Fair Partitioning of Public Resources: Redrawing District Boundary to Minimize Spatial Inequality in School Funding

Nuno Mota*
Max Planck Institute for Software Systems, Germany

Negar Mohammadi*
Tehran Institute for Advanced Studies, Iran

Palash Dey
Indian Institute of Technology Kharagpur, India

Krishna P. Gummadi
Max Planck Institute for Software Systems, Germany

Abhijnan Chakraborty†
Indian Institute of Technology Delhi, India

ABSTRACT

Public schools in the US offer tuition-free primary and secondary education to students, and are divided into school districts funded by the local and state governments. Although the primary source of school district revenue is public money, several studies have pointed to the inequality in funding across different school districts. In this paper, we focus on the spatial geometry/distribution of such inequality, i.e., how the highly funded and lesser funded school districts are located relative to each other. Due to major reliance on local property taxes for school funding, we find existing school district boundaries promoting financial segregation, with highly-funded school districts surrounded by lesser-funded districts and vice-versa.

To counter such issues, we formally propose Fair Partitioning problem to divide a given set of schools into k districts such that the spatial inequality in district-level funding is minimized. However, the Fair Partitioning problem turns out to be computationally challenging, and we formally show that it is strongly NP-complete. We further provide a greedy algorithm to offer practical solution to Fair Partitioning, and show its effectiveness in lowering spatial inequality in school district funding across different states in the US.

CCS CONCEPTS

• **Human-centered computing**;

ACM Reference Format:

Nuno Mota, Negar Mohammadi, Palash Dey, Krishna P. Gummadi, and Abhijnan Chakraborty. 2021. Fair Partitioning of Public Resources: Redrawing District Boundary to Minimize Spatial Inequality in School Funding. In *Proceedings of the Web Conference 2021 (WWW '21), April 19–23, 2021, Ljubljana, Slovenia.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3442381.3450041

1 INTRODUCTION

In both online and offline worlds, a population is often grouped for practical benefits. For example, city residents are grouped into municipal wards; faculty and students are grouped into university departments; voters are grouped into electoral constituencies. Similarly, users in large online systems are grouped into clusters to provide recommendations at scale. Importantly, people belonging to a particular group (or cluster) *share the same fate*, such as, students share the same departmental resources; ward residents experience similar municipal services; or users in a cluster receive similar recommendations. When it concerns the distribution of resources, the group membership effectively determines the amount of resources different individuals get, raising concerns about the fairness of the distribution. In this paper, we focus on a prime example of such scenario – distribution of government funding across the public school districts in the United States.

Public schools offer tuition-free elementary and secondary education to students from all financial backgrounds. Every school has a *school attendance zone* which determines the neighborhood from which children will attend the school. In most states, schools are grouped into school districts for better administration, and the geographical area of a district covers the corresponding school attendance zones [11]. School districts are mostly dependent on public money to manage their expenses, where the major sources of revenues are local government funding (collected mainly from property taxes) and funding from the state governments (general formula assistance). Although the primary source of school district revenue is public money, multiple studies have criticized the existing school district boundaries to be promoting racial and financial segregation (or gentrification) [16, 36, 38]. For example, a recent report by EdBuild, a non-profit organization focused on improving the public school finance system in the US, claimed that *"Non-white school districts get $23 Billion less than white districts, despite serving the same number of students"* [16]. Such funding disparities can play a major role in the long-term educational and economic outcomes of the students.

In a departure from past literature, in this work, we focus on the spatial geometry / distribution of such inequality – i.e., how the highly funded and lesser funded school districts are located relative to each other. For instance, it would seem more unfair if highly-funded school districts were islands surrounded by poorly funded school districts, as opposed to there being a more gradual spread of high to low-funded school districts. This notion of fairness matches with *people's perception of inequality*, where people tend to approximate the overall income distributions and their place in those distributions by

*Both authors contributed equally to this research.
†Author was at the Max Planck Institute for Software Systems during this work.

WWW '21, April 19–23, 2021, Ljubljana, Slovenia

ACM ISBN 978-1-4503-8312-7/21/04.
https://doi.org/10.1145/3442381.3450041

taking cue from their local neighborhood [18, 22, 26]. Such *spatial notion of fairness* can also be interpreted as a special case of the individual fairness notion proposed by Dwork et al. [15]: *similarly located districts (i.e., neighboring districts) should have similar per-student funding.*

In this paper, we propose *Spatial Inequality Index* to capture how different individuals in a population are treated compared to their immediate neighbors. Although it can be applied in any online or offline setting involving users and some connection between them, such as inequality in satisfaction from social recommendations, clustering algorithms, or even broadly comparing utilities of 'similar' users (given a similarity notion), our primary focus in this work is on educational inequality – how different a school district is funded compared to its neighbors.

Our analysis reveals high spatial inequality in the funding received by different school districts across the US. Essentially, the heavy reliance on local property taxes to fund public schools have resulted in schools in poor neighborhoods receiving significantly lower levels of funding compared to schools in adjacent rich neighborhoods. By distributing funds more equally across districts, policymakers can equalize access to quality education—which in turn can end the vicious cycle of poverty for many students, improving social mobility [23, 27].

Towards that, we propose to redraw the district boundaries to amalgamate wealthy and poor neighborhoods, to minimize the spatial inequality in funding. We formally propose Fair Partitioning problem to divide a given set of nodes into k partitions such that the spatial inequality in a partition-level property is minimized. However, the Fair Partitioning problem turns out to be computationally challenging, and we prove that it is strongly NP-complete. Thus, we provide Greedy Partitioning, a greedy algorithm to move nodes between different partitions, to offer a practical solution to Fair Partitioning. Extensive experiments show the effectiveness of Greedy Partitioning in lowering spatial inequality in school district funding across different states in the US.

Finally, while our focus has been on the computational aspects of solving the school redistricting problem, such proposals would ultimately require community participation and necessary legal support. To give policymakers and community leaders a visual idea of the promise redistricting holds, we have deployed a web-based visualization tool at `https://redistricting.mpi-sws.org`. We believe that wide adoptions of such online tools can help foster further discussions to lower educational inequality in our societies.

2 INEQUALITY IN SCHOOL FUNDING

Every public school in the United States (be it elementary, middle or high school) has a dedicated attendance zone – a geographical area from which it admits students. These schools are organized into school districts governed by local school boards, where the district boundary is formed by amalgamating their constituent school attendance zones. While some school districts are either *Primary* (catering only to elementary schools) or *Secondary* (catering only to middle and high schools), majority of the school districts are *Unified* (i.e., they cover all type of schools in a geographical area). Given the existing school district boundaries in the US and their corresponding finances, our focus in this section is to understand whether all students enrolled into public schools are funded equally, regardless of the school districts they belong to.

2.1 Dataset gathered

We gathered the data on school district revenues from the National Center for Education Statistics (NCES) [30], which presents the detailed breakdown of state, local and federal funding for the 2015-16 school year (latest information as of August, 2020). We also collected information about the schools falling under each of the districts and the number of students enrolled in them. Additionally, NCES provides the geographical boundary of the school districts, as well as the school attendance zone boundaries [31]. Overall, we got data about $43,976$ schools in $14,528$ districts throughout the US. After removing the districts (and schools) with missing entries, no boundary information or zero enrollments, we considered around $39,656$ schools in $10,461$ districts.

To enable meaningful funding comparisons across school districts in different states, we adjust absolute revenues with 'Comparable Wage Index for Teachers (CWIFT)' at the school district level [10], again collected from the NCES portal [29]. CWIFT enables normalization of dollar amounts and make them comparable, so that a school district getting higher funding in a costly neighborhood can be properly contrasted with another district getting lesser funding in a cheaper area.

2.2 Distribution of per student funding across school districts

To compare the funding received by different districts, we further normalize the CWIFT-adjusted revenue by the number of students in each district. Figure 1 shows boxplots of per-student funding different districts received in every state in the US. Considering the median funding, we observe that New York, Wyoming and Alaska have the highest per-student funding ($\sim$ \$25,000), whereas, Arizona, Idaho and Florida have the lowest ones ($\sim$ \$10,000). This means that the students in at least half of the districts in New York or Wyoming get an educational funding that is at least 2.5 times higher than what students in half of the districts in Arizona or Florida get. Such educational inequalities may translate into disparities in further opportunities, impacting the future livelihood of the students. As evident from Figure 1, there is no uniformity in the median funding among other states as well.

Moreover, even within a state, there is high variability in the funding different districts receive per-student. For example, *Friend Public School District* in Nebraska got \$26,048 funding per-student; whereas in the same state, *Elkhorn Public School District* received only \$13,021. Similarly, two districts in New York – *Pine Plains Central School District* and *Orchard Park Central School District* – received \$31,040 and \$21,211 funding respectively, revealing a gap of \$10,000 per-student.

