Network Design Analysis for Special Needs Student Services

Scott J. Mason, PhD (PI)
Edward A. Pohl, PhD (Co-PI)

MBTC 3019
June 2010

MBTC DOT 3019 Final Report:
Network Design Analysis for Special Needs Student Services

Faculty Investigators
Scott J. Mason, PhD (PI)
Edward A. Pohl, PhD (Co-PI)

Research Assistants
Behrooz Kamali, BSIE (GRA)
Elizabeth A. McCallion (UGRA)

Abstract

Population growth can lead to public school capacity issues, as well as increased school bus utilization. This increased utilization, in turn, can result in longer school bus transport times for both regular and special needs/medically fragile students. Special needs or medically fragile students are children with special health care needs who are at increased risk for a chronic physical, developmental, behavioral, or emotional condition. These students require health and related services of a type or amount beyond that required by typical children. It is common practice to provide special needs students with specially equipped buses and/or special classroom environments with specific facilities or services. However, the assignment of student services to schools is regularly made without regard to bus transportation considerations for special needs students. Considering the potentially negative impact of long school bus rides on these students, we present the first systematic, integrated analyses of special needs student busing and classroom assignments. We provide models and algorithms for maintaining administration-based transportation financial performance measures while simultaneously designing smarter transportation networks. The smarter networks produced by our models assign special needs services to schools in concert with considering both student geographical location and service needs. In the future, we hope to pilot our model results in local school districts to assess the efficacy of our proposed methods in practice.

Table of Contents

List of Figures

List of Tables

1 Introduction

As urban areas grow in population, some people choose to relocate to the suburbs, often for “more space”—to be more spread out across suburban neighborhood areas. One of the main public services that is impacted by these city-to-suburb moves are rural public education systems. When a school district grows both in terms of its number of schools and its geographic area, school capacity limitations and student bus transportation can become important challenges. Ineffectively making student-to-school assignments and/or inefficient bus routing plans can result in longer school bus rides for students. The magnitude of these inefficiencies is further magnified when one considers the transportation of special needs students.

According to McPherson *et al.* (1998), special needs or medically fragile students are “children with special health care needs…who have or are at increased risk for a chronic physical, developmental, behavioral, or emotional condition and who also require health and related services of a type or amount beyond that required by children generally.” Given this characterization, it follows that longer school bus rides caused by the planning inefficiencies described above can adversely impact special needs students.

Special needs students typically require special buses and/or special classroom environments with specific facilities or services. Based on the severity of their physical and emotional needs, special needs students are placed into a class containing a specific teacher-to-student ratio, such as a 1:6 class containing a maximum of six students and one teacher. Additionally, 1:10 and 1:15 classrooms are typically found in practice. Students in the latter classroom type typically have less or fewer needs for services than do special needs students in classrooms with a fewer number of students.

In terms of busing, special needs buses often have facilities for picking up and dropping off students in wheelchairs. Therefore, not all buses in a school district can be used for special needs student transport. In terms of service needs, the special services required by special needs students typically are not offered in all schools in a school district—often, they are offered in less than one half of the district's schools. It follows that these limited busing and services options can result in one or more special needs student being assigned to a school that is not necessarily close to his/her home. Finally, this fact leads to the aforementioned challenge of longer special needs student bus transportation times.

Interviews with school district officials suggest that current practice is for school administrators to assign special needs services to district schools based on either experience and/or principal requests, often with little or no consideration of where the special needs students reside. In one extreme case, we were told about a special needs student who rides her bus two hours each way to and from school, every day. As both the assignment of services and of students to schools is somewhat subjective and currently is not supported by any type of analytical models in the school districts we investigated, it is quite possible that model-supported assignment decisions can help impact current special needs student transportation practices by providing better transportation and special needs service assignments for school districts.

In this thesis, we investigate this important problem in three phases. First, we develop models and algorithms for assigning special needs students to schools such that all required special service needs and classroom capacity constraints are satisfied while minimizing the total distance all students live from their assigned school. The distances used in Phase 1 of this thesis relate to the Euclidean distance between each special needs student's home and their assigned school. Next, we integrate the assignment results from Phase 1 into a vehicle routing model in

Phase 2 in order to produce an actual, implementable special needs student bus routing plan that minimizes the total amount of time special needs students spend on their school bus each day. Given the complexity of the problem under study, we present three heuristics for analyzing this challenging problem in Phase 3 of this research effort and assess the efficacy of these decision rules on helping to minimize special needs student bus transportation times in practice.

2 Previous Research

A number of previous research studies investigate both assignment and transportation models. Unfortunately, only a small portion of the existing literature focuses on special needs students. Further, most special needs student-focused studies either present case study results or do not examine transportation-related issues. However, it is important to understand the current body of knowledge in order to effectively address the problem under study in this thesis. In the sections that follow, we review previous research related both to special needs students and to the research problem of interest in this thesis.

2.1. The Assignment Problem

The assignment problem for special needs students discussed earlier is similar to the generalized assignment problem in many ways. Generally, assignment problems can be thought of as having a number of agents and a number of tasks. Each agent should be assigned to one task under some conditions in order to accomplish some total job with minimal cost/maximal value. In the research problem of interest, the agents are special needs students and the tasks are available seats or positions in special needs classrooms at district schools of the previously defined types (i.e., 1:6, 1:10, or 1:15).

Pentico (2007) reviews many of the assignment models available in the literature and divides the models into three categories: 1) models with at most one task per agent, 2) models with multiple tasks per agent, and 3) multi-dimensional assignment models. Based on this classification scheme, the problem under study falls into the first category, at most one task per agent (i.e., at most one classroom assignment per special needs student). Further, Pentico (2007) divides all problems in this first category into single and multiple criteria problems.

Among the models he describes, the semi-assignment problem has the greatest similarity to the problem under study. In semi-assignment problems, 1) each agent should be assigned to exactly one task and 2) there are a limited numbers of task groups, each of which requires some number of agents. Clearly, 1) each special needs student must be assigned to a specific classroom and 2) special needs classes, which are limited in number across a school district, have a limited number of spaces as governed by the needs level. Pentico (2007) states that semi-assignment problems can be solved very quickly for large scale problems.

Lee and Schniederjans (1983) develop an assignment model for assigning teachers to schools. They investigate a multi-criteria problem using goal programming for two objectives: cost minimization and maximization of preference goals. Teachers express different preferences towards working in different schools and these preferences are integrated into the model as an objective function. Therefore, a priority is assigned to each goal in the model. Lee and Schniederjans (1983) solve the model under different priority ranking schemes and are able to find some solution combinations that satisfy a range of stated goals.

Ferland and Guenette (1990) develop a decision support system for school districts to assign groups of students to a school. They develop a student network and use heuristic procedures to assign the network's edges (i.e., students) to schools. In this study, the objective is to assign students to schools such that the total distance cost is minimized (i.e., student-to-school proximity is maximized).

