

An Analysis of Special Needs Student Busing

Behrooz Kamali, Virginia Tech

Scott J. Mason, Clemson University

Edward A. Pohl, University of Arkansas

Abstract

Population growth can lead to public school capacity issues as well as increased school bus utilization, which, in turn, can result in longer school bus transport times for regular and special needs students. Special needs or medically fragile students are children with special health care needs who are at increased health and safety risk. 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 considering both student geographical location and service needs.

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 is 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 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 6 students and 1 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.

In terms of busing, not all buses in a school district can be used for special needs student transport since they require special facilities. In terms of service needs, the special services required are not offered in all schools in a school district—often, they are offered in less than one half of the district’s schools. These special services do not necessarily refer to lifts and physical equipment; they could refer to a trained teacher for special needs students, for example. It follows that these limited busing and services options can result in one or more special needs students being assigned to a school that is not necessarily close to his/her home, resulting in longer 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 requests from a principal, often with little or no consideration of where the special needs students reside. In one extreme case, the authors learned about a special needs student who rides her bus two hours each way to and from school every day. Since the assignment process 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.

Previous Research

A number of previous research studies investigated 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 paper.

The assignment problem for special needs students 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 teacher-to-student ratios).

Among the different assignment models, the semi-assignment problem has the greatest similarity to the problem under study, because each agent should be assigned to exactly one task and, also, there are a limited numbers of task groups, each of which requires some number of agents (Pentico 2007). These problems can be solved very quickly for large-scale problems.

Lee and Schniederjans (1983) developed a multi-criteria assignment model for assigning teachers to schools using goal programming with two objectives: cost minimization and maximization of preference goals. They solved the model under different priority ranking schemes and were able to find some solution combinations that satisfy a range of stated goals. Ferland and Guenette (1990) developed a decision support system for school districts to assign groups of students to a school. They developed a student network and used heuristic procedures to assign the network's edges (i.e., students) to schools such that the total distance cost is minimized.

There exists some previous transportation literature 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 rout-

ing 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.

While the classical Vehicle Routing Problem considers only pickups or deliveries, the VRPPD assumes both pickups and deliveries can be performed on the same vehicle tour. Nagy and Salhi (2005) developed a heuristic transportation model for VRPPD. The main objective of their model was minimizing the total distance traveled using a four-step method that allows for weak feasibility/infeasibility of starting solutions.

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 in which students from different schools with different starting times are on the same school bus. There are some heuristic approaches developed to tackle VRPPDTW such as variable-depth algorithm, which has two steps of finding an initial solution and improving the solution (Bruggen et al. 1993). Ioachim et al. (1995) developed a clustering approach for the VRPPDTW problem. Their approach divides all requests into mini-clusters and then 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.

DARP is defined as requests for transportation that are submitted by users. This is a typical problem that applies to the transportation of older adults or persons with disabilities in urban areas. Requests are for transportation from a specific origin to a specific destination, and vehicles based at a common depot perform transportation. Since service is shared, typical objectives are to minimize user inconveniences and to minimize operation costs (Cordeau 2006).

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. Attanasio et al. (2004) proposed a more comprehensive version of a Tabu search for DARP that accommodates dynamic model data. The authors suggest that their problem can be solved using parallel computing techniques for real-time vehicle routing problems. Cordeau (2006) introduced a branch-and-cut algorithm and presents valid inequalities for the DARP. Although his algorithm is fast and efficient in comparison to other techniques, it cannot be used for large-scale problems.

Russell and Morrel (1986) presented one of the only papers to address special needs student bus routing. They developed a shuttle system to reduce bus rides and number of school visits. They identified two schools with the most students and visited them after picking up all the students; then they rerouted and dropped other students. Ripplinger (2005) focused on rural school vehicle routing and provides models and analysis for separating special needs student transportation from general students and generating single routes for both types of students. Braca et al. (1997) briefly mentions special needs students in one part of their research. 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.

Our review of the published literature to date reveals very little previous research on special needs student transportation. As our research problem contains many important decisions to be made, we employed a phased research approach as described above that contains two important subproblems: the student-to-school assignment problem and the student transportation/bus routing problem. We investigated these problems by developing assignment models and vehicle routing models to minimize the total amount of time students travel. Given the complexity of the problem under study, we also present heuristic approaches for analyzing this challenging problem.

