Optimizing schools' start time and bus routes

Dimitris Bertsimas[a,1], Arthur Delarue[a,1], and Sebastien Martin[a,1,2]

[a]Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139

Edited by Alan Murray, University of California, Santa Barbara, CA, and accepted by Editorial Board Member Susan Hanson February 11, 2019 (received for review July 3, 2018)

Maintaining a fleet of buses to transport students to school is a major expense for school districts. To reduce costs by reusing buses between schools, many districts spread start times across the morning. However, assigning each school a time involves estimating the impact on transportation costs and reconciling additional competing objectives. Facing this intricate optimization problem, school districts must resort to ad hoc approaches, which can be expensive, inequitable, and even detrimental to student health. For example, there is medical evidence that early high school starts are impacting the development of an entire generation of students and constitute a major public health crisis. We present an optimization model for the school time selection problem (STSP), which relies on a school bus routing algorithm that we call biobjective routing decomposition (BiRD). BiRD leverages a natural decomposition of the routing problem, computing and combining subproblem solutions via mixed integer optimization. It significantly outperforms state-of-the-art routing methods, and its implementation in Boston has led to $5 million in yearly savings, maintaining service quality for students despite a 50-bus fleet reduction. Using BiRD, we construct a tractable proxy to transportation costs, allowing the formulation of the STSP as a multiobjective generalized quadratic assignment problem. Local search methods provide high-quality solutions, allowing school districts to explore tradeoffs between competing priorities and choose times that best fulfill community needs. In December 2017, the development of this method led the Boston School Committee to unanimously approve the first school start time reform in 30 years.

optimization | education | transportation | public policy | fairness

In the 21st century, school districts across the United States face a wide array of challenging problems on a daily basis from adjusting to the digital age to educating an increasingly diverse and multicultural student body. Yet, perhaps the most complicated decision that administrators face is seemingly the most innocuous: determining what time each school in the district should start in the morning and end in the afternoon.

The issue of choosing appropriate school "bell times" has received increased attention in recent years, as too-early start times have been linked to a wide array of health issues among teenagers, including diminished academic achievement (1) and cognitive ability (2, 3) and increased rates of obesity (4), depression (5), and traffic accidents (6). Indeed, changes in the body's circadian clock during puberty effectively prevent adolescents from getting adequate sleep early in the night (7). While the American Academy of Pediatrics recommends that teenagers not start their school day before 8:30 AM, a recent CDC report found that just 17.7% of US high schools comply (8). Some experts estimate that, over the next 10 y, the dire public health implications of early high school start times could impact the US economy by over $80 billion (9).

Moreover, research suggests that these repercussions disproportionately affect the most economically disadvantaged students (10). As achievement gaps between students from different backgrounds remain stark (11), research has consistently found systematic biases, largely on racial lines (12), that partially explain these gaps. For example, school bell times can suffer from such biases, such as is the case in Boston (13).

For decades, school districts across America have considered ways to adjust their bell times and solve these issues in a fair way. However, the sheer complexity of the problem is a major obstacle to change. School districts typically struggle with balancing many competing objectives, including student health, special education programs, parent and staff schedules, state and federal regulations, and public externalities (14).

Perhaps the greatest obstacle to adjusting school bell times is the effect of changes on school transportation. Over 50% of US schoolchildren rely on an army of half a million yellow school buses to travel to and from school every day. In Boston, where specialized programs draw students from all over the city and traffic is often at a standstill, transportation accounts for over 10% of the district's $1 billion budget. To reduce transportation spending, school districts, such as Boston, stagger the start and end times of different schools, allowing vehicles to be reused several times throughout the day. Because many school districts construct bus routes by hand, it is exceedingly difficult for them to evaluate the impact of bell time changes on bus costs, let alone find a set of bell times that satisfies all of the district's objectives without inflating the budget. No matter how unpalatable, the status quo is often the only viable option. In addition, because of the impossibility of systemwide change, districts may experiment with a piecemeal approach to bell time change, where the most vocal and best-connected schools may benefit the most.

Significance

Spreading start times allows school districts to reduce transportation costs by reusing buses between schools. However, assigning each school a time involves both estimating the impact on transportation costs and reconciling additional competing objectives. These challenges force many school districts to make myopic decisions, leading to an expensive and inequitable status quo. For instance, most American teenagers start school before 8:00 AM, despite evidence of significant associated health issues. We propose an algorithm to jointly solve the school bus routing and bell time selection problems. Our application in Boston led to $5 million in yearly savings (maintaining service quality despite a 50-bus fleet reduction) and to the unanimous approval of the first school start time reform in 30 years.

Author contributions: D.B., A.D., and S.M. designed research; A.D. and S.M. performed research; D.B., A.D., and S.M. contributed new analytic tools; A.D. and S.M. analyzed data; and A.D. and S.M. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. A.M. is a guest editor invited by the Editorial Board.