2.2. The Transportation Problem

There exists some previous transportation literature that is related to the general school bus routing problem, such as the Vehicle Routing Problem with Pickups and Deliveries (VRPPD) and the Dial-A-Ride Problem (DARP). However, the bus routing problem for special needs students is different from the general student bus routing problem as special needs students often require door-to-door service. It is possible to consider each special needs student's home as an individual bus stop containing a single student. In this section, we overview these two similar transportation problem classes, then discuss previous research efforts that investigate topics related to special needs student bus routing.

2.2.1. The Vehicle Routing Problem with Pickups and Deliveries

While the classical Vehicle Routing Problem only considers either pickups or deliveries, the VRPPD assumes both pickups and deliveries are able to be performed on the same vehicle tour. Nagy and Salhi (2005) develop a heuristic transportation model which addresses both pickups and deliveries (i.e., the VRPPD). The main objective of their model is minimizing the total distance travelled. The proposed four-step method allows for weak feasibility/infeasibility of starting solutions. In each step, the infeasibility of the solution is decreased until a strong, feasible solution is produced in the last step which is optimal or near optimal.

The VRPPD can be extended to include time constraints. In a student transportation application, the Vehicle Routing Problem with Pickups and Deliveries and Time Windows (VRPPDTW) examines the case when students from different schools with different starting times are on the same school bus. Van Der Bruggen *et al.* (1993) develop a variable-depth algorithm for the VRPPDTW problem with one depot. Their model consists of two phases, finding an initial feasible solution (i.e., construction) and improving the solution (i.e., improvement). Phase 1's objective is to minimize solution infeasibility, while Phase 2 focuses on minimizing the sum of route durations.

Ioachim *et al.* (1995) develop a clustering approach for the VRPPDTW problem. Their approach divides all requests into mini-clusters. Their algorithm solves the problem for these mini-clusters using a column generation-based approach to improve upon an initial, existing solution. The authors also present a heuristic for minimizing the size of the mini-cluster network.

2.2.2. The Dial-A-Ride Problem

Cordeau (2006) defines the DARP as requests for transportation which are submitted by users. This is a typical problem which applies to the transportation of the elderly or disabled people in urban areas. Requests are for transportation from a specific origin to a specific destination, and transportation is performed by vehicles based at a common depot. Also, time windows are specified which bound the arrival and/or departure times deemed acceptable by the users. Since service is shared (several users may be in the vehicle at the same time), typical objectives are both to minimize user inconveniences (i.e., delays) and to minimize operation costs.

Cordeau and Laporte (2003) develop a Tabu search metaheuristic for the DARP. Their algorithm begins with an initial, feasible solution, and then moves to the best solution within the current solution's neighborhood. Neighborhood evaluation is based on minimum route duration and minimum ride times. Attanasio *et al.* (2004) propose a more comprehensive version of Tabu search for DARP that accommodates dynamic model data. It follows that the authors suggest their problem can be solved using parallel computing techniques for real-time vehicle routing problems.

Cordeau (2006) introduces a branch-and-cut algorithm and presents valid inequalities for the DARP. He shows that the branch-and-cut algorithm is faster and more efficient in terms of computer resource needs than existing solution procedures. Although this method cannot be used for large-scale problems, it can assess single routes or a small subset of routes in small- and medium-sized problem instances.

2.2.3. Special Needs Student Bus Routing

Russell and Morrel (1986) present one of the only papers to address special needs student bus routing. They mention that since students are diverse and each bus route has several destinations, the number of schools visited by a given bus tends to be large. They develop a shuttle system to reduce the number of schools visited by each bus and therefore, student bus riding time. They select two shuttle stations (depots) which are the two district schools with the most special needs students. Buses pick up students throughout the district and go to one of the two shuttle stations. After students destined for the shuttle station school are dropped off, the remaining students are again assigned to buses for subsequent outbound travel. The authors set the maximum number of schools visited by each bus to two and for their case, are successful in reducing the number of school visits per bus and providing short bus routes.

Ripplinger (2005) focuses on rural school vehicle routing. Although his main focus is on general students, he analyzes special needs student transportation briefly. He provides models and analysis for two alternatives: separating special needs student transportation from general students and generating single routes for both types of students. Braca *et al.* (1997) briefly mention special needs students in one part of their research. Their main focus is on the New York school bus system. The authors describe the difference between special needs students and general students, but did not develop any pertinent or applicable transportation models for the research problem under study.

3 Research Overview

3.1. Problem Statement

Our review of the published literature to date reveals very little previous research on special needs student transportation and no previous work focused on the research problem under study in this thesis: minimizing the total time (distance) special needs students travel from/to their residence to/from school through effective modeling and analysis of student-to-school assignment and bus (vehicle) routing. As our research problem contains many important decisions to be made, we employ a phased research approach as described above that contains two important subproblems:

1. The student-to-school assignment problem (Phase 1)
2. The student transportation/bus routing problem (Phase 2)

3.2. Research Plan

In the student-to-school assignment problem, students are assigned to district schools having some known classroom services and capacities such that total student-to-school distance is minimized. For this purpose, we will use existing service/classroom assignments in a local school district. We use distance as a surrogate measure for student bus riding time, because in Phase 1, the direct distance between each student's home and his/her school will be used in the model without any consideration of bus routing. Even though vehicle routing is not included in Phase 1, the result of this phase can estimate how much improvement may be possible under “smarter” assignment decisions.

The focus of our Phase 2 research will be on transportation and vehicle routing. As was the case with Phase 1, the current assignment of special needs services-to-school is used as input to the Phase 2 model. However, in Phase 2, we seek to simultaneously optimize student-to-school assignments (i.e., Phase 1) and bus routing plans such that total student travel time (i.e., on-the-road distance) is minimized. Finally, given the complexity of the problem under study, it will be necessary to develop and test heuristic solution approaches for the research problem under study.

All mathematical models developed in this thesis research will be coded in AMPL and solved using CPLEX’s mixed-integer programming solver. We will first validate each mathematical model with a variety of small, trivial sample problems that are easily solved by hand. Once model functionality is verified, we will use real world information furnished by our project sponsor, the Fort Smith Public School (FSPS) district, as a means of analyzing each model’s computational performance and solution quality under real world school district conditions. In addition, the heuristic solution methods developed in this thesis will all be verified and validated in similar manner.

3.3. Research Contribution

In this thesis research, we conduct what we believe to be the first systematic, analytical study of special needs student busing and produce models and algorithms to aid decision makers with this challenging, practically motivated problem. We develop the first monolithic solution approach for helping public school systems effectively 1) assign special needs students and their associated services to schools and 2) route transportation resources.

In addition to examining current special needs service-to-school assignment, our sensitivity analyses examine the impact of a school district's flexibility (or lack thereof) in making special needs service-to-school assignments, both in terms of the number of classes of each service type at each school and the specific district schools which can accommodate special needs students. We believe that both the sensitivity analysis and Phase 2 results have the potential to be used directly in practice, as all practical constraints will be incorporated into our models, including bus and classroom capacities and classroom service classifications.