Student-to-School Assignment

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 use existing service/classroom assignments in a local school district. We use distance as a surrogate measure for student bus riding time, because in this phase, the direct distance between each student's home and his/her school is used in the model without any consideration of bus routing. Even though vehicle routing is not included in this model, the result of this phase can estimate how much improvement may be possible under "smarter" assignment decisions.

The assignment model developed 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 define six parameters for use in our model:

- n_k Maximum number of students that 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 j should go to school level of l , otherwise 0
- $b_{i,l}$ 1 if level of school i is l , otherwise 0

The 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 model. We seek to minimize total direct distance (in miles) that students travel to their school.

$$\text{Minimize } \sum_i \sum_j x_{i,j} d_{i,j} + \frac{1}{M} \sum_i \sum_k y_{i,k} \tag{1}$$

$$\sum_i x_{i,j} = 1 \quad j \in T \tag{2}$$

$$x_{i,j} g_{j,k} \leq a_{i,k} \quad i \in S, j \in T, k \in C \tag{3}$$

$$\sum_j x_{i,j} g_{j,k} \leq a_{i,k} n_k \quad i \in S, k \in C \tag{4}$$

$$x_{i,j} e_{j,l} \leq b_{i,l} \quad i \in S, j \in T, l \in L \tag{5}$$

$$\sum_j x_{i,j} g_{j,k} \leq y_{i,k} \quad i \in S, k \in C \tag{6}$$

The objective function (1) minimizes total direct distances and ensures that bookkeeping variable $y_{i,k}$ does not become unnecessarily inflated. The very small coefficient of the second term makes sure the primary objective function term of interest is not adversely affected. Constraint set (2) requires that each student be assigned to exactly one school. Constraint set (3) verifies that each student is assigned to a school that offers his/her needed class/service type. Constraint set (4) guarantees that the number of students assigned to each class type at any school does not exceed the class's available capacity. Constraint set (5) ensures that each student is assigned to a school of his/her appropriate level (e.g., elementary school students should be assigned only to elementary schools). Finally, constraint set (6) is a valid inequality we introduced to update the bookkeeping variable $y_{i,k}$ according to the values of our primary decision variable of interest, $x_{i,j}$.

Bus Routing Problem

The vehicle routing problem (VRP) model developed 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 this model, 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 it returns to the depot after marking all of its appropriate student drop-offs.

As mentioned earlier, the assignment model presented in the previous section recommends the optimal assignment of students to schools based on their service needs, and its output will, in turn, be used as an input parameter in the routing model. Using assignment model output as an input of the routing model makes these two problems dependent. For example, in low-density networks, the quality of the routes becomes more dependent on the network shape. Based on the given description, we define five sets for our routing 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 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 primary decision variable in this model is x_{ijk} , which equals 1 if node i is immediately followed by node j on bus route k , otherwise $x_{ijk} = 0$. 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 routing model:

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

$$\sum_{j \in N} \sum_{k \in B} x_{ijk} = 1 \quad i \in T \tag{8}$$

$$\sum_{i \in N} \sum_{k \in B} x_{ijk} = 1 \quad j \in T \tag{9}$$

$$\sum_{i \in N} \sum_{k \in B} x_{ijk} \geq 1 \quad j \in S \tag{10}$$

$$\sum_{j \in N} x_{ijk} \leq 1 \quad i \in S, k \in B \tag{11}$$

$$\sum_{i \in T} \sum_{j \in N} x_{ijk} \leq c_k \quad k \in B \tag{12}$$

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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. 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.

