A Stable and Student-Optimal School Choice Mechanism with Transportation

Andrew Poile
University of Southampton
Southampton, United Kingdom
A.Poile@soton.ac.uk

Bahar Rastegari
University of Southampton
Southampton, United Kingdom
B.Rastegari@soton.ac.uk

Sebastian Stein
University of Southampton
Southampton, United Kingdom
ss2@ecs.soton.ac.uk

Vahid Yazdanpanah
University of Southampton
Southampton, United Kingdom
V.Yazdanpanah@soton.ac.uk

ABSTRACT

School choice mechanisms have increased school segregation across municipalities worldwide, with urban geography and proximity-based admissions priorities restricting disadvantaged families' feasible choices. Quota-based approaches for reducing segregation face computational intractability and can produce Pareto inferior outcomes. Existing research suggests transportation significantly influences school selection decisions. We propose a stable, student-optimal matching mechanism that incorporates transportation directly into the allocation process, enabling eligible students from disadvantaged areas to access expanded choice sets. Our mechanism extends deferred acceptance by allowing route-eligible students to rank route-school tuples alongside individual schools, with these students receiving priority equivalent to local students at served schools. We prove that our mechanism terminates, produces stable matchings, and returns the student-optimal stable matching when the student-oriented version is executed. Preliminary experiments demonstrate lower dissimilarity indices compared to standard deferred acceptance, suggesting potential for reducing proximity-induced segregation while maintaining desirable matching properties.

KEYWORDS

Matching, School Choice, Mechanism Design

ACM Reference Format:
Andrew Poile, Bahar Rastegari, Sebastian Stein, and Vahid Yazdanpanah. 2026. A Stable and Student-Optimal School Choice Mechanism with Transportation. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 8 pages.

1 INTRODUCTION

Municipalities across the world have implemented school choice mechanisms to allocate students to schools with modest improvements to school quality [11, 13, 15, 17, 22], but a significant increase in school segregation [25], which is a compositional difference between schools whereby schools are stratified by factors such as socio-economic status of attending students. While the source of the problem is context dependant, urban geography, and in particular induced demand for housing around popular schools, is a key factor [14, 21].

Particularly in England, school admissions have strong geographic features. This is evident in the over-subscription criteria, which is how a school prioritises seats if it receives more applications than it has capacity. Either implicitly or explicitly, these criteria increase the priority of students nearest to the school and restrict the feasible choices of disadvantaged families [7, 8]. This can then result in socio-economic segregation of schools, a compositional difference between schools whereby schools are stratified by factors such as socio-economic status of attending students. Many such municipalities exist in England, municipalities with established transport services that students might already use to commute to school. However, to the best of our knowledge, no municipality is currently utilising transport in allocation mechanisms as part of the selection process, indeed the government offers free school transport only to those eligible students who attend the nearest suitable school [1].

Consider an example city with two schools, under a system of school choice, with over-subscription criteria that increase priority with proximity, finite capacities much smaller than the total population of students in the city, and a robust transport network which connects the schools to many students across the city. Suppose one school has a very positive reputation, and the students with the means, who understand how the allocation mechanism works, have relocated to an area close to the school to maximise their chances of being selected by the mechanism. Disadvantaged students are displaced and live further away from this school as a consequence. Disadvantaged students can ostensibly select either school in the mechanism, but are de facto restricted from being accepted into the good school due to their distance, even if they could easily access the school via some sort of bus or train route. Under this system, the schools would be segregated by the allocation mechanism.

Currently, many modern school choice implementations follow the seminal work of Abdulkadiroğlu and Sönmez [3], who analysed school choice with a matching theory lens. They proposed the application of two celebrated mechanisms to the school choice problem: deferred acceptance (presented by Gale and Shapley [12]) and top trading cycles (presented by Shapley and Scarf [23]). The authors

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus.

identify key desiderata: stability, strategy-proofness, and Pareto-efficiency. Both mechanisms are strategy-proof, and if deferred acceptance mechanisms are run *student-oriented* they produce the *student optimal*, stable matching - which is to say it is the Pareto optimal *stable* matching outcome for student welfare.

While preserving stability, strategy-proofness, and student-optimality, we contribute a novel algorithm that reduces school segregation and expands the choice sets of disadvantaged students by integrating transportation into the allocation directly, allowing disadvantaged students to select schools connected to them by transport routes. Transport routes have their own capacities independent of the schools' capacities to reflect the limits of the transport system in reality. Selecting schools in this way allows students to overcome the effects of geography to attend better schools.

2 RELATED LITERATURE

We focus especially on the inequities caused by the interaction of the mechanism with municipal geography, as mentioned above. Since housing and education infrastructure is relatively permanent, it seems natural to consider how transportation interacts with school choice. Where school and house locations are static, transportation routes, stops and frequencies are dynamic and can be changed to serve the needs of students. Below, we review the ways that segregation and transportation have been approached in the school choice literature.

2.1 Reducing segregation in school choice

The most common approach to reduce segregation by school choice seems to be the use of quotas, pioneered by Abdulkadiroğlu and Sönmez in their seminal work (referred to as controlled choice). These techniques assign "types" to students that represent their socio-economic status, ethnicity, gender, or any other attribute subject to diversity considerations. Diversity is enforced through minimum and maximum type quotas; ensuring a certain minimum or maximum number of a given type of student is accepted, respectively.