4 Model and Solution Procedure Development

In this section, we present the mathematical models and heuristic solution approaches that were developed to analyze the challenging research problem under study. After initial model development, each solution approach was verified via manual calculations to ensure its accuracy and validity. After presenting all solution approaches used in this thesis research, we will present our experimental studies and results for this thesis research.

4.1. Phase 1 Assignment Model

The assignment model developed in Phase 1 is a mixed-integer model formulated to minimize the total direct distance that all students would travel in a straight line (without any regard to routing) from each of their houses to reach their school. First, we introduce the following set notation:

S Set of schools, indexed by *i*
T Set of students, indexed by *j*
C Set of class (service) types, indexed by *k*
L Set of school levels, indexed by *l*

In addition, we need to define six parameters for use in our Phase 1 model:

n_k Maximum number of students which can attend class type *k*
$d_{i,j}$ Distance from student *j* place of residence to school *i* (miles)
$g_{j,k}$ 1 if student *j* requires class/service type *k*, otherwise 0
$a_{i,k}$ Number of classes of type *k* available in school *i*
$e_{j,l}$ 1 if student *i* should go to school level of *l*, otherwise 0
$b_{i,l}$ 1 if level of school *i* is *l*, otherwise 0

The Phase 1 assignment model determines the student-to-school assignments that minimize the total direct distance between student homes and their schools. This decision is captured via the decision variable $x_{i,j}$ which equals 1 if student *j* is assigned to school *i*, otherwise, $x_{i,j}$ =0. Since it is possible that all currently available classes at a given school may not be used in any given assignment scheme recommended by the model, we define an additional integer bookkeeping variable to count the number of students assigned to a specific class (and its associated service type) at each school. Let $y_{i,k}$ denote the number of students assigned to class/service type *k* in school *i*.

Given this notation, we now present our preliminary model. We seek to minimize total direct distance (in miles) that students travel to their school. Based on the sets, parameters, and variables defined above, our objective function is as follows:

$$Minimize \sum_i \sum_j x_{i,j} d_{i,j} + 10^{-9} \sum_i \sum_k y_{i,k} \quad (1)$$

The second term in (1) insures that bookkeeping variable $y_{i,k}$ does not become unnecessarily inflated, as we scale the sum of all bookkeeping variables by a very small constant. In this way, we make sure that the second objective function term does not adversely affect the first, primary objective function term of interest.

We now define the sets of constraints that must be satisfied by any feasible solution to our assignment model. Constraint set (2) requires that each student be assigned to exactly one school:

$$\sum_i x_{i,j} = 1 \qquad j \in T \qquad (2)$$

Constraint set (3) verifies that each student is assigned to a school that offers his/her needed class/service type (recall our service types relate to teacher to student ratio, such as 1:6, 1:10, and 1:15):

$$x_{i,j} \times g_{j,k} \leq a_{i,k} \qquad i \in S, j \in T, k \in C \qquad (3)$$

Constraint set (4) guarantees that number of students assigned to each class type at any school does not exceed the class's available capacity:

$$\sum_j x_{i,j} g_{j,k} \leq a_{i,k} n_k \qquad i \in S, k \in C \qquad (4)$$

Next, constraint set (5) insures that the each student is assigned to a school of his/her appropriate level. For example, elementary school students should only be assigned to elementary schools. This constraint is important to include in the case where students of a wide range of ages destined for different special needs schools ride the same bus:

$$x_{i,j} e_{j,l} \leq b_{i,l} \qquad i \in S, j \in T, l \in L \qquad (5)$$

Finally, constraint set (6) is a valid inequality we introduce to update the bookkeeping variable $y_{i,k}$ according to the values of our primary decision variable of interest, $x_{i,j}$:

$$\sum_j x_{i,j} g_{j,k} \leq y_{i,k} \qquad i \in S, k \in C \qquad (6)$$

4.2. Phase 2 Vehicle Routing Model

The vehicle routing problem (VRP) model developed for Phase 2 of this research is now presented. It is a mixed-integer programming model that minimizes the total travel distance driven by all the buses when picking up all special needs students and delivering them to their intended school destinations. In the interest of improved tractability, we formulate this problem using a network-based approach in which students and schools are considered to be nodes and the different routes between students and schools are captured via arcs. Buses start their travel in the network from an origin node which represents a depot. Similarly, each bus's travel is deemed complete once they return to the depot after marking all of their appropriate student drop-offs.

As mentioned earlier, the Phase 1 assignment model presented in the previous section recommends the optimal assignment of students to schools, based on their service needs. This Phase 1 model output will, in turn, be used as an input parameter in the Phase 2 model. Based on the given description, we define five sets for our Phase 2 model:

- S Set of schools, indexed by i and j
- T Set of students, indexed by i and j
- D Set of depots, indexed by i and j
- N Set of nodes, which is union of S, T, and D, indexed by i and j
- B Set of buses, indexed by k

In addition, the following parameters are defined for use in our Phase 2 routing model:

- d_{ij} Distance from node i to node j
- c_k Capacity of bus k
- a_{ij} 1 if student i is assigned to school j, 0 otherwise

The Phase 2 model specifies the special needs bus routes that minimize the total distance traveled by all buses while delivering students to their destination schools. This model output prescribes the order in which 1) students should be picked up and 2) schools should be visited for student drop-off by each bus used in the transportation plan.

The primary decision variable in our Phase 2 model is x_{ijk}, which equals 1 if node i is immediately followed by node j on bus route k, otherwise $x_{ijk} = 0$. In order to formulate the model, some bookkeeping variables are required. First, we introduce bookkeeping variable y_{ik} to record which student is served by which bus: $y_{ik} = 1$ if student i is served by bus k; otherwise, $y_{ik} = 0$. Bookkeeping variable z_{ik} is similar to y_{ik}, but keeps track of which school is visited by which bus: $z_{ik} = 1$ if school i is visited by bus k; otherwise, $z_{ik} = 0$. Finally, bookkeeping variable w_{ik} shows the position of each node on each bus route. For example, if student A is the third student visited by bus Z and bus Z has not visited any schools yet, then $w_{AZ} = 4$, as the bus depot is always the first node to be visited by any bus. In addition to keeping track of the position of the nodes visited by each bus, bookkeeping variable w_{ik} also serves the purpose of eliminating any possible sub-tours traveled by each bus.

We now formally state our Phase 2 model. Given our goal of minimizing the total distance traveled by all buses, the objective function for the Phase 2 vehicle routing model is as follows:

$$Minimize \sum_{i \in N} \sum_{j \in N} \sum_{k \in B} d_{ij} x_{ijk} + \frac{1}{M} \sum_{i \in N} \sum_{k \in B} w_{ik} \qquad (7)$$

Objective function (7) has two terms. The first term models the primary objective of minimizing the total distance traveled by all buses, while the second term makes sure that bookkeeping variable w_{ik} is not unnecessarily inflated. Similar to our Phase 1 approach, we use a very small constant multiplier on our second objective function term so as to not adversely impact the value of the overall objective function.