[1]D.B., A.D., and S.M. contributed equally to this work.

[2]To whom correspondence should be addressed. Email: semartin@mit.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1811462116/-/DCSupplemental.

Published online March 12, 2019.

The problem of school bus routing has been addressed extensively (15, 16). It is typically decomposed into three main subproblems (Fig. 1 *D–F*): stop assignment (i.e., choosing locations where students will walk from their homes to get picked up); bus routing (i.e., linking stops together into bus trips); and bus scheduling (i.e., combining bus trips into a route that can be served by a single bus). State-of-the-art optimization algorithms exist for these subproblems in isolation (17, 18). However, the literature on optimally combining subproblem solutions is less extensive. Approaches typically involve formulating the school bus routing problem as a large combinatorial optimization problem, which can be solved using metaheuristics, including local search (19), simulated annealing (20), and special purpose vehicle routing heuristics (21, 22). Special purpose algorithms have also been designed to address variants of the school bus routing problem, allowing "mixed loads"—students from different schools riding the bus together (19, 22, 23), bus transfers (24), or arrival time windows (18–20, 23).

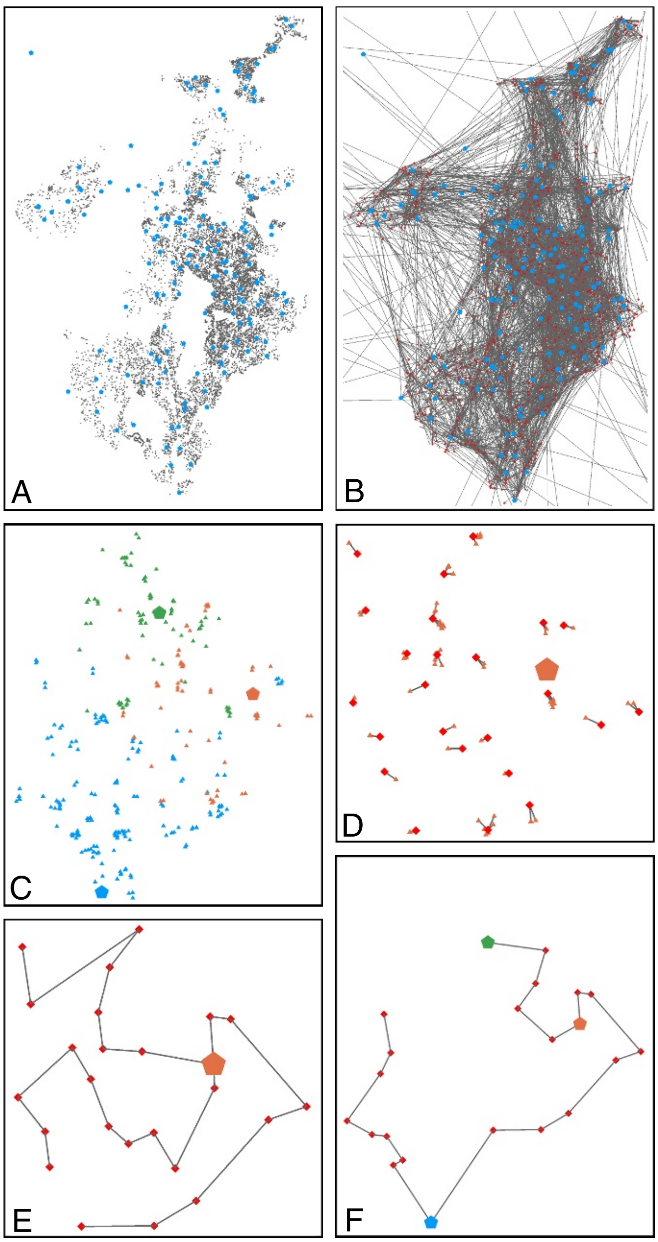

Fig. 1. Geographic visualization of the school bus routing problem (and subproblems). (*A*) BPS 2017–2018 data (anonymized), with gray triangles representing students and blue pentagons representing schools. (*B*) Sample BPS routing solution, with schools as blue pentagons and bus stops as red squares; lines connect bus stops that are served in sequence by the same bus, illustrating the complexity of Boston school transportation. (*C*) Small synthetic district (three schools); students (triangles) are the same color as their assigned schools (pentagons). (*D–F*) Examples of the three main routing steps in this district: stop assignment (*D*), where students (triangles) attending the orange school (pentagon) are shown connected to their assigned stops (red squares); one-school routing (*E*), where all bus stops for the orange school are connected into bus trips; and bus scheduling between multiple schools (*F*), where three trips (one from each school) are connected into a single bus itinerary.

Unfortunately, many tractable general purpose algorithms do not consider additional constraints (fleet heterogeneity and student-specific needs) and thus, lack portability. Although an optimization framework to the school time selection problem (STSP) has been proposed (25), no existing algorithms address bell time selection in conjunction with bus routing (18).

