



Transportation Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

A Bounded Formulation for The School Bus Scheduling Problem

Liwei Zeng, Sunil Chopra, Karen Smilowitz

To cite this article:

Liwei Zeng, Sunil Chopra, Karen Smilowitz (2022) A Bounded Formulation for The School Bus Scheduling Problem. Transportation Science 56(5):1148-1164. <https://doi.org/10.1287/trsc.2022.1130>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2022, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

A Bounded Formulation for The School Bus Scheduling Problem

Liwei Zeng,^{a,*} Sunil Chopra,^b Karen Smilowitz^c

^aThe Institute for Mathematics and Its Applications, University of Minnesota, Minneapolis, Minnesota 55455; ^bKellogg School of Management, Northwestern University, Evanston, Illinois 60208; ^cIndustrial Engineering and Management Sciences, Northwestern University, Evanston, Illinois 60208

*Corresponding author

Contact: zengx313@umn.edu,  <https://orcid.org/0000-0002-3197-8285> (LZ); s-chopra@kellogg.northwestern.edu (SC); ksmilowitz@northwestern.edu (KS)

Received: May 3, 2021

Revised: October 27, 2021; January 14, 2022

Accepted: January 23, 2022

Published Online in Articles in Advance:
March 9, 2022

<https://doi.org/10.1287/trsc.2022.1130>

Copyright: © 2022 INFORMS

Abstract. This paper proposes a new formulation for the school bus scheduling problem (SBSP), which optimizes school start times and bus operation times to minimize transportation cost. The goal is to minimize the number of buses to serve all bus routes such that each route arrives in a time window before school starts. We show that introducing context-specific features, common in many school districts, can lead to a new time-indexed integer linear programming (ILP) formulation. Based on a strengthened version of the linear relaxation of the ILP, we develop a dependent randomized rounding algorithm that yields near-optimal solutions for large-scale problem instances. The efficient formulation and solution approach enable quick generation of multiple solutions to facilitate strategic planning, which we demonstrate with data from two public school districts in the United States. We also generalize our methodologies to solve a robust version of the SBSP.

Funding: This work was supported by Division of Civil, Mechanical and Manufacturing Innovation, National Science Foundation [Grant CMMI-1727744].

Supplemental Material: The e-companion is available at <https://doi.org/10.1287/trsc.2022.1130>.

Keywords: school bus scheduling problem • time-indexed formulation • randomized rounding algorithm

1. Introduction

Public school districts across the United States face critical budget challenges (Bidwell 2015) and transportation is a common area to look for cost reduction. School districts can often lower transportation cost by staggering school start times, which allows the district to complete more than one route with a single bus. The marginal cost of reusing a bus for a second route is significantly lower than that of adding a new bus. In this paper, we study the problem of jointly determining school start times and bus route operation times to minimize total buses needed.

The school bus routing problem (SBRP) has been extensively studied (Newton and Thomas 1969, Desrosiers et al. 1980, Park and Kim 2010). The SBRP consists of two groups of subproblems: routing and scheduling. The routing subproblems select bus stops, assign students to stops, and connect stops into bus routes. This paper focuses on the scheduling subproblems, which determine school start times, bus route operation times, and route-to-bus assignment. We use the term “school bus scheduling problem” (SBSP) from Fügenschuh (2009) to represent the joint problem of determining school start times, bus route operation times, and route-to-bus assignment.

Problem 1 (School Bus Scheduling Problem). Given a set of schools and a set of routes for each school, the school bus scheduling problem determines the start times of schools and the arrival times of routes such that each route arrives in a time window before school starts. After school start times and route arrival times are determined, routes are assigned to buses such that routes assigned to the same bus operate on disjoint time intervals. The goal is to minimize the number of buses needed to complete all the routes.

The SBSP can be shown to be NP-hard by a reduction from the balanced partition problem (Garey and Johnson 2002). Prior work has identified exact and heuristic algorithms for the SBSP, such as column generation (Desrosiers et al. 1986), cutting planes (Fügenschuh 2009), and neighborhood search (Bertsimas et al. 2019). Although exact algorithms perform well on small instances, computational challenges arise as the problem scale increases, which can become an issue with large school districts. Heuristic algorithms scale well but typically lack theoretical guarantees.

Motivated by collaborations with two public school districts, we use context-specific features, common in many school districts in the United States, to develop a new formulation of the SBSP. These features, that appear in both collaborations, are as follows:

1. Bus routes are assumed to be given. Our collaborations focus on the impact of changing school start times and bus operation times on transportation costs. In contrast, Desrosiers et al. (1986), Spada et al. (2005), Bertsimas et al. (2019) consider changing both bus routes and school start times to potentially bring larger savings.

2. Each bus route is assumed to serve only one school. In contrast, several papers (Fügenschuh 2009; Park, Tae, and Kim 2012; Campbell, North, and Ellegood 2015; Ellegood, Campbell, and North 2015) study the “mixed load” problem where one bus route can serve students from multiple schools.

3. The transition time between routes is assumed to be constant. In contrast, several papers (Fügenschuh 2009; Kim, Kim, and Park 2012; Chen et al. 2015) allow transition times that depend on school locations. We also present a modification for location-dependent transition times.

As noted, more general versions of the SBSP exist in the literature. However, we show that incorporating context-specific features leads to a new approach to solve the SBSP that scales well and has a provable performance guarantee. Having a fast solution technique with a provable performance guarantee allows decision makers to more easily incorporate the impact on transportation cost into strategic planning decisions. Our approach provides decision makers with a set of low transportation cost solutions from which they can select one that best satisfies a variety of other requirements. As a result, our work aims to support strategic planning, in contrast to more operationally focused work designed to provide immediately actionable plans.

Several papers (Lenstra, Shmoys, Tardos 1990; Shmoys and Tardos 1993; Schulz 1996) in the scheduling literature have used LP relaxations of suitable ILP formulations to develop efficient approximation algorithms with good lower bounds. Following this approach, we formulate the SBSP as a time-indexed integer linear program (ILP) where binary variables indicate school start times and route arrival times. We strengthen the time-indexed ILP by adding inequalities that fully characterize the convex hulls of variable subspaces and show that the strengthened formulation has a bounded integrality gap. We then develop an efficient randomized rounding algorithm that is near-optimal for large-scale instances. The randomized algorithm provides multiple high-quality solutions, which can be used by decision makers to consider additional factors (e.g., Banerjee and Smilowitz (2019) which incorporates equity in start time decisions) when selecting the solution to implement.

The main contributions of this paper are as follows:

1. We develop a new time-indexed ILP formulation for the SBSP. Based on the LP relaxation of the ILP, we propose a randomized rounding algorithm with a

bounded performance guarantee that provides multiple high-quality solutions efficiently. Our approach enables school districts to quickly estimate the impact on transportation cost of different plans when considering school start time changes. Providing multiple solutions allows decision makers to incorporate external considerations while maintaining low transportation cost.

2. We are able to solve small SBSP instances to optimality using this ILP formulation. For larger instances, we prove that the randomized rounding algorithm is near-optimal. The near-optimality of our algorithm provides accurate estimates of potential cost-saving of different policies, allowing decision makers to focus on the most promising plans. With our context-specific features, this approach yields better solutions faster than existing approaches.

3. Using data from two public school districts, one moderate and one large, we demonstrate how decision makers can use this approach for strategic planning.

4. We generalize our findings for a robust version of the SBSP. Robustness is critical in school bus scheduling because school start times cannot be redesigned frequently, yet bus routes may vary from year to year. From a strategic planning perspective, districts are interested in school start time plans that remain consistent over time and still lower transportation cost.

The paper is organized as follows. In Section 2 we review related literature. In Section 3 we present a time-indexed ILP for the SBSP, which is strengthened with valid cuts. In Section 4, we develop a randomized rounding algorithm based on the strengthened LP-relaxation that is provably near-optimal for large-scale instances. Our methodology is generalized in Section 5 for a robust version of the SBSP. In Section 6, we present a case study with partnering school districts and a numerical analysis to complement our theoretical findings. We also generalize our approach for location-dependent transition times. We summarize results and discuss future work in Section 7.

2. Literature Review

We summarize related literature in school bus scheduling and link the SBSP to two other combinatorial optimization problems: the machine scheduling problem and the bin packing problem.

2.1. School Bus Scheduling Problem

As a part of the school bus routing problem, the SBSP specifies school start times and route arrival times to maximize bus usage. Special cases of the SBSP have been extensively studied in the literature. When route arrival times are fixed, the problem can be formulated as an assignment problem (Gavish, Schweitzer, and Shlifer 1978; Carraraesi and Gallo 1984). When arrival

times are not fixed, Orloff (1976) proves that the problem is NP-complete. When school start times are fixed, Swersey and Ballard (1984) present a mixed-integer programming (MIP) formulation to determine optimal route arrival times. The MIP formulation is further simplified to an integer programming (IP) formulation by discretizing the timeline into unit intervals. Kim, Kim, and Park (2012) studies the SBSP in which schools have fixed time windows within which routes must be finished. The authors show that when route operation times are fixed, the problem is equivalent to an assignment problem, and can be solved to optimality in polynomial time.