Building on this work, Kojima [19] proves that in a controlled choice paradigm, no stable mechanism "respects the spirit of quota-based affirmative action" [1] by showing that a stronger quota-based affirmative action policy (reducing type-specific capacity for the majority) can leave not just minority students, but all students worse off. Similarly, they prove that there is no stable mechanism that "respects the spirit of priority-based affirmative action"[2] by showing that a stronger priority-based affirmative action policy (leaving quotas alone but re-ordering school priority lists to increase minority students' rankings) can also result in a Pareto inferior outcome for minority students.

Hafalir et al. [16] reinforce this finding by running simulations which reveal that the scenarios described above are more than just edge cases. They contribute a "minority reserves" policy that Pareto dominates any stable matching produced by controlled choice and is never *strictly* Pareto dominated by deferred acceptance without affirmative action. They show this by considering the case where a school reaches their majority quota under controlled choice and has unfilled places - these under-full schools form blocking pairs with students who are either unassigned or prefer that school to their assigned school. They run an improvement algorithm which assigns students to under-full schools until either no under-full schools remain or no student would prefer an under-full school to their assigned school (no more blocking pairs), and hence produce another stable matching outcome which Pareto dominates the starting outcome.

However, solving instances of controlled choice can be intractable. Huang [18] presents the following dichotomy: if the student types are nested[3], then the problem is solvable in polynomial time; however if there is a single student who has overlapping types then it is NP-complete to decide whether there is a stable matching, even with no lower type quotas. This is supported by Aziz et al. [5], who show that checking both feasibility and stability is intractable for an instance of controlled choice.

Chen et al. [9] take this a step further and show that in general, deciding if a stable outcome exists for an instance of controlled choice is Σ_2^P complete - harder than NP.

2.2 Transportation and school choice

Trajkovski et al. [24] examine the effect of school bus provision on the schools that parents choose. They find that school bus eligibility makes parents more likely to select high quality schools that they would not otherwise consider due to the disutility of the distance. In some cases, it had a greater effect on choice than the quality of the school.

Burdick-Will and Stein [6] analyse the effects of a natural experiment in Baltimore City: the introduction of a new transportation system with little public advertisement and how this affected residents' school choice behaviours. They find that students avoid choosing schools with longer commutes or that require an additional transfer. It suggests that transport accessibility, independent from distance alone, contributes significantly to a student's decision to select a school to attend.

Cordes and Schwartz [10], writing for the Urban Institute, a non-profit, use data provided by the New York City Department of Education to understand the role of school buses and public transportation on students' school choices. They find that students who attend "choice schools" over neighbourhood schools, tend to attend better schools, and those that do are more likely to do so by using either school buses or public transportation.

Research by Angrist et al. [4] questions the financial viability of maintaining subsidised transport for students in Boston and New York City. They take advantage of random tie-breaking implemented in the matching mechanisms in both Boston and New York to estimate causal effects of non-neighbourhood enrolment. They find non-neighbourhood enrolment significantly reduces school

[1]"A matching mechanism ϕ is said to **respect the spirit of quota-based affirmative action** if there are no markets G and G' such that G' has a stronger quota-based affirmative action policy than G and ϕ(G') is Pareto inferior to ϕ(G) for the minority."
[2]"A matching mechanism ϕ is said to **respect the spirit of priority-based affirmative action** if there are no markets G and G' such that G' has a stronger priority-based affirmative action policy than G and ϕ(G') is Pareto inferior to ϕ(G) for the minority."

[3]e.g. out of a total capacity of 100, 50 must be type W and 50 type X, and given they are type X, 25 must be type Y and 25 type Z, and given they are type W, 10 must be type Y and 40 type Z BUT students who are both type X and Z do not impinge on the quota for students who are type W and Z, which is to say, student types do not overlap.

segregation but has small effects on test results and university attendance.

Finally, Michailidis et al. [20] use an agent-based simulation to investigate how school segregation under a central allocation mechanism is affected by interventions in transport systems. They find reducing segregation without transport is difficult, and that transport interventions can reduce school segregation even if agents have a moderate preference for schools with populations of same agent types.

3 MODEL

Quota-based approaches to reducing segregation are computationally hard [5, 9, 18]. In the UK, positive discrimination, such as quotas, are prohibited under the Equality Act 2010 [?]. The literature suggests transport strongly affects students' choice-making, therefore we attempt to incorporate transportation directly into the allocation mechanism as a way to combat proximity-induced segregation effects, without using quotas.

The mechanism takes disadvantaged zones, areas defined by socioeconomic metrics such as high free school meal eligibility, as one of its inputs. Students who live in these zones may have the option to additionally select a combined transportation-school tuple which bumps them into the same over-subscription priority bracket as students who live in the actual or de-facto catchment areas of schools served by the associated transport.

3.1 School choice

Our notation is inspired by Aziz et al. [5] where they describe:

Definition 3.1. An instance of school choice I^S consists of the tuple: $(S, C, q_C, \chi, \succ_S, \succ_C)$.