Next, we introduce the constraint sets required in our Phase 2 model. Constraint set (8) forces each bus to visit exactly one node immediately after visiting a student node. This is necessary, as it is not possible to pick up a student at the end of a bus's routing plan travel—at a minimum, this student must be delivered to his/her school. Therefore, it follows that there should be either another student or a school visited after any student visit:

$$\sum_{j \in N} \sum_{k \in B} x_{ijk} = 1 \qquad i \in T \qquad (8)$$

Constraint set (9) makes sure that exactly one node is visited before each student visit. This constraint, in concert with constraint set (8), forces each student home to be visited once, with exactly one arc going into and exactly one arc going out of each student node:

$$\sum_{i \in N} \sum_{k \in B} x_{ijk} = 1 \qquad j \in T \qquad (9)$$

Constraint set (10) guarantees that there is at least one student visited before any school is visited. This insures that a school is not the first place to be visited in any bus route, as there would be no students onboard to be dropped off at the school. It should be noted that this constraint set allows more than one bus to visit each school:

$$\sum_{i \in N} \sum_{k \in B} x_{ijk} \geq 1 \qquad j \in S \qquad (10)$$

Constraint set (11) insures that at most one node is visited immediately after each school visit by any bus. The visited node can be either a student node, another school node, or the final depot destination node when all the students are dropped off:

$$\sum_{j \in N} x_{ijk} \leq 1 \qquad i \in S, k \in B \qquad (11)$$

Constraint set (12) makes sure that capacity of each bus is not exceeded:

$$\sum_{i \in T} \sum_{j \in N} x_{ijk} \leq c_k \qquad k \in B \qquad (12)$$

Constraint sets (13) and (14) are valid equalities that update bookkeeping variable y_{ik} by relating it to the main decision variable, x_{ijk}:

$$\sum_{j \in N} x_{ijk} = y_{ik} \qquad i \in T, k \in B \qquad (13)$$

$$\sum_{j \in N} x_{jik} = y_{ik} \qquad i \in T, k \in B \qquad (14)$$

Constraint set (15) forces all buses to start their daily trips from the origin depot node:

$$\sum_{i \in D} \sum_{j \in T} x_{ijk} = 1 \qquad k \in B \qquad (15)$$

Constraint set (16) insures all buses end their daily trips at the depot:

$$\sum_{i \in S} \sum_{j \in D} x_{ijk} = 1 \qquad k \in B \qquad (16)$$

Constraint set (17) guarantees that no bus goes directly from the depot to a school, as no student(s) would have been picked up:

$$\sum_{j \in S} \sum_{k \in B} x_{ijk} = 0 \qquad i \in D \qquad (17)$$

Constraint set (18) verifies that no bus goes directly from a student node to the depot, as no school drop-off would have occurred:

$$\sum_{i \in T} \sum_{k \in B} x_{ijk} = 0 \qquad j \in D \qquad (18)$$

Constraint sets (19) and (20) update bookkeeping variable z_{ik} by relating it to the main decision variable x_{ijk}:

$$\sum_{j \in N} x_{ijk} = z_{ik} \qquad i \in S, k \in B \qquad (19)$$

$$\sum_{j \in N} x_{jik} = z_{ik} \qquad i \in S, k \in B \qquad (20)$$

Next, constraint set (21) makes sure that each student is picked up by a bus that visits his/her assigned school. This is one of the constraint sets that make use of the results from our Phase 1 assignment model:

$$a_{ij} y_{jk} \le z_{ik} \qquad i \in S, j \in T, k \in B \qquad (21)$$

Constraint set (22) ensures that there is no return travel from a node to back itself:

$$x_{iik} = 0 \qquad i \in N, k \in B \qquad (22)$$

Constraint set (23) sets bookkeeping variable z_{ik} for the origin depot node = 1, thereby forcing the depot to be the first node visited by each bus:

$$w_{ik} = 1 \qquad i \in D, k \in B \qquad (23)$$

Constraint sets (24) and (25) update bookkeeping variable w_{ik} and together, disallow sub-tours in the Phase 2 routing model. Constraint set (25) guarantees that students are picked up by a bus before that same bus visits their destination school:

$$w_{ik} \ge w_{jk} + 1 - (1 - x_{jik}) * M \qquad i \in S \cap T, j \in N, k \in B \qquad (24)$$

$$a_{ji} w_{ik} \le w_{jk} \qquad i \in T, j \in S, k \in B \qquad (25)$$

Finally, constraint set (26) is another valid inequality that insures no student is on a bus that does not visit his/her destination school.

$$\sum_{j \in T} a_{ij} y_{jk} \ge z_{ik} \qquad i \in S, k \in B \qquad (26)$$

4.3. Heuristic Solution Methodologies

Given the well-established NP-hard complexity of vehicle routing models containing only a single vehicle (Nagy and Salhi, 2005), large, practically-motivated real world problem instances for the research problem under study will be unsolvable in any practical amount of computation time. Therefore, we now turn our focus to the development of (hopefully) practically implementable heuristic solution methods. First, we present a constructive heuristic based on a greedy approach that generates feasible solutions quickly. Next, we introduce two local search-based post-processing techniques designed to improve the constructive heuristic's initial solution.

4.3.1. Greedy Heuristic

We employ the greedy procedure *InitialSolution* below to construct an initial, feasible solution to the problem under study in our heuristic *Greedy*. This procedure requires the following group of input parameters: 1) number of students, 2) number of schools, 3) number of buses, 4) bus capacities, 5) from/to straight line distance matrix between all pairs of students and schools, and 6) the existing list of student-to-school assignments. By assuming that all buses start their respective trips from the depot, following the procedure below guarantees the creation of a feasible solution, as first, all students are assigned to buses, and then each bus is required to visit all required schools for student drop-off.

Procedure InitialSolution

1. MAIN STUDENT ASSIGNMENT LOOP:
   a. Let *S* denote the set of all current students assigned to a school. Initially, *S* is empty.
   b. Let *S'* denote all current students not yet assigned to a school. Initially, *S'* contains all students.
   c. Let N_b denote the last visited network node of bus *b*. Initially, N_b is set to the depot for all buses.
   d. If *S'* is empty, go to Step 2. Otherwise,
      i. Find the student *s* in *S'* that lives closest to any node N_b (the current location of each bus *b*) for each bus *b* that has remaining capacity to take on more students.
      ii. Assign student *s* to bus *b*. Update N_b to reflect the network node associated with student *s*'s house. Remove *s* from *S'*. Add *s* to *S.* Go to Step 1d.
2. MAIN BUS ASSIGNMENT LOOP:
   a. For each bus, determine the schools which need to be visited for dropping off each student assigned to the bus. Let D_b denote the set of destination schools to be visited by bus *b*. Initially, D_b contains all schools attended by the students on bus *b.*
   b. If D_b is empty for all buses, STOP. Otherwise,
      i. Find the school *e* in D_b that is closest to any node N_b (the current location of each bus *b*) for each bus *b*.
      ii. Assign bus *b* to travel to school *e* by updating N_b to reflect the network node associated with school *e*. Remove *e* from D_b. Go to Step 2b.

