

Multi-Objective Optimization for Politically Fair Districting: A Scalable Multilevel Approach

Rahul Swamy

Department of Industrial and Enterprise Systems Engineering, University of Illinois Urbana-Champaign, 104 S Mathews Ave,
Urbana, IL 61801, rahulswa@illinois.edu

Douglas M. King

Department of Industrial and Enterprise Systems Engineering, University of Illinois Urbana-Champaign, 104 S Mathews Ave,
Urbana, IL 61801, dmking@illinois.edu

Sheldon H. Jacobson

Department of Computer Science, University of Illinois Urbana-Champaign, 201 N Goodwin Ave, Urbana, IL 61801,
shj@illinois.edu

Political districting in the United States is a decennial process of redrawing the boundaries of congressional and state legislative districts. The notion of fairness in political districting has been an important topic of subjective debate, with district plans affecting a wide range of stakeholders, including the voters, candidates, and political parties. Even though districting as an optimization problem has been well-studied, existing models primarily rely on non-political fairness measures such as the compactness of districts. This paper presents Mixed Integer Linear Programming (MILP) models for districting with *political fairness* criteria based on fundamental fairness principles such as vote-seat proportionality (efficiency gap), partisan (a)symmetry, and competitiveness. A multilevel algorithm is presented to tackle the computational challenge of solving large practical instances of these MILPs. This algorithm coarsens a large graph input by a series of graph contractions and solves an exact bi-objective problem at the coarsest graph using the ϵ -constraint method. A case study on congressional districting in Wisconsin demonstrates that district plans constituting the approximate Pareto-front are geographically compact, as well as efficient (i.e., proportional), symmetric, or competitive. An algorithmically transparent districting process that incorporates the goals of multiple stakeholders requires a multi-objective approach like the one presented in this study. To promote transparency and facilitate future research, the data, code, and district plans are made publicly available.

Key words: Gerrymandering, Political redistricting, Fairness, Efficiency gap, Partisan asymmetry, Competitiveness, Multilevel algorithm, Pareto optimal

History: This paper was first submitted on February 21, 2019, and was accepted at *Operations Research* on May 2, 2022.

1. Introduction

Once every ten years, state and federal institutions in the United States undertake the task of *districting* (or *partitioning*) a state into political districts. The precise way these districts are drawn affects how the voters are spatially divided, which in turn influences electoral representation and the political landscape of a legislature. Given its potential impact, *gerrymandering*, the act of intentionally manipulating district boundaries to achieve particular political objectives that certain stakeholders may consider “unfair,” has become a common phenomenon. Since the 2010 districting cycle, gerrymandering cases have appeared in the U.S. Supreme Court disputing the district plans of Wisconsin, Maryland, and North Carolina (Royden and Li 2017).

The analysis of fairness in political districting has been a topic of subjective debate, with different stakeholders favoring different notions of what constitutes fairness. For example, even if political parties agree that fair districts should have majorities that are “proportional” to their share of voters, the voters themselves may prefer districts that are “competitive” so that their votes have a greater impact on representation. Additionally, state districting commissions have increasingly incorporated political fairness criteria into their map-drawing processes. In January 2020, the Governor of Wisconsin signed an Executive Order laying out the criteria for the 2021 districting process, which include a requirement that the proposed plans be “free from partisan bias and partisan advantage” (Office of the Governor 2019). Several states (such as Arizona, Colorado, New York, and Washington) encourage the creation of competitive districts (Mann 2005). Hence, there is a need for a fairness-based approach to districting that addresses the concerns of multiple stakeholders.

This paper presents a multi-objective optimization-based districting framework that models several political fairness objectives, extending existing compactness-based models. Three political fairness metrics — *efficiency gap*, *partisan asymmetry*, and *competitiveness* — are considered that explicitly incorporate electoral information. The first two metrics capture fairness among the political parties while, competitiveness captures fairness from a voter’s perspective. The first metric, the efficiency gap, measures the difference in the “wasted” votes for the two major parties. Minimizing the efficiency gap ensures that one party does not waste more voters than the other, thereby limiting partisan gerrymandering. The second notion of fairness is to ensure that if voter preferences shift in the future, the rate at which the two parties gain or lose seats is *symmetric* between the two parties (assuming the shift is uniform across all the districts). This outcome is accomplished by minimizing the extent of asymmetry, called partisan asymmetry. Third, promoting competitiveness of districts is essential to tackle incumbent gerrymandering and promote greater accountability of representatives to their voters. This outcome is achieved by minimizing the margin of victory in all

the districts. The advantages and disadvantages of these metrics have been widely debated in the literature and judicial proceedings. However, the focus of this paper is to provide an optimization-based algorithmic framework for a districting process that seeks to measure fairness using these metrics.

The framework presented in this paper generates (approximate) Pareto-optimal solutions that highlight the trade-off between compactness and each of the three political fairness measures, namely efficiency gap, partisan asymmetry, or competitiveness. Three bi-objective optimization problems are solved using the ϵ -constraint method, each with the compactness of districts as the primary objective and a political fairness metric as the secondary objective. To generate problem instance sizes that are solvable by an exact method, the solution framework first solves a series of graph contraction problems forming multiple *levels* of problem instances of increasing *coarseness*. The coarsest level in this multilevel scheme is used to solve the exact bi-objective problem to produce Pareto-optimal solutions. A case study based on congressional districting in Wisconsin with Census 2010 data is presented. The results demonstrate that district plans can be produced that are not only compact but also efficient, symmetric or competitive. These results imply that algorithms that optimize fairness measures are essential to designing a districting process that integrates the conflicting objectives of multiple stakeholders. In the interest of algorithmic transparency, the data, code and district plans are made publicly available at [Swamy \(2022\)](#).

The rest of the paper is organized as follows. Section 2 reviews the key literature in algorithmic districting and the political debate surrounding fairness in districting. Section 3 introduces the political fairness criteria modeled in the multi-objective problem, the optimization formulations, and the complexity analysis. Section 4 presents a multilevel solution framework that tackles the related computational challenges. Section 5 presents computational results based on a case study on congressional districting in Wisconsin, and Section 6 provides concluding remarks.

2. Related Literature

This paper links two bodies of literature from (i) political sciences and related fields, and (ii) operations research and management sciences. The former covers extensive work on the socio-political requirements and implications of districting, while the latter focuses on efficient methodologies for creating district plans. The literature in these areas is presented in two parts. Section 2.1 discusses how the different notions of political fairness in districting can be quantified, and Section 2.2 explores prior research in algorithmic approaches to create district plans.

2.1. Quantifying Political Fairness

Political fairness can be approached from the partisan perspective as well as the voters' perspective. Metrics that quantify partisan fairness typically consider *proportionality* and *symmetry*. Proportionality is the notion that a party gets more seats (i.e., wins elections in more districts) if more voters favor that party. An alternate notion of partisan fairness is symmetry, the idea that a district plan favors both the parties symmetrically when vote-shares for the two parties fluctuate in hypothetical future elections. On the other hand, incumbent gerrymandering is a commonly observed (bi)partisan effort at drawing districts that intentionally increases the likelihood that certain (incumbent) candidates win (Silverberg 1995). From the voters' perspective, the *competitiveness* of a district is a proxy measure against incumbency. Additionally, competition increases accountability of representatives to their constituents (Jones 2013).

The rest of this section reviews the key literature surrounding the three metrics of fairness considered in this paper. The advantages and disadvantages of these metrics have been widely debated in the literature and in judicial proceedings. While the focus of this paper is to provide an optimization-based algorithmic framework for a districting process that seeks to measure fairness using these three metrics, the choice of these metrics is not an endorsement of the metrics themselves, but merely driven by their extensive popularity and adoption in practical gerrymandering cases. The broader goal of this paper is to highlight that optimization models can enable a districting commission in achieving a fairness outcome they desire.

2.1.1. Efficiency Gap The *efficiency gap* (EG) is a popular measure of partisan fairness. Introduced by Stephanopoulos and McGhee (2015), EG measures the difference in the “wasted votes” between the two major parties. Here, wasted votes in each district are those votes that are either for the losing party or the surplus votes for the winning party in that district. A fair district plan according to EG is mathematically proportional, i.e., as one party's vote-share increases, the number of districts they win increases proportionally (with a slope of two) (Warrington 2018). Note that other authors who consider a stricter view on proportionality, where the number of districts won increases at the *same rate* as the vote-share, may consider EG to possess “double-proportionality” (Bernstein and Duchin 2017), or to be “quasi-proportional” (Chambers et al. 2017). Besides capturing proportionality, Veomett (2018) shows that EG has an additional term that is a function of the difference in voter turnouts in the different districts.

Despite its limitations, EG has emerged as a prominent measure of partisan fairness due to its elegant definition and simplicity in computation, and has even been used by the Supreme Court in Wisconsin's partisan gerrymandering case (Gill v. Whitford 2018). Recently, the state of Missouri passed a resolution that requires that the district plan drawn for the 2021 districting cycle must

have an EG less than 15% (Ballotpedia 2020). In their seminal work introducing the efficiency gap, (Stephanopoulos and McGhee 2015) propose an 8% threshold for EG beyond which a district plan can be considered “unfair” on partisan grounds.

The spotlight on EG has led to debate on its legitimacy as a measure of partisan fairness. Bernstein and Duchin (2017), Chambers et al. (2017) and Veomett (2018) provide valid criticisms on the theoretical efficacy of EG in assessing gerrymandering. Four of the key concerns are discussed here.