This work presents a model for the STSP, allowing the joint optimization of school bell times with school bus routes. We first develop a school bus routing algorithm called biobjective routing decomposition (BiRD), which bridges the gap between standard subproblems to find better solutions. We then propose a mathematical formulation of the STSP, a multiobjective approach that can model any number of community objectives as well as transportation costs using BiRD.

BiRD outperforms state-of-the-art methods by 4–12% on average on benchmark datasets, and it allowed Boston Public Schools (BPS) to take 50 buses off the road and save almost \$5 million in the fall of 2017 without increasing the average student's walking or riding times. Our modeling approach to the STSP along with the successful implementation of BiRD led the Boston School Committee to reconsider start time policies for the first time since 1990, unanimously approving a comprehensive reform prioritizing student health in December 2017. Our STSP model was used by BPS to evaluate the impact of many different scenarios and ultimately, propose bell times for all 125 BPS schools. These start times have not been implemented yet due to parents concern with the magnitude of the change, but our quantitative approach to evaluating policy tradeoffs has informed the conversation about school start times both in Boston and across the nation.

School Transportation: A BiRD's Eye View

Solving the school bus routing problem means assigning students to stops near their homes and selecting which bus will pick them up and in what order (keeping in mind that a bus only carries students for one school but can serve several schools in succession thanks to staggered bell times) in a way that minimizes the overall number of buses or another objective of interest. We show examples of a school district (BPS) in Fig. 1*A* and of a model school district that mimics the real setting in Fig. 1*C* and *SI Appendix*, Fig. S2.

The BiRD algorithm consists of several steps (Fig. 2), for which we develop optimization-based approaches implemented with modern software tools (26, 27) and tools available online (28). For clarity, we focus on the morning problem, but our algorithm generalizes to the afternoon (*SI Appendix*). Because problem details often vary between districts, it may be advantageous to adjust some steps to changes in the problem setting. BiRD's defining feature is thus the decomposition of the problem and in particular, the scenario selection step, which bridges the gap between the single-school and multischool subproblems.

Single-School Problem. To assign students to stops (Fig. 1*D*), we use an integer optimization formulation of the assignment problem with maximum walking distance constraints. We minimize

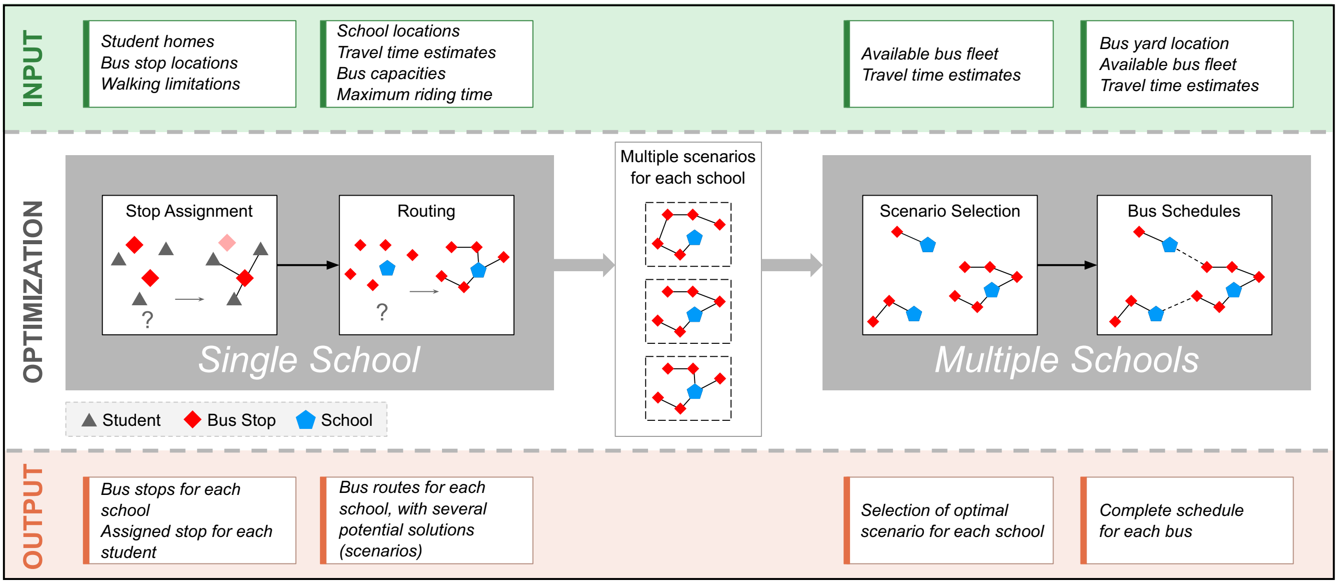

Fig. 2. Overview of BiRD algorithm. On the left, the single-school problem can be divided into the two subproblems of stop assignment and single-school routing; on the right, the multischool problem can be divided into the two subproblems of scenario selection and bus scheduling. The generation of not one but several routing scenarios for each school and the subsequent joint selection of a single scenario for each school bridge the divide between the single-school and multischool problems.

the overall number of stops, because (*i*) it simplifies bus trips and (*ii*) the minimum pickup time at a stop is typically high, even if the stop has few students. When long bus routes span the entire city, as in Boston (Fig. 1*B*), stop assignment has a negligible effect on the macroscopic quality of the routing solution. Our formulation can include additional objectives, such as the total student walking distance, and can exclude stop assignments that require students to cross major arteries or unsafe areas (*SI Appendix*).