To further investigate the funding inequality within states, we consider the interquartile range (IQR) (i.e., the difference between 25th and 75th percentile values) of the box-plots for each state. We can observe in Figure 1 that New Mexico,

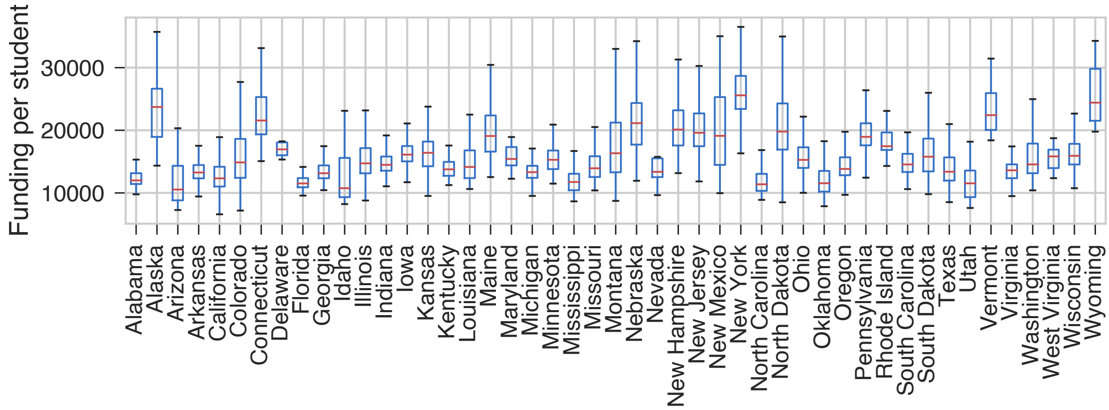

Figure 1: Box-plots showing funding per student across different school districts in different states. We see high variability in funding for school districts, both between different states and within individual states.

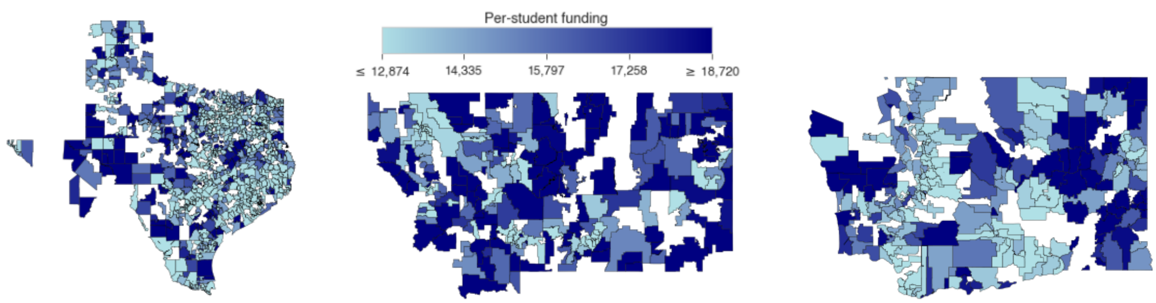

(a) Texas **(b) Montana** **(c) Washington**

Figure 2: In different states, lowly-funded school districts are co-located with highly-funded ones and vice versa (as evident from their colors in the maps). Due to missing entries (funding or boundary information) for some districts, the maps are not fully drawn.

Montana and North Dakota have the highest dispersion of per-student funding among their school districts, with IQR more than $\$7,000$. On the other hand, Alabama, Arkansas, Florida and Delaware have the least inequality of per-student funding across their districts, with IQR less than $\$2,000$.

2.3 Source of funding inequality

Since public schools offer tuition-free education, they rely on government funding to cover their expenditure. The funding primarily comes from local and state governments, with some supplementary funding from the federal government. On average, about 80% of local revenues for public school districts come from local property taxes [28]. Such reliance on property taxes creates huge disparity in the amount of funding going to different school districts, since the collection from property taxes is much higher in wealthier neighborhoods compared to poorer neighborhoods.

For multiple districts, the shortfall in local taxes is made up by the funding from state sources, with higher funding going to districts with lesser local revenue. Thus, state funding generally attempts to counterbalance the inequality created by the local revenue. However, this pattern is not consistent across states. In states like Arizona or Idaho, state funding is distributed similarly across school districts; whereas in states like West Virginia or Montana, few districts with higher local revenue end up getting more funding from state sources, thereby exacerbating the inequality. The contribution of federal sources is much lower compared to the state and local funding, and we observed that except few exceptions, federal funding is distributed more evenly over all school districts.

To summarize, while the revenue from local sources creates huge disparity, revenue from state governments tries to somewhat reduce the gap. However, since the overall contribution of federal sources is low, and there is no consistent direction in the distribution of state funding, the final funding per student closely mimic the trend in local revenues.

2.4 Spatial distribution of inequality

While the earlier analysis focused on the funding inequality across different school districts, it missed out on the *spatial geometry / distribution of inequality* – i.e., how the highly funded and lesser funded school districts are located relative

School District	Neighboring Districts
Colton School District, Washington (\$46, 137)	Clarkston School District (\$14, 483), Pullman School District (\$17, 521), Pomeroy School District (\$18, 428)
Doss Consolidated Common School District, Texas (\$48, 597)	Edinburg Consolidated Independent School District (\$11, 019), La Villa Independent School District (\$10, 908), Mercedes Independent School District (\$13, 354)
House Municipal Schools, New Mexico (\$31, 070)	South Conejos School District Re00 (\$15, 694), Archuleta County School District 50-Jt (\$14, 487), Centennial School District r0 (\$13, 166), Chama Valley Independent Schools (\$12, 270)
Greenport Union Free School District, New York (\$19, 795)	Sullivan West Central School District (\$59, 885), Deposit Central School District (\$37, 515), Downsville Central School District (\$30, 098)
Stamford School District, Connecticut (\$15, 356)	Regional High School District 19 (\$27, 881), Willington School District (\$29, 635), Coventry School District (\$23, 054), Ellington School District (\$26, 942)

Table 1: Examples of school districts with high disparity in per-student funding compared to their neighboring districts (funding amount is within parenthesis). In some places, a highly-funded school district is surrounded by lesser-funded districts; whereas in other places, a less-funded school district is surrounded by highly-funded districts.

to each other. For instance, for a given set of well-funded and poorly-funded school districts, it would seem more unfair if highly-funded school districts were islands surrounded by poorly funded school districts, as opposed to there being a more gradual spread of high to low-funded school districts.

Let's consider the *Greenport Union Free School District* in New York, which received \$20, 000 per-student funding and is surrounded by three other districts having per-student funding \$60, 000, \$38, 000 and \$30, 000 respectively. So, students in those three districts got 1.5 to 3 times more funding compared to the *Greenport* students, even though these four school districts are located next to each other. With such an island scenario, *Greenport* students would feel the inequality and unfairness more, compared to a situation where their neighboring districts had similar funding and higher funded districts were farther from their location. More examples of islands can be observed through existing school districts in the US and a few of them are listed in Table 1. In Figure 2, we can see similar funding inequalities in the district boundary maps for Texas, Montana and Washington states, and the pattern is similar in other states as well.

In summary, the existing school district boundaries create islands of highly-funded school districts surrounded by districts which are poorly-funded (and vice versa). The primary reliance on property taxes for public school education create such segregation of students based on their family's wealth, and the district boundaries tend to keep the wealth within the school district. This also has a circular effect of attracting more funding to the already high-funded schools, and the poorer districts being further impoverished. For instance, an article in *The Atlantic* [21], focusing on the educational inequalities in Connecticut, showed that high-poverty areas like Bridgeport have lower home values and thus the local government collects less taxes, whereas, homes in Darien are worth millions of dollars, resulting in high tax incomes. Schools in Darien have access to better facilities like school psychologists and personal laptops, which further attract wealthy homeowners, pushing up the local revenues further [21]. In this paper, we attempt to lower such spatial inequalities in school funding, but to do that, we need a concrete measure to compute the spatial inequality. Next, we propose one such measure.

3 QUANTIFYING SPATIAL INEQUALITY

A long line of works in welfare economics have proposed different *inequality indices* to quantify how unequally incomes/benefits are distributed over a population [3, 6, 19, 39, 42]. Formally, given a distribution/vector $\mathrm{y} = (y_1, \cdots, y_N) \in \mathcal{R}^N_{\geqslant 0}$, an inequality measure, $I : \bigcup_{n=1}^{\infty} \mathcal{R}^n_{\geqslant 0} \rightarrow \mathcal{R}_{\geqslant 0}$, maps any distribution/vector y to a non-negative real number $I(\mathrm{y})$. A distribution y is considered more equal than another distribution y' if and only if $I(\mathrm{y}) < I(\mathrm{y}')$. A popular inequality index is *Gini Index* [19], which captures the *relative mean absolute difference* of income between any two people in the population:

$$Gini = \frac{\sum_{i=1}^{N}\sum_{j=1}^{N}|y_i - y_j|}{2N^2\hat{y}} = \frac{\sum_{i=1}^{N}\sum_{j=1}^{N}|y_i - y_j|}{2N\sum_{i=1}^{N} y_i} \quad (1)$$

where $\hat{y}$ is the mean income of the population. *Gini* ranges between 0 to 1, with 0 denoting perfect equality. There are many other inequality indices as well, namely *Atkinson Index* [3], *Theil Index* [42], etc.