1. The computation of EG is deterministic, but it requires the use of uncertain and volatile election results. This factor is particularly critical in the presence of competitive districts, where a small change in voting behavior could “flip” the party holding the majority in those districts, thereby altering the value of EG drastically. To improve the robustness of computing EG, Stephanopoulos and McGhee (2015) suggest using a sensitivity analysis in a second phase of the computation.
2. A plan with a small value of EG (i.e., considered “fair” according to EG) could inadvertently favor districts that are highly non-competitive, or are highly “packed”. Such districts may be considered undesirable by a districting process that seeks to promote competition among the districts.
3. When districts have unequal voter turnouts, Veomett (2018) shows that even a district plan with $EG = 0$ could favor one party disproportionately more than the other. This is possible in a scenario where the turnout in a few districts is disproportionately larger than the turnouts in other districts, and so more votes are wasted in the districts with high turnout. This discrepancy benefits the party that wins in the districts with lesser turnout. To illustrate this effect, Veomett (2018) constructs a district plan with $EG = 0$, where a party with a 50% vote-share in the state disproportionately wins 60% of the districts. In this plan, the turnout in each district they lose is 1.5 times the turnout in each district they win.
4. The mathematical structure of EG exhibits certain undesirable properties when assessing district plans in states with highly polarized vote-shares. In particular, Bernstein and Duchin (2017) derive that in a state where one party has an overall vote-share greater than 80%, every district plan would have an EG value greater than 0.08, which is a threshold for EG suggested by Stephanopoulos and McGhee (2015) beyond which a district plan may be considered to be gerrymandered.

In practice, none of the 50 states have a lean greater than 80% based on the 2016 presidential election. Here, lean refers to the overall vote-share of the majority party in the state. Even for the states with the highest leans, it is possible to achieve EG values less than 0.08. Among the districts with six or more districts, California, Kentucky and Massachusetts have

the highest leans of 66.1%, 65.7% and 64.7%, respectively, while the EG values of their congressional district plans (as of 2016) are 0.004, 0.035 and 0.057, respectively (ESRI 2021). Further, EG correctly identifies partisan gerrymandering in practical instances. Michigan, North Carolina and Pennsylvania have the highest levels of partisan bias when measured using multiple quantitative measures (Royden and Li 2017). Their EG values are 0.132, 0.203 and 0.160, respectively, even though their vote-shares are all under 52% (ESRI 2021). Hence, despite the theoretical limitations of EG, its performance as a measure of partisan gerrymandering is reasonably good in practice.

In their conclusion, Bernstein and Duchin (2017) note that despite the flaws they raise, they “assess that EG can still be a useful component of a courtroom analysis”. In their rebuttal, Stephanopoulos and McGhee (2018) argue that despite these flaws, EG satisfies certain fundamental principles of a partisan fairness metric (such as *efficiency*, *distinctness*, and *breadth of scope*), which justifies its legitimacy as a measure of partisan fairness. See Stephanopoulos and McGhee (2018) for more details about these principles.

Given the extensive academic interest and adaptation of EG as a fairness measure by the courts, this paper considers the problem of minimizing EG and presents an optimization model to achieve this goal. The model presented in Section 3.3 is flexible enough to adapt to alternative ways to compute EG that address at least some of the concerns raised about EG. For example, the issue of uncertainty can be tackled in a robust optimization-based solution method, while the issue of inadvertently creating packed districts can be avoided by setting a hard constraint that enforces a packing limit. The effect of unequal voter turnouts as observed by Veomett (2018) can be minimized by setting a hard constraint that limits the difference in voter turnout between the districts.

2.1.2. Partisan Asymmetry Since proportionality is not a constitutional requirement in the U.S., an alternative school of thought seeks the broader notion of partisan symmetry. The symmetry standard requires “the electoral system to treat similarly-situated parties equally, so that each receives the same fraction of legislative seats for a particular vote percentage as the other party would receive if it had received the same percentage [of the vote]” (Grofman 1983). Several ways to quantify partisan asymmetry have been explored in the literature, with alternatives in the exact numerical function used (Jackman 1994, Duchin 2018). A survey by Grofman (1983) lists eight different ways to measure partisan asymmetry. This paper considers the measure called *Gini score* which computes a vote-seat curve for each party (i.e., a function of how many districts can be won by a party for every hypothetical vote-share it receives in the state) and then measures the area between the curves for the two parties (DeFord et al. 2020).

The symmetry standard as a measure of partisan fairness was put to test in *LULAC v. Perry* (2006), a key U.S. Supreme Court gerrymandering case involving the 2003 Texas’ legislative district

plan. An amicus brief filed by [King et al. \(2006\)](#) proposed that the gerrymandering test be based on partisan asymmetry since “a consensus exists about using the symmetry standard” to evaluate partisan bias. The proposal was discussed, and even though three of the opinions positively evaluated the symmetry standard, the majority opinion had reservations with adopting the test ([Grofman and King 2007](#)). Justice Kennedy expressed concerns with measuring symmetry through the *hypothetical* construction of vote-shares, noting that the court is “wary of adopting a constitutional standard that invalidates a plan based on unfair results that would occur in a hypothetical state of affairs” ([LULAC v. Perry 2006](#)). Justice Stevens (joined by Justice Breyer) endorsed the view that “the proponents of the symmetry standard have provided a helpful (though certainly not talismanic) tool” to measure partisan fairness, though Justice Kennedy (joined by Justices Souter and Ginsburg) noted that “asymmetry alone is not a reliable measure of unconstitutional partisanship” ([LULAC v. Perry 2006](#)). This view motivates the inclusion of partisan asymmetry as one objective in the multi-objective modeling framework for political districting presented in this paper.

A plan that emphasizes symmetry may not emphasize proportionality. Even though proportionality is not a constitutional requirement, some authors consider this lack of synergy between symmetry and proportionality to be a flaw of the symmetry standard itself. For example, a recent technical report by [DeFord et al. \(2020\)](#) finds that when randomly simulating a large number of district plans in Utah, Texas and North Carolina (with Republican vote-shares of 72%, 58% and 53%, respectively), the district plans with small partisan asymmetry values lean disproportionately in favor of Republicans in most cases. Hence, these plans may be viewed as gerrymandered from the proportionality standpoint. In their conclusion, [DeFord et al. \(2020\)](#) note that none of their findings “gives a theoretical reason for rejecting partisan symmetry as a definition of fairness” since a proponent of symmetry could argue that the symmetry standard reinforces “the legitimacy of district-based democracy by reassuring the voting public that the tables can yet turn in the future.” In the case study on Wisconsin (with a Republican vote-share of 48.4%) presented in Section 5, it is found that the majority of the plans align with proportionality, suggesting that there are cases where symmetry does synergize with proportionality.

Given the historic and extensive presence of symmetry in the debates of partisan fairness, this paper considers partisan asymmetry as one of the measures of fairness. While [LULAC v. Perry \(2006\)](#) and [DeFord et al. \(2020\)](#) point to aspects of the symmetry standard that prevent symmetry from being considered a talismanic measure for fairness, the symmetry standard still serves as a strong alternative to viewing fairness through proportionality. Hence, the multi-objective optimization model presented in this paper enables a districting process that seeks symmetry in tandem with other fairness measures, such as the compactness of districts.

2.1.3. Competitiveness While both efficiency gap and partisan asymmetry measure fairness with respect to the political parties, competitiveness captures the extent of advantage for incumbent candidates in future elections. In highly competitive districts, small changes in voter preferences lead to larger changes in the seat-share, and are considered desirable in preventing the creation of *safe seats*, or incumbent gerrymandering (Tapp 2018). Incumbent gerrymandering isolates politicians from being accountable to their constituents, and inhibits communities from receiving “meaningful and fair representation” (Kennedy et al. 2016). Hence, competitiveness captures fairness from the voters’ perspective.

Past studies have analyzed the relationship between competitive districts and other metrics. For example, McCarty et al. (2009) note that in the partisan gerrymandering strategy of “cracking” — where one party’s voters are divided among many districts where they lose close elections — would technically lead to the construction of competitive districts. Other works show unclear results on the trends in polarized (or non-competitive) districts based on congressional and presidential elections (Abramowitz et al. 2006, McCarty et al. 2009). Regardless, the general consensus is that incumbent gerrymandering can be prevented through the creation of competitive districts, thereby creating greater accountability of candidates to their voters (Friedman and Holden 2009). This paper seeks to maximize competitiveness by minimizing the maximum margin of victory in all the districts. The consequence of doing so is that “packing” of voters is minimized.

2.2. Algorithmic Approaches to Districting

This section reviews literature in algorithmic approaches to districting. The literature is broadly categorized based on methodology and intent, namely into simulation algorithms, exact optimization methods, and heuristic methods. The goal of a simulation algorithm is to assess whether a given district plan can be considered gerrymandered or not by comparing it with a large number of simulated district plans. On the other hand, an optimization algorithm creates a single district plan that is optimized with respect to a given fairness measure. The intent of a simulation algorithm is to *compare* or *assess*, while the intent of an optimization approach is to *create* or *optimize*. Optimization methods can further be classified into exact and heuristic methods. Exact optimization methods can achieve a guaranteed optimality, but are not scalable for large practical input sizes. Heuristic methods are computationally efficient in finding “good” district plans with respect to a particular fairness measure, but are not guaranteed to be globally optimal. Sections 2.2.1, 2.2.2, 2.2.3 and 2.2.4 review existing work in simulation algorithms, exact optimization methods, heuristic methods, and multilevel algorithms, respectively.

