle must be routed through a network of cities. The route taken by the driver must service every city exactly once, and return the vehicle to its origin, while minimizing travel distance (Miliotis 1976). Variations of the TSP can involve routing several vehicles from a single origin, commonly referred to as a depot. One step further is to define a problem with several depots, each of which are assigned a group of drivers. Other variants of the VRP add more real-world constraints and features. A natural extension assigns capacity to the vehicles, as delivery trucks can only carry so many items while servicing routes. This changes the definition of a location's demand from a simple visit to a quantity of resources the vehicle must supply. Planning periods can also be introduced, where demand varies over several days or weeks, and routes must be designed to accommodate the varying requests. Modeling parcel service industries requires that vehicles not only deliver goods, but pickup items along the route, adding another layer of complexity to vehicle capacity management (Laporte 1992). The VRP can also contain both soft and hard time

deadlines, or time-windows during by which a delivery must be made. These time constraints are used to model scenarios where customers have predefined appointments, frequently used in home health care, or when immediate deliveries must be made like in emergencies (Steeg and Schröder 2008).

Once a problem definition is completed, including an objective function and operational constraints, an optimization method is then applied to construct solutions that represent vehicle routes. When the VRP is formulated as an Integer Program, exact algorithms and commercial solvers can be applied to find the best solution possible- also referred to as the optimal solution. However, variants of the VRP belong to a family of problems referred to as NP-Hard and suffer from algorithm run times that grow exponentially with the size of the problem (Bazgan, Hassin and Monnot 2005). Therefore, these methods can only solve instances much smaller than real MOW problems in a reasonable length of time.

Given the impractical run-times of exact methods, and the demand for fast and flexible algorithms from commercial industries, many heuristic approaches have been devised to solve the VRP(Gendreau, Potvin, et al. 2008). A heuristic is an algorithm or method that typically generates good solutions to a problem, but does not guarantee the optimum. The tradeoff for quality of solution is that solutions can be generated in significantly less time than exact methods, removing the ceiling on instance size considerations. This is accomplished by intelligently searching portions of all possible solutions to a problem.

This thesis will present a heuristic for the variant of the VRP related to MOW daily operations. Features that are unique to the work presented here include optimization methods traditionally used in Inventory VRPs to consolidate deliveries that are integrated into a problem oriented towards service industry. Additionally, the use of variable neighborhood structures

create a thorough search process that not only optimize routing, but manage client demand and delivery schedules. Finally, combined objectives of route time, consistent service, and wait-list reduction not usually studied together are integrated into an aggregated objective function.

The remainder of this paper is organized as follows: first, a literature review organizes pertinent research on topics including vehicle routing, workforce management, and optimization techniques. Next, a formal problem statement is presented, detailing the problem specific constraints and objective function. An adapted Tabu Search method is then proposed to solve the problem by generating delivery routes and entire planning horizon, with search methods and parameters detailed. The experiments performed to optimize the Tabu Search's performance are also described. Benchmark tests versus a comparable paper on consistent Vehicle Routing and MOW sized instances are then presented. Finally, experimental analysis is presented and future work is suggested.

2 Literature Review

The Traveling Salesman Problem is described as finding the shortest tour a salesman takes by leaving their home, visiting n required destinations exactly once, and returning to their home. To find the shortest possible tour through these destinations by enumeration would require $(n-1)!$ evaluations, driving the need for analytical approaches to the TSP (DANTZIG, FULKERSON and JOHNSON 1954). Given the restriction on computational resources, early methods of analysis were developed for manual calculations that solved "large" instances of up to $n = 50$ locations, (Flood 1956) –though obtaining optimum solutions were not guaranteed.

Thanks to modern computational power and sophisticated optimization techniques, TSPs with large sizes of n (upwards of 85,900) can be solved to optimality with acceptable computation times (Applegate, et al. 2009). Unfortunately, real-world applications require more problem specific constraints than seen in the TSP, potentially making the problems harder to solve. An intuitive extension of the TSP is the Vehicle Routing Problem, which attempts to find the shortest tour for multiple vehicles over a shared group of destinations (Laporte 1992). Common requirements of the VRP include that: each location may only be visited by a single vehicle, a vehicle starts and ends at a single location (just as in the TSP), and additional operational constraints. Common operational constraints arise from realistic issues related to the capacity of each vehicle, the total time a route takes to complete, delivery time-windows or scheduled appointments, multiple depots to route vehicles from, planning horizons with several days of varying demand, and any additional delivery precedence not related to travel time.

Vehicle capacity is a limiting factor in many VRP applications, as it is natural to assume that delivery trucks can only carry a finite number of parcels per trip. This constraint is also seen

as a target in the Inventory Routing Problem (IRP), as higher utilization of vehicles reduces the fixed transportation costs. In general, Inventory routing can be described as an adaptation to the VRP in which customers are generally resupplied with products continuously over a planning horizon, and the intervals at which deliveries are made and the quantities delivered are decided by the supplier (Schwarz, Ward and Zhai 2006). When solving IRPs, the use of product consolidation algorithms has been suggested to increase utilization and lower delivery cost (Karabuk 2005). IRP and general parcel delivery problems may also include multiple depots, which in practice lowers the travel time to customers in the region of the depot.

Given the large number of operational constraints that apply to the VRP, and all of their possible combinations to define a specific problem, many approaches have been developed to solve these problems. Since the VRP is NP-hard, the issue of instance size solvability and computational time still apply, and a variety of solution techniques have been developed that produce optimal or near optimal solutions in reasonable amounts times. Direct tree search methods and branching techniques have been applied to derive exact solutions. Laporte et al. (1986), utilize a special case of the TSP with m different routes as a relaxation of the VRP with m vehicles in a branch and bound algorithm. This methodology allowed instances containing up to 260 locations to be solved. Fisher (1994) implemented an approach using K-trees to route a fleet of k vehicles from a single depot. Lagrangian relaxations combined with constraint generation allowed for problems including up to 100 locations to be solved to optimality.

Many authors have focused their efforts on developing heuristics to solve the VRP, forgoing the guarantee of optimality and gaining the benefit of short computational times and comparably good solutions. An early attempt at the VRP containing a single depot and capacitated vehicles applied a linear relaxation to the common integer formulation to find near

optimum solutions (Dantzig and Ramser 1959). A correction algorithm is used to swap fractional decision variables with binary ones based on a relative route cost function. Clarke and Wright (1964) propose an algorithm for routing multiple vehicles from a single depot by first starting with as many vehicles as there are locations. The algorithm then uses arc deletion and route merging to decrease the number of vehicles based on a savings function. Fisher and Jaikumar (1981) proposed a heuristic solution for the VRP with time windows that combines the General Assignment Problem with travel costs producing good solutions. A master problem assigns locations to each vehicle, and then TSP sub problems are solved iteratively.

Recent work has focused on Metaheuristics, which are powerful search processes that can be to many optimization problems. Categories of metaheuristics include advanced local search, population search, and learning based algorithms (J.-F. Cordeau, et al. 2005). Local search methods look to continually improve solutions by making a single, or a small number, of changes the current solution. Simulated annealing is an adapted local search algorithm that allows degrading, probabilistic changes in the current solution earlier on in the search while a "temperature function" T has a higher value, and focuses the search in a more restricted region later on when T has become smaller (Gendreau, Potvin, et al. 2008). Variable Neighborhood Search (VNS) is another type of local search where multiple types of solution modifications are applied to the current solution at each iteration (Gendreau, Potvin, et al. 2008). For example, a modification of a current solution could be to swap vertices between two vehicle routes, and the process of attempting each possible swap creates a neighborhood of solutions to the original. VNS takes neighborhood search a step further by applying several modifications per iteration until pre-determined time or iteration limits are met (Hansen and Mladenovic 2001). This type of

search is advantageous to the research in this paper as it allows for the consideration of several actions, which can assist the decision making process when routing with several goals in mind.

Population based search tends to look at creating improved solutions by modifying a pool of current solutions referred to as the population (J.-F. Cordeau, et al. 2005). Genetic Algorithms are a popular type of population-based heuristic, that modify solutions populations, generating "offspring", with the goal of finding optimal solutions. Offspring are created by mixing components from solutions found in the populations and piecing them together while slight changes to the offspring solutions are performed stochastically to add diversification to the search. Learning based procedures, on the other hand, explore various solutions and use updating probability distributions to rank which regions may have the most promising solutions(Gendreau, Potvin, et al. 2008). Ant-colony optimization is a learning based algorithm that replicates the natural technique ants use to locate food sources, where certain paths taken by ants are traveled more frequently the more fruitful they are at producing food. This is replicated by assigning probabilities of finding fruitful regions- as the search produces improving solutions, higher probabilities are assigned to choosing that region for further search.

A popular metaheuristic applied to the VRP with large degrees of computational success compared to other metaheuristics and exact methods is the Tabu Search (TS) method, and is the main focus of this research. TS combines local neighborhood search and diversification methods to examine a large portion of a problem (Glover 1989). Given an initial solution to a problem s, the Tabu Search will modify s with the hopes of gaining an improved solution. The collection of all possible modifications of s can define the neighborhood $N(s)$. The best neighbor in $N(s)$ is chosen as the subsequent solution for the following iteration, regardless if the solution improves or degrades the solution. This allows the search process to explore more of the solution space, as

it does not terminate at local optima, or potentially find solutions that are not necessarily feasible. As the search iterates, the best overall solution is stored.

To search several regions of the problem-space diversification techniques can be added to the algorithm. These generally make the search process take larger steps away from the current solution with the hopes of finding improving solutions. In addition, the acceptance of degrading allows for cycling between a small set of neighbors. To help prevent this, a list of recently completed modifications is stored for several iterations. This list, often known as the Tabu List, prevents the use of modifications recently applied during the search. However, aspiration criteria can override tabu status if undoing the tabu modification leads to an improving solution. The TS algorithm usually terminates after a predetermined number of iterations have passed, or a specified time length has elapsed.

Given the general framework of TS, problem specific encoding is needed to apply the method to the VRP, including the objective function, the derivation of the neighborhoods, and the route length approximations. Cordeau et al. (1995) present a TS algorithm suited for both the multi-depot VRP and the periodic VRP with the traditional objective function of minimizing total travel time. The neighborhood solutions are constructed by either performing a swap between routes on the same day, or change the daily delivery patterns assigned a customer. The insertions are accomplished by using the GENI algorithm also presented by the authors (Gendreau, Hertz and Laporte 1992).

The diversification techniques used include penalizing frequent use of certain modifications and allowing interim infeasible solutions regarding vehicle capacity and route length. The length of the tabu list is set to a fixed number of iterations proportionate to the number of vertices in the problem and the number of days in the planning horizon. A follow up

paper deals with a heterogeneous fleet of vehicles, with the main algorithm differences being the use of a tabu list with a random-uniform length and the neighborhood building based on only a finite number of randomly selected vertices, helping limit the required computational efforts. An adaptive memory algorithm is also introduced to save and reuse aspects of attractive routes (Gendreau, Laporte and Musaraganyi, et al. 1999).

Additional work has been done by Rochat and Taillard (1995) to study the effects of adding variability to the search process. When proper probability distributions are applied to parameters, like the duration of tabu status, or the diversification process, the effectiveness of the TS increases. The authors do warn, though, that blind randomization can do more harm than good in some cases.

Another successful implementation of Tabu Search for the VRP was done by Parthanadee and Logandran (2006), who solved a periodic, multi-depot, mixed product distribution problem with constrained resources. The authors compared three different Tabu Search strategies, only using a short term Tabu List, using a Tabu List with solution diversification, and using a Tabu List with intensification only, all of which performed to near optimality consistently. Like the Gendreau application, neighborhoods were constructed by altering delivery schedules, but another modification investigated involved swapping depots associated with routes. Wassan (2006) proposes a TS algorithm for the VRP that utilizes two forms of vertex swapping to generate neighbors, but also utilizes a memory technique called Hashing Function search, which identifies the revisiting of routes. When this occurs frequently, a drastic neighborhood change is made by restarting the search process.

Traditional studies of the TSP and VRP measure the effectiveness of a solution with one objective- typically the travel time associated with the solution. It has already been seen in fields

such as home health care or parcel delivery though, that there are many conflicting objectives in real applications such as consistent service or workload balancing, which may not all be addressed with a single objective. Addressing these issues fall into the field of Multi-Objective Optimization (MOO), where several operational goals are integrated into a single objective function, or optimization techniques are adapted to address several goals simultaneously. Two of the more popular approaches to MOO are to adopt an aggregated function which weights the conflicting objectives by relevance to the problem, or through Pareto Multi-Objective optimization. The aggregated function is beneficial when the decision making process has clear priorities in terms of the importance of each objective. However, Pareto MO optimization avoids a priori weighting by generating many acceptable trade-off optimal solutions. The set of these solutions is known as the Pareto Front, which gives decision makers the ability to chose from several good solutions (Ngatchou, Zarei and El-Sharkawi 2005).

An example of a linearized MOO problem was presented by Bowerman et al. (1995). where a VRP was solved in the field of School Bus routing. The problem contained four conflicting objectives: overall bus route length, the walking distance of students from their bus stop to their home, route length balance among all busses, and travel distance balance among all busses. A heuristic solution was applied that places students in districts and meets their demand using set covering algorithm. The author goes on to briefly discuss parameter weighting and its effects on the solution outcomes with specific analysis on the intervals where the specific objectives influence the solutions generated. Smilowitz et al. (2011) also use a linearized MOO function in their VRP to combine route time goals with proposed workforce management objectives. The problem is solved using a TS adapted from Cordeau et al. (1995), which was also

adopted in this research. These author also described an experimentation process in which they evaluated appropriate values for weighting the objectives.

Unlike the linearized objective functions, search processes using Pareto Optimization do not need a prior knowledge of which goal is more important. Kulturel-Konak et al. (2006). present a Pareto Optimization approach that is integrated into a TS algorithm for solving several NP-hard problems. At each iteration, an objective function is selected to evaluate the current solution based on a probability distribution. The search iterates similar to traditional TS procedures, but a list of minimal solutions to all objectives and solution spaces is maintained to compile a Pareto front, or best set of solutions found. Another sophisticated application of Pareto Optimization was implemented by Figueira et al (2010). that seeds several points in the Pareto front. Then, master and slave algorithms search the frontier at several places simultaneously. Several NP-hard problems were solved with this algorithm, including a job scheduling problem that looked to minimize tardiness and the makespan. Though these techniques are effective in analyzing multi-objective problems, the research presented here favored the linearized objective function as the ability to explicitly weight certain objectives would allow the methodology presented to be applied under different management schemes.

3 Problem Statement

3.1 MOW Operations

MOW services a large number of clients on a daily basis. Within each community, the demand for meals is scattered over rural, suburban, and densely populated urban settings. Servicing these various geographical regions makes efficient routing a major logistics issue, as routes need to accommodate both clustered and farther spread out delivery locations. To make servicing the different territories easier, MOW often sets up satellite a depot location in several regions. Volunteers pick up a supply of meals and their daily route information from these depot locations, leave to service the route, and return the meal containers to the depot each day. The capacity of the meal containers puts an upper bound on the number of meals a volunteer can deliver each day. Due to size restrictions, volunteers are typically assigned one meal container per route.

Depending on their needs, clients of MOW receive deliveries any number of days per week. To lower operational costs, MOW has limited deliveries to weekdays. Deliveries only take place once per day, but may include enough food for multiple meals, sustaining clients over several days. For instance, clients requiring meal service over the weekends will generally have multiple meals delivered per day towards the end of the week, with the excess expected to cover weekend meal requirements. MOW accomplishes these consolidated deliveries by supplying their clients with both hot and cold meals. Though the hot meals need to be consumed soon after delivery, the cold meals are frozen and can be saved for re-heating at a later time. This lets MOW provide more versatile delivery options, and allows for potential efficiency gains (King 2005).

Clients receiving meals do not have an appointment or scheduled time slot specified for deliveries each day. However, they are given a time frame of approximately one hour in which they can expect deliveries- allowing for flexible route sequences each day. There is an upper bound on the duration of a route due to the volunteer's availability, and it can be as little as one hour to several hours. Volunteers may not be available every day of the week, meaning clients may have more than one volunteer that services them throughout the week. In practice, all volunteers will start their routes at the same time to ensure meals are delivered around noon. Though the time frame when deliveries take place is more affected by volunteer availability (since a significant portion of volunteers dedicate their lunch-breaks from a full time job), delivering during the same time frame each day has the benefit of giving the MOW clients consistency in their personal schedules.

Factors such as age, economic and health status of people within the community, or the availability of support from their family and friends play a role in the need for service from MOW (O'Shaughnessy 2004). However, if there is demand that exceeds the capacity or funding of MOW, a waiting list is formed for clients to remain on until MOW can serve them. Proper attention must be taken into account while building routes with regards to the management of this waiting list, as proper utilization of volunteer resources can reduce the number of clients on the wait list.

The goal of MOW is to provide as many hungry, homebound clients with meals as possible. With all travel (in terms of fuel or vehicle maintenance) affecting only the volunteers that help MOW, a major performance measurement of the delivery routes generated is the number of clients that reside on the waiting list. This measurement is easily translated from MOW mission statement to "end senior hunger (Meals On Wheels Association of America

2010).” However, there is a clear relation between the number of clients being served and the quality of routes constructed, with more effective routes servicing as many clients as volunteer availability or vehicle capacity would allow.

Another concern of MOW is the consistency of volunteer routing and the relationships developed with clients serviced (Howard County Association of Volunteer Administrators 2010). Volunteers act as the main channel of information between the clients and the agency. According to executive director of MOW Central Maryland Tom Grazio, the volunteers provide “much needed information not only for MOW but also in offering additional support services to clients in the form of social contact and interaction.” Because of this, volunteers are seen as a means of measuring how effective their organization is operating, and with consistent service, the quality of this information and the service provided will increase. As an external comparison, a computational study on a VRP related to parcel delivery performed by UPS cites operational and service benefits from providing consistent service, though the authors here look at both consistent driver assignments and delivery times (Groër, Golden and Wasil 2009). This study also notes the positive effects of drivers who “take ownership” of their routes and develop relationships with their customers- something similarly noted for MOW.

4 The Tabu Search Heuristic for Vehicle Routing with a Volunteer Workforce (TS-VRVW)

To address the issue of instance size versus solution time, a Tabu Search heuristic is proposed to solve problems with varying planning horizons, volunteer fleet size, and client pools (TS-VRVW). The algorithm combines techniques from algorithms for Periodic Vehicle Routing and Multi-Depot Vehicle Routing to address the route planning needs of Meals on Wheels. Adapted from the work of Cordeau et al. (1995), this Tabu Search algorithm combines local neighborhood search with traditional cycle prevention techniques to identify improving solutions over several iterations. The main adaptations to the TS algorithm occur in how each solution generated is evaluated, and in the solution generation process itself. To incorporate the goals of volunteer specific workforce management as well as the operational specific service goals, a weighted objective function is adopted in lieu of minimal route time goals. This objective function will seek to:

- minimize overall route time,
- minimize the number of clients seen by more than one volunteer,
- and minimize the number of wait-listed clients.

Finally, in addition to neighborhood building based on route swapping and delivery schedule modification, techniques that further alter routes and manage waiting lists are integrated.