Constraint set (8) forces each bus to visit exactly one node immediately after visiting a student node. This is necessary to make sure all picked up students are delivered. Constraint set (9) makes sure that exactly one node is visited before each student visit. Constraint set (10) guarantees that there is at least one student visited before any school is visited. Constraint set (11) ensures that, at most, one node is visited immediately after each school visit by any bus. The visited node can be a student node, another school node, or the final depot destination node when all the students are dropped off. Constraint set (12) makes sure that the capacity of each bus is not exceeded. Constraints sets (13) and (14) are valid equalities that update bookkeeping variable y_{ik} by relating it to the main decision variable, x_{ijk} . Constraint set (15) forces all buses to start their daily trips from the origin depot node. Constraint set (16) ensures that all buses end their daily trips at the depot. Constraint set (17) guarantees that no bus goes directly from the depot to a school. Constraint set (18) verifies that no bus goes directly from a student node to the depot. Constraint sets (19) and (20) update bookkeeping variable z_{ik} by relating it to the main decision variable x_{ijk} .

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 makes use of the results from our assignment model. Constraint set (22) ensures that there is no return travel from a node back to itself. Constraint set (23) sets bookkeeping variable w_{ik} for the origin depot node = 1, thereby forcing the depot to be the first node visited by each bus. Constraint sets (24) and (25) update bookkeeping variable w_{ik} and together disallow sub-tours in the routing model. Constraint set (25) guarantees that a bus picks up students before that same bus visits their destination school. Finally, constraint set (26) is another valid inequality that ensures no student is on a bus that does not visit his/her destination school.

Given the complexity of the problem under study, it was necessary to develop and test heuristic solution approaches for the research problem under study.

Heuristics

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 are unsolvable in any practical amount of computation time. Therefore, we 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 gener-

ates feasible solutions quickly. Next, we introduce two local search-based post-processing techniques designed to improve the constructive heuristic's initial solution.

Greedy Heuristic

We employed 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 a 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) 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. 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 led us to the following improvement methods in our heuristic development.

Improving the Greedy Solution

Next, we sought 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 (d) in the main student assignment loop. Now, instead of simultaneously assigning students to every bus during the main assignment loop, we 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 sought 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. To identify this opportunity, we performed an additional step after assigning all students to schools that identifies the earliest position that each school can be assigned in each bus's route. Then, when performing the main bus assignment loop, we could assess school placement in each bus route from this earliest point to the end of the bus's route. The two improvement steps were included in our first improvement heuristic, *IH1*.

A Potential Issue with *IH1*

Preliminary experiments uncovered a potential issue with *IH1*. Consider a problem instance of 21 students and 2 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,

it would put the last remaining student on the second bus. While, logically, there is no problem with this assignment, it is not practically attractive or reasonable. To address this potential problem, we considered 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 solved each problem for all valid temporary bus capacity values and selected the solution with the lowest objective function value.

Finally, we performed a local search operation in *IH1* after the best heuristic solution was found. We post-processed this “best” *IH1* 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 ensured that feasibility was maintained such that all students can still be delivered to their proper destination school. Finally, the “best” overall routing plan identified was reported once heuristic *IH2* terminated.

Experimental Results and Analysis

All mathematical models developed in this research were coded in AMPL and solved using CPLEX’s mixed-integer programming solver. CPLEX was run on a 2.93 GHz quad core, quad processor server with 128 GB of RAM. We first validated each mathematical model with a variety of small, trivial sample problems that are easily solved by hand. Once model functionality was verified, we used 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 paper were all verified and validated in similar manner. To validate and analyze performance of the model, the following set of experimental factors were used:

- 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

Our research sponsor verified these values 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 were randomly generated within the corresponding school district area. Given this random component of our experimental design, we generated 10 problem instances for each of the factor combinations, resulting in a total of 540 problem instances. However, close inspection revealed that 60 of these instances were infeasible: the cases wherein 60 students were to be bused with only 2 buses of capacity 20. As we focused only on feasible problem instances, a total of 480 feasible instances remained for analysis by our optimization models and heuristic solution methods.

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

Vehicle Routing Model Results

We set a maximum model run time limit of 1 hour and analyzed each of the 480 test instances. In terms of required solution time, while some instances solved to optimality in less than 1 minute, CPLEX could not find any solution to some other instances in 1 hour. Table 1 shows a summary of the overall CPLEX results.

Table 1. 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 1 confirm the need for a reliable, fast heuristic. Almost 75 percent 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 percent 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) was 23.6 percent. The summary results in Table 1 are further broken down by experimental factor level in Table 2.