However, the inequality indices proposed in the past literature quantify the inequality in the overall population, missing an important aspect: *how individuals perceive inequality*. Recent research efforts have ventured into understanding people's perception of inequalities in different countries. Knell and Stix [26] have argued that people's inequality perception depend on their social position, as individuals typically do not observe the entire income distribution. Furthermore, Hauser and Norton [22] showed that people often rely on cues from their local environment to guess the overall distributions of income, and their place in those distributions. They further showed that these perceived inequalities drive people's behavior and preferences for redistribution [22]. Similar observations have been echoed by Ricci [35] and Gimpelson and Treisman [18]. Some researchers have also tied this perception of inequality and unfairness with *envy* [9, 12].

In this work, we attempt to account for people's perception of inequality and propose a measure to compute *spatial inequality*, where instead of comparing the income between every pair of individuals, we only compare the income of the neighbors. Here the notion of space need not be restricted to physical geography – it can easily be extended for individuals put into a n-dimensional abstract space. Such *spatial notion*

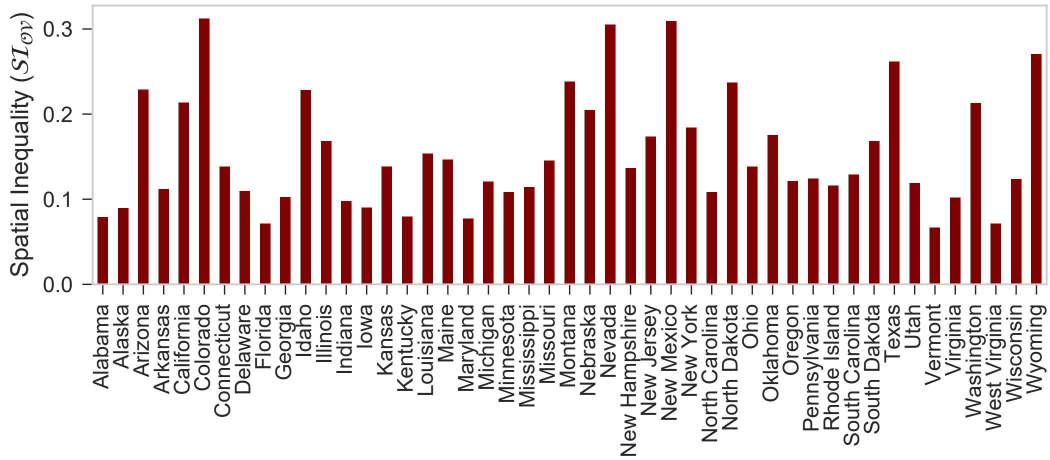

Figure 3: Spatial inequality in the per-student funding received by different school districts within the states.

of fairness can also be interpreted as a special case of the individual fairness notion proposed by Dwork et al. [15]: *similar individuals should have similar decision outcome.* In the context of school funding, we can interpret as: *the funding distribution would be fair if similarly located districts (i.e., neighboring districts) should have similar per-student funding.*

We formally express the inequality perceived by individual (school district) i as the average difference between the income (funding) y_i of i and its neighboring districts:

$$\mathcal{SI}_{\mathcal{IN}}(i) = \frac{1}{N_i}\sum_{j=1}^{N_i} |y_i - y_j| \tag{2}$$

where N_i is the number of neighbors of i.

To measure the overall inequality while taking into account the spatial distribution, we propose *Spatial Inequality Index* ($\mathcal{SI}_{\mathcal{OV}}$), adapted from Gini:

$$\mathcal{SI}_{\mathcal{OV}} = \frac{1}{N\hat{y}}\sum_{i=1}^{N} \mathcal{SI}_{\mathcal{IN}}(i) = \frac{\sum_{i=1}^{N} \frac{1}{N_i}\sum_{j=1}^{N_i} |y_i - y_j|}{\sum_{i=1}^{N} y_i} \tag{3}$$

where, instead of comparing the income of every pair of individuals, we only compare the neighbors, and then aggregate over all individuals. The value is further normalized by the number of individuals and the mean income in the population, allowing us to compare across different population groups. $\mathcal{SI}_{\mathcal{OV}}$ is zero when all individuals (school districts) have similar income (per-student funding). Higher the value of $\mathcal{SI}_{\mathcal{OV}}$, higher the spatial inequality.

Spatial inequality of per-student funding in different states: To compare the spatial inequality across different states, we compute $\mathcal{SI}_{\mathcal{OV}}$ over the per-student funding for different school districts within a state. We can see in Figure 3 that the spatial inequality in Colorado, Nevada and New Mexico are the highest among the states, implying that there is a large disparity in the funding students received in neighboring districts. On the other hand, Vermont, West Virginia, Florida and Alabama have relatively more equal distribution of per-student funding among co-located districts.

4 MINIMIZING SPATIAL INEQUALITY

In this work, we propose to reduce the existing spatial inequality in school funding by redistricting schools. We formulate the problem as a graph partitioning problem, where schools are the vertices and there is an edge between two schools if their school attendance zones share a common boundary.

4.1 Formal problem statement

We now formally write our problem definition. Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and two disjoint subsets $\mathcal{W}_1$ and $\mathcal{W}_2$ of $\mathcal{V}$, we say that $\mathcal{W}_1$ and $\mathcal{W}_2$ are neighbors of each other, and denote it by $\mathcal{W}_1 \sim \mathcal{W}_2$, if there exists an edge $\{u, v\} \in \mathcal{E}$ such that $u \in \mathcal{W}_1$ and $v \in \mathcal{W}_2$. For a subset $\mathcal{W} \subseteq \mathcal{V}$ of vertices, we denote the number of vertices in it by $|\mathcal{W}|$.

Definition 1 (Fair Partitioning). *Given an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, w : \mathcal{V} \longrightarrow \mathbb{R}^+, p : \mathcal{V} \longrightarrow \mathbb{R}^+)$ with two weights per vertex, an integer k denoting the number of districts, two integers L_{min} ($\geqslant 1$) and L_{max}, and a real number t, compute if there exits a partition $(\mathcal{V}_1, \dots, \mathcal{V}_k)$ of $\mathcal{V}$ such that (i) $L_{min} \leqslant |V_i| \leqslant L_{max}\ \forall i \in [k]$, (ii) the induced graph $\mathcal{G}[\mathcal{V}_i]$ is connected, and (iii) we have the following*

$$\sum_{i=1}^{k} \sum_{i<j\leqslant k, \mathcal{V}_i \sim \mathcal{V}_j} \left| \frac{\sum_{u\in\mathcal{V}_i} w_u}{\sum_{u\in\mathcal{V}_i} p_u} - \frac{\sum_{u\in\mathcal{V}_j} w_u}{\sum_{u\in\mathcal{V}_j} p_u} \right| \leqslant t$$

We denote an arbitrary instance of Fair Partitioning *by $(\mathcal{G}, k, L_{min}, L_{max}, t)$.*

4.2 Hardness results

We show that the Fair Partitioning problem is NP-complete even if the weights of every vertex is encoded in unary; that is the problem is strongly NP-complete. For that, we reduce from the 3-Partition problem, known to be strongly NP-complete [24]. The 3-Partition problem is defined as follows.

Definition 2 (3-Partition). *Given a multi-set $\mathcal{A} = \{a_i : i \in [3n]\}$ of $3n$ positive integers, compute if there exists a partition $(S_j)_{j\in[n]}$ of $\mathcal{A}$ into n sets such that each $S_j, j \in [n]$ contains exactly 3 elements from $\mathcal{A}$ and all the sets $S_j, j \in [m]$ have the same sum of its elements. We denote an arbitrary instance of 3-Partition by $\mathcal{A}$.*

Theorem 1. *The* Fair Partitioning *problem is strongly* NP*-complete even if $\mathcal{G}$ is a complete graph and every district contains exactly 3 vertices.*

Proof. The Fair Partitioning problem clearly belongs to NP. To show NP-hardness, we reduce from 3-Partition. Let $\mathcal{A} = \{a_i : i \in [3n]\}$ be an arbitrary instance of 3-Partition. We consider the following instance $(\mathcal{G} = (\mathcal{V}, \mathcal{E}, w), k, L_{\min}, L_{\max}, t)$ of Fair Partitioning.

$$\mathcal{V} = \{v_i : i \in [3n]\}$$
$$w_i = a_i, p_i = 1$$
$$k = n, L_{\min} = 3, L_{\max} = 3, t = 0$$

The graph $\mathcal{G}$ is a complete graph. We now claim that the two instances are equivalent.