4.1 Problem Parameters

The problem can be described as a graph with a vertex set $\boldsymbol{V} = \{v_1, v_2, \dots, v_{C_{total}}\}$ where there are C_{total} locations. Vertices v_1 through $v_{De_{max}}$ represent the set of depots **De**, where $De_{max} \geq 1$, and the client location set **C** is represented with vertices $v_{(De_{max}+1)}$ through $v_{C_{total}}$. The arc set Arc $= \{(v_i, v_j)\}$ between each pair of vertices *i* and *j* on the graph represent the travel time required to traverse from *i* to *j*. The travel time is dependent on the arc and is defined a $tt_{i,v}$. Additionally, each vertex *i* in the client set **C** has an associated service time $\mathrm{st_i}$ that represents the time required to deliver meals and provide general service to client *i*. The problem defines V_{total} as the number of volunteers that can service the client vertices over D_{max} days in a planning horizon. The maximum time spent traveling a delivery route and servicing clients is limited by volunteer *v*'s availability $A_{v,d}$ on day *d*, and the number of meals they can carry is restricted to their capacity C_v due to the meal containers.

Each volunteer services at most one route on each day where $A_{v,d}$ is greater than zero. A route is defined as a departure from a single depot in the set of depot vertices that services one or more client vertices *i* with positive demand $CD_{i,d}$ on day *d*. Each route concludes with a return trip to the depot of origin from the last client vertex serviced. A client can only belong to a single route on each day that service is required; if client i is not placed on a route when $CD_{i,d}$ is positive, the service instance is placed in the wait-list set Wl. Finally, the overall goal is to minimize the route time associated with all routes in the problem, volunteer inconsistency that occurs throughout the planning horizon, and the number of wait-listed clients.

4.2 Solution Description and Evaluation

The TS-VRVW represents each solution *s* visited throughout the search with the attribute set $B(s)$. This attribute set contains information used to build routes based on each client-volunteer-day relationship (*i,v,d*) within the solution and is expressed as $B(s) = \{(i, v, d): \forall\ i \in C \cup De, v \in V, d \in D\}$. Each attribute relationship can be interpreted as location *i* is visited by volunteer *v* on day *d*. An entire route serviced by a volunteer *v* on day *d* will be represented *(v,d)*. These relationships traditionally define all locations, vehicles, and days considered on the delivery routes, but must also incorporate clients residing on the Waiting List. Since these un-serviced clients do not have a volunteer, their attribute relationship is defined (i, Wl, d)as.

The route relationship set B(*s*) is independent of any specific ordering in terms of what sequence a volunteer services their assigned clients. A specific route order is defined a $RO_{v,d} = \{v_{De}, v_i, v_j, \dots, v_{De}\}$, where volunteer *v* services the vertices $v_{De}, v_i, v_j, \dots, v_{De}$ in order on day *d*. This route ordering can then be used to calculate the time associated with servicing the route (*v,d*) as $RT_{v,d} = \sum_{r \epsilon RO_{v,d}} tt_{r \in RO_{v,d},(r+1) \in RO_{v,d}} + st_{r \in RO_{v,d}}$. It is assumed that the service time at any depot is negligible and equal to zero.

Each solution *s* is evaluated by an aggregated objective function and its components:

1. $F(s) = \underbrace{RT(s)}_{Route\ Time} + \left(\underbrace{M_1 * INC(s)}_{Inconsistent\ Service}\right) + \left(\underbrace{M_2 * WL(s)}_{Wait\ Listing}\right) + \left(\underbrace{\alpha * CAP(s)}_{Capacity\ Overage}\right) + \left(\underbrace{\beta * OT(s)}_{Route\ Length\ Overage}\right),$

2. $RT(s) = \sum_{(\boldsymbol{v},\boldsymbol{d})} RT_{\boldsymbol{v},\boldsymbol{d}},$

3. $INC(s) = \sum_{iC} \sum_{v \in V} \sum_{d \in D} E_{i,v,d}$,

4. $WL(s) = \sum_{i \in WL} \sum_{d \in D} WLD_{i,d}$,

5. $CAP(s) = \sum_{(v,d) \in B(s)} max\{0, \sum_{(i,v,d) \in B(s)} CD_{i,d} - C_v\}$,

6. $OT(s) = \sum_{(v,d) \in B(s)} max\{0, RT_{v,d} - A_{v,d}\}$.

The solution found with the minimal objective found is denoted s *, with corresponding objective value $F(s^*)$. In (2) RT is defined as the sum of all travel time for each volunteer route (v,d) over the entire planning horizon. In (5) CAP is defined as the sum of all capacity violations for each volunteer route (v,d), and in (6) OT is the cumulative route length violation of all volunteer routes (v,d). Any solutions that contain CAP or OT with values larger than zero are considered infeasible and, although allowed as interim search steps, will not be accepted as s^*. To steer the search process away from the interim infeasible solutions, the sufficiently large weights α and β make moving away from infeasible solutions attractive.

The terms INC and WL are incorporated into (1) to penalize solutions containing inconsistent service and wait-listed clients. Although it may seem attractive to have these terms equate to the number of clients on the wait list, or the number of clients who are serviced by more than one volunteer, doing so would require mixing units within the objective function. Therefore, approximations were developed that force the tabu search to address these issues while maintaining continuity of units in the objective function at the expense of operational assumptions.

In (3), the term INC is an approximation used to lower the number of inconsistent volunteer-client pairings within a solution. A solution contains inconsistent service when more than one volunteer serves a client in the planning horizon. To address this, each volunteer is

assigned a penalty for providing inconsistent service proportionate to the amount of effort required to service the client. The effort for volunteer *v* to serve client *i* on day *d* is quantified as

$$7.\, E_{i,v,d} = \left(tt_{(r-1)\in RO_{v,d}, r\in RO_{v,d}} + tt_{r\in RO_{v,d},(r+1)\in RO_{v,d}} + st_{r\in RO_{v,d}}\right) * \left(1 - \left(\frac{VV_{i,v}}{TV_i}\right)\right)$$

where *r* is the position of client *i* on route (*v*,*d*), and TV_i is defined as the total number of visits client *i* receives in a planning horizon, while $VV_{i,v}$ is the total number of visits client *i* receives from volunteer *v*. The effort term captures the time required for a volunteer to travel to and from, and service an inconsistently seen client. Figure 1 shows an example of inconsistent service provided to client *i*, where Volunteer 1 is penalized with an additional effort of $\left(tt_{3,i} + tt_{i,d} + st_i\right) * \left(1 - \left(\frac{1}{2}\right)\right)$. Similarly, Volunteer 2 is penalized with $\left(tt_{5,i} + tt_{i,6} + st_i\right) * \left(1 - \left(\frac{1}{2}\right)\right)$.

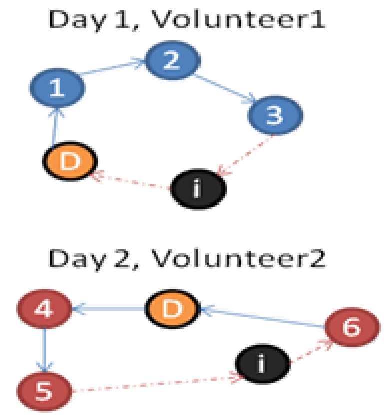

Figure 1: Example of inconsistent service to client i.

In (4), the second term WL looks to give incentive to servicing clients who are currently on the Waiting List. To penalize leaving clients without service, the travel time associated with an out and back delivery between the central depot and any wait-listed client is defined as

$$8.\, WLD_{i,d} = tt_{De,i} + tt_{i,De}.$$

Since many of the clients served by Meals on Wheels rely on the service as a main nutritional supplement, this approximation assumes that in a worst-case scenario (reflective of clients in poor health condition) an emergency meal would be delivered to the wait-listed client, regardless of the volunteer fleet capacity. This type of penalty is

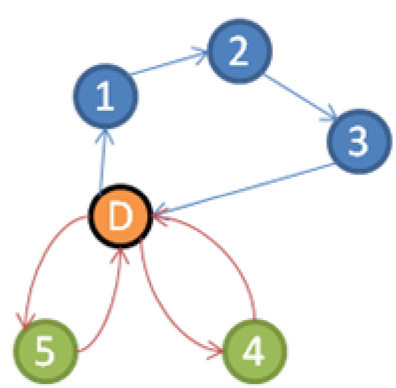

Figure 2: Example of wait listed clients 4 and 5.

frequently used in stochastic VRP problems, where recourse penalties are applied to not meeting demand under stochastic conditions (Gendreau, Laporte and Séguin 1996). Figure 2 shows an example of wait listed clients 4 and 5.

The objective function evaluates each solution based on the travel time and the volunteer workforce management goals-which also rely on some routing information from the solution. However, determining this information requires solving a TSP for each vehicle, with the addition of VRP constraints, making evaluating a single solution excessively time consuming. Spending excessive time solving these TSPs, or even approximating them would make TS-VRVW algorithm ineffective and no more attractive than solving the problem exactly. Therefore, the route-building approximation algorithm GENI is used to construct sub-optimal routes with comparatively low travel time and minimal computational time compared to exact algorithms. GENI inserts new vertices on a route either between two consecutive vertices, or through methods that swap the current arc arrangements (Gendreau, Hertz and Laporte 1992). This method integrates 3 and 4-opt interchanges to optimize the route built. To limit the total number of vertices considered on a route, GENI uses a p-neighborhood, or a subset of all vertices including only the p-closest vertices, of each vertex. This allows the user to reduce the algorithm's choices from n^4 vertices for insertion points to a much smaller number.

4.3 Neighborhood Building

To generate a new solution at each search iteration, the TS-VRVW applies up to seven neighborhood structures to alter the current solution. The attribute relation set $\mathrm{B}(s)$ will have one or two relationships altered by each modifier to build the neighborhood of the current solution $\mathrm{N}(s)$. The result of applying a single modifier to s generates a solution s' which may improve or degrade the current solution, even at the expense of feasibility. Though not explicitly stated in each modification algorithm listed below, it is assumed that only modifications of the route relationship (*i,v,d*) where $CD_{i,d} > 0$ are considered. Similarly, only route changes are considered when $\mathrm{A}_{\mathrm{v,d}} > 0$. All vertex insertions or removals are handled with the GENI algorithm, which will optimally insert the vertex in the route- implying that $RO_{v,d}$ may change significantly with each use of GENI.

4.3.1 Solution Modifier 1: Intra-Route Swap

The first modifier applied looks to reduce a volunteer's route time by swapping the order in which two clients receive service on the same day. This modifier is designed to replicate the "simple" insertion used in GENI, where a new vertex is added to a route simply by placing a vertex by one of its p-neighbors (Gendreau, Hertz and Laporte 1992). Since both vertices are currently on the route, the modifier only considers switching the two if the resulting move will place both vertices by a p-neighbor. Figure 3 shows an example MOD1 generating a neighbor solution.

Algorithm 1: Intra Route Swap (MOD1)

1) For each combination (i,v,d) ∈ B($\boldsymbol{s}$) where i ∈ **De**∪**C**

   a) For each (j,v,d) ∈ B($\boldsymbol{s}$) where $j \neq i$ and j ∈ **De**∪**C**:

      i) Exchange the locations of i and j on route (v,d).

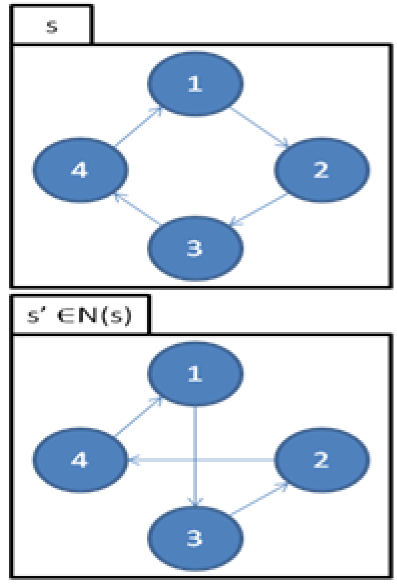

Figure 3: MOD 1 Example.

4.3.2 Solution Modifier 2: Depot Change

The next modifier looks to relocate the starting point of the route by exchanging the current depot assignment to another. Parthanadee and Logendran (2006) used this type of neighborhood scheme, where instead of changing the volunteer or route associated with a depot, clients were associated with different depots to alter clustering schemes. Though a seemingly drastic change, as the search process restructures the makeup of routes, the re-centering of the route may unveil shorter routing times. Although the GENI algorithm is used to insert the new depot on the route being modified, an extended neighborhood of $(Exn) * p$ is used when considering insertion points. Figure 4 shows an example MOD2 generating a neighbor solution by swapping Depot 1 and Depot 2.

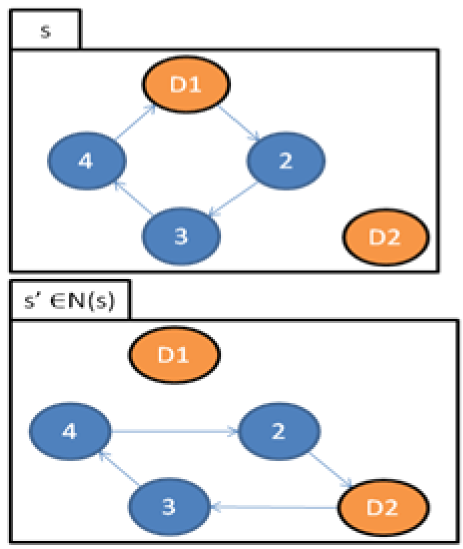

Figure 4: MOD 2 Example.

Algorithm 2: Depot Change (MOD2)

1) For each (i,v,d) ∈ B($\boldsymbol{s}$) where i ∈ **De**:

   a) For each j ∈ **De** where $j \neq i$:

i) Drop *i* from current route (*v*,*d*).

ii) Insert *j* on route (*v*,*d*) using the < GENI > algorithm.

4.3.3 Solution Modifier 3: Inter-Route Change

The third modification considered uses a more traditional VRP Tabu Search technique where a client is removed from their current route and is placed on another route on the same day with the goal of improving route time. Although used in many papers, Cordeau et al. (1995) and Barbarosoglu and Ozgur (1999) used this technique as one of their main neighborhood constructs. This insertion process uses GENI with the traditional p-neighborhood, but when volunteer workforce management goals are active in the objective function, can degrade or improve the volunteer consistency and route balance terms. Figure 5 shows MOD 3 generating a neighbor solution where client 2 is removed from the first route and placed on the second.

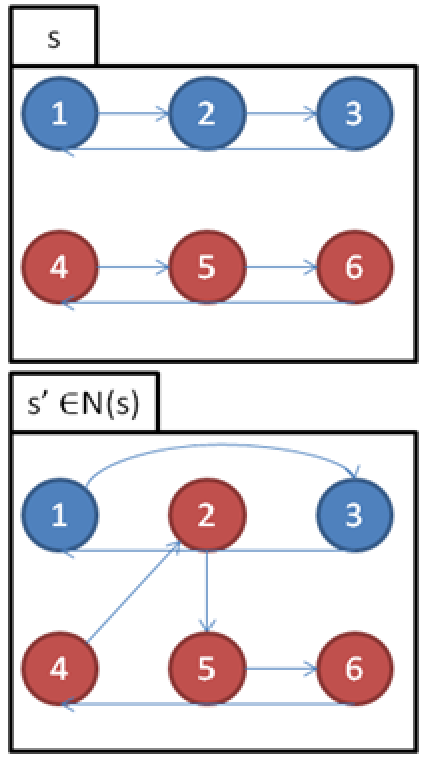

Figure 5: MOD 3 Example.

Algorithm 3: Inter Route Change (MOD3)

1) <u>For each (*i*,*v*,*d*) ∈B (***s***) where *i* ∈ **C**:</u>

   a) <u>For each $w \in \mathbf{V}$ where $w \neq v$:</u>

      i) Remove *i* from current route (*v*,*d*) and restructure the route using the < GENI > algorithm.

      ii) Insert *i* on route (*w*,*d*) using the < GENI > algorithm.

4.3.4 Solution Modifier 4: Drop Service to Client

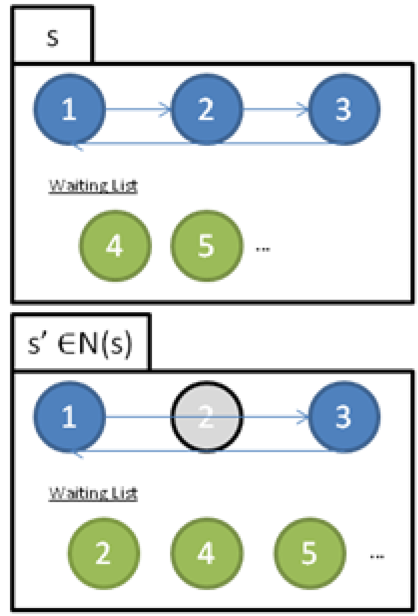

Figure 6: MOD 4 Example.

The fourth and fifth modifications primarily deal with wait list management- where client demand is either added or dropped from routes. A technique of adding or removing delivery service was used by Archetti et al. (2011) for managing inventory levels of customers in an IRP setting, where service is required to prevent stock-outs. However, the fourth modifier violates an operational standard held by MOW by dropping a client who is currently receiving services from a volunteer route on a single day. This modification can help remove the interim infeasible status of a solution, reduce route time, or reduce volunteer inconsistency. To prevent revoking service from clients, a solution will only be accepted as s^* if all client demand wait-listed in s was not met initially. This set of demand occurrences refers to the initial wait-listed demand in $\mathbf{Wl}^{\circ}$ defined previously. Additionally, if wait listing is not part of the objective there is nothing preventing the algorithm from dropping all clients from services, therefore in this situation only clients in $\mathbf{Wl}^{\circ}$ may be considered in this neighborhood structure. Figure 6 shows MOD 4 generating a neighbor solution where client 2 is placed on the waiting list.

Algorithm 4: Drop Service to Client (MOD4)

1) <u>If M2 > 0:</u>
    a) <u>For each $(i,v,d) \in$ B(s) where $i \in \mathbf{C}$:</u>
        i) Remove i from current route (v,d) and restructure the route using the < GENI > algorithm.

ii) Place (*i,Wl,d*) ∈ B(*s*).

Else

b) <u>For each (*i,v,d*) ∈ B(*s*) & ∈**Wl**° where *i* ∈ **C**:</u>

i) Remove *i* from current route (*v,d*) and restructure the route using the < GENI > algorithm.

ii) Place (*i,Wl,d*) ∈ B(*s*).

4.3.5 Solution Modifier Five: Add Client Service

The fifth solution modifier takes the opposite action of the fourth, and tries to place a client i from the Waiting List onto a route each day *d* where $CD_{i,d} > 0$. To increase the likelihood that a Wait Listed client will be serviced, the extended neighborhood of $(Exn) * p$ is applied while performing GENI insertions during modifier 5 insertions. This modification guarantees to reduce the term WL(s), but may adversely affect volunteer consistency, imbalance, and route time in a solution.

Algorithm 5: Add Client Service (MOD5)

1) <u>For each (*i,Wl,d*) ∈ B(*s*):</u>

a) <u>For each (*v,d*) ∈ B(*s*):</u>

i) Insert *i* on route (*v,d*) using the < GENI > algorithm.

4.3.6 Solution Modifier 6: Change Delivery Day

The sixth solution modifier restructures the current delivery schedule with the hopes of utilizing volunteers more efficiently. Authors Cordeau et al. (1995) and Parthanadee et al. (2006) implement neighborhood structures that modify the delivery schedule of clients by using

templates of delivery patterns, while modifier six looks at delivery days independently on a per week basis. It is assumed that changes made to the delivery schedule are valid as long as there is a volunteer available to service the client being affected. However, in reality there may be more stringent constraints related to client nutrition requirements limiting the effectiveness of this modifier. When performed, the day in which a client receives service will be changed to an alternative day that they do not already receive service. This modification does not change the volunteer servicing the client to avoid degrading the volunteer consistency objective. Figure 7 shows MOD 6 generating a neighbor solution by moving a clients delivery schedule from Mon., Wed., Fri, to Tue., Wed., Fri.

