The Soft Neighborhood Model: A Dynamic Enrollment-Balancing Framework

Brooke Cowan
Matt Marjanović
brooke@softneighborhoodmodel.org
maddog@softneighborhoodmodel.org
http://www.softneighborhoodmodel.org/

March 7, 2016, v3.1

. . . Before I built a wall I'd ask to know
What I was walling in or walling out,
And to whom I was like to give offense.
Something there is that doesn't love a wall,
That wants it down. . .

The Mending Wall, Robert Frost

Prologue: A Public Comment

The following public comment was delivered at the PPS Board meeting held on October 5, 2015.

Matt and I developed a solution to the enrollment balancing problem and attempted to present it to DBRAC in a series of six 2-minute public comments in the spring. This Soft Neighborhood Model eliminates the need for frequent, manual boundary changes — a critical piece of the current DBRAC proposal. We think the larger community is interested in understanding more about how our model works. Media reports and DBRAC summaries have unfortunately mischaracterized the model and its intents, so we have collected the six original comments into a single, cohesive document which we emailed to each of you this morning. It is also available online at soft–neighborhood–model–dot–org. We would like to take this opportunity to introduce the Soft Neighborhood framework to you — the Board — and to the larger community.

The Soft Neighborhood Model has been referred to in DBRAC documents and in media reports as a "soft boundary" model, but in fact it is not a boundary model at all:

- It is a *neighborhood* model that places students in schools close to their homes.
- It is a *dynamic enrollment balancing* model that enables long-term enrollment stability for all schools.
- It is an *equitable* model that encourages the mixing of populations and curtails the ability of families with more resources to buy into the public school of their choice.

Within our framework, children are assigned to one of the several schools that are closest to their home. The criteria used for making a final assignment are (1) proximity and (2) capacity. The *proximity* part means that if School A is closer than School B, then assignment to A is more likely than assignment to B. The *capacity* part means that the model will assign the right number of students to each school according to target capacities provided by the district. In other words, the model assigns children to schools using an algorithm that satisfies each school's enrollment target, but is otherwise more likely to place them closer to their home.

Its performance — in fact, the performance of any enrollment system considered by the district — should be measured by concrete, objective criteria, such as how balanced the enrollments are at each school; how far students have to travel to get to school; how diverse the assignments are.

The Soft Neighborhood framework is not a School Choice model. The Soft Neighborhood Model does not ask families to rank their preferred schools. Like any traditional neighborhood school model, the only input a family has is its home address.

The Soft Neighborhood Model emerged from the observation that there is a tightly-coupled relationship between boundaries, enrollment instability, and inequity. While the relationship between enrollment balancing and equity has been widely acknowledged, the critical role that boundaries play has not. School boundaries create tangible walls within our city that make it possible to say, "These kids will never go to school with those kids." Boundaries encourage wealthier families to move away from "bad" geographic regions and cluster into "good" ones. By tying school assignment to real estate, boundaries drive housing prices up in some regions and down in others. Thus, boundaries facilitate enrollment imbalance; they facilitate socio-economic stratification, and they turn public school assignment into a commodity. Boundaries are a historical artifact that have co-evolved with and reinforced the same racial and socio-economic inequities that we as a community are struggling to overcome. Indeed, these boundaries are the defining feature of the enrollment system which you are being asked to endorse tonight. Whether those boundaries are nudged around year-to-year or once a decade, the underlying system is the same and the end result will be the same.

The Soft Neighborhood Model is a powerful alternative to the PPS Hard Boundary system. This model emphasizes enrollment balancing and a sense of neighborhood, as well as mixing populations that would otherwise become stratified and segregated.

We have implemented and tested the Soft Neighborhood Model against seven years of historical PPS enrollment data. We have defined three objective metrics for evaluating our model on the basis of balanced enrollments, travel distance, and school assignment diversity. Our experimental results suggest that the Soft Neighborhood Model outperforms the historical Hard Boundary system with respect to enrollment balancing; that the Soft Neighborhood Model still assigns students to schools reasonably close to home, and that the Soft Neighborhood Model greatly improves assignment diversity. These empirical results demonstrate the potential of the Soft Neighborhood Model as a robust and equitable solution to the enrollment balancing problem in PPS.

We encourage you and the larger community to seriously consider the Soft Neighborhood Model for PPS. Moreover, we would like to see compelling empirical evidence that the DBRAC proposal can actually achieve its stated balancing and equity goals, and we request an objective evaluation of the district's proposed solution in terms of the metrics we have defined. Any such evaluations must be implemented with transparency to the public — it should use a data set that has been properly anonymized and released to the public, and its entire methodology should be made available to the public for independent replication, verification, and validation.

Contents

1 The Problem With Boundaries

1.1 Boundaries are Like Fences

The Soft Neighborhood Model is a novel framework for enrollment balancing that is designed to place the right number of students at each school while keeping families in schools close to their homes. The elimination of boundaries is a critical component of the model. By its very nature, a system of hard boundaries drives our school district into a state of enrollment imbalance and reinforces segregation and inequity.

You can think about a boundary like a fence. In the physical world, the lines on the district's boundary maps form fences. We all know where these fences are because they determine to which school community we belong. Redrawing the boundaries means digging up the fence posts and moving the fences. Boundary change is a painful process because these fences, once we erect them, become a strong part of our identity, and we have a lot at stake in them. We resort to boundary change because there are pressing problems that require immediate attention, and a boundary change appears unavoidable. Following a boundary change, we may get some short-term relief; but there has never been a boundary change that has solved enrollment problems in the long-term. This is because:

1. Families move. PPS can decide where to put its fences, but it can't decide where to put families.
2. Families with more resources and more privilege have more freedom to move when and where they want.
3. Moving is not just a matter of economics but also of culture, so some families are less empowered to move than others due to racism and classism.

In short, a district that maintains fences promotes a system of imbalance, choice, and privilege. No matter where you move the fences today to even out enrollment and inequity, people will continue to move around and the inequities and imbalances between one side of a fence and the other will return. That's just what any system of fences always does.

The Soft Neighborhood Model starts by acknowledging that boundaries are a social construct that drive the district toward imbalance. Removing boundaries facilitates enrollment balancing by allowing the district to share capacity across nearby schools. It encourages the mixing rather than the segregation of populations, and it curtails the ability of families with more resources to buy into a public school of their choosing. Lastly, eliminating boundaries eliminates the need for boundary change. There are no boundary change events in the Soft Neighborhood Model because there are no boundaries.

1.2 A Boundary-Free Solution

The enrollment balancing solution that the district is currently developing involves regular and frequent boundary change: whenever boundary adjustments need to occur, the district wants to make them in a way that is consistent with its equity goals. The trouble is, boundaries are inherently incompatible with equity and balance. PPS boundaries turn school assignment into a commodity: they allow people with means to buy assignment to the public school of their choice when they purchase a house (Figure 1). Boundaries partition Portland geographically and intensify stratification. Aligning school boundaries with any equity policy is deeply contradictory, and any framework that claims it can accomplish this should be viewed with skepticism.

The Soft Neighborhood Model is an alternative to a boundary-based solution. The framework is designed to dynamically adapt to shifts in population. When there are more kids for a few years in one area of the city, nearby schools will accommodate them. If there are fewer kids in some area, nearby kids will fill in the gap. The model's ability to fully balance the enrollments at PPS schools is not unlimited. It is constrained by proximity, and so it cannot assign kids to schools clear across town. In other words, the Soft Neighborhood Model makes a necessary compromise between proximity and enrollment balancing. However, the model is designed to be as elastic and fluid as possible to accommodate population shifts, which are both inevitable and hard to predict. Also, the model is intended to provide an assignment mechanism for new assignees — students who need a new public school assignment, for example because they have moved or because they are newly entering the system. This implies that the model is not at liberty to re-assign students for the sake of balancing enrollments. Because of these restrictions, there may still be local enrollment imbalances from time to time, but the Soft Neighborhood Model should reduce the incidence of enrollment crises significantly in the long run.

The Soft Neighborhood Model was designed with increased equitable access to our public schools in mind. Boundaries extend a guarantee and a special privilege to people with more resources. The Soft Neighborhood Model dilutes that privilege. Boundary adjustment — even if done frequently — does not. Frequent boundary change will engender zones of predictability (close to schools) and zones of uncertainty (in the margins between schools). Those who have the means will avoid the zones of uncertainty. The Soft Neighborhood Model has no such zones. It creates a system of overlapping neighborhoods in which public schools are viewed as communities of families who all live nearby.

In a Hard Boundary model, the geographic space is partitioned into non-overlapping regions, and those regions are inextricably tied to a school community. In the Soft Neighborhood Model, school communities overlap geographically. This change provides the district with a very flexible means for achieving enrollment stability across schools from year to year. This is because the model distributes the students in a geographic region over more than just a single school. In the current system, students who live across the street from one another, but on opposite sides of a cluster boundary, will almost never be assigned to the same school (Figure 2, top). This school assignment model makes it possible to say, "*These* kids will never go to school with *those* kids." In contrast, in the Soft neighborhood model, any two children who live close to each other might be assigned to the same school, no matter what sides of the street they live on (Figure 2, bottom). Schools are populated with a mix of nearby students.

1.3 Values in the Soft Neighborhood Model

Central to the Soft Neighborhood Model are the notions of *neighborhood*, *family*, *equity*, *stability*, and *consistency*:

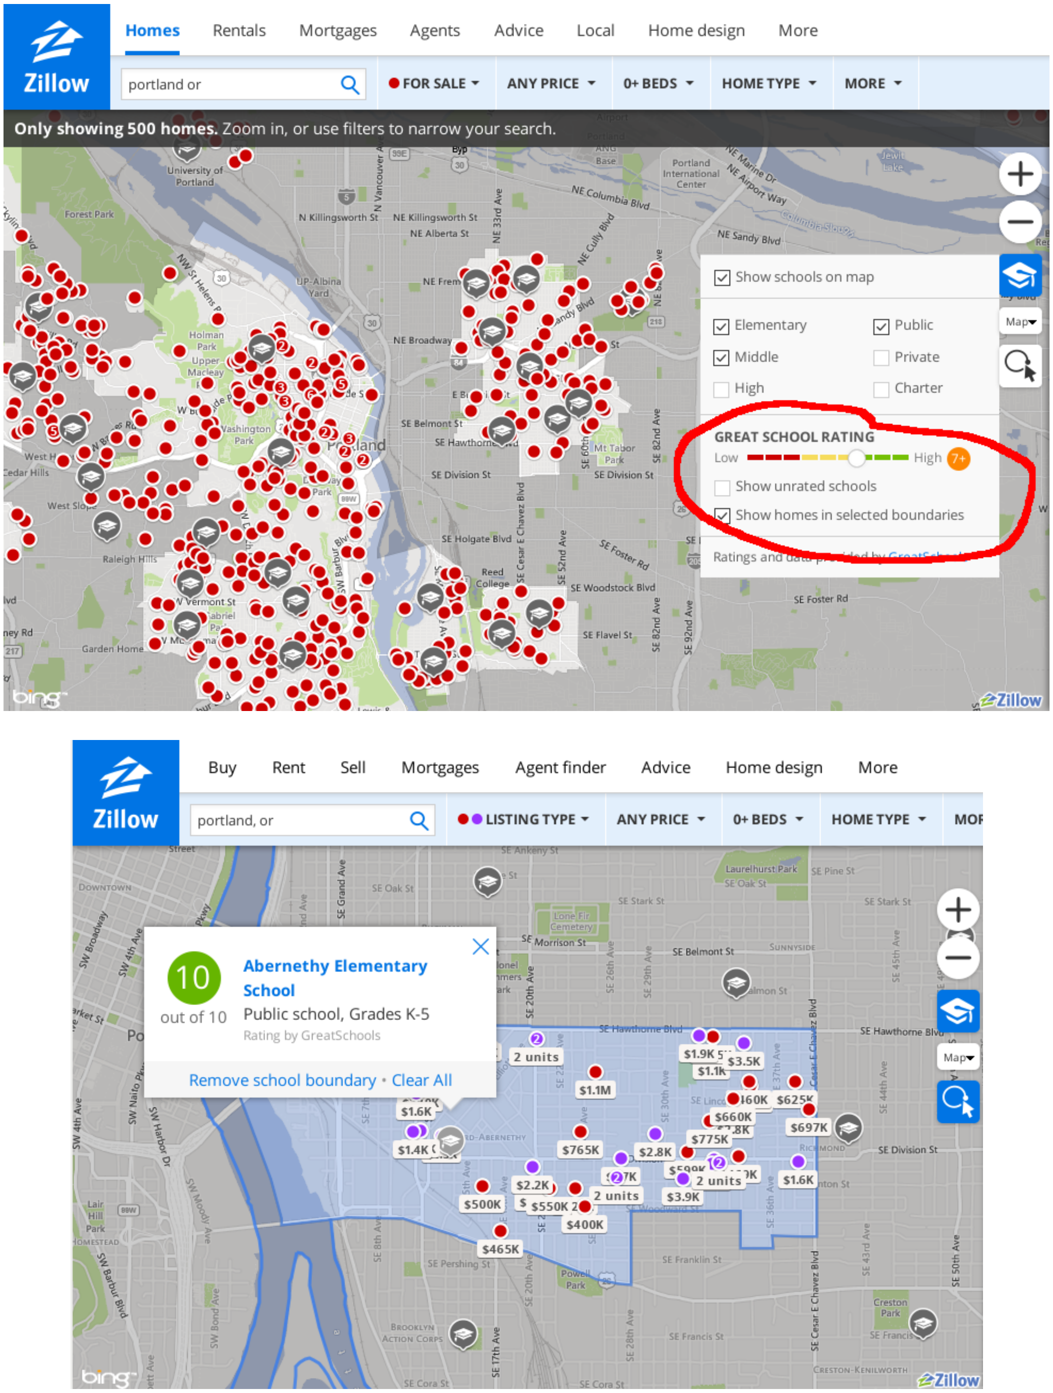

Figure 1: School assignment is a commodity in Portland. People with means can buy guaranteed assignment to the school of their choice when they purchase a house. This is a deeply-rooted source of inequity in our city.

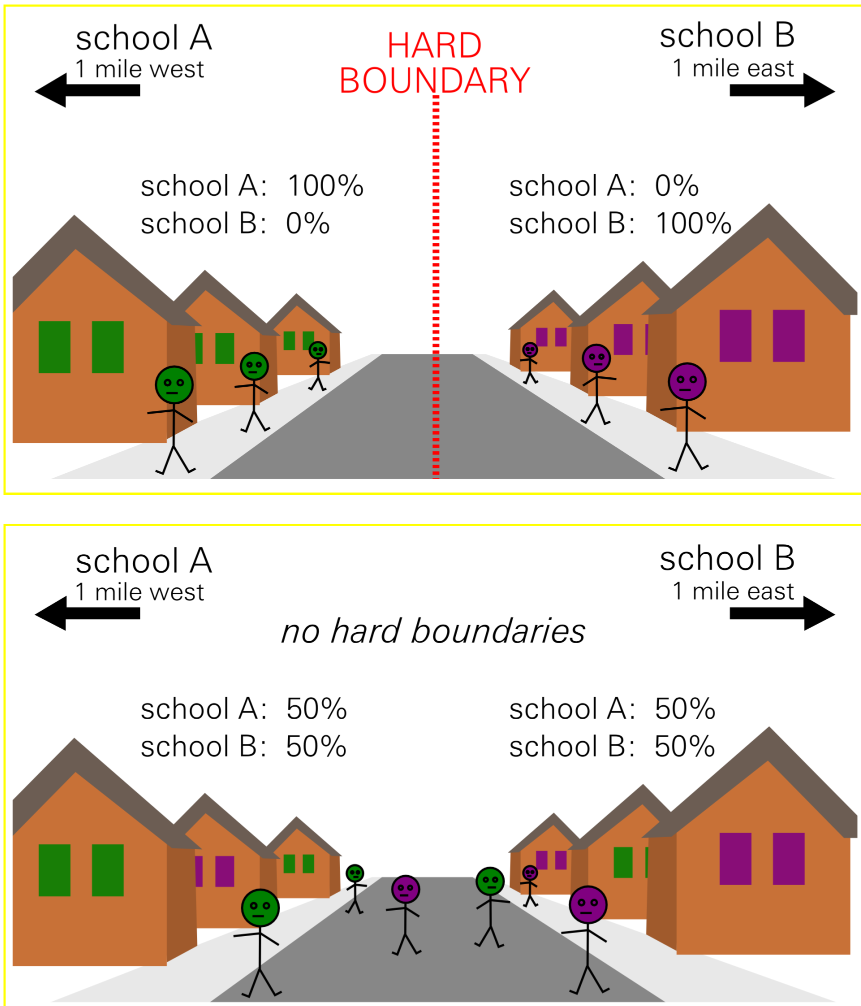

Figure 2: (top) A hard boundary will segregate students on one side of the street from those on the other. (bottom) In the Soft Neighborhood Model, hard boundaries don't exist, and schools are populated with a mix of nearby students.

Neighborhood In the Soft Neighborhood Model, what is close to your home is your neighborhood, and children are assigned to nearby schools. The neighborhood does not change suddenly by simply crossing the street; families that live close to each other are basically in the same neighborhood. This definition is consistent with an intuitive sense of "neighborhood" and "proximity".

Family Schools are treated as communities of families in the Soft Neighborhood Model. The option for families to co-enroll younger siblings with older ones, and the emphasis on proximity between home and school are two concrete ways in which the model expresses this value. Proximity of home and school is critical because it makes it easier for families to participate in school life.

Equity It seems unlikely that there exists a solution to the enrollment balancing problem that can also fix all of the equity problems in our school district. Where possible, the Soft Neighborhood Model aims to improve equity of access to and equity of opportunity within our public schools. By removing boundaries, the model reduces the ability of people with more resources to buy into a public school. By addressing enrollment balancing directly, the model avoids chronic under- and over-enrollment, which can become an equity problem. By encouraging mixing from overlapping school communities, the model mitigates the social and economic stratification that can exacerbate equity problems.

Stability The Soft Neighborhood Model enables long-term enrollment stability for all schools. By sharing capacity between nearby schools, the model grants the district far more flexibility than it has today to meet capacity-based enrollment targets in the face of unknown population changes. This property guards against both under- and over-enrollment. The underlying assumption is that greater stability in school enrollments leads to better educational environments for everyone.

Consistency Long-term consistency and predictability is a driving feature of the Soft Neighborhood Model. The model provides a consistent and stable set of expectations to parents around assignment. In contrast, the current PPS assignment model might appear predictable, except whenever the district needs to make changes in response to crises. These cataclysmic events have unpredictable outcomes (sometimes even forcing families to change schools mid-course, or to decide whether to send siblings to different schools or move an older sibling to a new school) and breed feelings of fear and insecurity. In the Soft Neighborhood framework, the intent is to provide an enrollment balancing mechanism that is robust enough to enable the district to guarantee that once a family is enrolled at a school, their children can remain at that school until they finish.

1.4 Overview of this Document

This document serves two related purposes: it is a proof-of-concept for the Soft Neighborhood Model, and it is a proof-of-concept for the development process of any solution to the

PPS enrollment balancing problem. As a proof-of-concept for the model, it addresses three important questions about the model:

1. How far do students have to travel?
2. How well is the model able to stabilize enrollments?
3. To what degree does it result in the mixing of populations?

This document provides answers to these questions that strongly position the Soft Neighborhood Model as a viable and robust enrollment balancing solution for PPS.

As a proof-of-concept for a rational process for developing solutions to the enrollment balancing problem, this document defines objective metrics and criteria that can be used to test any proposal, and to evaluate its performance once it's in place. Any proposed solution that is under serious consideration — whether it comes from the community, from the district, or from DBRAC — should be thoroughly evaluated before it is implemented, and should continue to be evaluated after it is implemented. For example, we would like to see the same questions we've asked of the Soft Neighborhood Model applied to the Frequent-Boundary-Change solution that the district is promoting. If this solution were implemented on the same data set supplied to us by PPS, how stable would enrollments be? How far would students have to travel? And how diverse would school assignments be? What would the boundaries look like in that seven-year time frame? How much work goes into making those boundary adjustments? Answering these questions objectively, and making predictions and projections ahead of time of what the community can expect to see under some proposed model, are critical steps in the development process.

The remainder of this document is structured as follows: in Section 2, we give a qualitative overview of the Soft Neighborhood Model. Section 3 uses a synthetic data set to demonstrate how the model works in action. In Section 6, we run the model on a data set containing seven years of historical PPS data and make three key empirical observations:

1. The Soft Neighborhood Model does a far better job of controlling over- and under-enrollment at the kindergarten level.
2. Both models — Soft Neighborhood and historical — assign most students to a school within a reasonable distance of their home.
3. Soft Neighborhood assignments exhibit a higher degree of school assignment diversity compared with historical assignments.

Section 4 explains how the Soft Neighborhood Model fits into the larger body of school assignment literature, and in particular how it differs from School Choice models. Section 9 answers a number of frequently-asked questions about the model. Section 10 concludes the work and outlines what we see as the next steps.

2 How Does the Soft Neighborhood Model Work?

The Soft Neighborhood Model is a framework for assigning students to schools. The model makes assignments using an algorithm that satisfies schools' capacity constraints, and otherwise is more likely to place students in schools closer to their homes. This algorithm is used to place *new assignees*: students who need a new assignment — for example, because they have moved or because they are entering the system for the first time. School assignment in the Soft Neighborhood Model is not something that is meant to be done to every student every year for the purposes of keeping enrollments balanced. Students matriculating up from an earlier grade are pre-assigned by the model, and are not placed using the assignment algorithm. Similarly, when a younger sibling enters the system, families can choose to co-enroll at the same school as the older sibling(s).

In any given year, the Soft Neighborhood Model will satisfy the constraints imposed by target capacities at each school to the extent possible. Before assignments are made, the model uses a *proximity function* to compute a set of "nearby" schools for each student based on his or her home address. This set contains the schools to which he or she can be assigned. Since assignment can only be made to one of the schools in this set, the Soft Neighborhood Model's ability to fully satisfy capacity constraints is not unlimited; however, in practice, it does a very good job of balancing enrollments even when subject to these restrictions (see Section 6).

One of the open questions that requires further investigation on an improved PPS data set is how to best define the proximity function. For the purposes of this proof-of-concept, we use a *3-School* rule: for each student, assignment can be made to any school within

- 110% of the distance to the 3rd closest school, or
- 1.25 miles,

whichever is larger. This definition has the desirable property that it scales itself to the school-density near each student. Students living in areas of the city with more schools nearby (i.e., more densely-populated regions) will be assigned to closer schools than students living in less densely-populated areas, but every student will have at least three available schools.

2.1 Examples: The Soft Neighborhood Model on the East and West Sides

Let's consider a couple of examples to help illustrate how student assignment actually works under this model. It's instructive to consider one example on the east side of the river, and one on the west side, since the densities in the two regions are so different, and what defines a "nearby" school depends on density.

We'll start with an incoming kindergartner who lives at NE 22nd Ave and Mason, pretty much equidistant from Sabin and Alameda, who would be assigned to Sabin under the current system. Under the Soft Neighborhood Model, that student will be assigned to one of several nearby schools. Using the 3-School rule defined above, she can be assigned to Sabin (0.3 mi), Alameda (0.4 mi), Vernon (0.7 mi), King (0.95 mi), Irvington (1.0 mi), or

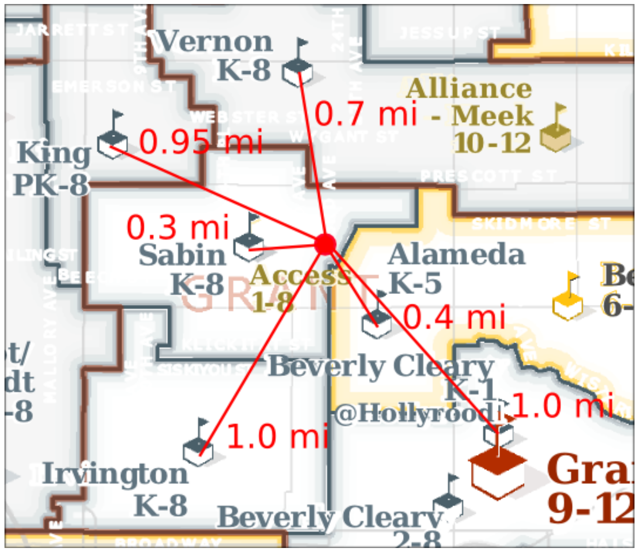

Figure 3: Possible school assignments for a student living at NE 22nd and Mason, with Cartesian distances to each school.

Beverly Cleary Hollyrood Campus (1.0 mi) (Figure 3).[1] Other kids on the same block may be assigned to different schools, but every child will have classmates that live nearby, no matter what school she is assigned to. Families are guaranteed placement at one of these nearby schools, but not a particular school.

Now let's see what happens in an example on the west side of the river. Let's consider an incoming student who lives near SW 57th and Taylor near East Sylvan School. His family currently resides within in the Chapman boundary. In this case, the closest three schools are Bridlemile (1.6 mi), Ainsworth (1.8 mi), and Chapman (2.1 mi).[2] Because the west side is more spread out, the closest three schools are farther away than on the east side. None is within 1.25 miles, and so this student has just those three schools available for assignment. The Soft Neighborhood Model will assign the child to one of these three schools, and will optimize the overall assignment process so as to avoid overcrowding or underpopulating any one school.

2.2 The Soft Neighborhood Model in Four Steps

We've seen that the Soft Neighborhood Model assigns students to one school in a set of nearby schools. Now let's take a closer look at how it actually picks one of those schools. For most kids (all except for those who are pre-assigned, like siblings and students moving up from the preceding grade), assignment is based on two factors: how far they'd have to travel to get to school from their house, and target enrollments at the schools near their

[1] These are Cartesian distances. Walking distances according to Google maps are 0.4 mi (Sabin), 0.4 mi (Alameda), 0.8 mi (Vernon), 1.3 mi (King), 1.4 mi (Irvington), and 1.4 mi (Beverly Cleary at Hollyrood). The proximity function and the metric used to estimate distance both have a strong impact on how the model actually works when implemented. What works best for Portland needs to be researched and validated against a proper data set.

[2] Driving distances are 3.2 mi (Bridlemile), 3.4 mi (Ainsworth), and 3.0 mi (Chapman).

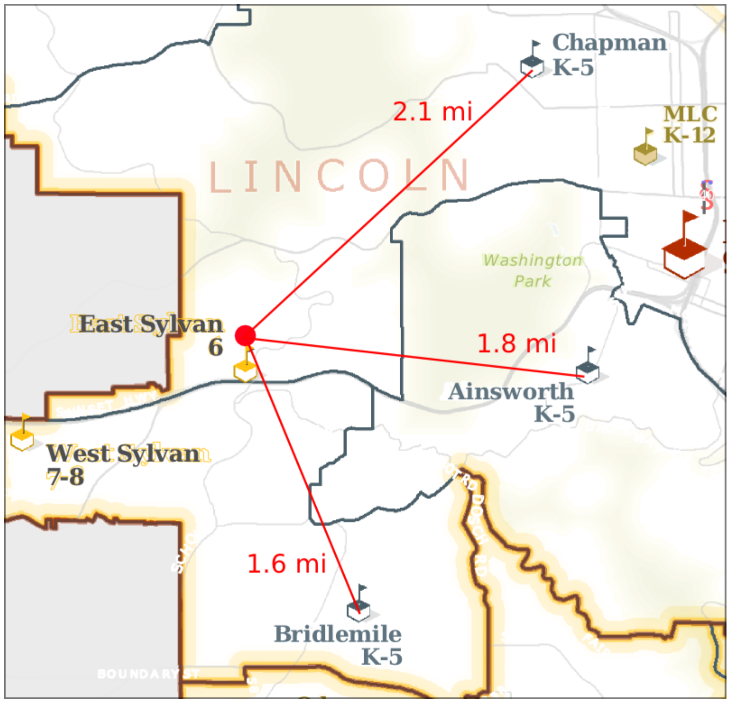

Figure 4: Possible school assignments for a student living near SW 57th and Taylor, with Cartesian distances to each school.

house. The assignment process is an algorithm in which each student is assigned to a nearby school, and the assignment is made by rolling a weighted die. Each side of the die represents one nearby school, and the probability of rolling a school (i.e., being assigned to a school) depends on both proximity and capacity. To the extent possible, the assignment algorithm ensures that once the assignment process is complete and all students have been assigned a school, schools are neither under-enrolled nor over-enrolled. The probabilities will tend to assign kids to schools that are closer to their houses. Living closer to a particular school makes it more likely to be assigned to that school but does not guarantee assignment there, and each school is populated by students who all live nearby.

We can think about the Soft Neighborhood Model as a sequence of four steps:

1. **Set Capacity Constraints**: Establish school capacities by grade level,[3] after assigning children needing pre-assignment.

2. **Seed Probabilities by Proximity**: For each new assignee, calculate the proximity-based probability of attending each nearby school.

3. **Balance Probabilities with Capacity Constraints**: Modify the seed probabilities from Step 2 so that the expected number of students at each school equals its target capacity from Step 1.

4. **Assign Students**: Assign children to schools using an algorithm that enforces each

[3]The *capacities* or *target enrollments* are the number of open seats in each grade level at each school.

school's capacity constraint, but otherwise is more likely to place them closer to their home.

Note that it is expected that enrollments will shift over time as families move, etc., and that these shifts will introduce imbalances. The framework keeps the system balanced to the extent possible given the actual students entering the system each year, where they live, and how much space there is at each school.

Now we will look at each step of the model in more detail.

Step 1: Set Capacity Constraints For each school, the district must set per-grade target enrollments, or capacity constraints. For example, the district might decide it wants 3 sections of kindergarten with 25 students each at Duniway Elementary; then the capacity constraint for kindergarten at that school would be 75. Certain students, for example those continuing on from the previous grade, and co-enrolled siblings entering the system for the first time,[4] are pre-assigned to their respective schools, so targets need only account for brand new enrollments.

Step 2: Seed Probabilities by Proximity Given the set of students needing to be assigned to a school, the next step is to seed, or initialize, the probability that each student is assigned to a certain school. It is during this step that we formalize the notion that families should be assigned to schools that are closer to their homes whenever possible. Let's revisit the example from Section 2.1 in which an incoming kindergartner living near NE 22nd and Mason needs to be assigned to one of Sabin, Alameda, Vernon, King, Irvington, or Beverly Cleary. Since Sabin is the closest school (0.3 mi), it will have the highest initial probability. Alameda, which is slightly farther away (0.4 mi), will have a smaller initial probability. Irvington and Beverly Cleary, which are farthest away (1.0 mi), will have the smallest initial probabilities.

Step 3: Balance Probabilities with Capacity Constraints The key idea here is to modify the probabilities from Step 2 to account for the capacity constraints from Step 1. At the end of Step 3, we will be able to run the assignment algorithm in Step 4 and avoid over-enrollment at any particular school, while still preserving the proximity-based preferences from the preceding step as much as possible.

A simple example illustrates why we need this step and how it works. Figure 5 depicts a city with two schools, Green School and Purple School, and 200 kindergarten students waiting to be enrolled. Green School has room for 80 kindergartners; Purple School has space for 120. The students all live in one of two large apartment buildings. 100 students live in a building right in the middle between the two schools. The other 100 students live in the second building located east of Purple School.

If we use the proximity-based probabilities alone to do assignments, then we are likely to overenroll Green School and underenroll Purple School. The kids who live in the middle

[4]Families should be allowed to decline the automatic assignment of co-enrolled siblings and allow a younger sibling to be randomly assigned if they want to for some reason; no harm in allowing that. Note, however, that the intent is not to allow older siblings to join younger siblings at a newly-assigned school. That would give an unfair statistical advantage to families with more children.

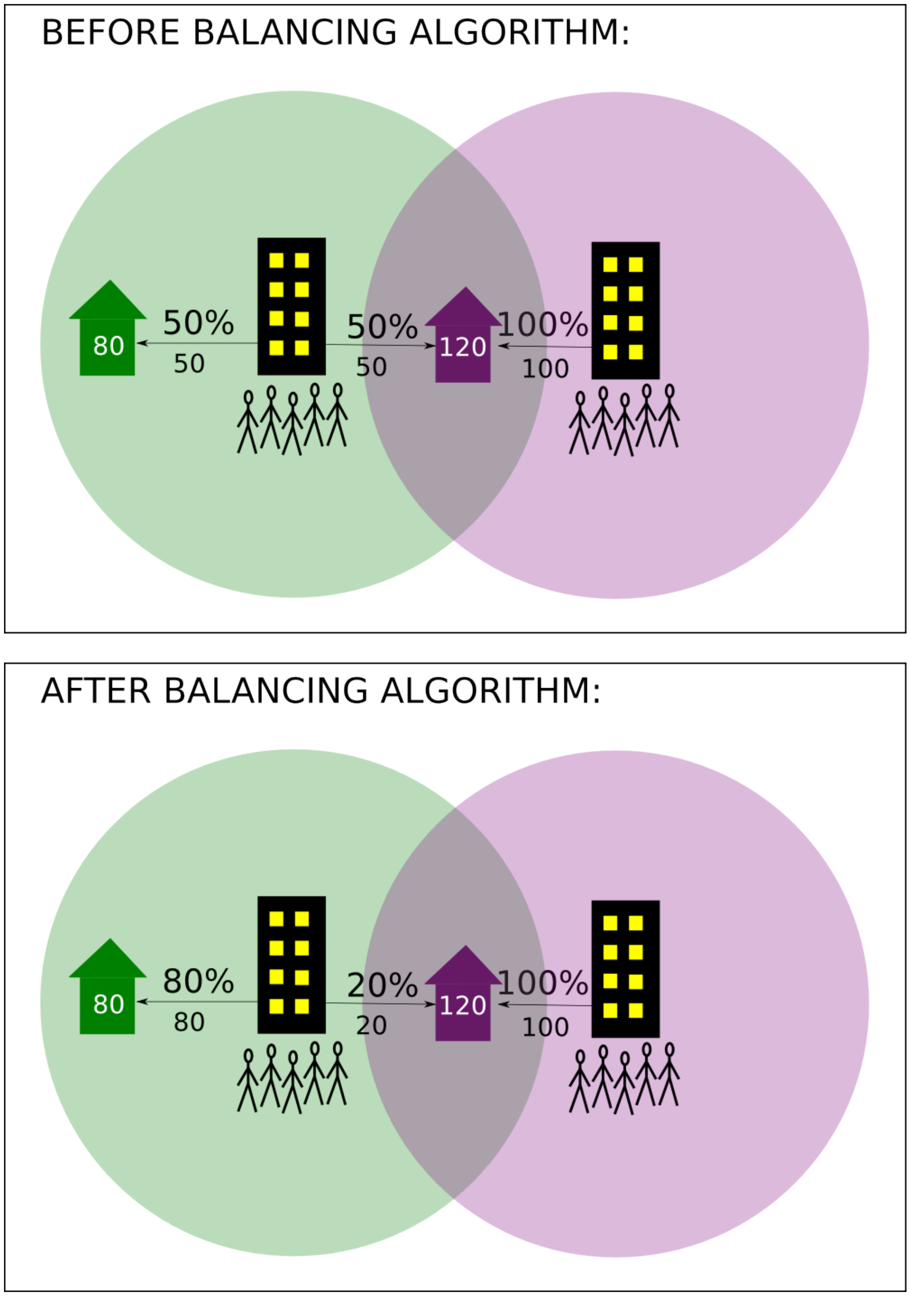

Figure 5: (top) Using raw proximity-based probabilities, the expected number of students at each school does not match target capacities. (bottom) After running the balancing step, expectations match capacities.

are equally likely to attend either school, so we expect 50 to be assigned to Green and 50 to Purple. The kids living east of Purple School are too far away from Green School, thus they all must attend Purple School. Purple School is likely to be assigned 150 students, Green only 50 (Figure 5 top). This is precisely the kind of imbalance we are trying to avoid. The balancing step nudges the probabilities in the right direction, so that we're likely to get enrollments that are aligned with capacity-based targets.

After running the balancing step, we get the modified distribution at the bottom of Figure 5 (the algorithm used to do this is described in detail in Appendix A). Now the kids who live between Green and Purple have an 80% chance of going to Green, and a 20% chance of going to Purple. This means that the expected number of kids at Green and Purple matches the target enrollments of 80 and 120, respectively.

Step 4: Assign Students The final step is to assign students to schools in a way that respects the capacity constraints as much as possible. The balanced probabilities from Step 3 are used to make the assignments. Though the idea behind the assignment algorithm is fairly simple, the algorithm itself is actually fairly complicated. We direct the interested reader to the detailed description in Appendix A.

3 Simulation: An Illustrative Example

To better understand how Soft Neighborhood assignment operates and what it accomplishes, it helps to step through an example. In this section, we demonstrate how the model behaves in an illustrative scenario. The following simulation shows the assignment of kindergarten students in a four-school district over two years. While the population shifts from Year 1 to Year 2, the capacities of the schools are assumed to be constant. In a hard boundary framework, these shifts in population put stress on the district, causing some schools to be over-enrolled, and others to be under-enrolled. In contrast, the Soft Neighborhood framework is a dynamic system that automatically assigns students to schools in a way that maintains appropriate sizing relative to target enrollments. Note that while the simulation shows a sample kindergarten assignment, the intent of the framework is to handle all new assignees in grades K–8[5].

Recall that the Soft Neighborhood Model can be described as a sequence of four steps. The pictures in Figures 6–9 show what happens at each step in this four-stage process as kindergartners enter the fictional Colorville Public School (CPS) system over two years. There are four schools in CPS: PinkA, OrangeB, GreenC, and BlueD. BlueD has room for 2 kindergarten classes (50 students);[6] PinkA and GreenC have room for 3 (75 students each school); and OrangeB has room for 4 (100 students). For the purposes of comparing assignments in Year 1 and Year 2, both the total capacity (300) and the number of students entering the system (290) have been held constant over the two years. This allows us to focus on how the system adapts to changes in population density.

We'll now walk through how the Soft Neighborhood Model behaves at each step in the process. Figures 7–9 show how each stage in the framework contributes to the final outcome, where students are assigned to schools subject to capacity constraints imposed by the target enrollments, and are more likely to be assigned to schools closer to their homes.

3.1 Step 1: Set Capacity Constraints

In Year 1, students are more concentrated in the southeast quadrant of the district, while in Year 2, the population shifts toward the northwest. Figure 6 shows how the population of incoming kindergartners is distributed for each year. This represents the state of the system at Step 1 when we set capacity constraints for a set of incoming new assignees.

3.2 Step 2: Seed Probabilities by Proximity

Figure 7 illustrates what the system does during Step 2 in Year 1 when it seeds the probabilities by proximity. The proximity-based probabilities are represented in the picture using colors corresponding to the name of each school. Each colored square in the picture corresponds to one student (i.e., one dot). The intensity of a particular color in a student's square corresponds to the probability of attending the school with that name. Figures 7(a)

[5]Recall that *new assignees* are students who need to be placed at a PPS school, for example, because they moved or because they are entering the system for the first time.

[6]Throughout this work, we use a target of 25 students per kindergarten section — a reasonable target for kindergarten based on conversations with PPS staff.