2.2.1. Simulation Algorithms Among the districting algorithms that explicitly capture political fairness metrics, simulation algorithms form the predominant approach. A simulation algorithm generates a large set of district plans using a suitable sampling method with the intent of studying patterns in political objectives, while also serving as a benchmark for existing district plans. For example, [Chen and Rodden \(2013\)](#) simulate district plans for all U.S. states and conclude that Democratic voters being concentrated in urban areas leads to district plans that strongly favor Republicans. [Fifield et al. \(2015\)](#) provide the first step toward a theoretical framework based on Markov chain Monte-Carlo (MCMC) simulations while incorporating feasibility constraints such as contiguity and population balance. In an MCMC, starting from an initial district plan, a large set of district plans is created by making a series of perturbations. In every step, the conventional method of perturbing a district plan, called “flip”, is to re-assign one unit from its current district to another district ([Fifield et al. 2015](#)). The work by [DeFord et al. \(2019\)](#) provide a method called “recombination” for making a more drastic perturbation to a district plan by merging two neighboring districts and completely redrawing their shared boundary. Recombination has a computational advantage since it enables a faster exploration of the solution space of district plans compared to using a flip operation.

Advances in computational power have enabled technologies that can generate a large number of random district plans. [Liu et al. \(2016\)](#) highlight how high-performance parallel computing can be used to generate and analyze more than a million district plans. Even though these algorithms play a vital role in detecting gerrymandering, it is unclear how these methods can be used in the active creation of a district plan that seeks a particular fairness-based outcome.

2.2.2. Exact Optimization Methods Exact optimization methods have traditionally been used to create district plans using non-political objectives such as compactness and population balance ([Ricca et al. 2013](#)). Political fairness metrics (e.g., the metrics discussed in Section 2.1) have not been explicitly captured in exact models and solution methods. Modeling the political districting problem as a Mixed Integer Program (MIP) with a compactness objective has early roots since the work by [Hess et al. \(1965\)](#). Due to the intractability of the districting problem, finding optimal solutions for large practical instances is a challenge. In particular, contiguity enforcement remains a major challenge. A method used in practice has been to solve the problem without contiguity constraints, and then to use a heuristic to alter a discontinuous solution to satisfy contiguity ([Gentry et al. 2015](#), [Salazar-Aguilar et al. 2011](#)). [Drexler and Haase \(1999\)](#) are the first to provide a set of constraints (albeit exponential in number) that ensure contiguity in districts when using an MIP-based model, while [Shirabe \(2009\)](#) provides a flow-based model for contiguity enforcement in a unit-allocation problem (a special case of districting). The work by [Haase and Müller \(2014\)](#)

extends the formulation to a sales-force districting problem. More recently, [Oehrlein and Haunert \(2017\)](#) provide a cutting-plane approach to contiguity constraints and empirically demonstrate its computational advantage over existing formulations, and [Validi et al. \(2021\)](#) provide a stronger branch-and-cut implementation that solves several previously-unsolved instances of compact districting at the census tract level.

In addition to the traditional branch-and-cut based methods, several alternative exact approaches have also been explored. [Garfinkel and Nemhauser \(1970\)](#) propose a two-stage strategy, where the first stage enumerates all feasible (compact, balanced and contiguous) districts, and the second stage solves a set covering problem to construct an optimal district plan. [Mehrotra et al. \(1998\)](#) present a modified approach: the set covering problem is solved using branch-and-price, and a post-processing stage performs local improvements to the solution.

Despite the extensive research in exact models and methods for districting, political fairness objectives such as efficiency gap, partisan asymmetry and competitiveness have not been captured in MIP models for political districting. It is also unclear if existing computational methods that optimize compactness will be effective when applied to other fairness-based objectives. For example, solving a compactness-seeking model without explicit contiguity constraints typically produces a solution that is either contiguous or “near-contiguous” ([Validi et al. 2021](#)). This reinforcing relationship between contiguity and compactness is advantageous when using a cutting planes approach by [Oehrlein and Haunert \(2017\)](#) or [Validi et al. \(2021\)](#) since “few inequalities are needed to prove optimality” ([Validi et al. 2021](#)). However, when solving a MIP with political fairness-based objectives/constraints, it is unclear whether these approaches will be as computationally effective.

The issue of fairness in districting has been studied in designing game theoretical protocols. In such a game, two or more political parties/candidates strategize to maximize their desired outcomes such as the numbers of districts they win ([Landau and Su 2014](#), [Pegden et al. 2017](#), [De Silva et al. 2018](#), [Benadè et al. 2021](#), [Ludden et al. 2022](#)). The goal of these works is to design a game such that the district plan produced is “fair” according to some measure(s). This goal is fundamentally different from the goal of this paper, which is to optimize a specific fairness objective(s) where that a single decision maker (such as an independent redistricting commission) draws the district plan.

2.2.3. Heuristic Methods Heuristic algorithms are popular in large-scale practical districting applications ([Ricca and Simeone 2008](#)). The heuristics explored in the literature predominantly cater to the goal of optimizing non-political objectives such as compactness and population balance. Some of these methods include simulated annealing ([Browdy 1990](#), [D’Amico et al. 2002](#)), tabu search ([Bozkaya et al. 2003](#)), local search ([Ricca and Simeone 2008](#)), genetic algorithm ([Bacao et al. 2005](#)), and polygonal clustering ([Joshi et al. 2012](#)), among others. Computational geometry-based techniques using Voronoi regions have also been explored ([Ricca et al. 2008](#)).

Some heuristic approaches have explored political objectives. For example, [Chatterjee et al. \(2020\)](#) and [King et al. \(2015\)](#) propose local search methods that optimize the efficiency gap and competitiveness metrics, respectively. These methods improve a given district plan using a sequence of neighborhood perturbations. While these political objectives have been explored in heuristic optimization, they have not been explored in exact optimization algorithms. Additionally, the partisan asymmetry metric has not been incorporated in an optimization setting. Hence, there is an opportunity for research in the integration of optimization methods in political fairness-based districting.

2.2.4. Multilevel Algorithms A multilevel algorithm is a framework for integrating an exact optimization method within a practical heuristic capable of solving large instances. They are considered the most effective algorithms for graph partitioning problems with applications in VLSI design and image processing ([Buluç et al. 2016](#)). A multilevel algorithm creates multiple hierarchical levels of graphs by an iterative coarsening procedure (where adjacent nodes are merged to create a smaller graph) and then plans an optimal district solution from the coarsest graph instance back to the original graph ([Kernighan and Lin 1970](#), [Hendrickson and Leland 1995](#)). Note that the optimal solution found using the coarsest instance when mapped back to the original instance may not be optimal to the original instance, and hence this algorithm is a heuristic method to solve an optimization problem. This framework has also been adopted to other combinatorial optimization problems such as the traveling salesman problem ([Walshaw 2002](#)) and graph-drawing ([Walshaw 2000](#)). While districting problems share many features with graph partitioning problems, they also differ from graph partitioning problems (whose objective typically is to minimize the total weight of the edges cut by the partition) in that districting involves more constraints, including contiguity and tighter population balance constraints, in addition to a broader set of fairness objectives. Recent work by [Magleby and Mosesson \(2018\)](#) presents a multilevel algorithm to generate random district plans where the coarsening stage randomly chooses adjacent units to be merged. This paper investigates how a multilevel algorithm can be adapted to solve deterministic large-scale districting problems, especially when considering multiple objectives.

3. Multi-objective Politically Fair Districting

This paper considers districting as an optimization problem with four objectives: compactness, efficiency gap, partisan asymmetry and the competitiveness of districts. The general multi-objective districting model is introduced in Section 3.1. The background, formulation and analysis of each of the single-objective optimization problems are presented in Sections 3.2—3.5.

3.1. The Districting Problem

A mathematical model for districting for political representation is now formalized. A *Districting Problem* (DP) partitions a geographical region into a finite number of districts where each district satisfies certain requirements. The input to DP consists of a set of discrete geographical units V (or simply *units*), their adjacency information, and the number of districts K . Each unit represents an area of administrative significance such as a county, census tract, or census block. Two units $i, j \in V$ are said to be *adjacent* if they share a boundary curve which contains a continuous segment (as opposed to being a set of singleton points). An *adjacency graph* $G = (V, E)$ is an undirected graph where the set of nodes V represents the set of units, and an edge (i, j) exists in the edge set E if and only if units i and j are adjacent. Note that an edge (i, j) here denotes an un-ordered pair of nodes, i.e., (i, j) is equivalent to (j, i) . For each unit $i \in V$, let $N(i) := \{j \in V : (i, j) \in E\}$ denote the set of *neighbors* of i . Additionally, each unit $i \in V$ contains a certain population of residents, denoted by $p_i \geq 0$. A *district* is a subset of V that satisfies certain conditions, and a *district plan* is a partition of V into K districts, i.e., $z : [K] \rightarrow 2^{|V|}$ denotes a district plan where for each $k \in [K]$, $z(k) \subset V$ is a district. A district plan z is said to be feasible for DP if for each district $k \in [K]$, the subgraph induced by $z(k)$ in G is connected, and the population $\sum_{i \in z(k)} p_i$ belongs to $[lb, ub]$, where lb and ub are the lower and upper bounds on the population required in each district, respectively. Typically, these bounds are formulated as $[\bar{P}(1 - \tau), \bar{P}(1 + \tau)]$, where $\bar{P} := (\sum_{i \in V} p_i)/K$ is an ideal population in each district, and $\tau \geq 0$ is the population deviation tolerance. Given an objective function $\phi : [K] \times 2^{|V|} \rightarrow \mathbb{R}$, DP finds a feasible district plan that minimizes ϕ .