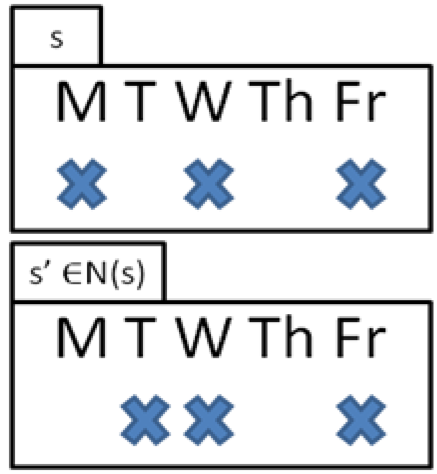

Figure 7: MOD 6 Example.

Algorithm 6: Change Delivery Day (MOD6)

1) For each $(i,v,d) \in B(s)$ where $i \in \mathbf{C}$:
   a) For each $e \in \mathbf{D}$ where $d \neq e$:
      i) For each $e \in \mathbf{D}$ where $d \neq e$:
         (1) Remove i from current route (v,d) and restructure the route using the < GENI > algorithm.
         (2) Insert i on route (v,e) using the < GENI > algorithm.

4.3.7 Solution Modifier 7: Consolidate Demand

The final modification attempts to conserve volunteer availability through the consolidation of deliveries. This type of modifier has been used by Shiguemoto and Armentano (2010) to

consolidate deliveries to customers in a production and distribution model. Neighboring solutions are built by taking unassigned customer demand and either placing them on new routes, routes not visiting the specific customer, or routes already visiting the customer by consolidating the products being shipped. An adaptation of this method is useful for MOW as they currently reduce weekly deliveries via consolidation efforts. To simplify this neighborhood structure, modifier seven only moves client demand to another day if that client is currently being serviced on that day. This move disregards volunteer consistency, and can result in a client being placed on a different volunteer's route. If a client is wait listed on one day, but receives service another, their un-serviced demand may also be allocated to a serviced day. This modification is useful for reducing the frequency of client visits, but may create infeasible solutions when volunteer capacity is violated. Figure 8 shows a neighbor solution being generated by MOD 7 where the deliveries for a client are consolidated from three days to two.

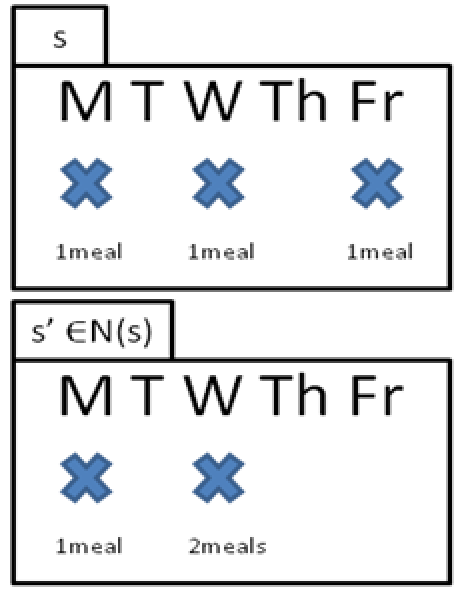

Figure 8: MOD 7 Example.

Algorithm 7: Consolidate Demand (MOD7)

1) <u>For each (*i*,*v*,*d*) ∈ B(*s*):</u>

   a) <u>For each (*i*,*w*,*e*) ∈ B(*s*) where $e \neq d$:</u>

      i) Set $CD_{i,e} = CD_{i,e} + CD_{i,d}$.

      ii) Remove *i* from current route (*v*,*d*) and restructure the route using the < GENI > algorithm.

4.4 Tabu Status & Diversification

Once the entire neighborhood of s is evaluated, the best neighbor solution in terms of $\mathrm{F}(s')$ is chosen to be the current s for the following iteration. To ensure that cycles do not occur in which the same attribute relationships are being altered back and forth, attributes of the newly chosen neighbor solution will be given a Tabu status. The Tabu Status $\tau_{i,v,d}$ takes on a value of the current iteration during $n + \theta$ iterations, preventing the acceptance of any attribute (i, v, d) until the iteration $\tau_{i,v,d}$ is reached when choosing neighbor solutions. This Tabu Status can be overridden if modifying the attribute (i, v, d) produces an objective value smaller than that of its associated aspiration value $\sigma_{i,v,d}$. Throughout the search, $\sigma_{i,v,d}$ will take the smaller value of either the $\mathrm{F}(s)$ associated with the last accepted solution or the minimum $\mathrm{F}(s)$ the attribute relationship (i, v, d) was assigned previously.

Solution diversification is implemented to steer the search away from local minima. To accomplish this, the objective function values for solutions containing route relationships that are frequently modified are degraded. If a neighbor solution s' does not improve upon its initial solution s, this penalty is added to F(s'). The penalty term $G(s)$, adopted from (Gendreau, Hertz and Laporte 1994), uses the parameter γ, a parameter used to scale the diversification intensity, and information specific to the problem size and current solution taking the form

$$G(s) = \gamma * \left(\sqrt{(Number\ of\ Cleints) * (Number\ of\ Volunteers) * (Planning\ Horizon)}\right) *$$

$$(RT(s) + M1 * INC(s) + M2 * WL(s)) * \frac{\sum_{(i,v,d) \in B(s') \setminus B(s)} ¢_{(i,v,d)}}{Current\ Iteration}.$$

Frequent modifications are penalized by summing $¢_{i,v,d}$, the modification counter for each attribute, across all new attributes within a solution and dividing it by the current search iteration.

Additional diversification is allowed by accepting interim infeasible solutions. The weights α and β are used to penalize the terms (3) and (4), but only do so effectively when they have large values. To encourage solution diversification via infeasible interim steps, but return the search towards feasible solutions, the values of α and β decrease each iteration where the current solution is feasible. Similarly, every iteration where the solution is infeasible the values of α and β increase. The values of α and β are increased or decreased by the constant parameter δ.

4.5 Initialization

The search process is initialized one of two ways: by providing an initial set of routes to be optimized or by using the <u>Full Utilization Initialization</u> (FUI) algorithm. The FUI algorithm generates an initial set of routes by arbitrarily assigning a depot to each volunteer for the entire planning horizon on each day that the volunteer is available. The algorithm then assigns the closest client to the volunteers last inserted vertex, and repeats this process until either the volunteers capacity or availability are fully utilized, or there is no more client demand for that day. This algorithm is adapted from a popular TSP insertion algorithm know as the Nearest Neighbor algorithm (Gutin and Punnen 2002). If all capacity is utilized and there is still unfulfilled client demand, those demand instances are placed on the waiting list. The FUI algorithm is detailed as follows:

Algorithm 8: Full Utilization Initialization

1) <u>For each $v \in \mathbf{V}$:</u>

   a) Arbitrarily assign $(i, v, d) \in \mathrm{B}(s)$ for each d where $i \in \mathbf{De}$.

b) Set $Depot = i$;

c) <u>For each $d \in \mathbf{D}$ with $A_{v,d} > 0$:</u>

i) <u>While $RT_{v,d} \leq A_{v,d}$ and $Ut_{v,d} \leq C_{v,d}$:</u>

(1) Choose $j \in \mathbf{C}$ that minimizes $tt_{i,j}$ with $CD_{j,d} > 0$.

(2) Insert $(j, v, d) \in \mathrm{B}(s)$.

(3) <u>If $i \neq Depot$.</u>

(a) Set $RT_{v,d} = (RT_{v,d} + tt_{i,j} + tt_{j,Depot} + st_j) - tt_{i,Depot}$.

<u>Else</u>

(b) Set $RT_{v,d} = (RT_{v,d} + tt_{i,j} + tt_{j,i} + st_j)$

(4) Set $Ut_{v,d} = Ut_{v,d} + CD_{j,d}$.

(5) Set $i = j$.

2) <u>For each client-day combination $(i, d) \notin B(\boldsymbol{s})$ where $CD_{i,d} > 0$</u>

a) Insert $(i, Wl, d) \in \mathrm{B}(s)$.

b) Insert $i \in \mathbf{Wl}^{\circ}$.

4.6 The TS-VRVW Algorithm

The Tabu Search for Vehicle Routing with a Volunteer Workforce algorithm is presented below. A key including all notations and their respective meaning accompanies the algorithm.

Algorithm 9: TS-VRVW

1) If initial solution of routes provided:

    a) Define solution s via the attribute set B(s).

    b) Evaluate F(s).

    c) If s is feasible:

        i) Set $s^* = s$.

        ii) F(s^*) = F(s).

Else:

    d) Generate initial solution s using the < FUI > algorithm.

    e) Evaluate F(s).

    f) Set s^*=s.

    g) F(s^*) = F(s).

2) <u>For all (i,v,d), initialize</u>:

    a) $¢_{i,v,d} = 0$.

    b) $\tau_{i,v,d} = 0$.

    c) If s is feasible and (i,v,d) ∈ B(s):

        i) $\sigma_{i,v,d} = \mathrm{F}(s)$.

        Else:

        i) $\sigma_{i,v,d} = \infty$.

    b) Initialize:

i) Set $t = 0$.

2) <u>While $n \leq N$ and $t \leq \check{T}$</u>

a) Set $\mathrm{F}(s^n) = \infty$.

b) Execute MOD_1 through MOD_7 to generate all possible $s' \in N(s)$.

c) <u>For each $s' \in N(s)$:</u>

i) Calculate $\mathrm{F}(s')$.

ii) <u>For each $(i, v, d) \in \mathrm{B}(s')\setminus\mathrm{B}(s)$ or each (i, v, d) swapped using MOD_1:</u>

(1) If ($\tau_{i,v,d} < n$) or ($\mathrm{F}(s') \leq \sigma_{i,v,d}$ and ($\mathrm{CAP}(s')+\mathrm{OT}(s') = 0$)):

(a) Place $s' \in M(s)$.

(b) If $\mathrm{F}(s') > F(s)$:

(i) Set $\mathrm{F}(s') = \mathrm{F}(s') + \mathrm{G}(s')$.

d) Choose $s^n \in \mathrm{M}(s)$ minimizing F.

e) <u>For each $(i,v,d) \in \mathrm{B}(s)\setminus \mathrm{B}(s^n)$:</u>

i) Set $\tau_{i,v,d} = n + \theta$.

f) <u>For each $(i,v,d) \in \mathrm{B}(s^n)\setminus \mathrm{B}(s)$ or each (i, v, d) swapped using MOD_1:</u>

i) Set $\phi_{i,v,d} = \phi_{i,v,d} + 1$.

g) <u>For each $(i,v,d) \in \mathrm{B}(s^n)$:</u>

i) Set $\sigma_{i,v,d} = min\,\{F(s^n),\ \sigma_{i,v,d}\}$

h) Set $s = s^n$.

i) Set $\mathrm{F}(s) = \mathrm{F}(s^n)$.

j) If $\mathrm{CAP}(s) = 0$:

i) Set $\alpha = \frac{\alpha}{(1+\delta)}$.

Else:

ii) Set $\alpha = \alpha * (1 + \delta)$.

k) If $\mathrm{OT}(s) = 0$:

i) Set $\beta = \frac{\beta}{(1+\delta)}$.

Else:

ii) Set $\beta = \beta * (1 + \delta)$.

l) If $\mathrm{F}(s) < \mathrm{F}(s^*)$, $\mathrm{OT}(s) = 0$, $\mathrm{CAP}(s) = 0$, and all $(i, Wl, d) \in B(s)$ are $\in \mathbf{Wl}^{\circ}$:

i) Set $s^* = s$.

ii) $\mathrm{F}(s^*) = \mathrm{F}(s)$.

m) Set $n = n + 1$.

n) Update *t*.

Table 1: TS-VRVW Parameters	
s	The current solution being modified and evaluated in the search process.
$s*$	The current minimum solution found thus far in the search process.
s'	The current neighbor solution generated by modifying s.
s^n	The current minimum neighbor solution found.
$¢_{i,v,d}$	Number of times the route relationship (i,v,d) was modified throughout the search.
$\tau_{i,v,d}$	Iteration at which the route relationship (i,v,d) loses its tabu status.
θ	The number of iterations a route relationship's tabu status is increased.
$\sigma_{i,v,d}$	Aspiration criteria of route relationship (i,v,d)
$RT(\boldsymbol{s})$	Total route time associated with solution s.
$CAP(\boldsymbol{s})$	Capacity violations associated with solution s.
$OT(\boldsymbol{s})$	Total route length violations associated with solution s.
$F(\boldsymbol{s})$	Objective Function.
$G(\boldsymbol{s})$	Diversification penalty term.
$N(s)$	All neighboring solutions of s created through the use of a solution modifier.
$M(s)$	A subset of N(s), where the solutions are either non-Tabu or meet aspiration conditions.
M_1	Weight associated with volunteer consistency.
M_2	Weight associated with clients on the waiting list.
α	Weight associated with capacity violations.
β	Weight associated with route length violations.
δ	Infeasibility weight update parameter.
n	The current search iteration.
N	The max number of iterations in the search process.
t	The current elapsed time of the search process.
$\check{T}$	The time limit at which the search process is ended.
γ	Factor used to adjust the diversification intensity.
$\mathbf{Wl}^{\circ}$	Initial set of wait listed clients.

5 TS-VRVW Computational Experiments

The computational effectiveness of the TS-VRVW is studied next. Because it is built on the work of Cordeau et al., the TS-VRVW relies on several parameters to execute. After initial testing, almost all parameters from Cordeau et al. that performed well and were adopted. However, there are two parameters that could not be adopted from the literature, or required further analysis: the extended p-neighborhood Exn and the infeasibility modifier δ. The extended p-neighborhood is useful for testing vertex insertions into routes at locations that may not yield the best route times, but may decrease other objectives such as $WL(s)$.

However, the effective value of Exn needs be determined to avoid excessive computational time caused by unfruitful GENI insertion attempts. The modifier δ affects how the search process views infeasible solutions which, if considered, may lead to improving solutions. Setting δ to too large of a value can restrict the amount of time the search spends in the infeasible region, while too small of a value can take an excessive amount of iterations to recover from infeasibility. Therefore, an appropriate value must be determined.

The volunteer consistency objective attempts to approximate the perceived penalty of increased effort for providing inconsistent service to clients. Likewise, the objective for wait-listing clients estimates the penalty of not servicing a client by an out-and-back trip. Since neither of these penalties are experienced in the real world, there is no evidence that they will encourage either consistent service or wait list reduction. Therefore, the effectiveness of these approximations to guide the TS-VRVW towards improving the routes with their respective goals must be tested.

The values of the objective weights M1 and M2 will have drastic effects on the solutions produced by the TS-VRVW and require thorough analysis as well. The volunteer consistency and wait listing objectives will encourage improving solution selection in their own respects, but may have adverse affects on other objectives. It is easily observed that the wait listing objective opposes the route time objective, as without the requirement of providing service the route time objective would choose to wait list every client. Wait listing and consistency can also compete as in some scenarios it may not be possible to eliminate the wait list and serve each client consistently. At each iteration, depending on the clients or routes in question and the values of the weights, the search will favor one of competing objectives. Therefore, an appropriate set of weights should be found that ensure the TS-VRVW can be used to solve several instances.

To quantify the effectiveness of the TS-VRVW, it should be tested in a manner that will require exploitation of its multi-objective features. Additionally, the applicability of the TS-VRVW needs to be tested for situations specific to MOW. To accomplish this, the TS-VRVW will be benchmarked against work from the Consistent VRP (ConVRP) presented by Groër et al., which finds good solutions in terms of consistency and fully meeting demand (Groër, Golden and Wasil 2009). The application of the TS-VRVW to larger scale VRP problems is also analyzed, and the exploitation of the "free" vehicle resources of MOW is explored.

5.1 Test Problems

The TS-VRVW was tested on both randomly generated instances and instances from the literature to evaluate the algorithm parameters and objective weights. The randomly generated test instances contained service requests from 30 to 200 clients. Delivery requests were generated

uniformly over the length of the planning horizon of one through five days. The number of meals per request was generated via a geometric distribution to create instances with fewer meals per service request on average, but with a small probability of requiring more. With a probability of .8, offset by a value of one, this distribution consistently generates values between one and three. Travel and service time was generated from a truncated normal distribution, with a lower bound of zero. Varying average and standard deviations were used for each instance, while the service time was always generated with a mean of 2 minutes and a standard deviation of 1 minute. In instances containing more than one depot the travel time between depots was explicitly to ensure specific spacing, while travel time to each client vertex was randomly generated obeying the triangle inequality. The number of volunteers in each instance ranged from two to 30, while their weekly availability was generated in the same manner as client demand. Volunteer availability per day was generated uniformly between 60 and 90 minutes to encourage route building that could realistically represent delivery runs completed over a lunch or work break. Each volunteer was assigned a constant capacity of 12 meals, replicating the limited capacity of a meal container. A summary of the instances generated is listed in Table 2. For these instances, not all clients can be served on all days as vehicle capacity is limited. The instances were designed this with such high demand to challenge the TS-VRVW in terms of minimizing all three objectives.

Table 2: Randomly Generated Instances

Instance	# Depots	#Clients	#Volunteers	Avg Travel Time	Std Dev. Travel Time	Inter Depot Travel Time	Avg Vol. Availability (min)	Avg Vol. Days Availabile	Avg Meals Per Delivery
1	2	30	2	20	8	30	70.25	4	1.15
2	2	100	7	25	10	30	75.35	3.7	1.24
3	2	200	16	30	10	30	74.89	3.65	1.27
4	5	30	2	20	8	30	80.6	4	1.32
5	5	100	7	25	10	30	74.8	3.85	1.24
6	5	200	16	30	10	30	74.36	3.3	1.25
7	2	30	3	20	8	30	67.45	4.33	1.38
8	3	100	11	25	10	30	73.46	3.81	1.26
9	3	200	23	30	10	30	73.74	3.52	1.25
10	3	100	16	25	10	30	72.93	3.5	1.23
11	3	200	30	30	10	30	74.38	3.33	1.21

Groër et al. (2009) solve a VRP with the goal of providing consistent driver assignments and delivery times, named the ConVRP. Their instances are used to gauge the effectiveness of the TS-VRVW. The authors present 12 problems adapted from previous literature that contain between 50 and 199 client vertices with daily requests varying over a 5 day planning horizon. The instances contain between 5 and 18 vehicles, vehicle capacity is limited to between 140 and 200 units, while route length is limited between 160 and 230 minutes (though in some instances route length was unbound).Client locations were spread on a $[-50, 50]^2$ grid, with service time set at either zero or ten minutes.

Table 3: Adopted Parameters.

θ	7.5log10(Number of Clients + Days in Planning Horizon)
α	500
β	500
γ	.015

The authors assumed travel time could be equated to the Euclidean distance between locations- this research did the same. It must be noted though, that the ConVRP accommodated all client demand by adding vehicles to the solution as needed. Since the number of vehicles is considered input to the TS-VRVW, the average number of vehicles utilized per instance in the ConVRP solutions were adopted as a parameter for the TS-VRVW. This implies that on some instances wait lists will be generated when all demand is not met, but should not be immediately associated with poor performance of the TS-VRVW.

5.2 Parameter δ and Exn

The value $\delta = .5$ proposed from the original authors was adopted at first, but initial testing showed that solution selection on some instances were moving between feasible and infeasible solutions too aggressively. On the other hand, the term of Exn was newly proposed in this work to allow GENI to investigate more insertion points, and its initial test values were estimated. To analyze the effect of varying these parameters across multiple instances and objectives a term to compare the percent difference in objective values was used. First, the minimum solution for each instance ins and objective Obj, $Min(Obj_{ins})$, was determined by solving the TS-VRVW using parameters from literature (Exn was inactive while testing δ). This minimum was then used to calculate the percent difference between each tested parameter value's solution and the minimum found quantified as $Diff(Obj_{ins}) = \frac{Obj_{ins} - Min(Obj_{ins})}{Obj_{ins}}$. The average of $Diff(Obj_{ins})$ was then taken across each parameter value tested to estimate it's performance.