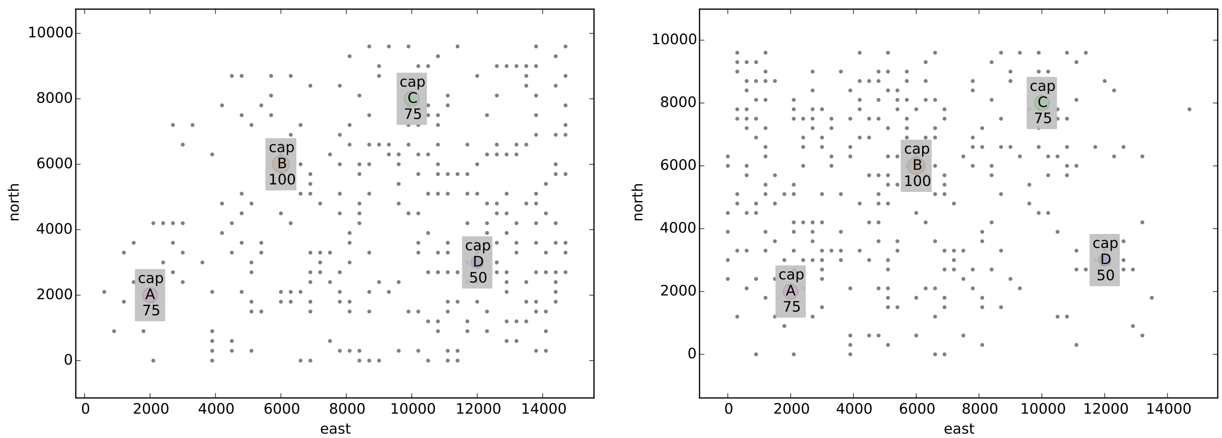

(a) Year 1: Population skewed to SE. (b) Year 2: Population skewed to NW.

Figure 6: *Set Capacity Constraints.* This depicts the state of the Soft Neighborhood system in the first step of the model. The gray rectangles indicate the location of the four schools, annotated with their target kindergarten capacities. PinkA and GreenC have seats for 75 kids; OrangeB has room for 100 kids; BlueD has room for 50 kids. The total capacity of this district is thus 300 students.

Each gray dot represents a student needing to be assigned to a school, and both Year 1 and Year 2 have 290 incoming students to place. However, the students are not uniformly distributed, and the distribution changes from one year to the next. In Year 1, the students are concentrated toward the SE; in Year 2, they are concentrated toward the NW.

and (b) show the blended probabilities for all four schools at once. Squares that are more monochromatic (blue, pink, orange, or green) signify that there is one school with very high probability in that student's neighborhood. Squares that are combinations of colors signify that there are two or more schools with significant probability. Based on distance alone, a student living between OrangeB, GreenC, and BlueD is equally likely to attend any of those three schools, independent of their capacity or of the distribution of the population.

For clarity, in Figures 7(b) and (c) we show only the probability of attending BlueD. The squares of the students in the district's southeast quadrant, closest to BlueD, are more intensely blue. As we move toward the center of the district, the blue fades. This gradient signifies the likelihood of attending BlueD School, according to the proximity-based probabilities: students who live close to BlueD are more likely to attend that school.

The numbers in the gray rectangles in Figure 7 tell us the number of students expected to attend each school according to the seed probabilities (top), and how over- or under-enrolled each school is expected to be (bottom). Note that in Year 1, because there are more students living near BlueD, we expect BlueD to be overcrowded according to the proximity probabilities. Likewise, because there are fewer students living near BlueD in Year 2, we expect it to be under-enrolled. Under- and over-enrollments for both years will be corrected in the next step.

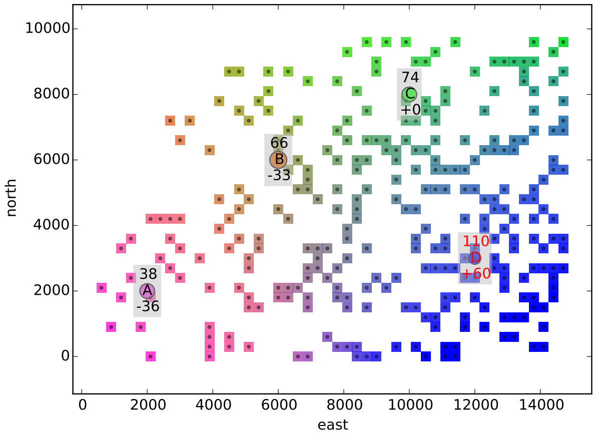

(a) Year 1, all schools

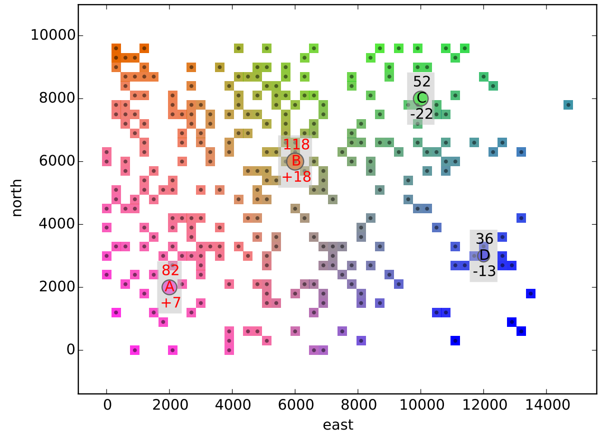

(b) Year 2, all schools

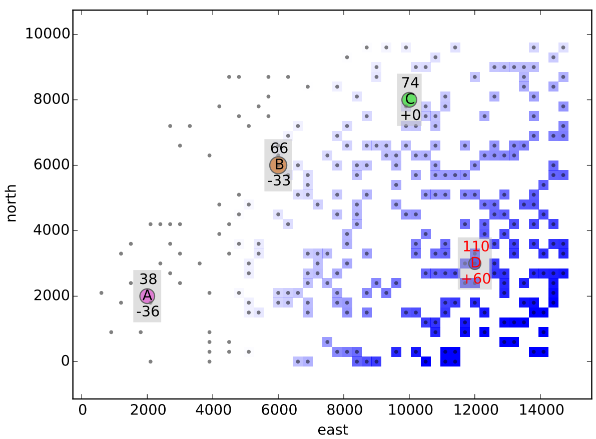

(c) Year 1, BlueD only

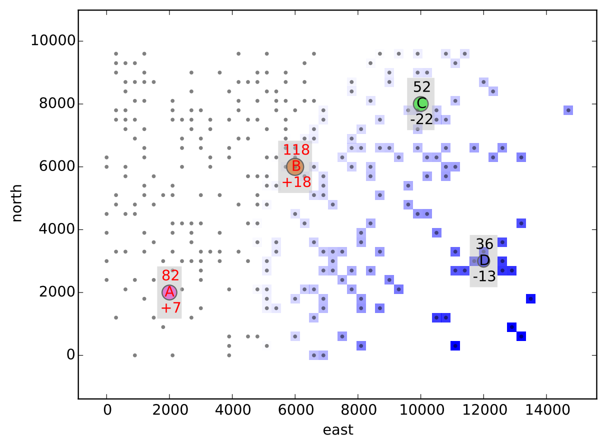

(d) Year 2, BlueD only

Figure 7: *Seed Probabilities by Proximity.* This is the result of Step 2. In (a) and (b), each colored square represents one student in Year 1, and the color is a blend of the schools' colors, weighted by the probabilities. Students in the SE corner of the district (intensely blue) are highly likely be assigned to BlueD, whereas those living between schools (blended colors) are equally likely to be assigned to one of several schools. In (c) and (d), for clarity we show only the probabilities of attending BlueD.

The schools are now annotated with two numbers. The top number is how many students we expect to assign to the school given the proximity-based probabilities. The bottom number is the expected over- or under-enrollment relative to the original target capacity. In Year 1, BlueD is expected to receive 110 students, when it only has room for 50. PinkA and OrangeB are expected to be severely under-enrolled. In Year 2, BlueD and GreenC are expected to be under-capacity, and PinkA and OrangeB are expected to be over-capacity. This will be corrected for both years in the next step.

3.3 Step 3: Balance Probabilities with Capacity

In Step 3, we bring the probabilities into alignment with target capacities. Figure 8 depicts the changes to the student probabilities that result from executing Step 3. Comparing Figures 8(a) and (b) with Figures 7(a) and (b), we see that the schools that were over- or under-filled in the previous step now have expected population sizes that are well-aligned with their target capacities. In Figure 8(c), we see that in Year 1, BlueD is at capacity with an expectation of 49. It is now much less likely (though still possible) for students who live between OrangeB, GreenC, and BlueD to be assigned to BlueD. In contrast, in Year 2 (Figure 8(d)), we see that the balancing step has slightly increased the likelihood of attendance at BlueD for students living near the center of the district. This slight increase in probability brings the expected enrollment at BlueD into alignment with its target capacity.

3.4 Step 4: Assign Students

Finally, we see what happens after Step 4, when we make assignments. In Figure 9, each square is now a single, solid color corresponding to the school to which that student has been assigned. The numbers in the gray rectangles now denote the size of the kindergarten class at each school (top) and the number of unfilled seats (bottom). On the left-hand side of the figure, we see assignments over all schools. On the right-hand side of the figure, we see how Blue's student population is distributed. Year 1 results are on the top, and Year 2 results are on the bottom. Crucially, no school is over-enrolled, and empty seats are distributed over all schools. Furthermore, the student populations are blended: there are no dividing lines that separate Blue students from Orange students or Pink students from Green students.

4 Comparison with Other Models

Many US school districts have stopped using the kind of deterministic neighborhood-based framework that Portland uses today (e.g., [11, 9, 10, 13]). Proponents of alternative frameworks such as *School Choice* often cite equity as a primary reason for rejecting hard-boundary neighborhood assignment:

> ...neighborhood-based assignment eventually leads to socioeconomically segregated neighborhoods, as wealthy parents move to the neighborhoods of their school of choice. Parents without such means have to continue to send their children to their neighborhood schools, regardless of the quality or appropriateness of those schools for their children. [3]

The School Choice movement is an alternative that attempts to counteract the socio-economic and racial segregation that has resulted from discriminative social, educational, and housing policies [8, 7, 5]. Like School Choice, the Soft Neighborhood framework represents an alternative to deterministic neighborhood models, though it is not a School Choice model. Understanding a little about School Choice will help to clarify the differences between the two approaches.

In this section, we will describe several models: the hard-boundary neighborhood model, a pure lottery model, and several variants of School Choice models. We will also consider

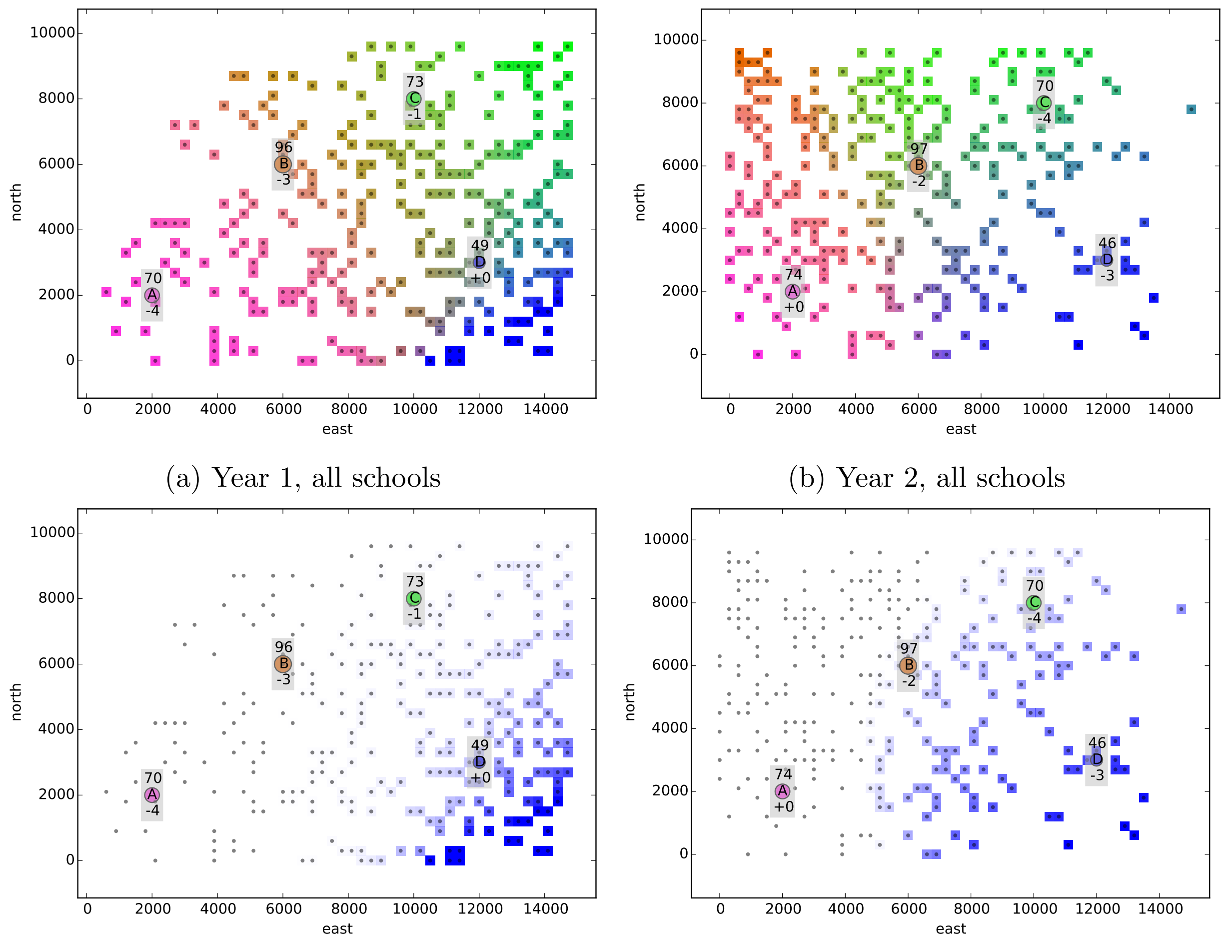

(a) Year 1, all schools

(b) Year 2, all schools

(c) Year 1, BlueD only

(d) Year 2, BlueD only

Figure 8: *Balance Probabilities With Capacity.* This is the result of the capacity balancing step. As in Figure 7, the colors reflect each student's probability of being assigned to the four schools, and (c) and (d) show only BlueD, for clarity. The balancing step shifts the probabilities around, minimally, until no school is expected to be over-capacity. Compared to Figure 7(c), students in the far southeast corner are still very likely to be assigned to BlueD, but students living between BlueD, OrangeB, and GreenC are now less likely to be sent to BlueD. The numbers on the schools reflect these shifts. BlueD is now expected to be at-capacity, and the other three schools are expected to be very slightly under-capacity — which is reasonable since there are only 290 students to fill 300 seats. The changes in Year 2 (c) are less dramatic, but the balancing step has increased the overall probability of attendance at BlueD enough to bring the expected number of students into alignment with its capacity.

how two hypothetical students would be assigned to kindergarten in each case. The two students in our example live across the street from one another, one at 4541 NE 22nd Ave. and the other at 4540 NE 22nd Ave. These addresses lie on opposite sides of a current PPS

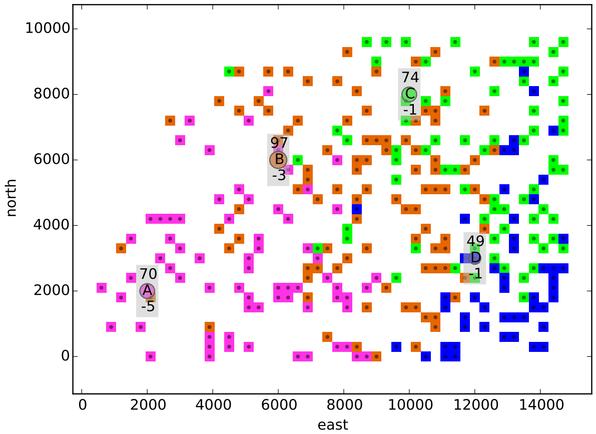

(a) Year 1, all schools

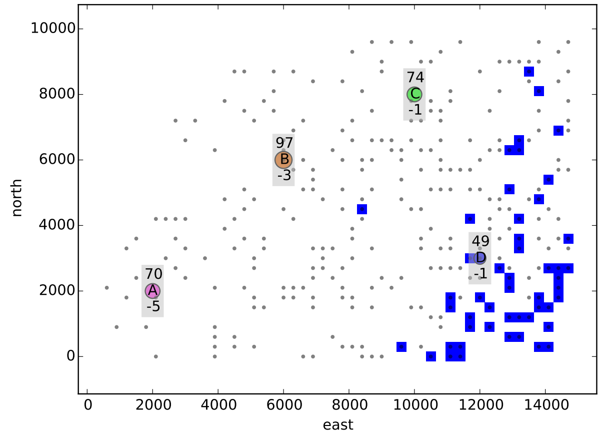

(b) Year 1, BlueD only

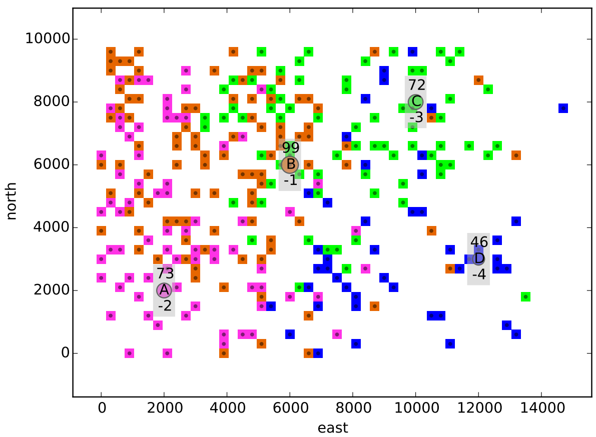

(c) Year 2, all schools

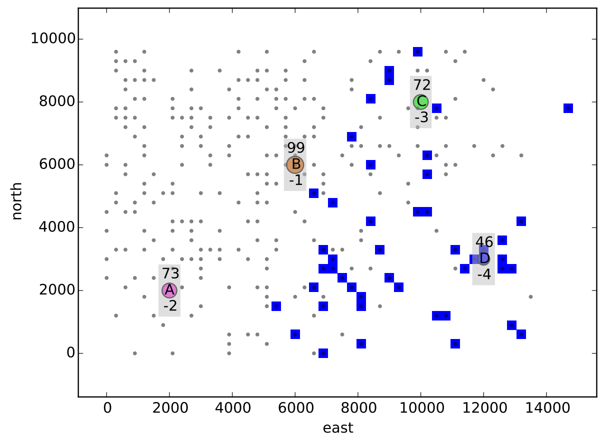

(d) Year 2, BlueD only

Figure 9: *Assignment.* This is the result of the final assignment step. The color of each square indicates the school to which that student has been assigned. In the final tally, no school is over-capacity; every school is very slightly under-capacity, consistent with having fewer students than seats. Students tend to live near other students assigned to the same school. Also, note the high degree of mixing: students from any given area are generally assigned to a mix of different schools, with no clear boundary isolating them from one another.

cluster boundary (Grant and Jefferson/Madison dual assignment).

In the Soft Neighborhood Model, the set of nearby schools for these two children children is the same because they live so close to each other (Figure 10). If we populate this set using the 3-School rule from Section 2, then the candidate schools are Sabin, Vernon, Alameda, King, and Irvington. Since they live so close to each other, the two students have nearly the same odds of attending each of these schools. That is, the likelihood of the model assigning student A to Sabin is the same as the likelihood of assigning student B to Sabin. Because the set of nearby schools is similar for all children living near each other, there will always be children living nearby who are assigned to the same school.

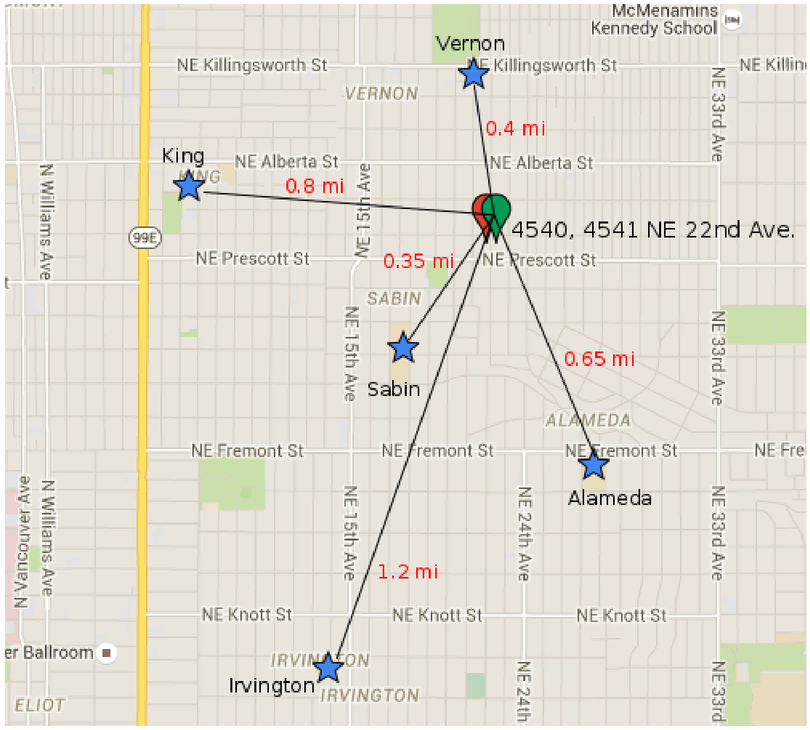

Figure 10: In the Soft Neighborhood Model, the two students living across the street from each other will both be assigned to one of five nearby schools.

4.1 Hard-Boundary Neighborhood Model

In a Hard Boundary model like the one PPS currently uses, boundaries are used to designate neighborhood attendance areas. For a given grade range, everyone residing within a certain boundary pretty much goes to the same school.[7] The two students living across the street from one another in Figure 11 will never be assigned to the same schools because they span a cluster boundary. Another way to think about it is that the probability that they go to the same school is zero. For the most part, everyone who lives on the same side of the boundary goes to the same schools with 100% probability.

4.2 Pure Lottery Model

A Pure Lottery model is geographically unconstrained and may assign children to any school in the district (Figure 12). The two children in our example have some chance of assignment to the same school; however, because of the number of possible schools, the likelihood of this happening is very small. The same holds for the children nearby: everyone will tend to go to a different school. Also, the students may have to travel very far away from where they live to get to school. A Pure Lottery model eliminates hard boundaries and can allow districts to fill schools with populations that match the district-wide demographic averages. It also offers the district a high degree of flexibility for balancing enrollments, since a child can be placed at any school. However, these advantages come at the expense of neighborhood

[7] There are exceptions, but this is the norm.

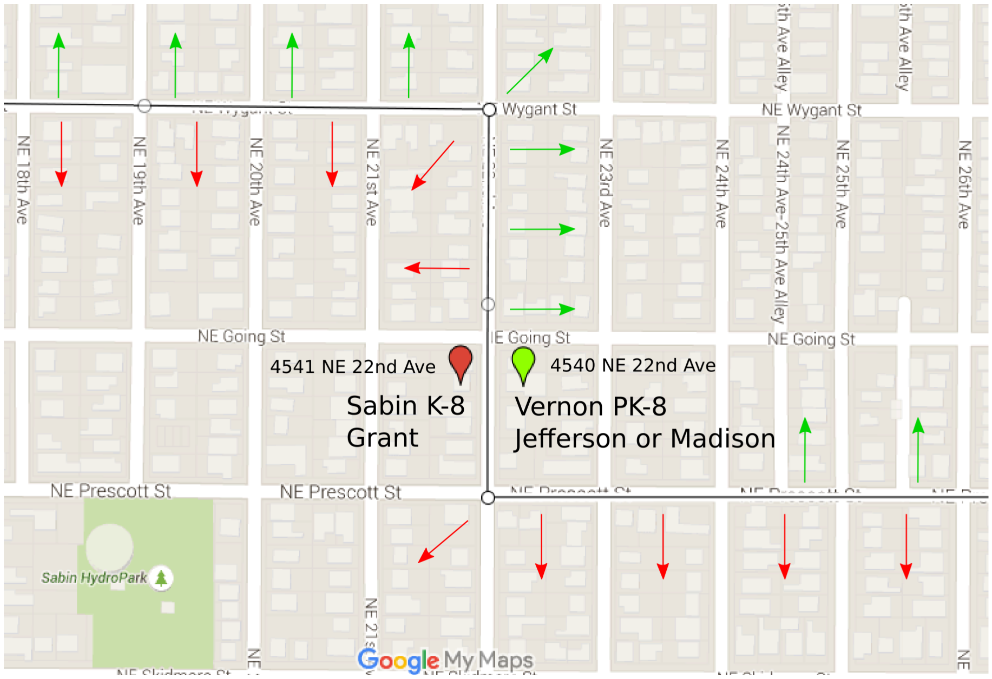

Figure 11: In a Hard Boundary model, the two students are always assigned to different schools because they live across a boundary. Only students on the same side of the boundary can go to the same school.

and proximity. Travel times make a Pure Lottery model impractical in Portland. However, we include it in the discussion as a point of contrast with PPS's Hard Boundary model. The two models together (Pure Lottery and Hard Boundary) represent two extremes of the assignment spectrum.

4.3 School Choice

The School Choice movement gained momentum as an alternative to both Hard Boundary models and desegregation plans (see [6] for an overview). There are many variants of School Choice, but in the canonical version, families rank order their preferred schools, and the district places students in schools, taking into account these preferences as well as the district's priorities.

There exists a large body of academic literature, particularly in the field of economics, that has emerged from and bolstered the School Choice movement. The academic literature tends to focus on "mechanism design" — the algorithms used for making assignments — and the properties of proposed mechanisms. Researchers usually use three criteria to evaluate assignment mechanisms:

- Pareto efficiency: A Pareto efficient mechanism results in assignments that cannot be improved without making at least one student worse off.
- Stability: In a stable mechanism, there are no assignment outcomes with any *blocking pairs*. A blocking pair is a ⟨Student A, School A⟩ pair where Student A prefers School

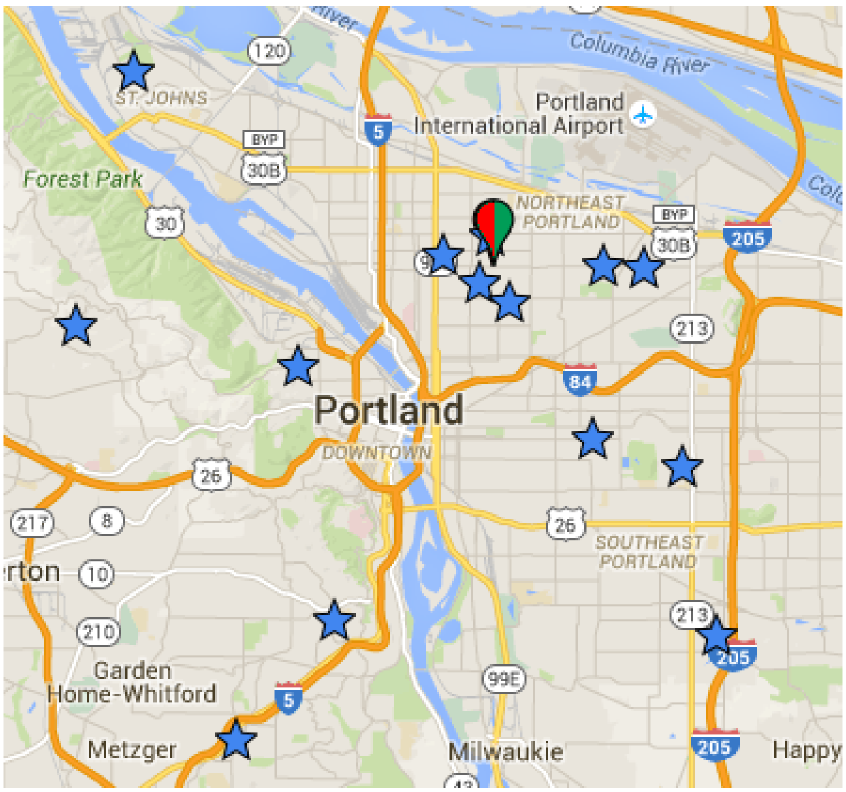

Figure 12: In a pure lottery model, students may be placed anywhere in the district. Some (but not all) district schools are shown in the picture.

A to her assigned school, and Student B, who has lower priority than Student A at School A, has been assigned to School A.

- Strategyproofness: Strategyproof mechanisms encourage truthful rankings from all students.

Pathak [6] summarizes very well the three primary mechanisms associated with School Choice: the student-proposing deferred acceptance mechanism [4], the top trading cycles mechanism [1, 14], and the Boston mechanism [1]. It's important to understand that each assignment mechanism is different: student-proposing deferred acceptance is stable and strategyproof; top trading cycles is Pareto efficient and strategyproof; and the Boston mechanism is Pareto efficient but neither stable nor strategyproof. Choice-based programs have been characterized as conferring strategic advantages to better-educated and wealthier families, or even to resegregating school districts [15], but each individual program's chosen assignment mechanism and implementation is instrumental in how it functions.

The Soft Neighborhood Model differs from School Choice models along several dimensions. The most obvious difference is that in the Soft Neighborhood Model, families do not rank preferences. In the Soft Neighborhood Model, proximity is used to seed an initial probability distribution, and then that distribution is modified so that the expected number of students at each school equals its capacity. Other differences are that the probabilities used in the Soft Neighborhood Model are real numbers, whereas in School Choice models, usually only ordinal ranking are considered. Another major difference is in the math underpinning the two models. Choice models use an optimal-matching algorithm like the ones we described above. Soft Neighborhood assignment uses probabilities in order to eliminate determinism between an address and a particular school, but does so in a way that maintains a strong notion of neighborhood.

5 The Soft Neighborhood Model in PPS

In a previous section, we demonstrated how the Soft Neighborhood Model adapts to population shifts in a synthetic data set. In this section, we establish a framework for applying the Soft Neighborhood Model to seven years of historical PPS data provided by the district. The data set contains two files, the STUDENTS data file and the SCHOOLS data file.[8] The STUDENTS data file contains student data from 2008–2015. Among other things, this file contains the anonymized home addresses of K–8 PPS students during that time frame, enabling us to explore how variations of the Soft Neighborhood Model would have assigned those students, and compare those assignments to historical Hard Boundary assignments. Unfortunately, there are flaws in the data that prevent us from running the Soft Neighborhood Model on grades other than kindergarten. Thus, it is unclear how well the model maintains balanced enrollments as each kindergarten class matriculates through the primary and middle grades. It is crystal clear, however, that the model nails enrollment balancing at the kindergarten level, balancing section sizes across all years and at all schools.

The rest of this section lays out exactly how we've run our experiments, how we configure each assignment model, and how we evaluate the models. In Section 6, we report our experimental results; and in Section 7, we wrap up and provide student-school assignment graphs which help to visualize what these models do.

5.1 Models

In these experiments, we consider eight assignment models: a Hard Boundary model, a Single Closest School[9] model, and six tunings of the Soft Neighborhood Model. We run each of these models in a *with transfers* and a *no transfers* setting. The *with transfers* setting assigns historic transfer students to their historic transfer school. The *no transfers* setting ignores the transfer school assignment of historic transfer students, instead assigning them just like any other student. We consider both variants because the number of neighborhood transfer students is expected to decrease going forward as a result of the elimination of the transfer lottery beginning with the 2015–16 school year [2].[10] Each of the models under consideration assigns students differently:

Hard Boundary Model In the *no transfers* setting, the Hard Boundary model assigns all students to their historic PPS neighborhood schools as listed in the STUDENTS data file, regardless of their actual historic assignment. In the *with transfers* variant, all students are given their actual historic assignments, which includes transfers to other neighborhood schools.

[8] See Appendix B for details. This data is publicly available on our website: `http://www.softneighborhoodmodel.org/`.

[9] In Single Closest School, each student is assigned to whichever school is closest to their home.

[10] PPS continues to grant transfers via petition. The STUDENTS dataset provided to us does not distinguish between lottery and petition transfer students.

Single Closest School Model The Single Closest School model assigns each student to the school which is closest to her home, as measured by driving distance.[11] In the *no transfers* setting, all assignments are made in this way. In the *with transfers* variant, historic transfer students are assigned to their historic transfer school, and all remaining students to the closest school.

Soft Neighborhood Model We consider six tunings of the Soft Neighborhood Model (SNM). Each version corresponds to a different configuration of the proximity function. In early versions of this paper, we used the *3-School Rule* defined in Section 2 to configure the Soft Neighborhood Model for experiments on PPS data. As of Version 3.0 of this document, we have replaced that *3-School Rule* with a simplified and more general function.[12] Using the new proximity function, we can vary (1) the number of nearby candidate schools considered for assignment, and (2) the intensity of the model's initial preference for closer schools. We consider models with a *strong* initial preference for closer schools, a *moderate* initial preference for closer schools, and an initial preference that is *equal* for all candidate schools. We experiment with candidate sets of size 2 and 3:

Model	**Number of Candidate Schools**	**Initial Preference for Closer Schools over Farther Schools**
SNM, 3-School, Strong	3	strong
SNM, 3-School, Moderate	3	moderate
SNM, 3-School, Equal	3	equal
SNM, 2-School, Strong	2	strong
SNM, 2-School, Moderate	2	moderate
SNM, 2-School, Equal	2	equal

We run each of these SNM versions in a *with transfers* and *no transfers* setting, where in the *no transfer* setting, historic transfer students get their assignment from their historic transfer school instead of the Soft Neighborhood Model.

5.2 Metrics

In this section, we define five metrics for evaluating the different assignment models. Three are fairly self-explanatory: enrollment balancing, travel distance from home to school (using

[11] Ties are resolved by randomly selecting from equidistant closest schools.

[12] The new proximity function is parameterized by (1) a limit of n_c closest schools, and (2) the maximum possible ratio ρ_{max} between the seed weights of the closest and farthest candidate school for any student. Under this function, the n_c schools closest to a student become the candidate schools. The total number of candidates may be greater than n_c if there are ties for n^{th} place. ρ_{max} is the ratio between the weight of a school with zero distance and the weight of the farthest candidate school. Thus, a higher ratio yields a stronger seed preference for closer schools. In our experiments, the *Strong* SNM versions set $\rho_{max} = 1000 : 1$; the *Moderate* SNM versions use $\rho_{max} = 10 : 1$, and the *Equal* SNM versions have $\rho_{max} = 1 : 1$.

driving distances), and distance to closest classmate. The remaining two — local assignment diversity and capture linkage — measure social network effects and require more explanation.

Enrollment Balancing This measures how well a model hits enrollment targets. We are generally aiming for *even* enrollment across the district, trying to get equal sizes for each section in the district, and SNM has been configured to seek this goal in these experiments. For a given year (e.g., 2014), we assess this by looking at the probability mass function (PMF)[13] of actual enrollments (sizes) of each kindergarten section in the district. Section-level enrollments are precisely determined at the balancing step, so we don't need to perform any actual assignments to measure this.

The ideal balanced enrollment target is basically the mean — the number of students divided by the number of sections — and over multiple years this drifts as the total population (and section count) changes. For cumulative results for 2008–2015, we instead look at deviation (difference) of the actual section sizes from the mean for their respective years.

Travel Distance This measures how far students have to travel from their home to get to school. In this document, the Travel Distance metric has been implemented using driving distances.[14] To obtain results for the Soft Neighborhood Model using this metric, we calculate the PMF of expected travel distances as follows: first, we look at the probabilities after the balancing step (*Balance Probabilities with Capacities*) and collect ⟨*assignment-probability, distance*⟩ pairs over all students and all schools, where *assignment-probability* is the probability of a student being assigned to a particular school. The PMF is derived by weighting each distance using its paired probability. As with *Enrollment Balancing*, this metric is computed at the balancing step and does not require making actual assignments.

Closest Classmate This measures the Cartesian distance[15] between each student's home and the home of her closest classmate.[16] Results are reported by calculating the PMF over the set of ⟨*student, closest-classmate*⟩ distances. This needs to be performed on actual assignments (not just expectations like the *Travel Distance* metric), so we perform fifteen Soft Neighborhood Model assignments for each SNM version and aggregate the distances.

Local Assignment Diversity This measures how many schools are well-represented among all students living within 1000 feet (Cartesian) of a given student's home, after school assignments have been made. In other words, how diverse are the school assignments near any student? If there were exactly 5 students from School-A and 5 from School-B living within some 1000-foot radius, the diversity would yield exactly 2, meaning that two schools were effectively well-represented in that area. If instead there were 2 School-A students and

[13] I.e., the distribution, represented as a histogram.

[14] Note that prior to Version 3.0 of this document, we used Cartesian distances to compute this metric. The original tables using Cartesian distances are available in Appendix C.

[15] We use Cartesian distances for the Closest Classmate metric because it is impractical for us to compute driving or walking distances between all pairs of students.

[16] "Classmate" is defined here as another student assigned to the same school and in the same grade, since we can only analyze kindergarten.

8 School-B students, the result would be ∼1.65, showing that more than one school was represented, but two schools were not equally-well represented.

Intuitively, this metric is associated with system gameability, where lower assignment diversity scores indicate more susceptibility to gaming. If the assignment diversity score in some region is very close to 1.0, then the probability, in that region, of being assigned to the one well-represented school is effectively 100%. If the score is two, and two schools are equally well-represented, then the probability of attending each school is effectively 50%. Assignment diversity is also an indicator for communication between school communities: higher diversity around a family's home means that a family is more likely to discuss (with neighbors) the goings-on at more than one nearby school. Thus, higher assignment diversity helps foster less insular and less isolated neighborhoods.

To implement this metric, we calculate the ${}^1D$ diversity index[17] of the set of schools assigned to all students living within a circle of radius 1000 ft of each student. Results are reported by calculating the PMF over the set of diversity values for all students. Like the *Closest Classmate* metric, this one also needs to be applied to actual assignments (not just expectations). We again aggregate the values from a set of fifteen sample assignments.

Capture Linkage Capture linkage (Figure 13) is a metric applied to each school, indicating how many other schools share some part of that school's capture population. *Capture population* is the set of students that may be assigned to a school. If a model may assign some student to either School-A or School-B, then School-A is *linked* to School-B, and vice-versa. The graph of connections between students and schools tells us a lot about how school communities are woven together (or not) by an assignment model.

Here, we boil that structure down to a single number per school. As a number, capture linkage is a diversity metric. For a given school, it collects all the weighted student-school links from that school's capture population, and counts how many schools are well-represented in that set.[18]

Imagine a school with a capture linkage of zero (Figure 14): it draws its students from a population that feeds to only that school and no other. This system is gameable, since anyone entering that population (e.g., by moving to a particular region of the city) is guaranteed to be assigned to that school. Conversely, no other school has any students from that population; as far as the assignment network goes, that school is an island unto itself. Whatever happens to that population — how its demographics change over time — is not felt at any other school. Likewise, that school is insulated from the changes happening to the populations and schools which surround it.

Compare that to a school with a capture linkage of nearly 5.0 (Figure 15): it draws its students from a population that also feeds at least five other schools. Some small regions might be guaranteed assignment to that school, but *most* of the population may also be assigned to at least one of the five other schools. No subpopulation gaming the system can

[17] Also known as exponentiated Shannon entropy; see https://en.wikipedia.org/wiki/Diversity_index.