There are a set of students $S = \{s_1, ..., s_{n_S}\}$ and schools $C = \{c_1, ..., c_{n_C}\}$ with each school's capacity given by q_c, held in a capacity vector $q_C = (q_c)_{c \in C}$. Let $\chi = (S \times C) \cup (S \times \{\emptyset\})$ denote the set of possible matching outcomes, where $S \times \{\emptyset\}$ denotes the outcome where a student remains unmatched or takes an outside option like a private school. If a student ranks $(s, \{\emptyset\})$ above any other preference it indicates that any preference below it is unacceptable to them, effectively submitting an incomplete list. Finally $\succ_S = \{\succ_{s_1}, ..., \succ_{s_{n_S}}\}$ and $\succ_C = \{\succ_{c_1}, ..., \succ_{c_{n_C}}\}$ denote the set of strict student preferences over $C \cup \{\emptyset\}$ and strict school priorities over S, respectively. A feasible matching $\mu \in \chi$ is one which respects the preferences and priorities of the students and schools, respectively, and respects the capacities of schools. This means the matching will not assign students to schools that are unacceptable to them, and vice versa, and will not assign more students to a school than the school's capacity permits.

As mentioned in the introduction, Abdulkadiroğlu and Sönmez [3] identify desirable properties of their proposed algorithms; deferred acceptance is strategy-proof, and returns a stable matching which is student-optimal if run student-oriented, and top trading cycles is strategy proof, and returns a Pareto optimal matching. We focus on the properties of deferred acceptance, since our mechanism extends this algorithm. We draw on [12] for the definition:

Definition 3.2. A matching is stable if there exist no two students a and b who are assigned to schools A and B respectively, even though b prefers A to B and A prefers b to a.

The definition says that a matching is stable if no student and schools who are unmatched mutually prefer each other to their assigned partners, and is important in preventing a matching from unravelling.

Definition 3.3. A matching is Pareto optimal if no participant can improve their assignment without making at least one other participant's assignment worse. Similarly, a matching Pareto dominates another if no participant receives a worse assignment and at least one participant receives a better assignment.

Definition 3.4. A stable matching is student-optimal if there is no other stable matching in which at least one student is matched with a school they prefer more and no student is matched with a school they prefer less. Alternatively, a matching is student-optimal if it Pareto dominates all other stable matchings from the students' perspective.

Definition 3.5. A mechanism is strategy-proof if no participating student can improve their outcome by misrepresenting their preferences.

We make the assumption that schools cannot misrepresent their priorities because they are legally required to publish them, and thus their priorities are publicly known by all students. Eliciting a truthful preference from students is desirable for many reasons [2] including easier explanation of the mechanism to participants, and hence less chance of disadvantaged students being penalised by information asymmetries.

3.2 School choice with transportation

In this section we will build up our problem instance.

Definition 3.6. A disadvantaged area is an area of a city within which the relative rates of depravation are high, as given by appropriate socioeconomic metrics, e.g. eligibility for free school meals.

Definition 3.7. A route is a transport service with a finite capacity that connects a disadvantaged area to a school. Each route serves exactly one school and one disadvantaged area.

Definition 3.8. A student is eligible for a route if they live in a disadvantaged area that is served by that route.

We now define what we mean by an instance of school choice with transportation:

Definition 3.9. An instance I^T of school choice with transportation consists of the tuple:
$(S, C, D, D_D, f : S \to D, R, q_C, q_R, \chi, \succ_S, \succ_C)$.

- S and C are finite and disjoint sets of students and schools, respectively: $S = \{s_1, ..., s_{n_S}\}, C = \{c_1, ..., c_{n_C}\}$.
- D is the set of districts in the city $D = \{d_1, ..., d_{n_D}\}$.
- $D_D \subset D$ is the set of disadvantaged areas $D_D = \{\delta_1, ..., \delta_{n_{D_D}}\}$.
- $f : S \to D$ is a function that maps each student s to the district that they live in.
- $R \subseteq C \times D_D$ is the set of routes. A route $r_i = (c_j, \delta_\ell)$ connects school c_j to the disadvantaged areas δ_ℓ.

- Each school c has a capacity q_c and each route r_i has a capacity p_i, with the capacity vector of all school capacities being $q_C = (q_c)_{c \in C}$ and the capacity vector of all route capacities being $p_R = (p_i)_{r_i \in R}$.
- Each match x can be between a student and a school directly $x = (s, c)$, between a student and a school via the route that joins them $x = (s, c, r)$, or between a student and the outside option $x = (s, \{\emptyset\})$. Let $\chi = (S \times C) \cup (S \times C \times R) \cup (S \times \{\emptyset\})$ denote the set of available matches, where $(s, \emptyset)$ represents being unmatched (or selecting an outside option like homeschooling or private school).
- Let $\succ_S = \{\succ_{s_1}, \succ_{s_2}, ..., \succ_{s_{n_S}}\}$ denote the set of all student preferences. Each route-eligible student $s \rightarrow \delta$ has a strict preference ordering $\succ_{s \rightarrow \delta}$ over $(R \times C) \cup C \cup \{\emptyset\}$, that is, over schools and (route, school) tuples. Each route-ineligible student $s \rightarrow d \notin D_D$ has a strict preference ordering $\succ_{s \rightarrow d \notin D_D}$ over $C \cup \{\emptyset\}$. An option is acceptable to s if they prefer it to $\{\emptyset\}$.
- Let $\succ_C = \{\succ_{c_1}, \succ_{c_2}, ..., \succ_{c_{n_C}}\}$ denote the set of all school priorities. Each school c has a strict priority ordering $\succ_c$ over $(R \times S) \cup S$, that is, over students and (route, student) tuples.