To test these parameters nine test problems were solved including the six test instances highlighted in Table 2 and the instances 1, 2, and 3 solved by the ConVRP. The experiments were broken down by objective solved, so instead of F(s) in its entirety the objectives were either RT, INC, or WL with CAP and OT. The parameter δ was tested at the values of .1, .2, .3, .4, and .5 with the assumption that a number lower than .5 would yield a less aggressive reaction to infeasibility, as Cordeau et al. (1995) showed higher values had the opposite effect. The parameter Exn was tested at 1, 2, 3, 4, and 50- any value larger would consist of evaluating all possible insertions in the smaller instances tested and excessive computationally. All other parameters remained constant, with the adopted parameters set to the values in Table 3, with the parameter Exn set to one while testing δ and the optimal value of δ was used while testing Exn. Each instance was solved either with all solution modifiers active, or without the consolidating

moves provided by Mod 7. The test were given a maximum 30 minutes of run time and were unrestricted by iterations completed.

The results of these experiments showed that a slight decrease in the value of δ down to.4 produced the minimum objective value on average over all scenarios. However, when consolidation was in effect, δ at a setting of .4 did well, but was usually not the best setting. Analysis of parameter Exn showed surprising results, as the overall best performing value was at one. The extended p-neighborhood relaxed the criteria for “good” insertion points, thus allowing GENI to accept insertions at less desirable locations in tours. On average though, this lead the search towards poorer solutions. Exn set to the value of 3 performed comparatively well, but investigation on the average $Diff(Iterations_{ins})$ showed that this setting had completed fewer iteration- which can potentially stifle the search process. Therefore, the favored setting is at the value of one, simply using the initial p-neighborhood proposed for insertions.

5.3 Modification Use

The TS-VRVW uses seven different neighborhoods to search for an improving neighbor solution at each iteration. However, the frequency at which certain modifications are used will vary on the instance being solved, or the objectives being modeled, and which modifications the TS-VRVW has access to. For example, solving an instance with a single depot could not use the second neighborhood structure to modify a solution. Additionally, the inclusion of consolidation is a strong tool in finding improving solutions, and can be more beneficial for the search than other modifiers independent of the objective.

A sample taken from first six test instances used for parameter testing is analyzed to highlight the effectiveness of each modification. The iterations analyzed include those leading to

the minimum solution found for the instance. Over this search process, both improving and interim solutions are chosen, eventually leading to minimum solution found for the instance. The percentage of which modifier was used during these iterations is calculated, and averaged over each of the instances. This information is broken down by objective solved, and whether consolidation was allowed during the search (NC- no consolidation, C- consolidation).

It is observed in Figure 9 that when a viable option, consolidation (Mod 7) is a popular neighborhood construct, used 89% of the search iterations before a minimum is found under the consistency objective and 78% while optimizing route time. The wait-list reduction objective WL not surprisingly preferred utilizing Mod 5, as it placed unmet demand onto delivery routes. However, the objectives INC and RT also utilized Mod 5, regardless of no need to meet additional demand. This can be attributed to the improvements in route makeup when GENI inserts clients and reorganizes routes- improving routes developed by the FUI algorithm. The remainder of the modifiers chosen were scattered amongst Mods 1, 2, 3,4, and 6, where Mod 1 was the next frequently chosen. Surprisingly though, was how infrequent the use of Mod 3 was given the INC objective. This can best be explained by analyzing the initial routes developed, as there are few cases that the routes were not 100% utilized in either capacity or availability constraints are violated. This implies that changing the route a client is serviced by tended to incur a penalty for infeasibility.

When consolidation was not an option, the search frequently used Mod 1 for objectives INC and RT, 50% and 40% of the time respectively. Again, objective WL heavily favored Mod 5 throughout the search, but surprisingly so did objectives RT and INC. This again can be attributed to the GENI insertion algorithms ability to improve the overall route travel time. The use of Mod 4 slightly increased, as it becomes a more attractive option to relieve infeasible conditions when the stronger consolidation modifier is not available.

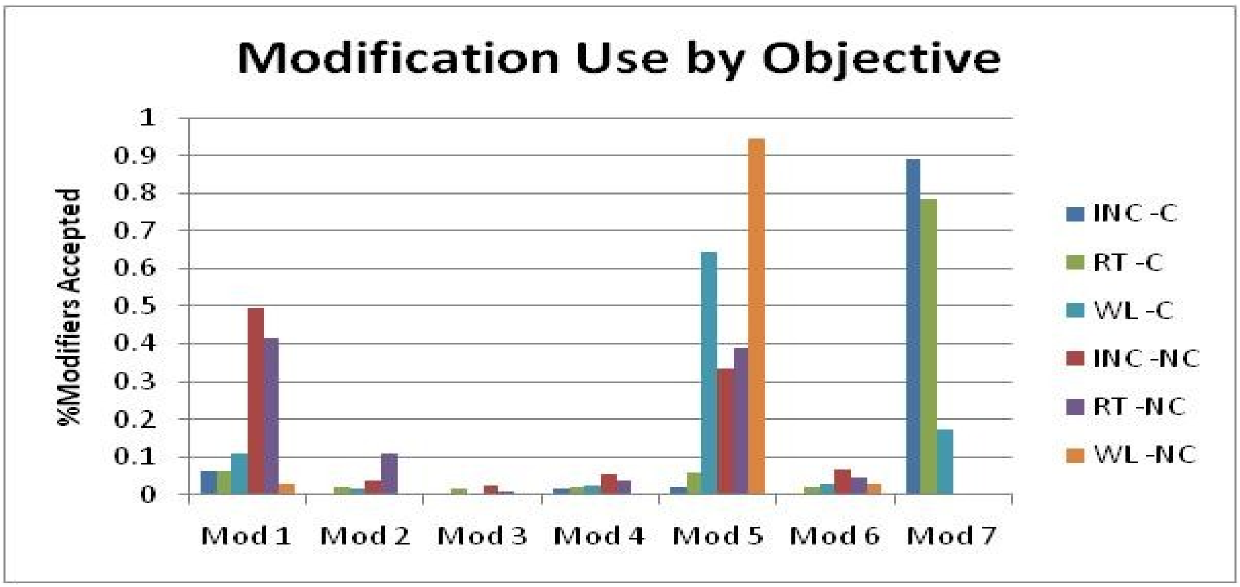

Figure 9: Modification Usage Leading to a Minimum Solution.

5.4 Objective INC and WL Validation

The approximations of volunteer-client consistency and client wait-listing are mechanisms designed to guide the TS-VRVW towards meeting their respective goals. Since these methods do not represent the actual cost incurred by MOW for wait-listing clients or pairing volunteers with clients inconsistently, the effectiveness of these approximations are studied by analyzing their effects on the solutions generated. To analyze the INC objective, three models are solved: where the main objective is RT(s), where the main objective is INC(s), and where the objective is

RT(s)+M1*INC(s). Similarly, the WL objective is analyzed with the three models: RT(s), WL(s), and the combination of RT(s)+M2*WL(s). Experiments are run for 30 minutes with unlimited iterations, and either with or without consolidation. Five randomly generated instances (numbers 1, 2, 3, 10, and 11) were used for comparing the objectives effectiveness.

The results pertaining to the volunteer consistency objective INC show that on average the

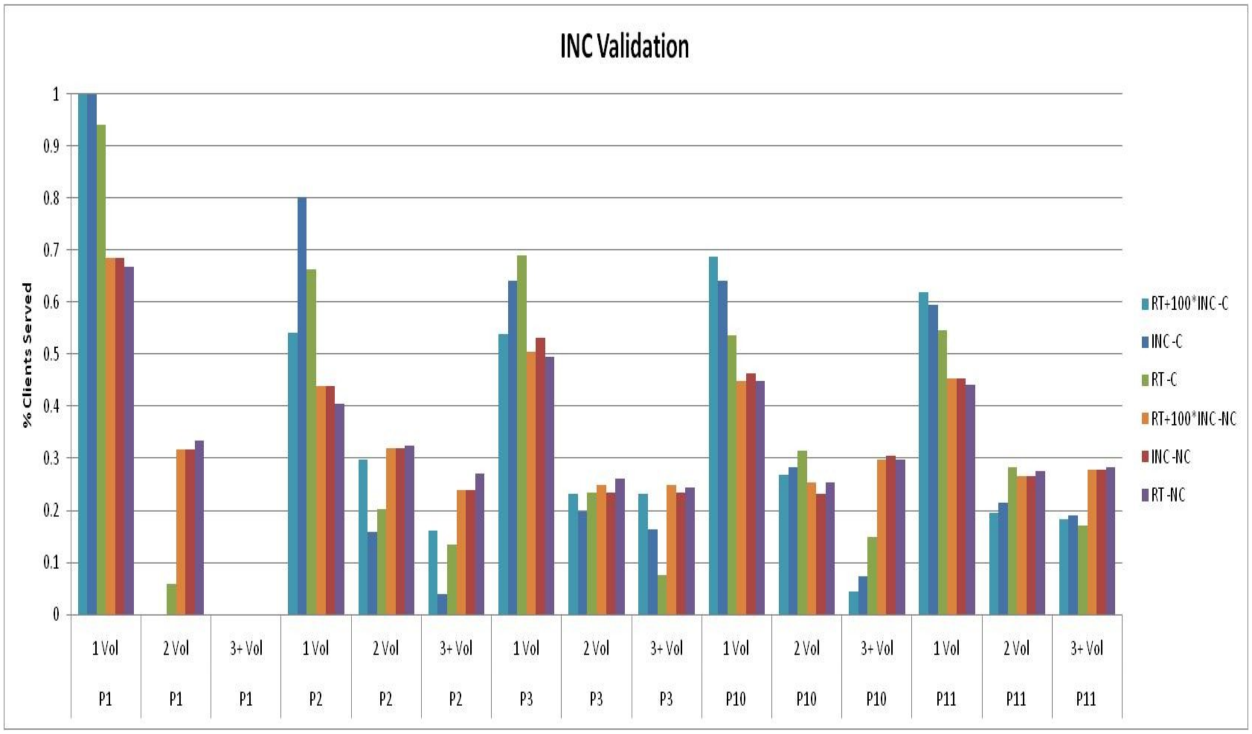

Figure 10: INC Validation by Objective.

INC objective will outperform RT based on the percentage of clients served by only a single volunteer. The INC objective increased the percentage of clients seen by a single volunteer to 74% across all of the instances, while RT only averaged 68%. However, RT was able to outperform INC when applied to instance 5 when consolidation was allowed during the search- potentially due to the higher average demand and larger travel times of the instance where route optimization lead more easily to consistent pairings. Additionally, the INC objective tends to reduce the number of clients with two, three, or more volunteers servicing them over the

planning horizon. Solving these instances with both RT and INC in the objective (with M1 set to 100) tended to produce consistent pairings on average slightly better than only RT, but not as strong of solutions as INC. The combined objective tended to fall short on instances where travel time was larger, demand was higher, and consolidation was allowed. This is attributed to the competing objectives allowing the search to spend more iterations at infeasible interim steps. Figure 10 shows the performance of each objective, labeled with C if consolidation was used and NC otherwise, broken down by instance solved.

The wait listing objective WL is measured by how effectively it promotes meeting all demand in a given instance. Across almost all five instances, the WL objective steered the search towards meeting more of the client demand over the planning horizon over the RT objective. The magnitude of improvement WL realized over RT was large when consolidation was a factor in the search process. This is an obvious result of the increased utilization of volunteer capacity. RT did surpass WL in terms of meeting demand when solving instance 11, however, as adding client service from the waiting list helped optimize the routing configuration through reorganizing the initial routes and allowing for further client consolidation across all volunteers. This was not realized when WL was the sole objective, as capacity along the routes weren't as efficiently optimized. The combination of RT and WL into a single objective (with M2 set to 100) also showed improvements on the percent of demand met across the instances, but when compared to RT alone only improved this metric marginally. Figure 11 shows the performance of each objective, labeled with C if consolidation was used and NC otherwise, broken down by instance solved.

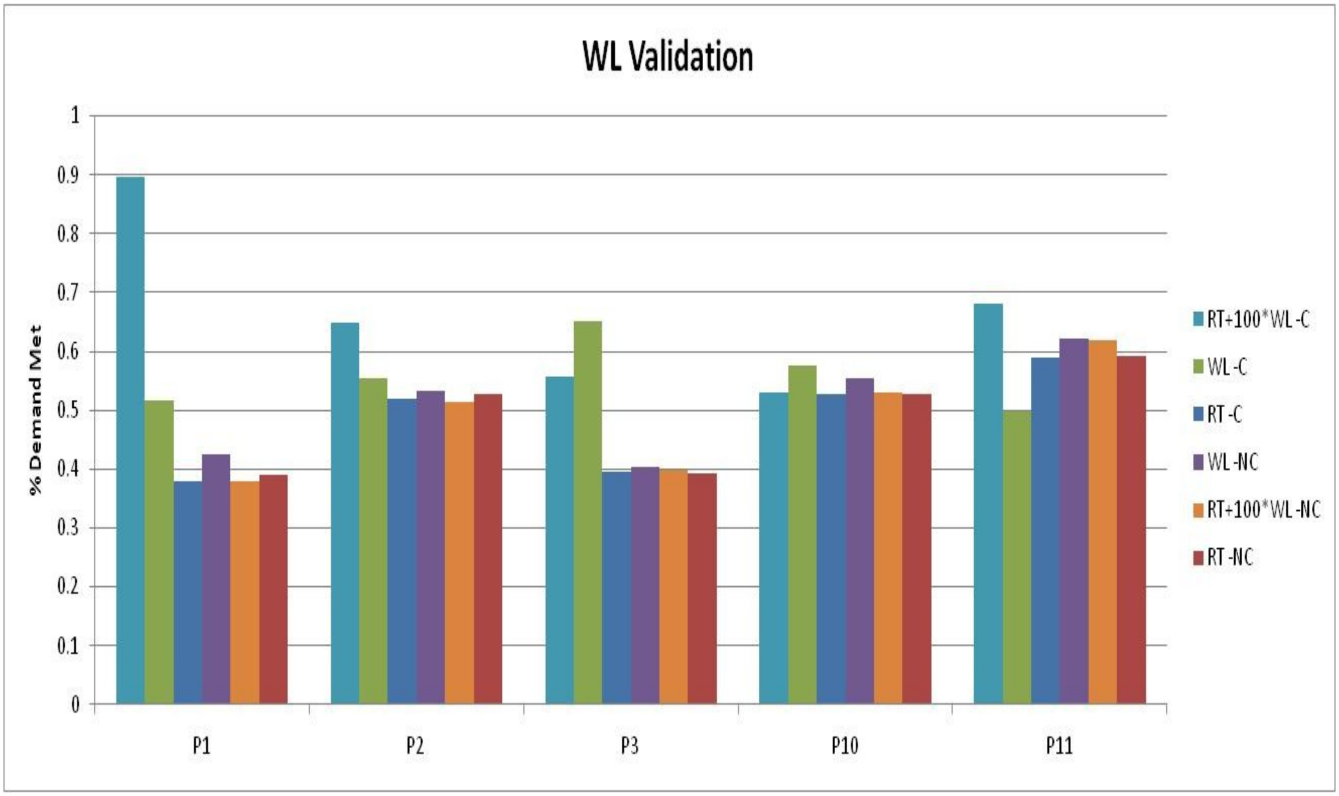

Figure 11: WL Validation by Objective.

5.5 Analyzing the Effect of Weight Values on the Objectives

To understand the effect the weights M1 and M2 have on each objective value used in the TS-VRVW sensitivity analysis is completed on both terms. This analysis is concerned with the effect the weight wt associated with a primary objective P on the values of the secondary or tertiary objectives in the aggregated objective function. When analyzing the interactions between two active objectives, there will be four Primary and Secondary objective pairings: M1*INC-RT, M1*INC-WL, M2*WL-RT, M2*WL-INC. In this analysis, P can be either the objective INC or WL as the effect of their weights is of concern. Additionally, the effects of weights will be conducted when all three objectives are active, varying either M1 or M2.

To compare the magnitude of effect the weights have across several test problems on the RT objective a differential term is adapted from Smilowitz et al. (2011), comparing F^{RT} to the overall minimum observed F^{RT^*} for the instance being solved across the varying weight levels. This differential takes the form

$$\Delta_{F^{RT}}(P_{wt}) = \frac{F^{RT}(P_{wt}) - F^{RT^*}}{F^{RT^*}}.$$

If the primary objective was INC, then the differential would be $\Delta_{F^{RT}}(INC_{M1}) = \frac{F^{RT}(INC_{M1}) - F^{RT^*}}{F^{RT^*}}$, showing the impact of INC with weight M1 on the objective of RT. Values of $\Delta_{F^{RT}}(P_{wt})$ closer to zero imply that the weight wt has a minor degrading effect on the RT objective, while large values of $\Delta_{F^{RT}}(P_{wt})$ imply a larger degrading effect.

To analyze the effect of weights on the consistency of a solution, the average number of volunteers assigned per client AVC over the planning horizon will be used to summarize each instance. The metric $AVC^{INC}(P_{wt})$ will show how a primary objective weight will affect the performance of the consistency in the final solution. If considering WL as the primary objective, then the metric of average volunteers per client is represented as $AVC^{INC}(WL_{M2})$. It is clear that a value of AVC closer to one implies the primary weight promotes consistent service, while the opposite is true for larger values of AVC.

While analyzing the differential of WL as a secondary objective, the main metric will be the actual demand met DM in comparison to the total demand across the planning horizon, TD. The TD is defined by the sum of all meal requests over the planning horizon of an instance, while DM is defined as the number of meals delivered over the planning horizon. This differential is expressed as

$$\Delta_{DM^{WL}}(P_{wt}) = \frac{TD - DM^{WL}(P_{wt})}{TD}.$$

A value of $\Delta_{DM^{WL}}(P_{wt})$ closer to zero implies that the primary weight promotes meeting demand in the problem, while a large differential shows that the weight has a negative impact on the objective.

To replicate the evaluation of the differentials presented by Smilowitz et al. (2011), four ConVRP instances and three randomly generated instances were chosen to run the experiments over. The values of M1 and M2 were tested at values ranging from 1 through 1000 and averaged across all instances. The effects the weights have on the objectives are presented graphically, comparing $\Delta_{AVC^{INC}}(INC_{wt})$ versus either $\Delta_{F^{RT}}(INC_{wt})$ or $\Delta_{DM^{WL}}(INC_{wt})$, and $\Delta_{DM^{WL}}(WL_{wt})$ versus $\Delta_{F^{RT}}(WL_{wt})$ or $\Delta_{AVC^{INC}}(WL_{wt})$. This analysis looks at all three statistics, even if there are only two active objectives experimental runs. If this is the case, the third statistic is just a resultant of decisions based on the first two objectives. Each instance was solved either with all solution modifiers active, or without the consolidating moves provided by Mod 7. The test were given a maximum 10 minutes, as shown in previous work of Smilowitz that a TS algorithm only needed a few minutes to solve the instances, but extra time was allotted to compensate for the extra neighborhoods being searched.