[18] Specifically, it's the ${}^1D$ diversity (exponentiated Shannon entropy) measured over all edges emanating from students who are connected to a given school, where edges are weighted by the post-balancing assignment probability. We then subtract 1 to discount the connections back to the given school, yielding the linkage to *other* schools.

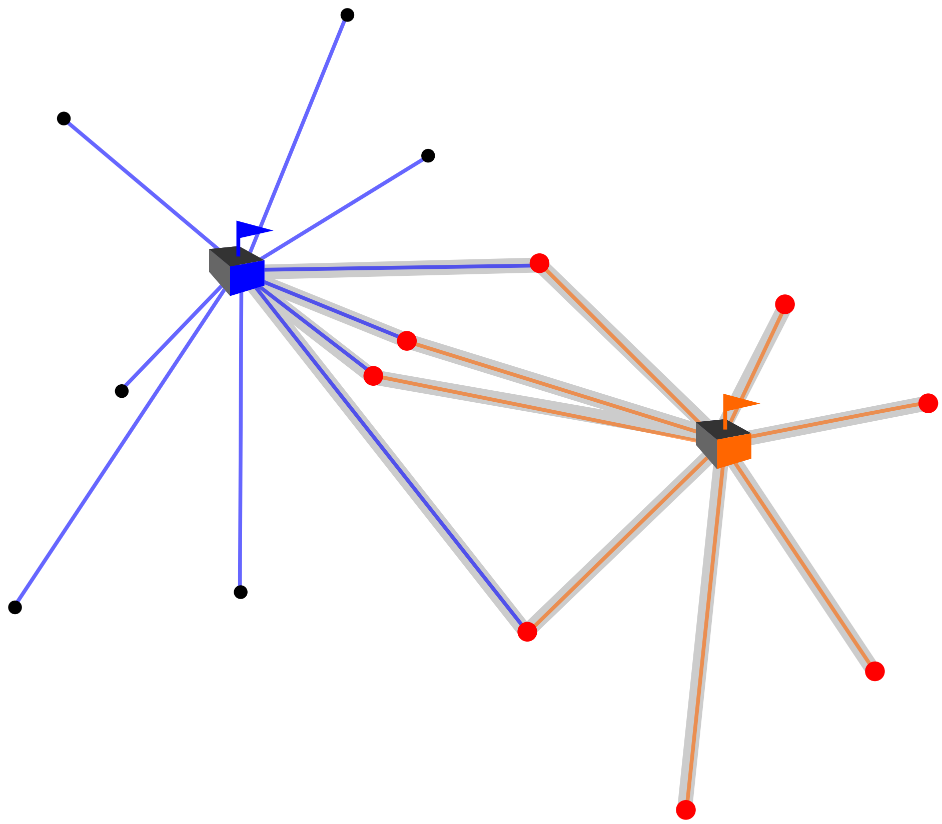

Figure 13: Illustration of *capture linkage*. Dots represent students; edges (lines) show possible assignments to schools. The red dots indicate the *capture population* of Orange School, all the students who might attend that school. Capture linkage is calculated by counting the edges connected to the capture population (shaded lines). Orange School is *linked* to Blue School by the four students that might be assigned to either. If those four students each have a 50-50 chance of attending either school, the capture linkage of Orange School works out to 0.75 — i.e., it is linked to one other school, but not as well as it could be.

dominate the school, and there cannot be a dominant population that is insulated from those other schools. The school shares its fate with those five other schools.

For Hard Boundary models, the capture linkage for any school is basically zero, by definition. A student's fate is predetermined by his address; there is only one school to which he will be assigned. Every school draws its students from a population which is independent of every other school.[19] As such, all Hard Boundary models are obviously terrible by this metric. Beyond making this crystal clear, the purpose of the capture linkage numerical results is to compare the other assignment models to one another.

[19] There may be some cross-pollination that occurs due to petition or lottery transfers — but that is not a property of the *neighborhood school model* itself. And, since PPS is specifically trying to reduce or eliminate neighborhood transfers, it is not something to count on.

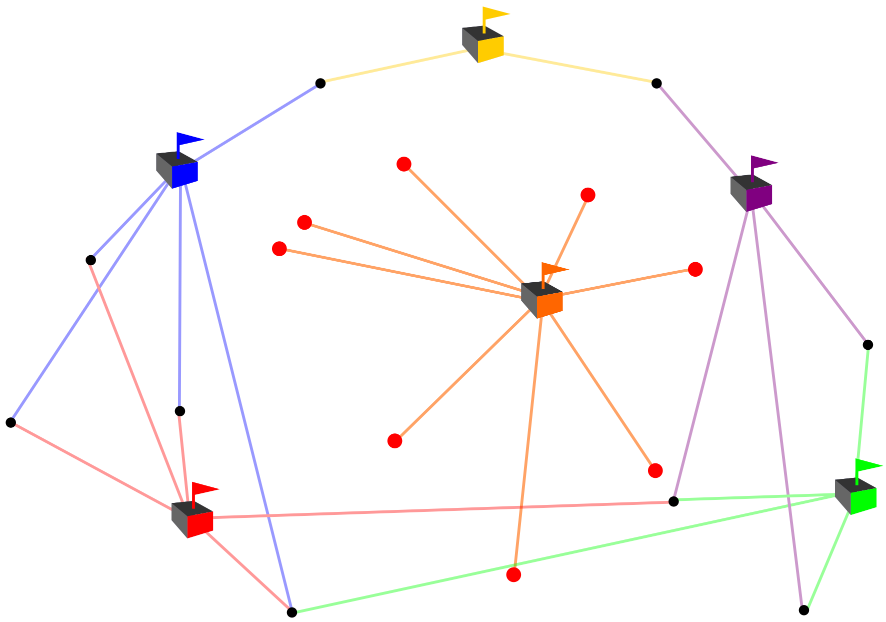

Figure 14: Example of zero capture linkage. Students who may attend Orange School cannot be assigned anywhere else. The *capture linkage* of Orange School is precisely zero. It is an island in the assignment network.

5.3 Setting Capacity Targets

Configuring the Soft Neighborhood Model requires setting target capacities for each school. A target is specified by two quantities: the desired number of sections (*section-count*) and ideal number of students in each section (*section-size*). The SNM seeks to balance out assignments to achieve these targets. If the sum of *section-count* $\times$ *section-size* is more or less than the actual population of the district, it will be impossible to do this exactly, but the SNM will still find a balance that respects the relative sizes of the targets.

In practice, we expect the district to establish these targets. However, it seems pretty reasonable to aim for having uniform section sizes across the entire district, so that is what we have done here. By assigning the same nominal *section-size* target to each school (e.g. 27 students per section), the SNM will try to balance all sections across all schools to the same size. What is left is to establish the target *section-count* for each school.

We've considered several different methods for doing this, including many which were ultimately suboptimal. For the record, we describe a few that we *did not* use first. Then, we describe the two we have used for the experiments in Section 6.

5.3.1 Section-Counts We *Are Not* Using

SCHOOLS Data File In the SCHOOLS data file, PPS included a field called `KG_HR_CNT`, or *Kindergarten Home Room Count*, which is meant to represent the number of kindergarten home rooms at each school in 2014. We spot-checked a few of these using Internet Archive snapshots of school staff pages from the 2014–15 school year and found the number of

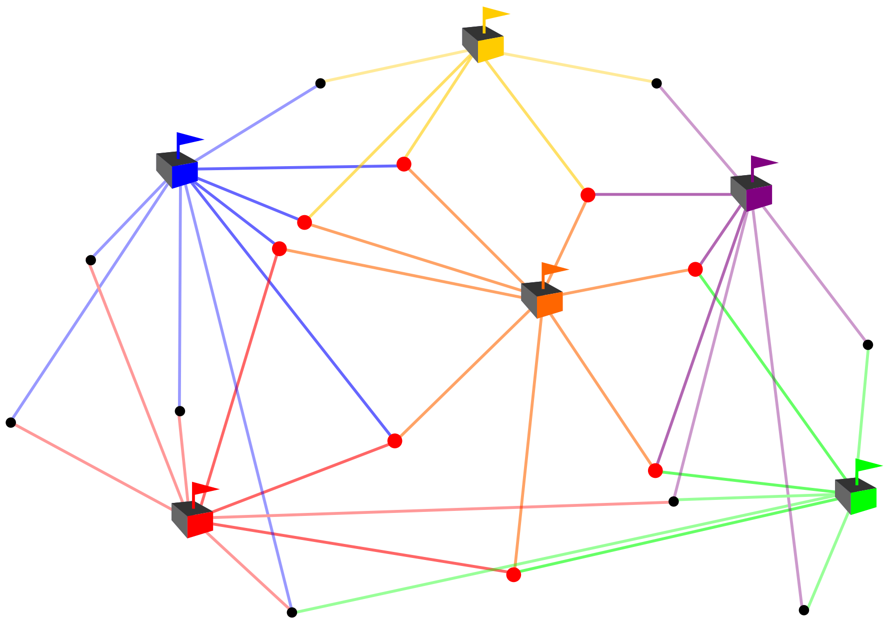

Figure 15: An example of capture linkage close to 5. Every student who may attend Orange School might also be assigned to one of two other schools. If assignment chances are 1-in-3 for all those students to each of those schools, the *capture linkage* of Orange School is 4.5. It is well-connected in the assignment network.

sections listed is often inaccurate. For example, a Grout snapshot dated Oct. 12, 2014[20] lists four kindergarten teachers on its staff page, while the `KG_HR_CNT` field lists five sections. Due to these potential inaccuracies, we have not used the `KG_HR_CNT` counts for any of our experiments.

Proportional to Historical Enrollments In early versions of this document (prior to version 3.0), we estimated section-counts from the actual historical enrollments at each school for each year. First, we defined a uniform target size of 25 students per section for all schools, a reasonable target based on conversations with PPS staff. Then, to derive a target number of sections at each school, we took the historical total kindergarten enrollment for a given year, divided by 25, and rounded to the nearest integer. For example, a school with 38 kindergartners was given a target of 2 sections since $38/25 = 1.52$. Unfortunately, the results reflect neither the actual historical section-counts nor the optimal section-counts (considering building capacity). This method produces counts which are overly optimistic, both for the historic Hard Boundary assignments and the Soft Neighborhood Model assignments. For the Hard Boundary assignments, it optimistically estimates the per-section deviation, precisely

[20] `https://web.archive.org/web/20141012214216/http://www.pps.k12.or.us/schools/grout/160.htm`, accessed 2014-01-03.

because the number of sections is calculated so as to minimize the deviation.[21] For the Soft Neighborhood Model, the method scales the number of sections to fit the nearby student population, thus making it easier for the model to get nearby students to nearby schools. Due to these problems, we are no longer using this method.

5.3.2 Section-Counts We *Do* Use

True 2014 Using a combination of community responses to our public request for accurate 2014–15 section counts at PPS schools, and the Internet Archive for the remaining schools, we have collected what is, to the best of our knowledge, the true 2014–15 kindergarten section counts at each school. Note that our working definition of a "section" is any homeroom/class led by a kindergarten teacher, including both neighborhood and focus option sections where applicable.[22] When searching the Internet Archive, we encountered four K/1 combination sections, and we counted each as one half of a section. We did not count any other types of sections as kindergarten sections, even if they contained some kindergarten students (e.g., communication behavior classrooms). When using the Internet Archive, we accessed each school's staff page from a snapshot taken between 11/2014 and 3/2015, and counted the number of kindergarten teachers. We were able to recover section counts in this manner for every school except for Irvington, which had no meaningful content on its staff page. Instead, we used the `KG_HR_CNT` number for Irvington.

We call the resulting set of community-plus-Internet-Archive counts the *True-2014* section counts and use them for our 2014–15 experiments in Section 6.1. The True-2014 section counts for each school, along with historical enrollment and resulting historical section-sizes are listed in Table 2. Without help from PPS, this method seems to be the best way we have to recover accurate section-counts for a particular year. However, it is not a practical method for obtaining section-counts over the full seven years of data. Also, it suffers from the same problem we saw with the "proportional to historical enrollments" method described above: it distorts the number of sections to fit the nearby population, and is not necessarily faithful to the actual capacity of the district's buildings. That said, these counts give us the best sense of how the various models behave in the actual 2014–15 setting people experienced.

Once we have the target section-count at each school, we estimate the target section-sizes for the district, the east side, and for the west side by dividing the total number of students by the total number of sections for each geographic region, and then rounding (Table 1).

Conservative-Optimal In the summer of 2015, PPS undertook a survey of the district's building capacities [12]. The result was a chart indicating a range of possible section counts for various grade configurations at each school. For example, Laurelhurst School, with 30 classrooms, is said to be able to comfortably house a 2- or 3-section K-5. Alternatively, it could comfortably house a 2-section K-8 or a middle school. Using this PPS-generated chart, we derive target section-counts by looking for the optimal number of sections (2, 3, or 4) for a particular school, given its historical grade configuration for a particular year. By

[21] The actual historic section-sizes vary far more wildly than we had expected. For instance, in 2014–15, we see actual section sizes as small as 14 and as large as 30 (see Table 2).

[22] We include both because the PPS-provided data does not distinguish neighborhood students from co-located focus option students but does contain both.

	# Sections (T14)	# Students	Target
All PPS	173.0	3681	21
East Side	139.5	2878	21
West Side	33.5	803	24

Table 1: Target section sizes for 2014-15 using the True-2014 (T14) section counts. Targets were obtained by dividing the total number of students by the total number of sections, and rounding the result.

School	T14	# Stu	Size	School	T14	# Stu	Size
Abernethy	4	79	19.8	Laurelhurst	3	73	24.3
Ainsworth	4	89	22.2	Lee	3	48	16.0
Alameda	5	117	23.4	Lent	3	57	19.0
Arleta	3	61	20.3	Lewis	3	57	19.0
Astor	2	57	28.5	Llewellyn	3	61	20.3
Atkinson	3	61	20.3	Maplewood	3	56	18.7
Beach	3	81	27.0	Markham	3	70	23.3
Beverly Cleary	4	90	22.5	Marysville	3	53	17.7
Boise-Eliot/Hum	4	88	22.0	CJOG	4	73	18.2
Bridger	3	77	25.7	Peninsula	3	54	18.0
Bridlemile	3	82	27.3	Rieke	3	64	21.3
Buckman	3	61	20.3	Rigler	4	102	25.5
Capitol Hill	3	89	29.7	Rosa Parks	3	54	18.0
Chapman	5	123	24.6	Roseway Hts	5	87	17.4
Creston	2	44	22.0	Sabin	3	84	28.0
César Chávez	3	56	18.7	Scott	4	65	16.2
Duniway	4	95	23.8	Sitton	3	73	24.3
Faubion	2.5	75	30.0	Skyline	1.5	38	25.3
Forest Park	3	65	21.7	Stephenson	2	48	24.0
Glencoe	3	69	23.0	Sunnyside	3	61	20.3
Grout	4	64	16.0	Vernon	3	55	18.3
Harrison Park	4	73	18.2	Vestal	3	48	16.0
Hayhurst	3	79	26.3	Whitman	2	38	19.0
Irvington	2	39	19.5	Woodlawn	3	59	19.7
James John	4	79	19.8	Woodmere	3	42	14.0
Kelly	5	116	23.2	Woodstock	4	87	21.8
King	4	65	16.2				

Table 2: Historical average section size for each school in the district in 2014–15, using actual historical enrollments and *True-2014* section counts. These same section sizes are shown in the assignment graph in Section 7, Figure 26.

"optimal number of sections," we mean that we look for the right-most "green dot" in the chart (a green dot seems to represent a configuration that will *definitely* fit the building). We only use the "yellow dots," which seem to represent a configuration that *might* fit the building, when no green dots exist. Based on anecdotal discussions in the community, even PPS's "green dot" numbers appear to be conservative estimates of the actual capacity of each building, erring on the side of fewer sections. Hence, we designate this set of section-counts *conservative-optimal*, or *Con-Opt* for short.

Table 4 shows the section-counts for each school using the *conservative-optimal* approach, compared to the *True-2014* numbers. Table 3 shows the total district-wide *conservative-optimal* section-counts for each year in 2008–2015. The numbers are mostly constant, since buildings are pretty inflexible; the changes are due to a couple of school configuration changes in that period.[23] Dividing the total population for each year by the total section-count yields target section-sizes for the district as a whole, and for the east and west sides independently. These are the sizes one would get if one had absolute freedom to balance enrollment across the respective regions. Overall, the sizes seem large, but that is expected if the section-counts are conservatively small. The *conservative-optimal* section-counts give us the best sense of how well the various models respect the actual capacity of buildings. We use the *conservative-optimal* section-counts for the cumulative experiments aggregated over all seven years of data (Section 6.2).

5.4 Experimental Methodology

We've made a few methodological decisions along the way that should be taken into consideration when interpreting our results:

Kindergarten only The experiments in Section 6 have been run on kindergarten data only. Given the aforementioned limitations in the STUDENTS data set (see Appendix B), there is no way to run meaningful simulations for any grades beyond kindergarten.

Neighborhood programs only In our historical analysis, we consider every student attending a neighborhood school to be enrolled in a neighborhood program. The data set contains anonymized information about all students who attend each neighborhood school, but it does not provide any way to distinguish between those who attend neighborhood programs and those who attend co-located focus option/immersion programs. Also, the data set does not contain any information about students attending district-wide focus option schools. This means that the simulations are unable to account for the students who attend these schools. For instance, experiments in a *no transfers* setting are missing all of the students attending lottery-only, stand-alone schools.

No siblings Because sibling relationships are not available in the data set (see Appendix B), every kindergarten student to be enrolled is considered a non-sibling in our simulations, i.e., no students are pre-assigned to specific schools due to sibling preference.

[23]Humboldt closed in 2012, and Chief Joseph merged with Ockley Green in 2013.

Year	# Sections (C-O)	# Students	Target
	All PPS		
Year	**# Sections (C-O)**	**# Students**	**Target**
2008–09	130	3601	28
2009–10	130	3726	29
2010–11	130	3657	28
2011–12	130	3715	29
2012–13	129	3884	30
2013–14	128	3843	30
2014–15	128	3681	29
	East Side		
Year	**# Sections (C-O)**	**# Students**	**Target**
2008–09	105	2850	27
2009–10	105	2958	28
2010–11	105	3027	29
2011–12	105	2955	28
2012–13	104	3193	31
2013–14	103	3141	31
2014–15	103	2878	28
	West Side		
Year	**# Sections (C-O)**	**# Students**	**Target**
2008–09	25	751	30
2009–10	25	768	31
2010–11	25	630	25
2011–12	25	760	30
2012–13	25	691	28
2013–14	25	702	28
2014–15	25	803	32

Table 3: Target sections sizes for each year, using the *conservative-optimal* (C-O) section counts derived from PPS's building capacity study [12]. Targets were obtained by dividing the total number of students by the total number of sections, and rounding the result. This method of deriving section counts is conservative, meaning that the target section sizes are large. The total number of sections on the east side decreases over time due to the closure of Humboldt in 2012 and the merging of Chief Joseph and Ockley Green in 2013.

School	T14	C-O	School	T14	C-O	School	T14	C-O
Abernethy	4	2	Duniway	4	3	Ockley Green	N/A	2
Ainsworth	4	3	Forest Park	3	2	Peninsula	3	2
Alameda	5	4	Glencoe	3	3	Rieke	3	2
Arleta	3	2	Grout	4	3	Rigler	4	3
Astor	2	2	Harrison Park	4	3	Rosa Parks	3	3
Atkinson	3	3	Hayhurst	3	2	Roseway Hts	5	3
Beach	3	2	Humboldt	N/A	1	Sabin	3	2
Beverly Cleary	4	2	Irvington	2	2	Scott	4	2
Boise-Eliot/Hum	4	3	James John	4	3	Sitton	3	2
Bridger	3	2	Kelly	5	4	Skyline	1.5	1
Bridlemile	3	3	King	4	3	Stephenson	2	2
Buckman	3	4	Laurelhurst	3	2	Sunnyside	3	2
Capitol Hill	3	2	Lee	3	2	Vernon	3	2
Chapman	5	3	Lent	3	2	Vestal	3	2
Chief Joseph	N/A	2	Lewis	3	2	Whitman	2	2
CJOG	4	3	Llewellyn	3	2	Woodlawn	3	2
Creston	2	2	Maplewood	3	2	Woodmere	3	2
César Chávez	3	2	Markham	3	3	Woodstock	4	3
Faubion	2.5	2	Marysville	3	2			

Table 4: *True-2014* (T14) and *conservative-optimal* (C-O) section counts for each school. We use True-2014 section counts in our 2014–15 experiments (Section 6.1), and conservative-optimal in our 2008–15 experiments (Section 6.2). CJOG's conservative-optimal kindergarten section-count of 3 was derived by taking the Con-Opt count for Chief Joseph K-5 (2), and dividing the total number of sections (12) by the number of grades attending CJOG at Chief Joseph K-3 (4). A few schools' grade configurations lack green and yellow dots in the chart [12]: Skyline K-8's conservative-optimal section-count is set at 1; Hayhurst K-8's conservative-optimal count is set at 2 because the neighborhood program is K-5; and Humboldt's conservative-optimal count is set at 1.

Distances Home-to-school distances are driving (by-car) distances, computed using the OSRM routing engine[24] applied to OSM maps[25]. When spot-checked against Google maps, the output from the OSRM routing engine seems comparable. There are a small number of peculiar distances that show up in the raw output, but these appear to be an artifact of the PPS-provided student location data: home locations were randomly perturbed to anonymize them, but this has placed a few students' homes in odd places (e.g., in the middle of an I-5 northbound lane). To filter these out, if any routed distance is greater than 3 times the Manhattan distance,[26] we use the Manhattan distance instead. Ideally, travel distances would be computed using the same system as the PPS Transportation Department, e.g., to match the methodology used by the district to determine busing eligibility.

No Cross-River Assignments When we run experiments using the Single Closest School and Soft Neighborhood Model in Section 6, we effectively disallow cross-river assignments by adding a very large penalty to the routed distances between each student's home and any school on the other side of the river.[27]

There is nothing about the Soft Neighborhood Model that inherently disallows cross-river assignments. In fact, SNM performs better if it has the freedom to balance enrollments over appropriate river crossings. However, the historical boundaries in our data do not cross the river, and the district has tended to treat the schools on opposite sides of the river independently from one another. We separate the two sides in the bulk of our discussion here in order to preserve the reference frames that people are familiar with, to make it easier to compare SNM against the historical hard boundaries. It also makes it clear that the success of the Soft Neighborhood Model does not *depend* on river-crossings.

Because we effectively separate the east and west sides of the river, we will see in the enrollment balancing experiments of Sections 6.1.1 and 6.2.1 that the mean section-size is different on each side. In Section 7, we include two sample assignments that allow cross-river assignments for comparison (Figures 34 and 35). Allowing cross-river assignments allows the SNM to equalize section-sizes uniformly across the entire district.

Instead of the artificial penalty, we would prefer to see a travel distance metric that realistically accounts for the difficulties of specific river-crossings — traffic congestion, bus routing, incidence of bridge closure, risk of meteor attack, etc. This would cause SNM to preclude cross-river assignments that don't make sense, and allow the ones that do.

[24] "Open Source Routing Engine," `http://project-osrm.org/`

[25] "Open Street Map," `http://www.openstreetmap.org/`

[26] *Manhattan distance* is simply the east-west displacement summed with the north-south displacement. It's the distance you would travel on a rectangular grid, i.e., Manhattan.

[27] Historical Hard Boundary transfer assignments are preserved as is, even if they cross the river.

6 Results on PPS Data

6.1 2014–2015 Results

In this section, we present the results of applying the different assignment models to the 2014–15 kindergarten population, using *True-2014* section counts described in Section 5.3.2, which to our knowledge represent the actual school configurations for the 2014–15 school year. We use these section counts here because they reflect the school configurations actually experienced by PPS families in the 2014–15 school year. Note, however, that these section counts are "tuned" to student populations near the schools: PPS established these section counts — squeezing in extra kindergarten sections in some schools, and dropping sections in others — to attempt to account for the number of students actually being captured by the inflexible hard boundaries. Thus, the *True-2014* counts give all models a boost in enrollment balancing performance. Even with that boost, the existing Hard Boundary system has horribly unbalanced enrollment.

For all metrics, we show the *No Transfers* setup, which simulates complete elimination of neighborhood school transfers — all students are assigned using a neighborhood school model. We do this because:

1. the district has actively reduced the number of transfers, abolishing the neighborhood transfer lottery beginning in 2015–16, leaving only need-based petition transfers;

2. for most metrics, transfers have the same effect on both Hard Boundaries and Soft Neighborhood assignments and just introduce noise that interferes with seeing what the neighborhood assignment models are doing.

As the one exception, we also show the *With Transfers* results for the *Enrollment Balancing* metric. Transfers actually help all models with enrollment balancing, by shifting more students into schools with excess capacity. The *With Transfers* setup uses the actual PPS Hard Boundary assignments for 2014–15 for all students, and shows that the Soft Neighborhood Models still do a far better job of enrollment balancing.

6.1.1 Enrollment Balancing

Year(s)	Model, East Side No Transfers	Kindergarten Section Sizes				
		≤17	18 to 19	20 to 21	22 to 23	≥24
2014-15	Hard Boundaries	26.5%	16.5%	16.5%	10.8%	29.7%
	Single Closest School	40.9%	6.1%	18.6%	10.8%	23.7%
	SNM, 3-school, Strong	0.0%	0.7%	99.3%	0.0%	0.0%
	SNM, 3-school, Moderate	0.0%	1.4%	97.8%	0.7%	0.0%
	SNM, 3-school, Equal	0.0%	2.2%	97.1%	0.7%	0.0%
	SNM, 2-school, Strong	0.0%	3.6%	87.8%	8.6%	0.0%
	SNM, 2-school, Moderate	0.0%	4.3%	87.1%	8.6%	0.0%
	SNM, 2-school, Equal	0.0%	5.0%	86.4%	8.6%	0.0%

Year(s)	Model, West Side No Transfers	Kindergarten Section Sizes				
		≤20	21 to 22	23 to 24	25 to 26	≥27
2014-15	Hard Boundaries	11.9%	22.4%	14.9%	14.9%	35.8%
	Single Closest School	34.3%	3.0%	20.9%	6.0%	35.8%
	SNM, 3-school, Strong	0.0%	4.5%	95.5%	0.0%	0.0%
	SNM, 3-school, Moderate	0.0%	4.5%	95.5%	0.0%	0.0%
	SNM, 3-school, Equal	0.0%	4.5%	95.5%	0.0%	0.0%
	SNM, 2-school, Strong	4.5%	0.0%	83.6%	11.9%	0.0%
	SNM, 2-school, Moderate	4.5%	0.0%	83.6%	11.9%	0.0%
	SNM, 2-school, Equal	4.5%	0.0%	83.6%	11.9%	0.0%

Table 5: *Kinder Section Sizes, No Transfers, True-2014 Section Counts, 2014–15.* The Soft Neighborhood Model is much better at hitting the targets section sizes than the Hard Boundary and Single Closest School models, on both sides of the river. In general, the SNM 3-School versions have more options for placing kids, and therefore do a little better at hitting the targets than the 2-School versions. Because we disallow cross-river assignments for the purposes of these experiments (Section 5.4), the model balances the east and west sides independently, and the target section sizes differ considerably on the east (21) and west sides (24). The SNM can do a better job of balancing enrollments uniformly across the whole district if it is allowed to make cross-river assignments (see Section 7, Figures 34 and 35).

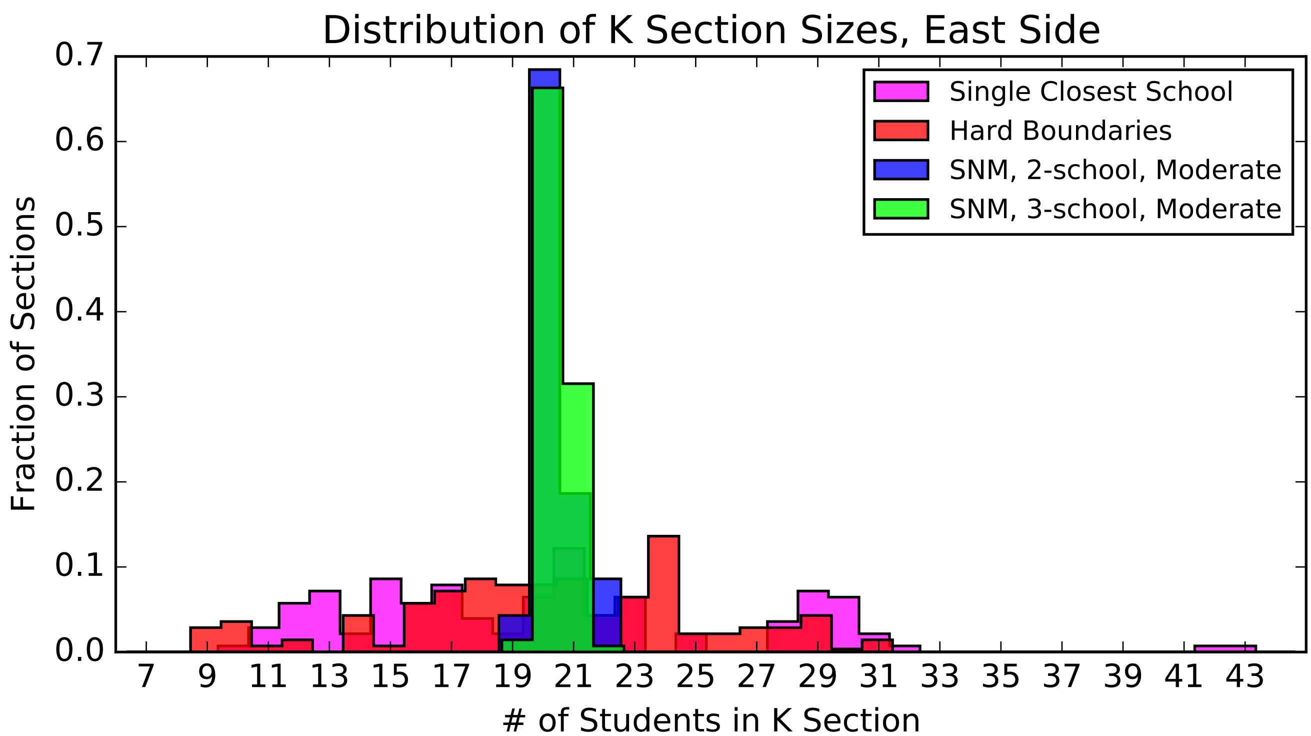

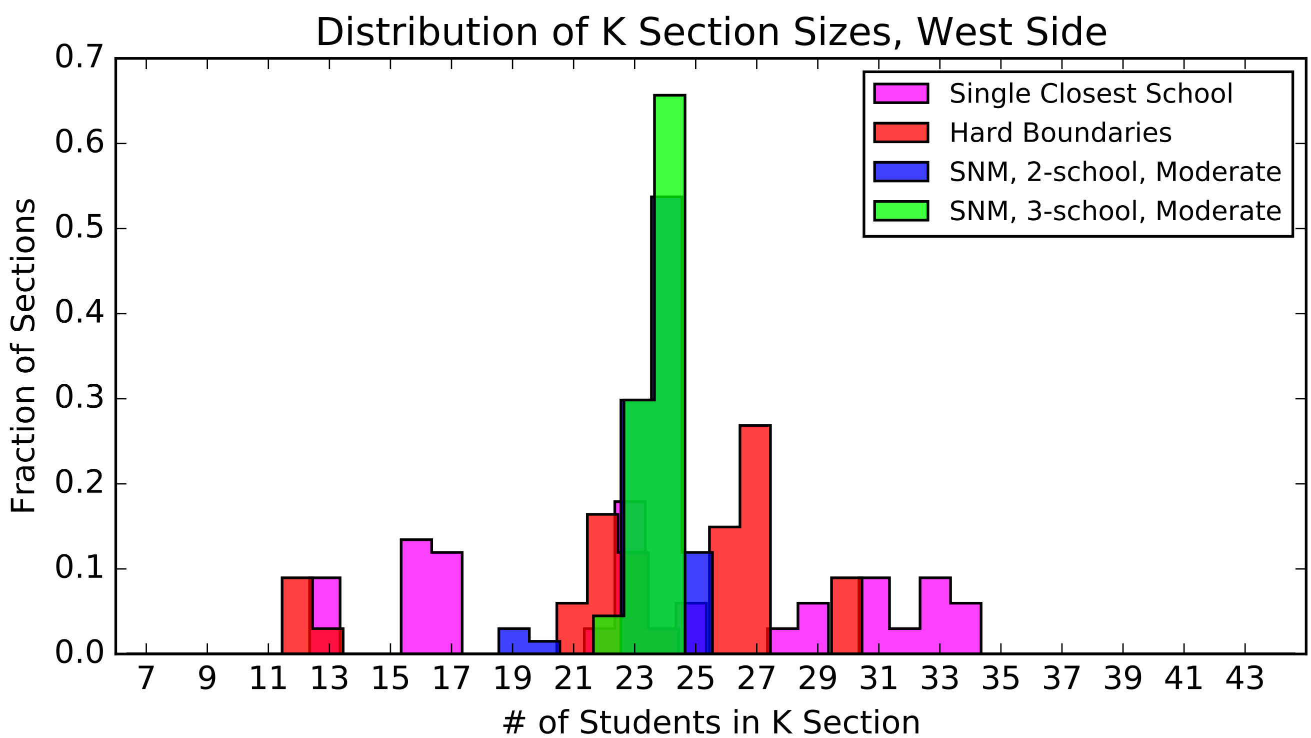

Figure 16: *Kinder Section Sizes, No Transfers, True-2014 Section Counts, 2014–15.* (Same data as Table 5.) SNM 2-School and SNM 3-School kindergarten section sizes are grouped very tightly to the targets. In contrast, Hard Boundary and Single Closest School sections vary wildly in size — on the east side, some schools have sections that are *3 times* as big as other schools.

Year(s)	Model, East Side With Transfers	Kindergarten Section Sizes				
		≤17	18 to 19	20 to 21	22 to 23	≥24
2014-15	Hard Boundaries	22.2%	27.2%	17.2%	12.9%	20.4%
	Single Closest School	38.7%	8.6%	11.8%	11.5%	29.4%
	SNM, 3-school, Strong	0.0%	1.4%	98.6%	0.0%	0.0%
	SNM, 3-school, Moderate	0.0%	2.2%	97.8%	0.0%	0.0%
	SNM, 3-school, Equal	0.0%	1.4%	98.6%	0.0%	0.0%
	SNM, 2-school, Strong	0.0%	10.8%	89.2%	0.0%	0.0%
	SNM, 2-school, Moderate	0.0%	10.0%	90.0%	0.0%	0.0%
	SNM, 2-school, Equal	0.0%	10.8%	89.2%	0.0%	0.0%

Year(s)	Model, West Side With Transfers	Kindergarten Section Sizes				
		≤20	21 to 22	23 to 24	25 to 26	≥27
2014-15	Hard Boundaries	9.0%	26.9%	23.9%	19.4%	20.9%
	Single Closest School	22.4%	20.9%	6.0%	23.9%	26.9%
	SNM, 3-school, Strong	0.0%	0.0%	100.0%	0.0%	0.0%
	SNM, 3-school, Moderate	0.0%	0.0%	100.0%	0.0%	0.0%
	SNM, 3-school, Equal	0.0%	0.0%	97.0%	3.0%	0.0%
	SNM, 2-school, Strong	0.0%	1.5%	98.5%	0.0%	0.0%
	SNM, 2-school, Moderate	0.0%	1.5%	98.5%	0.0%	0.0%
	SNM, 2-school, Equal	0.0%	1.5%	98.5%	0.0%	0.0%

Table 6: *Kinder Section Sizes, With Transfers, True-2014 Section Counts, 2014–15.* By including historic transfers, the Hard Boundaries values here reflect the actual historic assignments made in 2014–15. Note that the *With Transfers* setting makes enrollment balancing easier for all models by shifting students to schools with excess capacity.

6.1.2 Distance

Year(s)	Model, East Side No Transfers	Distance from Home to School			
		<0.5mi	<1.0mi	<1.5mi	>1.5mi
2014-15	Hard Boundaries	35.2%	79.3%	94.5%	5.5%
	Single Closest School	39.7%	88.5%	97.9%	2.1%
	SNM, 3-school, Strong	27.8%	75.7%	95.7%	4.3%
	SNM, 3-school, Moderate	24.4%	68.4%	93.1%	6.9%
	SNM, 3-school, Equal	14.0%	52.6%	88.7%	11.3%
	SNM, 2-school, Strong	34.1%	80.7%	96.6%	3.4%
	SNM, 2-school, Moderate	31.2%	78.1%	96.2%	3.8%
	SNM, 2-school, Equal	22.7%	71.7%	95.5%	4.5%
Year(s)	**Model, West Side No Transfers**	**Distance from Home to School**			
		<0.5mi	<1.0mi	<1.5mi	>1.5mi
2014-15	Hard Boundaries	19.1%	47.9%	73.3%	26.7%
	Single Closest School	20.0%	56.6%	81.8%	18.2%
	SNM, 3-school, Strong	15.2%	42.6%	69.5%	30.5%
	SNM, 3-school, Moderate	13.5%	38.5%	62.8%	37.2%
	SNM, 3-school, Equal	8.2%	25.3%	46.9%	53.1%
	SNM, 2-school, Strong	19.2%	51.0%	77.7%	22.3%
	SNM, 2-school, Moderate	17.5%	46.6%	73.4%	26.6%
	SNM, 2-school, Equal	11.4%	34.2%	63.5%	36.5%

Table 7: *Driving Distance, No Transfers, True-2014 Section Counts, 2014–15.* The Single Closest School model has the shortest driving distances. SNM 2-School Strong and SNM 2-School Moderate distances are very similar to the Hard Boundary model. Distances are longer for SNM 3-School Strong, SNM 3-School Moderate, and SNM 2-School Equal. SNM 3-School Equal has the longest driving distances. When we include transfers (not shown), distances lengthen for all models, but the relative distance trends are the same as in the *No Transfers* setting.