For the general SBSP, Desrosiers et al. (1986) determine school start times and route arrival times sequentially. They formulate the problem as a min-max binary program that minimizes the maximum number of routes operating during the same time period. The formulation is solved using column generation for small-sized instances. For large-scale instances, they present a heuristic that alternately updates school start times and route arrival times. Fügenschuh (2009) studies a general model for school bus scheduling where bus routes are allowed to serve multiple schools. The author uses a branch-and-cut algorithm to solve an IP formulation that determines school start times and route arrival times simultaneously. More recently, Köksal Ahmed et al. (2020) develop a reinforcement learning based genetic algorithm to provide fast solution to the SBSP. Banerjee and Smilowitz (2019) use a variant of the time-indexed formulation developed in this paper to incorporate equity (defined by the disutility associated with changing school start times).

Spada, Bierlaire, and Liebling (2005) consider bus route generation and bus scheduling, providing an initial feasible solution by solving an IP, which is then improved using heuristics. Bertsimas, Delarue, and Martin (2019) construct a set of routing scenarios for each school using a biobjective routing decomposition algorithm, and then formulate the school bell selection problem as a generalized quadratic assignment problem (the formulation is introduced in Wenzel (2016)). They implement their algorithms in collaboration with the Boston Public Schools, leading to \$5 million savings per year.

In contrast to previous work which covered more general SBSP settings, we focus on leveraging context-specific features common in many school districts to develop efficient approaches with provable performance guarantees to support strategic decision making.

2.2. Machine Scheduling and Bin Packing

Our work is related to machine scheduling problems where jobs with higher priority must arrive (or end) earlier than others. Ikura and Gimple (1986) study the batched scheduling problem where jobs in the same

batch have the same priority. A simplified version of the SBSP (where each route arrives exactly at the school's start time) can be restated as the batched machine scheduling problem where routes in the same school have the same priority. The SBSP is also similar in spirit to the one-dimensional bin packing problem (De La Vega and Lueker 1981; Scholl, Klein, and Jürgens 1997) where objects are bus routes, and bins are buses. The constraints on route arrival times can be transformed to constraints on objects' relative locations.

3. Integer Linear Programming Formulation for the SBSP

We introduce a time-indexed SBSP formulation, which is then strengthened to support a randomized rounding algorithm, presented in Section 4.

3.1. Preliminaries

We use the following notation throughout the paper. Let $\mathcal{S}$ be a set of schools and let $\mathcal{R} = \cup_{s \in \mathcal{S}} \mathcal{R}_s$ be a set of routes, where $\mathcal{R}_s$ is the set of routes for school s . We assume that each route only serves one school so that $\mathcal{R}_s \cap \mathcal{R}_{s'} = \emptyset$, $\forall s \neq s'$. We discretize the timeline into T unit intervals and assume that each school (route) chooses a start (arrival) time from $[T] = \{1, 2, \dots, T\}$. For each school s , $l_s \in \mathbb{N}$ represents the length of the time window during which its routes must arrive. Specifically, if school s starts at time t_s , all routes for this school must arrive in the interval $[t_s - l_s, t_s]$ ($[1, t_s]$ if $t_s - l_s \leq 0$). Finally, we use $r_i \in \mathbb{N}$ to denote the travel time of route i .

We assume constant transition times between routes, which allow us to leverage structural information and design fast algorithms with provable bounds. In Section 6, we generalize our approach, albeit without provable bounds, for more general settings without constant transition times.

The moderately sized school district with whom we collaborate is fairly compact such that constant transit time between routes is reasonable. The larger school district also assumes constant transition times in their analysis of school start times. They do so because their aim to limit buses to operate within compact subregions of the geographically dispersed district results in significantly smaller ranges of bus transition times compared with route travel and stopping times.

In what follows, the constant transition time is added to the beginning of each route (e.g., a 30-minute route with 10-minute transition becomes a 40-minute route) and we assume that there is no transition time between routes. School bell times are not affected because they are only related to the arrival time of routes. In practice, the transition time should not be added to the first route taken by each bus. However, adding the same time to

each route (including the first route for each bus) will not affect the optimal solution. The time added to the first route can be viewed as the travel time from the depot to the first stop, which does not affect the route connection pattern.

3.2. A Time-Indexed Formulation for the SBSP

Time-indexed formulations have been widely applied to scheduling problems to obtain strong lower bounds (Dyer and Wolsey 1990, Sousa and Wolsey 1992, Queyranne and Schulz 1994). The time horizon is divided into unit time periods, and binary variables indicate job start/end times.

Similar to Desrosiers et al. (1986), we develop a time-indexed SBSP formulation where binary variables indicate school start times and route arrival times. The novelty of our formulation is a set of constraints that captures the relation between school start times and route arrival times. We start with an observation from Desrosiers et al. (1986) that transforms the minimum number of buses to an easy-to-compute quantity.

Proposition 1 (Desrosiers et al. 1986). *Given a set of routes with fixed operation times, the minimum number of buses to complete all the routes is equal to the maximum number of routes in operation during the same time period.*

Given the arrival time of each route, Proposition 1 states that the minimum number of buses required can be computed by dividing the timeline into unit intervals and counting the number of routes in operation during each interval. From Proposition 1, the minimum number of buses required to serve all routes is equal to the maximum number recorded over all intervals. Moreover, with fixed route operation times, the route-to-bus assignment problem is equivalent to an interval graph coloring problem, which can be solved in polynomial time (see Olariu (1991) for details). Motivated by this observation, we introduce a time-indexed formulation in which the maximum number of routes in operations during the same unit interval is minimized.

For route $i \in \mathcal{R}$ and time $t \in [T]$, let $x_{i,t}$ be a binary variable such that $x_{i,t} = 1$ if route i arrives at time t . For school $s \in \mathcal{S}$ and time $t \in [T]$, we introduce a binary variable $y_{s,t}$ such that $y_{s,t} = 1$ if school s starts at time t . The variables are defined based on route arrival times and school start times for notational convenience. We extend the definition of x, y variables to $t \notin [T]$ and define that $x_{i,t} = 0, y_{i,t} = 0$ if $t \notin [T]$. The SBSP can be formulated as follows (formulation ILP1).

$$\begin{aligned} \min \quad & z & (\text{ILP1}) \\ \text{s.t.} \quad & \sum_{t=1}^T x_{i,t} = 1 & \forall i \in \mathcal{R}, s \in \mathcal{S} \quad (1a) \end{aligned}$$

$$\sum_{t=1}^T y_{s,t} = 1 \quad \forall s \in \mathcal{S} \quad (1b)$$

$$x_{i,t} \leq \sum_{t'=t}^{\min\{t+l_i, T\}} y_{s,t'} \quad \forall i \in \mathcal{R}, s \in \mathcal{S}, t \in [T] \quad (1c)$$

$$\sum_{i \in \mathcal{R}} \sum_{t'=t}^{\min\{t+r_i-1, T\}} x_{i,t'} \leq z \quad \forall t \in [T] \quad (1d)$$

$$x_{i,t} \in \{0, 1\} \quad \forall i \in \mathcal{R}, s \in \mathcal{S}, t \in [T] \quad (1e)$$

$$y_{s,t} \in \{0, 1\} \quad \forall s \in \mathcal{S}, t \in [T] \quad (1f)$$


Assignment constraints (1a) and (1b) ensure that each route (school) is assigned to one arrival (start) time. Time-window constraints (1c) enforce that each route arrives in a time window before school starts. Constraints (1d) link the decision variables to the objective by Proposition 1. For each $t \in [T]$, the left-hand side of (1d) computes the number of routes in operation during unit interval $[t-1, t]$. By introducing a constraint for each $t \in [T]$, (1d) implies that the objective z is equal to the maximum number of routes in operation during the same unit time interval.

Formulation ILP1 can incorporate restrictions on school start times and route operation times by fixing some decision variables. For example, if school s cannot (must) start at time t , we can add a constraint as $y_{s,t} = 0$ ($y_{s,t} = 1$). As shown in Section 4, incorporating these restrictions does not affect the theoretical performance guarantee of our solution approach.

Although the time-indexed formulation is known to provide better bounds, its large size makes it hard to solve directly to optimality. There are $(|\mathcal{R}| + |\mathcal{S}|)T$ binary variables in ILP1. For a district with 20 schools, 100 routes and $T = 120$ (two-hour time horizon divided into one-minute intervals), this requires solving an ILP with 14,400 variables (and even more constraints). Even with school start times restricted to multiples of five minutes, the number of binary variables still exceeds 12,000. Thus, we first strengthen the LP-relaxation of ILP1 by adding valid cuts. As shown in Section 4, although the problem size increases after adding valid cuts, we obtain high-quality solutions from the strengthened LP-relaxation of ILP1. In Section 6, we use a numerical study to show that the strengthened formulation provides tight lower bounds in practice.

3.3. A Strong Integer Linear Programming Formulation

Observe that ILP1 has a semidecomposable structure in that all its constraints, except for (1d), only contain variables for a single school. By suitably reordering the variables and constraints, ILP1 exhibits the block-structure

Figure 1. Block-Structure of ILP1

shown in Figure 1. For each school $s \in \mathcal{S}$ we define

$$V_s = \{(x_{i,t}, y_{s,t}) \mid i \in \mathcal{R}_s, (x_{i,t}, y_{s,t}) \text{ satisfies } (1a), (1b), (1c), (1e), \text{ and } (1f)\}.$$

V_s can be interpreted as the set of variables satisfying the block of constraints for school s in Figure 1. Hence, ILP1 can be rewritten as

$$\begin{aligned} \min \quad & z \\ \text{s.t.} \quad & \sum_{i \in \mathcal{R}} \sum_{t'=t}^{\min\{t+r_i-1, T\}} x_{i,t'} \leq z \quad \forall t \in [T] \\ & (x_{i,t}, y_{s,t}) \in V_s \quad \forall s \in \mathcal{S} \end{aligned}$$

A natural way to strengthen the LP-relaxation of ILP1 is by completely defining $\text{Conv}(V_s)$, the convex hull of V_s . The following theorem provides a complete characterization of $\text{Conv}(V_s)$.