5.5.1 General Weight Analysis Results

In general, weighting INC with more importance than other objectives tends to have a degrading effect on the RT of a solution with a no consolidation policy. On the other hand, the inclusion of consolidation shows that as AVC decreases, the RT of a solution decreases. Additionally, there is a correlation between a solution having lower AVC and meeting more demand with or without

using consolidation. Weighting WL with more importance showed that the RT of a solution increased as more demand was met without consolidation. When consolidation was allowed the RT of a solution initially decreases as more demand is met, and then increases again in a parabolic manner (Figure 12). The initial decrease in RT is attributed to Mod 7 consolidating demand in a solution. The inflection occurs when consolidation is no longer viable and WL was weighted strong enough that more demand was met from the Waiting List at the expense of RT. No relationship could be determined between the demand differential of a solution and the AVC under a no consolidation policy, while using consolidation showed a slight decrease in AVC as more demand was met. Detailed weighting analysis can be found in the Appendix, Chapter 9.1.

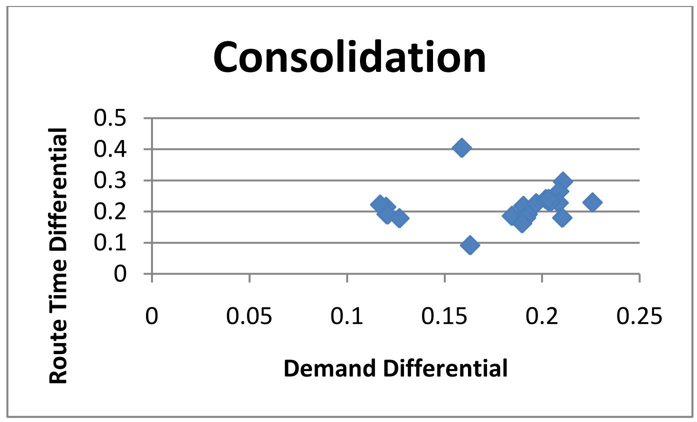

Figure 12: Effect of WL on RT (3obj-C).

5.5.2 Weight Selection

To successfully apply the TS-VRVW to instances with the goal of having an overall good solution, the appropriate values for M1 and M2 must be chosen. A value for each weight is determined to balance its effect on the rest of the objective F(s). A normalizing function is adopted from Smilowitz et al. (2011) where the objective value of RT, in proportion to the minimal RT objective found for each instance, the average volunteer per client, and proportion of

total demand to demand met are aggregated based on the corresponding objective weight. The normalization for the primary objective weight is expressed as

$$\mathrm{Norm_{Pwt}} = \frac{\mathrm{F^{RT}(P_{wt})}}{\mathrm{F^{RT^*}}} + \mathrm{AVC^{INC}(P_{wt})} + \frac{\mathrm{TD}}{\mathrm{DM^{WL}(P_{wt})}}.$$

The $\mathrm{Norm_{INCwt}}$ and $\mathrm{Norm_{WLwt}}$are graphed versus weight value in for M1 in Figure 13 and for M2 in Figure 15 over instances solved with only two objectives. Minimizing the normalization function yields a balanced value for the respective norm. For the volunteer consistency objective, INC the weight M1 should be set to 1 if the TS-VRVW will be using consolidation, and 0 otherwise. For the wait-listing objective WL the weight M2 should be given the value of 75 when the TS-VRVW will be using consolidation, and zero otherwise. When three objectives are included in the objective function, the term M1 should be set to zero when consolidation is

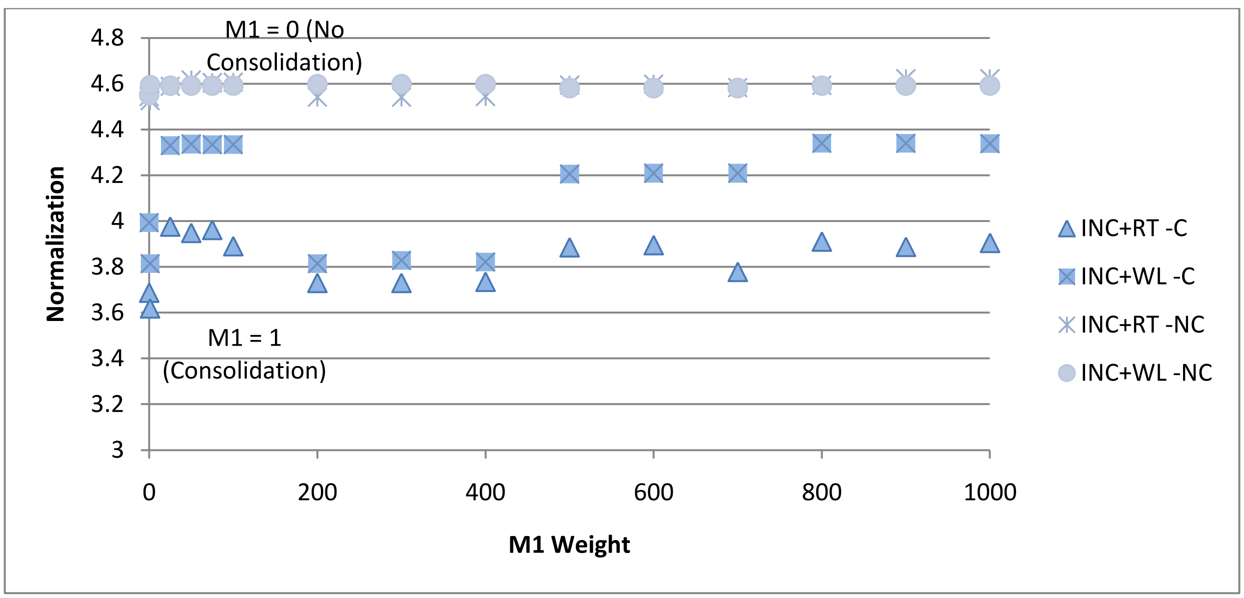

Figure 13: Normalization of M1 (Two Objectives).

active, and to zero when consolidation is not considered- implying the inclusion of INC degrades the overall objective drastically (Figure 14). Minimizing the normalization for M2 shows that the weight should be set to 400 with consolidation, and 150 otherwise (Figure 16). This balanced

weighting scheme implies that very low importance weights can have enough effect on the solution to encourage the objectives INC and WL. However, it also shows that the secondary objectives have too strong of an adverse effect on the overall objective function to include them.

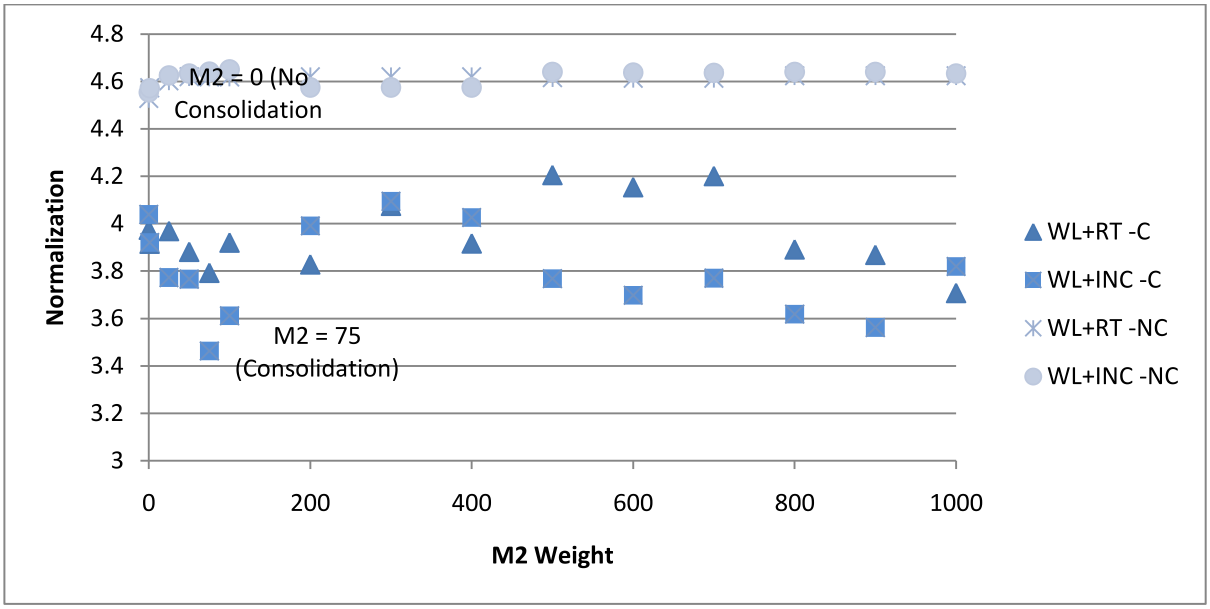

Figure 15: Normalization of M2 (Two objectives).

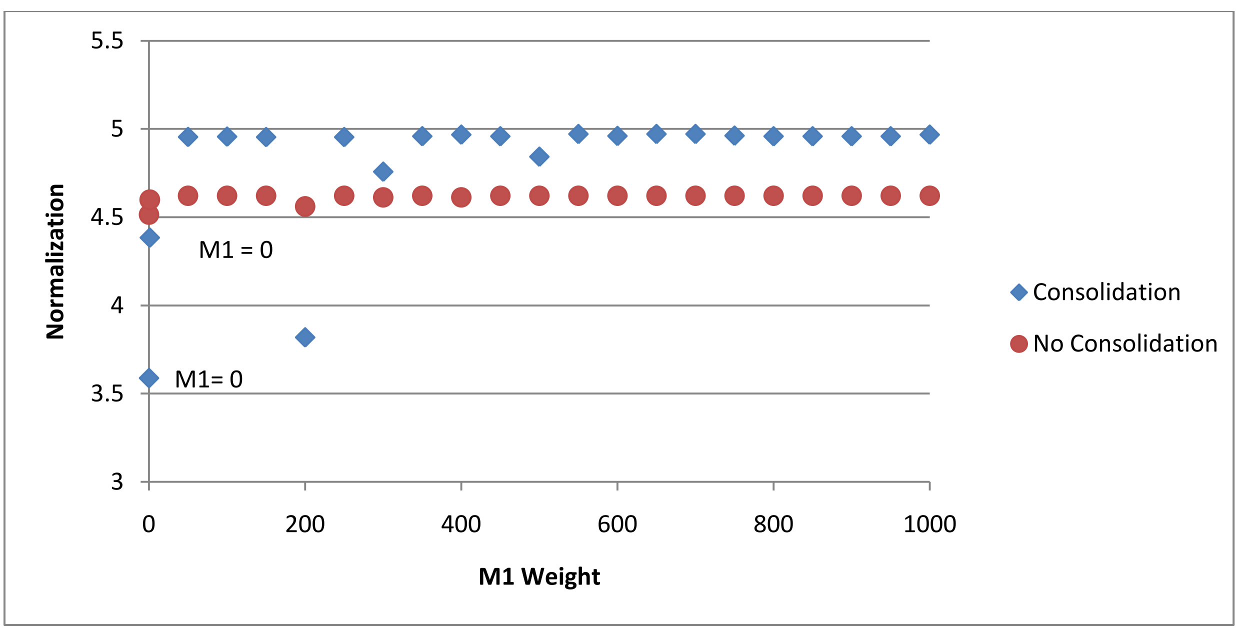

Figure 14: Normalization of M1 (Three Objectives).

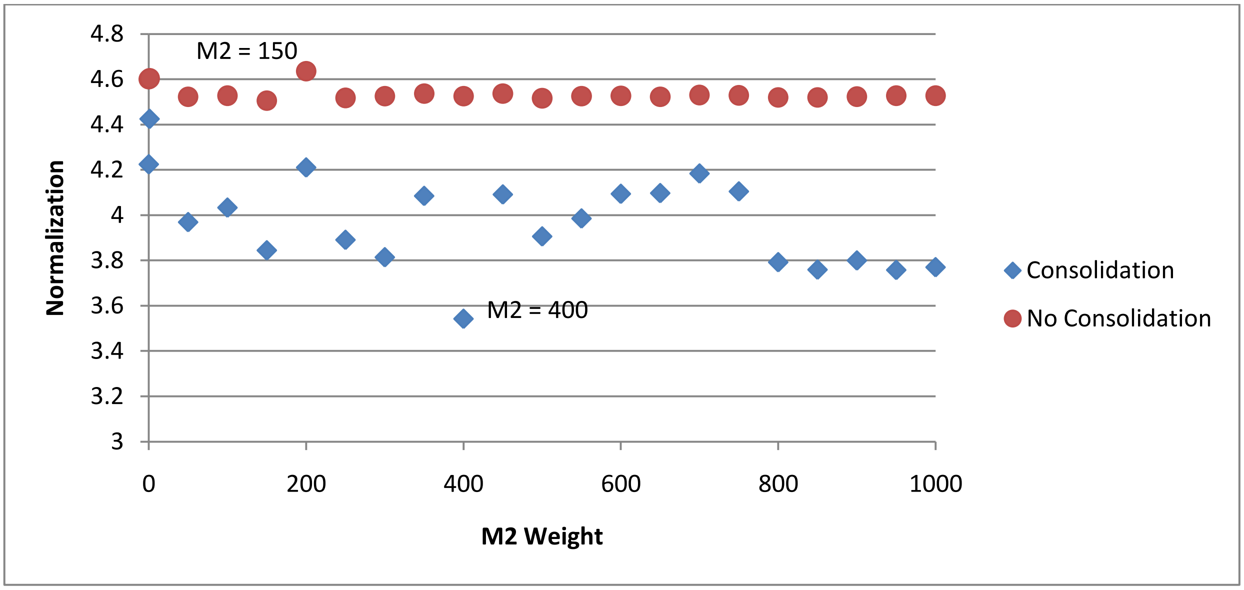

Figure 16: Normalization of M2 (Three Objectives).

5.6 Benchmarking the ConVRP

Table 4: TS-VRVW Average Time until Minimum Solution Found (ConVRP Instances).

Instance	Time Until Found (minutes)
C1	1.35
C2	5.99
C3	2.28
C4	14.56
C5	16.61
C6	0.22
C7	1.76
C8	2.68
C9	7.04
C10	15.48
C11	11.61
C12	2.03

With parameter and weight tuning of the TS-VRVW completed, the performance of the TS-VRVW can be compared to ConVRP in a comparative study. The solutions produced by the authors Groër et al. (2009) were 100% consistent in terms of drivers visiting the same clients over the planning horizon. Additionally, the solutions generated accommodated all client demand. Therefore, at best the objectives INC and WL will meet the

ConVRP at zero. The values found for objectives RT are expected to be comparable to those reported, and even surpass the ConVRP results when consolidation is allowed in the routing development process.

The benchmark tests were run for 30 minutes with unlimited iterations on the 12 ConVRP problems, though Table 4 shows that many of the instances averaged times much shorter time than allotted until a minimum solution was found. All parameters were adopted from the results of previous tests, and four objectives were tested: RT+INC+WL with balanced weights and RT alone, each under both consolidation strategies. The balanced weighting scheme derived from the normalization analysis when three objectives are active suggests setting M1 to one, and M2 to 400 under a consolidation strategy, while M1 is set to one and M2 should be set to 150 when consolidation is not allowed.

The solutions to the benchmarking problems are listed in Table 5. The instances solved are listed by consolidation policy as Consolidation or No Consolidation, ConVRP instance, and the weight values of M1 and M2. The algorithm output is listed under RT for minimal route time found by the TS-VRVW, ConRTR, and RTR algorithms proposed by Groër, et al. (2009), while a gap between TS-VRVW and RTR quantified as

$$\frac{(RT_{TS-VRVW} - RT_{RTR\ Algorithm})}{RT_{RTR\ Algorithm}}$$

is listed. Additionally, unmet demand in quantities of meals (or deliverable goods in the case of the ConVRP) are listed, while consistency information is provided for the TS-VRVW with the AVC statistic and the number of clients serviced by a single volunteer, two volunteers, or three or more volunteers over the planning horizon.

When consolidation was allowed in the search, the TS-VRVW was able to surpass the RT objective of the RTR algorithm over 11 of 24 experiments. In addition to these

improvements, two solutions were found to surpass the RT objective of the RTR without consolidation modifiers. There were also very few instances solved where demand was wait-listed for either consolidation policy. Though, when there was unmet demand it never surpassed that of a single delivery request. An unexpected result was that there was only one instance solved where volunteers were fully consistently serving clients. Though the results are not strong from a consistency perspective, comparing the solutions derived with only RT and the balanced objectives shows that the inclusion of INC drove the TS-VRVW towards assigning volunteers to clients more consistently.

			Route Time								
Instance	M1	M2	ConRTR	RTR	TS-VRVW	RT Gap	Unmet Demand	AVC	1 Vol.	2 Vol.	3+ Vol.
Cons.											
C1	0	0	2282.14	1963.39	1806	-0.08	0.00	1.66	24	19	7
	1	400	2282.14	1963.39	1766	-0.10	0.00	1.06	47	3	0
C2	0	0	3872.86	3182.31	3862	0.21	0.00	2.20	18	29	28
	1	400	3872.86	3182.31	3496	0.10	0.00	1.17	62	13	0
C3	0	0	3628.22	3127.77	2838	-0.09	0.00	1.20	82	16	2
	1	400	3628.22	3127.77	3204	0.02	0.00	1.06	95	4	1
C4	0	0	4952.91	4121.73	4246	0.03	0.00	1.54	80	60	9
	1	400	4952.91	4121.73	4996	0.21	0.00	1.19	123	23	3
C5	0	0	6416.77	5108.19	6570	0.29	0.00	2.58	20	74	104
	1	400	6416.77	5108.19	7406	0.45	0.00	2.57	36	58	104
C6	0	0	4084.24	3954.32	1508	-0.62	0.00	2.43	38	9	2
	1	400	4084.24	3954.32	1754	-0.56	0.00	2.43	49	0	0
C7	0	0	7126.07	6325.39	4752	-0.25	13.00	1.64	36	31	8
	1	400	7126.07	6325.39	5056	-0.20	0.00	1.08	69	6	0
C8	0	0	7456.19	6902.1	4330	-0.37	4.00	1.41	65	30	5
	1	400	7456.19	6902.1	4084	-0.41	0.00	1.00	100	0	0
C9	0	0	11033.54	9932.9	5744	-0.42	13.00	1.11	133	17	0
	1	400	11033.54	9932.9	7818	-0.21	0.00	1.09	136	14	0
C10	0	0	13916.8	12399.4	13164	0.06	4.00	2.48	32	69	97
	1	400	13916.8	12399.4	15052	0.21	0.00	2.61	38	47	113
C11	0	0	4753.89	4244.48	4504	0.06	0.00	1.12	105	14	0
	1	400	4753.89	4244.48	4604	0.08	0.00	1.02	117	2	0
C12	0	0	3861.35	3209.88	3898	0.21	0.00	1.60	51	38	11
	1	400	3861.35	3209.88	3598	0.12	0.00	1.15	86	13	1
No Cons.											
C1	0	0	2282.14	1963.39	2822	0.44	6.00	2.38	6	19	25
	1	150	2282.14	1963.39	3126	0.59	10.00	2.14	10	23	17
C2	0	0	3872.86	3182.31	4520	0.42	0.00	2.60	7	25	43
	1	150	3872.86	3182.31	4884	0.53	0.00	2.44	9	32	34
C3	0	0	3628.22	3127.77	4708	0.51	0.00	2.51	9	45	46
	1	150	3628.22	3127.77	5018	0.60	0.00	2.29	15	50	35
C4	0	0	4952.91	4121.73	6246	0.52	0.00	2.66	12	55	82
	1	150	4952.91	4121.73	6700	0.63	0.00	2.52	18	60	71
C5	0	0	6416.77	5108.19	7744	0.52	0.00	2.99	9	51	138
	1	150	6416.77	5108.19	7912	0.55	0.00	2.99	10	58	130
C6	0	0	4084.24	3954.32	3064	-0.23	0.00	2.43	6	22	21
	1	150	4084.24	3954.32	3210	-0.19	0.00	2.43	16	18	15
C7	0	0	7126.07	6325.39	8258	0.31	13.00	2.51	8	30	37
	1	150	7126.07	6325.39	8590	0.36	12.00	2.52	8	31	36
C8	0	0	7456.19	6902.1	9142	0.32	4.00	2.85	6	28	66
	1	150	7456.19	6902.1	9290	0.35	0.00	2.81	8	29	63
C9	0	0	11033.54	9932.9	12716	0.28	13.00	2.75	14	44	92
	1	150	11033.54	9932.9	13018	0.31	13.00	2.77	14	43	93
C10	0	0	13916.8	12399.4	16444	0.33	4.00	2.95	9	49	140
	1	150	13916.8	12399.4	16634	0.34	0.00	2.94	10	47	141
C11	0	0	4753.89	4244.48	6000	0.41	0.00	1.73	43	61	15
	1	150	4753.89	4244.48	6328	0.49	0.00	1.73	51	59	9
C12	0	0	3861.35	3209.88	4980	0.55	0.00	1.95	37	35	28
	1	150	3861.35	3209.88	4836	0.51	0.00	1.91	40	33	27

Table 5: ConVRP Benchmark Solutions

Further investigation into the effectiveness of the INC term shows that although it may not have performed well against the ConVRP instances, the objective does reduce inconsistent service in a solution significantly. Table 6 shows a comparison between the final solution of the TS-VRVW allowing consolidation, with balanced objective weights, and the initial solution generated from the FUI algorithm. The FUI algorithm creates an initial solution with very poor quality in regards to route time and consistent service. The magnitude of improvement across most of the solutions in terms of RT or AVC is rather significant in computationally short search times, even if full consistency is rarely reached. These results not only affirm the improvement of the solution based on the objective function, but also implies that search process may be hindered by the construction of the initial solutions.