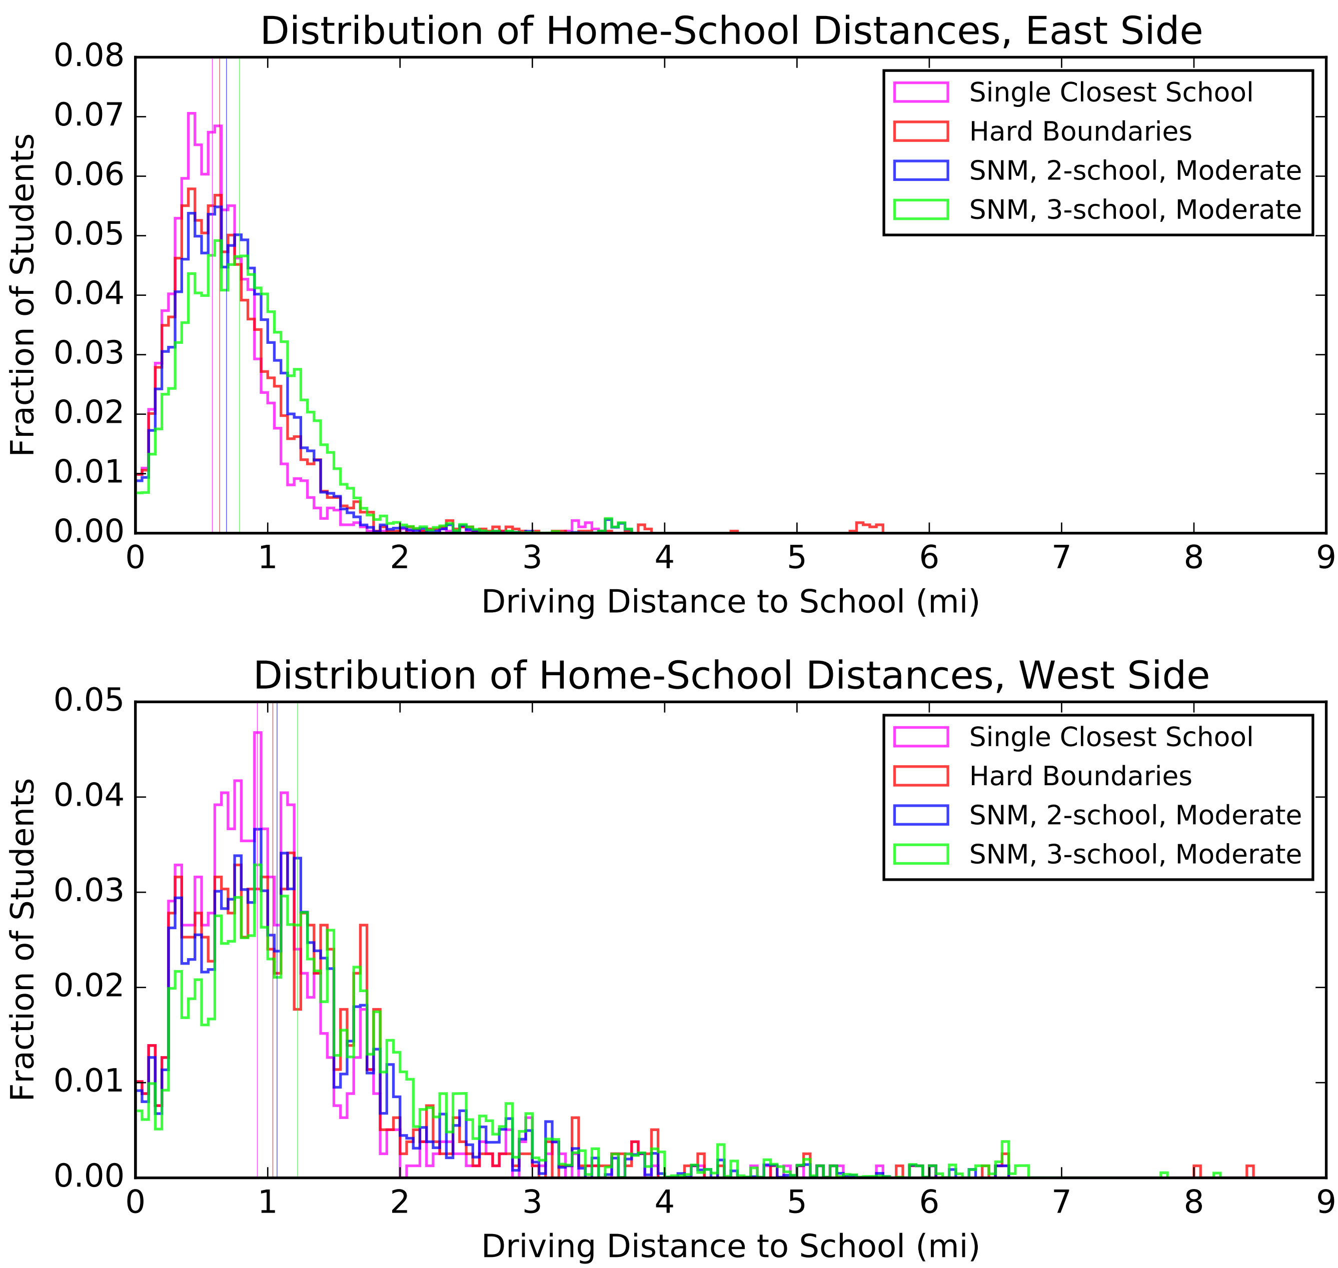

Figure 17: *Driving Distance, No Transfers, True-2014 Section Counts, 2014–15.* (Same data as Table 7. Vertical lines indicate median distances.) The Single Closest School model has the shortest driving distances. Hard Boundary and SNM 2-School Moderate distances are very similar. SNM 3-School Moderate distances are longest of the four.

6.1.3 Closest Classmate

Year(s)	Model, East Side No Transfers	Closest Classmate, ft			
		≤500.0	≤1000.0	≤1500.0	≤2000.0
2014-15	Hard Boundaries	77.9%	98.1%	99.6%	99.9%
	Single Closest School	81.4%	98.7%	99.5%	99.8%
	SNM, 3-school, Strong	69.1%	94.7%	98.4%	99.4%
	SNM, 3-school, Moderate	58.3%	90.2%	96.9%	98.8%
	SNM, 3-school, Equal	51.9%	88.1%	96.7%	98.8%
	SNM, 2-school, Strong	76.0%	96.6%	99.0%	99.6%
	SNM, 2-school, Moderate	71.2%	95.4%	98.7%	99.5%
	SNM, 2-school, Equal	67.8%	95.0%	98.7%	99.5%
Year(s)	**Model, West Side No Transfers**	**Closest Classmate, ft**			
		≤500.0	≤1000.0	≤1500.0	≤2000.0
2014-15	Hard Boundaries	57.5%	87.7%	93.8%	96.7%
	Single Closest School	59.2%	88.2%	94.3%	97.0%
	SNM, 3-school, Strong	47.2%	81.7%	91.2%	95.1%
	SNM, 3-school, Moderate	41.2%	76.4%	87.9%	93.0%
	SNM, 3-school, Equal	36.7%	72.6%	87.1%	92.7%
	SNM, 2-school, Strong	56.2%	86.7%	93.6%	96.2%
	SNM, 2-school, Moderate	50.2%	82.4%	91.1%	94.8%
	SNM, 2-school, Equal	45.5%	80.9%	91.3%	95.1%

Table 8: *Closest Classmate, No Transfers, True-2014 Section Counts, 2014–15.* As a point of reference, 500 ft is approximately the length of one "long" block in NE Portland. Closest classmate distances are very similar for Hard Boundaries and Single Closest School models. SNM 2-School distances are a little longer, and SNM 3-School distances are a little longer still (though 2-School Equal and 3-School Strong are very similar). The models are all pretty similar by around 1500 ft (around 3 standard east-side long blocks). The proportion of students who have to travel very long distances (>2000 ft, or around 4 east-side long blocks) to reach the home of the closest classmate is small in all models. As with driving distances, distances lengthen for all models in the *With Transfers* setting (not shown), but the relative trends are the same.

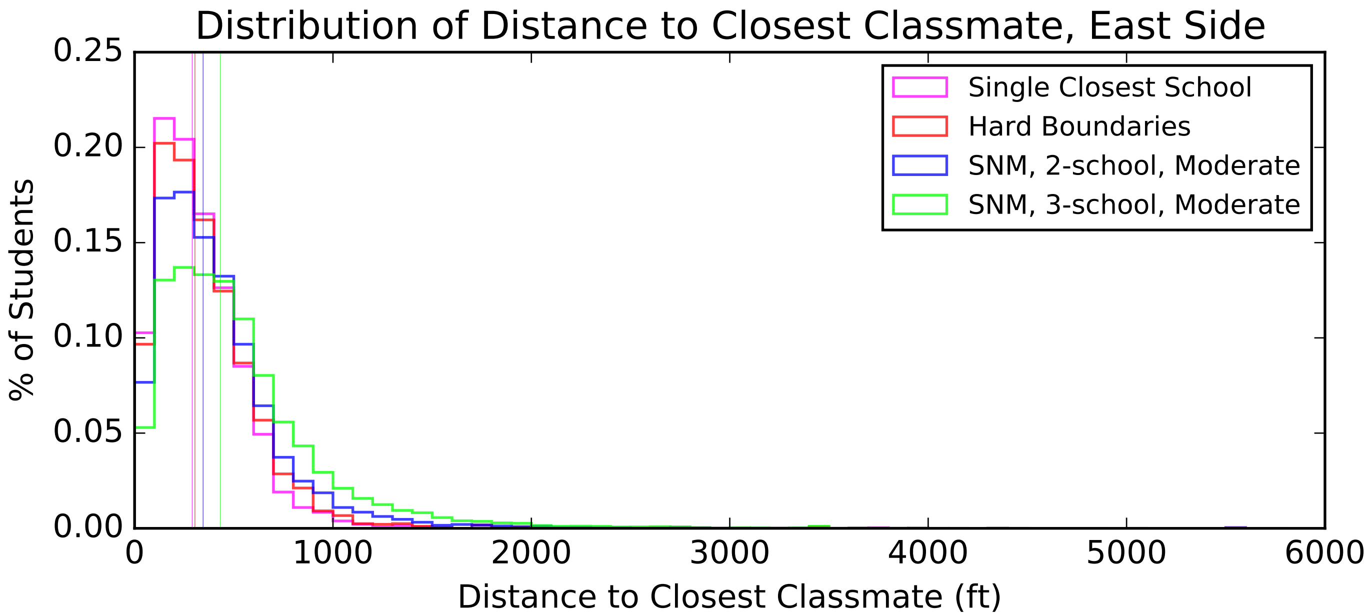

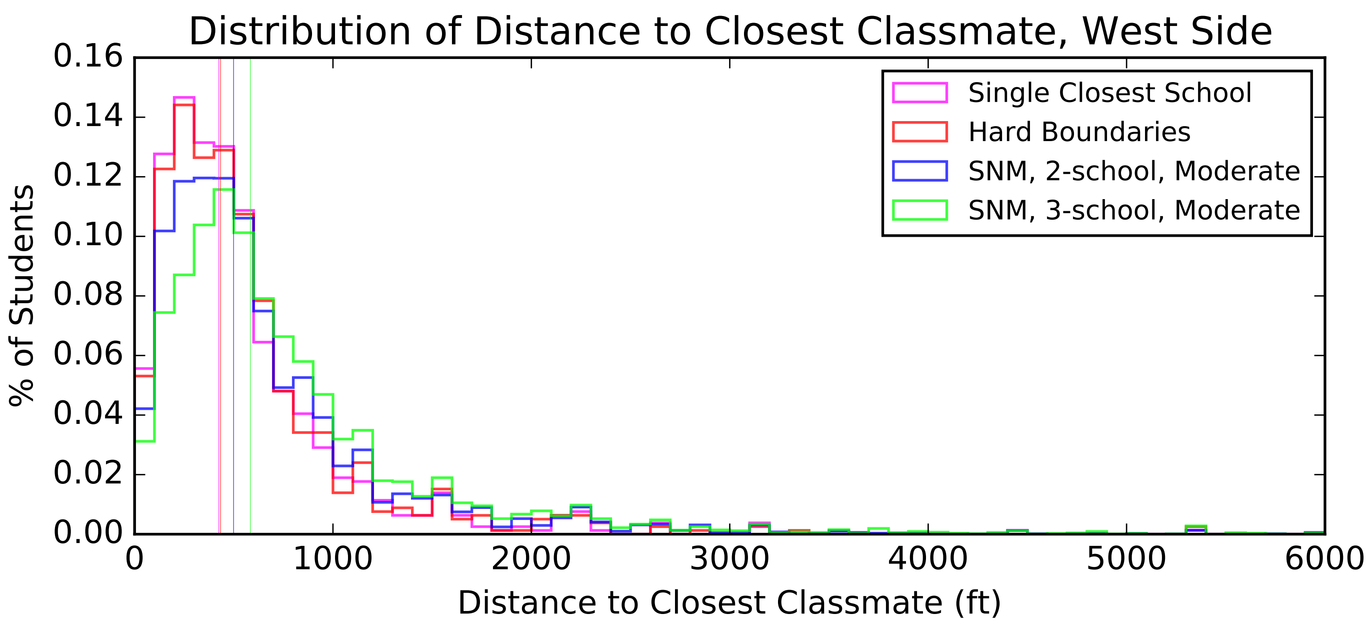

Figure 18: *Closest Classmate, No Transfers, True-2014 Section Counts, 2014–15.* (Same data as Table 8. Vertical lines indicate median distances.) Single Closest School and Hard Boundary distances are the shortest. The SNM 2-School moderate distances are slightly longer, and the SNM 3-School moderate distances are the longest of the four. The most marked differences between the models are at distances of <500 ft, or one standard east-side long block. The proportion of students who have to travel very long distances (>2000 ft, or around 4 east-side long blocks) to reach the home of the closest classmate is small in all models.

6.1.4 Local Assignment Diversity

Year(s)	**Model, East Side No Transfers**	**Assignment Diversity** [1.0, 1.5)	[1.5, 2.5)	[2.5, 3.5)	≥3.5
2014-15	Hard Boundaries	65.2%	30.8%	4.0%	0.1%
	Single Closest School	67.7%	29.1%	3.1%	0.1%
	SNM, 3-school, Strong	29.3%	57.7%	12.5%	0.5%
	SNM, 3-school, Moderate	12.2%	49.0%	35.1%	3.7%
	SNM, 3-school, Equal	5.2%	31.0%	49.9%	13.9%
	SNM, 2-school, Strong	50.0%	45.1%	4.8%	0.1%
	SNM, 2-school, Moderate	33.5%	57.7%	8.5%	0.4%
	SNM, 2-school, Equal	21.2%	60.9%	16.7%	1.2%

Year(s)	**Model, West Side No Transfers**	**Assignment Diversity** [1.0, 1.5)	[1.5, 2.5)	[2.5, 3.5)	≥3.5
2014-15	Hard Boundaries	80.4%	19.3%	0.3%	0.0%
	Single Closest School	86.3%	12.8%	0.9%	0.0%
	SNM, 3-school, Strong	45.7%	49.9%	4.3%	0.1%
	SNM, 3-school, Moderate	31.7%	49.6%	17.9%	0.9%
	SNM, 3-school, Equal	24.5%	36.9%	33.9%	4.7%
	SNM, 2-school, Strong	75.6%	23.4%	1.0%	0.0%
	SNM, 2-school, Moderate	56.3%	41.0%	2.7%	0.0%
	SNM, 2-school, Equal	36.9%	55.8%	7.2%	0.1%

Table 9: *Local Assignment Diversity, No Transfers, True-2014 Section Counts, 2014–15.* All Soft Neighborhood models are more diverse than Hard Boundary and Single Closest School models. SNM 3-School models have greater diversity than SNM 2-School models. As discussed in Section 5.2, more assignment diversity suggests less gameability and more communication among families attending different nearby schools.

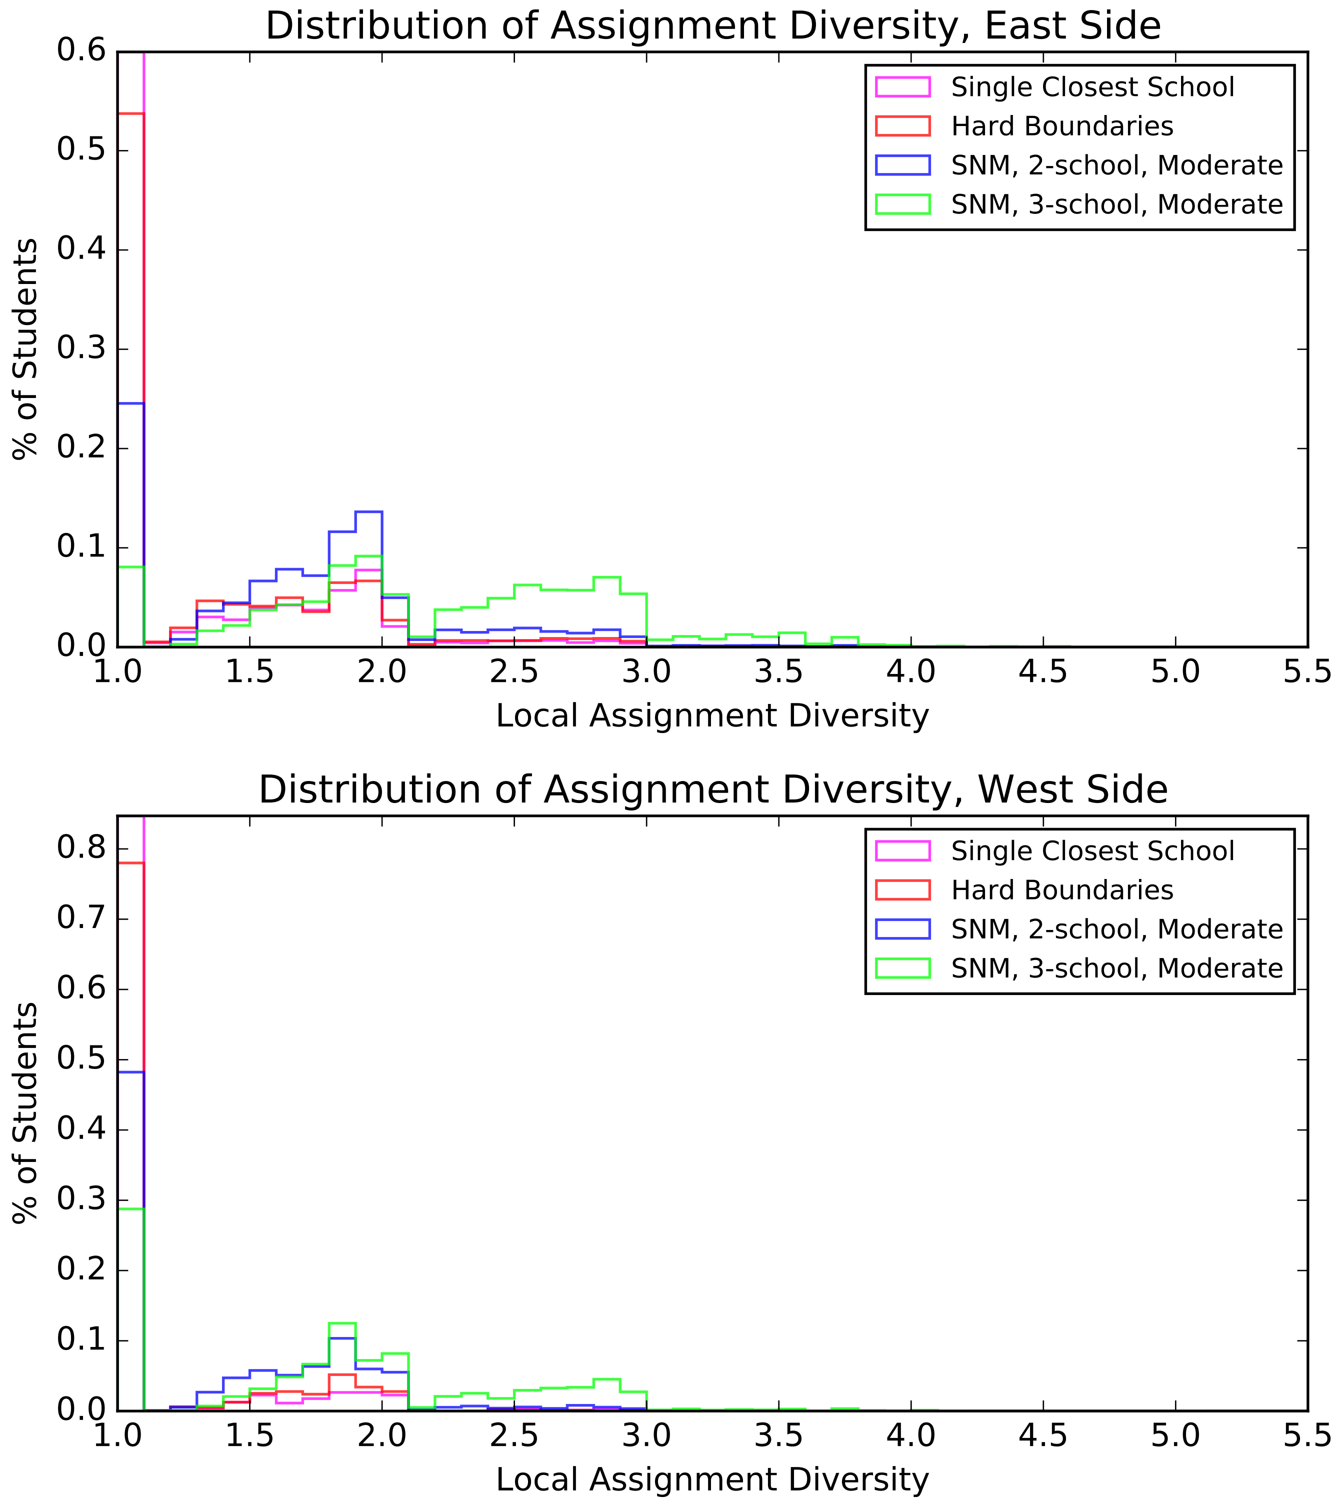

Figure 19: *Local Assignment Diversity, No Transfers, True-2014 Section Counts, 2014–15.* (Same data as Table 9.) In Single Closest School and Hard Boundary assignments, the vast majority of students live in areas where only one school is well represented. In the SNM, substantially fewer students live in such areas. More students live in areas where 1–2 schools are well represented, with good representation of 2–3 schools in a 3-school SNM.

6.1.5 Capture Linkage

Model, East Side No Transfers	**Capture Linkage, weighted balanced matrix**					
	<1.0	[1.0, 2.0)	[2.0, 3.0)	[3.0, 4.0)	[4.0, 5.0)	≥5.0
Hard Boundaries	100.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Single Closest School	100.0%	0.0%	0.0%	0.0%	0.0%	0.0%
SNM, 3-school, Strong	0.0%	16.7%	26.2%	23.8%	21.4%	11.9%
SNM, 3-school, Moderate	0.0%	7.1%	23.8%	26.2%	16.7%	26.2%
SNM, 3-school, Equal	0.0%	2.4%	23.8%	19.0%	23.8%	31.0%
SNM, 2-school, Strong	28.6%	35.7%	31.0%	4.8%	0.0%	0.0%
SNM, 2-school, Moderate	21.4%	45.2%	26.2%	4.8%	2.4%	0.0%
SNM, 2-school, Equal	19.0%	42.9%	33.3%	2.4%	2.4%	0.0%
Model, West Side No Transfers	**Capture Linkage, weighted balanced matrix**					
	<1.0	[1.0, 2.0)	[2.0, 3.0)	[3.0, 4.0)	[4.0, 5.0)	≥5.0
Hard Boundaries	100.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Single Closest School	100.0%	0.0%	0.0%	0.0%	0.0%	0.0%
SNM, 3-school, Strong	18.2%	9.1%	18.2%	45.5%	9.1%	0.0%
SNM, 3-school, Moderate	18.2%	9.1%	18.2%	27.3%	27.3%	0.0%
SNM, 3-school, Equal	18.2%	9.1%	18.2%	18.2%	18.2%	18.2%
SNM, 2-school, Strong	36.4%	45.5%	18.2%	0.0%	0.0%	0.0%
SNM, 2-school, Moderate	27.3%	45.5%	27.3%	0.0%	0.0%	0.0%
SNM, 2-school, Equal	27.3%	45.5%	27.3%	0.0%	0.0%	0.0%

Table 10: *Capture Linkage, No Transfers, True-2014 Section Counts, 2014–15.* Hard Boundary and Single Closest School models have no flexibility in assignment, and thus always have zero[28] capture linkage (Section 5.2). In SNM 3-School models, capture linkage is more skewed towards higher linkage than in SNM 2-School models. Because of their higher linkage, schools in SNM 3-School models are less influenced by gaming than in SNM 2-School models. Higher linkage also means nearby schools "share" populations more, helping to diffuse and stabilize demographic trends.

[28]Single Closest School's capture linkage is not exactly zero because of the handful of students that live equidistant from two equally-closest schools.

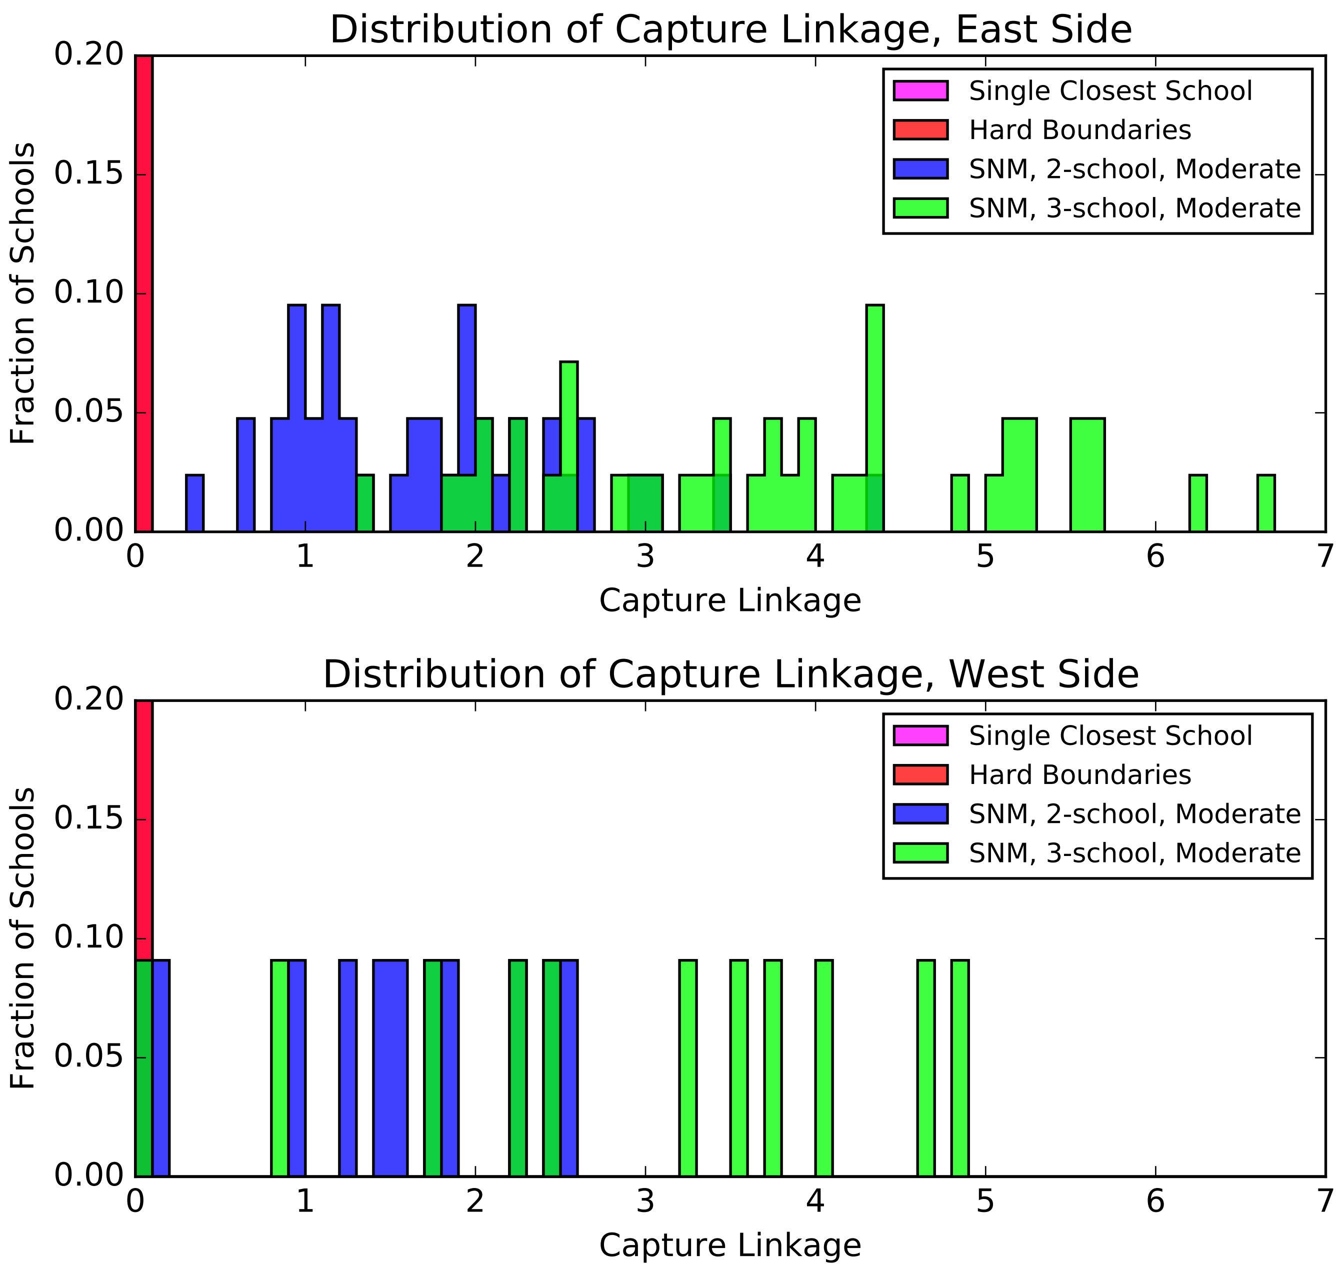

Figure 20: *Capture Linkage, No Transfers, True-2014 Section Counts, 2014–15.* (Same data as Table 10.) In Single Closest School and Hard Boundary models, school linkages are exactly[29] zero for all schools. Soft Neighborhood Model capture linkages are higher, and 3-School is quite a bit higher than 2-School. *Fractions for Hard Boundaries and Single Closest School at zero-linkage are 1.0, but have been clipped to 0.20 in the plots above.*

[29] Almost: see Footnote 28.

6.2 Cumulative (2008–2015) Results

In this section, we present cumulative results of applying the different assignment models to all seven kindergarten classes in the 2008–2015 school years. Recall that in Section 6.1, we used *True-2014* section counts, which effectively have been tuned in response to population imbalances throughout the district. Here, we are instead using the *Conservative-Optimal* (*Con-Opt*) section counts as defined in Section 5.3.2. These counts are based on a PPS report which is nominally an assessment of the physical capacity of its buildings, and thus should be independent of the nearby student populations.[30] They should let us see what the various models would have done over 2008–2015 if the schools had been optimally configured (according to the definition of "optimal" PPS used for its report).

Again, we focus on the *No Transfers* setup, as explained previously in Section 6.1.

[30] We do account for the three schools which experienced grade reconfigurations in this time period.

6.2.1 Enrollment Balancing

Year(s)	Model, East Side No Transfers	Kindergarten Enrollment Deviations				
		≤-4	-3 to -2	-1 to 1	2 to 3	≥4
2008-15	Hard Boundaries	37.1%	5.6%	12.9%	8.5%	36.0%
	Single Closest School	41.5%	5.1%	8.8%	7.4%	37.2%
	SNM, 3-school, Strong	0.0%	0.6%	98.1%	1.4%	0.0%
	SNM, 3-school, Moderate	0.0%	0.6%	98.2%	1.3%	0.0%
	SNM, 3-school, Equal	0.0%	0.3%	98.5%	1.2%	0.0%
	SNM, 2-school, Strong	3.8%	14.0%	75.1%	4.6%	2.5%
	SNM, 2-school, Moderate	3.8%	13.9%	75.3%	4.6%	2.5%
	SNM, 2-school, Equal	3.8%	14.3%	74.7%	4.7%	2.5%

Year(s)	Model, West Side No Transfers	Kindergarten Enrollment Deviations				
		≤-4	-3 to -2	-1 to 1	2 to 3	≥4
2008-15	Hard Boundaries	34.3%	5.1%	14.3%	14.9%	31.4%
	Single Closest School	37.7%	8.6%	15.4%	6.9%	31.4%
	SNM, 3-school, Strong	1.7%	1.1%	97.1%	0.0%	0.0%
	SNM, 3-school, Moderate	1.7%	1.1%	96.6%	0.6%	0.0%
	SNM, 3-school, Equal	1.7%	1.1%	97.1%	0.0%	0.0%
	SNM, 2-school, Strong	2.9%	0.6%	94.9%	1.7%	0.0%
	SNM, 2-school, Moderate	2.9%	0.6%	94.9%	1.7%	0.0%
	SNM, 2-school, Equal	2.9%	0.6%	94.9%	1.7%	0.0%

Table 11: *Kinder Enrollment Deviations, No Transfers, Con-Opt Section Counts, 2008–15.* This table shows under and over-enrollment of kindergarten sections with respect to the target section size, aggregated over seven years. The *Con-Opt* section counts are mostly constant, which implies that the target section size changes from year to year as the total population increases or decreases. Deviations are the differences between the number of students assigned to each kindergarten section, and that year's target section size. Because we disallow cross-river assignments (Section 5.4), we treat the east and west sides independently. SNM outperforms Hard Boundaries and Single Closest School, as it did in the 2014–15 only *True-2014* scenario (Section 6.1.1). However, without the advantage of having the section counts pre-tuned to local populations, 2-School SNM does noticeably worse on the east side than 3-School SNM.

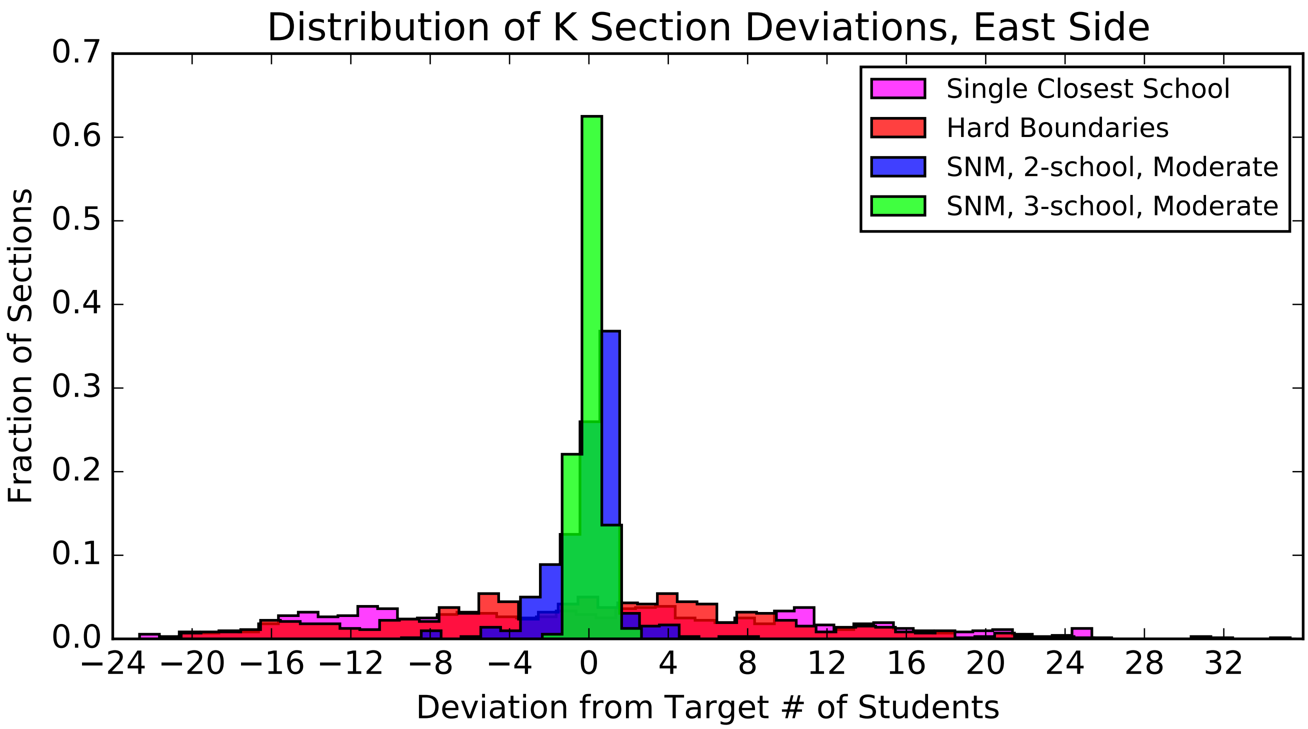

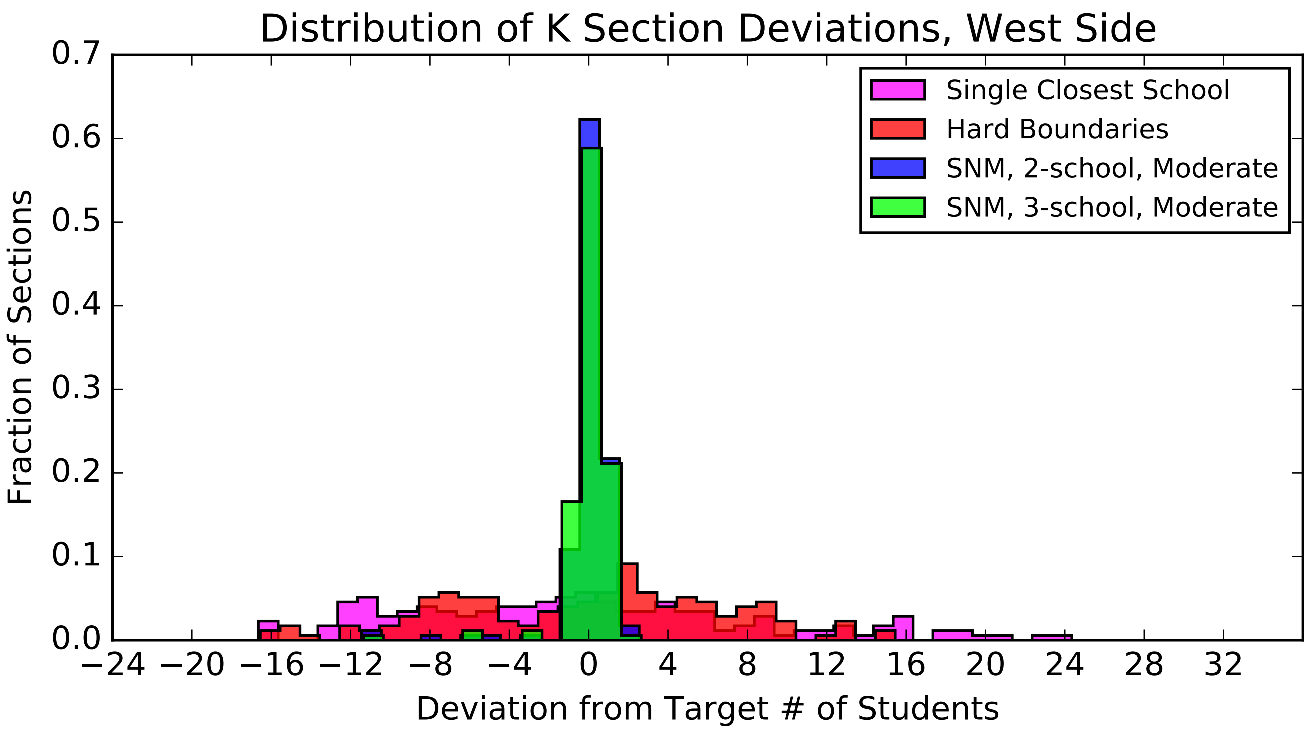

Figure 21: *Kinder Section Sizes, No Transfers, Con-Opt Section Counts, 2008–15.* (Same data as Table 11.) Over the entire 7 year period, SNM 2-School and SNM 3-School kindergarten section sizes are consistently grouped very tightly around the targets for each side of the river. In contrast, Hard Boundary and Single Closest School sections vary wildly in size.