Theorem 1 (Convex Hull of V_s). *For any $s \in \mathcal{S}$, $\text{Conv}(V_s) = C_s$, where C_s is a convex polytope defined by the following constraints:*

- (C1). $\sum_{t=1}^T x_{i,t} = 1, \forall i \in \mathcal{R}_s$
- (C2). $\sum_{t=1}^T y_{s,t} = 1, \forall s \in \mathcal{S}$
- (C3). $\sum_{t'=1}^t x_{i,t'} \leq \sum_{t'=1}^{\min\{t+l_s, T\}} y_{s,t'}, \forall i \in \mathcal{R}_s, t \in [T]$
- (C4). $\sum_{t'=1}^t y_{s,t'} \leq \sum_{t'=1}^t x_{i,t'}, \forall i \in \mathcal{R}_s, t \in [T]$
- (C5). $0 \leq x_{i,t} \leq 1, \forall i \in \mathcal{R}_s, t \in [T]$
- (C6). $0 \leq y_{s,t} \leq 1, \forall s \in \mathcal{S}, t \in [T]$.

(C1) – (C2) are the assignment constraints, and (C5) – (C6) are relaxations of the integrality constraints. (C3) – (C4) combined can be viewed as a stronger form of the time-window constraint (1c). Specifically, for any school s and route $i \in \mathcal{R}_s$, (C3) implies that the start time of s is no later than the arrival time of i plus a time-window length l_s . (C4) implies that the arrival time of i is no later than the start time of s .

Proof of Theorem 1. We first prove that $\text{Conv}(V_s) \subseteq C_s, \forall s \in \mathcal{S}$. Because the set of extreme points of $\text{Conv}(V_s)$ is exactly V_s , it suffices to show that $V_s \subseteq C_s$.

This is equivalent to showing that any $(x_s, y_s) \in V_s$ satisfies (C1) – (C6). It is easy to see that (C1), (C2), (C5) and (C6) follow from the assignment constraints (1a)–(1b), and the binary constraints (1e)–(1f). It remains to show that (x_s, y_s) satisfies (C3) and (C4) as well.

For any $(x_s, y_s) \in V_s$ and $i \in \mathcal{R}_s$, let $t_i, t_s \in [T]$ be indices such that $x_{i,t_i} = 1$ and $y_{s,t_s} = 1$. In constraint (1c), let $t = t_i$. We have $1 \leq \sum_{t'=t_i}^{\min\{t_i+l_s, T\}} y_{s,t'}$, implying that $t_i \leq t_s \leq t_i + l_s$. Note (C3) and (C4) can be derived from $t_s \leq t_i + l_s$ and $t_i \leq t_s$, respectively. Thus, (x_s, y_s) satisfies (C1) – (C6).

We next prove that $C_s \subseteq \text{Conv}(V_s), \forall s \in \mathcal{S}$. As a first step, we show that every integral point in C_s is also in V_s . For any integral $(x_s, y_s) \in C_s$ and $i \in \mathcal{R}_s$, let $t_i, t_s \in [T]$ be indices such that $x_{i,t_i} = 1$ and $y_{s,t_s} = 1$. Constraints (C3) and (C4) imply that $t_s \leq t_i + l_s$ and $t_i \leq t_s$, which together indicate that (x_s, y_s) satisfies (1c). Because (1a), (1b), (1e), and (1f) naturally follow from (C1), (C2), (C5), and (C6), we have $(x_s, y_s) \in V_s$.

Now consider the linear projection $f: C_s \rightarrow P_s$ defined by

$$f(x_s, y_s) = (X_s, Y_s) \text{ where } X_{i,t} = \sum_{t'=1}^t x_{i,t'}, \text{ and } Y_{s,t} = \sum_{t'=1}^t y_{s,t'},$$

and P_s is the polytope defined by the following linear inequalities:

- (P1) $X_{i,T} = 1, \forall i \in \mathcal{R}_s$
- (P2) $Y_{s,T} = 1, \forall s \in \mathcal{S}$
- (P3) $X_{i,t} \leq Y_{s, \min\{t+l_s, T\}}, \forall i \in \mathcal{R}_s, t \in [T]$
- (P4) $Y_{s,t} \leq X_{i,t}, \forall i \in \mathcal{R}_s, t \in [T]$
- (P5) $0 \leq X_{i,t} \leq X_{i,t+1}, \forall i \in \mathcal{R}_s, t \in [T-1]$
- (P6) $0 \leq Y_{s,t} \leq Y_{s,t+1}, \forall s \in \mathcal{S}, t \in [T-1]$.

The linear constraints (P1) – (P6) correspond to (C1) – (C6) by transforming (x, y) variables to (X, Y) variables. It is easy to verify that the inverse projection

$f^{-1} : P_s \rightarrow C_s$ exists, and is linear:

$$\begin{aligned} f^{-1}(X_s, Y_s) &= (x_s, y_s) \text{ where } x_{i,t} = X_{i,t} - X_{i,t-1} \\ &\text{and } y_{s,t} = Y_{s,t} - Y_{s,t-1}. \end{aligned} \quad (1)$$

Here we define $X_{i,0} = 0, Y_{s,0} = 0$.

Moreover, the matrix containing (P1) – (P6) is totally unimodular because it has no more than two nonzero entries in each column (constraint) and every column (constraint) that contains two nonzero entries has exactly one +1 and one –1 (Keller and Tompkins 1956). Because the matrix defined by (P1) – (P6) is totally unimodular, all extreme points of P_s are integral.

For any $(\tilde{x}_s, \tilde{y}_s) \in C_s$, note that $f(\tilde{x}_s, \tilde{y}_s) \in P_s$. By representing $f(\tilde{x}_s, \tilde{y}_s)$ as a linear combination of extreme points of P_s , we have

$$\begin{aligned} (\tilde{x}_s, \tilde{y}_s) &= f^{-1}(f(\tilde{x}_s, \tilde{y}_s)) \\ &= f^{-1}\left(\sum_{k=1}^K \lambda_k p_k\right) \\ &= \sum_{k=1}^K \lambda_k f^{-1}(p_k) \end{aligned}$$


where each λ_k is nonnegative, $\sum_{k=1}^K \lambda_k = 1$, and all p_k are integral extreme points of P_s . The last step applies the linearity of f^{-1} . Figure 2 visualizes each step of the projection.

From Equation (1), $f^{-1}(p_k) \in C_s$ is also integral given that p_k is integral for every k . Thus $(\tilde{x}_s, \tilde{y}_s)$ can be written as a convex combination of integral points in C_s . Because these integral points are also in V_s , $(\tilde{x}_s, \tilde{y}_s)$ must be in the convex hull of V_s , which immediately implies $C_s \subseteq \text{Conv}(V_s)$. This completes the proof of Theorem 1. $\square$

With Theorem 1, we obtain formulation ILP2 by replacing (1c) in ILP1 with (C3) and (C4).

$$\begin{aligned} \min z & \quad (ILP2) \\ \text{s.t. } \sum_{t=1}^T x_{i,t} &= 1 \quad \forall i \in \mathcal{R}_s, s \in \mathcal{S} \quad (2a) \end{aligned}$$

Figure 2. Projections Between Polytopes C_s and P_s



$$\sum_{t=1}^T y_{s,t} = 1 \quad \forall s \in \mathcal{S} \quad (2b)$$

$$\sum_{t'=1}^t x_{i,t'} \leq \sum_{t'=1}^{\min\{t+l_i, T\}} y_{s,t'} \quad \forall i \in \mathcal{R}_s, s \in \mathcal{S}, t \in [T] \quad (2c)$$

$$\sum_{t'=1}^t y_{s,t'} \leq \sum_{t'=1}^t x_{i,t'} \quad \forall i \in \mathcal{R}_s, s \in \mathcal{S}, t \in [T] \quad (2d)$$

$$\sum_{i \in \mathcal{R}} \sum_{t'=t}^{\min\{t+r_i-1, T\}} x_{i,t'} \leq z \quad \forall t \in [T] \quad (2e)$$

$$x_{i,t} \in \{0, 1\} \quad \forall i \in \mathcal{R}_s, s \in \mathcal{S}, t \in [T] \quad (2f)$$

$$y_{s,t} \in \{0, 1\} \quad \forall s \in \mathcal{S}, t \in [T] \quad (2g)$$

As a remark on Theorem 1, we note that for instances with external constraints on start time of schools and arrival time of routes, we can redefine set V_s by adding constraints (corresponding to the external constraints) in the form of: $x_{i,t} = 0$ or $y_{s,t} = 0$. Polytope P_s can be redefined in a similar way, by adding constraints in the form of $X_{i,t-1} = X_{i,t}, Y_{i,t-1} = Y_{i,t}$. Incorporating these equality constraints does not affect the totally unimodularity of the matrix that defines P_s . Each equality constraint has one term with coefficient +1 on the left-hand side and one term with coefficient +1 on the right-hand side, which satisfies the condition for totally unimodular matrix in Keller and Tompkins (1956). For problem instances with external constraints on school start times and route arrival times, ILP2 (plus linear constraints corresponding to external constraints) remains a valid strengthened formulation.