The *Politically Fair Districting Problem* (PFDP) is a multi-objective generalization of DP, with each objective representing a particular notion of fairness. In addition to the inputs for DP, for district plans to be evaluated on their political fairness, PFDP also requires information on the political leanings of voters. Consider a two-party electoral system, where parties A and B have a distribution of voters spread among the units in V . Let $p_i^A, p_i^B \geq 0$ (with $p_i^A + p_i^B \leq p_i$) be the number of voters for parties A and B respectively residing in unit $i \in V$. For a given district plan z , in each district $k \in [K]$, let $P_k^r(z) := \sum_{i \in z(k)} p_i^r$ be the number of voters for party $r \in \{A, B\}$. In a plurality-based election (e.g., in U.S. congressional elections), a party A (B) is said to *win* (*lose*) district k if $P_k^A(z) > P_k^B(z)$, i.e., if A has more voters in district k than party B . Note that ties (i.e., $P_k^A(z) = P_k^B(z)$ for some $k \in [K]$) rarely occur in practice when the number of voters is large and the units are discrete with a heterogeneous distribution of voters. For a given district plan, based on the political leanings of the voters in each district and the districts won/lost by the two parties, one can assess how politically *fair* the district plan is. There are several ways to assess political fairness (e.g., the efficiency gap and partisan asymmetry) as discussed in Section

2.1. Let $\{\phi_q\}_{q=1}^Q$ be a set of $Q \geq 1$ political fairness objective functions. Here, for each $q \in [Q]$, $\phi_q : [K] \times 2^{|V|} \rightarrow \mathbb{R}$ represents a real-valued function that measures the fairness of a district plan z . PFDP is a Q -objective optimization problem that minimizes the set of objective functions $\{\phi_q\}_{q=1}^Q$ within the solution space of DP.

This section presents a mathematical model for a PFDP with $Q = 4$ objectives, namely compactness, efficiency gap, partisan asymmetry, and competitiveness. Computational experiments applying this framework to instances with select pairs of these objectives (i.e., $Q = 2$) are presented in Section 5.

PFDP can be formulated as follows. The primary decision variables that assign units to districts are given by,

$$\begin{aligned} x_{ij} &= \begin{cases} 1, & \text{if } j (\neq i) \in V \text{ is assigned to the district with center } i \in V. \\ 0, & \text{otherwise.} \end{cases} \\ x_{ii} &= \begin{cases} 1, & \text{if unit } i \in V \text{ is chosen as a district center.} \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

Here, a *district center* is a special unit in a district that serves as the source of flow within its district in the flow-based contiguity constraints discussed below. Originally introduced by Shirabe (2009), flow-based constraints for contiguity have recently become popular for contiguity enforcement in MIPs. Originating from every unit $i \in V$, let $f_{ijv} \geq 0$ be a decision variable that denotes the amount of flow from unit $j \in V$ to its neighbor $v \in N(j)$, whose value is positive only if i is a district center, and units v and j are both assigned to the district with center i . Adapting from the model presented in seminal work by Hess et al. (1965), PFDP can be formulated as,

$$\text{(PFDP) Minimize } \{\phi_1(x), \phi_2(x), \dots, \phi_Q(x)\} \quad (1)$$

$$\text{subject to } (1 - \tau)\bar{P} x_{ii} \leq \sum_{j \in V} x_{ij} p_j \leq (1 + \tau)\bar{P} x_{ii} \quad \forall i \in V, \quad (2)$$

$$\sum_{i \in V} x_{ii} = K, \quad (3)$$

$$\sum_{i \in V} x_{ij} = 1 \quad \forall j \in V, \quad (4)$$

$$x_{ij} \leq x_{ii} \quad \forall i, j \in V, \quad (5)$$

$$x_{ij} + \sum_{v \in N(j)} (f_{ijv} - f_{ivj}) = 0 \quad \forall i, (\neq)j \in V, \quad (6)$$

$$x_{ii} + \sum_{v \in N(i)} (f_{iiv} - f_{ivi}) - \sum_{v \in V} x_{iv} = 0 \quad \forall i \in V, \quad (7)$$

$$|V| x_{ij} - \sum_{v \in N(j)} f_{ivj} \geq 0 \quad \forall i, j \in V, \quad (8)$$

$$x_{ij} \in \{0, 1\}, f_{ivj} \geq 0 \quad \forall i, j \in V, v \in N(j). \quad (9)$$

The tuple of objectives in (1) are minimized in a PFDP. Constraints (2) define the upper and lower bounds on district populations. Constraint (3) ensures that there are exactly K districts. Constraints (4) assign a unique district to each unit. Constraints (5) ensure that a unit can be assigned to a district with unit i as its center only if i is district center, i.e, $x_{ii} = 1$. Constraints (6)-(8) have been adapted from the flow-conservation constraints introduced in Haase and Müller (2014) for a sales force deployment problem. Here, flow originates from a district center i , and every unit j assigned to the district with center i acts as a sink for one unit of flow. These constraints ensure that overall flow is conserved if and only if all the districts are contiguous. Note that Oehrlein and Haunert (2017) provide a more elegant way to express the flow constraints, where constraints (7) are omitted. However, the computational study in this paper includes constraints (7) since preliminary investigations revealed that the MIP can be solved faster with their inclusion than without them. Constraints (9) define the variable domains. This formulation can be improved with the inclusion of the constraints $\sum_{v \in N(i)} f_{ivi} = 0$ for all $i \in V$. Additionally, an alternative formulation with an x_{ik} variable that assigns unit $i \in V$ to district $k \in [K]$ could reduce the number of variables, though preliminary analysis indicates that this formulation has weaker relaxation bounds (compared to using the x_{ij} variable), which lead to a slower convergence.

3.2. Compact Districting Problem

Compactness evaluates a district's geometrical shape and can be quantified in several ways (Young 1988). A popular metric is the *moment of inertia*, measured by the squared sum of distances from the units to their geographical district centers, weighted by the unit populations. Here, a district center is a unit with the smallest sum of weighted squared distances to the rest of the units in its district, and whose location is determined within the optimization algorithm. Note that this district center still serves as the source of flow in constraints (6)-(8). Moment of inertia has been used in compact districting since early work by Hess et al. (1965). Let $d_{ij} \geq 0$ be the given distances between every pair of units $i, j \in V$. The *Compact Districting Problem*, or Compact-DP, finds a district plan that minimizes this compactness objective, and can be formulated as,

$$\begin{aligned} \text{(Compact-DP) Minimize } \phi_{comp}(x) &:= \sum_{i,j \in V} p_j d_{ij}^2 x_{ij} \\ &\text{subject to (2) -- (9).} \end{aligned} \tag{10}$$

The compactness objective ϕ_{comp} in (10) minimizes the moment-of-inertia function. Compact-DP is a well studied classical problem that generalizes the NP-complete p -median problem (Megiddo and Supowit 1984).

Even though compact districts are generally perceived to be *fair*, optimally compact districts could be considered unfair upon consideration of voter information. Consider an example with




Figure 1 (a) Optimal district plan for Compact-DP (b) Districts with proportional seat-share
Example with four districts on a grid graph suggesting that an optimally compact solution may not yield a proportional seat-share to the two parties.

units forming a 10×10 grid (Figure 1), where each 1×1 square is a unit. Assume that the shaded units are densely populated with party A voters, whereas the unshaded units are populated with party B voters and non-voters, i.e., the tuple (p_i, p_i^A, p_i^B) for each unit i is: $(40, 40, 0)$ for the shaded units and $(40, 0, 10)$ for the unshaded units. Here, the total number of voters for A and B are 1000 and 750, respectively. This example illustrates a scenario of spatial polarization of voters in a real-life setting where different areas could have different concentrations of party support (Chen and Rodden 2013). Consider the four districts defined by the four equal-sized quadrants of the grid in Figure 1a. Clearly, these districts satisfy the contiguity and population balance constraints, and are also optimally compact. However, the seat-share for B is 0.75 (i.e., B wins three out of the four districts), while their vote-share is 0.43 (i.e., B receives 750 out of the 1,750 voters). From a proportionality stand-point, this district plan is not considered fair since B wins three times the number of districts as A does, even though B 's vote-share is less than A 's. In contrast, in the district plan in Figure 1b, the seat-share for B is 0.5, which is closer to B 's vote-share than that of the previous district plan despite having less compact districts. Hence, if proportionality is sought in districting, there is a need to explicitly consider voter information in an optimization model in addition to the compactness objective.

3.3. Efficient Districting Problem

In a scenario where one party (A) gerrymanders (i.e., intentionally draws the districts to gain artificial advantage in the district elections), the phenomenon of *packing* and *cracking* is observed. Packing occurs when a large number of the other party's (B) voters are assigned to a small number of districts in which party B wins by large margins, while cracking occurs when the remaining

party B voters are *diluted* among the rest of the districts in which party A holds a majority. Even though this phenomenon has been widely discussed in decades of gerrymandering research, it was not until recently that a metric to quantify packing and cracking has been formulated and adopted in practice. The *efficiency gap* introduced by Stephanopoulos and McGhee (2015) captures both packing and cracking by counting the difference in wasted votes between the two parties. A vote in a certain district is said to be *wasted* if it is either a surplus vote for the winning party, or a vote to the losing party. In each district $k \in [K]$, let $w_k^A(z)$ be the wasted votes for party A given by,

$$w_k^A(z) = \begin{cases} P_k^A(z) - \frac{P_k^A(z) + P_k^B(z)}{2}, & \text{if } P_k^A(z) > P_k^B(z). \\ P_k^B(z), & \text{if } P_k^A(z) < P_k^B(z). \end{cases}$$

Exactly half the votes in every election are wasted. However, if one party disproportionately wastes more votes than the other, it is considered to be an indication of partisan gerrymandering in practical districting scenarios (Stephanopoulos and McGhee 2015), unless one party achieves an overwhelming majority of the total votes (Chambers et al. 2017). Among the wasted votes, the surplus votes for the winning party quantify packing, while all the votes for the losing party quantify cracking. Note that when there is a tie in a district k (i.e., when $P_k^A(z) = P_k^B(z)$), the notion of wasted votes is not defined in Stephanopoulos and McGhee (2015) since there is no clear winner in that district. Even though ties rarely occur in practice as noted in Section 3.1, the model discussed later in this section explicitly deals with ties. Assuming there are no ties, the difference in wasted votes between the two parties in each district k is given by,

$$w_k^A(z) - w_k^B(z) = \begin{cases} \frac{P_k^A(z) - 3P_k^B(z)}{2}, & \text{if } P_k^A(z) > P_k^B(z). \\ \frac{3P_k^A(z) - P_k^B(z)}{2}, & \text{if } P_k^A(z) < P_k^B(z). \end{cases} \quad (11)$$