Year(s)	**Model, East Side With Transfers**	**Kindergarten Enrollment Deviations**				
		≤-4	-3 to -2	-1 to 1	2 to 3	≥4
2008-15	Hard Boundaries	35.3%	11.2%	15.1%	13.3%	25.0%
	Single Closest School	32.0%	14.6%	17.9%	8.2%	27.2%
	SNM, 3-school, Strong	0.0%	1.7%	98.2%	0.1%	0.0%
	SNM, 3-school, Moderate	0.0%	1.8%	97.9%	0.3%	0.0%
	SNM, 3-school, Equal	0.0%	2.1%	97.9%	0.0%	0.0%
	SNM, 2-school, Strong	1.8%	6.7%	79.3%	11.7%	0.6%
	SNM, 2-school, Moderate	1.8%	6.9%	79.0%	11.7%	0.6%
	SNM, 2-school, Equal	1.8%	7.2%	78.6%	11.8%	0.6%

Year(s)	**Model, West Side With Transfers**	**Kindergarten Enrollment Deviations**				
		≤-4	-3 to -2	-1 to 1	2 to 3	≥4
2008-15	Hard Boundaries	29.7%	8.6%	25.1%	11.4%	25.1%
	Single Closest School	37.1%	4.7%	10.0%	12.4%	35.8%
	SNM, 3-school, Strong	1.1%	1.1%	97.7%	0.0%	0.0%
	SNM, 3-school, Moderate	1.1%	0.6%	97.7%	0.6%	0.0%
	SNM, 3-school, Equal	1.1%	1.1%	97.7%	0.0%	0.0%
	SNM, 2-school, Strong	1.7%	1.7%	93.1%	3.4%	0.0%
	SNM, 2-school, Moderate	1.7%	1.7%	93.1%	3.4%	0.0%
	SNM, 2-school, Equal	1.7%	1.7%	93.1%	3.4%	0.0%

Table 12: *Kinder Enrollment Deviations, With Transfers, Con-Opt Section Counts, 2008–15.* The *With Transfers* setting makes enrollment balancing easier for all models, although the effect is greater for the Hard Boundary and Single Closest School models than the SNM. However, the Soft Neighborhood Model is still consistently better at balancing enrollments across all schools, across the span of seven years, and on the east and west sides. Going forward, we expect fewer transfers since the district eliminated the neighborhood-to-neighborhood transfer lottery in 2015.

6.2.2 Distance

Year(s)	Model, East Side No Transfers	Distance from Home to School			
		<0.5mi	<1.0mi	<1.5mi	>1.5mi
2008-15	Hard Boundaries	36.4%	80.8%	95.6%	4.4%
	Single Closest School	40.3%	88.8%	98.5%	1.5%
	SNM, 3-school, Strong	26.5%	73.0%	95.4%	4.6%
	SNM, 3-school, Moderate	23.7%	66.5%	92.9%	7.2%
	SNM, 3-school, Equal	14.2%	52.4%	88.8%	11.3%
	SNM, 2-school, Strong	33.3%	79.5%	96.7%	3.3%
	SNM, 2-school, Moderate	30.7%	77.3%	96.4%	3.6%
	SNM, 2-school, Equal	22.9%	71.8%	95.7%	4.3%
Year(s)	**Model, West Side No Transfers**	**Distance from Home to School**			
		<0.5mi	<1.0mi	<1.5mi	>1.5mi
2008-15	Hard Boundaries	20.2%	48.8%	72.2%	27.9%
	Single Closest School	20.8%	56.5%	80.4%	19.6%
	SNM, 3-school, Strong	15.5%	41.3%	66.3%	33.7%
	SNM, 3-school, Moderate	13.8%	37.8%	60.5%	39.5%
	SNM, 3-school, Equal	8.4%	24.8%	44.5%	55.5%
	SNM, 2-school, Strong	19.5%	48.4%	72.6%	27.4%
	SNM, 2-school, Moderate	18.1%	45.0%	69.3%	30.7%
	SNM, 2-school, Equal	12.1%	33.4%	59.0%	41.0%

Table 13: *Driving Distance, No Transfers, Con-Opt Section Counts, 2008–15.* The trend for driving distances over seven years is similar to what we see for 2014–15. The Single Closest School model has the shortest driving distances, followed by SNM 2-School Strong/Moderate and the Hard Boundary Model, the three of which are fairly comparable. The SNM 3-School Strong/Moderate and the SNM 2-School Equal versions have longer driving distances. The SNM 3-School Equal version has the lengthiest driving distances.

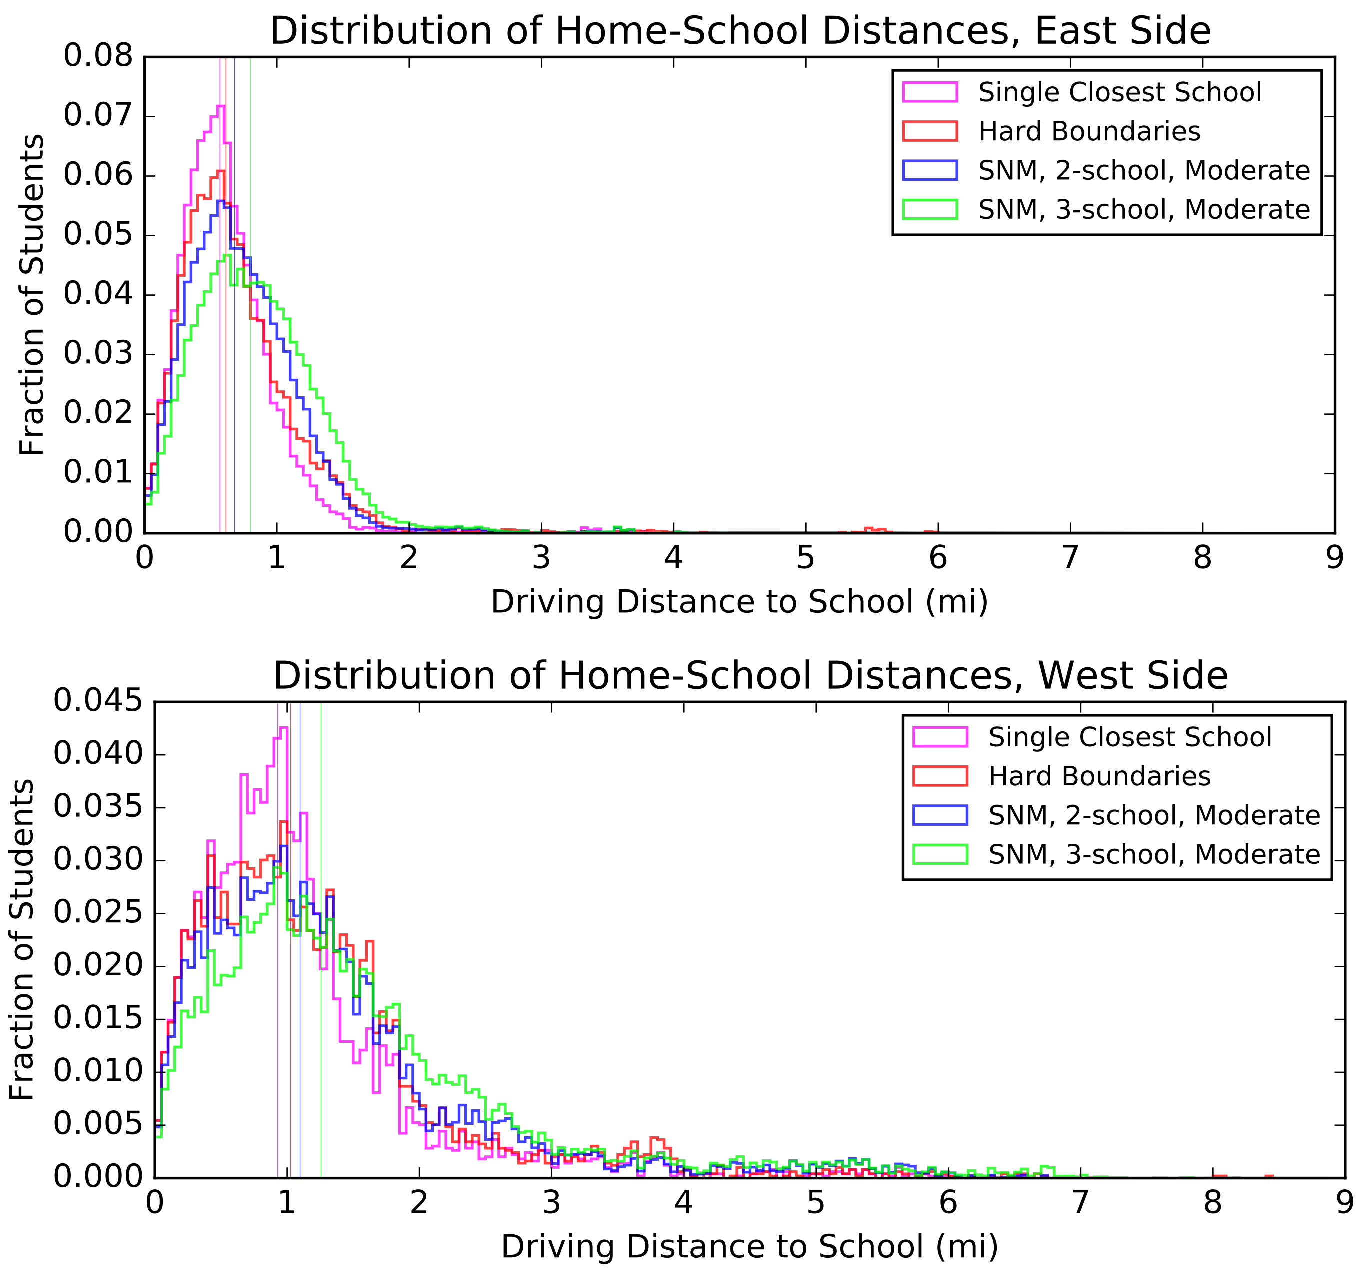

Figure 22: *Driving Distance, No Transfers, Con-Opt Section Counts, 2008–15.* (Same data as Table 13. Vertical lines indicate median distances.) The Single Closest School model has the shortest driving distances. Hard Boundary and SNM 2-School Moderate distances are similar. SNM 3-School Moderate distances are the longest of the four.

6.2.3 Closest Classmate

Year(s)	**Model, East Side No Transfers**	**Closest Classmate, ft**			
		$\leq$500.0	$\leq$1000.0	$\leq$1500.0	$\leq$2000.0
2008-15	Hard Boundaries	79.1%	98.1%	99.7%	99.9%
	Single Closest School	81.9%	98.5%	99.7%	99.8%
	SNM, 3-school, Strong	70.8%	95.0%	98.5%	99.3%
	SNM, 3-school, Moderate	61.7%	91.3%	97.1%	98.7%
	SNM, 3-school, Equal	55.9%	89.4%	96.8%	98.8%
	SNM, 2-school, Strong	77.0%	97.0%	99.2%	99.6%
	SNM, 2-school, Moderate	72.8%	95.7%	98.8%	99.5%
	SNM, 2-school, Equal	69.8%	95.2%	98.7%	99.4%
Year(s)	**Model, West Side No Transfers**	**Closest Classmate, ft**			
		$\leq$500.0	$\leq$1000.0	$\leq$1500.0	$\leq$2000.0
2008-15	Hard Boundaries	55.5%	88.9%	95.5%	97.5%
	Single Closest School	58.3%	89.5%	95.6%	97.4%
	SNM, 3-school, Strong	43.7%	79.9%	91.2%	95.1%
	SNM, 3-school, Moderate	36.7%	73.0%	87.1%	92.6%
	SNM, 3-school, Equal	31.5%	68.2%	85.1%	92.2%
	SNM, 2-school, Strong	51.6%	85.5%	93.8%	96.4%
	SNM, 2-school, Moderate	46.4%	81.4%	91.7%	95.4%
	SNM, 2-school, Equal	41.6%	79.2%	91.7%	95.7%

Table 14: *Closest Classmate, No Transfers, Con-Opt Section Counts, 2008–15.* Like driving distances, closest classmate statistics over all seven years are similar to what we see for 2014–15. Closest classmate distances are almost identical for Hard Boundaries and Single Closest School models. The SNM 2-School Strong closest classmate distances are not far behind. All SNM 2-School versions have shorter distances than the SNM 3-School versions (though 2-School Equal and 3-School Strong are very similar). As with the 2014–15 results, most models are pretty similar by around 1500 ft. (∼3 east-side long blocks), and in all models, only a small proportion of students have to travel very long distances to reach their closest classmate.

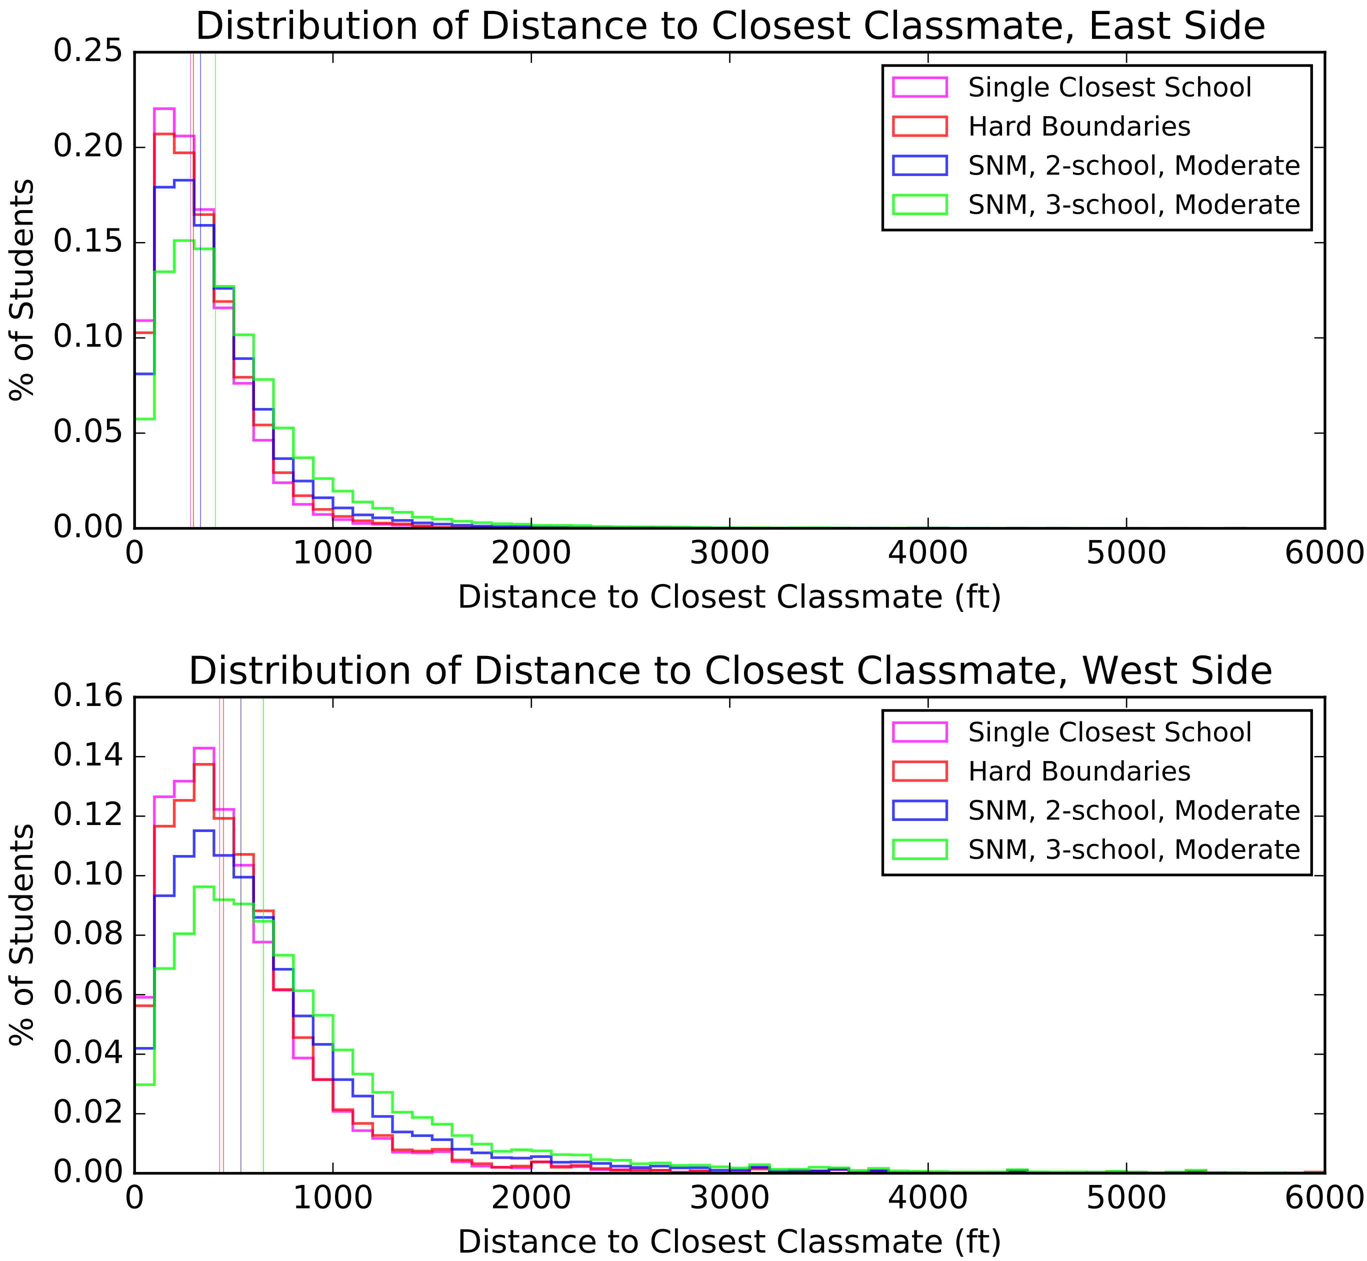

Figure 23: *Closest Classmate, No Transfers, Con-Opt Section Counts, 2008–15.* (Same data as Table 14. Vertical lines indicate median distances.) Single Closest School and Hard Boundary distances are the shortest. The SNM 2-School/moderate distances are slightly longer, and the SNM 3-School/moderate distances are the longest of the four. The most noticeable differences between the models are at distances of less than around 500–750 ft (∼1–1.5 east-side long blocks).

6.2.4 Local Assignment Diversity

Year(s)	Model, East Side No Transfers	Assignment Diversity			
		[1.0, 1.5)	[1.5, 2.5)	[2.5, 3.5)	≥3.5
2008-15	Hard Boundaries	64.8%	31.4%	3.7%	0.2%
	Single Closest School	68.4%	28.4%	3.1%	0.1%
	SNM, 3-school, Strong	30.3%	55.6%	13.3%	0.7%
	SNM, 3-school, Moderate	14.1%	48.9%	33.2%	3.9%
	SNM, 3-school, Equal	7.2%	33.0%	45.9%	13.8%
	SNM, 2-school, Strong	49.5%	45.0%	5.3%	0.3%
	SNM, 2-school, Moderate	34.9%	55.4%	9.1%	0.6%
	SNM, 2-school, Equal	23.5%	59.8%	15.5%	1.2%
Year(s)	**Model, West Side No Transfers**	**Assignment Diversity**			
		[1.0, 1.5)	[1.5, 2.5)	[2.5, 3.5)	≥3.5
2008-15	Hard Boundaries	82.3%	16.7%	0.9%	0.0%
	Single Closest School	87.8%	11.8%	0.4%	0.0%
	SNM, 3-school, Strong	37.9%	57.1%	4.8%	0.2%
	SNM, 3-school, Moderate	25.0%	53.0%	20.8%	1.1%
	SNM, 3-school, Equal	18.2%	40.8%	36.2%	4.8%
	SNM, 2-school, Strong	64.1%	34.2%	1.7%	0.0%
	SNM, 2-school, Moderate	47.0%	49.7%	3.2%	0.1%
	SNM, 2-school, Equal	29.5%	63.9%	6.3%	0.3%

Table 15: *Local Assignment Diversity, No Transfers, Con-Opt Section Counts, 2008–15.* Local assignment diversity over all seven years looks similar to 2014–15 by itself. As we noted earlier, all Soft Neighborhood models are more diverse than Hard Boundary and Single Closest School models. 3-School SNM versions have greater diversity than 2-School, but even SNM 2-School versions are an improvement over the Hard Boundary and Single Closest School models. Section 5.2 discusses the relationship between higher diversity, less gameability, and more networking between the communities of neighboring schools.

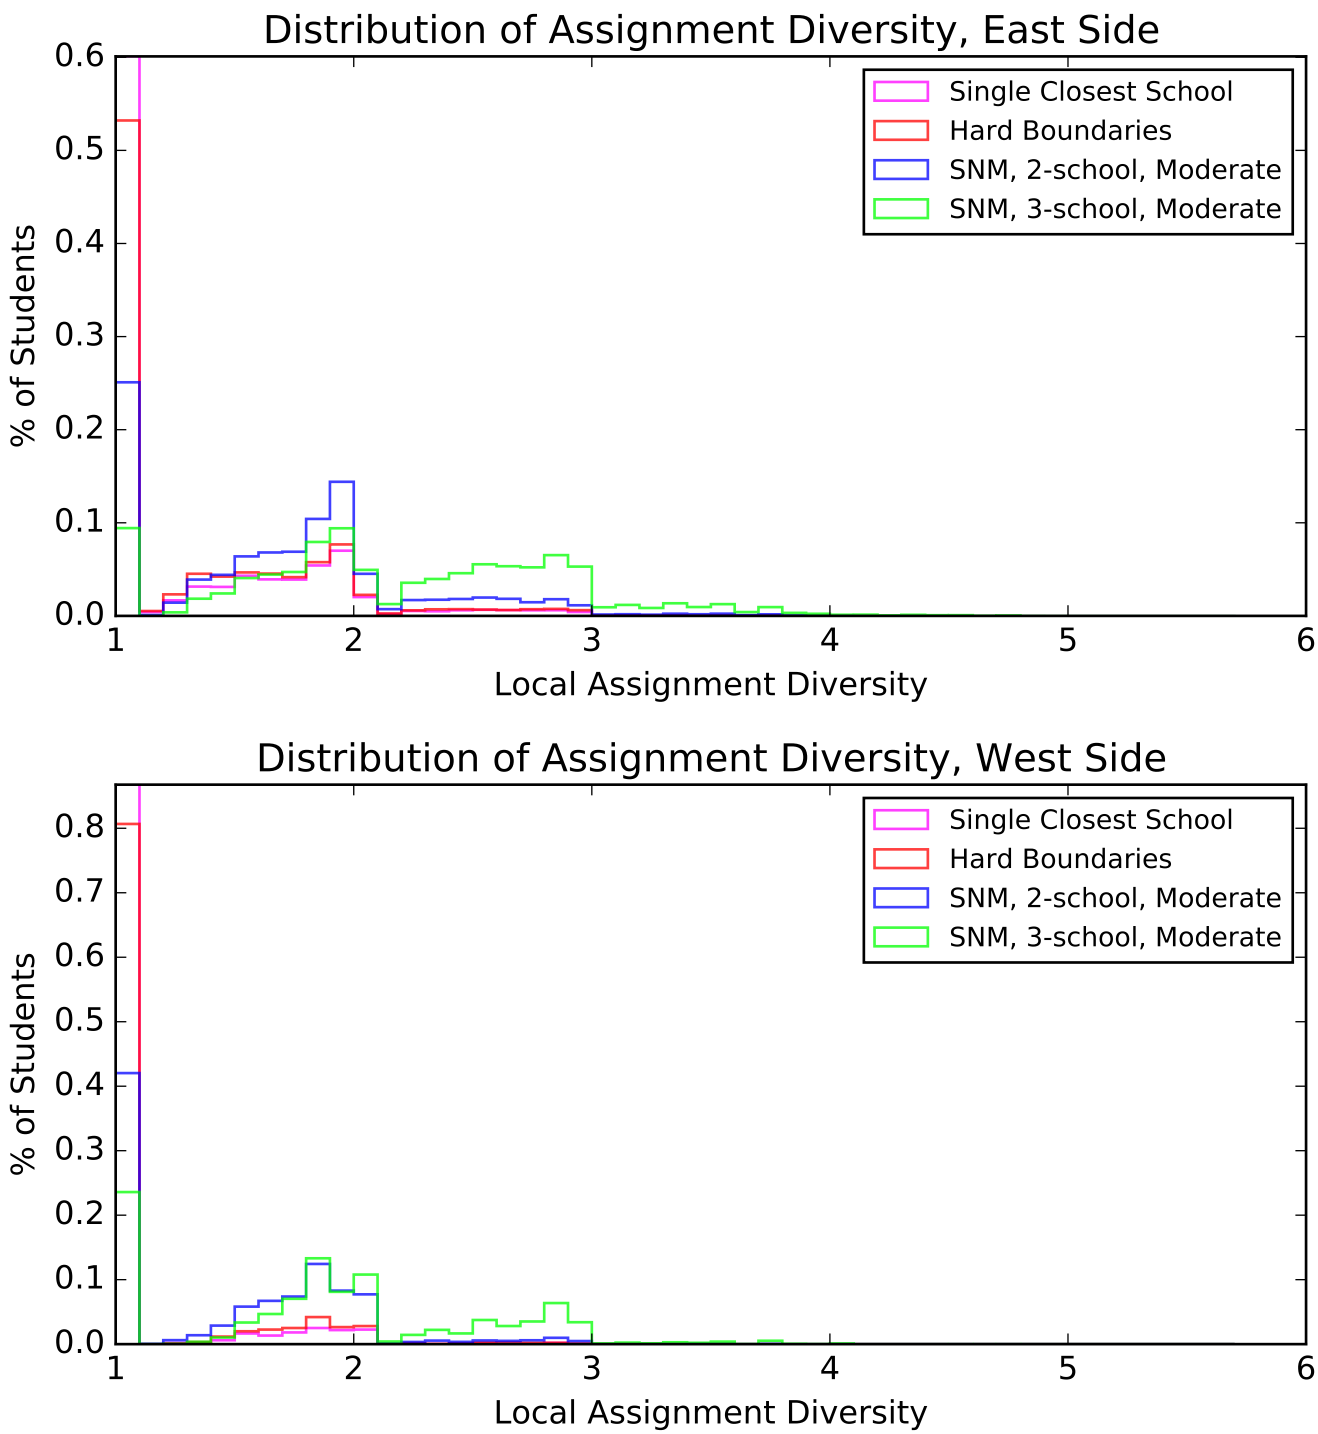

Figure 24: *Local Assignment Diversity, No Transfers, Con-Opt Section Counts, 2008–15.* (Same data as Table 15.) In Single Closest School and Hard Boundary assignments, the majority of students live in areas where only one school is well represented. Both SNM 2-School and SNM 3-School configurations result in a substantial reduction in the number of these areas, and an increase in the number of areas where 1–2 schools are well represented. In the SNM 3-school configuration, we see more areas where 2–3 schools are well-represented.

6.2.5 Capture Linkage

Model, East Side No Transfers	Capture Linkage, weighted balanced matrix					
	<1.0	[1.0, 2.0)	[2.0, 3.0)	[3.0, 4.0)	[4.0, 5.0)	≥5.0
Hard Boundaries	100.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Single Closest School	100.0%	0.0%	0.0%	0.0%	0.0%	0.0%
SNM, 3-school, Strong	3.4%	11.4%	27.9%	35.9%	16.8%	4.7%
SNM, 3-school, Moderate	3.4%	6.4%	20.1%	32.6%	26.8%	10.7%
SNM, 3-school, Equal	3.4%	2.7%	18.1%	27.2%	29.5%	19.1%
SNM, 2-school, Strong	34.6%	41.6%	19.8%	3.7%	0.3%	0.0%
SNM, 2-school, Moderate	30.5%	38.9%	24.2%	5.7%	0.7%	0.0%
SNM, 2-school, Equal	29.2%	39.9%	24.5%	5.7%	0.7%	0.0%

Model, West Side No Transfers	Capture Linkage, weighted balanced matrix					
	<1.0	[1.0, 2.0)	[2.0, 3.0)	[3.0, 4.0)	[4.0, 5.0)	≥5.0
Hard Boundaries	100.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Single Closest School	100.0%	0.0%	0.0%	0.0%	0.0%	0.0%
SNM, 3-school, Strong	13.0%	11.7%	19.5%	28.6%	26.0%	1.3%
SNM, 3-school, Moderate	11.7%	10.4%	20.8%	15.6%	27.3%	14.3%
SNM, 3-school, Equal	11.7%	5.2%	23.4%	15.6%	19.5%	24.7%
SNM, 2-school, Strong	29.9%	48.1%	20.8%	1.3%	0.0%	0.0%
SNM, 2-school, Moderate	24.7%	41.6%	29.9%	3.9%	0.0%	0.0%
SNM, 2-school, Equal	26.0%	39.0%	29.9%	5.2%	0.0%	0.0%

Table 16: *Capture Linkage, No Transfers, Con-Opt Section Counts, 2008–15.* As we noted earlier, Hard Boundary and Single Closest School models always by definition have zero capture linkage (Section 5.2). Similar to what we saw with the 2014–15 results, in SNM 3-School versions, capture linkage is more skewed towards higher numbers than 2-School versions, and the east side has higher capture linkage than the west side due to its higher density, though the differences aren't as stark over all seven years as they were in 2014–15. In terms of gameability, SNM 3-School versions are less susceptible to gaming. And, SNM 3-School versions will drive more demographic similarity and population sharing than 2-School versions. However, though weaker than 3-School, SNM 2-School versions are still a major improvement over a Hard Boundary model.

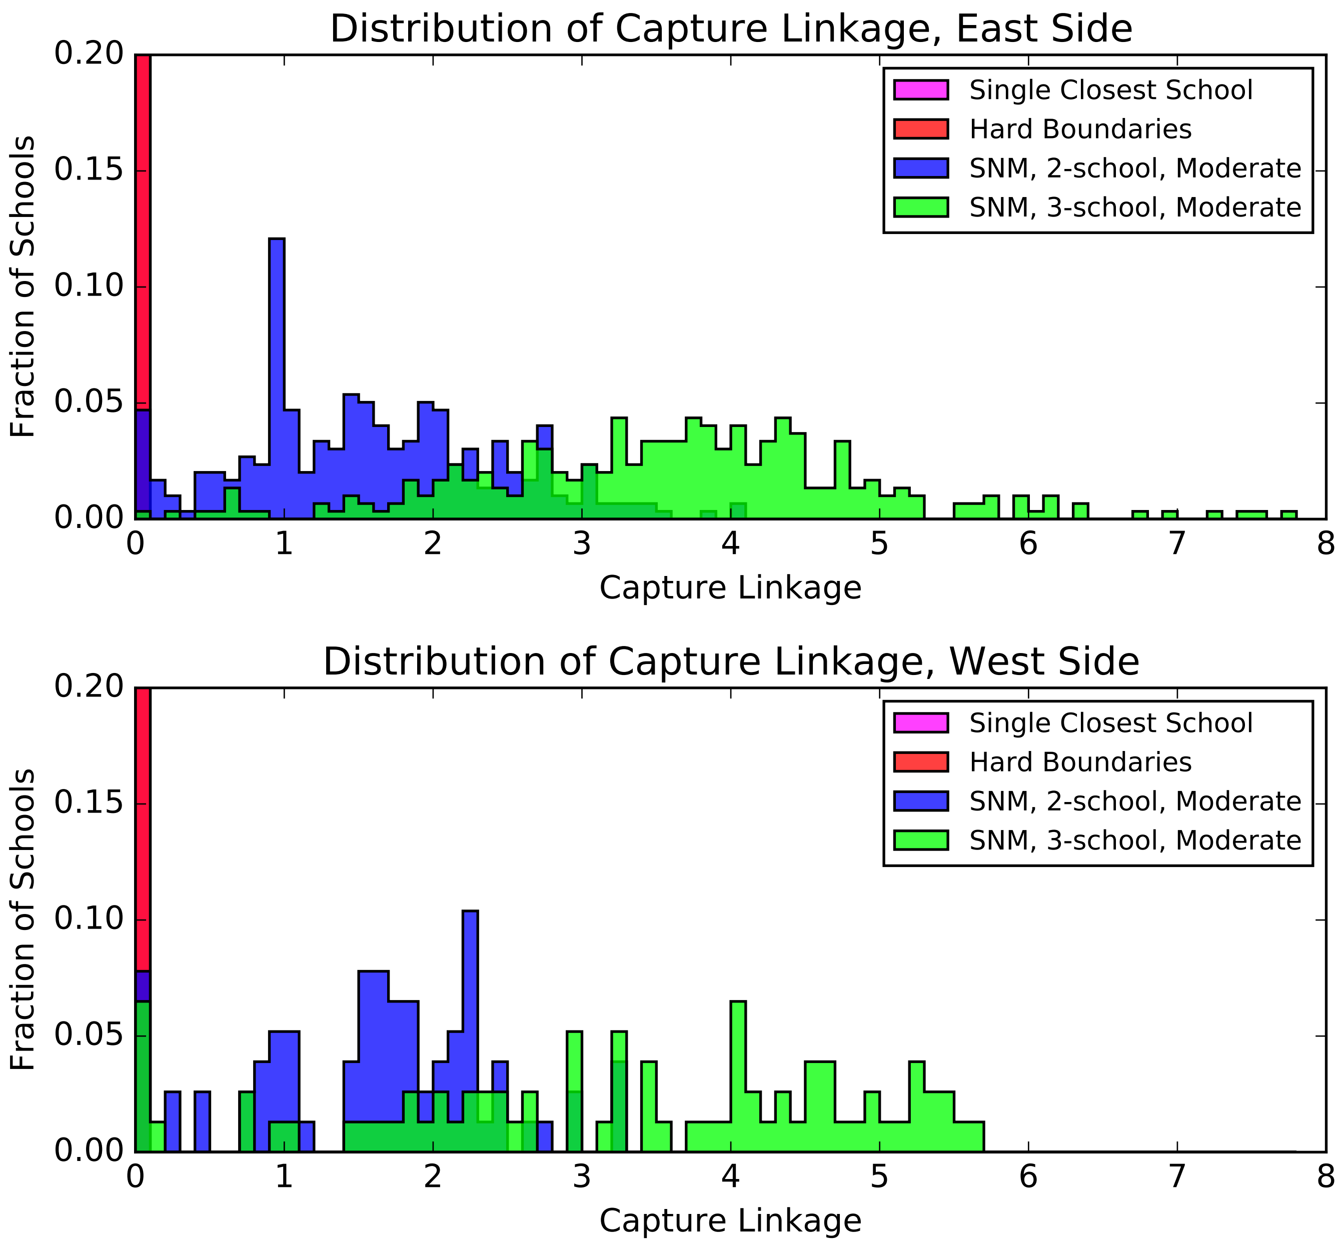

Figure 25: *Capture Linkage, No Transfers, Con-Opt Section Counts, 2008–15.* (Same data as Table 16.) In Single Closest School and Hard Boundary models, school linkages are zero for all schools. Soft Neighborhood Model capture linkages are higher, and 3-School is quite a bit higher than 2-School. *Fractions for Hard Boundaries and Single Closest School at zero-linkage are 1.0, but have been clipped to 0.20 in the plots above.*

7 Discussion and Sample Assignments

First, we would like to make a few general observations, based on the results presented in Section 6. Following that, we will present a number of assignment graphs, to give the reader a visual sense of what the Soft Neighborhood Model could do in PPS, and, conversely, what the PPS hard boundaries have already been doing (or not doing, as the case may be).

1. **The Soft Neighborhood Model nails enrollment balancing.** The results in Sections 6.1.1 and 6.2.1 are pretty clear: enrollment balancing is a major strength of the Soft Neighborhood Model. On the east side, on the west side, in a single year, over many years — the Soft Neighborhood Model does what it has been designed to do: fill sections in a school district as uniformly as possible, so that no school has too many or too few students relative to other schools, while keeping students in schools close to home. The Hard Boundary and Single Closest School models, on the other hand, effectively guarantee over- and underenrollment.

2. **The Soft Neighborhood Model can be configured to have a distance profile that is very similar to the Hard Boundary model. . . at the expense of population mixing.** Using two schools instead of three, and a moderate to strong preference for the closest school, the Soft Neighborhood Model's distance profile is just about the same as the Hard Boundary model's. But when we move from a three-school model to a two-school model, population mixing (measured by Local Assignment Diversity and Capture Linkage) takes a hit. Population mixing is how the Soft Neighborhood Model counters the concentration of wealth and poverty. It's a trade-off: to get more population mixing, we incur slightly longer commutes to school. But even a 3-School Soft Neighborhood Model comes nowhere close to busing kids all over town.

3. **A Soft Neighborhood Model configured with more mixing probably costs more in busing...** To get more mixing, kids do travel a little farther to school, but it's a matter of degrees, not a matter of everyone having to criss-cross around town. When we crank up the mixing, more kids technically become bus-eligible, but we (the authors) don't have any information about how eligibility translates into ridership, nor do we know how the Transportation Department does its route planning/optimization, nor we do know how they do their budgeting.

4. **...but, the Soft Neighborhood Model saves dollars in space efficiency.** The SNM *nails* the enrollment balancing problem — with consequent gains from space-efficiency and right-sizing each classroom. Any estimate of changes in transportation costs needs to be accompanied by an estimate of the value of both the improved educational outcomes and the real-dollars-saved by filling schools with the right numbers of students and teachers. (Not to mention the value of avoiding regular boundary adjustments.)

5. **The Soft Neighborhood Model does not guarantee equitable programming, but it is consistent with that goal.** Having the right number of students in each section at each school supports the goal of equitable programming across the district

— it may even be a prerequisite. Right-sizing sections (a) more equitably distributes funds to all schools, and (b) enables the economies of scale necessary to offer a diverse array of classes.

The rest of this section presents a series of *assignment graphs*, illustrating the action of different assignment models on the 2014 neighborhood-school kindergarten population of PPS. These graphs are made up of three elements:

- Each student is represented by a tiny gray dot located at her home address.
- Each school is represented by a gray box with an abbreviated name for the school. Each box has two numbers — a mean section size on top, and section count on the bottom. For example, "27.3 / Bri / x3" means *Bridlemile, with 3 sections, averaging 27.3 students per section* (i.e., section counts of 27, 27, 28).
- Each student is connected to her assigned school by a colored line.[31] Some graphs show *possible* assignments, in which case a student may be connected to multiple schools.

Don't print these out! You'll just waste paper. The best way to view them is on a computer screen, where you can zoom in and scroll around in the details. Pay attention to how section sizes vary from school to school. In any case, the goal here is not to analyze individual student placements, but rather to get a sense of the *systems* engendered by each assignment model.