			Route Time			AVC			
Instance	M1	M2	TS-VRVW	FUI	RT Gap	TS-VRVW	FUI	AVC Gap	Time Until Found (Min)
C1	1.00	400.00	1766.00	3238.00	0.83	1.06	2.36	1.23	4.95
C2	1.00	400.00	3496.00	16560.00	3.74	1.17	2.95	1.52	0.25
C3	1.00	400.00	3204.00	6792.00	1.12	1.06	1.73	0.63	0.70
C4	1.00	400.00	4996.00	5408.00	0.08	1.19	1.96	0.64	2.68
C5	1.00	400.00	7406.00	4840.00	-0.35	2.57	2.60	0.01	3.50
C6	1.00	400.00	1754.00	5196.00	1.96	2.43	2.51	0.03	0.03
C7	1.00	400.00	5056.00	6678.00	0.32	1.08	2.66	1.47	0.50
C8	1.00	400.00	4084.00	8068.00	0.98	1.00	2.98	1.98	0.18
C9	1.00	400.00	7818.00	3216.00	-0.59	1.09	2.42	1.21	1.03
C10	1.00	400.00	15052.00	8550.00	-0.43	2.61	2.52	-0.03	0.98
C11	1.00	400.00	4604.00	9240.00	1.01	1.02	2.85	1.80	3.20
C12	1.00	400.00	3598.00	13018.00	2.62	1.15	2.77	1.41	0.25

Table 6: TS-VRVW Improvements Over FUI.

5.7 The TS-VRVW and MOW

5.7.1 TS-VRVW Applied to a Large

To evaluate the applicability of the TS-VRVW MOW or other volunteer delivery organizations, an instance similar in size to one handled by a regional MOW facility was solved. The instance contains 750 clients requiring service over a period of five weekdays. The region in consideration has three depots located 30 minutes from one another. Travel time between clients is distributed randomly with a mean of 20 minutes and a standard deviation of 10 minutes. There are 100 volunteers MOW has access to for servicing the clients over the five days, but each has specific availability throughout the week. On average, the volunteers have 74.14 minutes available per day for service, with actual availability ranging from 60 to 90 minutes. The average number of days available for each volunteer is 3.31, though the actual days available vary between only once per week to each day of the week.

The Large Scale problem was solved using the balanced weight scheme including all three objectives and only RT as an objective. Both consolidation and non-consolidation is considered when building the delivery routes. The TS-VRVW was run for 120 minutes due to the size of the

			Route Time			AVC			Unmet Demand			
	M1	M2	TS-VRVW	FUI	RT Gap	TS-VRVW	FUI	AVC Gap	TS-VRVW	FUI	AVC Gap	Time Until Found (Min)
Cons.												
	1.00	75.00	15791.00	15744.00	0.00	2.87	2.87	0.00	3.00	6.00	1.00	104.48
	0.00	0.00	15429.00	15744.00	0.02	2.87	2.87	0.00	6.00	6.00	0.00	107.42
No Cons.												
	1.00	75.00	15791.00	15744.00	0.00	2.87	2.87	0.00	3.00	6.00	1.00	101.60
	0.00	0.00	15335.00	15744.00	0.03	2.87	2.87	0.00	6.00	6.00	0.00	110.75

Table 7: TS-VRVW Performance on a Large Instance.

neighborhood structures. shows that compared to the FUI algorithm, the TS-VRVW has little effect on solution generated. Over the two hour span of search time, the algorithm only completed an astonishingly low number of three iterations. This small number of iterations implies that the search did not have adequate time to improve the initial solution, and can be attributed to size of the neighborhoods considered at each iteration.

5.7.2 Utilizing Free Resources

Volunteer organizations such as MOW are unique in the aspect that they do not incur fixed or variable costs for routing volunteers. This leads to the question of whether the objective RT is required for the TS-VRVW to produce attractive results. Since the RT objective leads the search algorithm towards creating better routes, a natural assumption would be that while this occurs, the search algorithm benefits from freed capacity by minimizing the waiting list and making volunteer assignments more consistent.

To test this, the TS-VRVW is run with a balanced weight scheme for M1 and M2, but with no RT in the objective function over the 12 ConVRP instances. Experimental procedures replicated those in section 5.6. Table 8 shows the comparison between the TS-VRVW results without RT versus its results only considering RT. While consolidation was active, very strong reductions in AVC and wait-listed demand occurred under no RT consideration. These metrics were only beat on one instance against pure RT regarding AVC. The results also show some rather large increases in overall RT when it is not considered. The same trends are observed when consolidation was not allowed; though the improvements of AVC or wait listed demand was marginal.

These results prove that TS-VRVW can in fact be applied with no consideration to RT in the objective function, and still obtain reasonable results. How reasonable these results are, however, are specific to the application of the VRP, as an organization that pays for routing resources may not be willing to accept a 51% increase in total route time. Running without RT met practically all demand (excluding the No Consolidation trials as their balanced weight scheme did not include WL), which is an attractive result service based organization. Additionally, consolidating allowed for 8 of 12 instances to be solved with AVSs within 10% of complete consistent service.

	No-RT Weights		RT		Gap	AVC		Gap	Uns	
					No-RT from			RT-Only		
Instance	M1	M2	RT-Only	No-RT	RT-Only	RT-Only	No-RT	from No-RT	RT-Only	No-RT
Cons.										
C1	1	75	1806	2656	0.47	1.66	1.02	0.63	0	0
C2	1	75	13164	15536	0.18	2.48	2.55	-0.03	4	0
C3	1	75	4504	6328	0.40	1.12	1.02	0.10	0	0
C4	1	75	3898	5170	0.33	1.60	1.18	0.36	0	0
C5	1	75	3862	3738	-0.03	2.20	1.08	1.04	0	0
C6	1	75	2838	4062	0.43	1.20	1.03	0.17	0	0
C7	1	75	4246	5892	0.39	1.54	1.14	0.35	0	1
C8	1	75	6570	7828	0.19	2.58	2.54	0.01	0	0
C9	1	75	1508	2284	0.51	1.27	1.04	0.22	0	0
C10	1	75	4752	5798	0.22	1.64	1.07	0.54	13	0
C11	1	75	4330	5886	0.36	1.41	1.00	0.41	4	0
C12	1	75	5744	8262	0.44	1.11	1.00	0.11	13	0
No Cons.										
C1	1	0	2822	4118	0.46	2.38	2.28	0.04	0	0
C2	1	0	16444	16578	0.01	2.95455	2.93939	0.01	4	4
C3	1	0	6000	8214	0.37	1.78151	1.66387	0.07	0	0
C4	1	0	4980	6280	0.26	1.95	1.89	0.03	0	0
C5	1	0	4520	5792	0.28	2.6	2.50667	0.04	0	0
C6	1	0	4708	5234	0.11	2.51	2.42	0.04	0	0
C7	1	0	6246	7306	0.17	2.66443	2.56376	0.04	0	0
C8	1	0	7744	8090	0.04	3	2.93434	0.02	0	0
C9	1	0	3064	3348	0.09	2.42857	2.32653	0.04	0	0
C10	1	0	8258	8730	0.06	2.50667	2.28	0.10	13	13
C11	1	0	9142	9450	0.03	2.85	2.79	0.02	4	4
C12	1	0	12716	13018	0.02	2.75333	2.76	0.00	13	13

Table 8: ConVRP Instances- No RT.

6 Conclusions

The TS-VRVW is an effective metaheuristic for solving the periodic, multi-depot VRP with minimal route time solutions in comparable computational times. The solution evaluation framework of the TS-VRVW allowed for the incorporation of competing objectives related to wait-listing and consistent service via an aggregated objective function. Though these objectives were approximations of true cost, they encouraged solution building towards meeting their respective objectives. Many of the test instances did show, however, that the approximation of consistency has a close relationship with the objective of route time and that to achieve minor improvements in consistent service usually meant accepting much poorer route times.

The TS-VRVW is able to solve instances with several hundred clients easily in a small amount of time. Unfortunately, the neighborhood structures examined during the search were too large to handle instances upwards of 1000 locations quickly. If, for the purposes of horizon planning, run times of several hours is acceptable for instances these sizes, then the algorithm will perform with some success. In addition, the use of a consolidation neighborhood structure allowed the search to make large improvements to the current routing schemes. This method should be taken as a holistic implementation though, as allowing the modifier to blindly consolidate demand as done in this work will create solutions that are unfavorable from a client service perspective.

7 Future Work

To further this research, adaptations to the methodology presented can be made to closer model MOW operations. The impacts of delivery consolidation is analyzed in this paper, but when left unrestricted may produce solutions that are not feasible from a service perspective. Unrestrained delivery consolidation has the ability to change solutions drastically by assigning a week's worth of client demand to a single day, given capacity constraints are not violated. This is just an example of an unfavorable result from a client perspective as a result of the consolidation neighborhood structure. Similarly, reorganizing the delivery schedule had positive impacts on the quality of the solutions generated, but at the expense of potentially assembling poorly balanced meal deliveries over the planning horizon. To better utilize these functions, rules governing how much demand may be consolidated, or the days in which certain deliveries must be made would benefit this research.

This research looked at meal delivery as a single product demand to simplify the problem. In reality, client demand comes in the form of both hot and cold meals. Modeling this problem as a mixed-commodity VRP can add more accuracy to the routes generated, as capacity constraints become more stringent in terms of hot and cold box utilization. Additionally, at the completion of this research data was unable to be obtained directly from MOW. In future testing of the TS-VRVW, and work adapted from its framework, it would be beneficial to examine the heuristics performance on realistic instances.

Improvements to the TS-VRVW can be made in the area of diversification by testing methods in addition to frequent modification penalization. Authors proposed techniques such as restarts based on a predefined time span in which no improvement occur. Drastic changes to the current solution may be beneficial for the TS-VRVW as it tends to find the best seen value early

on in the search process. Additionally, intensification is a technique not incorporated into this algorithm that may add value to the search process.

It would be beneficial to enhance the effectiveness of the TS-VRVW by improving its ability to solve large problems in a timely manner. The seven solution modifiers proposed in this work effectively reduce the objective function, but at a significant cost of time in comparison to other Tabu Search methods. Reducing the size of these neighborhood structures or the amount of them effectively will reduce run time. Methods such as partial random neighborhood building, or learning methods, could improve the TS-VRVW in terms of this goal.

The approximation terms INC and WL are incorporated into the objective function of the TS-VRVW to guide the search towards eliminating volunteer inconsistency and reducing the size of the waiting list MOW maintains. It may be worthwhile to further examine and improve the effectiveness of these approximations to improve the overall performance of the heuristic. Additionally, if there is a way to link actual costs of not providing service, or the impact inconsistent service, their incorporation into the objective function of the TS-VRVW could be meaningful outside of the context of a search algorithm and easily translate the operational goals to operational expenses.

8 Bibliography

Applegate, David L., et al. "Certification of an optimal TSP tour through 85,900 cities." *Operations Research Letters* 37, no. 1 (2009): 11-15 .

Archetti, Claudia, Luca Bertazzi, Alain Hertz, and M. Grazia Speranza. "A Hybrid Heuristic for an Inventory Routing Problem." *INFORMS JOURNAL ON COMPUTING*, 2011.

Barbarosoglu, Gulay, and Demet Ozgur. "A tabu search algorithm for the vehicle routing problem." *Computers & Operations Research* 26 (1999): 255-270.

Bazgan, Cristina, Refael Hassin, and J'erome Monnot. "Approximation algorithms for some vehicle routing problems." *Discrete Applied Mathematics* 146, no. 1 (2005): 27-42.

Bowerman, Robert, Brent Hall, and Paul Calami. "A Multiobjective Optimization Approach to Urban School Bus Routing : Formulation and Solution Method." *Transportation Research* 29 (1995): 107-123.

Clarke, G., and J. W. Wright. "Scheduling of Vehicles from a Central Depot to a Number of Delivery Points." *OPERATIONS RESEARCH* 12, no. 4 (1964): 568-581.

Cordeau, Jean-François, Michel Gendreau, Alain Hertz, Gilbert Laporte, and Jean-Sylvain Sormany. "NEW HEURISTICS FOR THE VEHICLE." *Logistics Systems: Design and Optimization*, 2005: 279-297.

Cordeau, Jean-Franois, Michel Gendreau, and Gilbert Laporte. "A Tabu Search Heuristic for Periodic and Multi-Depot Vehicle Routing Problems." *Networks* (Networks) 30, no. 2 (1995): 105-119.

DANTZIG, G, R FULKERSON, and S JOHNSON. "SOLUTION OF A LARGE-SCALE TRAVELING-SALESMAN." *OPERATIONS RESEARCH* 2, no. 4 (1954): 393-410.

Dantzig, G. B., and J. H. Ramser. "The Truck Dispatching Problem." *Management Science* 6, no. 1 (1959): 80-91.

Devlin, Rob. "Volunteer's perceptiveness saves quality of a life." *McClatchy - Tribune Business* , December 10, 2010.

Figueira, J.R., A. Liefooghe, E.-G. Talbi, and A.P. Wierzbicki. "A parallel multiple reference point approach for multi-objective optimization." *European Journal of Operational Research* 205, no. 2 (2010): 390–400.

Fisher, Marshall L. "Optimal Solution of Vehicle Routing Problems Using Minimum K-Trees." *Operations Research*, Julu 1994: 626-642.

Fisher, Marshall L., and Ramchandran Jaikumar. "A generalized assignment heuristic for vehicle routing." *Networks* 11, no. 2 (1981).

Flood, Merrill. "The Traveling-Salesman Problem." *OPERATIONS RESEARCH* 4, no. 1 (1956): 61-75 .

Gendreau, Michel, Alain Hertz, and Gilbert Laporte. "A Tabu Search Heuristic for the Vehicle Routing Problem." *Management Science* 40, no. 10 (1994): 1276-1290.

Gendreau, Michel, Alain Hertz, and Gilbert Laporte. "New Insertion and Postotimization Procedures for the Traveling Salesman Problem." (Operations Research) 40, no. 6 (1992).

Gendreau, Michel, Gilbert Laporte, and René Séguin. "Stochastic Vehicle Routing." *European Journal of Operational Research* 88, no. 1 (1996): 3-12.

Gendreau, Michel, Gilbert Laporte, Christophe Musaraganyi, and Éric D. Taillard. "A Tabu Search Heuristic for the Heterogeneous Fleet Vehicle." *Computers and Operations Research* 26, no. 10 (1999): 140-151.

Gendreau, Michel, Jean-Yves Potvin, Olli Braysy, Geir Hasle, and Arne Lokketangen. "Metaheuristics for the Vehicle Routing Problem and its Extensions: A Categorized Bibliography." In *The Vehicle Routing Problem: Latest Advances and New Challenges*, 143-169. Springer US, 2008.

Glover, Fred. "Tabu Search- Part I." *Operations Research Society of America* 1, no. 3 (1989): 190-206.

Groër, Chris, Bruce Golden, and Edward Wasil. "The Consistent Vehicle Routing Problem." *MANUFACTURING & SERVICE OPERATIONS MANAGEMENT* 11, no. 4 (2009): 630-643.

Gutin, G., and A.P. Punnen. *The Traveling Salesman Problem and Its Variations.* Boston: Kluwer Academic Publishers, 2002.

Hansen, Pierre, and Nenad Mladenovic. "Variable Neighborhood Search: Principles and Applications." *European Journal of operational Research* 130, no. 3 (2001): 449-467.

Howard County Association of Volunteer Administrators. "Get Staff Support for your Volunteer Program." *HCAVA.* December 29, 2010. http://howardcountyava.org/2010/12/get-staff-support-for-your-volunteer-program/ (accessed March 15, 2011).

Karabuk, Suleyman. "Modeling and optimizing transportation decisions." *Transportation Research Part E: Logistics and Transportation Review* 43, no. 4 (2005): 321-337.

King, Paul. *Feeding Seniors at Home.* April 18, 2005. http://www.nrn.com/ Keyword: Meals on Wheels (accessed January 5, 2011).

Kulturel-Konak, Sadan, Alice E. Smith, and Bryan A. Norman. "Multi-objective tabu search using a multinomial probability mass function." *European Journal of Operational Research* 169, no. 3 (2006): 918-931.

Laporte, Gilbert. "The Vehicle Routing Problem: An overview." *European Journal of Operational Research* 59, no. 3 (1992): 345-358 .

Laporte, Gilbert. "The Vehicle Routing Problem: An overview of exact and approximate algorithms." *European Journal of Operational Research*, no. 59 (1992): 354-358.

Laporte, Gilbert, Hélène Mercure, and Yves Nobert. "An exact algorithm for the asymmetrical capacitated vehicle routing problem." *Networks* 16, no. 1 (1986): 33-46.

Meals On Wheels Association of America. *2009 Annual Report.* Alexandria: Meals On Wheels Association of America, Inc., 2010.

Miliotis, P. "Integer Programming Approaches to the Traveling Salesman Problem." *Mathematical Programming* 10, no. 1 (1976): 367-378.

Ngatchou, Patrick, Anahita Zarei, and M. A. El-Sharkawi. "Pareto Multi Objective Optimization." *Proceedings of the 13th International Conference on Intelligent Systems Application to Power Systems.* IEEE, 2005. 84-91.

O'Shaughnessy, Carol. *Older Americans Act Nutrition Program.* Domestic Social Policy Division, 2004.

Parthanadee, Parthana, and Rasartnam Logendran. "Periodic product distribution from multi-depots under limited supplies." *IIE Transactions* 38, no. 11 (2006).

Rardin, Ronald L. *Optimization in Operations Research.* Upper Saddle River, NJ: Prentice Hall Inc., 1998.