Procedure *InitialSolution* produces two main outputs: the total distance traveled by all buses and the order in which nodes are visited by each bus. Consider an example problem containing seven students (labeled "1", "2", etc.), two schools ("A" and "B"), and three buses (labeled route "1", "2", and "3"). Table 1 displays an example set of output from the procedure in which each bus begins at origin O. Next, all seven students are picked up, and then subsequently delivered to their destination schools (note that only route 1 is required to visit both schools A and B).

Table 1: Sample Output of Constructive Heuristic

Route	Position 1	Position 2	Position 3	Position 4	Position 5	Position 6
1	O	1	2	5	B	A
2	O	7	6	A		
3	O	4	3	B		

It is quite possible that procedure *InitialSolution* might produce inefficient solutions in terms of minimum total student distances to the assigned schools, given its greedy approach. It is this reality that leads us to the following improvement methods in our heuristic development.

4.3.2. Improving the Greedy Solution

Next, we seek to improve our initial, greedy solution by focusing on 1) the way students are assigned to buses from the unassigned student pool *S'* and 2) the placement of school visits in the bus route. In procedure *InitialSolution*, a student is added to each bus during every iteration of 1d in the main student assignment loop. Now, instead of simultaneously assigning students to *every* bus during the main assignment loop, we will assign students to only one single bus at a time. When the number of students on the bus reaches capacity, the bus is removed from further consideration and the next empty bus is used for student assignment.

Considering the main bus assignment loop in procedure *InitialSolution*, we also seek to improve the placement of school visits on each bus's route. In the greedy heuristic, schools are visited at the end of each bus's route, regardless of when and where the last student is picked up—this could lead to a missed opportunity for earlier student drop-off. For example, consider the case of 20 students and two schools (A and B) being served by a single bus. If only three students are destined for school A, and these same students are picked up at the beginning of the route, there might be a chance that school A can be visited at some point earlier in the route in a way that reduces the total distance traveled. In order to identify this opportunity, we perform an additional step after assigning all students to schools which identifies the earliest position that each school can be assigned in each bus's route. Then, when performing the main bus assignment loop, we can assess school placement in each bus route from this earliest point to the end of the bus's route. The two improvement steps are included in our first improvement heuristic, *IH1*.

4.3.3. A Potential Issue with IH1

Preliminary experiments uncovered a potential issue with *IH1*. Consider a problem instance of 21 students and two buses, each with capacity for 20 students. Our *IH1* would assign the first 20 students to the first bus and then, as this bus is at capacity, would put the last remaining student on the second bus. While logically there is no problem with this assignment, it is practically not attractive or reasonable. To address this potential problem, we consider different combinations of assigning students to buses by establishing and analyzing temporary bus capacities. Consider a problem instance containing n students and b buses. While we keep the upper bound on bus capacity at 20 (its true value), we set a lower bound bus capacity value of

$\lceil n/b \rceil$ and analyze the same problem for all bus capacity values from $\lceil n/b \rceil$ up to 20. For example, in the case n = 21 students and b = 2 buses, we now examine temporary bus capacities from 11 to 20 in our second improvement heuristic, *IH2*. In *IH2*, we solve each problem for all valid temporary bus capacity values and select the solution with the lowest objective function value.

Finally, we perform a local search operation in IH2 after the best heuristic solution is found. We post-process this "best" *IH2* solution via adjacent pairwise interchanges within each bus's route to see if an improved (i.e., less distance), feasible solution exists. The interchanges are made starting from the head of each bus's route, after the depot visit. We insure feasibility is maintained such that all students can still be delivered to their proper destination school. Finally, the "best" overall routing plan identified is reported once heuristic *IH2* terminates.

5 Experimental Results and Analyses

As mentioned previously, all mathematical models and heuristic solution approaches developed in this thesis were verified for proper functionality and calculations using small, trivial problem instances with solutions that could be verified manually. Now, we use the following set of experimental factors and their associated levels to analyze the performance of our competing solution approaches for the special needs busing problem under study:

- Number of buses (3 levels): 2, 3, 4
- Number of special needs students (3 levels): 20, 40, 60
- Number of special needs schools (3 levels): 2, 4, 6
- School district area (2 levels): 10 miles x 10 miles, 20 miles x 20 miles
- Bus capacity (1 level): 20 students

These values for our experimental design were verified by our research sponsor to be valid in terms of typical school district size and complexity with regards to special needs student busing.

In each problem instance, student home and school locations are randomly generated within the corresponding school district area. Given this random component of our experimental design, we generate 10 problem instances for each of the factor combinations—this results in a total of 540 problem instances. However, close inspection reveals that 60 of these instances are infeasible: the cases wherein 60 students are to be bused with only two buses of capacity 20. As we will focus only on feasible problem instances, a total of 480 feasible instances remain for analysis by our optimization models and heuristic solution methods.

As mentioned previously, our Phase 1 assignment model solves quickly and optimally for all cases, due to its structure. Therefore, we present results below pertaining to the more complex optimization model, our Phase 2 vehicle routing model. This is appropriate in that the Phase 1 model's outputs are used as input in the Phase 2 model and it is the Phase 2 model that lends itself to direct comparison with our heuristic solution methodologies.

5.1. Phase 2 Vehicle Routing Model Results

We implemented the Phase 2 model in AMPL and analyzed it using CPLEX on a 2.93 GHz quad core, quad processor server with 128 GB of RAM. We set a maximum model run time limit of one hour and analyzed each of the 480 test instances. In terms of required solution time, while some instances solved to optimality in less than one minute, CPLEX could not find any solution to some other instances in one hour. Table 2 shows a summary of the overall CPLEX results.

Table 2: Overall Status of CPLEX Results

CPLEX Solution Type	Optimal	Time Limit	No Solution
Number	121	241	118
Percentage	25.21%	50.21%	24.58%

Results from Table 2 confirm the need for a reliable, fast heuristic. Almost 75% of the problem instances were not solved to optimality within the 60 minute time limit. In addition, CPLEX could not produce any solution for almost 25% of the instances. However, for the cases in which CPLEX could find a solution, the average gap between CPLEX's best solution and the problem's lower bound (i.e., the optimality gap) is 23.6%. The summary results in Table 2 are further broken down by experimental factor level in Table 3.

Table 3: Analysis of the Solutions of the Test Problem using CPLEX

	Instance	CPLEX		
		Optimal	Time Limit	No Solution
Number of Buses	2	46	74	0
	3	41	90	49
	4	34	77	69
Number of Students	20	121	59	0
	40	0	154	26
	60	0	28	92
Number of Schools	2	57	90	13
	4	43	72	45
	6	21	79	60
District Area	10x10	67	116	57
	20x20	54	125	61

Table 3 confirms that increasing either the number of buses, students, and/or schools makes the problem under study more difficult to solve. It appears that the number of students has the biggest effect on CPLEX's ability to achieve optimal solutions. While 67% of the solutions are optimal in the 20 student case, CPLEX found no optimal solutions for the 40 and 60 student cases. In fact, 77% of the 60 student cases resulted in no solution after the one hour time limit had elapsed. However, school district area has little to no effect on solution optimality. Again, these results confirm the need for our heuristic solution methodology, given the complexity of the problem under study.