School over-subscription priorities are divided into 2 overall priority brackets:

(1) Route *ineligible* students within a certain distance of a school selecting that school **without** transport AND route *eligible* students in a disadvantaged area selecting that school **with** transport,
(2) all other students.

Ties within the brackets are broken by proximity to the school. A matching $\mu \in \chi$ for an instance of school choice with transportation is feasible if:

- it respects the school and route capacities q_C and q_R respectively. This means no school and no route is assigned more students than its stated capacity.
- $(s_i, c_j) \in \mu$ only if s_i and c_j are mutually acceptable to each other;
- $(s_i, c_j, r_k) \in \mu$ only if $r_k = (c_j, \delta_\ell)$ where $f(s_i) = \delta_\ell$; which is to say, student s_i is mapped to the disadvantaged area that the route r_k connects to school c_j.

Consider again the example from the introduction. With the above framework, students who live close to the school will receive priority, but distant students who are eligible for route-selection receive equivalent priority.

3.3 Deferred acceptance with Transportation (DAT)

Our mechanism is an extension of the deferred acceptance mechanism [12] that takes an instance of school choice with transportation and outputs a matching $\mu \subseteq \chi$. We present the algorithm used in deferred acceptance with transportation mechanism below.

Route-school tuples are given by the notation $a_{i,j} = (r_i, c_j)$, where r_i is a route which serves c_j. Each tuple can contain exactly one route and school. Route-student tuples are given by the notation $b_{i,k} = (r_i, s_k)$ where r_i is the route that serves the disadvantaged area where s_k lives. Each tuple can contain exactly one route and student.

We assume that students only rank (route, school) tuples where they are eligible for that route. This can be enforced through a polynomial-time pre-processing algorithm that scans the provided preference lists and removes all ineligible tuples.

Algorithm 1 Deferred Acceptance with Transportation (DAT). Takes as input an instance of school choice with transportation $I^S = (S, C, D, D_D, f : S \rightarrow D, R, q_C, q_R, \chi, \succ_S, \succ_C)$ and returns a matching μ.

$\mu \leftarrow \emptyset$
$free \leftarrow S$
while some student $s_n \in free$ **do**
  if $\succ_{s_n}$ is empty **then** $free \backslash s_n$
  end if
  s_n proposes to their top preference (c_j or $a_{i,j}$)
  if proposing to $a_{i,j}$ **then** $\mu = \mu \cup \{b_{i,n}, a_{i,j}\}$ and $free \backslash s_n$
    if route r_i at capacity **then** $\mu = \mu \backslash \{b_{i,k}, a_{i,j}\}$ and $\succ_{s_k} = \succ_{s_k} \backslash a_{i,j}$ and $free \leftarrow s_k$ ▷ $b_{i,k} = (r_i, s_k)$ is lowest priority
    else if school c_j at capacity **then**
      if lowest priority is some $b_{\ell,m}$ **then** $\mu = \mu \backslash \{b_{\ell,m}, a_{\ell,j}\}$ and $\succ_{s_m} = \succ_{s_m} \backslash a_{i,j}$ and $free \leftarrow s_m$
      else if lowest priority is some s_m **then** $\mu = \mu \backslash \{s_m, c_j\}$ and $\succ_{s_m} = \succ_{s_m} \backslash c_j$ and $free \leftarrow s_m$
      end if
    end if
  else if proposing to c_j **then** $\mu = \mu \cup \{s_n, c_j\}$
    if school c_j at capacity **then**
      if lowest priority is some $b_{\ell,m}$ **then** $\mu = \mu \backslash \{b_{\ell,m}, a_{\ell,j}\}$ and $\succ_{s_m} = \succ_{s_m} \backslash a_{i,j}$ and $free \leftarrow s_m$
      else if lowest priority is some s_m **then** $\mu = \mu \backslash \{s_m, c_j\}$ and $\succ_{s_m} = \succ_{s_m} \backslash c_j$ and $free \leftarrow s_m$
      end if
    end if
  end if
end while

The algorithm starts with an empty matching. While there is a student s who is free and has not proposed to every school or (route,school) tuple in their preference list, s proposes to the next school or (route,school) tuple in their list. After proposing, s is tentatively matched with their top preference, and removed from the list of free students. The algorithm first checks if the student proposed with a route, and if they did, it then checks the route capacity. If the route has reached capacity, then the algorithm finds the lowest priority student who applied with the same route to the same school and removes them from the matching. The removed student then removes that (route,school) tuple from their preference list and will not apply to it in future iterations. If the route has capacity but the school does not, the lowest student or (route,student) tuple is removed from the matching, and removes the appropriate school or (route,school) tuple from its preference list, respectively.

If s applied without a route, the algorithm checks the school's capacity and if the school is at capacity then the lowest student or (route,student) tuple is removed from the matching, and removes the appropriate school or (route,school) tuple from its preference list, respectively, same as above. If a student is matched or

In our running example, let us say that the students in the good school's catchment area choose to apply to that school as before, but the distant, and now route-eligible, disadvantaged students also choose to apply to the good school, via the transport tuple: the mechanism ensures that those distant students will have an equal chance of getting in - as long as they select the transport option.