4. Randomized Rounding Algorithm for the SBSP


We develop a dependent randomized rounding algorithm for the SBSP based on the LP relaxation of ILP2. Banerjee et al. (2021) observe that such algorithms work well on routing problems. Our algorithm constructs a feasible integral solution to ILP2 based on a fractional optimal solution of the LP relaxation of ILP2. We show in Theorem 2 that the optimality gap of the resulting integral solution is bounded by the square root of the optimal solution value. We show in Corollary 1 that this also gives a bounded integrality gap of the LP relaxation of ILP2.

Algorithm 1 (Randomized Rounding Algorithm for the SBSP)

Step 1: solve the LP relaxation of ILP2 to get $\{(x_s^*, y_s^*)\}_{s \in \mathcal{S}}$

Step 2: for each school $s \in \mathcal{S}$, draw a uniform random number $r_s \sim U[0, 1]$

Step 3: for each route $i \in \mathcal{R}_s$, let $t_i = \arg \min\{\sum_{t=1}^{t_i} x_{i,t}^* \geq r_s\}$ and set $x_{i,t_i} = 1$

Figure 3. Randomized Rounding Algorithm

Step 4: for each school $s \in \mathcal{S}$, let $t_s = \operatorname{argmin}\{\sum_{t=1}^{t_s} y_{s,t}^* \geq r_s\}$ and set $y_{s,t_s} = 1$

Step 5: solve the route-to-bus assignment problem using the scheduling from Steps 3 and 4

We illustrate the algorithm with an example of one school and two routes (see Figure 3). For each school and route, we divide a $[0, 1]$ segment into T intervals (some may have width 0 if the corresponding variable has value 0) based on the fractional optimal solution $\{(x_s^*, y_s^*)\}$. We then draw a random number $r_s \sim U[0, 1]$ (Step 2) and cut all $[0, 1]$ segments at r_s (Steps 3 and 4). For each school (route), the index of the interval where the random number r_s falls is selected as its start (arrival) time.

Let $\{(x_s, y_s)\}_{s \in \mathcal{S}}$ be the integral solution generated by Algorithm 1. In Proposition 2, we show that the integral solution is always feasible. In Proposition 3, we show that the fractional optimal solution $\{(x_s^*, y_s^*)\}_{s \in \mathcal{S}}$ defines the probability that each corresponding variable takes the value 1 in the integer solution generated by Algorithm 1. Finally, we prove in Theorem 2 that the integral solution is near-optimal with constant probability.

Proposition 2 $\{(x_s, y_s)\}_{s \in \mathcal{S}}$ satisfies constraints (2a)–(2d) and (2f)–(2g).

Proof. In Steps 3 and 4 of Algorithm 1, each route (school) is assigned to exactly one arrival (start) time. Hence, the assignment constraints (2a)–(2b), and the binary constraints (2f)–(2g) naturally hold. For the time-window constraints (2c)–(2d) to also hold, it suffices to show that $t_i \leq t_s \leq t_i + l_s$ for any school $s \in \mathcal{S}$ and route $i \in \mathcal{R}_s$.

Recall that t_i and t_s are determined by the random variable r_s generated in Step 2. From the definition of t_i (Step 3) and t_s (Step 4), we have

$$\sum_{t=1}^{t_i-1} x_{i,t}^* < r_s \leq \sum_{t=1}^{t_i} x_{i,t}^* \quad (2)$$

and

$$\sum_{t=1}^{t_s-1} y_{s,t}^* < r_s \leq \sum_{t=1}^{t_s} y_{s,t}^*. \quad (3)$$

Because the fractional solution $\{(x_s^*, y_s^*)\}_{s \in \mathcal{S}}$ satisfies the time-window constraints (2c)–(2d), we have

$$\sum_{t=1}^{t_s-1} y_{s,t}^* < r_s \leq \sum_{t=1}^{t_i} x_{i,t}^* \leq \sum_{t=1}^{t_i+l_s} y_{s,t}^*. \quad (4)$$

Comparing the left-hand and right-hand sides of (4), $t_s - 1 < t_i + l_s$, that is, $t_s \leq t_i + l_s$. Similarly, from

$$\sum_{t=1}^{t_i-1} x_{i,t}^* < r_s \leq \sum_{t=1}^{t_s} y_{s,t}^* \leq \sum_{t=1}^{t_s} x_{i,t}^* \quad (5)$$

we have $t_i - 1 < t_s$, that is, $t_s \geq t_i$. This completes the proof of Proposition 2. $\square$

Proposition 3 $\Pr[x_{i,t} = 1] = x_{i,t}^* \quad \forall i \in \mathcal{R}, t \in [T]; \Pr[y_{s,t} = 1] = y_{s,t}^* \quad \forall s \in \mathcal{S}, t \in [T]$.

Proof. $x_{i,t} = 1$ if and only if the random number r_s falls within the interval $[\sum_{t'=1}^{t-1} x_{i,t'}^*, \sum_{t'=1}^t x_{i,t'}^*]$. Because $r_s \sim U[0, 1]$,

$$\Pr[x_{i,t} = 1] = \sum_{t'=1}^t x_{i,t'}^* - \sum_{t'=1}^{t-1} x_{i,t'}^* = x_{i,t}^*.$$

The proof for y variables follows the same argument. $\square$

Theorem 2 (Optimality Gap of Algorithm 1). Let OPT be the optimal solution value to the SBSP and let z_{round} be the solution value obtained from Algorithm 1. We have

$$z_{\text{round}} \leq OPT + \sqrt{2R_{\max} \log(2T) OPT} + R_{\max} \log(2T)$$

with probability at least $\frac{1}{2}$, where $R_{\max}$ is the maximum number of routes in one school.

Thus, the optimality gap of Algorithm 1 is bounded by $\sqrt{2R_{\max} \log(2T) OPT} + R_{\max} \log(2T)$. In practice, $R_{\max}$

is bounded by the maximum number of students in a school (divided by bus capacity) and the term $\log(2T)$ can be treated as a constant. If the number of routes for each school is uniformly bounded and $\log(2T)$ is (or can be bounded by) a constant, the optimality gap is of order $\sqrt{OPT}$. Probability $\frac{1}{2}$ can be boosted to $(1 - \varepsilon)$ for any $\varepsilon > 0$ by running Algorithm 1 $\lceil \log(1/\varepsilon) \rceil$ times and selecting the best solution.

Proof of Theorem 2. Let $\{(x_s^*, y_s^*)\}_{s \in S}$ be an optimal fractional solution to the LP relaxation of ILP2 and let $\{(x_s, y_s)\}_{s \in S}$ be the integral solution obtained from Algorithm 1. Let z_{LP}^* be the optimal objective value of the LP relaxation of ILP2 and let z_{round} be the objective value of ILP2 using the integral solution $\{(x_s, y_s)\}_{s \in S}$.

For any $t \in [T]$, let $z_t = \sum_{i \in \mathcal{R}} \sum_{t'=t}^{\min\{t+r_i-1, T\}} x_{i,t'}$. From Proposition 3,

$$\mathbb{E}[z_t] = \mathbb{E}\left[\sum_{i \in \mathcal{R}} \sum_{t'=t}^{\min\{t+r_i-1, T\}} x_{i,t'}\right] = \sum_{i \in \mathcal{R}} \sum_{t'=t}^{\min\{t+r_i-1, T\}} x_{i,t'}^* \leq z_{LP}^*.$$

Note that $z_t = \sum_{i \in \mathcal{R}} \sum_{t'=t}^{\min\{t+r_i-1, T\}} x_{i,t'} = \sum_{s \in S} (\sum_{i \in \mathcal{R}_s} \sum_{t'=t}^{\min\{t+r_i-1, T\}} x_{i,t'})$ is the sum of $|S|$ independent random variables (x variables for different schools are independent) where each is in $[0, R_{max}]$. The next step is to bound the deviation of z_t from its mean with high probability for each $t \in [T]$, which can be derived from the Chernoff bound (Chernoff 1952).