In one direction, let us assume that the 3-Partition instance is a yes instance. Let $(S_j)_{j\in[n]}$ of $\mathcal{A}$ into n sets such that each set contains exactly 3 elements and all the sets $S_j, j \in [m]$ have the sum of its elements. We define a partition $(\mathcal{V}_j)_{j\in[n]}$ of $\mathcal{V}$ into $k = n$ districts as: $\mathcal{V}_j = \{v_i : i \in [3n], a_i \in S_j\}$ for $j \in [n]$. Since the sum of the weights of the vertices of every district is the same and every district has exactly 3 vertices, we have the following for every $i, j \in [n]$.

$$\frac{\sum_{u\in\mathcal{V}_i} w_u}{|\mathcal{V}_i|} - \frac{\sum_{u\in\mathcal{V}_j} w_u}{|\mathcal{V}_j|} = 0$$

Hence the Fair Partitioning instance is also a yes instance.

On the other direction, let us assume that the Fair Partitioning instance is a yes instance. Let $(\mathcal{V}_j)_{j\in[n]}$ be a partition of $\mathcal{V}$ into $k = n$ districts such that we have the following for every $i, j \in [n]$ (since $t = 0$).

$$\frac{\sum_{u\in\mathcal{V}_i} w_u}{|\mathcal{V}_i|} - \frac{\sum_{u\in\mathcal{V}_j} w_u}{|\mathcal{V}_j|} = 0$$

We define a partition $(S_j)_{j\in[n]}$ of $\mathcal{A}$ into n sets as follows: $S_j = \{a_i : i \in [3n], v_i \in \mathcal{V}_j\}$ for $j \in [n]$. Since we have $L_{\min} = L_{\max} = 3$, we have $|S_j| = 3$. Also if there exists two indices $r, s \in [n]$ such that these two sets have a different sum of its elements, then we have the following.

$$\frac{\sum_{u\in\mathcal{V}_r} w_u}{|\mathcal{V}_r|} - \frac{\sum_{u\in\mathcal{V}_s} w_u}{|\mathcal{V}_s|} \neq 0$$

This is a contradiction since we have $t = 0$. □

The proof of Theorem 1 shows that the Fair Partitioning problem is also inapproximable in polynomial time within any factor $\alpha(\cdot)$ where $\alpha(\cdot)$ is any computable function if $P \neq$ NP.

Corollary 1. *Let $\alpha(\cdot)$ is any computable function. Then there does not exist any polynomial time approximation algorithm for the* Fair Partitioning *problem if $P \neq$ NP.*

We show next that Fair Partitioning is NP-complete even when we wish to partition the graph into 2 districts (that is $k = 2$) and the graph is planar. For that, we reduce from the NP-complete problem Planar 2-Disjoint Connected Partitioning [20] which is defined as follows.

Definition 3 (Planar 2-Disjoint Connected Partitioning). *Given a connected planar graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and two disjoint nonempty sets $\mathcal{Z}_1, \mathcal{Z}_2 \subset \mathcal{V}$, compute if there exists a partition $(\mathcal{V}_1, \mathcal{V}_2)$ of $\mathcal{V}$ such that $\mathcal{Z}_1 \subseteq \mathcal{V}_1, \mathcal{Z}_2 \subseteq \mathcal{V}_2, \mathcal{G}[\mathcal{V}_1]$ and $\mathcal{G}[\mathcal{V}_2]$ are both connected. We denote an arbitrary instance of* Planar 2-Disjoint Connected Partitioning *by $(\mathcal{G}, \mathcal{Z}_1, \mathcal{Z}_2)$.*

Theorem 2. *The* Fair Partitioning *problem is* NP*-complete even if we need to partition the graph into exactly 2 connected districts and the underlying graph is planar.*

Proof. The Fair Partitioning problem clearly belongs to NP. To show NP-hardness, we reduce from the Planar 2-Disjoint Connected Partitioning problem. Let $(\mathcal{G} = (\mathcal{V}, \mathcal{E}), \mathcal{Z}_1, \mathcal{Z}_2)$ be an arbitrary instance of Planar 2-Disjoint Connected Partitioning. Let us assume without loss of generality that $|\mathcal{Z}_1| \geqslant |\mathcal{Z}_2| \geqslant 2$, and the number n of vertices in $\mathcal{G}$ is at least 50 since the Planar 2-Disjoint Connected Partitioning problem is known to be NP-complete even under this restriction [43, Theorem 1]. We consider the following instance $(\mathcal{G}' = (\mathcal{V}', \mathcal{E}', w), k, L_{\min}, L_{\max}, t)$ of Fair Partitioning. Let $z_1 \in \mathcal{Z}_1$ and $z_2 \in \mathcal{Z}_2$ be two arbitrary vertices from $\mathcal{Z}_1$ and $\mathcal{Z}_2$.

$$\mathcal{V}' = \{v_u : u \in \mathcal{V}\} \cup \mathcal{D}_1 \cup \mathcal{D}_2, \text{ where}$$
$$\mathcal{D}_1 = \left\{d_i^1 : i \in \left[|\mathcal{Z}_1|\, n^5\right]\right\}$$
$$\mathcal{D}_2 = \left\{d_i^2 : i \in \left[|\mathcal{Z}_2| n^2\right]\right\}$$
$$\mathcal{E}' = \{\{v_a, v_b\} : \{a, b\} \in \mathcal{E}\}$$
$$\cup \{\{d_i^1, d_j^1\} : i, j \in \left[|\mathcal{Z}_1|\, n^5\right], j = i + 1\} \cup \{\{z_1, d_1^1\}\}$$
$$\cup \{\{d_i^2, d_j^2\} : i, j \in \left[|\mathcal{Z}_2| n^2\right], j = i + 1\} \cup \{\{z_2, d_1^2\}\}$$
$$w(x) = \begin{cases} n^5 & \text{if } x = v_u \text{ for some } u \in \mathcal{Z}_1 \\ n^2 & \text{if } x = v_u \text{ for some } u \in \mathcal{Z}_2 \\ 1 & \text{otherwise} \end{cases}$$
$$p(x) = 1 \forall x \in \mathcal{V}'$$
$$k = 2, L_{\min} = |\mathcal{Z}_2| n^2, L_{\max} = |\mathcal{V}'|, t = \frac{6}{|\mathcal{Z}_2| n}$$

We now claim that the two instances are equivalent. In one direction, let us assume that the Planar 2-Disjoint Connected Partitioning instance is a yes instance. Let $(\mathcal{V}_1, \mathcal{V}_2)$ be a partition of $\mathcal{V}$ such that (i) $\mathcal{G}[\mathcal{V}_i]$ is connected and $\mathcal{Z}_i \subseteq \mathcal{V}_i$ for $i \in [2]$. Let us define $\ell_i = |\mathcal{V}_i \setminus \mathcal{Z}_i|$ for $i \in [2]$. We consider the partition $(\mathcal{V}_1', \mathcal{V}_2')$ of $\mathcal{V}'$ where $\mathcal{V}_1' = \{v_u : u \in \mathcal{V}_1\} \cup \mathcal{D}_1$ and $\mathcal{V}_2' = \mathcal{V}' \setminus \mathcal{V}_1'$. Since $\mathcal{G}[\mathcal{V}_i]$ is connected, it follows that $\mathcal{G}'[\mathcal{V}_i']$ is also connected for $i \in [2]$. Also, for the district $\mathcal{V}_1'$, we have

$$\left|\frac{\sum_{u\in\mathcal{V}_1'} w_u}{|\mathcal{V}_1'|} - \frac{\sum_{u\in\mathcal{V}_2'} w_u}{|\mathcal{V}_2'|}\right|$$
$$= \left|\frac{2|\mathcal{Z}_1| n^5 + \ell_1}{|\mathcal{Z}_1| n^5 + |\mathcal{Z}_1| + \ell_1} - \frac{2|\mathcal{Z}_2| n^2 + \ell_2}{|\mathcal{Z}_2| n^2 + 2 + \ell_2}\right|$$

$$\leqslant \max_{\ell_1+\ell_2\leqslant n}\left|\frac{2|\mathcal{Z}_1|n^5+\ell_1}{|\mathcal{Z}_1|n^5+|\mathcal{Z}_1|+\ell_1}-\frac{2|\mathcal{Z}_2|n^2+\ell_2}{|\mathcal{Z}_2|n^2+2+\ell_2}\right|$$
$$=\max_{\ell_1+\ell_2\leqslant n}\left|\frac{2|\mathcal{Z}_1|+\ell_1}{|\mathcal{Z}_1|n^5+|\mathcal{Z}_1|+\ell_1}-\frac{\ell_2+4}{|\mathcal{Z}_2|n^2+2+\ell_2}\right|$$
$$\leqslant \frac{n+4}{|\mathcal{Z}_2|n^2+n+2}-\frac{2}{n^5+1}\leqslant\frac{5}{|\mathcal{Z}_2|n}<t$$