Party A wastes more votes than party B when $w_k^A(z) - w_k^B(z)$ is positive; vice versa if it is negative. For a given district plan z , the *efficiency gap* (ϕ_{EG}) measures the magnitude of the net difference in wasted votes between two political parties A and B across all the districts normalized by the total number of voters, defined as,

$$\phi_{EG}(z) := \frac{\left| \sum_{k=1}^K (w_k^A(z) - w_k^B(z)) \right|}{\sum_{k=1}^K (P_k^A(z) + P_k^B(z))}. \quad (12)$$

For any given district plan z , $\phi_{EG}(z)$ is a value between 0 and 0.5, where a value of 0 implies that both parties have wasted an equal number of votes. In this paper, a district plan is said to be *perfectly efficient* if both the parties waste the same number of votes, i.e., with efficiency gap equal to zero. The *Efficient Districting Problem*, or Efficient-DP, is a districting problem whose objective is to minimize the efficiency gap, $\phi_{EG}(z)$. Recent work by Chatterjee et al. (2020) shows that minimizing the efficiency gap with population balance constraints is strongly NP-complete

when the instance given by the adjacency graph is planar (which is typically the case for practical districting problems).

This section presents an optimization model for Efficient-DP. In this model, a key assumption is that in the rare case of a tie in a district k (i.e., when $P_k^A(z) = P_k^B(z)$), the model arbitrarily chooses a winner in that district such that this choice minimizes the efficiency gap. Efficient-DP can be formulated as a mixed integer program (MIP). Let $a_i \in \mathbb{R}$ for unit $i \in V$ be the difference in wasted votes between parties A and B in a district with center i ; 0 if i is not a district center. Additional decision variables supporting this formulation are given by,

$$y_i^A = \begin{cases} 1, & \text{if } i \in V \text{ is a district center and party } A \text{ wins that district.} \\ 0, & \text{otherwise.} \end{cases}$$

$$v_{ij}^A = \begin{cases} 1, & \text{if party } A \text{ wins the district with center } i, \text{ and } j \in V \text{ is assigned to that district.} \\ 0, & \text{otherwise.} \end{cases}$$

Efficient-DP can be formulated as,

$$\text{(Efficient-DP) Minimize } \phi_{EG}(x) = \frac{|\sum_{i \in V} a_i|}{\sum_{i \in V} (p_i^A + p_i^B)} \quad (13)$$

subject to (2) – (9),

$$-M \leq \sum_{j \in V} (p_j^A - p_j^B) x_{ij} - M y_i^A \leq 0 \quad \forall i \in V, \quad (14)$$

$$v_{ij}^A \leq x_{ij} \quad \forall i, j \in V, \quad (15)$$

$$v_{ij}^A \leq y_i^A \quad \forall i, j \in V, \quad (16)$$

$$v_{ij}^A \geq x_{ij} + y_i^A - x_{ii} \quad \forall i, j \in V, \quad (17)$$

$$y_i^A \leq x_{ii} \quad \forall i \in V, \quad (18)$$

$$a_i = \sum_{j \in V} \left(\frac{3p_j^A - p_j^B}{2} \right) x_{ij} - \sum_{j \in V} (p_j^A + p_j^B) v_{ij}^A \quad \forall i \in V, \quad (19)$$

$$v_{ij}^A, y_i^A \in \{0, 1\}, \quad a_i \in \mathbb{R} \quad \forall i, j \in V. \quad (20)$$

The efficiency gap objective ϕ_{EG} in (13) minimizes the absolute value of the net difference in wasted votes normalized by the total number of voters. The absolute value function in (13) is linearized in constraints (35) in Section 4.2. Note that the denominator in (13) is a constant which can be removed in order to make the model more numerically stable for an MIP solver. Constraints (14) define $\{y_i^A\}_{i \in V}$ in relation to $\{x_{ij}\}_{i, j \in V}$, where M can be any value greater than $\bar{P}(1 + \tau)$. Here, the first inequality ensures that $y_i^A = 1 \Rightarrow \sum_{j \in V} (p_j^A - p_j^B) x_{ij} \geq 0$. The second inequality ensures that $\sum_{j \in V} (p_j^A - p_j^B) x_{ij} \geq 0 \Rightarrow y_i^A = 1$. Constraints (15)–(17) linearize the quadratic constraints $v_{ij}^A = x_{ij} y_i^A$ that ensure that $v_{ij}^A = 1$ if and only if $x_{ij} = y_i^A = 1$ using standard linearization techniques (Adams and Sherali 1990). Constraints (18) ensure that $y_i^A = 1$ only if unit i is a district center. Constraints (19) define the difference in wasted votes $\{a_i\}_{i \in V}$. If $i \in V$ is a district center (i.e.,

$x_{ii} = 1$), let $V^i \subset V$ be the set of units assigned to that district (i.e., $V^i := \{j \in V : x_{ij} = 1\}$). For each $i \in V$, constraint (19) ensures that $a_i = \sum_{j \in V^i} (p_j^A - 3p_j^B)/2$ if A wins the district with center i , or $a_i = \sum_{j \in V^i} (3p_j^A - p_j^B)/2$ if B wins that district. Hence, in accordance with (11), a_i stores the difference in wasted votes between the two parties in the district with center i . If i is not a district center, then a_i is 0. Constraints (20) define the variable domains.

The way in which the model picks the winner in each district with center $i \in V$ is detailed below. First consider the case when $\sum_{j \in V} (p_j^A - p_j^B)x_{ij} \neq 0$. Then, i is a district center and there is no tie in that district. In this case, there are two possibilities: (i) $\sum_{j \in V} (p_j^A - p_j^B)x_{ij} > 0$ (where party A wins that district) or (ii) $\sum_{j \in V} (p_j^A - p_j^B)x_{ij} < 0$ (where party B wins that district). Then, constraint (14) ensures that $y_i^A = 1$ in case (i) and $y_i^A = 0$ in case (ii), satisfying the definition of y_i^A . Second, consider the case when $\sum_{j \in V} (p_j^A - p_j^B)x_{ij} = 0$. Then, either i is not a district center or there is a tie in the district with center i . If i is not a district center, $x_{ii} = 0$ and hence constraint (18) ensures that $y_i^A = 0$. In the case when there is a tie, according to constraint (14), y_i^A can either be 0 or 1. The rest of the constraints (15)-(19) define the wasted votes a_i in district i . Hence, the optimization model selects the value of y_i^A appropriately such that the wasted votes are attributed to that party that minimizes the EG objective. In particular, when the model sets y_i^A to be 1(0), party A (B) wins that district and all of B (A)'s voters are set to be wasted. The end result is that the solver will break ties in a way that minimizes the overall efficiency gap.

3.4. Symmetric Districting Problem

A well-established notion of partisan fairness in a two-party system is *partisan asymmetry*, which captures the extent of advantage one party has over the other when voter preferences fluctuate. To understand the idea behind partisan asymmetry, a *vote-seat curve* is first defined. For a given district plan, let $\omega_k \in [0, 1]$ be the district vote-share in district $k \in [K]$ for an arbitrary party A . Let $\bar{\omega} := (\sum_{k \in [K]} \omega_k)/K$ be the *average vote-share* for A across all the districts. For every given average vote-share $\bar{\omega} \in [0, 1]$, the *vote-seat curve* is defined to be the corresponding fraction of districts won by A , denoted by $s(\bar{\omega}) \in [0, 1]$. To compute the vote-seat curve for a given district plan, starting from A 's actual vote-share (based on past elections), its vote-shares across all the districts are shifted by adding and subtracting votes. A key assumption here is that the shifting is done *uniformly* across all the districts, as discussed in Duchin (2018) and Katz et al. (2020). Then, the vote-seat curve for party A is a step function that gains a step when A wins a district. Party B 's vote-seat curve is given by $1 - s(1 - \bar{\omega})$. Note that prior work in defining vote-seat curves consider each party's *overall vote-share*, which is the fraction of votes for a party across the entire state, instead of the average vote-share. However, it was empirically observed that the average vote-share is close to the overall vote-share in the district plans obtained from real-world

instances. For example, in the 2010 congressional district plan for Wisconsin, Democrats (D) have an average vote-share of 0.513, whereas their overall vote-share in the state is 0.516. Across the U.S. states with at least two congressional districts in each, D's average vote-share in each state's 2010 congressional district plan and D's overall vote-share in that state differ by 0.004 on average. Additionally, considering the average vote-share has computational advantages since it is a linear function of the district vote-shares.