We then use an insertion-based algorithm to connect sequences of stops into feasible bus trips (Fig. 1*E*). We use integer optimization to combine these feasible trips with a minimum number of buses, with a set cover formulation reminiscent of crew scheduling problems (29) (*SI Appendix*). Our method has the flexibility to handle practical modifications in the routing problem from vehicles with different capacities to student–bus compatibility restrictions (e.g., students in a wheelchair need a bus with a special ramp/lift). In principle, the modularity of the overall algorithm means that the single-school routing algorithm can be replaced with any state-of-the-art vehicle routing method.

Routing Multiple Schools. We use the single-school routing method to generate not one but several varied optimized routing scenarios for each school to select the best one for the system. In particular, we consider several scenarios on the Pareto frontier of two objectives (hence, the name of BiRD): number of buses and average riding time. This tradeoff makes sense, because shorter routes are more easily connected into bus schedules.

Then, we jointly select one scenario for each school in a way that favors maximal reuse of buses from school to school (Fig. 2) by formulating an integer optimization problem with network flow structure that seeks to minimize the number of buses at the scale of the entire district (*SI Appendix*). Given one routing scenario for each school, we can then solve another integer optimization problem to identify a trip-by-trip itinerary for each bus in the fleet (Fig. 1*F*). In this final subproblem, we optimize the number of buses jointly in the morning and in the afternoon (*SI Appendix*).

Evaluating the Routing Algorithm. We compare BiRD's ability to minimize the total number of buses with existing methods (20, 22) on 32 published benchmarks (23) and on 20 of our own synthetically generated examples. We outperform all other methods on all but one instance, with an average improvement of 4% on the instances from ref. 23 and 12% on our instances. The scenario selection step is key to this improvement: computational experiments (*SI Appendix*) indicate that BiRD's performance improves by 20% when we compute two different routing scenarios for each school and select the best one by considering the whole system as opposed to using the best scenario for each school. Intuitively, what is optimal for one school may not be optimal for the entire system, motivating the biobjective decomposition approach.

Application in Boston. BPS has the highest transportation expenditure per student in the United States, with rising costs due in part to narrow streets and infamous rush hour traffic, a large fraction of special education students, and a complicated history of school desegregation. In addition, over the last decade, BPS has adopted a "controlled choice" approach to school selection, which gives parents greater latitude in selecting a public school while promoting fairness across the district (30, 31). As a result of this policy, some schools may draw students from far across the city, further complicating the school transportation problem and driving up costs.

Before we started working with BPS, bus routes for 125 public schools and over 80 private and charter schools were computed and maintained manually. BiRD's ability to incorporate district-specific constraints (including four different bus types and only one compatible with wheelchairs) was essential in producing a practical solution. In the end, we solved the Boston school bus routing problem using only 530 buses vs. 650 for the manual solution. This represents an 18% reduction, with estimated cost savings in the range of \$10–15 million. To ensure a smooth transition, BPS decided to only take 50 buses off the road in the first year of implementation, still amounting to a hefty \$5 million in cost savings (32). Despite the smaller number of buses, the average student ride time stayed constant from 2016 to 2017 (around 23 min).

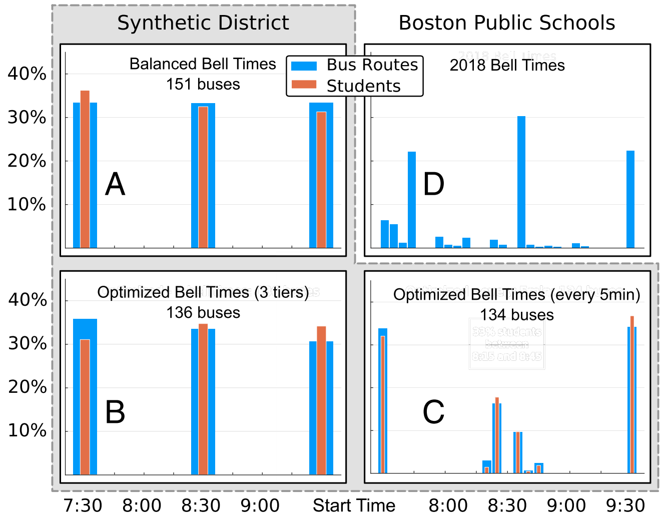

Fig. 3. Bell time optimization. Comparison of three bell time optimization strategies on a synthetic district. When only three bell times are allowed, balancing the number of bus routes across bell times (*A*) works well but is typically beaten by routing compatibility optimization (*B*). Even better solutions can be obtained by allowing more bell times in the middle tier (*C*). In comparison, BPS bell times are not even balanced (*D*).