5.2. Heuristic Solution Results

All three heuristics (*Greedy*, *IH1*, and *IH2*) were coded in C# using Microsoft Visual Studio. Each heuristic easily solved every one of the 480 test instances in less than five seconds, which compares favorably to the optimization model's 60 minute maximum solution time (not to mention that the optimization model some times did not produce *any* feasible solution within this one hour time limit). We must assess the quality of our heuristic solutions as compared to the optimization model in order to determine whether they implementation in practice is justifiable.

Let *PR(H,I)* be the performance ratio computed by dividing the problem instance solution produced by heuristic *H* for problem instance *I* by the solution produced by the Phase 2 optimization model for the same problem instance. Table 4 displays both the average and standard deviation of the *PR* ratios for each heuristic across the experimental design space. The results are separated according to whether or not the optimization model was able to produce the optimal solution or if the one hour time limit was reached.

Table 4: Comparison of Performance Ratios for Heuristic Methods

	IH2				***IH1***				***Greedy***			
	Optimal		**Time Limit**		**Optimal**		**Time Limit**		**Optimal**		**Time Limit**	
	AVG	STD	AVG	STD	AVG	STD	AVG	STD	AVG	STD	AVG	STD
2 buses	1.20	0.09	1.01	0.14	1.27	0.10	1.07	0.16	1.68	0.26	1.37	0.19
3 buses	1.43	0.13	1.14	0.23	1.58	0.15	1.22	0.27	1.91	0.33	1.51	0.33
4 buses	1.53	0.16	1.29	0.27	1.67	0.23	1.43	0.33	2.08	0.40	1.78	0.42
20 stds	1.37	0.19	1.41	0.19	1.49	0.24	1.59	0.25	1.87	0.37	1.95	0.31
40 stds	-	-	1.11	0.18	-	-	1.18	0.18	-	-	1.49	0.25
60 stds	-	-	0.80	0.09	-	-	0.82	0.10	-	-	1.07	0.13
2 schools	1.35	0.16	1.11	0.24	1.42	0.19	1.15	0.25	1.71	0.32	1.44	0.33
4 schools	1.40	0.20	1.16	0.22	1.55	0.27	1.25	0.27	1.99	0.37	1.57	0.30
6 schools	1.38	0.22	1.19	0.26	1.54	0.25	1.34	0.34	2.05	0.29	1.67	0.42
100 sq mi	1.37	0.18	1.14	0.25	1.48	0.23	1.22	0.29	1.82	0.32	1.53	0.37
400 sq mi	1.38	0.19	1.16	0.25	1.50	0.25	1.25	0.30	1.93	0.41	1.57	0.37
Overall	**1.37**	**0.19**	**1.15**	**0.25**	**1.49**	**0.24**	**1.24**	**0.30**	**1.87**	**0.37**	**1.55**	**0.37**

The overall performance ratio of the original *Greedy* constructive heuristic is 1.87 for the cases in which the Phase 2 vehicle routing model gave optimal solution. This ratio improves to 1.49 for *IH1* and 1.37 for *IH2*. This trend confirms that the proposed improvements to the original constructive heuristic help to produce better solutions. In the cases where CPLEX found a solution but not the optimal solution, all three heuristics again show superior performance as expected. Table 5 presents the 95% confidence intervals for each of the sets of heuristic results described in Table 4.

Table 5: Performance Ratio 95% Confidence Intervals

	IH2				*IH1*				*Greedy*			
	Optimal		**Time Limit**		**Optimal**		**Time Limit**		**Optimal**		**Time Limit**	
	5%	95%	5%	95%	5%	95%	5%	95%	5%	95%	5%	95%
2 buses	1.05	1.35	0.78	1.24	1.11	1.43	0.81	1.33	1.25	2.11	1.06	1.68
3 buses	1.22	1.64	0.76	1.52	1.33	1.83	0.78	1.66	1.37	2.45	0.97	2.05
4 buses	1.27	1.79	0.85	1.73	1.29	2.05	0.89	1.97	1.42	2.74	1.09	2.47
20 stds	1.06	1.68	1.10	1.72	1.10	1.88	1.18	2.00	1.26	2.48	1.44	2.46
40 stds	-	-	0.81	1.41	-	-	0.88	1.48	-	-	1.08	1.90
60 stds	-	-	0.65	0.95	-	-	0.66	0.98	-	-	0.86	1.28
2 schools	1.09	1.61	0.72	1.50	1.11	1.73	0.74	1.56	1.18	2.24	0.90	1.98
4 schools	1.07	1.73	0.80	1.52	1.11	1.99	0.81	1.69	1.38	2.60	1.08	2.06
6 schools	1.02	1.74	0.76	1.62	1.13	1.95	0.78	1.90	1.57	2.53	0.98	2.36
100 sq mi	1.07	1.67	0.73	1.55	1.10	1.86	0.74	1.70	1.29	2.35	0.92	2.14
400 sq mi	1.07	1.69	0.75	1.57	1.09	1.91	0.76	1.74	1.26	2.60	0.96	2.18
Overall	**1.06**	**1.68**	**0.74**	**1.56**	**1.10**	**1.88**	**0.75**	**1.73**	**1.26**	**2.48**	**0.94**	**2.16**

It is interesting to observe how increasing the number of students affects heuristic solution performance. While this increase negatively impacts our Phase 2 vehicle routing model's performance, it *improves* the performance of each heuristic. However, increasing the number of buses or the number of schools slightly decreases heuristic performance. Finally, as was the case with the optimization model, school district area has no noticeable effect on our performance ratios. Overall, our results tables confirm that *IH2* produces the best overall performance. Based on these findings, we now turn our final research efforts to investigating a real world case study of a local school district to assess the ability of our heuristics to perform well in practice.

6 Fort Smith Public Schools (FSPS) Case Study

Fort Smith is the second largest city in Arkansas and has a population of approximately 100,000 people. Fort Smith is approximately 53 square miles in area and is located on the border of Arkansas and Oklahoma. Currently, nine Fort Smith schools serve 111 special needs students

(Figure 1). There are three types of classes/service levels with different capacities offered for special needs students in Fort Smith: 1:6, 1:10, and 1:15 teacher to student ratios. Further, there exist three levels of schools which offer services to special needs students in Fort Smith: elementary, junior high, and senior high school.