In the other direction, let us assume that the Fair Partitioning instance is a yes instance. Let $(\mathcal{V}_1', \mathcal{V}_2')$ be a partition of $\mathcal{V}'$ such that $\mathcal{G}[\mathcal{V}_i']$ is connected for $i \in [2]$ and

$$\left|\frac{\sum_{u\in\mathcal{V}_1'} w_u}{|\mathcal{V}_1'|}-\frac{\sum_{u\in\mathcal{V}_2'} w_u}{|\mathcal{V}_2'|}\right|\leqslant t$$

Let us define a partition $(\mathcal{V}_1, \mathcal{V}_2)$ of $\mathcal{V}$ as $\mathcal{V}_1 = \{u \in \mathcal{V} : v_u \in \mathcal{V}_1'\}$ and $\mathcal{V}_2 = \mathcal{V} \setminus \mathcal{V}_1$. We claim that $\mathcal{Z}_1 \subseteq \mathcal{V}_1$. Suppose not, then we consider two cases. In the first case, suppose we have $\mathcal{D}_1 \subseteq \mathcal{V}_1'$, Then we have

$$\frac{\sum_{u\in\mathcal{V}_2'} w_u}{|\mathcal{V}_2'|}\geqslant n^2 \text{ and } \frac{\sum_{u\in\mathcal{V}_1'} w_u}{|\mathcal{V}_1'|}\leqslant 2-\frac{1}{n^5}$$

and thus we have

$$\left|\frac{\sum_{u\in\mathcal{V}_1'} w_u}{|\mathcal{V}_1'|}-\frac{\sum_{u\in\mathcal{V}_2'} w_u}{|\mathcal{V}_2'|}\right|\geqslant n^2-2+\frac{1}{n^5}>t$$

which is a contradiction. In the second case, we have $\mathcal{V}_1' \subseteq \mathcal{D}_1$. Then we have

$$\frac{\sum_{u\in\mathcal{V}_1'} w_u}{|\mathcal{V}_1'|}=1 \text{ and } \frac{\sum_{u\in\mathcal{V}_2'} w_u}{|\mathcal{V}_2'|}\geqslant\frac{2|\mathcal{Z}_1|n^5+2n^2}{|\mathcal{Z}_1|n^5+n}>2$$

and thus we have

$$\left|\frac{\sum_{u\in\mathcal{V}_1'} w_u}{|\mathcal{V}_1'|}-\frac{\sum_{u\in\mathcal{V}_2'} w_u}{|\mathcal{V}_2'|}\right|>1>t$$

which is a contradiction. We now claim that $\mathcal{Z}_2 \subseteq \mathcal{V}_2$. Suppose not, then we have the following which is a contradiction.

$$\left|\frac{\sum_{u\in\mathcal{V}_1'} w_u}{|\mathcal{V}_1'|}-\frac{\sum_{u\in\mathcal{V}_2'} w_u}{|\mathcal{V}_2'|}\right|$$
$$\geqslant \min_{\ell_1+\ell_2\leqslant n}\left|\frac{2|\mathcal{Z}_1|n^5+n^2+\ell_1}{|\mathcal{Z}_1|n^5+|\mathcal{Z}_1|+\ell_1}-\frac{(2|\mathcal{Z}_2|-1)n^2+\ell_2}{|\mathcal{Z}_2|n^2+2+\ell_2}\right|$$
$$=\min_{\ell_1+\ell_2\leqslant n}\left|\frac{|\mathcal{Z}_1|n^5+n^2-|\mathcal{Z}_1|}{|\mathcal{Z}_1|n^5+|\mathcal{Z}_1|+\ell_1}-\frac{(|\mathcal{Z}_2|-1)n^2-2}{|\mathcal{Z}_2|n^2+2+\ell_2}\right|$$
$$\geqslant\frac{1}{2|\mathcal{Z}_2|}$$
$$>t$$

Hence we have $\mathcal{Z}_i \subseteq \mathcal{V}_i$ for $i \in [2]$. Moreover since, for $i \in [2]$, $\mathcal{G}'[\mathcal{V}_i']$ is connected, $\mathcal{G}[\mathcal{V}]$ is also connected. Hence the Planar 2-Disjoint Connected Partitioning instance is a yes instance.

By proving Fair Partitioning is NP-complete under the aforementioned special settings, we conclude that it is NP-complete in its general form. Intuitively, if the conclusion is not true and there is a polynomial time algorithm for the generic problem, then one can apply the generic algorithm to the special setting as well, contradicting its NP-hardness result. Formally, there is a reduction from the special setting to the general form which simply ignores the additional structures that the special setting assumes. □

Input: *Districts* = Set of existing districts in a state.
Result: A partition into *Districts* with potentially reduced spatial inequality.
Initialize $\mathcal{S} \in \mathcal{R}$ as the state's aggregated per-student funding, $\mathcal{D} \in$ *Districts* as an arbitrary district, and $\mathcal{F}$: *Districts* $\rightarrow \mathcal{R}$ as a function that, given a district, returns its funding per-student.

```
for district ∈ Districts do
    if |F(district) − S| > |F(D) − S| then
        Set D ← district
    end
end
for school ∈ D do
    if school is bordering D′ ∈ Districts then
        Set diff_before ← |F(D) − F(D′)|
        Set diff_after ←
          |F(D \ {school}) − F(D′ ∪ {school})|
        if diff_after < diff_before then
            Set D ← D \ {school}
            Set D′ ← D′ ∪ {school}
        end
    end
end
```

Algorithm 1: Greedy Partitioning algorithm.

4.3 Greedy algorithm

Given the high complexity of the Fair Partitioning problem, we propose Greedy Partitioning– a greedy heuristic algorithm to minimize spatial inequality, by redistributing existing schools into districts, and thus redefining district boundaries. Greedy Partitioning starts with an initial arbitrary (or the existing) partitioning and then greedily moves nodes between adjacent partitions such that spatial inequality is minimized. More precisely, since minimal $\mathcal{SI}_{\mathcal{OV}}$ is achieved when all districts within a state have equal per-student funding (i.e., equal to the overall state's per-student funding), the algorithm greedily selects the district whose per-student funding most deviates from the whole state's. Then, since $\mathcal{SI}_{\mathcal{IN}}$ is minimized when all neighbors have the same per-student funding, the algorithm iterates over all schools at selected district's border (i.e., all schools which can *reasonably* be redistricted) and immediately redistricts any that would bring two districts' funding closer. This process (illustrated in Algo. 1) is then repeated iteratively until either no more schools can be redistricted to reduce spatial inequality, or a maximum threshold of iterations has been reached without significant improvement.

Despite its simplicity, the algorithm requires several considerations when adapting to this use-case. First, only schools at a district's border (with an adjacent district) can be selected as equalization candidates – as opposed to picking any arbitrary school. Second, despite spatial inequality easily being minimized by merging all districts, we enforced that no district can cease to exist (i.e., they are forced to retain at least one school at all times), as to preserve a necessary set of K partitions for which spatial inequality should be minimized. And third, every schools' number of students (and initial total funding) will be treated as constant, even after redistricting.

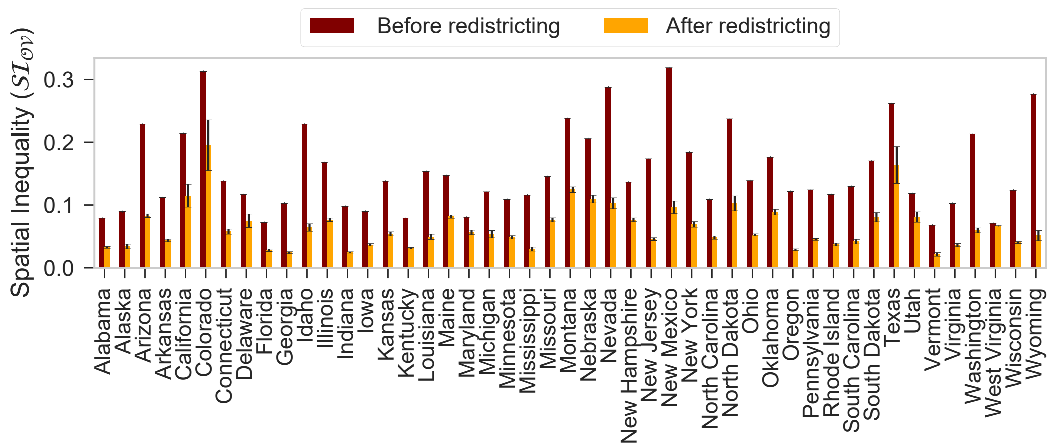

Figure 4: Though the effectiveness in minimizing spatial inequality varies across states, it gets significantly reduced for every state (while maintaining same number of districts). States like New Mexico and Wyoming demonstrate substantial improvements – becoming similar to other states with much lower initial inequality values.