Formulating the STSP

Selecting bell times is a complex policy problem with many stakeholders. We first focus on the interplay with transportation, since computing school bus routes is a necessary component of bell time selection. For instance, it is of interest to evaluate transportation costs when each school S is assigned a particular bell time t_S. However, there are too many possibilities to explore in practice (exponential in the number of schools). Instead, we develop a general formulation for the STSP, which contains a tractable proxy for transportation cost constructed using BiRD. We show how to include other community objectives in the next section.

Transportation Costs. A key factor in an optimized school bus routing solution is the "compatibility" of pairs of trips (i.e., how easy it is for a single bus to serve them with minimum idle time in between). We define a trip compatibility cost that trades off (*i*) the feasibility of a bus serving the two trips sequentially and (*ii*) the amount of idle or empty driving time involved, with tradeoff parameters that depend on characteristics of the school district and can be found using cross-validation. Then, for any pair of schools S and S', we can define a routing pairwise affinity cost $c^{\text{routing}}_{S,t,S',t'}$ that is the sum of the compatibility costs between every trip in every routing scenario for S at time t and S' at t' (*SI Appendix*).

Optimizing. Because its objective function only includes pairwise affinity costs, our model of the STSP is a special case of the generalized quadratic assignment problem (GQAP) (33). When different GQAP formulations for the STSP were investigated in ref. 25, even small instances could be intractable. We, therefore, develop a simple local improvement heuristic that works well in practice. Given initial bell times, we select a random subset of schools. The problem of finding the optimal start times for this subset while fixing all other schools' start times is also a GQAP.

We can then solve this restricted GQAP problem using mixed integer optimization to obtain a new set of bell times in seconds for small-enough subsets. We repeat the operation with new random subsets until convergence. Results on synthetic data suggest that a subset size of one gives near-optimal results if the local improvement heuristic is run several times with random starting points. We note that the heuristic is interpretable: with a subset size of n, a solution obtained after convergence can only be improved by changing the bell times of at least $n+1$ schools.

Evaluating Three-Tier Systems. In many districts, such as Boston (Fig. 3*D*), start times are separated into three equally spaced "tiers" (e.g., 7:30, 8:30, and 9:30 AM). Such a system allows each bus to serve up to three schools every morning (34), and therefore, districts will typically try to balance the number of bus trips across all three tiers. Our method allows us to quantify the empirical behavior of this intuitive idea.

Simulations suggest that optimizing three-tier bell times using our algorithm (Fig. 3*A*) yields an 11% cost improvement over simply balancing the number of bus routes across tiers (Fig. 3*B*), which is already better than what school districts typically do (Fig. 3*D*). Distributing schools across tiers without accounting for geography/routing compatibility is suboptimal.

Furthermore, a three-tier system is not necessarily the right answer per se. Fig. 3*C* shows that allowing many possible start times for the middle tier (5-min intervals between 8:00 and 9:00 AM) can yield a 1–2% improvement over the standard three-tier optimized solution (Fig. 3*B*). Interestingly, no school starts at 8:30 AM in this system. Although tiered bell times are popular because of their simplicity, algorithmic tools, such as our tool, suggest that better solutions exist. For instance, in Boston, we can find a bell time solution that requires just 450 buses, which represents a 15% improvement over the number of buses obtained without changing the bell times and a 31% improvement over the number of buses used by BPS in the 2016–2017 school year.