As we have already mentioned, school choice mechanisms can segregate schools through their over-subscription priorities. We have presented a mechanism which attempts to address this problem by allowing distant, disadvantaged students to increase their chances of attending a desirable school. In the following sections we present first some theoretical results for the properties of our mechanism, and then some preliminary experimental results from simulations.

3.4 Theoretical results

Theorem 3.10. *Deferred acceptance with transportation terminates in polynomial time.*

Proof. Consider an instance of school matching with transportation. We refer to schools or route-school tuples as objects:

(1) The number of students and schools, the length of the preference and priority lists, and the school and route capacities are all finite.
(2) Students propose to objects in order of preference. If they are accepted by an object they will not propose again unless they are unmatched with that object at a later stage. If they are unmatched later, they will not propose to objects that they have been unmatched with, so their preference lists will be weakly decreasing with every iteration of the algorithm.

□

Theorem 3.11. *Deferred acceptance with transportation returns a stable matching.*

Proof. Take any pair of student s and [school c OR tuple (r, c)] who are acceptable to each other and not matched together in μ after termination of the algorithm. In the case of tuple (r, c) suppose that r connects s's district to school c, and thus s is eligible for r.

Suppose that s prefers [c OR (r, c)] to $\mu(s)$, and hence s must have applied to [c OR (r, c)] before applying to $\mu(s)$. In the case where $\mu(s)$ is empty, the student s must be unmatched, and hence must have exhausted their entire preference list before ending up unmatched.

Since s is not matched to [c OR (r, c)] when the algorithm terminates:

- IF s is NOT eligible for r: either s applied to c, was accepted and then later rejected by c in favour of a higher priority student/routed student, OR c was already at capacity at the point of application and was rejected as the lowest priority student/routed student.
- ELSE IF s IS eligible for r and ONLY applied to (r, c); either:
  - s applied to (r, c), was accepted and was later rejected from r (and since r only serves c is therefore also rejected from c) in favour of a higher priority student routed to c, OR r was already at capacity at the point of application and s was rejected as the lowest priority student routed to r,
  - OR s applied to (r, c), was accepted and later rejected from c in favour of a higher priority student/routed student, OR c was already at capacity at the point of application and was rejected as the lowest priority student/routed student.
- ELSE IF s IS eligible for r and applied to BOTH (r, c) AND c; BOTH:
  - s applied to (r, c), was accepted and was later rejected from r (and therefore also c) in favour of a higher priority routed student, OR r was already at capacity at the point of application and s was rejected as the lowest priority routed student,
  - AND s applied to c, was accepted and was later rejected by c in favour of a higher priority student/routed student, OR c was already at capacity at the point of application and was rejected as the lowest priority student/routed student.

□

A key aspect of the stability of our algorithm is the order in which the over-subscription is resolved in the case where both school and route are oversubscribed: as long as route over-subscription is resolved first, school over-subscription is resolved simultaneously and the number of remaining seats at a school *can only weakly decrease with each step of the algorithm.*

Because of this, *the lowest priority student in a school's preference list can **only** weakly improve.* Hence, the cases above hold not only at the point of rejection, but at the termination of the algorithm.

Below are examples that showcase the above statements:

Example 3.12. Consider the case:

School	Priorities
c_1	$(r_1, s_4), (r_1, s_5), (r_1, s_6), s_3, (r_1, s_7)$

where c_1 has a capacity of 3 and r_1 has a capacity of 2.

If the students propose in the reverse order of school c_1's priority list, then we see that:

(1) $(r_1, s_6), s_3, (r_1, s_7)$ are initially assigned to c_1,
(2) (r_1, s_5) applies and displaces (r_1, s_7) since r_1 has reached capacity (as well as c_1) and (r_1, s_7) is the lowest priority:
   - r_1 and c_1 reach capacity simultaneously, but since the lowest priority item, (r_1, s_7), is simultaneously the lowest priority item in c_1's priority list *and* the lowest priority routed student, the order in which the overcapacity problem is solved does not matter.
(3) (r_1, s_4) applies and now there is a problem: both r_1 and c_1 are already at capacity, and the order in which the over-subscription is resolved is important:
   - If c_1's over-subscription is resolved first, then s_3 will be rejected from the school, after which (r_1, s_6) will be rejected from both route and school, resulting in the final allocation $c_1 : (r_1, s_4), (r_1, s_5)$, leaving the school *under-subscribed*, but importantly violating one of the principles that guarantees termination and stability: that schools only "trade up."
   - On the contrary, if the route over-subscription is resolved first, then only (r_1, s_6) will be rejected from both route and

school, which satisfies both capacity conditions simultaneously, and the final allocation will be: $c_1 : (r_1, s_4), (r_1, s_5), s_3$.

Example 3.13. Consider the next case where the priority list is prepended with s_1, s_2, who are local students in the upper priority bracket.

School	Priorities
c_1	$s_1, s_2, (r_1, s_4), (r_1, s_5), (r_1, s_6), s_3, (r_1, s_7)$

After following the same procedure as above, s_1, s_2 now apply to c_1 in reverse order as before, displacing first s_3, then (r_1, s_5). This will leave r_1 *undersubscribed*. However, this is not an issue here because there are no students, routed or otherwise who have a higher priority than the assigned students, and that would be able to change the assignment, as all other acceptable students have already been rejected, and they cannot re-apply later.