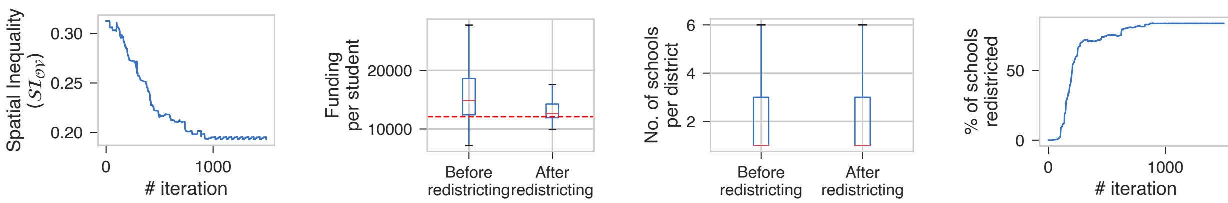

(a) (b) (c) (d)

Figure 5: Analysis of Greedy Partitioning applied on the state of Colorado: (a) we observe spatial inequality being successfully minimized, from 0.31 to below 0.20; (b) the average per-student funding (across districts) comes much closer to the whole state's – marked as a dashed red line. We notice in (c) that there is no general trend for districts to either become disproportionately larger (or smaller) as a byproduct of Greedy Partitioning. But (d) shows that over 80% of schools are still needed to be redistricted for the observed inequality reduction.

Closely following the formal definition for Fair Partitioning, Greedy Partitioning's parameterization requires three variables: (i) $\mathcal{G}$, a school neighborhood graph partitioned into k districts; (ii) l_{min}, the minimum number of schools per district (optional, but required to prevent an initial set of k partitions from merging); and (iii) l_{max}, the maximum number of schools per district (optional, but required to prevent certain districts from siphoning too many schools). Unless stated otherwise, throughout our analysis, we assume $\mathcal{G}$ to be the existing school district assignment, $l_{min} = 1$ and $l_{max} = \infty$.

4.4 Complexity of Greedy Partitioning

Greedy Partitioning yields a computational complexity of $O(I \times D^2 \times S)$, where I, D and S represent the number of iterations, districts and schools respectively. A single iteration of the algorithm includes two computational loops. The first iterates over all districts, with each step trivially completing in $O(1)$. The second loop iterates over all schools, with each step completing in $O(S)$ as each school's neighbors would have to be tested for assignment. This would lead to an overall complexity of $O(I \times D \times S^2)$. By utilizing dynamic programming, i.e., maintaining/updating necessary information for the later loop requires $O(D)$, which leads to a final complexity of $O(I \times D^2 \times S)$.[1] Empirically, the algorithm took 16.36 seconds on average across all states in a quad-core 2.4GHz CPU and 16GB of RAM, compared to an exhaustive brute-force algorithm taking weeks even for smaller states.

5 EXPERIMENTAL EVALUATION

Since ours is the first attempt at formalizing spatial inequality and proposing to minimize it, there is no prior baseline. Hence, we use the existing district assignment as a baseline in our experiments. In Figure 4, we can observe spatial inequality indices for all states in the US, before and after applying our proposed Greedy Partitioning algorithm. Figure 4 shows that our approach can significantly reduce spatial inequality for every state. This effect is especially noticeable for states like New Mexico and Wyoming, among the highest spatial inequality states. In this section, we further investigate the impact of Greedy Partitioning along several dimensions.

[1] The full algorithm and source codes can be found at `https://github.com/nunomota/spatial-inequality`.

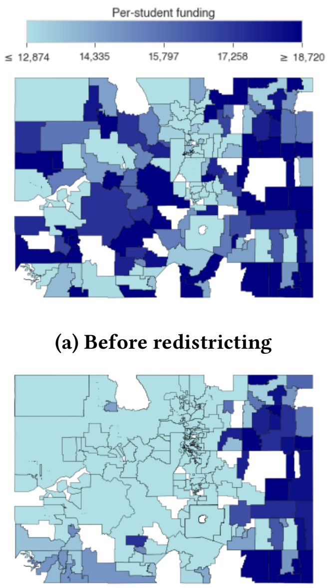

(a) Before redistricting

(b) After redistricting

Figure 6: District maps and their per-student funding before and after applying Greedy Partitioning, for the state of Colorado.

5.1 The case of Colorado

As an illustrative example of our algorithm's effect on a state's landscape, we now look into the specific case of Colorado – one of the two states with highest initial spatial inequality. In Figure 5a, we observe that Greedy Partitioning was indeed able to continuously reduce inequality during most of its iterations, starting from an initial spatial inequality of 0.31 to below 0.20 by the 1000th iteration. Moreover, Figure 5b shows that the per-student funding observed across all districts within the sate, become much closer to the state's average after the algorithm was applied. In Figure 5c, we see that districts maintained a similar number of schools before and after the algorithm was applied – signifying that there is neither a general tendency for small districts to absorb neighboring schools, nor for large districts to either fragment or grow. Interestingly, Figure 5d indicates that over 80% of all schools still needed to be redistricted. Taking into consideration that most schools had already been redistricted before iteration #300, and that at this stage spatial inequality had only reached its halfway point (Fig. 5a), it seems unlikely that an optimal result could have been achieved with a much lower percentage of schools redistricted. Lastly, in Figure 6 we can see the maps for the state of Colorado before and after applying Greedy Partitioning. As expected, we observe a much smoother gradient between districts' per-student funding.

5.2 Cross-state similarities

Similar to the artifacts observed for Colorado, even at the country level, Greedy Partitioning (i) is able to significantly reduce the variance in per-student funding distributions for most states (Figure 7); (ii) does not significantly increase sizes of districts, instead decreasing it for states with larger districts (Figure 8); however (iii) requires a substantial amount of schools to be redistricted in order to minimize spatial inequality. Figure 9 highlights this, where states like Washington or New Mexico require more than 75% of their schools being redistricted. But, we also observe that mitigating the problem for certain states requires much less change than others. For example, redistricting only 10% of schools in Nevada would already originate 75% of the final minimization achieved by our algorithm, whereas states like Oklahoma would require redistricting over 40% of their schools before observing even 25% of the achieved spatial inequality reduction.

5.3 Minimizing inequality further

Up to this point, we enforced that our algorithm preserved at least one school per district, as to prevent them from merging. If one's goal is to minimize spatial inequality between districts, the trivial solution for the problem would be to merge all districts into one. However, it would be an administrative nightmare to have too many schools under the same governance. As such, preserving the initial set of districts would allow for a more realistic setting. Alternatively, one could also allow districts to merge but constrain that no district can surpass a predefined amount of schools. Figure 10 shows the average inequality (country-wide) under this setting, with different thresholds for the maximum number of schools. We see that by increasing the maximum number of allowed schools per district – whilst simultaneously allowing districts to merge – spatial inequality can be reduced further, tending towards 0 as all districts merge within a state.

Though our focus in this paper is on minimizing spatial inequality, it is important to note that since minimizing spatial inequality ensures lesser disparity among neighbors, it would also lower any inequality index computed over the overall population (e.g., *Gini*).

6 RELATED WORK

In this section, we briefly review the related research efforts in graph partitioning, fair partitioning and school redistricting.

6.1 Graph partitioning

The field of graph partitioning covers a multidisciplinary spectrum of research efforts. Dhillon [14] proposed methods for similarity-based clustering of words and documents; Karypis et al. [25] applied similar methods for efficient transistor placement on electronic chips. Despite approximation methods being available for particular subsets of the problem (e.g., Andreev and Racke [2]), due to its generally high complexity, many algorithms start from an initial partition and then gradually refine it towards an objective. Predari and Esnard [33] leverage this to maximize parallel computing throughput - by assigning instructions over k distributed processors. Also, Qian et al. [34] demonstrated the benefits of distilling social networks' partitions with added constraints. We adopted this approach of starting with an initial partition, and greedily move schools to lower funding inequality.

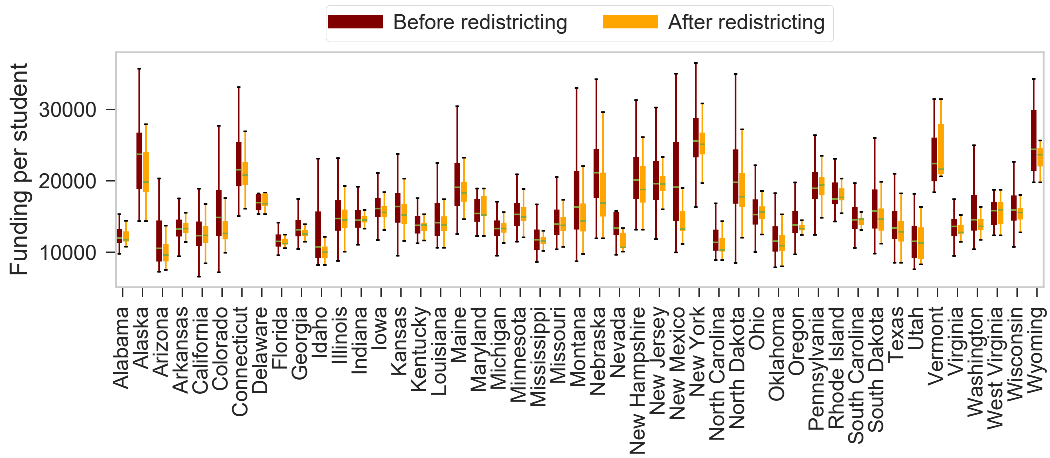

Figure 7: By comparing per student funding before and after redistricting, we observe that we significantly reduce variance in the later distributions whilst minimizing spatial inequality. This difference is especially noticeable for states like Colorado, Idaho, New Mexico and Wyoming.