Partisan asymmetry quantifies the functional difference between the vote-seat curves for the two parties, given by $|s(\bar{w}) - (1 - s(1 - \bar{w}))|$. As proposed by Grofman (1983) as “Measure 7”, this paper measures partisan asymmetry for a district plan z by the area between the two curves. This measure is also referred to as the *partisan Gini score* in DeFord et al. (2020) given by,

$$\phi_{PA}(z) := \frac{1}{K} \int_0^1 |s(\bar{w}) - (1 - s(1 - \bar{w}))| d\bar{w}. \quad (21)$$

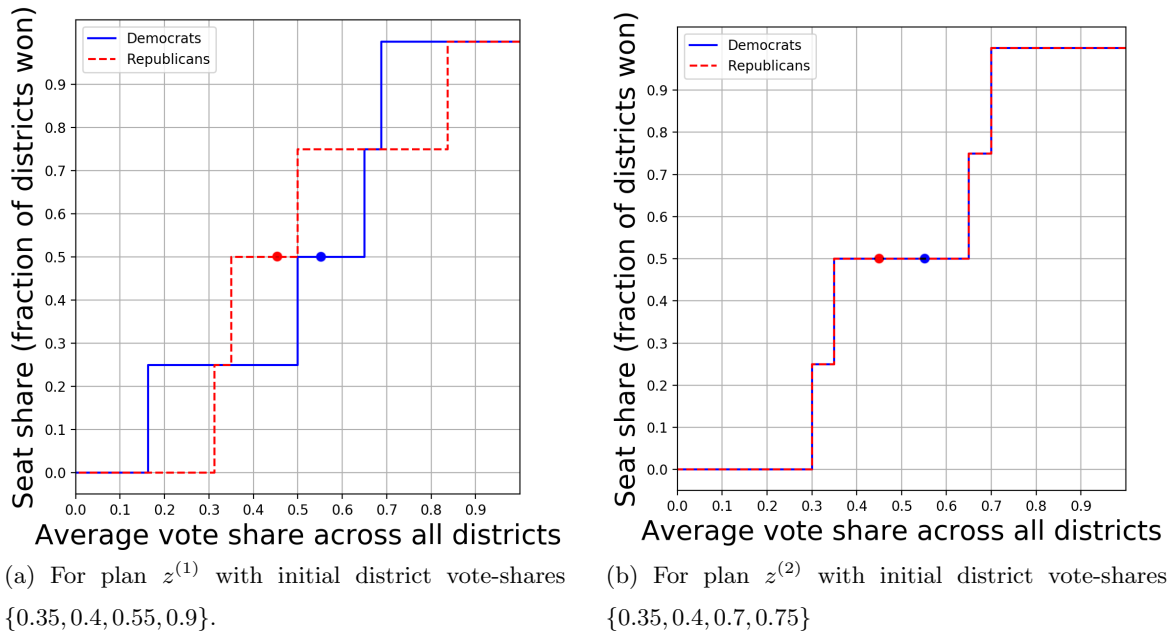


Figure 2 The vote-seat curves for an example with four districts and Democrats' average vote-share of 0.55.

To illustrate the computation of partisan asymmetry, consider an example with four districts and two parties — the Democrats (D) and the Republicans (R). For a given district plan $z^{(1)}$, let D's actual vote-shares in the four districts be $\{0.35, 0.4, 0.55, 0.9\}$, where D wins two districts. Here, D's actual average vote-share is $\bar{w} = 0.55$. To start computing D's vote-seat curve, their district vote-shares are incremented uniformly across all the districts until D *just* wins one more district. Doing so brings D's district vote-shares to $\{0.45, 0.5, 0.65, 1\}$ and the corresponding average vote-share to $\bar{w} = 0.65$. Similarly, as votes are uniformly first incremented (and then decremented) across

the districts until D wins (loses) all the districts, the points at which D just wins (and loses) seats are recorded to obtain its vote-seat curve $\{s(\bar{\omega}) : \bar{\omega} \in [0, 1]\}$. Note that when a district vote-share reaches 1 (0), votes in that district are no longer incremented (decremented). Figure 2a depicts the vote-seat curves for D and R. The area between the two curves gives plan $z^{(1)}$ a partisan asymmetry of $\phi_{PA}(z^{(1)}) = 0.04$. Consider another district plan $z^{(2)}$, where D's actual district vote-shares in the four districts are $\{0.35, 0.4, 0.7, 0.75\}$, where D wins two districts. The vote-seat curves for D and R are depicted in Figure 2b and this plan has a partisan asymmetry of $\phi_{PA}(z^{(2)}) = 0$.

The asymmetry in plan $z^{(1)}$ can be interpreted as follows. If R's average vote-share increases (from its actual value of 0.45) to 0.55, R wins three districts; whereas D wins only two for the same average vote-share. In general, in the regions in the vote-seat space where one party's seat-share is higher than the other party's, the difference in the seats is considered to be asymmetry in outcomes between the two parties. This asymmetry is intrinsic to how the voters are distributed in the districts, and is measured by the area between the curves. On the other hand, the vote-seat curves for plan $z^{(2)}$ indicate that both parties win the same number of districts for the same hypothetical vote-shares and is hence perfectly symmetric.

The *Symmetric Districting Problem*, or Symmetric-DP, is a districting problem with two parties A and B whose objective minimizes the partisan asymmetry ϕ_{PA} of a district plan. The following variables are introduced to formalize Symmetric-DP as a Mixed Integer Non-linear Program.

ω_i = district vote-share for party A in the district with center $i \in V$; 0 if i is not a district center.

α_k = the k^{th} largest district vote-share for party A , for $k \in [K]$

μ_{km} = shifted district vote-share in district $m \in [K]$ when party A wins $k \in [K]$ districts

$\hat{\omega}_k$ = minimum average vote-share for A to win $k \in [K]$ districts.

Note that since the two vote-seat curves are step functions, the objective ϕ_{PA} can be expressed as a finite sum over the regions defined by the K breakpoints of each curve. Here, a breakpoint in the step function corresponds to the average vote-share when a party *just* wins one more district. Further, for all $k \in [K]$, let $k \max\{\omega_i\}_{i \in V}$ denote a function that returns the k^{th} largest value from the set of values $\{\omega_i\}_{i \in V}$. Then, Symmetric-DP can be formalized as,

$$\text{(Symmetric-DP) Minimize } \phi_{PA}(x) = \frac{1}{K} \sum_{k=1}^K |\hat{\omega}_k - (1 - \hat{\omega}_{K-k+1})| \quad (22)$$

subject to (2) – (9),

$$\omega_i = \begin{cases} \frac{\sum_{j \in V} p_j^A x_{ij}}{\sum_{j \in V} (p_j^A + p_j^B) x_{ij}}, & \text{if } x_{ii} = 1 \\ 0, & \text{otherwise} \end{cases} \quad \forall i \in V, \quad (23)$$

$$\alpha_k = k \max\{\omega_i\}_{i \in V} \quad \forall k \in [K], \quad (24)$$

$$\mu_{km} = \begin{cases} 0, & \text{if } \alpha_m + \frac{1}{2} - \alpha_k \leq 0 \\ \alpha_m + \frac{1}{2} - \alpha_k, & \text{if } \alpha_m + \frac{1}{2} - \alpha_k \in (0, 1) \\ 1, & \text{if } \alpha_m + \frac{1}{2} - \alpha_k \geq 1 \end{cases} \quad \forall k, m \in [K], \quad (25)$$

$$\hat{\omega}_k = \frac{1}{K} \sum_{m=1}^K \mu_{km} \quad \forall k \in [K], \quad (26)$$

$$\mu_{km}, \omega_i, \alpha_k, \hat{\omega}_k \in [0, 1] \quad \forall i \in V, k, m \in [K]. \quad (27)$$

The objective in (22) minimizes the partisan asymmetry function. Each constraint in (23) defines ω_i to be A 's district vote-share if $i \in V$ is a district center; 0 otherwise. Constraints (24) define the k^{th} largest vote-share α_k for each $k \in [K]$, where the $k \max$ function returns the k -th largest value of a given set. Each constraint in (25) defines the shifted district vote-share μ_{km} in district m when party A wins k districts. Here, the district vote-share α_m in each district m is altered by a value of $\alpha_k - \frac{1}{2}$, which is the fraction of voters needed to *just* win k districts. This piecewise linear relationship ensures that the shifted district vote-share $\left(\alpha_m - (\alpha_k - \frac{1}{2})\right)$ does not exceed the $[0, 1]$ interval. Each constraint in (26) defines the minimum average vote-share $\hat{\omega}_k$ needed for A to win $k \in [K]$ districts. Constraints (27) restrict the variables to be non-negative. The non-linear constraints (23)-(25) can be linearized, as described in Proposition 1.

PROPOSITION 1. *Symmetric-DP can be formulated as a Mixed Integer Linear Program (MILP).*

The proof given in Section EC.2 of the online appendix linearizes the constraints (23)-(25) by the introduction of additional variables and constraints. The proof also presents more details on the intuition behind constraints (23)-(25). Note that the objective function in (22) can be linearized by standard linearization techniques and is implemented in constraint form as discussed in Section 4.2.

3.5. Competitive Districting Problem

Incumbent gerrymandering (or bipartisan gerrymandering) is the act of intentionally protecting incumbent candidates by drawing districts that can be considered “safe” for them, i.e., highly packed in the incumbents’ favor (Mann 2005). Avoiding the creation of safe districts, or in turn creating *competitive* districts, is vital to ensure that the representatives are responsive to their constituents. In the U.S., Arizona and Colorado have explicitly encoded in their Constitutions that districts be competitive “to the extent practicable” without significant detriment to other districting goals, while Washington and New York require that competition is “encouraged” or at the least “not discouraged”; eleven other states prohibit the creation of districts that intentionally favor incumbent candidates (National Conference of State Legislatures 2020).