THEOREM 3.14. *Deferred acceptance with transportation is student optimal.*

PROOF. We show that no student is rejected by an achievable object and so the algorithm returns the student optimal stable matching.

Suppose the algorithm is at iteration i and no student has been unmatched with an object that is achievable to them:

(1) At this iteration i suppose route r, which serves only school c is at capacity and student s who applied to c with route r (henceforth called (r, s)), is rejected from both r and c:
  - If (r, s) is rejected in favour of (r, s'), then (r, s) is lower in c's priority list than (r, s').
  - Since there have been no rejections up until this iteration, and students apply to their top preferences first in the algorithm, we know that (r, s') prefers c over all other schools.
  - If we consider a matching μ where (r, s) is matched to c and all other students are matched to achievable objects, we know it will be unstable since (r, s') prefers c to its current partner, and c prefers (r, s') to (r, s).
  - Hence, there is no stable matching where (r, s) is matched to c, and so they are unachievable for each other.

(2) Else at this iteration suppose school c is at capacity and rejects student s or (r, s).
  - If s or (r, s) is rejected in favour of s' or (r, s'), then s or (r, s) is lower in c's priority list than either of s' or (r, s').
  - Clearly s' or (r, s') prefer c over all other schools since students apply to their top preferences first in the algorithm, and no student has been rejected as of iteration i.
  - If we consider a matching μ' where s or (r, s) is matched to c and all other students are matched to achievable objects, we know it will be unstable since s' or (r, s') prefers c to its current partner, and c prefers s' or (r, s') to s or (r, s).
  - Hence, there is no stable matching where s or (r, s) is matched to c, and so they are unachievable for each other.

□

4 EXPERIMENTAL RESULTS

We hypothesise that in a school choice setting with transportation, our mechanism will reduce school segregation caused by proximity-based over-subscription. Of course if a setting allows many (school, district) routes, e.g. between every good school and every disadvantaged area as an extreme case, segregation will reduce. However, transportation is costly and adds to congestion and pollution. Therefore, our ultimate aim is to investigate the relationship between a limited number of (school, district) routes, reduction in segregation, and the number of car trips displaced by a single route.

In this section we provide our preliminary results. We design a simulation that replicates the environment in which these problems occur. As in our running example, in this environment advantaged students are clustered around the best schools, and disadvantaged students are forced further out to the periphery. We assume that the environment contains transportation infrastructure that can be utilised by the students.

Experiments were carried out on a two dimensional grid environment, with distances between points calculated by the Manhattan distance:

$$d_M(a, b) = \sum_{i=1}^{n} |a_i - b_i|$$

where n = 2, and a and b are two points on the grid. We choose to use a grid with Manhattan distance for simplicity of implementation. schools were placed randomly on the grid with a minimum distance condition to stop them from overlapping. Each school has a quality attribute, and a catchment area, as well as a maximum capacity. Students have a socio-economic attribute that divides them into high and low types to represent advantaged and disadvantaged students, respectively. They were placed on the grid according to their socio-economic attribute, with high type students being highly likely to appear in the catchment areas of high quality schools. Disadvantaged zones were placed such that they did not overlap with the catchment areas of schools with a quality attribute above 0.5 and routes were assigned to connect each zone to each school.

Students rank schools proportional to school quality, and inversely proportional to the travel distance to the school. If a student lives in the catchment area of a school, it ranks that school higher in their preference list. If a student is eligible for a route that connects them to a school, they will rank that route-school tuple higher than the same school without a route, and also higher overall. Schools rank students based on the priority brackets introduced in section 3.3, and ties are broken based on proximity.

Using the above as input, we applied both deferred acceptance and deferred acceptance with transport to match 200 students to 10 schools with 3 disadvantaged areas. Each disadvantaged area provides a potential route to each school for a total of 30 possible routes. All colleges have their capacities set to 30, the minimum class size per teacher for primary schools in the UK, and all routes have a capacity of 30. Fig 1 shows the difference between mechanisms for the same instance. From the figure we see that almost all of the students in disadvantaged zones took advantage of the routes, and those that did all ended up matched to higher quality schools. To attend those schools, the students collectively only used 6 routes even though 30 were available for students to use.

We use the dissimilarity index as a measure of school integration to compare algorithms:

$$D = \frac{1}{2} \sum_{i=1}^{N} \left| \frac{l_i}{L} - \frac{h_i}{H} \right| \qquad (1)$$