Table 2. 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 2 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 percent 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 percent of the 60-student cases resulted in no solution after the 1-hour time limit had elapsed. However, school district area has little to no effect on solution optimality. Again, these results confirmed the need for our heuristic solution methodology, given the complexity of the problem under study.

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 5 seconds, which compares favorably to the optimization model’s 60-minute maximum solution time. We assessed the quality of our heuristic solutions as compared to the optimization model in order to determine whether their implementation in practice was 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 routing model for the same problem instance. Table 3 displays both the average and standard deviation of the PR ratios for each heuristic across the experi-

mental 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 3. Comparison of Performance Ratios for Heuristic Methods

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

The overall performance ratio of the original *Greedy* constructive heuristic was 1.87 for the cases in which the vehicle routing model gave optimal solution. This ratio improved to 1.49 for *IH1* and 1.37 for *IH2*. This trend confirmed that the proposed improvements to the original constructive heuristic helped to produce better solutions. In the cases where CPLEX found a solution but not the optimal solution, all three heuristics again showed superior performance as expected. Table 4 presents the 95% confidence intervals for each of the sets of heuristic results described in Table 3.

Table 4. Performance Ratio 95% Confidence Intervals

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

It is interesting to observe how increasing the number of students affects heuristic solution performance. While this increase negatively impacted our vehicle routing model’s performance, it *improves* the performance of each heuristic. However, increasing the number of buses or the number of schools slightly decreased heuristic performance. Finally, as was the case with the optimization model, school district area had no noticeable effect on our performance ratios. Overall, our results tables confirm that *IH2* produced 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.

Case Study

Fort Smith is the second largest city in Arkansas and has a population of approximately 100,000 people. It is approximately 53 square miles in area and is located on the border of Arkansas and Oklahoma. Currently, 9 Fort Smith schools serve 111 special needs students. There are 3 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 that offer ser-

VICES to special needs students in Fort Smith: elementary, junior high, and senior high. 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.

Assignment Model Results

Our 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 5 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 5. 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

In solving the assignment model, the current baseline conditions in FSPS were kept the same with respect to the number of available classes of each type in each school. Table 6 displays the results from assignment model.

Using the assignment model to make student-to-school assignments, the total directed distance between student residences and their school was reduced by 13.2 percent, from 467.2 to 405.7 miles. 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.

Table 6. Optimal 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

Upon comparing the results in Table 6 to the original Table 5 baseline case, the only changes that occurred were for the 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 8 out of the possible 9 schools with no need to increase the number of teachers or classes.

Vehicle Routing Model Results

We now sought to produce a practical solution for implementation in practice by creating a bus routing strategy to accompany our assignment decisions. Based on FSPS’s stated bus capacity of 20 special needs students per bus on average, we assumed this value for all buses. Student-to-school assignment data was obtained from the results of our assignment model. One important consideration was that because our vehicle routing model forces all available buses to be used in its solution, we chose to examine each problem instance with varying numbers of available buses. For this case study, no maximum CPLEX solution time limit was specified. Therefore, the solution process finished either by finding the optimal solution or by exceeding the memory resources available to CPLEX. Results of the vehicle routing model for our FSPS case study are shown in Table 7 by type of school in terms of distance traveled, optimality gap, and model computation time.

Table 7. Vehicle Routing 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 was 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 were assigned to two schools. Although increasing the number of buses increases model solution time, all results for the junior high cases were optimal.

None of the senior high cases, which each contain 33 students and 2 schools, were solved to optimality. Although the average optimality gap was approximately 20 percent, the required model solve time was 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 5 schools were not easily analyzed by the vehicle routing model.

Table 7 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 3 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.

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 8 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 9 displays the ratio of each heuristic’s results to the vehicle routing 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 8. 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 9. Comparison of Heuristics to Vehicle Routing 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	-	-	-

Conclusion and Future Research

In this research, we conducted what we believe to be the first systematic, analytical study of special needs student busing and produced models and algorithms to aid decision makers with this challenging, practically-motivated problem. We developed the first monolithic solution approach for helping public school systems to effectively 1) assign special needs students and their associated services to schools and 2) route transportation resources. We investigated 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 approach developed in this study can also be applied to other problems containing both assignment and vehicle routing phases.