Reflecting Section 6, the SNM scenarios here have river-crossing penalized to avoid cross-river assignments. However, in the last two graphs, we have relaxed that constraint to give our readers an idea of what cross-river assignments could look like. Please keep in mind that these assignments are made considering only *travel distance*, and not any other factors such as travel time or reliability.

[31]The colored line represents an *assignment*, not a measured distance. Distances are measured using actual driving routes.

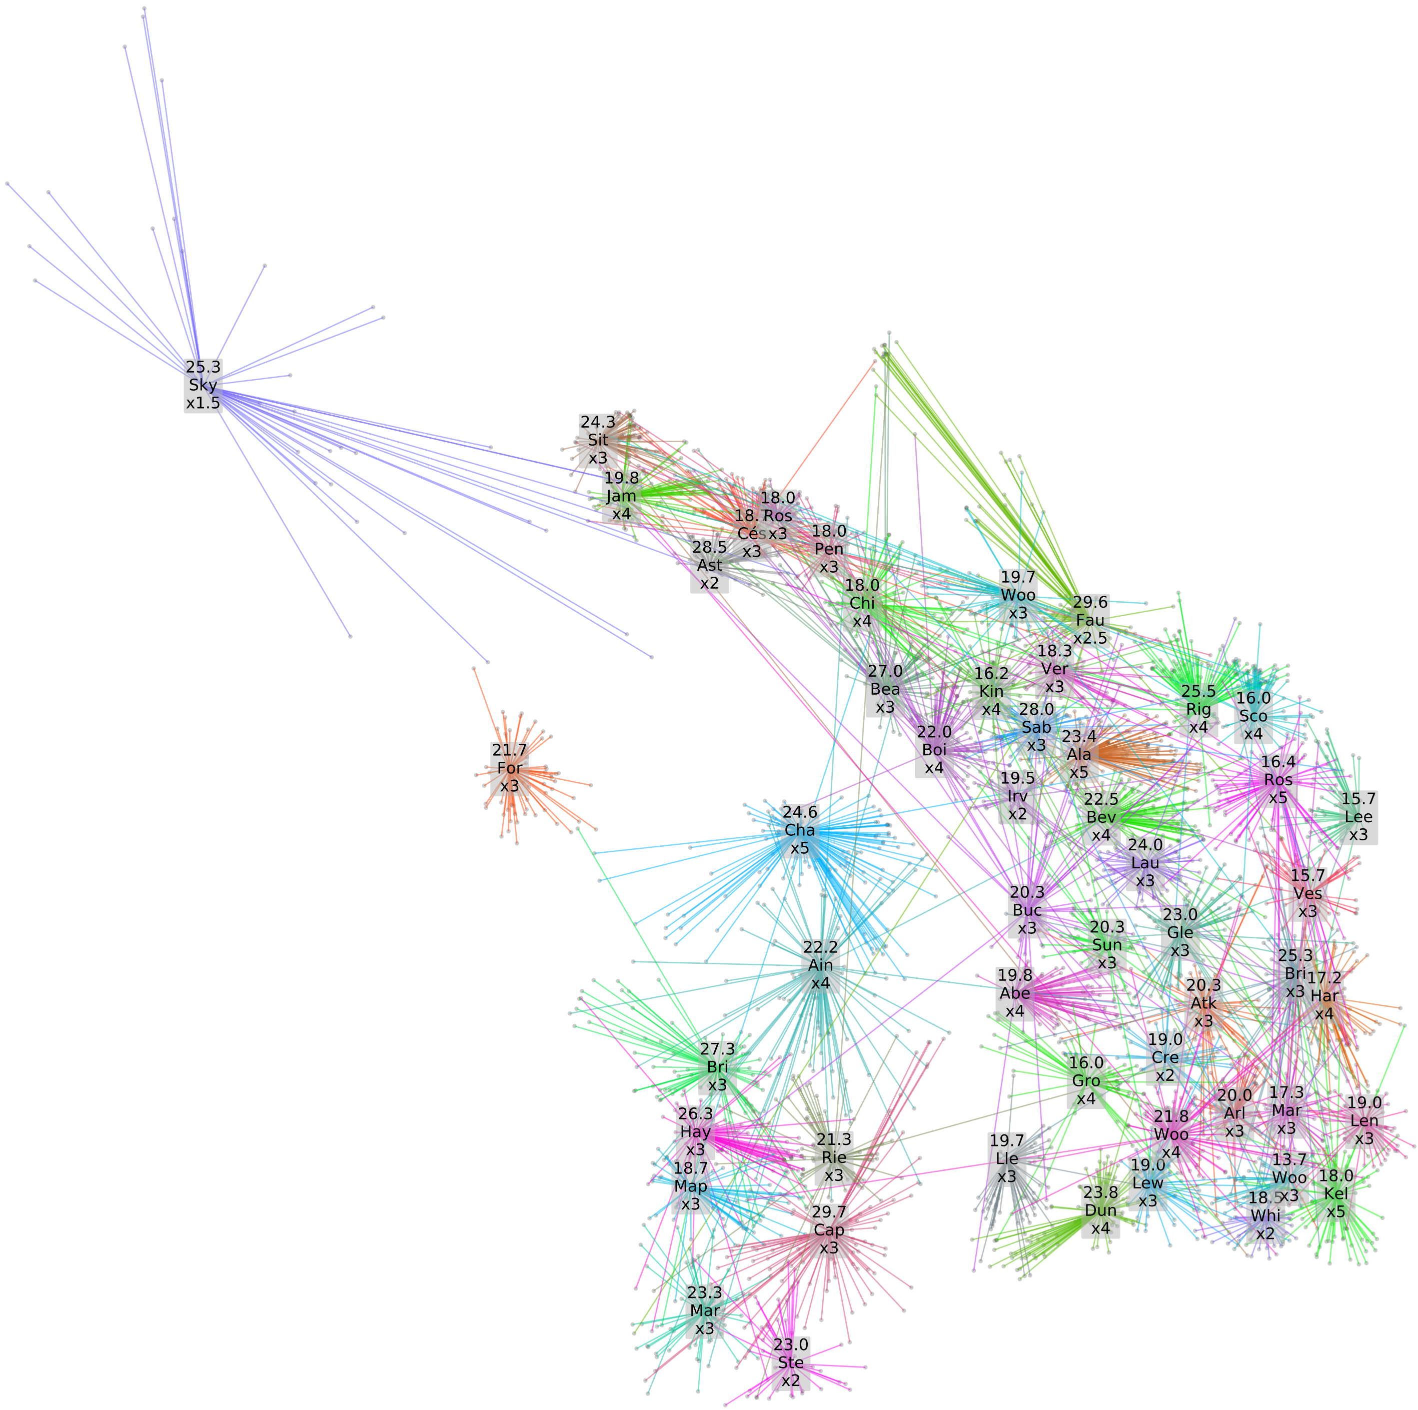

Figure 26: *Historic Hard Boundary assignment with transfers, True-2104.* This is how the existing hard boundaries actually assigned students in 2014–15 — with the help of the neighborhood lottery and petition transfers. To the best of our knowledge, these (“xN”) are the actual section counts (and corresponding section sizes) for that year. When looking at section sizes, keep in mind: the district has had to add/remove/combine sections in some schools to deal with population imbalances; transfers also act to shift kids from overcrowded schools to underenrolled ones.

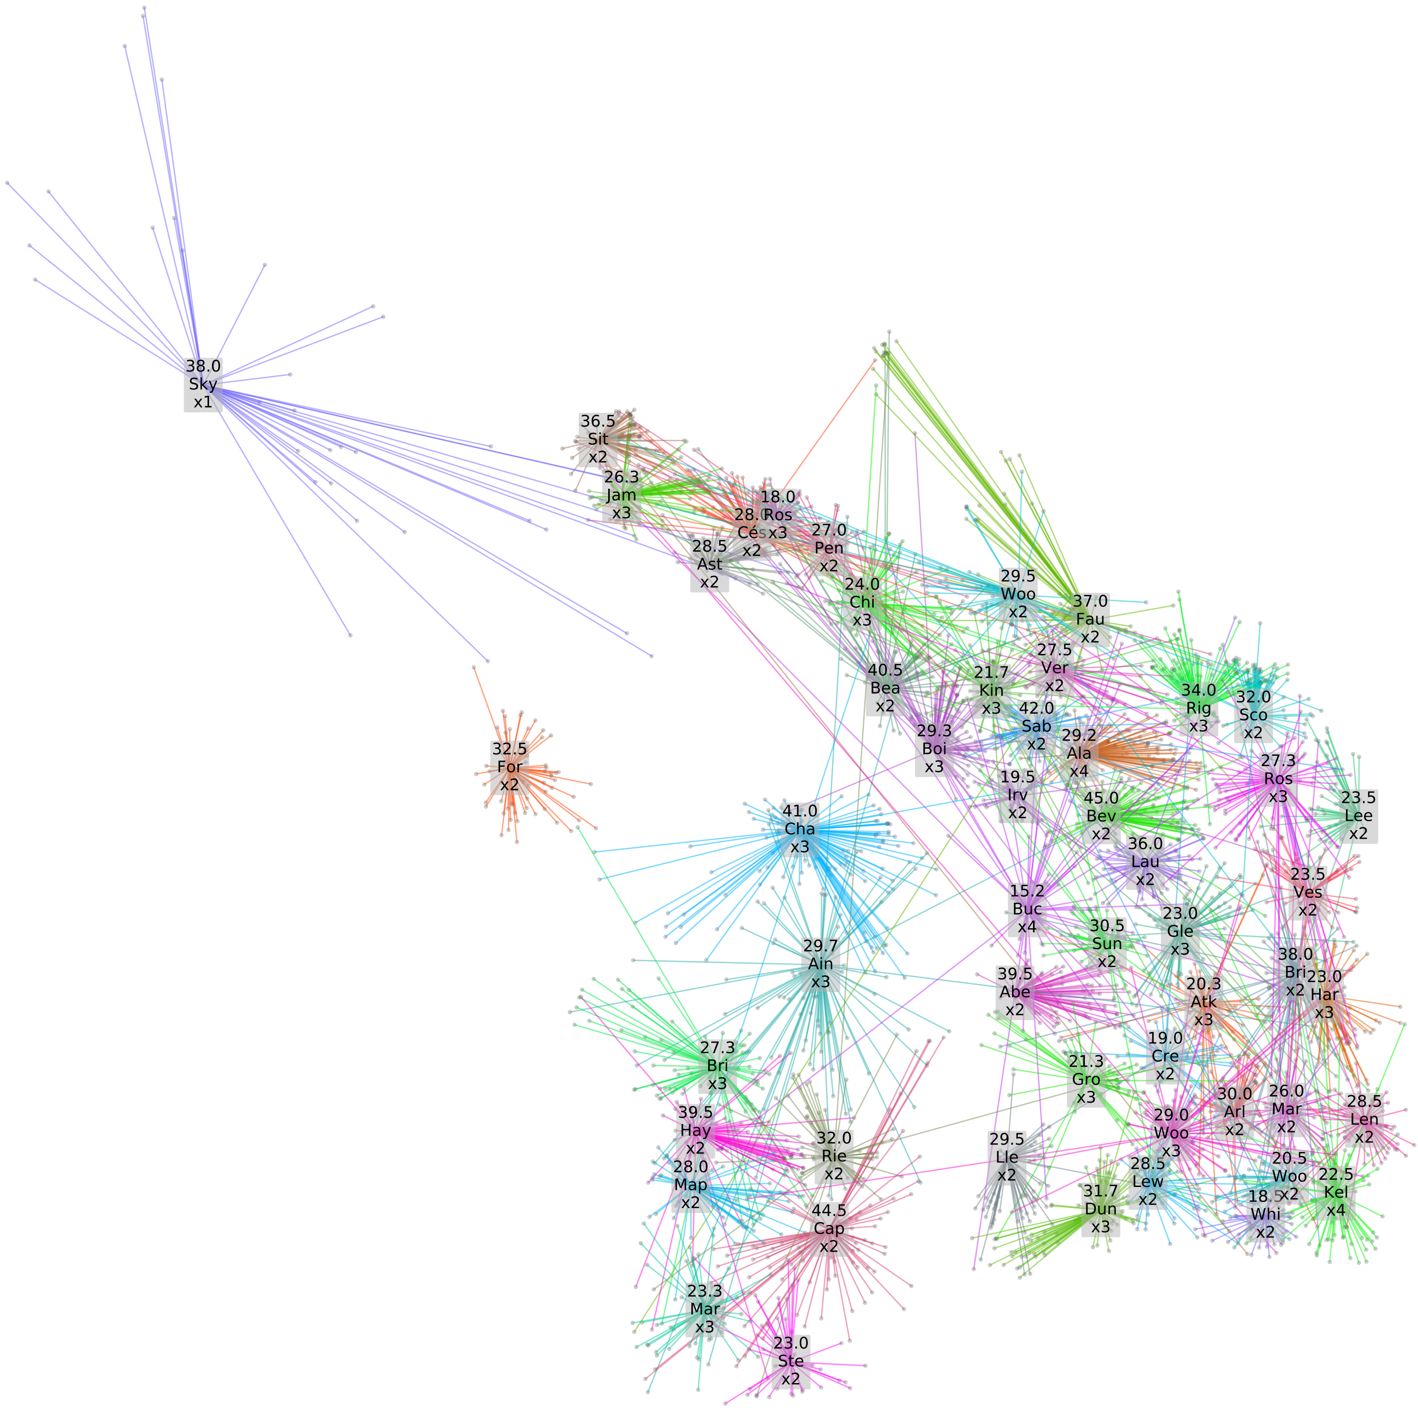

Figure 27: *Historic Hard Boundary assignment with transfers, Con-Opt.* Here are the same assignments as the previous page, but using *Conservative-Optimal* section counts — only the section counts and sizes have changed. Without finding space for kids in boiler rooms in some schools, or leaving rooms vacant in others, section sizes have become unbearable in many *(but not all)* schools. However, this graph retains all lottery and petition transfers, some of which would not be allowed if a school were already overcrowded.

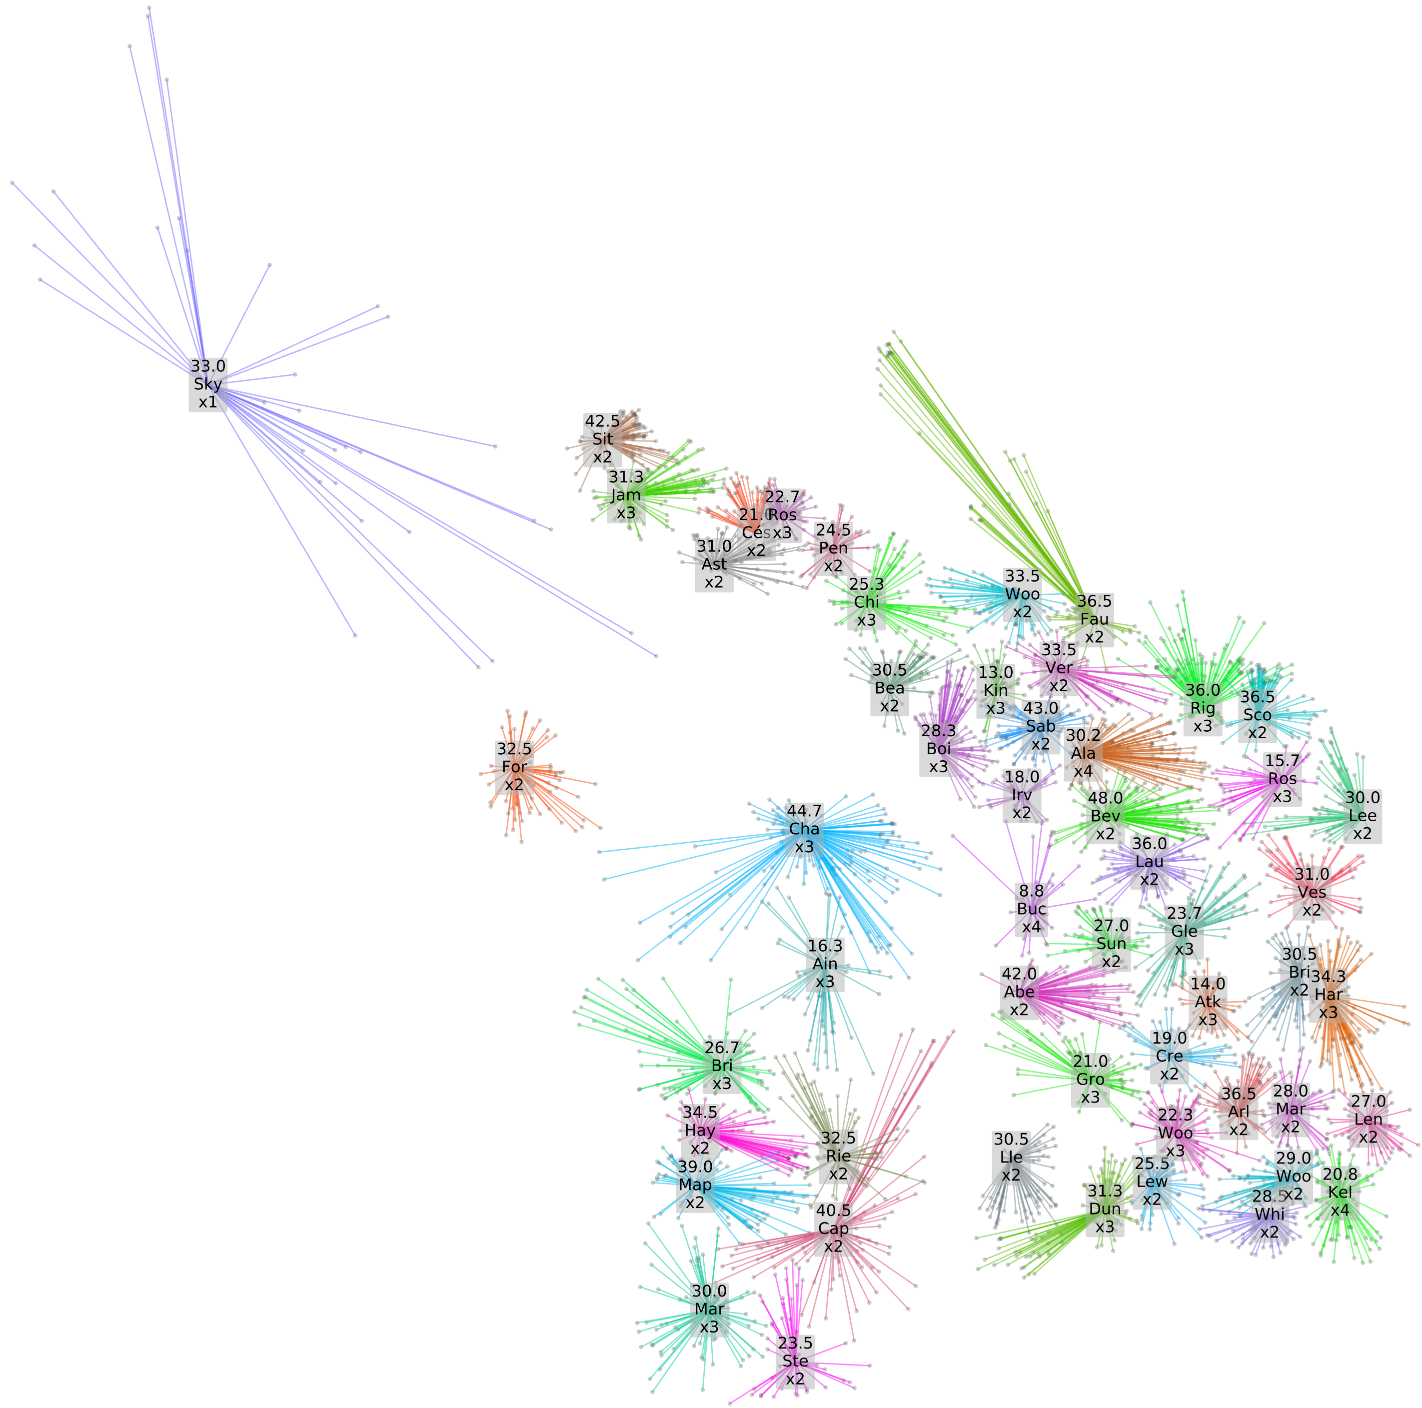

Figure 28: *Historic Hard Boundary assignment with no transfers, Con-Opt.* Now we have eliminated all transfers, in line with the SACET recommendations eliminating neighborhood-to-neighborhood lottery transfers starting in 2015–16. This assignment graph is what the hard boundaries produce when they aren't helped by distorted section counts or shifting from transfers; it is a mess of unbalanced enrollment. Note also how isolated each school is. The neighborhood schools are a collection of geographically distinct fiefdoms.

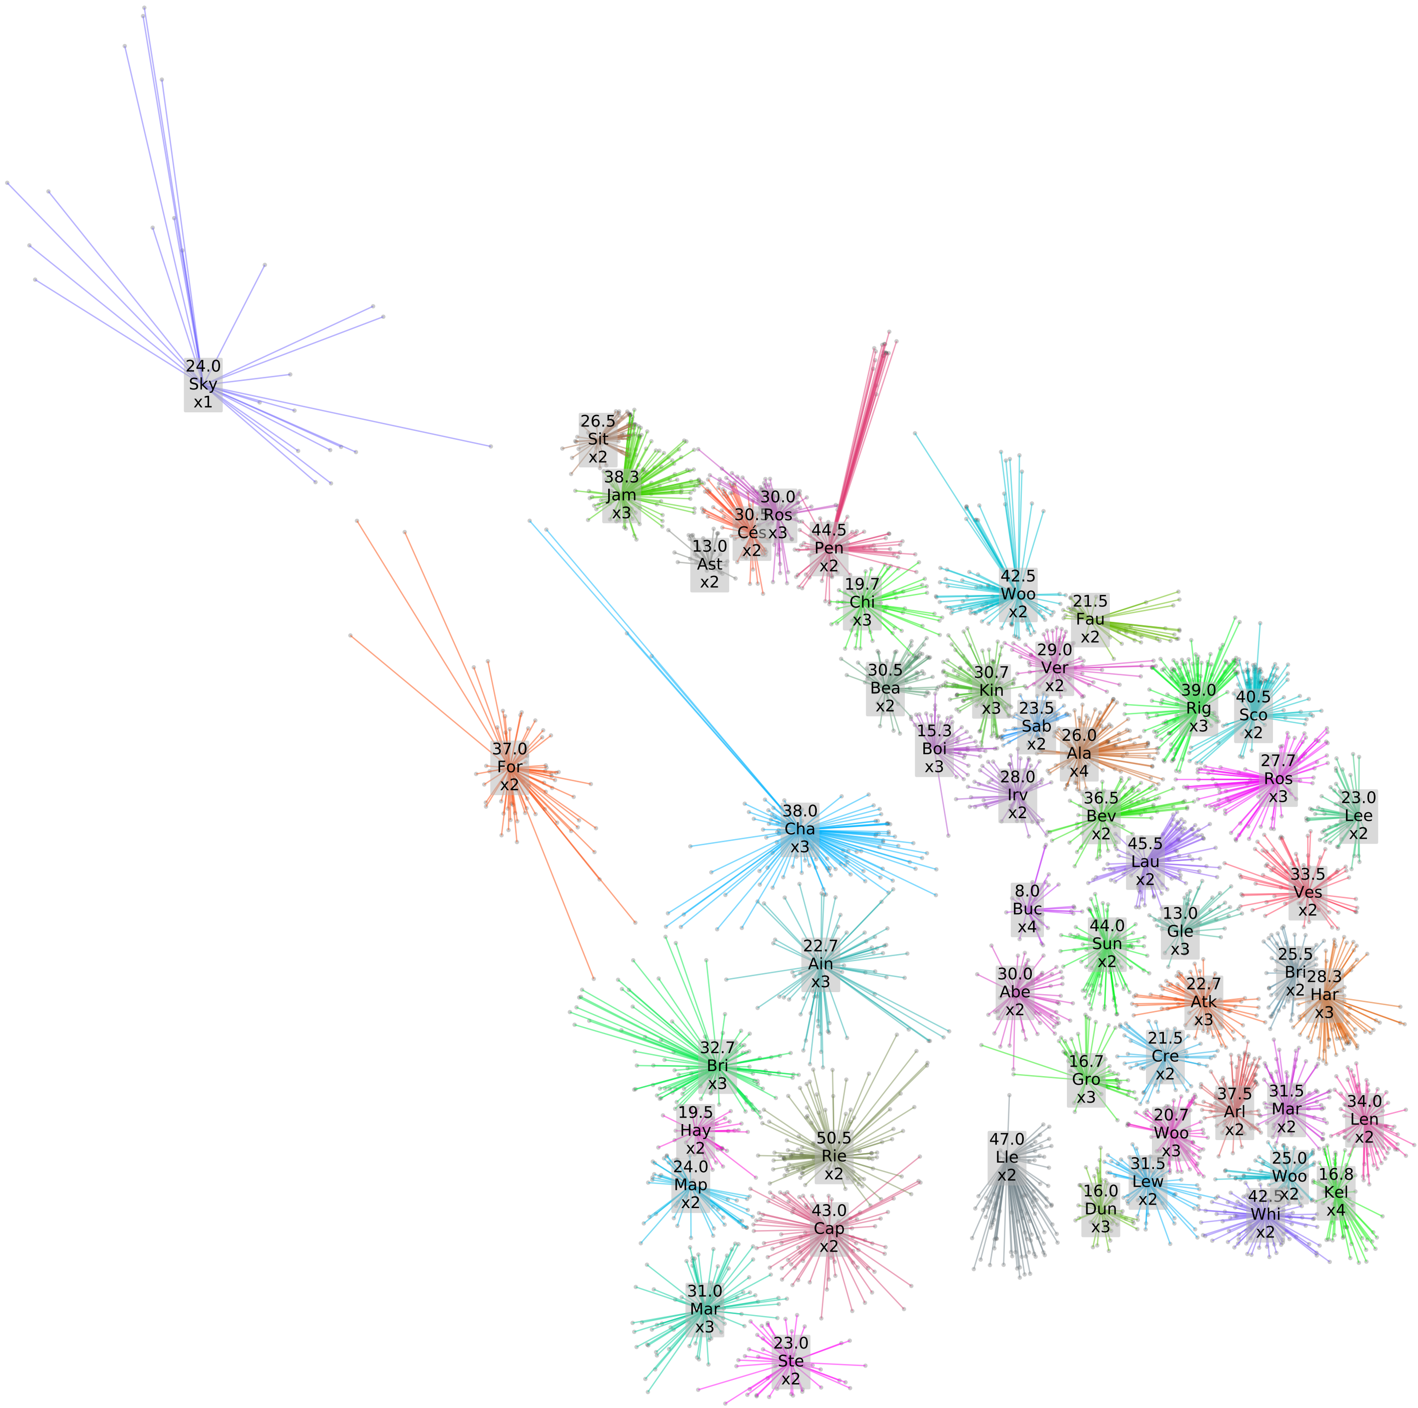

Figure 29: *Single Closest School assignment, Con-Opt, no transfers.* This is what would happen if PPS ditched its hard boundaries and simply assigned every student to her single closest school (by driving distance). It doesn't look very different from the previous picture: the sections are all out of balance, and each school's population is still a distinct, isolated blob of color. On the plus side, travel distances are minimized, and no one has to spend any time or money arguing about the boundaries!

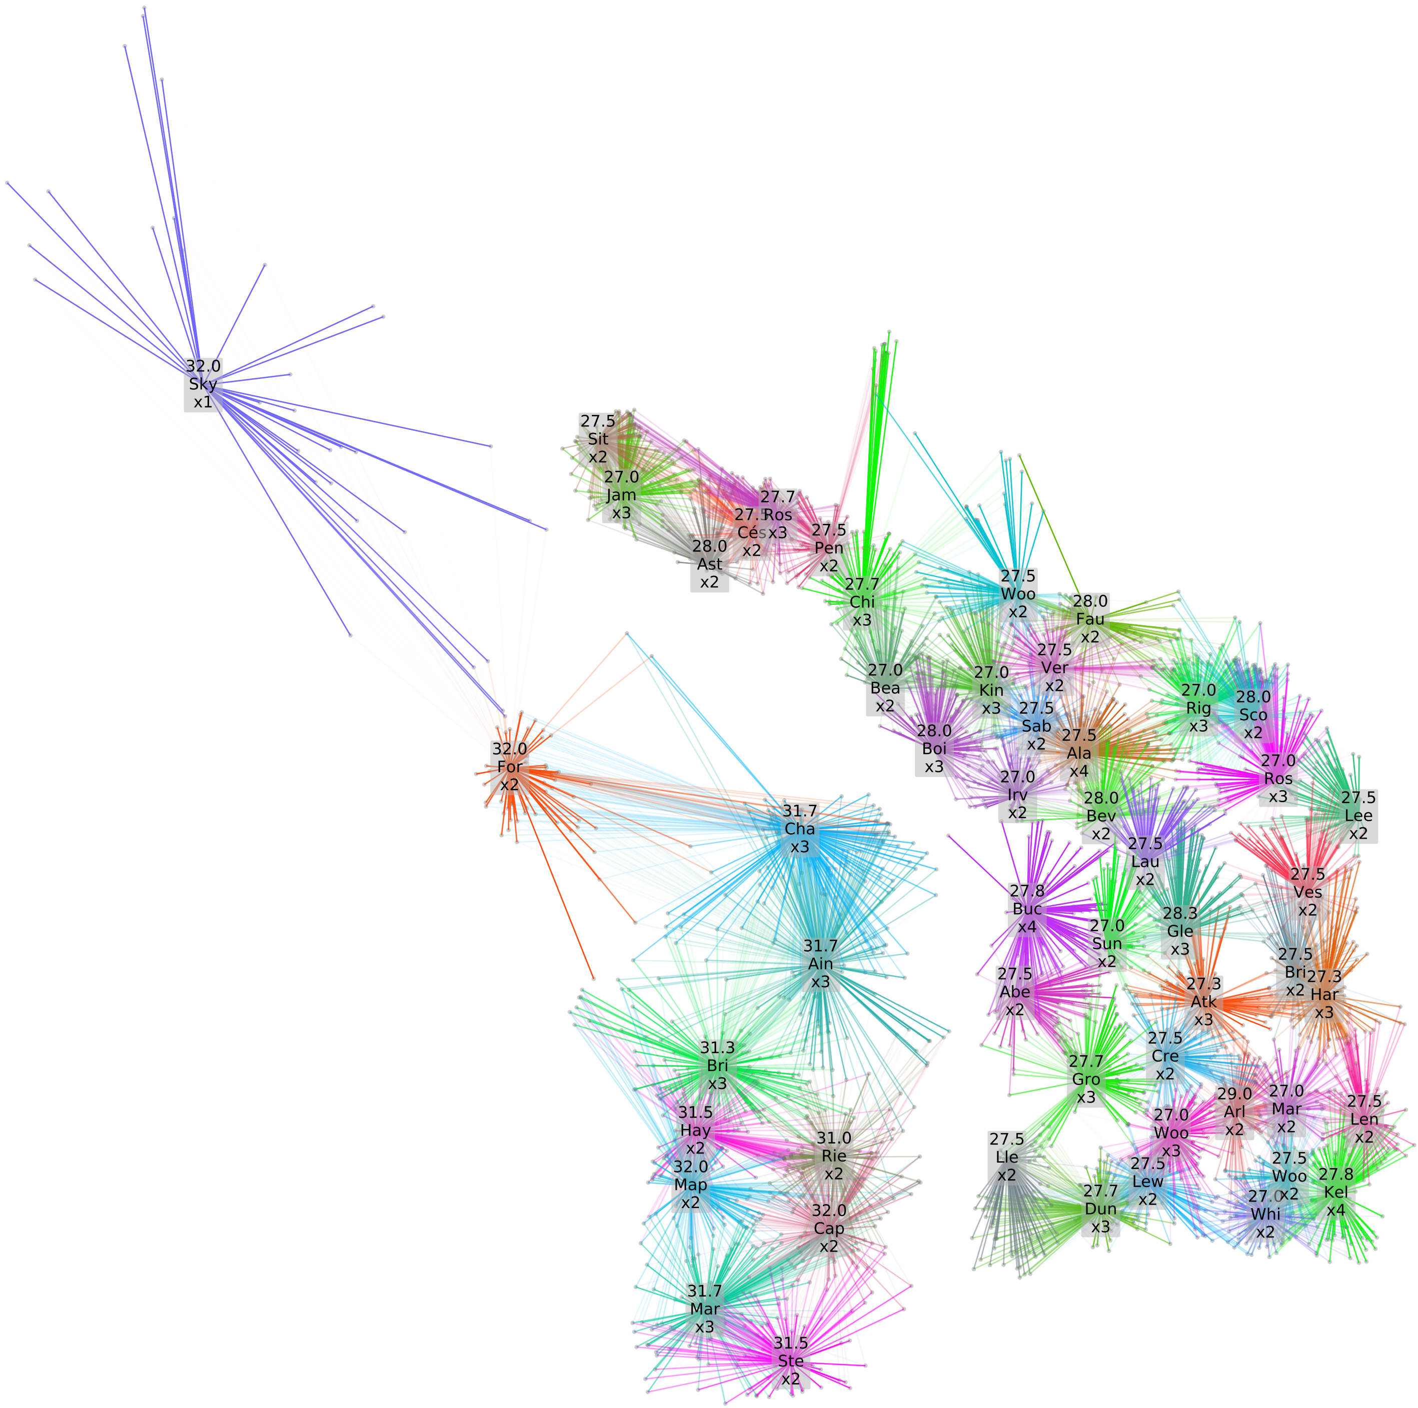

Figure 30: *SNM with 3-Schools/moderate-proximity, balanced matrix, Con-Opt, no transfers.* This graph is a SNM balanced probability matrix, showing all *possible* assignments for the given configuration of students and schools. Actual assignments are random for each student; darker lines are more likely, lighter lines are less likely. Despite the randomness, the section sizes displayed with each school are guaranteed for any assignment generated by the process, and they are all extremely well balanced. Here, every student began with the 3 closest schools as possibilities, but the balancing step has adjusted the probabilities to ensure balanced enrollments.

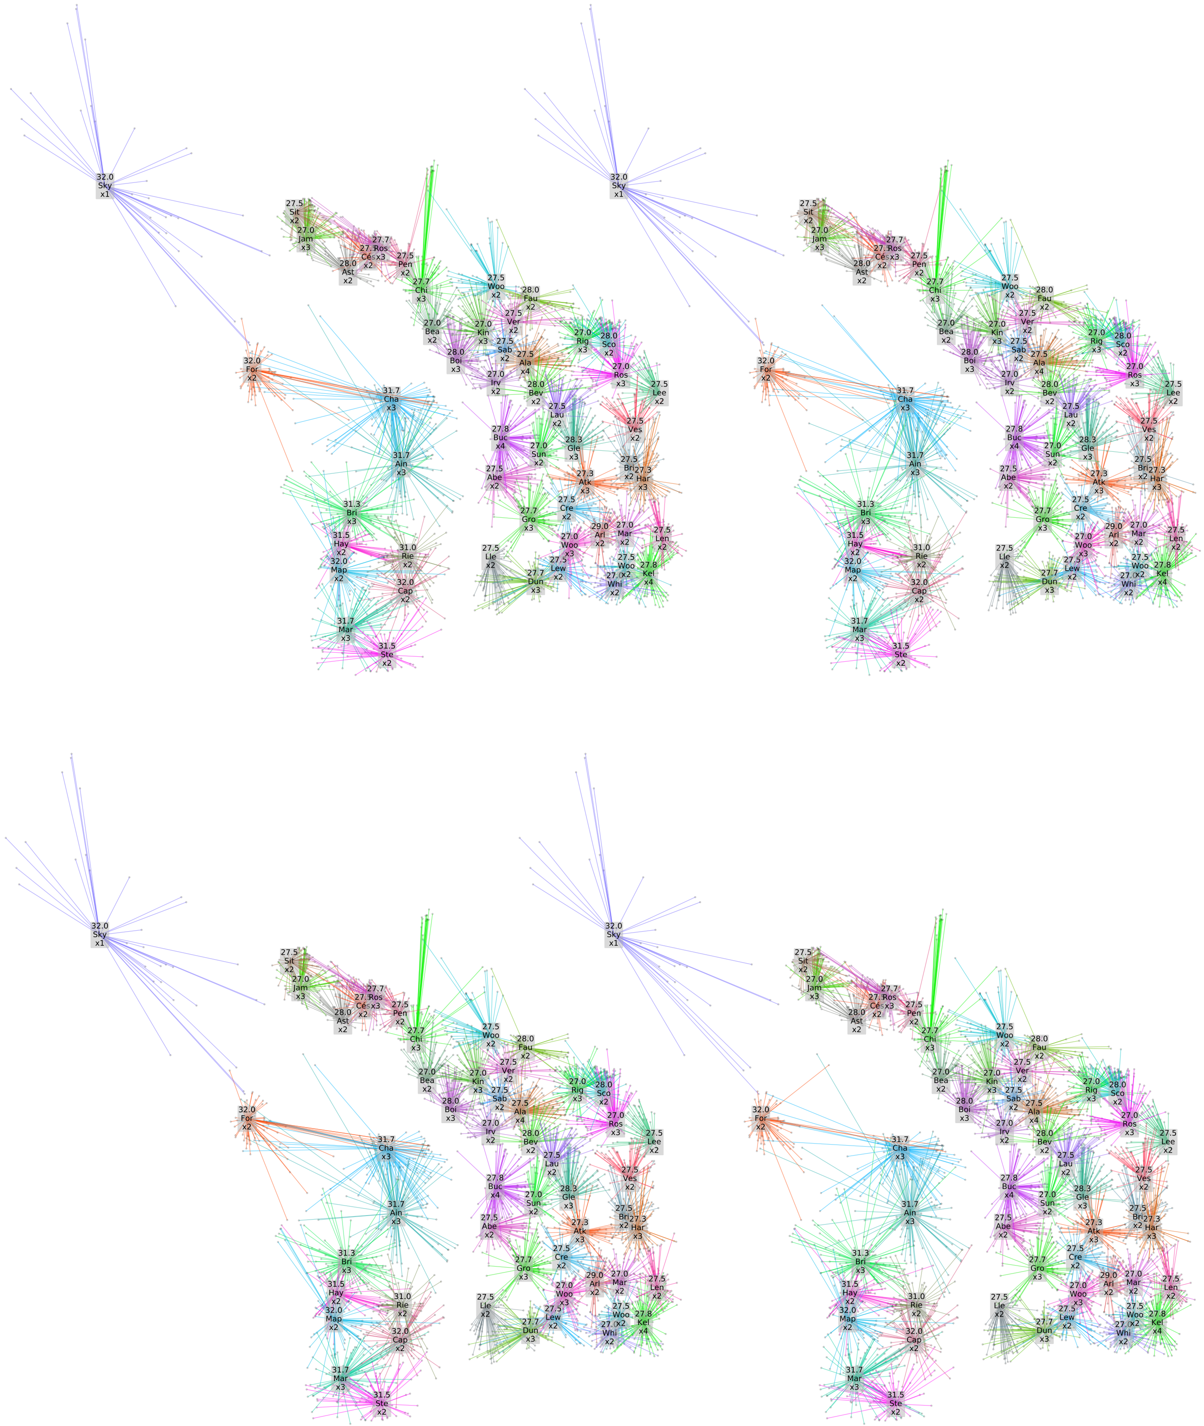

Figure 31: *Four assignments from the same balanced matrix.* Here are four possible SNM assignments, out of many, made by randomly drawing from the balanced matrix in Figure 30. Each outcome is guaranteed to have the same enrollment numbers as decided by the balanced matrix, but if you look closely you will see that different students may be assigned to each school to achieve those results.