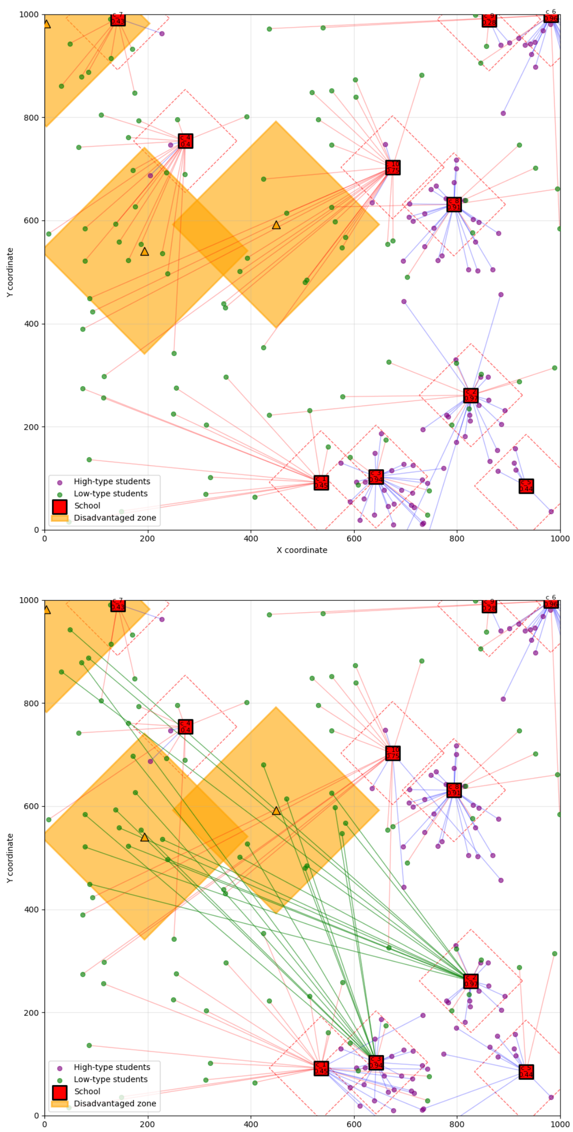

Figure 1: Map of a simulated municipal area. Numbers below the school name display the quality attribute. Green lines indicate allocations made with routes, blue lines indicate high type student allocations and red lines indicate low-type student allocations.

where l_i denotes the population of low-type students matched to school i, L is the total population of low-type students, h_i denotes the population of high-type students matched to school i, and H is the total population of high-type students. A dissimilarity index of 0 indicates perfect integration, and an index of 1 indicates perfect segregation. The simulation was run 100 times and the mean dissimilarity index for deferred acceptance with routes was 3.78 percentage points lower than for deferred acceptance, indicating a lower average dissimilarity across schools, shown in fig.2.

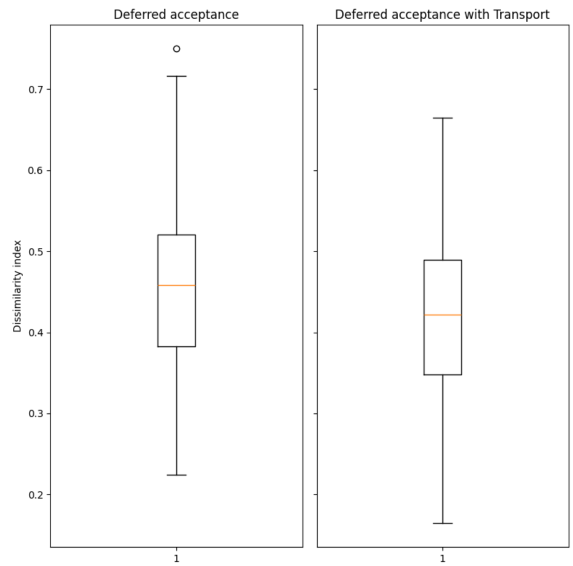

Figure 2: Dissimilarity index of Deferred acceptance vs Deferred acceptance with Transport. Results gathered over 100 iterations.

5 FUTURE WORK

First and foremost, we plan to enrich our simulations and experiments to investigate the relationship between the number of routes, reduction in segregation, and the number of car trips displaced by a single route (to estimate the effect our mechanism may have on urban emissions and congestion). To extend this idea further, the problem could be recast as a graph optimisation problem, where the optimal placement of transport routes could be explored under the mechanism to see how it affects school segregation, and urban emissions and congestion.

We plan to apply the deferred acceptance with transport mechanism to a set of inputs informed by real municipal areas, and data collected on preferences made there, to improve the accuracy of our simulation. The problem instance can be extended to include routes that serve multiple schools or disadvantaged areas with a shared route capacity, and we plan to investigate the effects on the mechanism's computational complexity. The mechanism does not take into account the three choice limit that is often imposed on school choice in England, which may induce agents to misrepresent their preferences. It would be interesting to study how it affects our mechanism's properties, and if the mechanism remains strategy-proof.

ACKNOWLEDGMENTS

This work was supported by the UKRI AI Centre for Doctoral Training in AI for Sustainability [EP/Y030702/1].

REFERENCES

[1] [n.d.]. Free School Transport. https://www.gov.uk/free-school-transport.
[2] Atila Abdulkadiroglu, Parag Pathak, Alvin E. Roth, and Tayfun Sonmez. 2006. Changing the Boston School Choice Mechanism. https://doi.org/10.3386/w11965 arXiv:11965

[3] Atila Abdulkadiroğlu and Tayfun Sönmez. [n.d.]. School Choice: A Mechanism Design Approach. 93, 3 ([n. d.]), 729–747. https://doi.org/10.1257/000282803322157061

[4] Joshua Angrist, Guthrie Gray-Lobe, Clemence M. Idoux, and Parag A. Pathak. [n.d.]. *Still Worth the Trip? School Busing Effects in Boston and New York.* https://doi.org/10.3386/w30308 arXiv:30308