Rochat, Yves, and Éric D. Taillard. "PROBABILISTIC DIVERSIFICATION AND INTENSIFICATION IN LOCAL SEARCH FOR VEHICLE ROUTING." *Journal of Heuristics* 1, no. 1 (1995): 147-167.

Schwarz, Leroy B., James E. Ward, and Xin Zhai. "On the Interactions Between Routing and Inventory-Management Policies in a One-Warehouse N-Retailer Distribution System." *Manufacturing & Service Operations Management* 8, no. 3 (2006): 253-272.

Shiguemoto, Andre´ Luı´s, and Vinı´cius Amaral Armentano. "A tabu search procedure for coordinating production, inventory and distribution routing problems." *International Transactions in Operational Research* 17 (2010): 179–195.

Smilowitz, Karen, Maciek Nowak, and Tingting Jiang. "Workforce management in periodic delivery." *Transportation Science (Accepted)*, 2011.

Steeg, Jörg, and Michael Schröder. *A Hybrid Approach to Solve the Periodic Home Health Care Problem.* Vol. 2007, in *Operations Research Proceedings*, 297-302. 2008.

Wassan, NA. "A reactive tabu search for the vehicle routing." *Operational Research Society* 15, no. 4 (2006): 368-386.

9 Appendix

9.1 Detailed Weight Analysis Results

9.1.1 Impact of INC on RT or WL

Figure 17 shows the relationship between RT and INC when consolidation was not allowed during the search while Figure 18 shows the relationship when consolidation is used. When consolidation is not allowed in the search, Figure 17 shows the degradation of RT as AVC is minimized. An interesting observation is that the search produced all around poor solutions when INC was weighted too heavily. When consolidation is allowed a less expected relationship is noticed. As the value of AVC lowers, so does the differential of route time. This can mainly be attributed to consolidation decisions that group client demand from several days to a single volunteer on fewer trips.

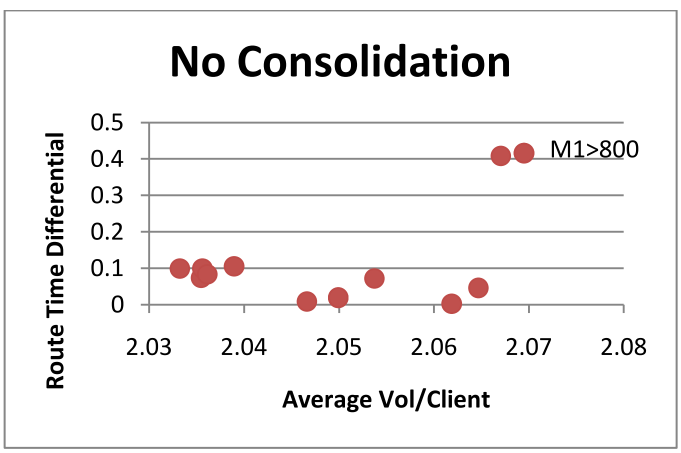

Figure 17: Effect of INC on RT (2obj-NC).

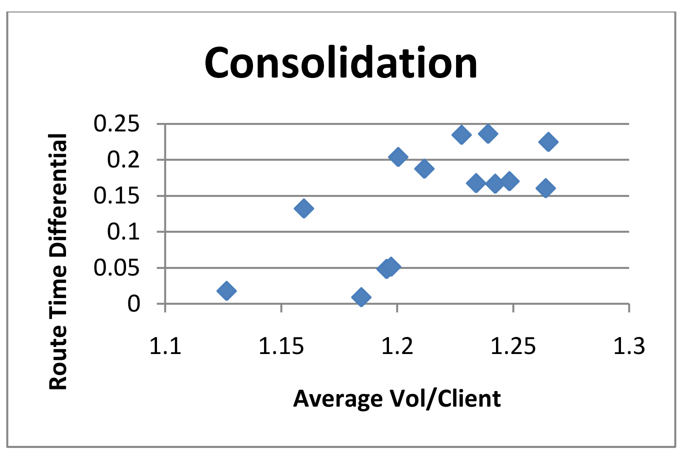

Figure 18: Effects of INC on RT (2obj-C).

Examining the relationship between INC and WL shows that there are no drastic decreases in AVC given the weight assigned (Figure 19). Allowing consolidation also does not produce dramatic differences in ACV (Figure 20). An interesting relationship produced, though, is the trend between the reduction of AVC and the reduction of the waiting list. This is resultant of taking demand from the waiting list and immediately consolidating it one volunteer at a time.

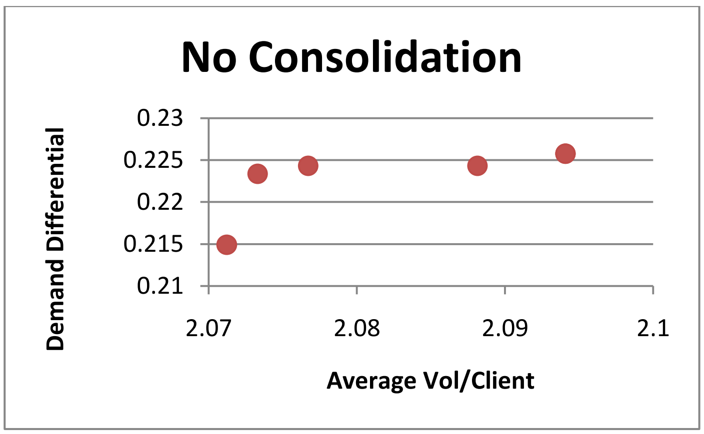

Figure 19: Effect of INC on WL (2obj-NC).

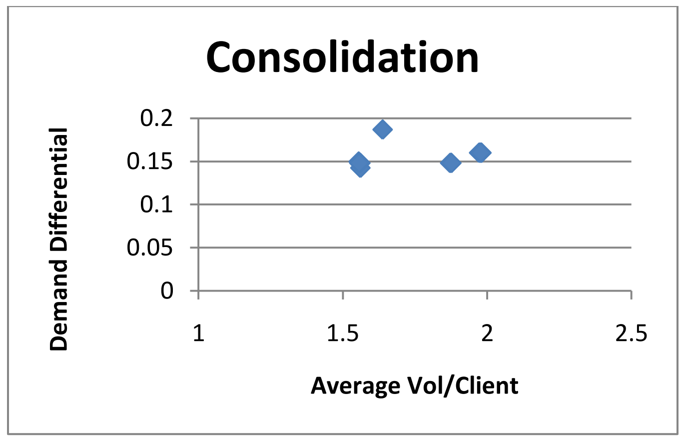

Figure 20: Effect of INC on WL (2obj-C).

9.1.2 Impact of WL on RT or INC

The relationship between WL and RT when no consolidation is allowed implies that when WL largely degrades RT in a solution. The data points labeled in Figure 21 show the drastic change in RT differential once the term is included. Allowing consolidation in the search algorithm produced results that aligned with traditional routing methods. As more emphasis was placed on servicing wait listed clients, the differential of RT grew (Figure 22). However, once a large preference was placed on WL, a reduction in both wait listed demand and RT differential was noticed. Comparing WL to INC shows that the objective of WL has no substantial effect on

AVC when there is no consolidation in the search, as it only improved .05 while the demand differential suffered by over .2. In general, heavily weighting WL only worsened the solution found by the TS-VRVW (Figure 23). If the search was allowed consolidation, it is observed that there is not a clear relationship between WL objective and INC. Varying levels of weights did not produce a consistent cost tradeoff between meeting more demand and providing consistent service (Figure 24).

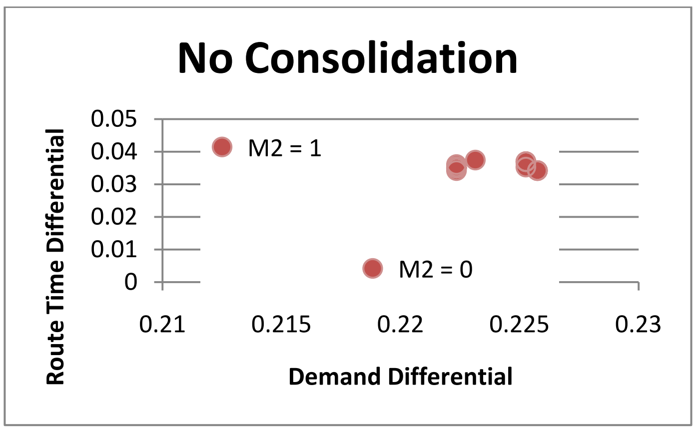

Figure 21: Effect of WL on RT (2obj-NC).

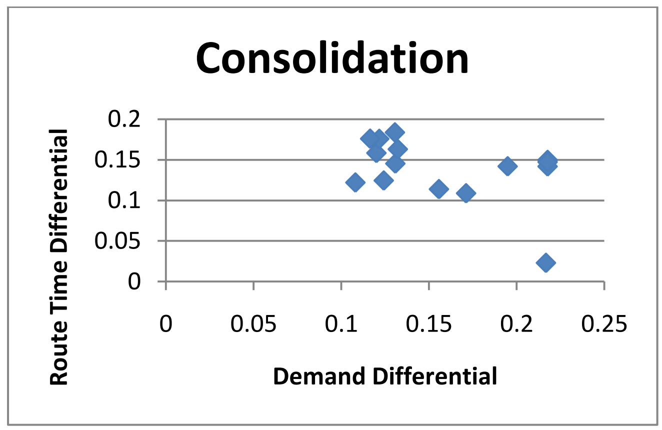

Figure 22: Effect of WL on RT (2obj-NC).

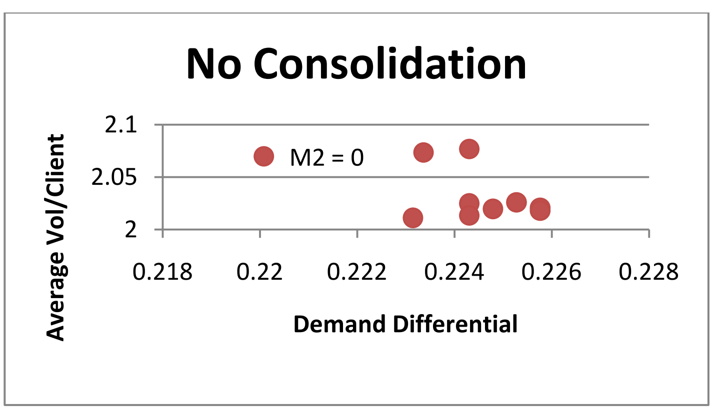

Figure 23: Effect of WL on INC (2obj-NC).

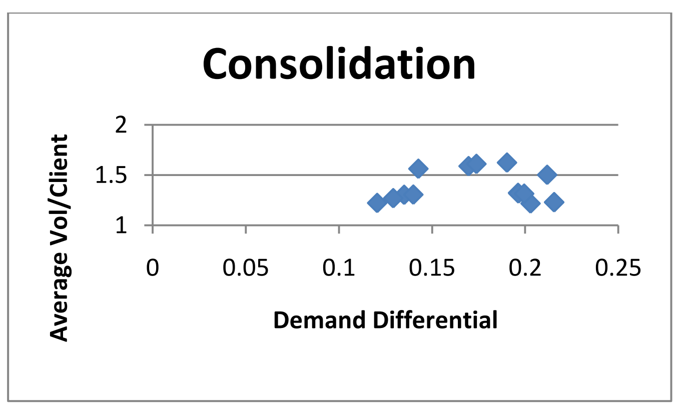

Figure 24: Effect of WL on INC (2obj-C).

9.1.3 Impact of INC on RT and WL

Comparing the effect INC has on RT shows that there is a positive correlation between AVC and route time differential. While consolidation is in effect, the RT differential rose steadily as the AVC increased (Figure 25). This can be attributed to the strength of the consolidation neighborhood structures ability to satisfy both objectives concurrently. Without consolidation, a similar relationship was observed between AVC and the RT differential, but consistency degraded much faster than route time (Figure 26). This trend is broken at large values of M1 as both objectives degrade substantially. This correlation implies that the AVC on average is reflective of the quality of the routes a solution produced. Increasing the weight of INC drastically only seems to degrade the overall quality in this regard.

The effect INC has on the WL objective show similar trends in comparison to the RT objective. As the AVC degrades, so does the demand differential in the case for both consolidation and consolidation free search experiments. When consolidation was allowed, the AVC degrading at by a larger magnitude ranging from 1.2 to 2 AVC, while the demand differential only varied between .175 and .25 (Figure 28). Without consolidation the values for both AVC and the demand differential were poorer, but did not vary largely (Figure 27). In both cases though, the introduction of INC to the objective degraded the wait-listing objective greatly.

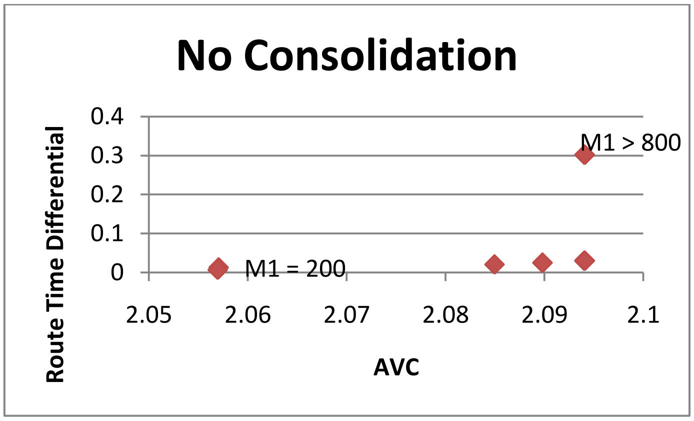

Figure 26: Effect of INC on RT (3obj-NC).

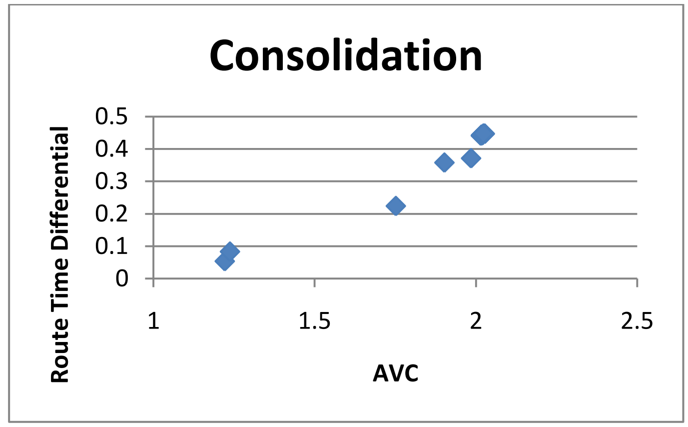

Figure 25: Effect of INC on RT (3obj-C).

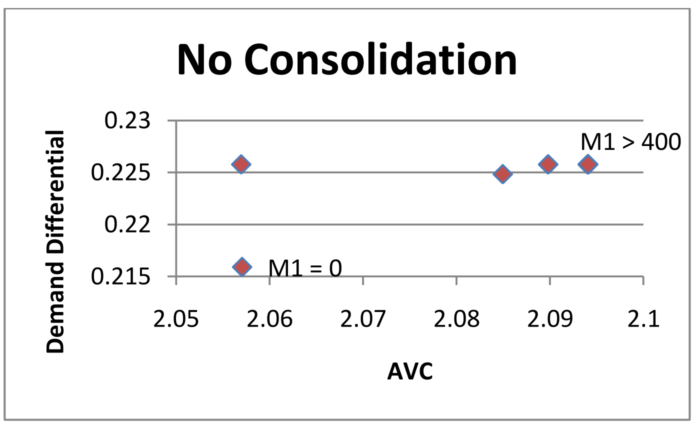

Figure 27:Effect of INC on WL (3obj-NC).

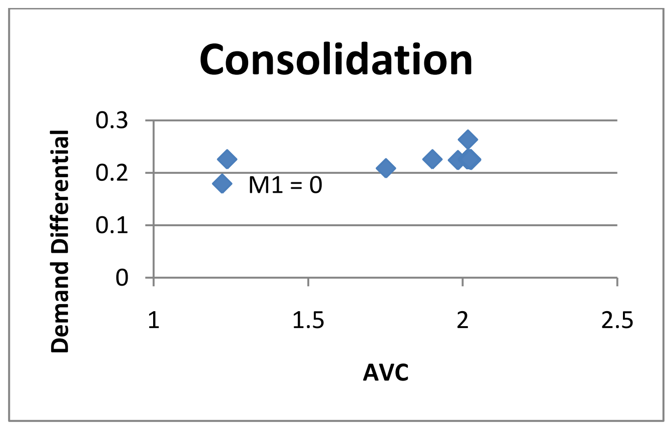

Figure 28: Effect of INC on WL (3obj-C).

9.1.4 Impact of WL on RT and INC

Observing the interaction between WL and RT when consolidation was active produced a surprising result. As the demand differential increased, the route time decreased, but then increased again in a parabolic manner (Figure 29). This can be explained by the improvements realized by reconfiguring routes with previously unmet demand. However, these improvements only occur until the used capacity is fully optimized, then route time is again degraded as more demand is met. Without consolidation a more traditional pattern is seen where route time degrades as more demand is met (Figure 30).

When comparing WL and INC with consolidation there was a positive correlation observed between AVC and the demand differential (Figure 31). As before, this is attributed to the ease of meeting both objectives concurrently with consolidation efforts. Without consolidation, there was not a clear relationship between the effect WL had on AVC (Figure 32). More often, though, AVC tended to be lower as less demand was met.

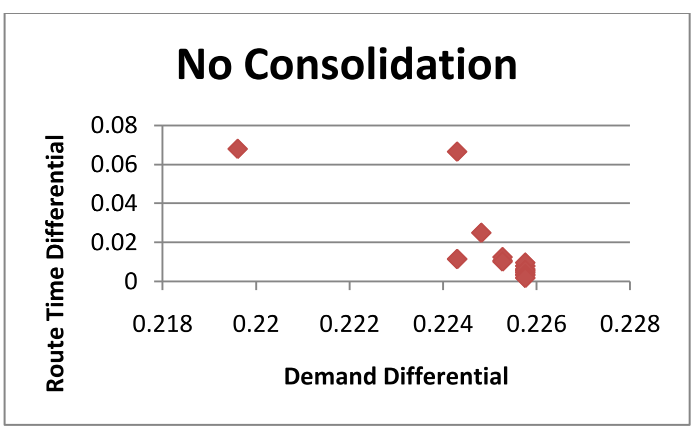

Figure 30: Effect of WL on RT (3obj-NC).

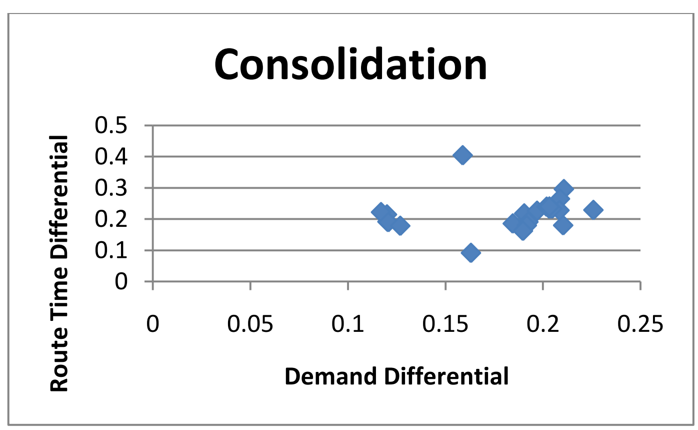

Figure 29: Effect of WL on RT (3obj-C).