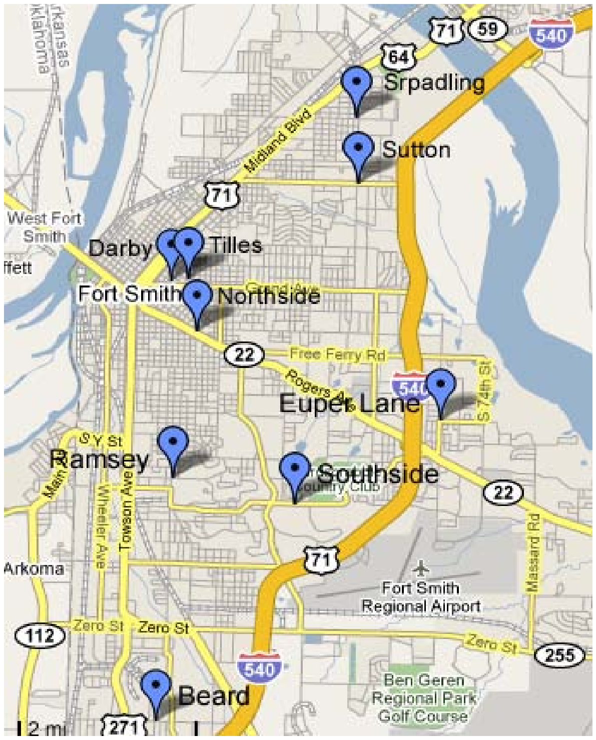

Figure 1: Fort Smith Schools Currently Accommodating Special Needs Students

After gathering all pertinent data from FSPS personnel for our models, we converted the data into an appropriate format for each of our solution methodologies. One of the big challenges during our data collection process was the determination of all possible travel distances between students and schools. This was accomplished by writing programming code to interface with Google Maps' API, thereby alleviating the need to capture all of this information manually. This distance information was then properly formatted for inclusion into the experimental data files.

6.1. Phase 1 Assignment Model Results

Our Phase 1 assignment model was used to ascertain the total distance between special needs student homes and their currently assigned schools in Fort Smith in order to make future comparisons to some known, existing baseline. This calculation resulted in a total of 467.2 miles of direct distance for the current FSPS solution of today. Again, this distance does not account for bus routing. To further describe the current conditions, Table 6 shows the number of students currently assigned to each class type in each FSPS school today (i.e., our baseline case). School names have been changed to numbers for ease of reference.

Table 6: Current Special Needs Student Assignments in FSPS by Class Type

School	1:6	1:10	1:15
1		3	
2		1	
3	9	6	34
4		3	
5	3	8	7
6	6	24	
7		1	
8		1	
9		5	
Total	**18**	**52**	**41**

Three different scenarios were investigated with the Phase 1 assignment model:

- Case 1: The current, baseline conditions in FSPS with respect to the number of available classes of each type in each school
- Case 2: We assume an infinite number of classes are available at each school, but the type of classes that each school can offer mirrors the current conditions in FSPS.
- Case 3: We assume both an infinite number of classes are available at each school and that all schools are allowed to offer all types of classes.

Analyzing Case 1 will help to assess the optimality of the current FSPS student assignments in terms of direct distance measurements or student-to-school proximity. Similarly, analyzing Cases 2 and 3 will reveal how much improvement may be possible if either the number of special needs classes or special needs class type restrictions are relaxed along with the number of classes available, respectively. Table 7 displays the results from Phase 1 assignment model runs for Case 1.

Table 7: Optimal Case 1 FSPS Student Assignment for Directed Student Distance

School	1:6	1:10	1:15
1		1	
2		3	
3	9	4	34
4		10	
5	3	6	7
6	6	17	
7		1	
8		2	
9		8	
Total	**18**	**52**	**41**

Using the assignment model to make student-to-school assignments, the total directed distance between student residences and their school is reduced by 13.2%, from 467.2 to 405.7 miles. (Although not directly tied to bus routing, it is important to note that this reduction would occur both for student transport to school and back home). If a similar amount of mileage savings (in terms of percentage) can be realized from our bus routing analysis, this would prove to be a significant savings for FSPS.

Upon comparing the results in Table 7 to the original Table 6 baseline case, the only changes that occurred were for 1:10 classes. Therefore, it appears that under the current service assignment and capacities, FSPS has optimally assigned both 1:6 and 1:15 classes in terms of directed student distance to their respective schools. However, student assignments for the 1:10 classes change in nine out of the possible 10 schools *with no need to increase the number of teachers or classes.*

Table 8 presents the CPLEX results for Case 2 wherein the number of available classes is assumed to be infinite, but existing school-to-class type restrictions are still enforced. Case 2 results reveal a 15.5% savings in total directed distance as compared to the baseline. Case 2 resulted in approximately 11 less directed miles as compared to Case 1 due to changes associated with schools 4 and 6, as seven students changed their school assignment from school 6 to school 4 under Case 2. Finally, Case 3 results (Table 9) reveal a total directed distance of 261.2 miles. However, this "ideal" case would be quite difficult to implement in reality, given the number of required changes that would need to be made in the current system. Under these optimal Case 3 results, almost every school is required to provide all three types of services—an unlikely reality given budget, teacher, and space constraints.

Table 8: Optimal Case 2 FSPS Student Assignment for Directed Student Distance

School	1:6	1:10	1:15
1		1	
2		3	
3	9	4	34
4		17	
5	3	6	7
6	6	10	
7		1	
8		2	
9		8	
Total	**18**	**52**	**41**

Table 9: Optimal Case 3 FSPS Student Assignment for Directed Student Distance

School	1:6	1:10	1:15
1	2	1	7
2	2	3	5
3	3	4	9
4		17	
5	1	6	2
6	6	10	
7	1	1	6
8	1	2	
9	2	8	12
Total	**18**	**52**	**41**

6.2. Phase 2 Vehicle Routing Model Results

We now seek to produce a practical solution for implementation in practice by creating a bus routing strategy to accompany our assignment decisions. The FSPS-supplied data for our Phase 1 case study is used again in this portion of our experimentation. Based on FSPS's stated bus capacity of 20 special needs students per bus on average, we assume this value for all buses. Student-to-school assignment data is obtained from the results of our Phase 1 assignment model. One important consideration is that because our Phase 2 vehicle routing model forces all available buses to be used in its solution, we choose to examine each problem instance with varying numbers of available buses. For this case study, no maximum CPLEX solution time limit is specified. Therefore, the solution process finishes either by finding the optimal solution or by exceeding the memory resources available to CPLEX. Results of the Phase 2 model for our FSPS case study are shown in Table 10 by type of school in terms of distance traveled, optimality gap, and model computation time.

Table 10: Phase 2 Optimization Model Results for FSPS Case Study

Level	# of Buses	Distance Traveled (mi)	Optimality / Gap	Solve Time (s)
Junior High School	1	36.9	Optimal	3.6
	2	33.3	Optimal	2.1
	3	33.0	Optimal	6.0
	4	34.0	Optimal	114.7
Senior High School	2	50.5	22.7%	6,036.2
	3	44.2	13.7%	10,697.0
	4	45.8	15.7%	8,149.4
	5	50.5	24.7%	13,758.3
Elementary School	3	100.6	53.6%	11,596.3
	4	127.0	67.7%	15,379.8
	5	-	-	-
	6	-	-	-

As expected, the cases with fewer buses were solved optimally in a short amount of time. But as instance size grows, more time is required to solve the problem—this results in even poorer solution quality. This is evident when one considers that all junior high instances were solved optimally. In the junior high cases, 19 students are assigned to two schools. Although increasing the number of buses increases model solution time, all results for the junior high cases are optimal.