Lemma 1 (Chernoff Bound). *Let $a_1, a_2, \dots, a_n$ be n independent random variables in $[0, M]$ and $A = \sum_{i=1}^n a_i$ with mean $\mathbb{E}[A] \leq \mu$. For any $\varepsilon > 0$,*

$$\Pr[A > \mu + \varepsilon] \leq \exp\left(-\frac{\varepsilon^2}{(2\mu + \varepsilon)M}\right).$$

For any $t \in [T]$, note that

$$z_t = \sum_{i \in \mathcal{R}} \sum_{t'=t}^{\min\{t+r_i-1, T\}} x_{i,t'} = \sum_{s \in S} \left(\sum_{i \in \mathcal{R}_s} \sum_{t'=t}^{\min\{t+r_i-1, T\}} x_{i,t'} \right)$$

is the sum of $|S|$ independent random variables where each of them lies in $[0, R_{max}]$. Further,

$$\mathbb{E}[z_t] = \sum_{i \in \mathcal{R}} \sum_{t'=t}^{\min\{t+r_i-1, T\}} x_{i,t'}^* \leq z_{LP}^*.$$

Let ε^* be the positive root of $\exp\left(-\frac{(\varepsilon^*)^2}{(2z_{LP}^* + \varepsilon^*)R_{max}}\right) = \frac{1}{2T}$. Using Lemma 1, for all $t \in [T]$

$$\Pr[z_t > z_{LP}^* + \varepsilon^*] \leq \frac{1}{2T}.$$

Applying a union bound over all $t \in [T]$, we have

$$\Pr[\max_{t \in [T]} \{z_t\} > z_{LP}^* + \varepsilon^*] \leq T \cdot \frac{1}{2T} = \frac{1}{2}.$$

From constraint (2d) in ILP2,

$$z_{round} = \max_{t \in [T]} \left\{ \sum_{i \in \mathcal{R}} \sum_{t'=t}^{\min\{t+r_i-1, T\}} x_{i,t'} \right\} = \max_{t \in [T]} \{z_t\}.$$

Therefore, with probability at least $\frac{1}{2}$, $z_{round} \leq z_{LP}^* + \varepsilon^*$. It is easy to calculate that

$$\begin{aligned} \varepsilon^* &= \frac{R_{max} \log(2T) + \sqrt{(R_{max} \log(2T))^2 + 8R_{max} \log(2T) z_{LP}^*}}{2} \\ &\leq \frac{R_{max} \log(2T) + \sqrt{(R_{max} \log(2T))^2} + \sqrt{8R_{max} \log(2T) z_{LP}^*}}{2} \\ &= R_{max} \log(2T) + \sqrt{2R_{max} \log(2T) z_{LP}^*}. \end{aligned}$$

Note that $z_{LP}^* \leq OPT$, this completes the proof of Theorem 2. $\square$

In fact, the proof of Theorem 2 implies the stronger result that with probability at least $\frac{1}{2}$,

$$z_{round} \leq z_{LP}^* + R_{max} \log(2T) + \sqrt{2R_{max} \log(2T) z_{LP}^*}.$$

Because z_{round} is always greater than or equal to OPT , we have the following corollary.

Corollary 1 (Integrality Gap of ILP2).

$$z_{LP}^* \leq OPT \leq z_{LP}^* + R_{max} \log(2T) + \sqrt{2R_{max} \log(2T) z_{LP}^*}.$$

Because z_{LP}^* and OPT are deterministic values for a given instance, the probabilistic condition in Theorem 2 can be dropped in Corollary 1.

The randomized rounding algorithm and its theoretical guarantee do not rely on specific values of the optimal solution to ILP2; that is, $\{(x_s^*, y_s^*)\}_{s \in S}$. Adding external constraints on school start and route arrival times does not affect the correctness of the algorithm and its theoretical guarantee. Further, the optimality gap proved in Theorem 2 is the gap between the upper bound obtained from Algorithm 1 and z_{LP}^* - a lower bound of the optimal solution OPT . In practice, the optimality gap of Algorithm 1 is considerably smaller than the one proved in Theorem 2.

5. Robust SBSP

The route set, $\mathcal{R}$, may vary over time due to enrollment changes and route redesigns. School start times, however, cannot be changed frequently. Decision makers are interested in plans for a long time period. We introduce a robust SBSP, generalizing results of Sections 3 and 4.

When start times are set, uncertainty may exist in the number and travel time of routes over the planning horizon. We model uncertainty in the route set $\mathcal{R}$ with scenarios where each scenario has a set of routes that may differ by number and travel time of routes.

Table 1. Comparison of the SBSP ($|\Pi|=1$) and the Robust SBSP ($|\Pi|>1$)

Number of scenarios	Single ($ \Pi =1$)	Multiple ($ \Pi >1$)
Number of ILP variables	$O(T(\mathcal{S} + \mathcal{R}))$	$O(T(\mathcal{S} + \mathcal{R} \cdot \Pi))$
Number of ILP constraints	$O(T(\mathcal{S} + \mathcal{R}))$	$O(T(\mathcal{S} + \mathcal{R} \cdot \Pi))$
Optimality gap of rounding algorithm	$O(\sqrt{R_{\max} \log(T) OPT})$	$O(\sqrt{R_{\max} \log(\Pi \cdot T) OPT})$

Scenarios can be derived from historical route data or from routes designed to accommodate projected future enrollment distributions. Our goal is to determine school start times and route schedules such that all scenarios share the same school start times, minimizing the maximum number of buses over all scenarios. Let Π be the set of scenarios. In scenario $\pi \in \Pi$, $\mathcal{R}^\pi$ is the route set, $\mathcal{R}_s^\pi$ is the route set for school s and r_i^π is the travel time of route i . The robust SBSP can be modeled as follows.

$$\min z_{\text{robust}} \quad (\text{ILP-robust})$$

$$\text{s.t.} \quad \sum_{t=1}^T x_{i,t}^\pi = 1 \quad \forall i \in \mathcal{R}, \pi \in \Pi \quad (3a)$$

$$\sum_{t=1}^T y_{s,t} = 1 \quad \forall s \in \mathcal{S} \quad (3b)$$

$$\sum_{t'=1}^t x_{i,t'}^\pi \leq \sum_{t'=1}^{\min\{t+l_s, T\}} y_{s,t'} \quad \forall i \in \mathcal{R}_s^\pi, s \in \mathcal{S}, t \in [T], \pi \in \Pi \quad (3c)$$

$$\sum_{t'=1}^t y_{s,t'} \leq \sum_{t'=1}^t x_{i,t'}^\pi \quad \forall i \in \mathcal{R}_s^\pi, s \in \mathcal{S}, t \in [T], \pi \in \Pi \quad (3d)$$

$$\sum_{i \in \mathcal{R}} \sum_{t'=t}^{\min\{t+r_i^\pi-1, T\}} x_{i,t'}^\pi \leq z_{\text{robust}} \quad \forall t \in [T], \pi \in \Pi \quad (3e)$$

$$x_{i,t}^\pi \in \{0, 1\} \quad \forall i \in \mathcal{R}^\pi, t \in [T], \pi \in \Pi \quad (3f)$$

$$y_{s,t} \in \{0, 1\} \quad \forall s \in \mathcal{S}, t \in [T] \quad (3g)$$

ILP-robust can be viewed as an extension of ILP2, minimizing the maximum number of buses across scenarios. Binary variable $x_{i,t}^\pi$ indicates the arrival time of route i in scenario π . Because all scenarios share the same school start times, binary variable $y_{s,t}$ indicates the start time of school s across all scenarios. Similar to ILP2, **ILP-robust** consists of assignment constraints (3a)–(3b), time-window constraints (3c)–(3d), linking constraints (3e) and binary constraints (3f)–(3g).

Similar to Algorithm 1, after solving the LP relaxation of **ILP-robust**, for school s , the start time of s and arrival times of routes in s over all scenarios are determined simultaneously through the rounding algorithm to ensure feasibility. We use the same probabilistic approach (detailed in Theorem 2) to provide an optimality gap of the rounding algorithm for the robust SBSP.

Table 1 compares the SBSP ($|\Pi|=1$) and the robust SBSP ($|\Pi|>1$) in terms of formulation size and theoretical optimality gap. The number of variables and constraints in **ILP-robust** increases linearly with the number of scenarios and the optimality gap of the rounding algorithm includes an additional term $\sqrt{\log(|\Pi|)}$. In practice, this additional term $\sqrt{\log(|\Pi|)}$ can be viewed as a constant; the rounding algorithm still provides near-optimal solutions for large-scale instances.

6. Numerical Studies

We present a case study based on two school districts of different size, followed by experiments with synthetic data to complement our theoretical findings.

6.1. Case Study

This work is motivated by collaborations with two public school districts: one moderately sized with ~20 schools and one large with ~200 schools. Both districts wished to explore transportation cost savings from changes to school start times. For the moderate district, our work focuses on 39 routes serving 13 schools. For the larger district, we focus on 380 bus routes for 127 schools.

The districts currently use commercial software to determine routes. With these routes, we optimize school and bus scheduling using our solution approach and evaluate potential cost reductions from more flexible start time options. Although both districts assume a constant transition time when determining routes to pair, we explore the impact of location-dependent transition times by extending our approach in Section 6.2.4.

6.1.1. District-Specific Modifications. Each district has specific operational considerations detailed, which we incorporate into our approach.

Arrival time windows. In Problem 1, routes arrive in a time window before school starts. If school s starts at time t_s with a time window l_s , routes arrive in interval $[t_s - l_s, t_s]$. Both districts use a fixed arrival time q_s prior to the start time rather than a window; routes arrive at $t_s - q_s$. For the moderate district, q_s is 15 minutes for two schools and 10 minutes for the others. For the larger district, q_s varies by school between 5 and 30 minutes. We simplify constraints (2c) and (2d) as

$$x_{i,t} = y_{s,t+q_s} \quad \forall i \in \mathcal{R}_s, s \in \mathcal{S}, t \in [T].$$

We also relax this constraint to assess the value of introducing arrival time windows.

Restriction on school start times. As is common, both districts have restricted times when schools may start. The moderate district operates two tiers of start time options:

- Tier 1: 8:05 a.m., 8:30 a.m.;
- Tier 2: 9:05 a.m.