In a two-party system, the competitiveness of a district is quantified by the *margin of victory*, the fraction of voters by which the winning party leads the other party (McDonald 2006). The smaller the margin, the more competitive the district is. Given a district plan, the competitiveness of the plan can also be measured by the number of “competitive districts”, where a district is competitive if its margin of victory is below a certain threshold (McDonald 2006). However, when an optimization algorithm maximizes the number of competitive districts, it frequently creates district plans that “pack” the non-competitive districts. For example, an optimal solution with six competitive districts out of eight often has the other two districts be highly packed. Hence, to ensure that all the districts are as competitive as possible, this paper measures competitiveness by minimizing the maximum margin across all the districts. For a district plan z , the competitiveness objective, $\phi_{cmpttv}(z)$, is defined to be the maximum margin among all the districts, i.e.,

$$\phi_{cmpttv}(z) := \max_{k \in [K]} \frac{|P_k^A(z) - P_k^B(z)|}{P_k^A(z) + P_k^B(z)}.$$

The *Competitive Districting Problem*, or Competitive-DP, is a districting problem whose objective is to minimize the competitiveness function $\phi_{cmpttv}(z)$. Minimizing the maximum margin encourages the creation of competitive districts by associating the quality of a district plan by its least-competitive district.

Even though the computational hardness of Competitive-DP is unclear, this paper shows hardness results for a closely related problem. The *Aggregate Competitive Districting Problem* (ACDP) finds a feasible district plan z that minimizes the maximum absolute difference in the two parties’ voters given by $\max_{k \in [K]} |P_k^A(z) - P_k^B(z)|$. ACDP is different from Competitive-DP in that ACDP’s objective does not contain the normalization factor in Competitive-DP’s objective given by the total number of voters in that district. Despite this difference, the ACDP’s objective captures the essence of creating competitive districts. It can be shown that ACDP is NP-complete if $K = 2$ and is strongly NP-complete if $K \geq 3$ with or without the population balance constraints (2).

THEOREM 1. *For a general graph $G = (V, E)$, unit populations $\{p_i, p_i^A, p_i^B\}_{i \in V}$ and number of districts $K \geq 2$, the Aggregate Competitive Districting Problem (ACDP) with or without the population balance constraints is NP-complete if $K = 2$ and is strongly NP-complete if $K \geq 3$.*

The proof (presented in Section EC.1 of the e-companion) proceeds by showing a reduction from the NP-complete 2-partitioning problem for $K = 2$, and from the strongly NP-complete 3-partitioning problem for $K \geq 3$ (Garey and Johnson 1978).

Competitive-DP can be formulated as the following MIP by linearizing the fractional and absolute value functions in the objective. Let $h \geq 0$ be the maximum margin across all the districts. For

every pair of units $i, j \in V$, let $b_{ij} := h x_{ij}$ be an artificial quadratic variable introduced to define h in the MIP where b_{ij} takes the value h when x_{ij} is 1; 0 otherwise. Competitive-DP is given by,

$$(\text{Competitive-DP}) \text{ Minimize } \phi_{cmpttv}(x) = h \quad (28)$$

subject to (2) – (9),

$$\sum_{j \in V} (p_j^A + p_j^B) b_{ij} \geq \sum_{j \in V} (p_j^A - p_j^B) x_{ij} \quad \forall i \in V, \quad (29)$$

$$\sum_{j \in V} (p_j^A + p_j^B) b_{ij} \geq - \sum_{j \in V} (p_j^A - p_j^B) x_{ij} \quad \forall i \in V, \quad (30)$$

$$b_{ij} \leq x_{ij} \quad \forall i, j \in V, \quad (31)$$

$$b_{ij} \leq h \quad \forall i, j \in V, \quad (32)$$

$$b_{ij} \geq h - (1 - x_{ij}) \quad \forall i, j \in V, \quad (33)$$

$$h, b_{ij} \geq 0 \quad \forall i, j \in V. \quad (34)$$

The competitiveness objective ϕ_{cmpttv} in (28) minimizes the maximum margin h . Constraints (29)-(30) define $\{b_{ij}\}_{i,j \in V}$ in relation to the margins in each district. Constraints (31)-(33) linearize the quadratic variables $\{b_{ij}\}_{i,j \in V}$, where $b_{ij} = h x_{ij}$ for all $i, j \in V$ using standard linearization techniques (Adams and Sherali 1990). Constraints (34) define the non-negativity of the variables.

Note that optimizing for the competitiveness objective may produce results that are not fair with respect to a *partisan* fairness metric such as the efficiency gap or partisan asymmetry. This is permissible since the idea behind ensuring competitive districts is to prevent incumbent gerrymandering, which is independent of the political affiliation of the winners. Consider an example with district vote-shares in three equipopulous districts exhibiting district vote-shares of $\{0.51, 0.53, .55\}$ for party A . Although this solution is highly inequitable with respect to the political parties (with an efficiency gap of 44%), all three districts are competitive within a 10% margin and can be considered fair in terms of their responsiveness to the voters.

4. Solution Method

This section presents a solution approach for solving bi-objective PFDPs. Here, each of the political fairness objectives (i.e., efficiency gap, partisan asymmetry, and competitiveness) is paired with the compactness objective. The goal is to create Pareto-optimal district plans that are compact and are also fair with respect to one of the three political fairness criteria considered in this paper. Compactness is considered a key criterion since it is generally regarded as an essential requirement in practical district plans (National Conference of State Legislatures 2020). Since optimizing the compactness objective is NP-complete (Megiddo and Supowit 1984), each of the bi-objective PFDPs are NP-complete. Hence, an optimization algorithm for districting must address

two main challenges. First, the practical instance sizes for PFDPs are typically too large to be solved by an exact method. For example, Wisconsin has 1,409 census tracts (units) and 3,857 edges between them, and the MIP (2)-(9) has 7.5×10^6 variables and 6×10^6 constraints. Second, the PFDPs have multiple conflicting objectives, and the solution approach must produce solutions that highlight the inherent trade-offs between the objectives. This section presents a *multilevel* optimization framework which tackles these challenges.

Algorithm 1: Schematic of a multilevel districting algorithm

Input : Level 0 adjacency graph $G_0 = (V_0, E_0)$, number of districts K , population $p_i \forall i \in V$, number of voters $p_i^A, p_i^B \forall i \in V$, population deviation threshold $\tau \geq 0$, number of levels L , epsilon-constraint parameter ϵ

Output: z_0^* , the district solution at level 0

- 1 $\{G_0, G_1, \dots, G_L\} \leftarrow$ Coarsen from level 0 through L ;
 - 2 $z_L^* \leftarrow$ Find a Pareto-optimal solution at the L -th level using the ϵ -constraint method;
 - 3 **for** level $l \in \{L, L-1, \dots, 1\}$ **do**
 - 4 $z_{l-1}^* \leftarrow$ Execute heuristic improvement at G_{l-1} with z_l^* as initial solution;
 - 5 **end**
-

In this multilevel approach, the problem instance is reduced in size in multiple levels, an exact method solves the reduced instance to (near-)optimality, and the solution is mapped back to the original instance with local search refinements. An outline of the multilevel approach is presented in Algorithm 1. There are three stages in the approach: (a) coarsening, (b) solving an exact optimization problem at the coarsest-level, and (c) un-coarsening. The coarsening stage contracts the graph instances in multiple steps, thereby creating levels of graphs that are progressively coarser (i.e., smaller). This is achieved by finding a matching in every graph level and merging every pair of matched units to create a coarsened unit in the next graph level. The number of levels of coarsening is predetermined such that the size of the coarsest level graph is within a desired range. The coarsest level of graph is used as an input to the exact solution stage, where the districting problem is solved to optimality (or near-optimality). In the un-coarsening stage, the optimal solution at the coarsest level is mapped backed to the finer levels, with heuristic local search improvements at each of the levels.

The multilevel algorithm presented in this paper differs from existing multilevel algorithms for graph partitioning in two aspects. First, the objective of a typical graph partitioning problem is to minimize the number of “cut” edges, i.e., edges that cross two partitions. In the coarsening procedure of a typical multilevel algorithm, a *maximal matching* is found in the current graph

level and the matched pair of units (vertices) are merged. Several types of matchings are proposed in the literature, ranging from random matchings to matchings that minimize the sum of edge weights (Karypis and Kumar 1998). However, for a districting problem where unit populations in the original graph are heterogeneous, the goal of the matching procedure used in this paper is to coarsen the graph such that the unit populations in the coarse graph are as homogeneous as possible. To achieve this, the edges are given a weight based on the unit populations of the end points and a matching is found such that the maximum edge weight is minimized, as described in Section 4.1. This method is essential since it was observed that the presence of highly populous units in the coarse graphs prevents the algorithm from finding district plans that satisfy practical population balance constraints.

Second, solving the inner-level MIP of the districting problem is computationally more challenging than solving a typical graph partitioning problem. This is because of the additional constraints and variables in the districting MIP, such as the flow-based contiguity constraints and the fairness-defining constraints introduced in Section 3. Hence, a larger computational resource is needed to solve the inner-level MIPs. This challenge is handled in the presented work by setting large time limits of up to 24 hours for solving each MIP in the optimization stage as discussed in Section 5.

The rest of this section details the multilevel algorithm presented in this paper, organized as follows. Section 4.1 discusses the matching-based coarsening scheme used to create a coarse level instance. Section 4.2 presents an ϵ -constraint method to find Pareto-optimal solutions at the coarsest level by solving a series of MIPs. Section 4.3 provides a heuristic feasible solution to serve as an initial solution when an optimization solver is used to solve each MIP. Section 4.4 discusses the un-coarsening procedure that maps the coarse-level solution back to the original instance level.

4.1. Coarsening

The coarsening stage creates multiple levels of progressively smaller graphs by *merging* a subset of units at every level. Starting from the given adjacency graph $G_0 = G$ and number of levels $L \geq 1$, graphs $\{G_l = (V_l, E_l)\}_{l=1}^L$ are constructed such that pairs of units in G_{l-1} are merged to form the units in G_l for every $l \in [L]$. When two units are merged, their populations and numbers of voters are added. The pairs of units to be merged are chosen by solving a matching problem. Given an undirected graph $G = (V, E)$, a *matching* is a subset of edges $\mathcal{M} \subseteq E$ such that no two edges share the same node (unit), i.e., $(i, j), (i, j') \in \mathcal{M} \Rightarrow j = j'$. Once a matching is obtained, the two units at the endpoints of every matched edge are merged to form a single unit in the next level. Note that merging a pair of units is equivalent to imposing a constraint that those two units have to be assigned to the same district in the districting problem solved at the coarse instance.