[5] Haris Aziz, Serge Gaspers, Zhaohong Sun, and Toby Walsh. 2020. From Matching with Diversity Constraints to Matching with Regional Quotas. https://doi.org/10.48550/arXiv.2002.06748 arXiv:2002.06748 [cs]

[6] Julia Burdick-Will and Marc L. Stein. [n.d.]. Transit Trade-Offs: Public Transportation Difficulty, Schedule Variation, and School Preferences. 128 ([n. d.]), 103160. https://doi.org/10.1016/j.ssresearch.2025.103160

[7] Simon Burgess, Ellen Greaves, and Anna Vignoles. [n.d.]. School Choice in England: Evidence from National Administrative Data. 45, 5 ([n. d.]), 690–710. https://doi.org/10.1080/03054985.2019.1604332

[8] Simon Burgess, Ellen Greaves, Anna Vignoles, and Deborah Wilson. [n.d.]. Parental Choice of Primary School in England: What Types of School Do Different Types of Family Really Have Available to Them? 32, 5 ([n. d.]), 531–547. https://doi.org/10.1080/01442872.2011.601215

[9] Jiehua Chen, Robert Ganian, and Thekla Hamm. 2020. Stable Matchings with Diversity Constraints: Affirmative Action Is beyond NP. In *Twenty-Ninth International Joint Conference on Artificial Intelligence*, Vol. 1. 146–152. https://doi.org/10.24963/ijcai.2020/21

[10] Sarah A. Cordes and Amy Ellen Schwartz. [n.d.]. Does Pupil Transportation Close the School Quality Gap? https://www.urban.org/research/publication/does-pupil-transportation-close-school-quality-gap

[11] Dennis Epple, Richard E. Romano, and Miguel Urquiola. [n.d.]. School Vouchers: A Survey of the Economics Literature. 55, 2 ([n. d.]), 441–492. https://doi.org/10.1257/jel.20150679

[12] D. Gale and L. S. Shapley. 1962. College Admissions and the Stability of Marriage. *The American Mathematical Monthly* 69, 1 (1962), 9–15. https://doi.org/10.2307/2312726 arXiv:2312726

[13] Stephen Gibbons, Stephen Machin, and Olmo Silva. [n.d.]. Choice, Competition, and Pupil Achievement. 6, 4 ([n. d.]), 912–947. https://doi.org/10.1162/JEEA.2008.6.4.912

[14] Stephen Gibbons, Stephen Machin, and Olmo Silva. 2013. Valuing School Quality Using Boundary Discontinuities. *Journal of Urban Economics* 75 (May 2013), 15–28. https://doi.org/10.1016/j.jue.2012.11.001

[15] Philip Gleason, Melissa Clark, Christina Clark Tuttle, and Emily Dwoyer. [n.d.]. The Evaluation of Charter School Impacts. ([n. d.]).

[16] Isa E. Hafalir, M. Bumin Yenmez, and Muhammed A. Yildirim. 2013. Effective Affirmative Action in School Choice. *Theoretical Economics* 8, 2 (2013), 325–363. https://doi.org/10.3982/TE1135

[17] Chang-Tai Hsieh and Miguel Urquiola. [n.d.]. The Effects of Generalized School Choice on Achievement and Stratification: Evidence from Chile's Voucher Program. 90, 8–9 ([n. d.]), 1477–1503. https://doi.org/10.1016/j.jpubeco.2005.11.002

[18] Chien-Chung Huang. 2010. Classified Stable Matching. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '10)*. Society for Industrial and Applied Mathematics, USA, 1235–1253.

[19] Fuhito Kojima. 2012. School Choice: Impossibilities for Affirmative Action. *Games and Economic Behavior* 75, 2 (July 2012), 685–693. https://doi.org/10.1016/j.geb.2012.03.003

[20] Dimitris Michailidis, Mayesha Tasnim, Sennay Ghebreab, and Fernando P. Santos. [n.d.]. Tackling School Segregation with Transportation Network Interventions: An Agent-Based Modelling Approach. 38, 1 ([n. d.]), 22. https://doi.org/10.1007/s10458-024-09652-x

[21] Phuong Nguyen-Hoang and John Yinger. 2011. The Capitalization of School Quality into House Values: A Review. *Journal of Housing Economics* 20, 1 (March 2011), 30–48. https://doi.org/10.1016/j.jhe.2011.02.001

[22] Camilo Quintero-Fragozo, Yasna Cortés, and Mauricio Sarrias. [n.d.]. Effects of Spatial Competition on Public Educational Efficiency: An Analysis for the Chilean Metropolitan Region. 38, 1 ([n. d.]), 115–140. https://doi.org/10.1080/02680939.2021.1966102

[23] Lloyd Shapley and Herbert Scarf. 1974. On Cores and Indivisibility. *Journal of Mathematical Economics* 1, 1 (March 1974), 23–37. https://doi.org/10.1016/0304-4068(74)90033-0

[24] Samantha Trajkovski, Jeffrey Zabel, and Amy Ellen Schwartz. [n.d.]. Do School Buses Make School Choice Work? 86 ([n. d.]), 103607. https://doi.org/10.1016/j.regsciurbeco.2020.103607

[25] Deborah Wilson and Gary Bridge. [n.d.]. School Choice and the City: Geographies of Allocation and Segregation. 56, 15 ([n. d.]), 3198–3215. https://doi.org/10.1177/0042098019843481