For the large district, these are:

- Tier 1: 7:30 a.m., 7:35 a.m., 7:40 a.m.;
- Tier 2: 8:05 a.m., 8:10 a.m., 8:15 a.m., 8:20 a.m., 8:25 a.m.;
- Tier 3: 8:45 a.m., 8:50 a.m., 8:55 a.m., 9:00 a.m., 9:05 a.m.

In ILP2, let $T_{possible}(s) \subseteq [T]$ be the set of possible start times for school s . We add a constraint for start time restrictions by tier options and school type.

$$y_{s,t} = 0 \quad \forall s \in \mathcal{S}, t \notin T_{possible}(s).$$

6.1.2. Improving Solutions and Expanding Flexibility.

Given current operating constraints and constant transition times, we examine the extent to which our approach improves on current practice. Table 2 compares the current number of buses (column 3) with solutions obtained using our approach for both constant and location dependent transition times. For each district, schools can begin in any tier option for that district (i.e., a common set of tiers for all schools, $T_{possible}$). For the moderate district, we also evaluated a set of tiers restricted by school type, $T_{possible}(s)$, and found that it was not possible to reduce the number of buses through improved route scheduling. This highlighted to the district the limits on cost reductions set by their current tiers, prompting a further study of where best to introduce flexibility.

With constant transition times, Table 2 shows a reduction in the number of buses (compared with current practice) by 36.4% for the moderate district and 14.5% for the large district. The rounding algorithm (best from 100 runs) provides the optimal solution for the moderate district and is within 5% of the lower bound for the large district.


The last two columns in Table 2 present solutions when the assumption of constant transition time is relaxed. We estimate transition times for each pair of routes based on geographical information of the start and end points of routes, plus an added three minutes of buffer time. Travel times are derived from the Distance Matrix API in Google Map API (Wang and Xu 2011). We apply the modified randomized rounding algorithm described in Section 6.2.4 to solve the problem with location dependent transition times. For both districts, the use of location dependent transition times reduces the number of buses required, but the solution obtained remains within 10% of the solution values with constant transition times. This shows that our approach works well with nonconstant transition times. Well-estimated constant transition times, however, provide good solutions in practice if they are representative of travel times between nearby schools and include a buffer that is independent of the actual distance between routes.

Although the results in Table 2 represent cost savings, they may also represent significant deviations from current operations. Many practical considerations can limit the set of feasible start times. To allow school administrators the opportunity to consider many low-cost transportation options when making strategic choices, we explore the ability of the randomized rounding algorithm to generate multiple near-optimal solutions. We generate 100 solutions using the randomized rounding algorithm for each district and record the number of unique solutions within 10% of the best solution value. For the moderate district, we obtain 52 solutions requiring no more than 23 buses. For the large district, we obtain 83 solutions requiring no more than 133 buses. This demonstrates the ability of our approach to provide multiple scheduling plans with near-optimal costs.

The discussion is based on a district-specific constraint that bus routes arrive at a fixed time before the school start time (i.e., the length of arrival time window for each school is equal to 0). Increasing the length of this time window can potentially reduce the number of buses because it adds flexibility to route

Table 2. SBSP Solutions for Fixed Route Arrival Time (No Arrival Windows)

		Constant transition		Location-dependent transition		
	Size	Current	Lower bound (CPU time)	Rounding (CPU time)	Lower bound (CPU time)	Rounding (CPU time)
Moderate	39 routes	33	21	21	17	19
	13 schools		(3.1)	(3.3)	(3.3)	(3.6)
Large	380 routes	144	117	123	82	116
	127 schools		(297.4)	(303.2)	(297.8)	(316.1)

Figure 4. Value of Arrival Time Window Expansion

Note. (a) Moderate district; (b) Large district.

arrival times, thus increasing the set of feasible route connections.

Figure 4 illustrates the value of flexible arrival time windows by comparing solutions with increasing time window lengths from 0 to 20 minutes. The dashed line represents solutions with constant transition times and the solid line represents solutions with location-dependent transition times.

Figure 4 shows that increasing the arrival time windows typically decreases the number of buses required. For a large enough time window (here, 20 minutes), the benefit of location-dependent transition times disappears. For a large enough time window, using a well-estimated constant transition time gives a solution that is as good as using location-dependent transition times.

6.2. Numerical Analysis of Formulations and Solution Methods

We extend the case study findings with synthetic data for a broader analysis of our approach. The case study shows the power of the strengthened formulation (ILP2) to provide tight lower bounds and the near-optimality of the rounding algorithm. We explore these results for a wider range of instances. We also explore our robust SBSP approach and location-dependent transition times.

6.2.1. SBSP Instances and Benchmarks. We generate ten SBSP instances representing a range of district sizes; the number of schools and routes are (10, 50), (20, 100), ..., (100, 500). Across all instances, each school has a fixed time window of 20 minutes. School start times and route arrival times can be chosen from $\{1, 2, \dots, 120\}$, which corresponds to minutes in a two-hour time window. School start times are further restricted to multiples of five minutes; that is, $\{5, 10, \dots, 120\}$. The route set for

each instance is generated with Algorithm 2. All generated data can be found at https://github.com/zenglw1206/school_bus_scheduling/.

Algorithm 2 (Route Generation for SBSP Instances).

Step 1: Generate a 100×100 grid and locate schools on integral points of the grid at random

Step 2: For each route, randomly select an integral point of the grid as its starting point and a school location as its ending point

Step 3: Route travel time is computed by the L_1 distance between starting and ending points, divided by constant travel speed v_1 , rounded to the nearest integer. The travel speed v_1 is determined such that the average route travel time is approximately 25 minutes.

For each instance, after solving the LP relaxation, we perform the randomized rounding algorithm 10 times and select the best solution. As benchmarks, we implement the integer programming approach from Fügenschuh (2009) and a local search heuristic. Our problem can be viewed as a special case of the formulation in Fügenschuh (2009) where mixed loads and non-constant transition times are allowed. As described in the Appendix, we tailor their IP approach to accommodate our context-specific features. Local search heuristics have been widely applied to scheduling problems due to their simplicity and scalability, including for school bus scheduling (Spada, Bierlaire, and Liebling 2005; Chen et al. 2015; Bertsimas, Delarue, and Martin 2019). In each iteration of these heuristics, a typical strategy is to optimize the objective for a small set of schools, keeping decisions for all other schools fixed. In our implementation, we select one school at random in each iteration and optimize its start time by enumerating all possible start time choices (and shift arrival times for routes in the selected school to maintain feasibility) and then select the one that requires the minimum number

of buses. We apply two policies to initialize the local search heuristic: first, we use the solution of the randomized rounding algorithm as a starting point; second, we draw multiple starting points at random and select the best one.

6.2.2. Comparing Formulations and Solution Approaches. We illustrate the strength of our ILP formulations by comparing the lower bounds obtained from the LP relaxations of ILP1 and ILP2 with the exact optimal solution value (obtained by solving ILP2) as well as the computational times (in seconds) to obtain these bounds and solutions. As shown in Table 3, the relaxation of ILP2 outperforms that of ILP1 and gives tight lower bound for all instances. In nine of the 10 instances, rounding up the lower bound from the LP-relaxation of ILP2 provides the optimal integer value. The largest instance can be solved to optimality within 30 minutes using formulation ILP2.

For larger instances where exact methods are intractable, we compare the lower bounds from ILP1 and ILP2 with the feasible solution from the rounding algorithm. We test the algorithms on three sets of instances varying the number of schools and routes (each with five instances) and report the average and maximum relative gap for each set. The relative gap of ILP1 (ILP2) is defined as:

$$\text{relative gap} = \left(\frac{\text{rounding solution}}{\text{lower bound from ILP1 (ILP2)}} - 1 \right) \times 100\%.$$

As shown in Table 4, the relative gap for ILP2 is smaller and shrinks faster as the instance size increases. This supports our theoretical findings that ILP2 is a strong formulation for the SBSP and that the randomized rounding algorithm provides high-quality solution for large instances.

We next evaluate the efficiency of our approach relative to other approaches within our context-specific features. Recall that the goal of this work is not to develop an approach for the general SBSP, but to show that leveraging specific features can lead to more efficient approaches to assist with strategic decision making. Table 5 compares the exact solution, the lower bound obtained from the LP relaxation of ILP2 and five upper bounds. The first three upper bounds are (i) rounding algorithm, (ii) rounding algorithm followed by local search, and (iii) local search algorithm. We set a two-hour time limit for local search. For the second policy, we use the first 30 minutes to select starting points, followed by a two-hour local search. The final upper bounds come from our implementation of the SBSP-IP from Fügenschuh (2009). Column “IP-2.5” (“IP-24”) corresponds to bounds obtained by running the IP for 2.5 (24) hours.