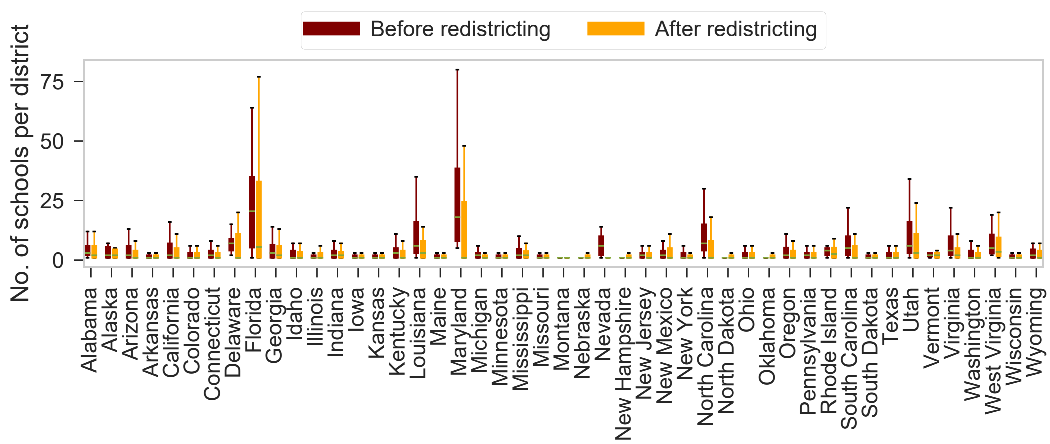

Figure 8: After redistricting, we observe that the median number of schools per district is generally reduced across all states. This is most noticeable in Florida, Maryland and Nevada. Moreover, variance in the distributions is broadly reduced. An interesting exception to this is Florida, where a few large districts have seemingly formed.

6.2 Fair Partitioning

The problem of fair partitioning has been explored in different domains. Chen et al. [7] proposed methods for proportionally fair wireless network resources' attribution; Devulapalli [13] analyzed methods for workload distribution based on land segmentation. Several works have also looked into fair division of assets. For instance, Abebe et al. [1] have proposed a a modification of the cake cutting problem to satisfy envy-freeness in nodes assigned to different graph partitions. Patro et al. [32] applied envy-free division in the context of recommendations. Stoica et al. [40] proposed an approach to minimize margin of victory in political partitions. Similarly, Bredereck et al. [4] proposed to satisfy value equity in resource allocation – considering a social network of mediated interactions among individuals. Chen and Shah [8] created a setting where imperfect knowledge could be represented as a partial graph coverage, to estimate perceived disparities among individuals. Although these works provide empirically proven methods to reduce some notion of inequity, in our work, we expand on them and provide a new measure for geospatial inequity – the spatial inequality index.

6.3 School funding and redistricting

Prior works have looked into different aspects of the current educational system in the US and their funding situation. EdBuild showed that school districts with majority of students of color receive $23 billion less in education funding than predominantly white school districts [5]. Moreover, several works, such as by Satz [37] and Swift [41], have shown that school

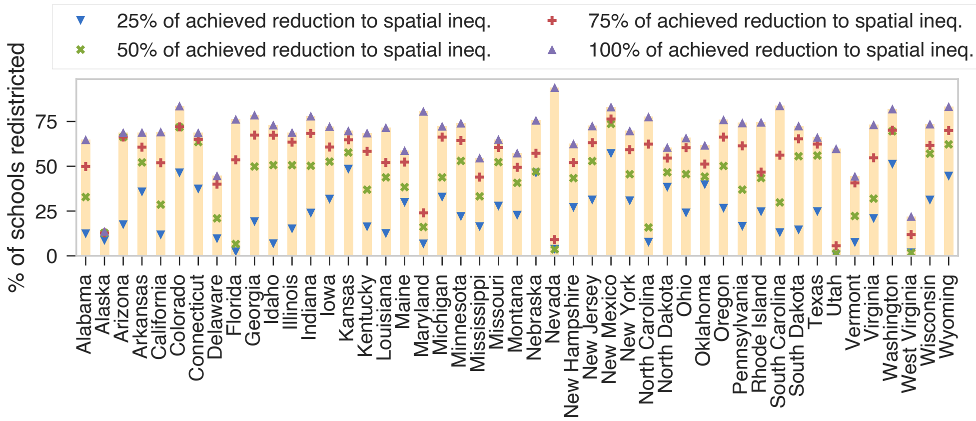

Figure 9: Despite being able to reduce inequality for most states, we now observe that in most cases over half of the schools in a state would have to be redistricted. Moreover, in states like Oklahoma, to observe just 25% of the observed inequality reduction (achieved by Greedy Partitioning), one would need to redistrict over 40% of its schools. On the other hand, states like Nevada would already observe 75% reduction by redistricting 10% of its schools.

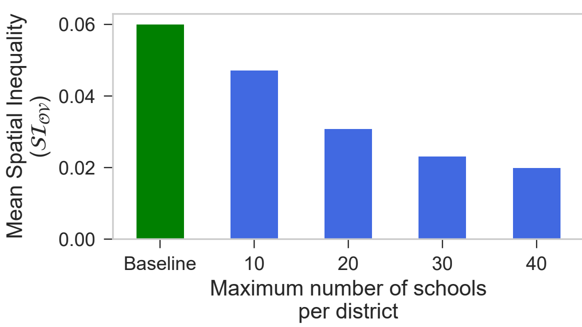

Figure 10: Country-wide mean spatial inequality, by enforcing different thresholds of maximum schools per district, $l_{max} \in \{\infty, 10, 20, 30, 40\}$. l_{min} is only enforced in the baseline (with the value of 1). The higher the number of schools allowed in merged districts, the lower the spatial inequality.

funding inequalities can significantly affect the quality of education for students of varying socio-economical backgrounds – arguing that such differences directly interfere with principles like *Equality of Opportunity.* Beyond school districting, Gentry et al. [17] highlighted geographical disparity for patients in need of liver transplant. Vickrey [44] showcased the implications of existing political district boundaries to control certain party's dominance (i.e., gerrymandering). Though we recognize the importance of redistricting as a multi-disciplinary effort for equitable treatment, in our study, we focus on its economical impact on equitability of school districts' funding.

7 CONCLUDING DISCUSSION

In this paper, we propose a new inequality index, named *Spatial Inequality Index* ($\mathcal{SI}_{OV}$). As opposed to other indices, this new proposal assumes a geospatial distribution of inequity as an important factor for perceived treatment. In other words, it considers groups of individuals in a population to have a neighborhood against which they assess their own treatment ($\mathcal{SI}_{IN}$). Minimizing this index is then analogous to a Fair Partitioning problem, where K neighboring partitions ultimately yield similar benefits – for some notion of "benefit". Due to this problem being NP-complete, we propose Greedy Partitioning, a heuristic algorithm that leverages local properties of our index to greedily minimize overall inequality.

To evaluate our proposal, we tackle the problem of school (re)districting. Prior works have shown the importance of school district boundaries in determining socio-economical outcomes of their students [5, 37, 41]. Moreover, this problem setting shares the same geospatial nature as $\mathcal{SI}_{OV}$, where a country is partitioned into K districts and each of which potentially being treated very differently from its immediate geographical neighbors. By focusing on per-student funding as a measure of benefit in our districts, our goal in this paper is to have schools in similar neighborhoods receiving similar per-student funding.

During our experiments, Greedy Partitioning was able to reduce spatial inequality for every state. Despite requiring a large percentage of schools to be redistricted, districts mostly maintained their dimensions across all states. As expected, after applying Greedy Partitioning, per-student funding in neighboring districts indeed became much closer. However, our algorithm relies on three separate assumptions: (i) only schools at a district's border can be redistricted; (ii) existing districts may not cease to exist (though their shapes may change); and (iii) all sources of funding assigned to a school will carry-over when it is redistricted.

While (i) and (ii) are more easily acceptable premises for our problem, assumption (iii) may not be entirely true in a real setting. After a school gets redistricted, state and/or federal criteria for funding attribution may change the amount of funds going to a particular school. Only funding from local sources could reasonably be expected to follow the school regardless of the district assignment (as the algorithm does not change a school's catchment area). One alternative to our proposal would be to first consider each schools' funds to be fully described by their local component, and similarly solve Fair Partitioning. However, this would require state and federal funds to still be assigned at the end – implying complete modeling over these components' criterion. Due to the intricacies of this alternative, we leave this for future work.