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 precluding 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. As a future step, we can also include administrative costs such as labor costs and bus financing costs in our analysis as shown in Ibeas et al. (2006, 2009). Finally, school district flexibility in terms of the 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.

References

- Attanasio, A., J. Cordeau, G. Ghiani, and G. Laporte. 2004. Parallel Tabu search heuristics for the dynamic multi-vehicle dial-a-ride problem. *Parallel Computing* 30(3): 377-387.
- Braca, J., J. Bramel, B. Posner, and D. Simchi-Levi. 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 G. Laporte. 2003. A Tabu search heuristic for the static multi-vehicle dial-a-ride problem. *Transportation Research Part B*, 37(6): 579-594.
- Ferland, J., and G. Gu enette. 1990. Decision support system for the school districting problem. *Operations Research* 38(1): 15-21.
- Ibeas, A., J. L. Moura, and L. Dell'Olio. 2009. Planning school transport: Design of routes with flexible school opening times. *Transportation Planning and Technology* 32(6): 527-544.
- Ibeas, A., J. L. Moura, L. Dell'Olio, and J. D. Ortuzar. 2006. Costing school transport in Spain. *Transportation Planning and Technology* 29(6): 483-501.
- Ioachim, I., J. Desrosiers, Y. Dumas, M. Solomon, and D. Villeneuve. 1995. A request clustering algorithm for door-to-door handicapped transportation. *Transportation Science* 29(1): 63-78.
- Lee, S., and M. Schniederjans. 1983. A multicriteria assignment problem: A goal programming approach. *Interfaces* 13(4): 75-81.

- McPherson, M., P. Arango, H. Fox, C. Lauver, M. McManus, J. Perrin, and others. 1998. A new definition of children with special health care needs. *Pediatrics* 102(1): 137-139.
- Nagy, G., and S. Salhi. 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 R. Morrel. 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., J. Lenstra, and P. Schuur. 1993. Variable-depth search for the single-vehicle pickup and delivery problem with time windows. *Transportation Science* 27(3): 298-311.

About the Authors

BEHROOZ KAMALI (*kamali@vt.edu*) is a Ph.D. student in Industrial and Systems Engineering at Virginia Polytechnic Institute and State University. He received a master's degree in Industrial Engineering from University of Arkansas and a bachelor's degree from Sharif University of Technology in Iran. His main research interests are optimization and mathematical modeling in the areas of transportation, logistics, and supply chain systems. He is a member of Institute of Industrial Engineers and INFORMS.

DR. SCOTT J. MASON (*mason@clemson.edu*) is the Fluor Endowed Chair in Supply Chain Optimization and Logistics and a Professor of Industrial Engineering at Clemson University. Prior to joining Clemson, he spent 10 years in the Department of Industrial Engineering at the University of Arkansas. He received a Ph. D. in Industrial Engineering from Arizona State University after earning bachelor's and master's degrees from The University of Texas at Austin. His areas of focus include operations planning, scheduling, and control of capital project supply chains and large-scale systems modeling, optimization, and algorithms, with domain expertise in semiconductor manufacturing. He is an Associate Editor for *IEEE Transactions on*

Electronics Packaging Manufacturing, a senior member of the Institute for Industrial Engineers, and a member of INFORMS.

EDWARD A. POHL, PH.D. (*epohl@uark.edu*) is an Associate Professor in the Department of Industrial Engineering at the University of Arkansas. He serves as the Director of the Operations Management Program. He received a Ph.D. in Systems and Industrial Engineering from the University of Arizona, an M.S. in Reliability Engineering from the University of Arizona, an M.S. in Systems Engineering from AFIT, an M.S. in Engineering Management from the University of Dayton, and a B.S.E.E. from Boston University. His primary research interests are in repairable systems modeling, reliability, decision making under uncertainty, engineering optimization, and probabilistic design. He is a senior member of IIE, a senior member of ASQ, and a senior member of IEEE and serves as an Associate Editor for both the *Journal of Military Operations Research* and the *Journal of Risk and Reliability*. He is a member of RAMS management committee and is a two-time winner of the Alan Plait award for Outstanding Tutorial at RAMS.