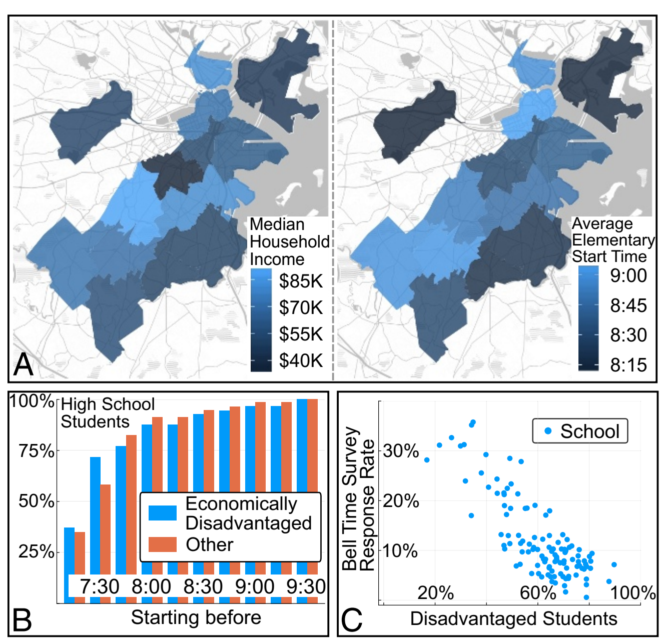

Fig. 4. Equity and current start times in Boston. (*A*) Maps of Boston with neighborhoods shaded by median household income (American Community Survey) and average elementary start time. Elementary students start later in wealthier neighborhoods (0.78 correlation between household income and start time). (*B*) Proportion of high school students starting before each time in the morning (comparing economically disadvantaged students with other students). Start times skew early for economically disadvantaged high school students (χ^2 homogeneity *P* value $< 10^{-5}$). (*C*) BPS Community Survey response rate by school shown against the fraction of disadvantaged students attending the school. Economically fragile populations have a lower bell time survey response rate.

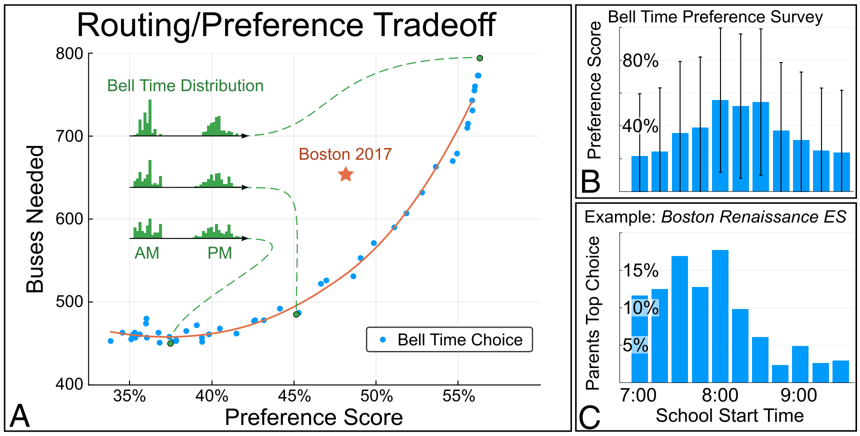

Fig. 5. Optimizing preferences is hard. (*A*) Tradeoff curve derived by our algorithm between preference score (metric of community satisfaction) (*SI Appendix*) and transportation cost along with three sets of bell times along the curve. Even a slight improvement in satisfaction comes at a high cost. (*B*) District-wide preference score of each bell time, showing that parents typically prefer from 8:00 to 8:30 AM start times, with high variance. (*C*) Distribution of parents' top bell time choice at a particular school (Boston Renaissance). Even within a single school, agreement is hard to come by: although the school's current start time is 7:30 AM, only 17% of parents list this time as their favorite.

Bell Times in Practice

In a real district, bell time selection goes far beyond minimizing the number of buses, which we found in our work with BPS. For context, Boston's existing bell time policy, enacted in 1990, split the public schools into three tiers with start times of 7:30, 8:30, and 9:30 AM, stipulating that tiers would rotate through the start times every 5 y. Unfortunately, this policy was never enforced, and the bell times assigned in 1990 mostly remain today.

These bell times are flawed. First, because they have remained static while school demographics have evolved, they have contributed to the steady rise of the BPS transportation budget over the last decade. Second, over 74% of high school students currently start school before 8:00 AM. Many studies have shown that the negative effects of early high school starts are magnified in economically fragile students (10). However, in Boston, such students have worse bell times, on average, than economically advantaged students (13). In Fig. 4, we see, for example, that economically disadvantaged high school students are more likely to start before 7:30 AM than other high school students.

Gridlock. The Boston status quo has persisted for decades despite its shortcomings. Indeed, bell time selection is intrinsically difficult, because stakeholders cannot agree on what is best for everyone. Fig. 5 *B* and *C* shows community preferences for different start times across all public schools obtained through a BPS survey. Although families and school staff tend to favor start times between 8:00 and 8:30 AM, the displayed preferences are mostly characterized by broad disagreement, even within a single school (Fig. 5*C*). Any bell time for any school is sure to have both fervent supporters and vehement critics.

School districts have no hope of satisfying all or even most of their constituents. Moreover, the cost of even trying to satisfy the individual preferences of parents and staff can be prohibitive: Fig. 5*A* shows that each additional point of community satisfaction in Boston can cost dozens of additional buses and tens of millions of taxpayer dollars.

For BPS, the tradeoff curve in Fig. 5*A* represented a paradigm shift, the first time that the district could visualize or even quantify any of the tradeoffs of bell time policy making. The curve illustrates our model's first use: providing a district the quantitative support necessary to understand the problem and make the best decision.