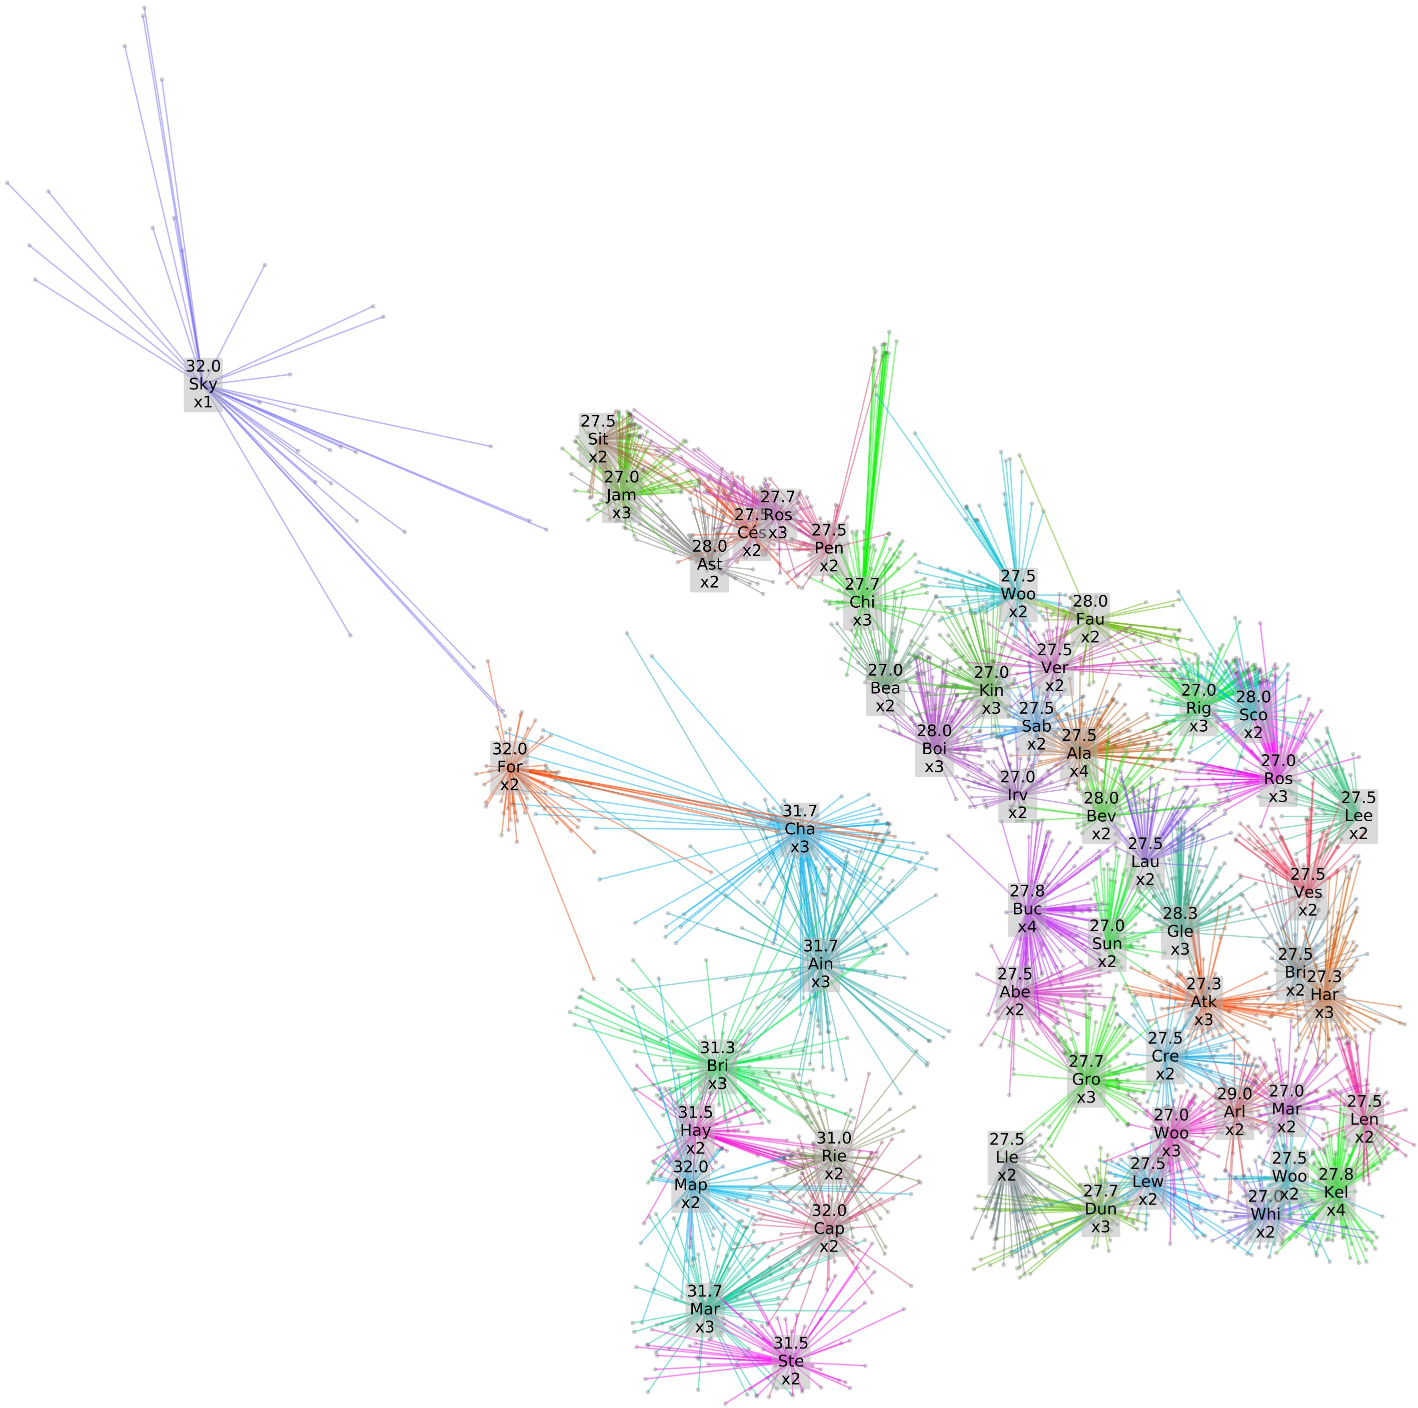

Figure 32: *SNM with 3-Schools/moderate-proximity Con-Opt, no transfers, assignment.* This is a graph of *one possible* SNM assignment for 2014–15 made from the balanced matrix in Figure 30, choosing from 3 closest schools with a moderate proximity preference. Compare with the Hard Boundary assignment of Figure 28. The SNM section sizes are incredibly well-balanced. The SNM school populations are geographically mixed and mingled.

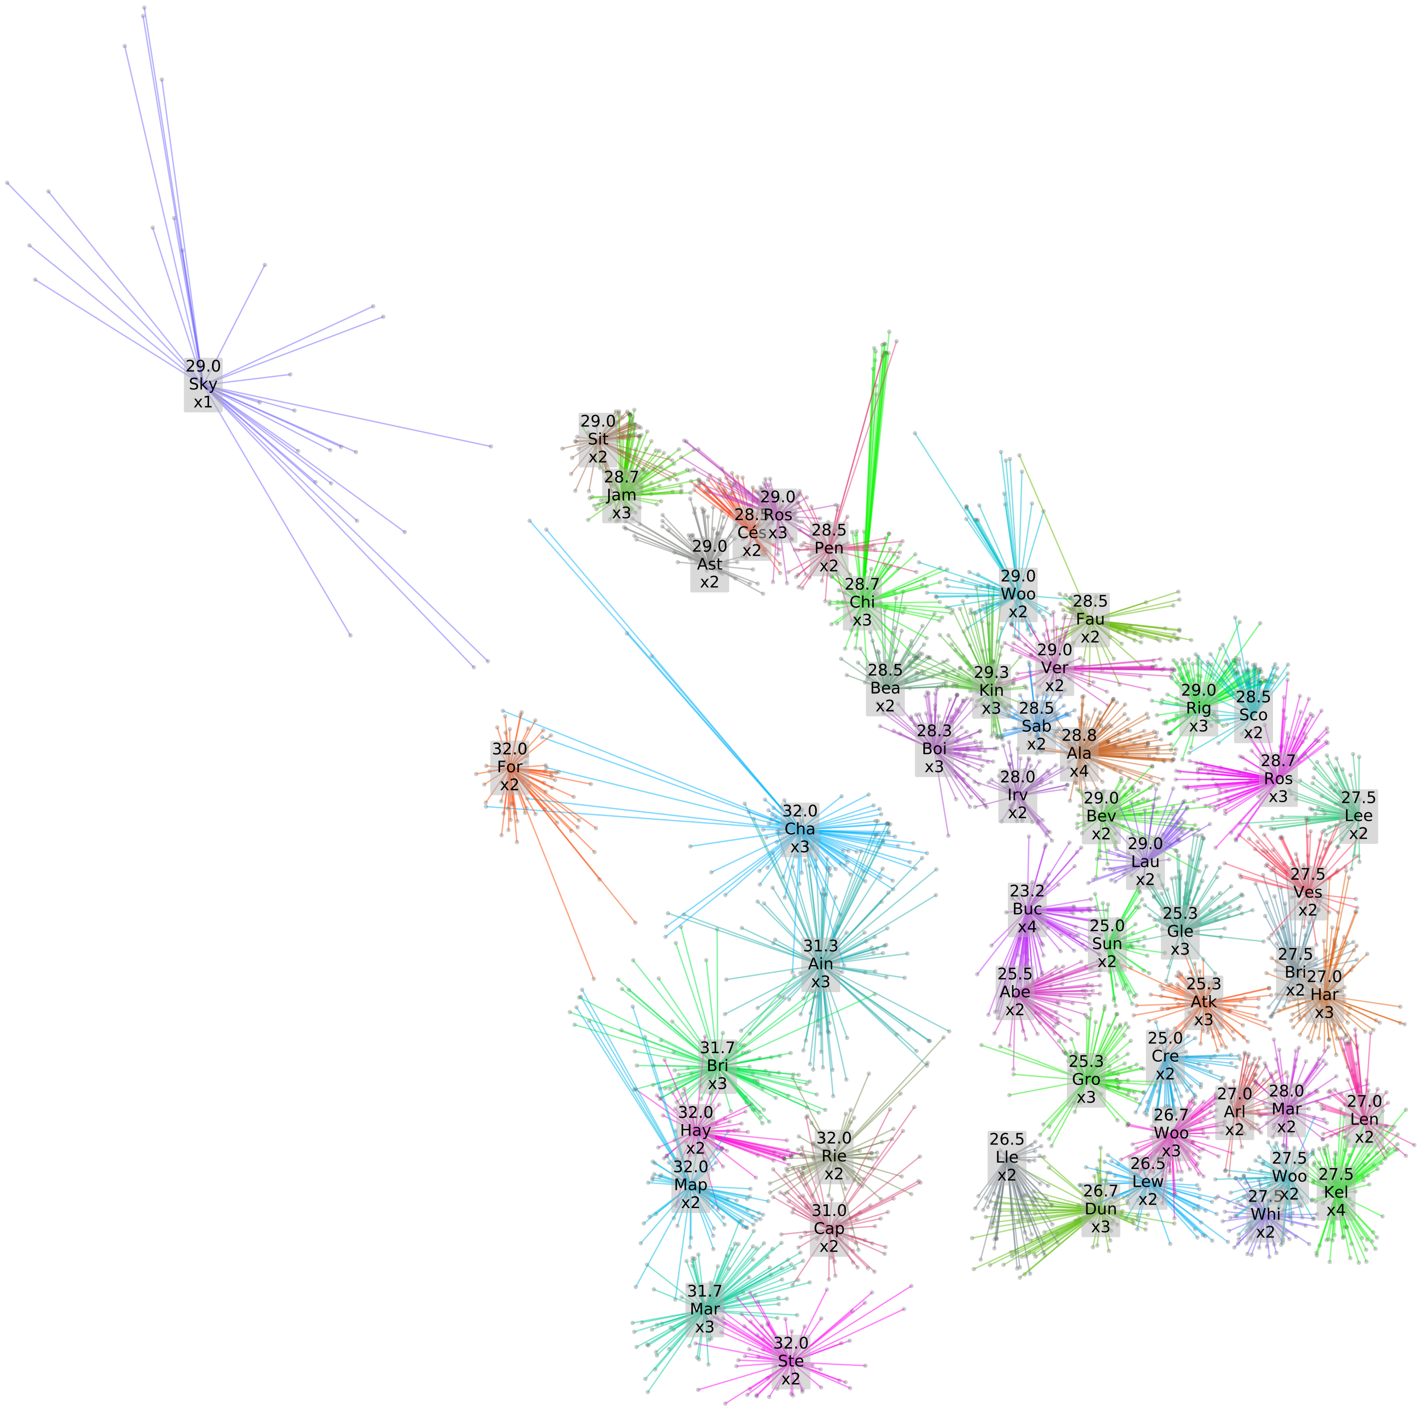

Figure 33: *SNM with 2-Schools/moderate-proximity Con-Opt, no transfers, assignment.* This is a graph of *one possible* assignment SNM assignment for 2014–14 made under the 2-Schools/moderate-proximity configuration. Compared to Hard Boundaries (Figure 28), the balancing is quite good, but not as good as the 3-Schools/moderate configuration (Figure 32). Likewise, the populations are visibly not as well blended as 3-Schools/moderate.

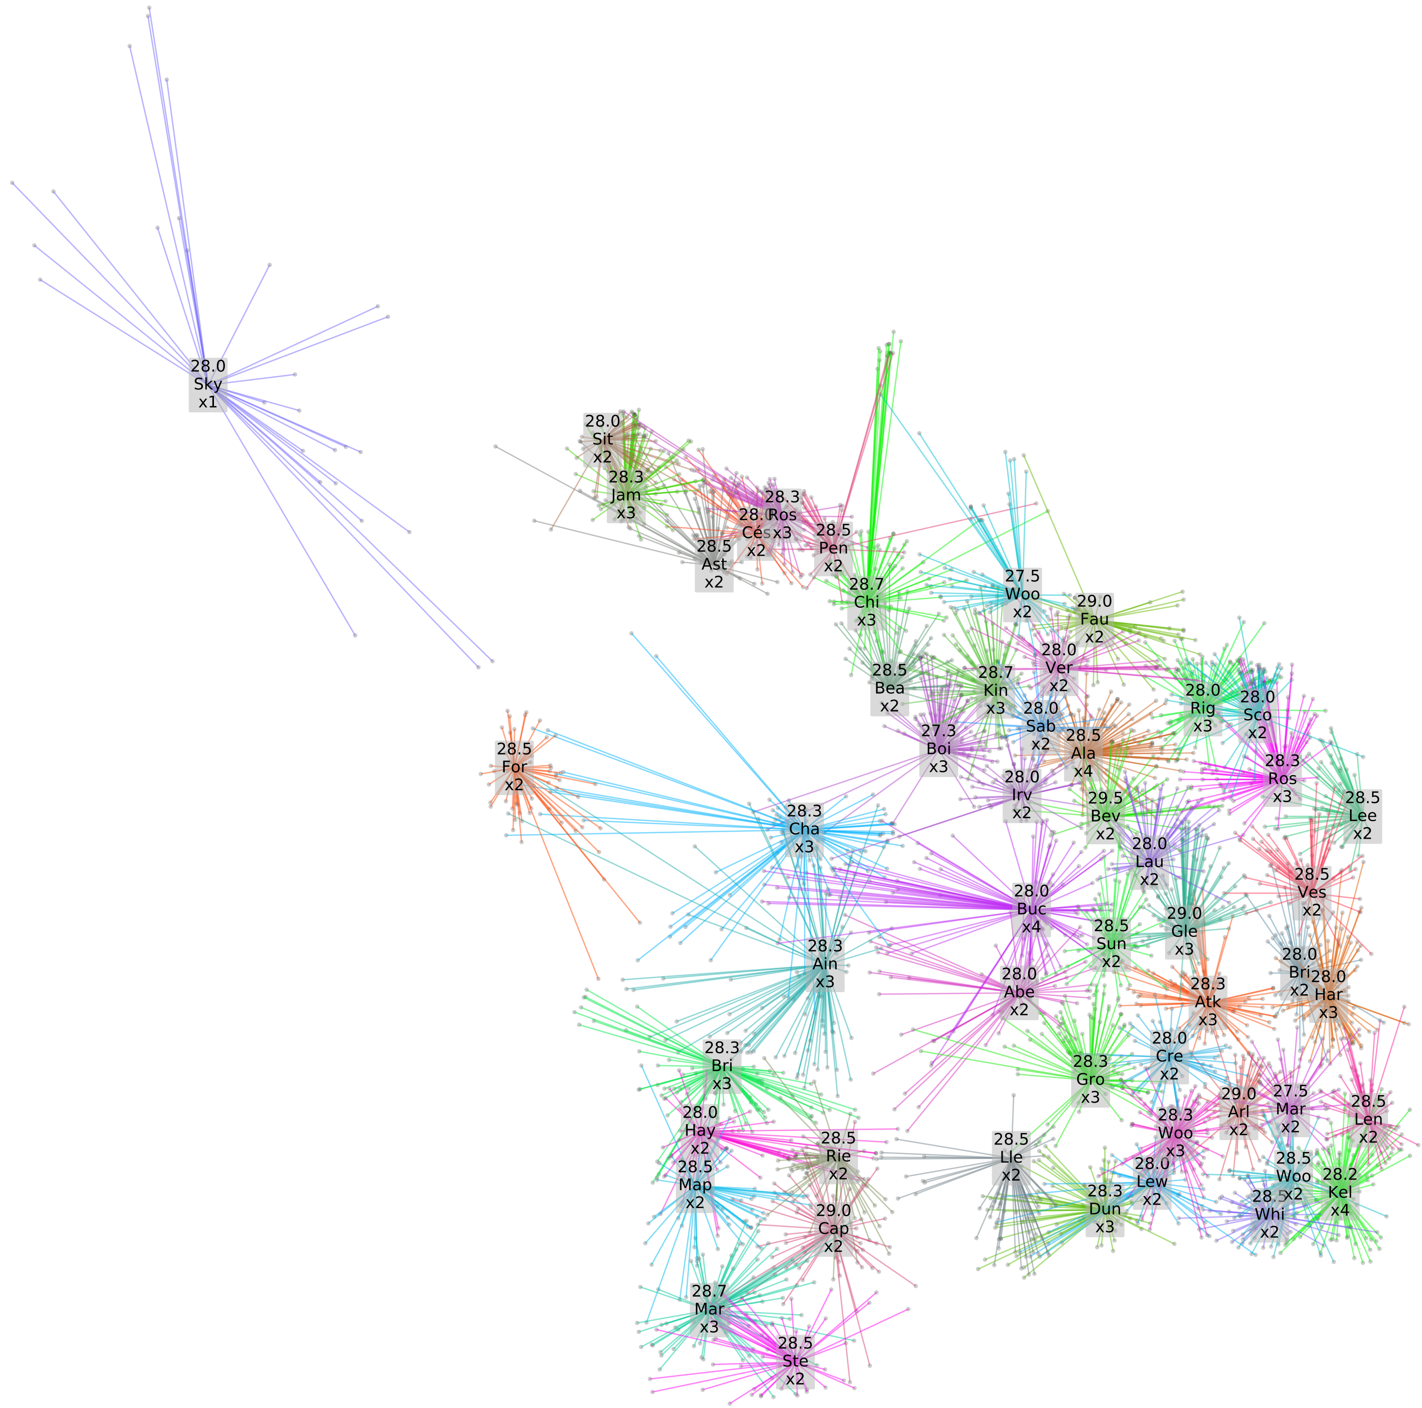

Figure 34: *SNM with 3-Schools/moderate-proximity Con-Opt, no transfers, with cross-river assignments.* Here is an example of *one possible* SNM assignment for 2014–15, with the 3-Schools/moderate-proximity function, with no restriction against cross-river assignments. Keep in mind that this example uses only travel distance, with no consideration for traffic patterns. Contrasting with Figure 32, since SNM can send kids across the river, it has managed to balance enrollments on *both* sides to the *same* target.

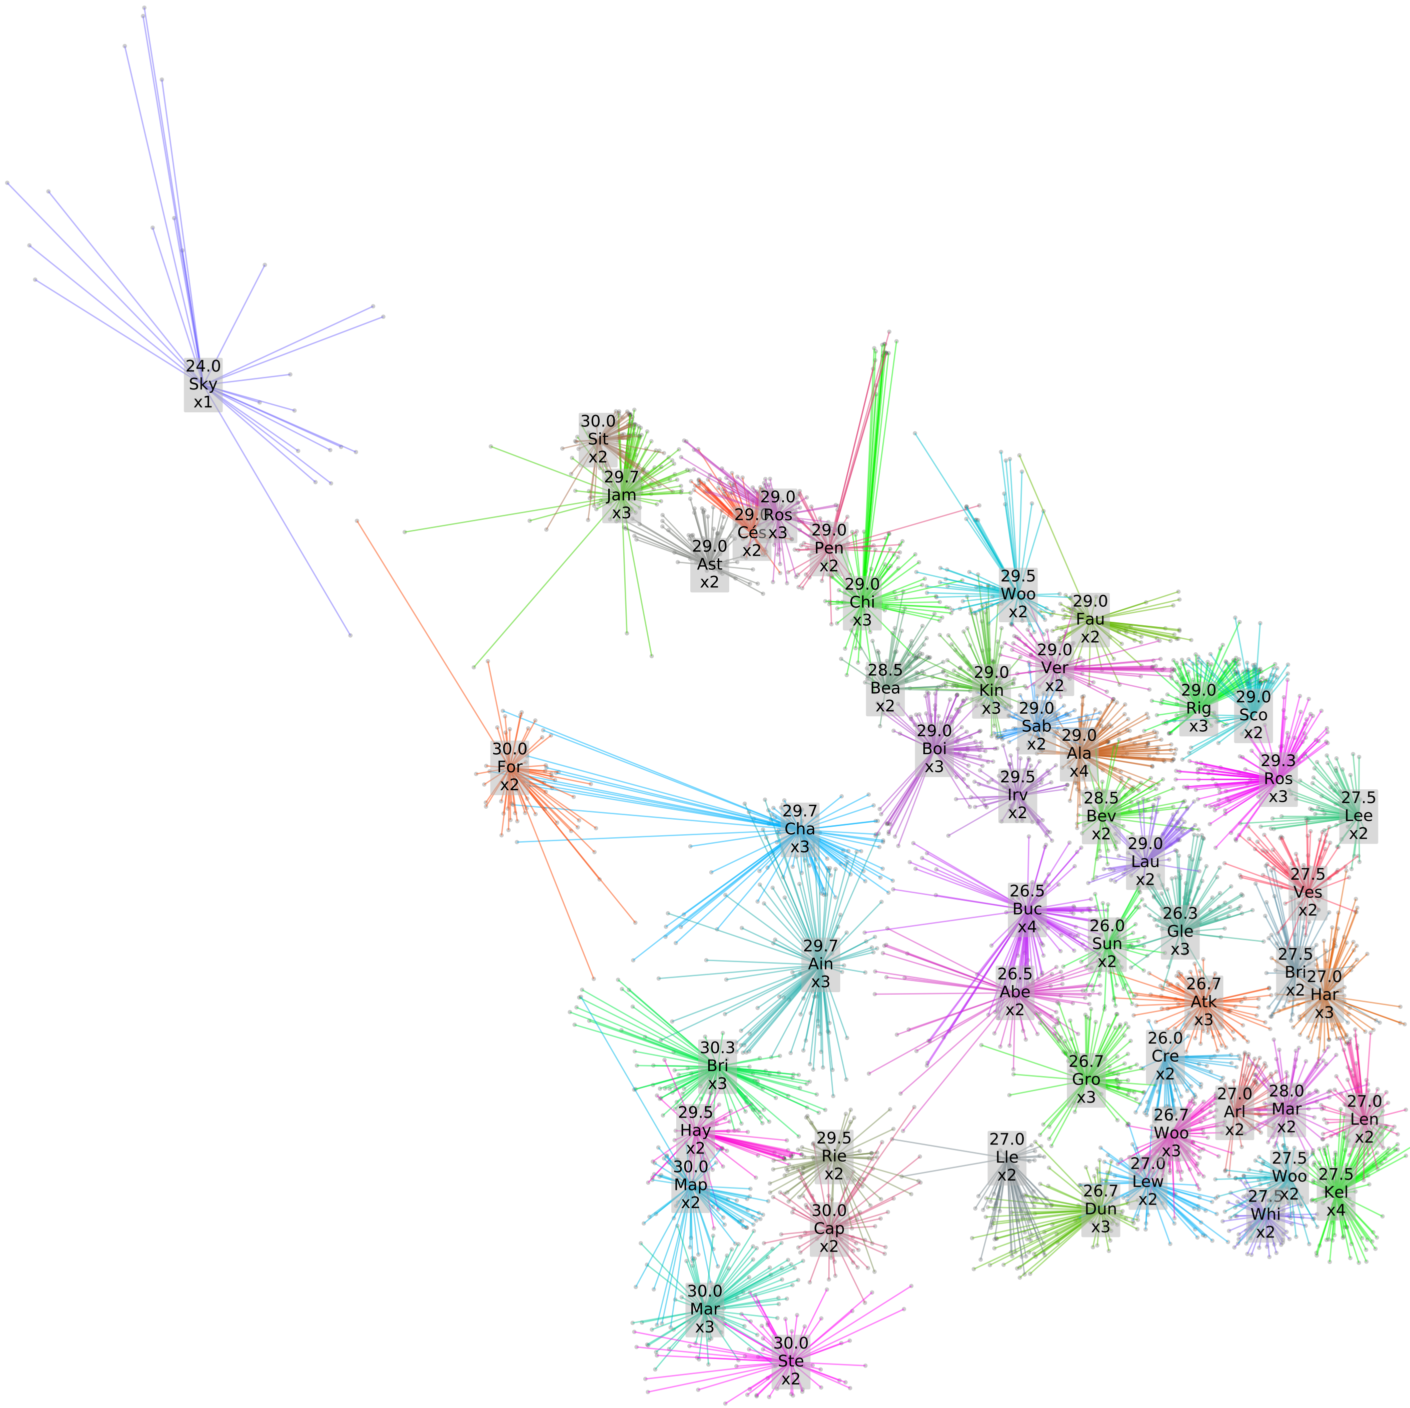

Figure 35: *SNM with 2-Schools/moderate-proximity Con-Opt, no transfers, with cross-river assignments.* This is an example of *one possible* SNM assignment for 2014–15, with the 2-Schools/moderate-proximity function, with no restriction on cross-river assignments. Again, keep in mind that this example uses only travel distance, with no consideration for traffic patterns. Here, the SNM manages to do a better job of balancing enrollments across the whole district when allowed to make cross-river assignments, but not as good as in Figure 34 with 3 schools.

8 Acute Enrollment Problems and Reassignment

The Soft Neighborhood Model is designed to avoid acute enrollment problems, and it does this by balancing the assignment of new students. However, a district using hard boundary assignments will inevitably face acute enrollment balancing problems, and may resort to reassignment of existing students to fix those problems. If a district has resigned itself to using reassignment, the Soft Neighborhood framework can perform the reassignment in a way that is efficient and fair. This same mechanism is compatible with grade reconfiguration within schools or the opening/closing of entire schools.

PPS is currently experiencing severe enrollment imbalances in certain grades in a number of schools, and the community is concerned about how quickly these enrollment problems can be fixed. The Soft Neighborhood Model can fix these problems as quickly as hard boundary reassignments can, with the following major advantages:

Targets specific grade levels The Soft Neighborhood approach targets specific over- and under-enrolled grades at specific schools. This constrains reassignments to the families with students in those grades. In contrast, hard boundary adjustments affect all families at all grades living in the reassigned regions, and continue to do so until the next boundary adjustment. The Soft Neighborhood approach will reassign just as many students as necessary to balance enrollments at all schools.

Produces a global solution from a simple specification To invoke the reassignment mechanism in the Soft Neighborhood process, the only new input is a list of the specific grades in specific schools experiencing critical enrollment problems. The mechanism then balances enrollment system-wide, considering all these problems at once.

Provides an impartial reassignment mechanism The Soft Neighborhood reassignment process provides an impartial mechanism for determining who will be impacted by reassignment: rather than placing the entire burden on families in specific geographic regions, the risk is distributed across the community.

In the remainder of this section, we describe how the Soft Neighborhood reassignment process works, and then we illustrate the process in simulation.

8.1 Reassignment with the Soft Neighborhood Model

Using the Soft Neighborhood Model to fix acute enrollment problems can be achieved with some simple modifications to the original assignment process. In the regular assignment process, the system only assigns students who have not been pre-assigned to a school (e.g., those continuing from an earlier grade). In the modified assignment process, the system is also allowed to reassign some students continuing from earlier grades in order to balance enrollments in critically over- and under-enrolled classes. Specifying which grade-levels in which schools are critically over- or under-enrolled yields a set of restrictions on reassignment. The system will reassign just as many students as necessary to balance enrollments. We'll review the four steps of the Soft Neighborhood process to see how this works.

Step 1: Set Capacity Constraints In the regular assignment process, this is where we establish the number of slots available in each grade level at each school, and where we determine which new entrants need a school assignment. In the modified process, we need only specify the *target total capacity* of each grade-level of a school. If a student has been pre-assigned to a specific school (e.g., she is continuing from the year before), then she is tagged with that school assignment. During this step, we also declare which grades at which schools are critically over- or under-enrolled. Doing so tells the system that it is allowed to perform reassignments in addition to simply assigning new entrants.

Step 2: Seed Probabilities by Proximity In the regular assignment process, we calculate proximity-based assignment probabilities for each entering student. The modified process does that, too, but for all students (new and continuing), and with two modifications.

One, if a student has been pre-assigned to a school in Step 1, then that school is given an extremely high weight in the student's set of seed probabilities. This tells the system to try to preserve that assignment.

Two, the system sets some of the probabilities to zero to enforce the following assignment restrictions:

- New students, with no pre-assigned school, have no restrictions.
- Students pre-assigned to critically over-enrolled schools have no restrictions.
- All other pre-assigned students are forbidden from being assigned to any school except the pre-assigned school or a nearby critically under-enrolled school.

What does this mean? New students are treated just like before; they could be assigned to any school deemed a candidate by the proximity function. Students in critically over-enrolled classes are treated much like new students; they could be assigned to any candidate school, however they do have a very strong bias to be kept where they are. All other pre-assigned students either stay where they are or might be reassigned to a critically under-enrolled candidate school, with a very strong bias to stay where they are.

Step 3: Balance Probabilities with Capacity Constraints This step is identical to the regular Soft Neighborhood assignment process. The seed probabilities for each student (incorporating the reassignment restrictions from Step 2) are modified to take into account available spaces in each grade at each school.

Step 4: Assign Students This step is also identical to the regular assignment process. New students get assigned to new schools. Most continuing students will stay assigned to their original school. However, some students in critically over-enrolled classes will be reassigned to nearby schools. And, some students will be reassigned to nearby schools with critically under-enrolled classes.[32] The end result is balanced enrollments at all schools, to the degree possible.

[32] It makes sense to give families with co-enrolled siblings the opportunity to move all of them to the new school if one of them is re-assigned.

8.2 Demonstration of Reassignment

We have implemented a demonstration of the Soft Neighborhood assignment process with reassignment, using the fourth grade at the fictitious Colorville Public School (CPS) system. There are four schools in the CPS system: PinkA, OrangeB, GreenC, and BlueD. CPS has maintained hard boundaries for as long as anyone can remember, but now it has decided to switch to the Soft Neighborhood Model. In Figure 36, we see that under the historic boundaries, OrangeB and BlueD's fourth grade classes have become severely under-enrolled (OrangeB) and over-enrolled (BlueD). Even without considering new fourth graders entering the district, OrangeB is 59 students under its target of 100, and BlueD is 22 students over its target of 50 for the entire grade.

After we declare OrangeB to be critically under-enrolled and BlueD to be critically over-enrolled, the Soft Neighborhood process both assigns new students and reassigns just as many continuing students as necessary to bring CPS schools into a balanced state. After all assignments have been made, enrollments across the district are much more balanced than they were before (see Figure 37). No school is dramatically under-enrolled, and no school is over-enrolled. In contrast, if only new students are assigned and no reassignment is performed, OrangeB remains severely under-enrolled, and BlueD remains severely over-enrolled (Figure 38). Table 17 provides a detailed overview of the outcome after running Soft Neighborhood assignment with reassignment on CPS' fourth grade classes.

We can imagine applying other changes to CPS before running the Soft Neighborhood process. For example, perhaps CPS decides to convert OrangeB to a middle school, with BlueD and GreenC feeding into it as K–5s (PinkA remains a K–8). The same Soft Neighborhood process can be run on the reconfigured CPS, and OrangeB's current K–5 students can be reassigned using the Soft Neighborhood reassignment mechanism. Similarly, we can imagine that CPS expects an influx of students in the coming decade, and decides to re-open the shuttered PurpleE school in the district's NW corner. PurpleE could be treated as a school with under-enrolled grades and populated with the Soft Neighborhood reassignment mechanism. To summarize, the Soft Neighborhood Model is compatible with grade reconfigurations and school openings or closures and, if a district resorts to it, student reassignment.

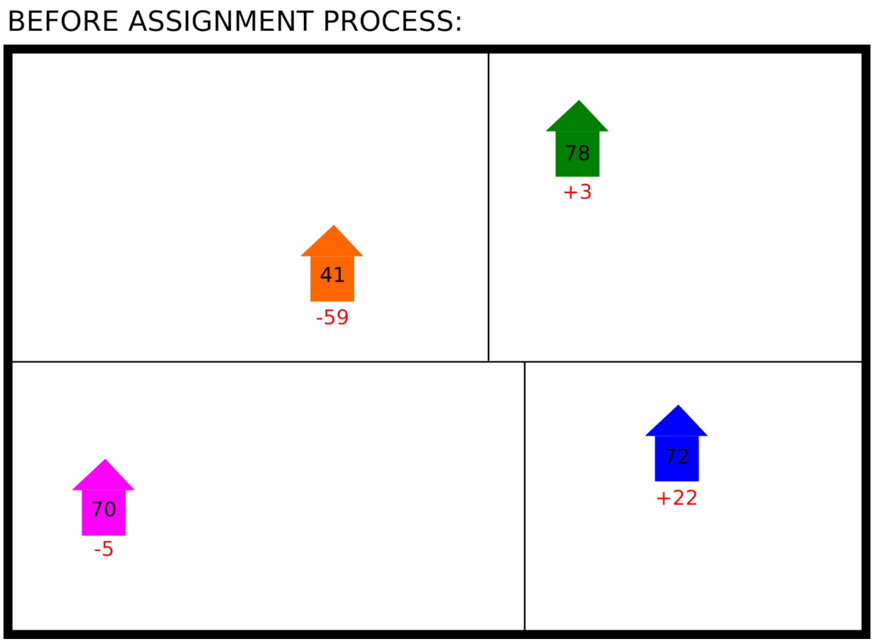

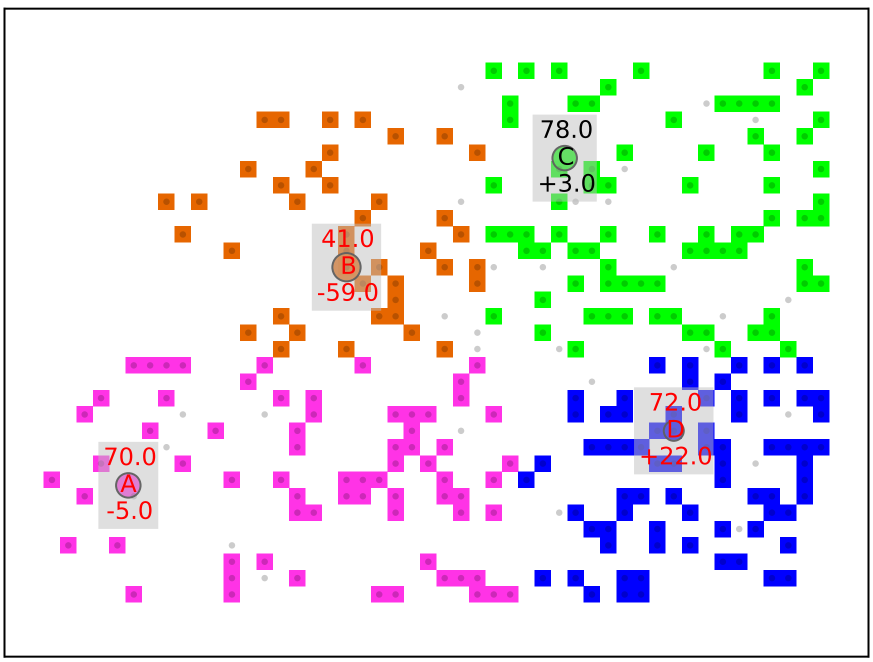

Figure 36: (top) Before running the reassignment process, two CPS fourth grade classes have severe enrollment problems under the existing boundaries: BlueD's fourth grade is over-enrolled by 22 students, and OrangeB's fourth grade is under-enrolled by 59 students. (bottom) Here we see where fourth grade students live and to which schools they are currently assigned. Each square represents a continuing student, and the color of the square corresponds to the school name to which she is assigned. There are also 29 new fourth graders, represented as gray dots, who haven't yet been assigned to any school.