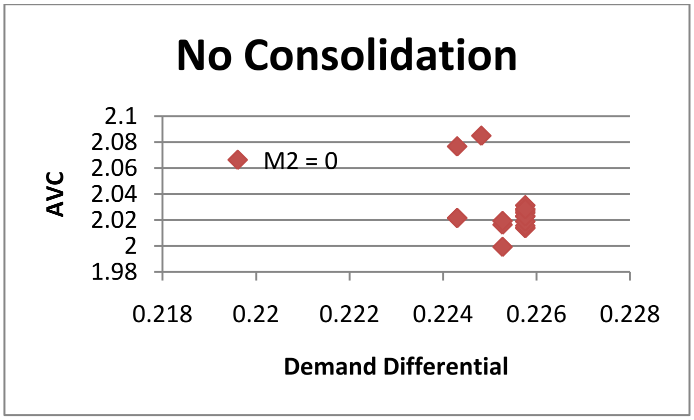

Figure 32: Effect of WL on INC (3obj-NC).

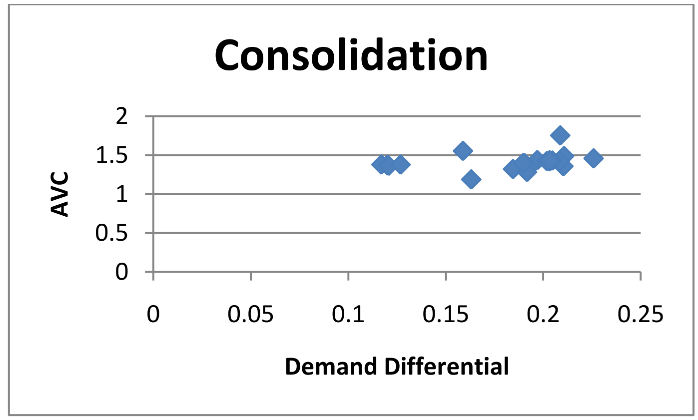

Figure 31: Effect of WL on INC (3obj-C).

9.2 Result Tables

Table 9: Parameter Testing on δ Over Random Instances 1-6, and ConVrp Problems 1-3.

Value	Avg Gap F(s)
0.1	0.396
0.2	0.369
0.3	0.363
0.4	0.289
0.5	0.297

Table 11: Parameter Testing on Exn Over Random Instances 1-6, and ConVrp Problems 1-3.

Value	Avg Gap F(s)	Avg Iterations
1	0.233	30329.02
2	0.253	28206.28
3	0.248	29763.24
4	0.248	29891.06
5	0.264	29673.94

Table 10: Average Performance of Wait Listing Objectives over Random Instances (1, 2, 3, 10, 11).

%Demand Met		
WL	C	0.61
	NC	0.43
RT	C	0.56
	NC	0.41
RT+100*WL	C	0.56
	NC	0.42

Table 13: Percentage of Modification Usage Broken Down by Objective.

	Objective	INC	RT	WL
Mod 1	Consolidation	0.06	0.06	0.11
Mod 2	Consolidation	0.00	0.02	0.02
Mod 3	Consolidation	0.00	0.02	0.01
Mod 4	Consolidation	0.02	0.02	0.03
Mod 5	Consolidation	0.02	0.06	0.64
Mod 6	Consolidation	0.01	0.02	0.03
Mod 7	Consolidation	0.89	0.78	0.17
Mod 1	No Consolidation	0.50	0.42	0.03
Mod 2	No Consolidation	0.04	0.11	0.00
Mod 3	No Consolidation	0.02	0.01	0.00
Mod 4	No Consolidation	0.06	0.04	0.00
Mod 5	No Consolidation	0.33	0.39	0.94
Mod 6	No Consolidation	0.07	0.05	0.03
Mod 7	No Consolidation	0.00	0.00	0.00

Table 12: Average Performance of Consistency Objectives over Random Instances (1, 2, 3, 10, 11).

% Clients Served by:		1 Volunteer	2 Volunte	3+ Volunteers
INC	C	0.74	0.17	0.09
	NC	0.51	0.27	0.21
RT	C	0.68	0.22	0.11
	NC	0.49	0.29	0.22
RT+100*INC	C	0.68	0.20	0.12
	NC	0.51	0.28	0.21

Table 14: Normalization of INC+RT and INC+WL (No Consolidation).

No Consolidation						
	Obj	Weight	Norm	RT/MinRT	AVC	TD/DM
M1	INC+RT	0	4.54	1.00	2.06	1.48
M1	INC+RT	1	4.53	1.01	2.05	1.47
M1	INC+RT	25	4.59	1.07	2.04	1.48
M1	INC+RT	50	4.62	1.10	2.04	1.47
M1	INC+RT	75	4.60	1.10	2.03	1.47
M1	INC+RT	100	4.61	1.10	2.04	1.47
M1	INC+RT	200	4.54	1.02	2.05	1.47
M1	INC+RT	300	4.54	1.02	2.05	1.47
M1	INC+RT	400	4.54	1.02	2.05	1.47
M1	INC+RT	500	4.59	1.08	2.04	1.47
M1	INC+RT	600	4.60	1.07	2.05	1.47
M1	INC+RT	700	4.58	1.05	2.06	1.47
M1	INC+RT	800	4.59	1.05	2.07	1.47
M1	INC+RT	900	4.62	1.06	2.07	1.49
M1	INC+RT	1000	4.62	1.06	2.07	1.49
M1	INC+WL	0	4.55	1.04	2.07	1.44
M1	INC+WL	1	4.59	1.03	2.07	1.49
M1	INC+WL	25	4.59	1.00	2.09	1.50
M1	INC+WL	50	4.59	1.00	2.09	1.50
M1	INC+WL	75	4.59	1.00	2.09	1.50
M1	INC+WL	100	4.59	1.00	2.09	1.50
M1	INC+WL	200	4.60	1.03	2.08	1.49
M1	INC+WL	300	4.60	1.03	2.08	1.49
M1	INC+WL	400	4.60	1.03	2.08	1.49
M1	INC+WL	500	4.58	1.00	2.09	1.49
M1	INC+WL	600	4.58	1.00	2.09	1.49
M1	INC+WL	700	4.58	1.00	2.09	1.49
M1	INC+WL	800	4.59	1.00	2.09	1.50
M1	INC+WL	900	4.59	1.00	2.09	1.50
M1	INC+WL	1000	4.59	1.00	2.09	1.50

Table 15: Normalization of INC+RT and INC+WL (Consolidation).

Consolidation						
	Obj	Weight	Norm	RT/MinRT	AVC	TD/DM
M1	INC+RT	0	3.69	1.01	1.18	1.49
M1	INC+RT	1	3.62	1.02	1.13	1.47
M1	INC+RT	25	3.98	1.22	1.27	1.49
M1	INC+RT	50	3.95	1.23	1.23	1.49
M1	INC+RT	75	3.96	1.24	1.24	1.49
M1	INC+RT	100	3.89	1.20	1.20	1.49
M1	INC+RT	200	3.73	1.05	1.20	1.49
M1	INC+RT	300	3.73	1.05	1.20	1.49
M1	INC+RT	400	3.74	1.05	1.20	1.49
M1	INC+RT	500	3.89	1.19	1.21	1.49
M1	INC+RT	600	3.90	1.17	1.24	1.49
M1	INC+RT	700	3.78	1.13	1.16	1.49
M1	INC+RT	800	3.91	1.16	1.26	1.49
M1	INC+RT	900	3.89	1.17	1.23	1.49
M1	INC+RT	1000	3.91	1.17	1.25	1.49
M1	INC+WL	0	3.99	1.02	1.64	1.33
M1	INC+WL	1	3.81	1.01	1.56	1.24
M1	INC+WL	25	4.33	1.07	1.97	1.29
M1	INC+WL	50	4.34	1.07	1.98	1.29
M1	INC+WL	75	4.33	1.07	1.98	1.29
M1	INC+WL	100	4.33	1.07	1.98	1.29
M1	INC+WL	200	3.81	1.00	1.56	1.26
M1	INC+WL	300	3.83	1.01	1.56	1.26
M1	INC+WL	400	3.82	1.01	1.56	1.26
M1	INC+WL	500	4.21	1.08	1.87	1.26
M1	INC+WL	600	4.21	1.08	1.88	1.26
M1	INC+WL	700	4.21	1.08	1.88	1.26
M1	INC+WL	800	4.34	1.07	1.98	1.29
M1	INC+WL	900	4.34	1.07	1.98	1.29
M1	INC+WL	1000	4.34	1.07	1.98	1.29

Table 16: Normalization of WL+RT and WL+INC (No Consolidation).

No Consolidation						
	Obj	wts	Norm	RT/MinRT	AVC	TD/DM
M2	WL+RT	0	4.53	1.00	2.06	1.47
M2	WL+RT	1	4.57	1.04	2.10	1.43
M2	WL+RT	25	4.60	1.04	2.08	1.48
M2	WL+RT	50	4.62	1.04	2.09	1.50
M2	WL+RT	75	4.62	1.04	2.09	1.50
M2	WL+RT	100	4.62	1.04	2.09	1.50
M2	WL+RT	200	4.62	1.04	2.10	1.49
M2	WL+RT	300	4.62	1.04	2.10	1.49
M2	WL+RT	400	4.62	1.04	2.10	1.49
M2	WL+RT	500	4.62	1.04	2.09	1.49
M2	WL+RT	600	4.61	1.04	2.09	1.49
M2	WL+RT	700	4.61	1.04	2.09	1.49
M2	WL+RT	800	4.62	1.03	2.09	1.50
M2	WL+RT	900	4.62	1.03	2.09	1.50
M2	WL+RT	1000	4.62	1.03	2.09	1.50
M2	WL+INC	0	4.55	1.02	2.07	1.47
M2	WL+INC	1	4.57	1.01	2.07	1.49
M2	WL+INC	25	4.63	1.13	2.01	1.48
M2	WL+INC	50	4.63	1.12	2.02	1.49
M2	WL+INC	75	4.64	1.12	2.03	1.50
M2	WL+INC	100	4.65	1.13	2.03	1.50
M2	WL+INC	200	4.57	1.01	2.08	1.49
M2	WL+INC	300	4.57	1.01	2.08	1.49
M2	WL+INC	400	4.57	1.01	2.08	1.49
M2	WL+INC	500	4.64	1.12	2.02	1.49
M2	WL+INC	600	4.64	1.13	2.01	1.49
M2	WL+INC	700	4.64	1.13	2.01	1.49
M2	WL+INC	800	4.64	1.12	2.02	1.50
M2	WL+INC	900	4.64	1.12	2.02	1.50
M2	WL+INC	1000	4.63	1.12	2.02	1.50

Table 17: Normalization of WL+RT and WL+INC (Consolidation).

Consolidation						
	Obj	wts	Norm	RT/MinRT	AVC	TD/DM
M2	WL+RT	0	3.97	1.02	1.48	1.47
M2	WL+RT	1	3.92	1.11	1.52	1.29
M2	WL+RT	25	3.97	1.18	1.57	1.22
M2	WL+RT	50	3.88	1.16	1.52	1.20
M2	WL+RT	75	3.79	1.15	1.44	1.21
M2	WL+RT	100	3.92	1.16	1.54	1.22
M2	WL+RT	200	3.83	1.12	1.51	1.19
M2	WL+RT	300	4.08	1.14	1.57	1.36
M2	WL+RT	400	3.92	1.11	1.55	1.25
M2	WL+RT	500	4.21	1.15	1.59	1.47
M2	WL+RT	600	4.15	1.14	1.54	1.47
M2	WL+RT	700	4.20	1.15	1.58	1.47
M2	WL+RT	800	3.89	1.15	1.54	1.20
M2	WL+RT	900	3.87	1.15	1.53	1.19
M2	WL+RT	1000	3.71	1.10	1.44	1.16
M2	WL+INC	0	4.04	1.09	1.50	1.45
M2	WL+INC	1	3.92	1.12	1.56	1.24
M2	WL+INC	25	3.77	1.09	1.32	1.37
M2	WL+INC	50	3.77	1.08	1.31	1.37
M2	WL+INC	75	3.46	1.05	1.22	1.19
M2	WL+INC	100	3.61	1.07	1.30	1.24
M2	WL+INC	200	3.99	1.12	1.59	1.28
M2	WL+INC	300	4.09	1.12	1.62	1.35
M2	WL+INC	400	4.03	1.12	1.61	1.29
M2	WL+INC	500	3.77	1.08	1.23	1.47
M2	WL+INC	600	3.70	1.06	1.22	1.42
M2	WL+INC	700	3.77	1.08	1.23	1.47
M2	WL+INC	800	3.62	1.07	1.30	1.24
M2	WL+INC	900	3.56	1.07	1.27	1.22
M2	WL+INC	1000	3.82	1.08	1.31	1.42

Table 18: Normalization of WL+RT+INC (No Consolidation).

No Consolidation				bc	bd	bh
	Obj	wts	Norm	RT/MinRT	AVC	TD/DM
M2	WL+RT+INC	0	4.60	1.07	2.07	1.47
M2	WL+RT+INC	1	4.60	1.02	2.08	1.49
M2	WL+RT+INC	50	4.52	1.01	2.02	1.50
M2	WL+RT+INC	100	4.53	1.01	2.02	1.50
M2	WL+RT+INC	150	4.51	1.01	2.00	1.50
M2	WL+RT+INC	200	4.64	1.07	2.08	1.49
M2	WL+RT+INC	250	4.52	1.01	2.01	1.50
M2	WL+RT+INC	300	4.53	1.01	2.02	1.49
M2	WL+RT+INC	350	4.54	1.01	2.03	1.50
M2	WL+RT+INC	400	4.53	1.01	2.02	1.49
M2	WL+RT+INC	450	4.54	1.01	2.03	1.50
M2	WL+RT+INC	500	4.52	1.00	2.01	1.50
M2	WL+RT+INC	550	4.52	1.00	2.02	1.50
M2	WL+RT+INC	600	4.53	1.01	2.02	1.50
M2	WL+RT+INC	650	4.52	1.01	2.02	1.50
M2	WL+RT+INC	700	4.53	1.01	2.03	1.50
M2	WL+RT+INC	750	4.53	1.01	2.03	1.50
M2	WL+RT+INC	800	4.52	1.00	2.02	1.50
M2	WL+RT+INC	850	4.52	1.00	2.02	1.50
M2	WL+RT+INC	900	4.52	1.00	2.02	1.50
M2	WL+RT+INC	950	4.53	1.00	2.03	1.50
M2	WL+RT+INC	1000	4.53	1.00	2.03	1.50

Table 19: Normalization of WL+RT+INC (Consolidation).

Consolidation						
	Obj	wts	Norm	RT/MinRT	AVC	TD/DM
M2	WL+RT+INC	0	4.22	1.30	1.49	1.44
M2	WL+RT+INC	1	4.42	1.26	1.75	1.41
M2	WL+RT+INC	50	3.97	1.22	1.40	1.35
M2	WL+RT+INC	100	4.03	1.23	1.43	1.37
M2	WL+RT+INC	150	3.84	1.19	1.32	1.34
M2	WL+RT+INC	200	4.21	1.40	1.55	1.25
M2	WL+RT+INC	250	3.89	1.19	1.34	1.36
M2	WL+RT+INC	300	3.81	1.18	1.28	1.35
M2	WL+RT+INC	350	4.08	1.23	1.42	1.43
M2	WL+RT+INC	400	3.54	1.09	1.19	1.26
M2	WL+RT+INC	450	4.09	1.24	1.42	1.43
M2	WL+RT+INC	500	3.91	1.16	1.40	1.35
M2	WL+RT+INC	550	3.99	1.18	1.36	1.45
M2	WL+RT+INC	600	4.09	1.23	1.42	1.44
M2	WL+RT+INC	650	4.10	1.23	1.43	1.44
M2	WL+RT+INC	700	4.18	1.23	1.46	1.50
M2	WL+RT+INC	750	4.10	1.24	1.43	1.43
M2	WL+RT+INC	800	3.79	1.21	1.37	1.20
M2	WL+RT+INC	850	3.76	1.19	1.36	1.20
M2	WL+RT+INC	900	3.80	1.22	1.38	1.20
M2	WL+RT+INC	950	3.76	1.19	1.36	1.20
M2	WL+RT+INC	1000	3.77	1.18	1.38	1.21

Table 20: Normalization of INC+RT+WL (No Consolidation).

No Consolidation						
	Obj	wts	Norm	RT/MinRT	AVC	TD/DM
M1	INC+RT+WL	0	4.51	1.01	2.06	1.44
M1	INC+RT+WL	1	4.60	1.02	2.08	1.49
M1	INC+RT+WL	50	4.62	1.03	2.09	1.50
M1	INC+RT+WL	100	4.62	1.03	2.09	1.50
M1	INC+RT+WL	150	4.62	1.03	2.09	1.50
M1	INC+RT+WL	200	4.56	1.01	2.06	1.50
M1	INC+RT+WL	250	4.62	1.03	2.09	1.50
M1	INC+RT+WL	300	4.61	1.02	2.09	1.50
M1	INC+RT+WL	350	4.62	1.03	2.09	1.50
M1	INC+RT+WL	400	4.61	1.02	2.09	1.50
M1	INC+RT+WL	450	4.62	1.03	2.09	1.50
M1	INC+RT+WL	500	4.62	1.03	2.09	1.50
M1	INC+RT+WL	550	4.62	1.03	2.09	1.50
M1	INC+RT+WL	600	4.62	1.03	2.09	1.50
M1	INC+RT+WL	650	4.62	1.03	2.09	1.50
M1	INC+RT+WL	700	4.62	1.03	2.09	1.50
M1	INC+RT+WL	750	4.62	1.03	2.09	1.50
M1	INC+RT+WL	800	4.62	1.03	2.09	1.50
M1	INC+RT+WL	850	4.62	1.03	2.09	1.50
M1	INC+RT+WL	900	4.62	1.03	2.09	1.50
M1	INC+RT+WL	950	4.62	1.03	2.09	1.50
M1	INC+RT+WL	1000	4.62	1.03	2.09	1.50

Table 21: Normalization of INC+RT+WL (Consolidation).

Consolidation						
	Obj	wts	Norm	RT/MinRT	AVC	TD/DM
M1	INC+RT+WL	0	3.59	1.05	1.22	1.31
M1	INC+RT+WL	1	4.38	1.22	1.75	1.41
	INC+RT+WL	50	4.95	1.44	2.02	1.50
M1	INC+RT+WL	100	4.96	1.44	2.02	1.50
M1	INC+RT+WL	150	4.95	1.44	2.02	1.50
M1	INC+RT+WL	200	3.82	1.08	1.24	1.50
M1	INC+RT+WL	250	4.95	1.44	2.02	1.50
M1	INC+RT+WL	300	4.76	1.36	1.90	1.50
M1	INC+RT+WL	350	4.96	1.44	2.02	1.50
M1	INC+RT+WL	400	4.97	1.45	2.02	1.50
M1	INC+RT+WL	450	4.96	1.44	2.02	1.50
M1	INC+RT+WL	500	4.84	1.37	1.98	1.49
M1	INC+RT+WL	550	4.97	1.45	2.03	1.50
M1	INC+RT+WL	600	4.96	1.45	2.03	1.49
M1	INC+RT+WL	650	4.97	1.45	2.03	1.50
M1	INC+RT+WL	700	4.97	1.45	2.03	1.50
M1	INC+RT+WL	750	4.96	1.44	2.02	1.50
M1	INC+RT+WL	800	4.96	1.44	2.02	1.50
M1	INC+RT+WL	850	4.96	1.44	2.02	1.50
M1	INC+RT+WL	900	4.96	1.44	2.02	1.50
M1	INC+RT+WL	950	4.96	1.44	2.02	1.50
M1	INC+RT+WL	1000	4.97	1.45	2.02	1.50