The Greater Good. Although stakeholders have many competing personal priorities, they often agree on broader goals, such as having fair and equitable bell times or reinvesting saved transportation costs into schools. Starting in 2016, BPS led an engagement process aimed at understanding broad community values. The results suggested four main objectives: to maximize how many high school students start after 8:00 AM, minimize how many elementary school students end after 4:00 PM, prioritize schools with high special education needs, and reinvest transportation savings into classrooms while achieving these objectives in an equitable manner.

In the general case, solving the STSP in practice means optimizing a set of several objectives, such as the ones outlined above. We call an objective GQAP representable if it can be represented using only single affinity costs $c_{S,t}$ (representing the aversion of school S for bell time t) and pairwise affinity costs $c_{S,t,S',t'}$. We find that the GQAP framework has sufficient modeling power to represent all of the objectives and constraints that interest school districts in general (*SI Appendix*) and Boston in particular.

Typically, school districts will wish to balance multiple GQAP-representable objectives, including transportation costs. As is usual in multiobjective optimization, we consider that the final cost function to optimize is a weighted average of the district's different (GQAP-representable) objectives, with weights indicating policy makers' priorities.

	Buses	Early HS	Late ES	Survey score	Bell time distribution
Current	650	74%	33%	48%	
NewRoutes	530	74%	33%	48%	
LowCost	450	43%	27%	37%	
MaxSurvey	934	0%	8%	56%	
Optimal	481	6%	15%	40%	

Fig. 6. Bell time selection tradeoffs. Sample of a few scenarios considered by BPS. Current start times (with or without new routes) have many high school students starting before 8:00 AM (Early HS) and elementary school students ending after 4:00 PM (Late ES), mediocre community satisfaction (survey score), and a suboptimal bell time distribution both in the morning and in the afternoon (histogram weighted by students: blue, morning; orange, afternoon). The three other scenarios present different tradeoffs between the bell time objectives—BPS chose the "Optimal" scenario.

COMPUTER SCIENCES

SOCIAL SCIENCES

We explored tens of thousands of tradeoffs for BPS, such as those presented in Fig. 6. We notice that, in Boston, reducing both the number of high school students starting too early and the number of elementary school students ending too late can be done at little to no cost.

Application in Boston. In December 2017, the Boston School Committee unanimously approved a new policy (35) stipulating that all future bell time solutions should optimize the verifiable criteria described above, paving the way for algorithmic bell time selection. Our flexible methodology allowed us to take into account a number of very specific constraints (e.g., preventing large neighboring high schools from dismissing at the same time, which could create unsafe situations at neighboring subway stations). In the end, the proposed bell times (Fig. 6) reduced the number of high school students starting before 8:00 AM from 74 to 6% and the number of elementary school students dismissing after 4:00 PM from 33 to 15%. The plan also led to an estimated reinvestment of up to $18 million into classrooms. Because of the significant amount of change under this new plan and in response to protests by some families, BPS delayed the plan's implementation to allow more time to adjust the objective weights and constraints. As BPS continues to gather community input, the legitimate concerns raised by these families can be modeled as objectives within our general formulation and integrated within our framework.

Ultimately, using an algorithm for bell time selection at the scale of a city allows leaders to thoroughly evaluate their options and empowers them to make decisions based not on the political whims of special interest groups but on an objective standard agreed on by the community.

ACKNOWLEDGMENTS. We express our admiration and respect for John Hanlon and Will Eger, whose single-minded devotion to public service should be an example to the world. We thank our partners at BPS for their tireless efforts on behalf of the children of Boston. Google Maps helped make this work possible by providing point-to-point travel time estimates for all locations in Boston. The Julia language and its JuMP extension for optimization significantly eased the implementation of our algorithms. We thank the editor and reviewers for many helpful suggestions.