None of the senior high cases, which each contain 33 students and two schools, were solved to optimality. Although the average optimality gap is approximately 20%, the required model solve time is much larger than that of the junior high instances. This example demonstrates how a small increase in problem size can affect solution times exponentially in NP-hard problems. Finally, the elementary school cases with 59 students and five schools were not easily analyzed by the Phase 2 optimization model.

Table 10 results suggest the optimal busing strategy for different school levels. However, practical considerations such as the available number of buses and bus drivers must be assessed in practice to see if these solutions can be implemented. Often, a small difference in total miles can be taken on in order to save requiring an additional bus. For example, while junior high results suggest three buses is best, an entire bus can be saved for the cost of only 0.3 additional miles each morning and afternoon. However, length of bus ride should also be analyzed for these recommended solutions as clear tradeoffs may exist between the available options.

6.3. Heuristic Results

As was the case previously, all three heuristic approaches can solve the FSPS case study models very quickly (e.g., in less than two seconds). As expected, Table 11 confirms that *IH2* generates the best solutions in all test cases when comparing the three heuristic approaches. The amount of improvement achievable by using *IH2* instead of the other two heuristic methods is much larger for the elementary school case that has the largest number of students. Table 12 displays the ratio of each heuristic's results to the Phase 2 optimization model for the FSPS case study problems. Again, improving performance is evident for *IH2*, especially in the cases where there is a larger number of available buses for student transport.

Table 11: Heuristics Results in Total Miles for FSPS Case Study

Level	Buses	*Greedy*	*IH1*	*IH2*
Junior High School	1	50.6	50.6	47.0
	2	65.3	42.4	40.4
	3	60.7	45.9	44.0
	4	76.7	51.0	44.3
Senior High School	2	72.8	50.7	48.2
	3	83.7	53.8	50.4
	4	86.0	68.6	64.6
	5	100.3	55.0	51.8
Elementary School	3	95.2	94.5	70.6
	4	115.8	95.4	74.9
	5	134.6	90.5	74.7
	6	154.4	94.9	80.1

Table 12: Comparison of Heuristics to Phase 2 Optimization Model for FSPS Case Study

Level	Buses	*Greedy*	*IH1*	*IH2*
Junior High School	1	1.37	1.37	1.27
	2	1.96	1.27	1.21
	3	1.84	1.39	1.33
	4	2.26	1.50	1.30
Senior High School	2	1.44	1.00	0.95
	3	1.89	1.22	1.14
	4	1.88	1.50	1.41
	5	1.99	1.09	1.03
Elementary School	3	0.95	0.94	0.70
	4	0.91	0.75	0.59
	5	-	-	-
	6	-	-	-

7 Conclusion and Future Research

In this thesis, we investigate the special needs student busing problem using a phased approach to assess both optimization- and heuristic-based solution approaches' ability to produce effective solutions to this challenging problem in a practically acceptable amount of time. The motivation for this research is to help school districts transport special needs students to their schools in a timely manner. Long bus ride times can be difficult on these often medically fragile students, so we seek to identify a systematic method for developing transportation routing plans for public school districts that can feasibly help to reduce this burden.

Experimental results demonstrated our proposed methods' abilities to develop transportation plans for both our experimental design dataset as well as for the data supplied by our research partner, the Fort Smith (Arkansas) Public School system. In the future, we hope to obtain the necessary permission/clearance to verify our case study results with current FSPS practice, as the school district's concerns for student privacy currently are stopping us from doing so. Also, as our heuristics shows promising results for problem instances with a large number of students and a few number of schools, further modifications can be made to *IH2* in the future to improve its performance over a wider range of school district scenarios. Finally, school district flexibility in terms of their offering of special needs services at different district schools should be investigated, as our Phase 1 model sensitivity cases suggest that some minor reassignments of special needs teachers and/or classrooms may result in a non-trivial decrease in transportation costs.

Although we tested our models and algorithms on a specific school district via our Fort Smith Public Schools case study, our solution approaches are viable for any type of school district, provided the necessary model data described above is readily available. As was shown, while our optimization models do experience computational limitations as school district size grows, our heuristics perform reasonably well, very quickly for any size of school district. The fact that we believe the performance of our heuristics improve with increasing school district size also gives us hope as to their viability in practice.

References

Attanasio, A., Cordeau, J., Ghiani, G., and Laporte, G. (2004). Parallel Tabu Search Heuristics for the Dynamic Multi-vehicle Dial-a-Ride Problem. *Parallel Computing*, 30(3), 377-387.

Braca, J., Bramel, J., Posner, B., Simchi-Levi, D., (1997). A computerized approach to the New York City school bus routing problem. *IIE Transactions*, 29, 693–702.

Cordeau, J. (2006). A Branch-and-Cut Algorithm For The Dial-a-Ride Problem. *Operations Research*, 54(3), 573-586.

Cordeau, J., and Laporte, G. (2003). A Tabu Search Heuristic for the Static Multi-vehicle Dial-a-Ride Problem. *Transportation Research: Part B*, 37(6), 579.

Ferland, J., and Guénette, G. (1990). Decision Support System For The School Districting Problem. *Operations Research*, 38(1), 15.

Ioachim, I., Desrosiers, J., Dumas, Y., Solomon, M., and Villeneuve, D. (1995). A Request Clustering Algorithm for Door-to-Door Handicapped Transportation. *Transportation Science*, 29(1), 63.

Lee, S., and Schniederjans, M. (1983). A Multicriteria Assignment Problem: A Goal Programming Approach. *Interfaces*, 13(4), 75-81.

McPherson, M., Arango, P., Fox, H., Lauver, C., McManus, M., Perrin, J., et al. (1998). A new definition of children with special health care needs. *Pediatrics*, 102(1), 137.

Nagy, G., and Salhi, S. (2005). Heuristic algorithms for single and multiple depot vehicle routing problems with pickups and deliveries. *European Journal of Operational Research*, 162(1), 126-141.

Pentico, D. (2007). Assignment problems: A Golden Anniversary Survey. *European Journal of Operational Research*, 176(2), 774-793.

Russell, R., and Morrel, R. (1986). Routing Special-Education School Buses. *Interfaces*, 16(5), 56-64.

Ripplinger, D., (2005). Rural school vehicle routing problem. *Transportation Research Record 1992*, 105–110.

Van Der Bruggen, L., Lenstra, J., and Schuur, P. (1993). Variable-Depth Search for the Single-Vehicle Pickup and Delivery Problem with Time Windows. *Transportation Science*, 27(3), 298.