The randomized rounding algorithm achieves an average optimality gap (optimality gap = {rounding solution value - optimal solution value}/optimal solution value) of 12.9%, which improves to 10.3% with local search. The local search heuristic yields an average optimality gap of 37.3% and the IP approach has an average optimality gap of 39.6% (in 2.5 hours) and 33.0% (in 24 hours). Moreover, the rounding

Table 3. Comparison Between LP Relaxations of ILP1 and ILP2

Instance size	Exact solution (CPU time)	Lower bound (relaxation of ILP1) (CPU time)	Lower bound (relaxation of ILP2) (CPU time)
10 schools	9	7.7	8.5
50 routes	(102.9)	(1.2)	(4.8)
20 schools	17	15.6	16.5
100 routes	(174.2)	(3.1)	(10.9)
30 schools	24	22.0	23.9
150 routes	(345.4)	(4.4)	(16.8)
40 schools	32	30.8	31.1
200 routes	(464.2)	(6.2)	4(29.6)
50 schools	42	38.1	40.1
250 routes	(548.2)	(10.1)	(31.5)
60 schools	51	48.1	50.9
300 routes	(864.6)	(11.0)	(42.8)
70 schools	61	58.2	60.5
350 routes	(937.1)	(15.4)	(55.1)
80 schools	65	61.5	64.5
400 routes	(1048.6)	(13.9)	(75.7)
90 schools	76	71.7	75.2
450 routes	(1361.8)	(19.5)	(97.5)
100 schools	84	80.6	83.8
500 routes	(1782.3)	(21.8)	(99.3)

Table 4. Comparison Between LP Relaxations of ILP1 and ILP2 for Large Instances (Over 5 Instances)

Instance size	Average/maximum relative gap of ILP1	Average/maximum relative gap of ILP2
200 schools 1,000 routes	10.7%/12.5%	5.7%/7.1%
500 schools 2,500 routes	8.2%/8.6%	3.7%/4.2%
1000 schools 5,000 routes	7.3%/7.8%	2.7%/3.1%

algorithm solves the largest instance within two minutes (for each one of the 10 runs). These results suggest that the randomized rounding algorithm is very fast and provides better solutions compared with benchmark algorithms for our context-specific features.

As in the case study, we evaluate the ability of our approach to generate multiple near-optimal solutions. In our experiments, we record all solutions obtained with the rounding algorithm, in addition to the best solutions reported in Table 5. We find that all generated solutions are within 10% of the best solution and 73.3% of solutions are within 5% of the best solution, further demonstrating that the rounding algorithm is able to generate multiple high quality solutions quickly.

6.2.3. Analysis of the Robust SBSP. We incorporate two sources of uncertainty into the route set $\mathcal{R}$: (i) uncertainty in the number of routes, and (ii) uncertainty in travel times. For uncertainty (i), we increase or decrease the number of routes in each school by 1, each with probability 15%. For uncertainty (ii), we

add a random integer from $[-5, 5]$ to the travel time of each route. The transition times are perturbed in a similar way. We construct three groups of instances with one or both types of uncertainty: in each set, we use the 10 combinations of routes and schools from 6.2.2 with five route set scenarios.

Table 6 presents results with different sources of uncertainty. Compared with Tables 4 and 5, the average relative gap of the rounding algorithm increases from 12.9% to 24.1% ~ 26.7%. The number of buses can be highly sensitive to small changes in route travel times and number of routes across scenarios (e.g., adding a route may require a new bus). Thus, it is harder to obtain tight lower bounds for the robust SBSP with multiple scenarios.

We observe that the relative gap between the upper and lower bounds shrinks with instance size, suggesting that formulation ILP-robust and the modified randomized rounding algorithm remain effective for the robust SBSP, especially for large-scale instances.

Table 5. Solution Comparisons with Benchmarks

Instance size	Exact solution	Lower bound (ILP2)	Rounding	Rounding + local search	Local search	IP-2.5	IP-24
10 schools 50 routes	9	9	11	11	12	12	11
20 schools 100 routes	17	17	21	19	24	24	23
30 schools 150 routes	24	24	27	27	32	32	32
40 schools 200 routes	32	32	36	35	45	44	44
50 schools 250 routes	42	41	46	46	56	55	55
60 schools 300 routes	51	51	57	55	66	65	65
70 schools 350 routes	61	61	66	65	89	89	83
80 schools 400 routes	65	65	70	69	92	94	88
90 schools 450 routes	76	76	81	80	101	111	102
100 schools 500 routes	84	84	94	92	116	128	113
Average gap			12.9%	10.3%	37.3%	39.6%	33.0%

Table 6. Robust Instances with Varying Sources of Uncertainty

Instance size	Source of uncertainty		Travel time		#Routes + route travel time	
	#Routes Lower bound	Rounding	Lower bound	Rounding	Lower bound	Rounding
10 schools	9	12	9	12	8	12
50 routes						
20 schools	16	22	17	23	17	25
100 routes						
30 schools	24	33	26	33	25	33
150 routes						
40 schools	33	45	34	43	32	38
200 routes						
50 schools	39	50	42	55	41	53
250 routes						
60 schools	48	59	51	64	46	60
300 routes						
70 schools	56	66	57	67	54	65
350 routes						
80 schools	66	78	65	77	66	76
400 routes						
90 schools	71	80	72	82	73	82
450 routes						
100 schools	81	91	82	92	84	94
500 routes						
Average gap		25.7%		24.1%		26.7%

6.2.4. Modifications for Location-Dependent Transition Times. Although the theoretical performance guarantee for the rounding algorithm relies on the assumption of constant transition times, we are able to modify the rounding algorithm to provide feasible solutions and cost lower bounds for instances with location-dependent transition times through the following steps.

- Step 1: Reduce the problem to a tractable instance
- Step 2: Determine route arrival times using the randomized rounding algorithm
- Step 3: Obtain a feasible solution for the original problem
- Step 4: Obtain a cost lower bound

Step 1: Reduction to a tractable instance
 We define a class of SBSP instances with location-dependent transition times that are tractable using the rounding algorithm.

Definition 1 (Tractable Instance). A SBSP instance is tractable if there exists t_i^{in} and t_i^{out} for each route i such that the transition time from route i to route j is equal to $t_i^{out} + t_j^{in}$ for all (i, j) .

For a tractable instance, the transition time t_{ij} can be decomposed into two parts—an “outbound transition time” t_i^{out} and an “inbound transition time” t_j^{in} . By adding these to the route travel time (i.e., a route with travel time r_i becomes one with travel time $t_i^{in} + r_i + t_i^{out}$), the instance is reduced to one with constant transition time. As in practice, we also add a constant buffer to all

transitions, as shown in the case study. The randomized rounding algorithm can then be directly applied to solve this tractable SBSP instance.

First, we seek to find a tractable instance that approximates the given instance with location-dependent transition times. In other words, we find t_i^{in} and t_i^{out} such that tran_{ij} (transition time in the given instance) can be approximated by $t_i^{out} + t_j^{in}$. To that end, we fit t_i^{in} and t_i^{out} into the following linear regression model.

$$\text{tran}_{ij} = t_i^{out} + t_j^{in} + \varepsilon_{ij} \quad (6)$$

In practice, we use least square regression (i.e., minimizing $\sum \varepsilon_{ij}^2$) to obtain t_i^{in} and t_i^{out} .

Step 2: Route arrival times

In Step 2, we apply the randomized rounding algorithm to solve the tractable instance obtained from Step 1 and determine school start times and route arrival times for the original instance.

Step 3: Feasible solution

After determining the arrival time for each route, we find the minimum number of buses to serve all the routes using the following proposition.

Proposition 4. Given a route set $\mathcal{R}$ and an arrival time of each route, consider a bipartite graph $G = (U, V; E)$ where $U = V = \mathcal{R}$. For any $r_1 \in U$ and $r_2 \in V$, $(r_1, r_2) \in E$ if and only if route r_2 can be served after route r_1 using the same bus. Let $|G_{BM}|$ be the size of the maximum bipartite matching of graph G . Then, the minimum number of buses to complete all routes is equal to $|\mathcal{R}| - |G_{BM}|$.

From Proposition 4, finding the minimum number of buses is equivalent to solving a maximum bipartite matching problem, which can be done in polynomial time. To see the correctness of Proposition 4, each route connection pattern corresponds to a bipartite matching of graph G where $(r_1, r_2) \in E$ if and only if route r_2 is served immediately after route r_1 using the same bus. Moreover, the size of this bipartite matching is equal to $|\mathcal{R}|$ minus the number of buses. Therefore, minimizing the number of buses is equivalent to finding a maximum bipartite matching on G .

Step 4: Cost lower bound

Enforcing the error term ε_{ij} in the regression model (6) to be nonnegative, such that the estimated transition time $t_i^{out} + t_j^{in}$ is less than or equal to tran_{ij} , provides a lower bound on the number of buses. By inserting the t_i^{in} and t_i^{out} , with the buffer time, into ILP2 and its LP relaxation, we are able to obtain a lower bound on the number of buses.

The transition time between routes is computed by the L_1 distance between the ending point of the first route and the starting point of the second route, divided by constant transition speed v_2 . The value of v_2 is chosen such that the average transition time is approximately 15 minutes.

In our analysis of location-dependent transition times, we explore the relative gap between the lower bound and the feasible solutions obtained from the rounding heuristic, as well as the impact of the buffer time estimation on solution values. Comparisons with benchmark approaches are included in the Appendix.

As in Table 4, we define the relative gap of ILP2 as:

$$\text{relative gap} = \left(\frac{\text{rounding solution}}{\text{lower bound from ILP2}} - 1 \right) \times 100\%.$$

To compare solutions with location-dependent transition times and varying buffer times relative to those with constant transition time, we define

change from constant

$$= \left(\frac{\text{location-dependent solution}}{\text{solution with constant transition}} - 1 \right) \times 100\%.$$

From Table 7, we observe that the relative gap of ILP2 is larger with the location-dependent SBSP. One possible explanation is that transition times are underestimated in the regression model in Step 1, which then weakens the lower bound, as compared with instances with constant transition time. These gaps decrease as the buffer increases from 0 to five minutes between routes. As with the constant transition time SBSP, we find that the relative gap of ILP2 decreases with instance size. Although we believe the larger gap mainly comes from the lower bound, it remains to be explored whether the location-dependent transition time can be better incorporated into the rounding algorithm, which is a potential topic for future research.