Despite its potential limitations, Greedy Partitioning proved successful in reducing the per-student funding disparities in public school districts. For reproducibility, we have made the source code publicly available at `https://github.com/nunomota/spatial-inequality`. Finally, the redistricting proposal would need community participation and policy support to have a real impact. To give policymakers and community leaders a visual idea of the promise redistricting holds, we have deployed a web-based visualization tool at `https://redistricting.mpi-sws.org`. We believe that such online tools can help foster further discussions and help move towards a society with lesser inequities.

REFERENCES

[1] Rediet Abebe, Jon Kleinberg, and David C Parkes. 2017. Fair Division via Social Comparison. In *AAMAS*. 281–289.

[2] Konstantin Andreev and Harald Racke. 2006. Balanced graph partitioning. *Theory of Computing Systems* 39, 6 (2006), 929–939.

[3] Anthony B Atkinson et al. 1970. On the measurement of inequality. *Journal of economic theory* 2, 3 (1970).

[4] Robert Bredereck, Andrzej Kaczmarczyk, and Rolf Niedermeier. 2018. Envy-free allocations respecting social networks. In *AAMAS*. 283–291.

[5] Lauren Camera. 2019. White Students Get More K-12 Funding Than Students of Color: Report. https://www.usnews.com/news/education-news/articles/2019-02-26/white-students-get-more-k-12-funding-than-students-of-color-report

[6] Abhijnan Chakraborty, Nuno Mota, Asia J Biega, Krishna P Gummadi, and Hoda Heidari. 2019. On the Impact of Choice Architectures on Inequality in Online Donation Platforms. In *WWW*. 2623–2629.

[7] Jiasi Chen, Mung Chiang, Jeffrey Erman, Guangzhi Li, KK Ramakrishnan, and Rakesh K Sinha. 2015. Fair and optimal resource allocation for LTE multicast (eMBMS): Group partitioning and dynamics. In *IEEE INFOCOM*.

[8] Yiling Chen and Nisarg Shah. 2017. *Ignorance is often bliss: Envy with incomplete information*. Technical Report.

[9] Yochi Cohen-Charash and Jennifer S Mueller. 2007. Does perceived unfairness exacerbate or mitigate interpersonal counterproductive work behaviors related to envy? *Journal of applied psychology* 92, 3 (2007), 666.

[10] Stephen Q. Cornman, Laura C. Nixon, Matthew J. Spence, Taylor Lori L., and Douglas E. Geverdt. 2019. *American Community Survey Comparable Wage Index for Teachers (ACS-CWIFT)*. Technical Report.

[11] Antonella Corsi-Bunker. 2015. Guide to the education system in the United States. *University of Minnesota* 23 (2015).

[12] Frank Cowell and Udo Ebert. 2008. Inequality and envy. In *Inequality and Opportunity: Papers From the Second Ecineq Society Meeting. Research on Economic Inequality*, Vol. 16. 37–47.

[13] Raghuveer Devulapalli. 2014. *Geometric partitioning algorithms for fair division of geographic resources*. Ph.D. Dissertation. University of Minnesota.

[14] Inderjit S Dhillon. 2001. Co-clustering documents and words using bipartite spectral graph partitioning. In *ACM KDD*. 269–274.

[15] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *ITCS*. 214–226.

[16] EdBuild. 2019. Non-white school districts get $23 Billion less than white districts, despite serving the same number of students. https://edbuild.org/content/23-billion

[17] Sommer E Gentry, Allan B Massie, Sidney W Cheek, Krista L Lentine, EH Chow, Corey E Wickliffe, Nino Dzebashvili, Paolo R Salvalaggio, Mark A Schnitzler, David A Axelrod, et al. 2013. Addressing geographic disparities in liver transplantation through redistricting. *American Journal of Transplantation* 13, 8 (2013), 2052–2058.

[18] Vladimir Gimpelson and Daniel Treisman. 2018. Misperceiving inequality. *Economics & Politics* 30, 1 (2018), 27–54.

[19] Corrado Gini. 1912. Variabilità e mutabilità. *Reprinted in Memorie di metodologica statistica (Ed. Pizetti E, Salvemini, T). Rome: Libreria Eredi Virgilio Veschi* (1912).

[20] Chris Gray, Frank Kammer, Maarten Löffler, and Rodrigo I Silveira. 2012. Removing local extrema from imprecise terrains. *Computational Geometry* 45, 7 (2012), 334–349.

[21] Adam Harris. 2019. The Whiter, Richer School District Right Next Door. https://www.theatlantic.com/education/archive/2019/08/segregated-school-districts-trapped-their-borders/595318/

[22] Oliver P Hauser and Michael I Norton. 2017. (Mis) perceptions of inequality. *Current Opinion in Psychology* 18 (2017), 21–25.

[23] C Kirabo Jackson, Rucker C Johnson, and Claudia Persico. 2015. The effects of school spending on educational and economic outcomes: Evidence from school finance reforms. *The Quarterly Journal of Economics* 131, 1 (2015).

[24] David S Johnson and Michael R Garey. 1979. *Computers and intractability: A guide to the theory of NP-completeness*. WH Freeman.

[25] George Karypis, Rajat Aggarwal, Vipin Kumar, and Shashi Shekhar. 1999. Multilevel hypergraph partitioning: applications in VLSI domain. *IEEE Trans. on Very Large Scale Integration Systems* 7, 1 (1999), 69–79.

[26] Markus Knell and Helmut Stix. 2020. Perceptions of inequality. *European Journal of Political Economy* 65 (2020), 101927.

[27] Julien Lafortune, Jesse Rothstein, and Diane Whitmore Schanzenbach. 2018. School finance reform and the distribution of student achievement. *American Economic Journal: Applied Economics* 10, 2 (2018), 1–26.

[28] Joel McFarland, Bill Hussar, Xiaolei Wang, Jijun Zhang, Ke Wang, Amy Rathbun, Amy Barmer, Emily Forrest Cataldi, and Farrah Bullock Mann. 2018. The Condition of Education 2018. NCES 2018-144. *National Center for Education Statistics* (2018).

[29] NCES. 2019. Comparable Wage Index for Teachers (CWIFT). https://nces.ed.gov/programs/edge/Economic/TeacherWage

[30] NCES. 2019. NCES Common Core of Data. https://nces.ed.gov/ccd

[31] NCES. 2019. The School Attendance Boundary Survey (SABS). https://nces.ed.gov/programs/edge/SABS

[32] Gourab K Patro, Arpita Biswas, Niloy Ganguly, Krishna P Gummadi, and Abhijnan Chakraborty. 2020. Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms. In *WWW*. 1194–1204.

[33] Maria Predari and Aurélien Esnard. 2016. A k-way greedy graph partitioning with initial fixed vertices for parallel applications. In *IEEE PDP*.

[34] Tieyun Qian, Yang Yang, and Shuo Wang. 2010. Refining graph partitioning for social network clustering. In *WISE*. 77–90.

[35] Chiara Assunta Ricci. 2016. Perceived Social Position and Objective Inequality: Do They Move Together? Evidence from Europe and the United States. *Italian Economic Journal* 2, 3 (2016), 281–303.

[36] Meredith P Richards. 2014. The gerrymandering of school attendance zones and the segregation of public schools: A geospatial analysis. *American Educational Research Journal* 51, 6 (2014).

[37] Debra Satz. 2007. Equality, adequacy, and education for citizenship. *Ethics* 117, 4 (2007), 623–648.

[38] Genevieve Siegel-Hawley. 2013. Educational gerrymandering? Race and attendance boundaries in a demographically changing suburb. *Harvard Educational Review* 83, 4 (2013).

[39] Till Speicher, Hoda Heidari, Nina Grgic-Hlaca, Krishna P. Gummadi, Adish Singla, Adrian Weller, and Muhammad Bilal Zafar. 2018. A Unified Approach to Quantifying Algorithmic Unfairness: Measuring Individual and Group Unfairness via Inequality Indices. In *ACM KDD*. 2239–2248.

[40] Ana-Andreea Stoica, Abhijnan Chakraborty, Palash Dey, and Krishna P Gummadi. 2020. Minimizing Margin of Victory for Fair Political and Educational Districting. In *AAMAS*. 1305–1313.

[41] Adam Swift. 2003. *How not to be a hypocrite: School choice for the morally perplexed parent*. Routledge.

[42] Henri Theil. 1967. *Economics and information theory*. Technical Report.

[43] Pim van 't Hof, Daniël Paulusma, and Gerhard J. Woeginger. 2009. Partitioning graphs into connected parts. *Theor. Comput. Sci.* 410, 47-49 (2009).

[44] William Vickrey. 1961. On the prevention of gerrymandering. *Political Science Quarterly* 76, 1 (1961), 105–110.