1. Carrell SE, Maghakian T, West JE (2011) A's from ZZZZ's? The causal effect of school start time on academic achievement of adolescents. *Am Econ J* 3:62–71.
2. Curcio G, Ferrara M, De Gennaro L (2006) Sleep loss, learning capacity and academic performance. *Sleep Med Rev* 10:323–337.
3. Owens JA, Belon K, Moss P (2010) Impact of delaying school start time on adolescent sleep, mood, and behavior. *Arch Pediatr Adolesc Med* 164:608–614.
4. Chen X, Beydoun MA, Wang Y (2008) Is sleep duration associated with childhood obesity? A systematic review and meta-analysis. *Obesity* 16:265–274.
5. Fredriksen K, Rhodes J, Reddy R, Way N (2004) Sleepless in Chicago: Tracking the effects of adolescent sleep loss during the middle school years. *Child Dev* 75: 84–95.
6. Danner F, Phillips B (2008) Adolescent sleep, school start times, and teen motor vehicle crashes. *J Clin Sleep Med* 4:533–535.
7. Crowley SJ, Acebo C, Carskadon MA (2007) Sleep, circadian rhythms, and delayed phase in adolescence. *Sleep Med* 8:602–612.
8. Wheaton AG, Ferro GA, Croft JB (2015) School start times for middle school and high school students—United States, 2011–12 school year. *Morb Mortal Wkly Rep* 64: 809–813.
9. Hafner M, Stepanek M, Troxel WM (2017) Later school start times in the US.: An economic analysis (RAND Corporation, Santa Monica, CA), Technical Report RR-2109-RE.
10. Edwards F (2012) Early to rise? The effect of daily start times on academic performance. *Econ Educ Rev* 31:970–983.
11. Wiggan G (2007) Race, school achievement, and educational inequality: Toward a student-based inquiry perspective. *Rev Educ Res* 77:310–333.
12. Cohen GL, Garcia J, Apfel N, Master A (2006) Reducing the racial achievement gap: A social-psychological intervention. *Science* 313:1307–1310.
13. Shuster R (2018) School bell times equity impact, 2017. Available at https://www.bostonpublicschools.org/cms/lib/MA01906464/Centricity/Domain/2389/Equity%20Analysis.pdf. Accessed February 27, 2019.
14. Malone SK, Ziporyn T, Buttenheim AM (2017) Applying behavioral insights to delay school start times. *Sleep Health J Natl Sleep Found* 3:483–485.
15. Desrosiers J, Dumas Y, Solomon MM, Soumis F (1995) Time constrained routing and scheduling. *Handbooks Oper Res Manag Sci* 8:35–139.
16. Park J, Kim B-I (2010) The school bus routing problem: A review. *Eur J Oper Res* 202:311–319.
17. Schittekat P, Sevaux M, Sorensen K (2006) A mathematical formulation for a school bus routing problem. *Proceedings of the International Conference on Service Systems and Service Management* (IEEE, Piscataway, NJ), Vol 2, pp 1552–1557.
18. Fügenschuh A (2009) Solving a school bus scheduling problem with integer programming. *Eur J Oper Res* 193:867–884.
19. Spada M, Bierlaire M, Liebling TM (2005) Decision-aiding methodology for the school bus routing and scheduling problem. *Transportation Sci* 39:477–490.
20. 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:e0132600.
21. Bramel J, Simchi-Levi D (1995) A location based heuristic for general routing problems. *Oper Res* 43:649–660.
22. Braca J, Bramel J, Posner B, Simchi-Levi D (1997) A computerized approach to the New York cityschool bus routing problem. *IIE Trans* 29:693–702.
23. Park J, Tae H, Kim B-I (2012) A post-improvement procedure for the mixed load school bus routing problem. *Eur J Oper Res* 217:204–213.
24. Bögl M, Doerner KF, Parragh SN (2015) The school bus routing and scheduling problem with transfers. *Networks* 65:180–203.
25. Wenzel C (2016) Optimale Schulanfangszeiten zur Entlastung des Nah verkehrs in der Stadt Nuernberg [Optimal times for the beginning of schools to reduce the pressure on public transport in the City of Nuremberg] (Universität der Bundeswehr München, Neubiberg Germany). No. AMOS 42.
26. Bezanson J, Edelman A, Karpinski S, Shah VB (2017) Julia: A fresh approach to numerical computing. *SIAM Rev* 59:65–98.
27. Dunning I, Huchette J, Lubin M (2017) Jump: A modeling language for mathematical optimization. *SIAM Rev* 59:295–320.
28. Arthur D, Martin S (2018) SchoolBusRouting repository. Available at https://github.com/mitschoolbus/SchoolBusRouting. Accessed February 27, 2019.
29. Smith BM, Wren A (1988) A bus crew scheduling system using a set covering formulation. *Transportation Res A Gen* 22:97–108.
30. Abdulkadiroğlu A, Pathak PA, Roth AE, Sönmez T (2005) The Boston public school match. *Am Econ Rev* 95:368–371.
31. Pathak PA, Shi P (2013) Simulating alternative school choice options in Boston–Technical appendix (MIT School Effectiveness and Inequality Initiative, Cambridge, MA).
32. O'Brien D (2017) MIT's "Quantum Team" wins first-ever BPS transportation challenge with revolutionary computer model. Available at https://www.bostonpublicschools.org/site/default.aspx?PageType=3&DomainID=4&ModuleInstanceID=14&ViewID=6446EE88-D30C-497E-9316-3F8874B3E108&FlexDataID=12431. Accessed February 27, 2019.
33. Hahn PM, Kim B-J, Guignard M, Smith JMG, Zhu Y-R (2008) An algorithm for the generalized quadratic assignment problem. *Comput Optim Appl* 40:351–372.
34. Owens J, Drobnich D, Baylor A, Lewin D (2014) School start time change: An in-depth examination of school districts in the United States. *Mind Brain Educ* 8:182–213.
35. Sullivan EA (2017) Order on school start time realignment. Available at https://www.bostonpublicschools.org/cms/lib/MA01906464/Centricity/Domain/162/starttimesmotion.pdf. Accessed February 27, 2019.