Note the negative gaps in columns 6 and 7 of Table 7. Without a buffer, the location-dependent transition times can be quite low, providing more opportunities to reuse buses. However, solutions may

Table 7. Instances with Location-Dependent Transition Time

Instance size	Relative gap of ILP2				Change from constant		
	Constant	Buffer = 0	Buffer = 3	Buffer = 5	Buffer = 0	Buffer = 3	Buffer = 5
10 schools	25%	50%	50%	36%	0%	0%	0%
50 routes							
20 schools	8%	37%	29%	27%	−8%	−4%	0%
100 routes							
30 schools	8%	42%	34%	30%	−5%	0%	0%
150 routes							
40 schools	4%	35%	24%	28%	−6%	−4%	4%
200 routes							
50 schools	8%	31%	29%	33%	−12%	−5%	3%
250 routes							
60 schools	4%	35%	30%	30%	−8%	−3%	3%
300 routes							
70 schools	6%	35%	30%	27%	−9%	−3%	0%
350 routes							
80 schools	5%	39%	33%	31%	−8%	−4%	1%
400 routes							
90 schools	5%	35%	30%	30%	−10%	−4%	0%
450 routes							
100 schools	3%	33%	32%	29%	−11%	−2%	2%
500 routes							
Average gap	8%	37%	32%	30%	−7%	−3%	1%

be infeasible in practice as routes can be scheduled in quick succession (i.e., if the transition time is simply the travel distance from the end of one route to the start of the other, the time between routes is insufficient). Thus, adding a buffer time is of practical importance. As the buffer time increases, the location-dependent transition time solutions come closer to the constant transition time solutions. We observe this in the case study solutions which used a buffer time of three minutes. Thus, if well-chosen, the constant transition time, adopted often in practice, can be a good approximation for location-dependent transition times.

7. Discussion

In this paper, we develop a time-indexed ILP formulation for the SBSP and use its LP relaxation to develop a randomized rounding algorithm that yields near-optimal solutions for large-scale instances in polynomial time. Our computational study suggests that the rounding algorithm is fast and effective to solve problems where the transition time between routes is a constant. Our approach is generalized to solve instances with non-constant transition times (by reducing them to tractable instances) as well as to obtain robust solutions for multiple scenarios. From a strategic planning perspective, our approach provides an efficient way to identify scheduling plans with sufficient potential for cost savings (and remove from consideration those without sufficient savings) before more detailed operational analysis.

In this work, we assume that bus routes are fixed in the scheduling problem. Future work can expand our analysis to incorporate route changes to increase the benefit of such models, building on Bertsimas, Delarue, and Martin (2019), which demonstrates the advantages of jointly solving route design problem and school bus scheduling problem, and Zeng, Chopra, and Smilowitz (2019), which proposes efficient methods to approximate bus route travel time. Another important extension is to integrate our results with other strategic decisions, such as school boundary design and school location selection. Our approach can serve as a toolbox to help decision makers understand the effect on school transportation when making policy changes. Currently we are exploring these generalizations with partner school districts to provide solution approaches that are efficient, robust and easy to implement.

Acknowledgments

The authors thank school district partners for their productive collaborations.

References

Banerjee D, Smilowitz K (2019) Incorporating equity into the school bus scheduling problem. *Transp. Res., Part E Logist. Trans. Rev.* 131:228–246.

- Banerjee S, Hssaine C, Périvier N, Samaranayake S (2021) Real-time approximate routing for smart transit systems. Preprint, submitted March 10, <https://arxiv.org/abs/2103.06212>.
- Bertsimas D, Delarue A, Martin S (2019) Optimizing schools' start time and bus routes. *Proc. Natl. Acad. Sci. USA* 116(13):5943–5948.
- Bidwell A (2015) Report: Federal education funding plummeting. www.usnews.com/news/blogs/data-mine/2015/06/24/report-federal-education-funding-cut-by-5-times-more-than-all-spending.
- Campbell JF, North JW, Ellegood WA (2015) *Modeling Mixed Load School Bus Routing. Quantitative Approaches in Logistics and Supply Chain Management* (Springer), 3–30.
- Carraresi P, Gallo G (1984) Network models for vehicle and crew scheduling. *Eur. J. Oper. Res.* 16(2):139–151.
- Chen X, Kong Y, Dang L, Hou Y, Ye X (2015) Exact and metaheuristic approaches for a bi-objective school bus scheduling problem. *PLoS One* 10(7).
- Chernoff H (1952) A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Ann. Math. Statist.* 23:493–507.
- De La Vega WF, Lueker GS (1981) Bin packing can be solved within $1 + \epsilon$ in linear time. *Combinatorica* 1(4):349–355.
- Desrosiers J, de Montréal Centre de recherche sur les transports U, de Montréal Département d'informatique et de recherche opérationnelle U (1980) *An Overview of School Busing System* (Montréal: Université de Montréal, Centre de recherche sur les transports).
- Desrosiers J, Ferland JA, Rousseau JM, Lapalme G, Chapleau L (1986) Transcol: A multi-period school bus routing and scheduling system. *Management Sci.* 22:47–71.
- Dyer ME, Wolsey LA (1990) Formulating the single machine sequencing problem with release dates as a mixed integer program. *Discrete Appl. Math.* 26(2–3):255–270.
- Ellegood WA, Campbell JF, North J (2015) Continuous approximation models for mixed load school bus routing. *Transportation Res. Part B: Methodological* 77:182–198.
- Fügenschuh A (2009) Solving a school bus scheduling problem with integer programming. *European J. Oper. Res.* 193(3):867–884.
- Garey MR, Johnson DS (2002) *Computers and Intractability*, vol. 29 (W. H. Freeman, New York).
- Gavish B, Schweitzer P, Shlifer E (1978) Assigning buses to schedules in a metropolitan area. *Comput. Oper. Res.* 5(2):129–138.
- Ikura Y, Gimple M (1986) Efficient scheduling algorithms for a single batch processing machine. *Oper. Res. Lett.* 5(2):61–65.
- Keller I, Tompkins C (1956) An extension of a theorem of Dantzig's. linear inequalities and related systems. *Ann. Math. Stud.* 38:247–254.
- Kim BI, Kim S, Park J (2012) A school bus scheduling problem. *Eur. J. Oper. Res.* 218(2):577–585.
- Köksal Ahmed E, Li Z, Veeravalli B, Ren S (2020) Reinforcement learning-enabled genetic algorithm for school bus scheduling. *J. Intelligent Transportation Syst.*, <https://www.sciencedirect.com/org/science/article/abs/pii/S1547245022003619>.
- Lenstra JK, Shmoys DB, Tardos E (1990) Approximation algorithms for scheduling unrelated parallel machines. *Math. Program.* 46(1–3):259–271.
- Newton RM, Thomas WH (1969) Design of school bus routes by computer. *Socio-Econom. Planning Sci.* 3(1):75–85.
- Olariu S (1991) An optimal greedy heuristic to color interval graphs. *Inform. Process. Lett.* 37(1):21–25.
- Orloff CS (1976) Route constrained fleet scheduling. *Transportation Sci.* 10(2):149–168.
- Park J, Kim BI (2010) The school bus routing problem: A review. *Eur. J. Oper. Res.* 202(2):311–319.
- Park J, Tae H, Kim BI (2012) A post-improvement procedure for the mixed load school bus routing problem. *Eur. J. Oper. Res.* 217(1):204–213.
- Queyranne M, Schulz AS (1994) *Polyhedral Approaches to Machine Scheduling* (TU, Fachbereich 3, Berlin).

- Scholl A, Klein R, Jürgens C (1997) Bison: A fast hybrid procedure for exactly solving the one-dimensional bin packing problem. *Comput. Oper. Res.* 24(7):627–645.
- Schulz AS (1996) Scheduling to minimize total weighted completion time: Performance guarantees of LP-based heuristics and lower bounds. *Internat. Conf. Integer Programming Combinatorial Optimization*, 301–315 (Springer).
- Shmoys DB, Tardos É (1993) An approximation algorithm for the generalized assignment problem. *Math. Programming* 62(1–3): 461–474.
- Sousa JP, Wolsey LA (1992) A time indexed formulation of non-preemptive single machine scheduling problems. *Math. Programming* 54(1):353–367.
- Spada M, Bierlaire M, Liebling TM (2005) Decision-aiding methodology for the school bus routing and scheduling problem. *Transportation Sci.* 39(4):477–490.
- Swersey AJ, Ballard W (1984) Scheduling school buses. *Management Sci.* 30(7):844–853.
- Wang F, Xu Y (2011) Estimating O–D travel time matrix by Google Maps API: implementation, advantages, and implications. *Ann. GIS.* 17(4):199–209.
- Wenzel C (2016) *Optimale schulanfangszeiten zur entlastung des nahverkehrs in der stadt nürnberg*. *Angewandte Mathematik und Optimierung Schriftenreihe (AMOS)*.
- Zeng L, Chopra S, Smilowitz K (2019) The covering path problem on a grid. *Transportation Sci.* 53(6):1656–1672.