AFTER ASSIGNMENT PROCESS:

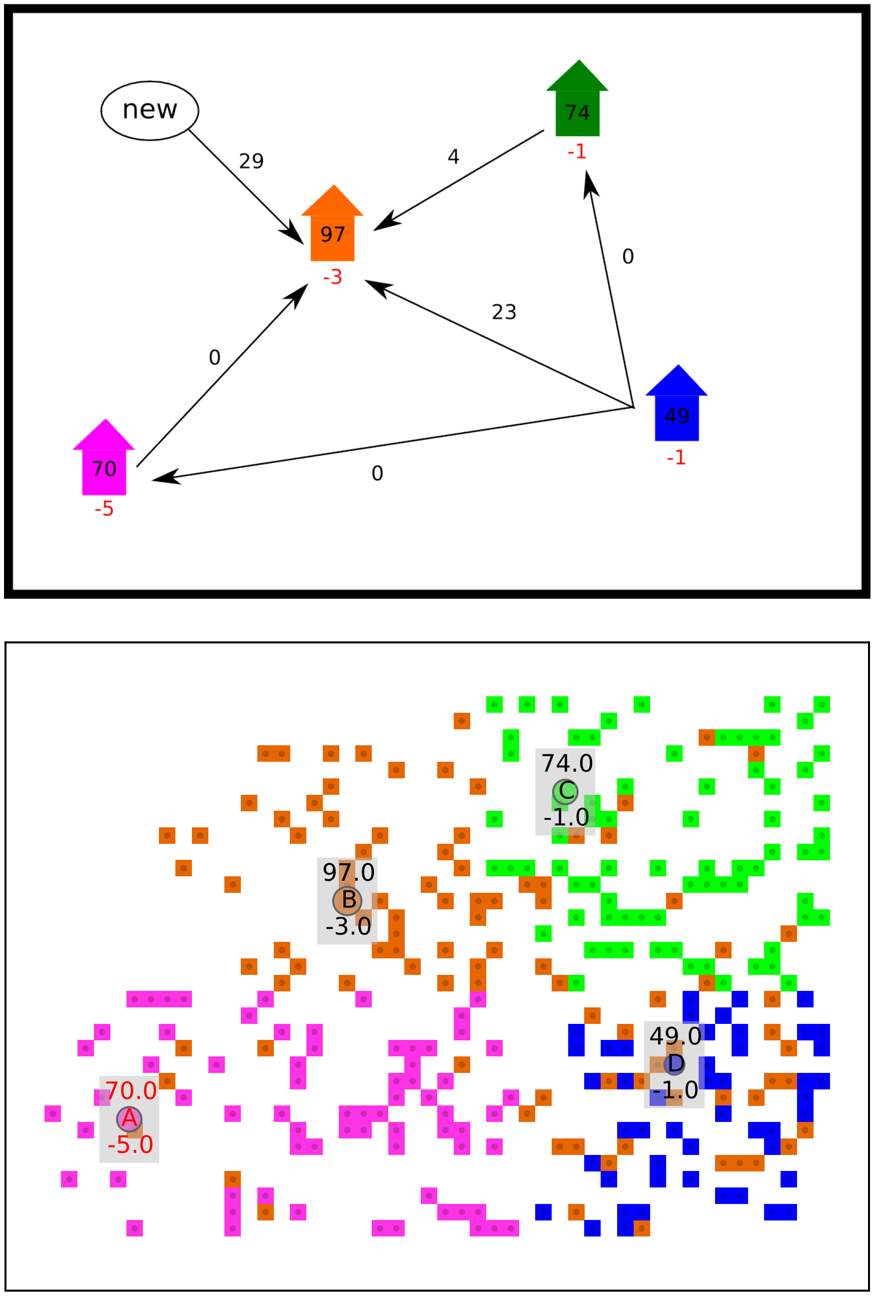

Figure 37: (top) The reassignment process moves some BlueD fourth graders to nearby OrangeB to relieve BlueD's over-enrollment; OrangeB gains 27 students total from BlueD and GreenC, and all 29 new students. PinkA is both well-enough balanced and distant enough from the other schools/students that it has no reassignments. (bottom) Here we see all fourth grade assignments after the process has completed. Enrollments at each school are well-balanced with respect to target enrollments. Since all the new students are allowed to attend OrangeB by the 3-School Rule (Section 2), assigning them all to OrangeB minimizes the reassignment of continuing students. All students are assigned to schools near their home address.

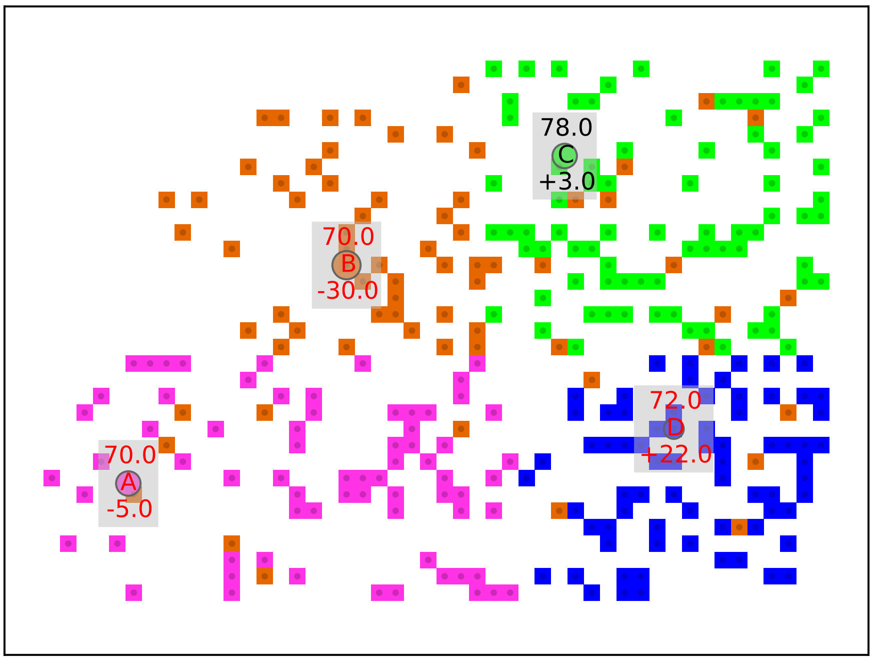

Figure 38: Assignments made if only new entrants are considered (i.e., without Soft Neighborhood reassignment). OrangeB is so under-enrolled that all new students are assigned there. However, even with 29 new students, OrangeB's fourth grade is still dramatically under-enrolled. And, nothing has been done to relieve BlueD's over-enrolled fourth grade class.

	original		now assigned to				
school	count	over/under	A	B	C	D	reassignments
A	70	-5	70	0	0	0	0
B	41	-59	0	41	0	0	0
C	78	+3	0	4	74	0	4
D	72	+22	0	23	0	49	23
(new students)	29	—	0	29	0	0	—
		result headcount	70	97	74	49	total: 27
		over/under	-5	-3	-1	-1	
			A	B	C	D	

Table 17: A detailed account of how fourth grade student assignments were made by the Soft Neighborhood process with reassignment. PinkA's enrollment stays the same. OrangeB gains 4 students from GreenC, 23 students from BlueD, and all 29 new students. A total of 27 reassignments have been made.

9 Frequently-Asked Questions

Are you suggesting that kids get reassigned to different schools every year in order to balance enrollments? No, we are definitely not saying that. School assignment in the Soft Neighborhood Model is not something that is meant to be done to every student every year for the purposes of keeping enrollments balanced. Assignments are only made when students need to be placed at a PPS school — for example, because they moved or because they are entering the system for the first time.

We believe strongly in stability and consistency. Involuntary school reassignment creates a lot of stress on families and school communities, and should be avoided whenever possible. That said, there are "transition points" in the typical K–12 grade sequence — at 6th and 9th grades when students enter middle school and high school, respectively. At these transition points, it might make sense to allow students to apply for reassignment by the Soft Neighborhood framework if they want.

What happens if a family moves to some other location within Portland? A family moving to a different location might trigger reassignment depending on the location of the new residence. If the original school is still one of the "nearby" schools according to the model, then the family should stay at the original school, and no reassignment occurs. If the original school is no longer a "nearby" school, then reassignment might make the most sense. The district could consider allowing students to stay at the original school even after a move — perhaps via petition, and without a transportation guarantee.

How does middle school and high school assignment work in this framework? The Soft Neighborhood Model provides a framework for assigning students to schools in a way that satisfies capacity constraints and otherwise is more likely to place students in schools closer to their homes. Keeping schools close to students is most important in the lower grades, when students are younger and less independent and require more day-to-day involvement from their families. There are several ways Soft Neighborhood assignment could be implemented across grade-levels within PPS. We like the idea of allowing students to re-apply for Soft Neighborhood assignment at critical transition points, e.g., at 6th and 9th grades. In the case that a student didn't want to apply for reassignment, he or she would continue on with the regular "feeder pattern" sequences for K–5 into 6–8, and from 8th grade into high school. Students in K–8 could also apply for reassignment at 6th grade if so desired.

This idea is compelling for a number of reasons. For one, not everyone wants to stay with their cohort all the way from K–12, and this would allow these students a chance to go to a different school. At the same time, students who would like to remain with the same cohort would be able to. Second, it provides the system an opportunity to restabilize and correct for enrollment "drift" — where enrollments shift away from a balanced state as students move, leave the system, etc. The Soft Neighborhood Model will have its best chance at balancing enrollments at the kindergarten level when most new students enter the system, and it can be used to assign new assignees at any grade level. But allowing students to be voluntarily reassigned at regular intervals gives the system a better chance to maintain a balanced state.

So you mean schools will get filled with students who live closest first? No... in the Soft Neighborhood Model, proximity makes it more likely that a student will get assigned to one school over another, but it doesn't make any guarantees. It figures out how to roll the dice to ensure that schools are filled to capacity with nearby students, and then it rolls the dice for everybody. Simply filling schools with the closest students first would solve the enrollment balancing problem, but with respect to equity and socioeconomic stratification, it would be just as bad (or worse) than classic hard boundaries.

Would this mean kids have to travel a lot farther to get to school? No, not necessarily. The Soft Neighborhood Model can be configured to have a distance profile that is very similar to the Hard Boundary model (at the expense of some population mixing). Using two schools instead of three, and a moderate to strong preference for the closest school, the Soft Neighborhood Model's distance profile is just about the same as the Hard Boundary model's (see Section 6, Tables 7 and 13, and Figures 17 and 22). But when we move from a 3-school model to a 2-school model, population mixing (measured by Local Assignment Diversity and Capture Linkage) takes a hit. Population mixing is how the Soft Neighborhood Model counters the concentration of wealth and poverty. It's a trade-off: to get more population mixing, we incur slightly longer commutes to school. But even a 3-School Soft Neighborhood Model comes nowhere close to busing kids all over town.

How would the Soft Neighborhood Model impact busing? A Soft Neighborhood Model configured with more mixing probably costs more in busing, but we imagine those costs would be offset by gains made elsewhere, e.g., savings through space efficiency and less-stressed educational environments. To get more mixing, kids do travel a little farther to school, but it's a matter of degrees, not a matter of everyone having to criss-cross around town. More kids technically become bus-eligible, but we (the authors) don't have any information about how eligibility translates into ridership, nor do we know how the Transportation Department does its route planning/optimization, nor we do know how they do their budgeting. On the other hand, the SNM *nails* the enrollment balancing problem — with consequent gains from space-efficiency and right-sizing each classroom. Any estimate of changes in transportation costs needs to be accompanied by an estimate of the value of both the improved educational outcomes and the real-dollars-saved by filling schools with the right numbers of students and teachers. (Not to mention the value of avoiding regular boundary adjustments.)

Why are you using Cartesian (straight-line) distances and not driving distances? We are no longer using Cartesian distance to approximate distances between home addresses and schools. As of Version 3.0 of this document, we use driving distances instead (details are in Section 5.4). For comparison, results from earlier versions using Cartesian distances can now be found in Appendix C.

Would this even work on the west side of the city? Yes. The Soft Neighborhood assignment algorithm is applied the same way to the west side of the city as the east. The experiments in Section 6 effectively separate the east and west sides for the purposes of comparison with a Hard Boundary model, and results show that the Soft Neighborhood Model balances enrollments as well on the west side as it does on the east side. And, despite the sparser schools and population, the median home-school driving distance (Figure 22) doesn't increase by so much: from 1.03mi with existing Hard Boundaries to 1.26mi with 3-School SNM.

Could kids be assigned to schools across the river in this framework? Yes, they could, just as with any framework. Hard Boundaries can assign kids across rivers, too; it's just a matter of drawing a boundary that crosses the water. As we demonstrated in Section 6, it is possible to force the Soft Neighborhood Model to *not* produce cross-river assignments by penalizing river-crossings. Unless this becomes a specific policy requirement for the district, we don't recommend this; an arbitrary boundary down the middle of a river does not contribute to balancing, proximity, or equity. It would be better to create a realistic home-to-school travel metric which incorporates estimates of traffic congestion, bridge closure risks, etc. Provided with that, SNM would only produce cross-river assignments if they were comparable travel-wise to same-side assignments.

Hmmm....this sounds interesting, but I want my kids to go to school with the kids next door. Well, you're in luck — they just might! They might also get to go to school with the kids across the street, or in the house behind you, or one block over. That's a feature of the Soft Neighborhood Model: if kids live close together, they might go to school together. There is no guarantee, though, but it's likely that some kid on your block will go to the same school. Under the current hard boundary system, if you live near an existing school boundary you don't have any guarantees either, since the district can move that boundary in unpredictable ways whenever it sees the need. Either way, if your kids like to play with the kids next door, that probably won't change just because of the schools they go to.

Isn't this kind of like the old transfer lottery? Nope, the Soft Neighborhood Model is a neighborhood model, where assignment is based only on the home address, and it must provide an assignment to everyone who applies to it. The old transfer lottery was built on top of the PPS Hard Boundary neighborhood model and was under no obligation to assign all applicants.

I've heard of this before in ⟨Cambridge, San Francisco, etc.⟩ — isn't it called Controlled Choice or something like that? The Soft Neighborhood Model is not a School Choice model (Controlled Choice is a variant of School Choice). School Choice systems involve eliciting ranked preferences over schools from families, and then making assignments based on those preferences. Choice-based programs may confer strategic advantages to better-educated and wealthier families and have been blamed for intensifying segregation in some cases [15]. We wanted to develop a solution to enrollment balancing in PPS that would counteract the tendency towards segregation. The Soft Neighborhood Model is tailored to achieve enrollment balancing (capacity) and retain a sense of neighborhood (proximity), while inducing a gentle mixing of student populations.

Shouldn't PPS ensure a baseline of equitable academic program offerings at every school first, before considering something like the Soft Neighborhood Model? PPS absolutely should ensure a baseline of equitable academic program offerings at every school, but we think the community shouldn't have to wait for this to happen before implementing a robust enrollment balancing solution. In fact, having predictable and optimal school enrollment will help the district distribute educational resources to the schools in an equitable way. We are skeptical that a solution involving boundary change, even frequent boundary change, is sufficiently robust to address enrollment balancing. Moreover, we see boundaries as a historical artifact that have coevolved with and reinforced the racial and socio-economic inequities the community is trying to counteract. In short, our school district needs a functional solution to enrollment balancing that doesn't directly conflict with our equity goals.

10 Conclusions: Where Do We Go From Here?

PPS is currently investing a substantial amount of time and energy trying to fix its enrollment balancing problem. It has convened an advisory committee (District-Wide Boundary Review Advisory Committee, or DBRAC) to assess the problem, and DBRAC has recommended a framework for regular boundary adjustments. As a means for solving enrollment balancing, we believe that boundary adjustment is inadequate. Any hard-boundary system will result in imbalanced enrollments and segregated populations. It will also confer an advantage to families with more resources, who can buy assignment to the public school of their choice.

We have developed a boundary-free solution to the enrollment balancing problem that focuses directly on satisfying school-specific target enrollments. Moreover, our model reinforces the district's stated equity goals by eliminating any determinism between home address and school assignment, and by encouraging mixing between populations. We designed the Soft Neighborhood Model to be a robust and equitable solution to enrollment balancing that retains the community ideal of neighborhood. Like any traditional neighborhood school model, the Soft Neighborhood Model aims to place students in schools close to their homes, and the only input into the system from a family is its home address.

We have implemented and tested the Soft Neighborhood Model against seven years of historical PPS enrollment data. We have defined three metrics for evaluating our model on the basis of balanced enrollments, travel distance, and assignment diversity. Our experimental results suggest that the Soft Neighborhood Model outperforms the historical Hard Boundary system with respect to enrollment balancing; that it improves assignment diversity; and that students travel comparable distances whether using Soft Neighborhood assignments or historical assignments. These empirical results demonstrate the potential of the Soft Neighborhood Model as a robust and equitable solution to the enrollment balancing problem in PPS. Moving forward, the model needs to be fully tuned and evaluated against a proper data set.

We encourage the PPS Board of Directors and the larger community to seriously consider the Soft Neighborhood Model for implementation in PPS. Moreover, we would like to see compelling empirical evidence that the DBRAC proposal can actually achieve its stated balancing and equity goals, and we request an objective evaluation of the district's proposed solution in terms of the metrics we have defined in this document. Any such evaluation must be implemented with transparency to the public — which means it should use a data set that has been properly anonymized and released to the public, and that it should make the entirety of its evaluation methodology (e.g., the source code, too) available to the public for independent verification and validation.

A Appendix: Soft Neighborhood Algorithms

This appendix delves into the algorithmic details of the Soft Neighborhood assignment model. As with any mathematical algorithms, one could carry out a school assignment using just pencil and paper, or a pen-knife and bark chips, but it would certainly be faster if executed by modern computing hardware.

The basic inputs to the system are:

1. A set of schools $S = \{s_i\}$ with known geographic locations (address).
2. A set of students $C = \{c_j\}$ with known geographic locations (home address).

A.1 Step 1: Set Capacity Constraints

The first step is to establish the enrollment targets t_i for each school s_i, also referred to as the capacity constraints. For the given grade level, the district must decide how many spots are available for new assignees. The general formula looks like this:

$$t_i = ((\textit{\# of sections in } s_i) \times (\textit{ideal section size})) - (\textit{\# of pre-assignees to } s_i)$$

The *# of sections* and the *ideal section size* are things for the district to decide, based on whatever factors already go into filling a school with students. The pre-assignees are all the students already assigned to the school and not subject to the random assignment process, including:

- students continuing on from an earlier grade;
- co-enrolling younger siblings of older students already assigned.

The pre-assignees are removed from the corpus of students C, as they have already been assigned to a school.

The capacity constraints can be expressed as a vector $\{t_i\}$ of targets, one per school s_i.

A.2 Step 2: Seed Probabilities by Proximity

In Step 2, we establish a probability matrix $P = (p_{ij})$ in which each element p_{ij} represents the probability with which student c_j should be assigned to school s_i. This matrix is initialized (seeded) using weights from a proximity function that favors schools that are closer to a student over schools that are farther away.

The precise form of this seed weighting function is important to the performance of the model overall, as it sets the limits on how far students can live from candidate schools.

The function used in our simulations is:

$$W_{\rho,n_c}(s,c) = \begin{cases} (1-\rho)\left(1 - \dfrac{d(s,c)}{D(c)}\right) + \rho & \text{if } d(s,c) <= D(c), \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

where

$$D(c) = \text{distance to } n_c\text{-th closest school to } c \tag{2}$$

Intuitively, this means: "For a student c_j, the candidate schools are the n_c closest schools, and the weight for the farthest school is ρ."

In early versions of this document, we used a simple Cartesian metric ("crow's flight") to measure home-school distance $d(s, c)$. As of Version 3.0, we are using driving distances routed over a map of Portland. As of yet, we do not incorporate elements such as traffic congestion patterns or driving-vs-walking trade-offs. Ideally, the distance measure would be compatible with methodologies already in use by the school district, e.g., to determine busing eligibility and schedules.

The seed weights are plugged into the probability matrix — though, they must be normalized for each student (over all schools) before they become actual probabilities:

$$p_{ij} = \frac{W(s_i, c_j)}{\sum_i W(s_i, c_j)}$$

At this point, the matrix P contains the probabilities p_{ij} that a student c_j should be assigned to school s_i based solely on her proximity to that school. For any given student, most p_{ij} will be zero, and only the schools with non-zero probability are candidates for that student.

A.3 Step 3: Balance Probabilities with Capacity Constraints

The Balancing Algorithm is a kind of iterative *relaxation algorithm* that repeatedly alternates two steps until it converges to a solution.

In a nutshell, the algorithm works like this:

1. For each school, sum the probabilities applying to that school to find the expected number of children assigned to that school, which may be more or less than the target capacity. Go back and multiply each of those probabilities by a correction factor *target-capacity*/*expected-assignments*.

2. For each child, renormalize his/her assignment probabilities: divide each probability by their sum, so that they add up to 1.

3. Repeat 1 and 2 until the expected assignments to each school converge to their target capacities, quitting if the convergence completes or stalls.

If there *is* a solution that exactly satisfies the target capacities, this algorithm will find it — but there must be a solution. A necessary condition for this is that the capacity of the district (total of target capacities of all schools) is equal to the number of children being assigned. Even if there is no exact solution, the algorithm will find something close. If the algorithm stalls leaving one or more schools in an large over- or under-capacity state, that indicates a grave capacity problem in some part of the district, e.g., there simply aren't enough nearby schools to handle the local student population.

A.4 Step 4: Assign Students

The complete assignment step will be described in a future revision of this document, but the gist of it is:

1. Partition the probability matrix and the sets of students into groups which have non-zero probabilities for the same subset of schools.

2. Nudge the probability submatrix for each of these groups to ensure an integer expectation for the number of students assigned to each of the candidate schools in the group.

3. (Optionally) propagate probability mass along the network formed by the groups to further balance the integer expectations with the capacity constraints.

4. For each group, divide its students among its schools according to the integer expectations in a way that honors the probability submatrix, using an approximation algorithm similar to the balancing algorithm.

B Appendix: Detailed Description of the Data Sets

Portland Public Schools provided two data sets for the purposes of running the simulations described in this document: the STUDENTS data set and the SCHOOLS data set. The STUDENTS data set contains historical and current information from the last seven school years. The SCHOOLS data set contains information about K–8, K–5, and 6–8 schools in PPS in the year 2014-15.

B.1 The Students Data Set

The STUDENTS data set consists of just over 200K records for all students in kindergarten through eighth grade over the seven-year period from 2008–2015. Table 18 shows the number of student records (K–8 and K only) each year.

Year	**K–8**	**K**
2008-09	28,760	3,601
2009-10	29,077	3,726
2010-11	29,211	3,657
2011-12	29,875	3,715
2012-13	30,056	3,884
2013-14	30,671	3,843
2014-15	30,744	3,681
2008-15	208,395	26,107

Table 18: The number of student records (K–8 and K only) per year in the STUDENTS database.

Each record in the data set represents a student and contains the following information:

- **School year**: The school year for which this record is valid (2008-09 through 2014-15).
- **Grade**: Student's grade level (K–8) during that school year.
- **Campus enrolled**: Student's assigned school during that school year. Schools included in the data set are K–8, K–5, and 6–8 schools that have a neighborhood program on their campus. District-only focus options, alternative, and charter schools are not included.
- **Campus code**: A numeric ID associated with the assigned school.
- **Neighborhood**: The neighborhood in which the student lives. Values for this attribute can be a PPS neighborhood for in-district students (205,113 records)), *out of district* (3,264 total records), or *not yet assigned* (18 total records).
- **Capture?**: TRUE if the enrolled campus is the same as the neighborhood in which the student lives, and FALSE otherwise (e.g., transfer students).

- **Home Coordinates**: The X, Y coordinates of the student's home address, randomly offset by the district for the purposes of anonymity. Addresses are represented in Oregon State Plane coordinates, North Zone, in feet.

B.2 The Schools Data Set

The SCHOOLS data set contains information about K–8, K–5, and 6–8 schools in PPS during the year 2014–15. Each record in the data set represents a school (or campus for schools with more than one building) and includes the following information:

- **Name**: The name of the school/campus
- **Configuration**: K–8, K–5, or 6–8
- **Address**: street, city, state, zip of the building
- **School Coordinates**: The X, Y coordinates of the schools address represented in Oregon State Plane coordinates, North Zone, in feet.
- **Kindergarten Homeroom Count**: Number of kindergarten homeroom classrooms for 2014-15.
- **Demographics**: total number enrolled, number free-or-reduced lunch status (by direct certification) enrolled, and a breakdown by racial designation, at each school

B.3 Flaws in the Data Sets

The data sets provided by the district are flawed in ways that limit the extent to which we have been able to validate the model via the simulations. It is important to understand these flaws. Fixing them is a prerequisite for full validation of the Soft Neighborhood framework.

1. **No sibling information**: Sibling relationships are not available in the data set, making it impossible to evaluate the effect of guaranteed placement of co-enrolled siblings at the same school. Basically, in this data set, there are no siblings.
2. **Lack of information about new vs continuing students**: The address information in the STUDENTS database has been perturbed so as to preserve student anonymity. This is a good thing. However, there is no address continuity within the data set. In other words, it is impossible to tell whether or not a 1st-grader has continued on from kindergarten or is newly-enrolled at a school. This makes it impossible to assess how well the Soft Neighborhood Model is able to maintain balanced enrollments as students move up from kindergarten through the primary and middle grades.
3. **Lack of sufficient information about focus options**: The data set lacks sufficient detail about focus option programs co-located at neighborhood schools, and focus option schools. There is no information at all about the latter (focus option schools). At schools with a neighborhood program and co-located focus option program(s), there is no way to distinguish which students are in which program. Basically, in this data

set, focus option schools and the students who attend them don't exist, and everyone at a neighborhood school is enrolled in a neighborhood program. This makes it impossible to estimate the effect of program placement.

There are also some limiting factors in the data sets that cannot be fixed but that should be acknowledged:

- **No student-specific socio-economic and racial/ethnic information**: This information cannot be released, and therefore any evaluation of the effect of the Soft Neighborhood Model on mixing and diversity is limited.

- **Lack of real information regarding target capacities**: In the Soft Neighborhood framework, it is important for the district to set school-specific target capacities based on the preferred configuration of the actual space available at each building. We don't have these target capacities. The SCHOOLS data set does report the number of kindergarten classrooms at each school in 2014-15. However, this information does not generalize well to earlier years, and it does not capture whether those numbers are an artifact of overcrowding or underenrollment (e.g., district was forced to convert some space into an extra classroom). See Section 5.4 for more details regarding how target enrollments were set in the simulations used to generate the results in this document.

B.4 How to Fix the Data Sets

The following corrections to the STUDENTS data set would allow for complete validation of the model without compromising student anonymity:

1. **Data for all programs, and complete program identification**: The data set should include data for the student population of all programs, e.g., focus option schools/programs as well as neighborhood schools. Furthermore, student data should identify the specific program in which a student is enrolled (e.g., *Mandarin Immersion* or *Odyssey*), not just the campus. This would allow one to tease apart the populations attending such programs, and also allow for analyzing how focus-option populations interact with the distribution of neighborhood school populations.

2. **Consistently anonymized addresses**: Ideally, when a given address (in this case, a coordinate pair of (feet-east, feet-north)) is randomly perturbed to anonymize it, the same perturbed result should be used for every instance of it. For instance, if the address (84757.1E, 33364.2N) is tweaked to (84790.2E, 33298.4N) the first time it is encountered, then it should be tweaked to that same value everywhere in the data set. This would allow one to make educated guesses as to which students are newly-enrolled at a school, and which students are siblings.

3. **Consistent student identifier**: Additionally, the data would be even more clear if each student were tagged with a unique-ID that was consistent from one year to the next. The actual PPS student-ID should certainly remain private, but a one-way hash of that ID could keep it private and serve the same purpose in the data set. This would make it very clear which students were continuing on from one year to the next.

4. **Explicit sibling references**: Consistent student ID's would also allow for an explicit reference to the next older sibling and/or same-year sibling (if any) — which would clarify which students should be handled as co-enrolled siblings and how to handle them.

C Appendix: Results Using Cartesian Distances

As of Version 3.0 of this document, we have been using driving distances to approximate distance between home addresses and schools. In earlier versions, we reported results using "as-the-crow-flies" Cartesian distances. While Cartesian distance is not the best approximation for real travel distances, we initially used it because it is simple and quick to implement. We report those earlier results here for the sake of posterity. Our newest results using driving distances can be found in Sections 5, 6, and 7.

Also, in earlier versions we ran the Soft Neighborhood Model in a single configuration using the following "3-School-Rule" proximity function:

$$W(s,c) = \begin{cases} D(c) - d(s,c) & \text{if } d(s,c) <= D(c) \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

where

$$D(c) = \text{maximum of } \begin{cases} 1.1 \times \text{distance to third closest school to } c \\ 1.25 \text{ miles} \end{cases} \tag{4}$$

Intuitively, this means: "For a student c_j, the candidate schools are: the three closest schools, any school within 110% of the third closest school, and any schools within 1.25mi distance."

C.1 Enrollment Balancing with Cartesian Distances

Table 19 shows results for the Enrollment Balancing metric. Looking at all seven years (third section in the table), we see that the Soft Neighborhood Model is much better at hitting enrollment targets than either historical model. It hits the 25-student enrollment target about 42% of the time, and is $+/-2$ (23–27) students 96% of the time. In contrast, even with transfers, historical assignments are either extremely under (<23 students) or over ($>$ 27 students) target around 42% of the time. Without transfers, under- and over-enrollment is much worse (61%). Similar observations hold in the two independent years shown (first two sections). Note that 2014–15 appears to have been a more difficult year for enrollment balancing: historically, even with transfers, only around 7% of sections were at the 25-student target (c.f. 40% for the Soft Neighborhood Model). However, even in a year when balancing is difficult, the Soft Neighborhood Model avoids extreme under- and over-enrollment, with 98% of sections within the range of 23–27 students.

Year(s)	Model	Kindergarten Section Sizes				
		<23	23–24	25	26–27	>27
2008-09	Hard Boundary, historical	21.7%	28.7%	14.0%	16.8%	18.9%
	Hard Boundary, no transfers	37.8%	12.6%	4.9%	7.7%	37.1%
	Soft Neighborhood	1.4%	24.5%	55.9%	17.5%	0.7%
2014-15	Hard Boundary, historical	28.4%	26.4%	6.8%	17.6%	20.9%
	Hard Boundary, no transfers	35.1%	19.6%	4.7%	12.2%	28.4%
	Soft Neighborhood	2.0%	42.6%	39.9%	15.5%	0.0%
2008-15	Hard Boundary, historical	22.4%	27.9%	11.0%	19.0%	19.7%
	Hard Boundary, no transfers	31.5%	16.7%	6.3%	15.6%	30.0%
	Soft Neighborhood	2.3%	32.8%	41.9%	21.3%	1.7%

Table 19: Under- and over-enrollment of kindergarten sections for two years (2008-09 and 2014-15), and over all 7 years, for the current system and the Soft Neighborhood Model. The Soft Neighborhood Model is much better at hitting enrollment targets than the current hard-boundary framework. If we look at the data over all seven years (third section of the table), the Soft Neighborhood Model comes within $+/-2$ of the 25-student target (23–27 students) 96% of the time. The historical models fall within this range only 58% of the time (with transfers) and 39% (without transfers).

C.2 Travel Distance with Cartesian Distances

Table 20 shows results for the Travel Distance metric. Here we include not just statistics for the district as a whole (top), but also for the east side of the river (middle) and west side of the river (bottom) independently. Overall, the Soft Neighborhood Model assigns students to a school within a reasonable distance from home. If we look at the district as a whole, over all seven years, we see that under the Soft Neighborhood Model, around 33% of students travel less than 0.5 mi, about 81% of students travel less than 1.0 mi, and about 96% travel less than 1.5 mi. Recall that in our simulations, the Soft Neighborhood Model does not maintain historical transfers, so it's most instructive to compare it against the historical no-transfer model. Here, we see that around 55% of students travel less than 0.5 mi, around 89% travel less than 1 mi, and around 97% travel less than 1.5 mi. More students travel very short distances under the current system, but both assignment models keep most students within reasonable distances of their homes.

Year(s)	Model, whole district	Distance from Home to School			
		<0.5mi	<1.0mi	<1.5mi	>1.5mi
2008-09	Hard Boundary, historical	47.7%	78.1%	87.6%	12.4%
	Hard Boundary, no transfers	55.9%	89.3%	97.2%	2.8%
	Soft Neighborhood	32.7%	80.6%	95.8%	4.2%
2014-15	Hard Boundary, historical	46.3%	78.5%	88.5%	11.5%
	Hard Boundary, no transfers	53.3%	87.8%	96.7%	3.3%
	Soft Neighborhood	31.2%	79.0%	95.7%	4.3%
2008-15	Hard Boundary, historical	47.1%	78.2%	88.4%	11.6%
	Hard Boundary, no transfers	55.4%	88.8%	97.3%	2.7%
	Soft Neighborhood	32.5%	81.0%	96.2%	3.8%
Year(s)	**Model, east of river**	**Distance from Home to School**			
		<0.5mi	<1.0mi	<1.5mi	>1.5mi
2008-09	Hard Boundary, historical	50.8%	80.7%	88.8%	11.2%
	Hard Boundary, no transfers	60.3%	93.0%	99.2%	0.8%
	Soft Neighborhood	33.8%	84.3%	98.8%	1.2%
2014-15	Hard Boundary, historical	49.2%	81.4%	89.8%	10.2%
	Hard Boundary, no transfers	57.6%	91.8%	98.6%	1.4%
	Soft Neighborhood	32.3%	82.3%	98.0%	2.0%
2008-15	Hard Boundary, historical	49.9%	80.4%	89.3%	10.7%
	Hard Boundary, no transfers	59.4%	92.2%	98.8%	1.2%
	Soft Neighborhood	33.6%	84.5%	98.7%	1.3%
Year(s)	**Model, west of river**	**Distance from Home to School**			
		<0.5mi	<1.0mi	<1.5mi	>1.5mi
2008-09	Hard Boundary, historical	35.4%	68.0%	83.2%	16.8%
	Hard Boundary, no transfers	38.8%	74.7%	89.5%	10.5%
	Soft Neighborhood	28.5%	66.5%	84.0%	15.9%
2014-15	Hard Boundary, historical	35.7%	68.0%	83.7%	16.3%
	Hard Boundary, no transfers	37.7%	73.5%	90.0%	10.0%
	Soft Neighborhood	27.2%	67.0%	87.5%	12.4%
2008-15	Hard Boundary, historical	35.6%	68.9%	84.6%	15.4%
	Hard Boundary, no transfers	38.6%	74.8%	90.7%	9.3%
	Soft Neighborhood	27.8%	66.2%	85.7%	14.3%

Table 20: Travel distance from home to school for two years (2008-09 and 2014-15), and over all 7 years, for the current system and the Soft Neighborhood Model. Overall, more students travel very short distances under the current system, but both assignment models keep most students within a reasonable distance of their home.

C.3 Assignment Diversity with Cartesian Distance

Table 21 shows results for the Assignment Diversity metric. Again, we see the whole district as well as the breakdown for the east and west sides. The Soft Neighborhood Model displays greater assignment diversity in all cases: the fraction of students who have only one school represented among all students within 1000 ft is only 10.8% for the district as a whole. The fraction with two or more schools well-represented is 64.2%, and the fraction with three or more schools is 24.3%. The distributions for the historical models without transfers are highly skewed toward only 1 school. This makes sense because in most regions, students are assigned to the same school. Only at boundaries will there be regions with even the semblance of mixing. For the historical model with transfers, the distributions show greater diversity, though not as much as the Soft Neighborhood Model. Note that we achieve this diversity in the Soft Neighborhood Model without any transfers.

Year(s)	Model, whole district	Assignment Diversity		
		≤1.0	≥2.0	≥3.0
2008-09	Hard Boundary, historical	22.7%	43.5%	15.3%
	Hard Boundary, no transfers	57.6%	7.6%	0.3%
	Soft Neighborhood	10.8%	64.2%	24.3%
2014-15	Hard Boundary, historical	26.3%	38.7%	12.3%
	Hard Boundary, no transfers	59.0%	7.8%	0.1%
	Soft Neighborhood	13.6%	61.4%	23.1%
2008-15	Hard Boundary, historical	23.7%	42.1%	16.1%
	Hard Boundary, no transfers	58.3%	7.6%	0.4%
	Soft Neighborhood	9.8%	64.4%	25.1%
Year(s)	**Model, east of river**	**Assignment Diversity**		
		≤1.0	≥2.0	≥3.0
2008-09	Hard Boundary, historical	14.4%	49.7%	18.5%
	Hard Boundary, no transfers	52.1%	8.5%	0.4%
	Soft Neighborhood	4.7%	72.1%	29.8%
2014-15	Hard Boundary, historical	18.9%	44.4%	15.1%
	Hard Boundary, no transfers	53.7%	9.0%	0.1%
	Soft Neighborhood	6.6%	70.7%	28.0%
2008-15	Hard Boundary, historical	16.4%	48.1%	19.4%
	Hard Boundary, no transfers	53.0%	8.4%	0.5%
	Soft Neighborhood	4.6%	72.3%	30.3%
Year(s)	**Model, west of river**	**Assignment Diversity**		
		≤1.0	≥2.0	≥3.0
2008-09	Hard Boundary, historical	54.7%	19.5%	2.6%
	Hard Boundary, no transfers	79.2%	4.0%	0.0%
	Soft Neighborhood	34.3%	33.3%	2.9%
2014-15	Hard Boundary, historical	52.7%	18.5%	2.0%
	Hard Boundary, no transfers	78.0%	3.3%	0.0%
	Soft Neighborhood	38.6%	27.9%	5.4%
2008-15	Hard Boundary, historical	54.7%	17.2%	2.5%
	Hard Boundary, no transfers	80.7%	4.1%	0.0%
	Soft Neighborhood	31.7%	31.0%	3.2%

Table 21: Assignment diversity for two years (2008-09 and 2014-15), and over all 7 years, for the current system and the Soft Neighborhood Model. Overall, the Soft Neighborhood Model shows greater assignment diversity in all scenarios.

References

[1] Atila Abdulkakiroğlu and Tayfun Sönmez. School choice: A mechanism design approach. *The American Economic Reivew*, 93(3):729–747, 2003. 27

[2] Judy Brennan. Preliminary impact of 2015 enrollment and transfer policy revisions. PPS Board meeting materials, pp. 69–74, June 23, 2015. 28

[3] Lars Ehlers, Isa E. Hafalir, M. Bumin Yenmez, and Muhammed A. Yildirim. School choice with controlled choice constraints: Hard bounds versus soft bounds. *Journal of Economic Theory*, 153:648–683, 2014. 21

[4] David Gale and Lloyd Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1), 1962. 27

[5] Ethan Johnson and Felicia Williams. Desegregation and multiculturalism in the Portland public schools. *Oregon Historical Quarterly*, 3(1), 2010. 21

[6] Parag A. Pathak. The mechanism design approach to student assignment. *Annual Review of Economics*, 3(1):513–536, 2011. 25, 27

[7] Richard Rothstein. The racial achievement gap, segregated schools, and segregated neighborhoods — a constitutional insult. *Race and Social Problems*, 6(4), Dec 2014. 21

[8] Richard Rothstein. Should we force integration on those who don't want it?, and other commonplace questions about race relations. Economic Policy Institute, Working Economics Blog, March 2015. 21

[9] Boston Public Schools. Student assignment policy. `http://www.bostonpublicschools.org/assignment`. Accessed: 2015-09-15. 21

[10] Cambridge Public Schools. About controlled choice. `http://www.cpsd.us/departments/frc/making_your_choices/about_controlled_choice`. Accessed: 2015-09-15. 21

[11] Jefferson County Public Schools. Student assignment. `http://www.jefferson.k12.ky.us/departments/StudentAssignment/Index.html`. Accessed: 2015-09-15. 21

[12] Portland Public Schools. `http://www.pps.k12.or.us/files/enrollment-transfer/Room_to_Optimal_Size_Analysis_2015-09-10.v2.pdf`. Accessed: 2016-01-03. 35, 38, 39

[13] San Francisco Public Schools. SFUSD: The assignment process. `http://www.sfusd.edu/en/enroll-in-sfusd-schools/how-student-assignment-works/the-assignment-process.html`. Accessed: 2015-09-15. 21

[14] Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, pages 23–37, 1974. 27

[15] Jeremy Adam Smith. As parents get more choice, S.F. schools resegregate. San Francisco Public Press, Feb 2015. 27, 86