Based on experimental investigations, it was observed that merging two highly populated units leads to an imbalanced population distribution among the units in the next level, $\{p_i\}_{i \in V_{l+1}}$, thereby making it difficult to find district plans that satisfy the population balance constraints. Hence, the objective of the matching-based coarsening considered in this paper is to ensure that the highly populous units are not matched together. For each edge $(i, j) \in E$, let $u_{i,j} = (p_i + p_j)$ be the edge weight that captures the total combined population of units i and j . The *minmax matching problem* finds a matching $\mathcal{M} \subset E$ of a particular size such that the maximum weight among the matched edges, $\max_{e \in \mathcal{M}} u_e$, is minimized. Two matching types are considered. A *maximum* (or maximum cardinality) matching is a matching of the largest size. When a maximum matching is used for coarsening, it leads to the most reduction in size of each graph level — by at most half the number of nodes (in which case it is said to be a *perfect* matching). A *maximal* matching \mathcal{M} is a matching such that no additional edges can be added to \mathcal{M} while \mathcal{M} remains a matching, i.e., $\forall e \in E \setminus \mathcal{M}$, $\mathcal{M} \cup \{e\}$ is not a matching. On one hand, coarsening using a maximum matching will yield the smallest possible graph in the next level, but may create highly populous units to achieve this reduction. In contrast, since there are at least as many maximal matchings as maximum matchings (i.e., since every maximum matching is also a maximal matching), the minmax maximal matching problem permits a larger set of matchings; hence, its optimal matching may be better (i.e., smaller weight) than that of the minmax maximum matching problem. Therefore, these two matching schemes provide a trade-off between achieving a coarse graph with the fewest units (i.e., maximum matching) and achieving a slightly larger coarse graph that avoids creating highly populous units (i.e., maximal matching). To explore this trade-off, this paper uses both maximum and maximal matchings to perform the coarsening.

4.1.1. Maximum Matchings The *minmax maximum matching problem* (MMP), also known as the *bottleneck matching problem* finds a maximum matching such that the weight of the heaviest edge is minimized. MMP can be solved using binary search as presented in [Gabow and Tarjan \(1988\)](#). The first step is to find the size of a maximum matching in G , which takes $O(n^2m)$ time using Edmonds’ blossom contraction algorithm ([Edmonds 1965](#)), where $n = |V|$ and $m = |E|$. Even though this paper uses Edmonds’ algorithm to find a maximum matching, note that [Micali and Vazirani \(1980\)](#)’s augmenting path algorithm is faster and takes $O(\sqrt{nm})$ time. For the instance of Wisconsin solved in this paper with $n = 1,409$ and $m = 3,857$, Edmonds’ algorithm implemented using [Eppstein \(2003\)](#) takes less than 5 seconds to find a maximum matching, so its computational requirements are not a limiting factor in the experiments presented in this study.

While Edmonds’ algorithm will find a maximum matching for G , it does not necessarily produce a minmax maximum matching, since it does not attempt to minimize the weight of the heaviest

edge in this matching. However, Edmonds' algorithms can form the basis of an algorithm to find a minmax maximum matching. Let $l^* \leq |V|/2$ be the size of the maximum matching in G . Next, the edges in G are sorted in the non-decreasing order of the edge weights, $\{e_1, e_2, \dots, e_m\}$, which takes $O(m \log n)$ time. To find the minmax maximum matching, the idea is to find the smallest subset of the sorted edges $\{e_1, e_2, \dots, e_t\}$, for $t \in [m]$, such that there exists a maximum matching of size l^* in the subgraph induced by this subset of edges. Note that the maximum edge weight in this subgraph is from edge e_t , and hence finding the smallest subset of edges is equivalent to minimizing the maximum edge weight (i.e., the objective of the MMP). This is achieved using binary search on $t \in [m]$, which takes $O(\log m)$ steps. Edmonds' algorithm is used as a subroutine to find the maximum matching for each subgraph induced by the subset of edges. Hence, the overall algorithm takes $O(n^2 m \log m)$ time.

4.1.2. Maximal Matchings The *minmax maximal matching problem* (MLP) finds a maximal matching such that the weight of the heaviest edge is minimized. While MMP can be solved in polynomial time, MLP is NP-complete via a reduction from the classical dominating set problem (Lavrov 2019). In this paper, a *greedy* algorithm is used to find a solution to MLP as follows. First, the edges are sorted in a non-decreasing order based on their weights, which takes $\mathcal{O}(m \log n)$ time. Starting with an empty matching $\mathcal{M} = \emptyset$, edges are visited in the sorted order. A visited edge e is added to \mathcal{M} if and only if e is not incident to any edge already in \mathcal{M} . This is verified by maintaining a binary array of size $|V|$ (where each element corresponds to a unit in V), that tracks whether each unit currently appears in the matching. Hence, it takes $\mathcal{O}(1)$ time to query the endpoints of an edge, to decide whether to add the edge to \mathcal{M} , and to update the array if the edge is added. The algorithm terminates when all the edges in the sorting are visited once, and there are $\mathcal{O}(m)$ steps in total. Overall, this algorithm runs in $\mathcal{O}(m \log n)$ time.

Even though the greedy algorithm may not solve MLP optimally, a theoretical guarantee on its approximation gap can be derived. Section EC.3 of the e-companion discusses the approximability of MLP for the population-based edge weights (i.e., $u_{i,j} = (p_i + p_j)$ for edge $(i, j) \in E$) considered in this paper assuming that the unit populations $\{p_i\}_{i \in V}$ are strictly positive. Theorem 2 shows that the objective value of any maximal matching (and in extension, the solution returned by the greedy algorithm) is at the most ρ times the objective value of the optimal maximal matching to MLP, where $\rho := \max_{(i,j) \in E} \max\{p_i/p_j, p_j/p_i\}$ is the maximum population ratio among neighboring units in G .

THEOREM 2. Consider a connected graph $G = (V, E)$, positive unit populations $p_i > 0$ for every unit $i \in V$, and edge weights $u_{i,j} = p_i + p_j$ for every edge $(i, j) \in E$. Let M^* be an optimal solution to the minmax maximal matching problem with inputs $(G, \{u_{i,j}\}_{(i,j) \in E})$, and let $u^* := \max_{e \in M^*} u_e$

be its maximum edge weight. Let M be an arbitrary maximal matching with maximum edge weight $u' := \max_{e \in M} u_e$. Then, $u'/\rho \leq u^* \leq u'$, where $\rho := \max_{(i,j) \in E} \max\{p_i/p_j, p_j/p_i\}$.

The proof is presented in Section EC.3 of the e-companion. Note that this result does not apply to instances which possess zero-population units. Furthermore, Section EC.3 demonstrates the tightness of this approximation by constructing an adversarial example where the greedy algorithm returns the worst possible maximal matching with an objective value that is ρ times the optimal objective value. Therefore, the approximability of MLP for the general class of edge weights remains an open problem.

4.1.3. Matching-Based Coarsening Once a matching \mathcal{M}_l is chosen at level $l \in [L]$, the matched edges are merged to obtain a coarse graph. The next graph level $G_{l+1} = (V_{l+1}, E_{l+1})$ is created by merging every matched edge from G_l into a single node in V_{l+1} . Let $\mathcal{C} : V_l \rightarrow V_{l+1}$ be the mapping of a node from level l to level $l+1$ such that for $i, j \in V_l$, $\mathcal{C}(i) = \mathcal{C}(j)$ if and only if $(i, j) \in \mathcal{M}_l$, where \mathcal{M}_l is a minmax weight matching in G_l . The matching \mathcal{M}_l is either the maximum matching found using Edmonds' algorithm described in Section 4.1.1, or the maximal matching found using the greedy algorithm described in Section 4.1.2. The edges of the coarse-level graph G_{l+1} are defined such that if and only if there is at least one edge $(i, j) \in E_l \setminus \mathcal{M}_l$ in G_l , the edge $(\mathcal{C}(i), \mathcal{C}(j)) \in E_{l+1}$ in G_{l+1} . The population of a node in V_{l+1} is the sum of the populations of the corresponding matched units in V_l . The number of voters are also similarly aggregated.

When the maximal/maximum matching is close to a perfect matching, the next level graph has approximately half the number of nodes of the previous level. Hence, the number of levels L can be chosen such that the size of the graph in the coarsest level, $|V_L|$, is within the capacity of an optimization solver to exactly solve the districting problem for the coarsened instance. For example, when starting from a graph of size $|V_0| = 1,409$ (i.e., the number of census tracts in Wisconsin), if an exact solution for PFDP can be obtained for a 200-node graph, assuming near-perfect matchings at each level, L can be set to $\lceil \log_2(\frac{1409}{200}) \rceil = 3$ levels.

4.2. Exact Method for PFDP

This subsection presents the inner stage of the multilevel algorithm (line 2 of Algorithm 1) comprising an exact solution strategy for solving a bi-objective PFDP. The coarsest level of graph, G_L , is used as the input. In a multi-objective optimization problem with $Q \geq 2$ minimization objective functions given by $\phi = \{\phi_1, \phi_2, \dots, \phi_Q\}$ and a solution space \mathcal{X} , two solutions $\hat{\mathbf{x}}, \mathbf{x} \in \mathcal{X}$ can be compared using dominance relationships. The solution $\hat{\mathbf{x}}$ is said to *dominate* \mathbf{x} if $\hat{\mathbf{x}}$ has a strictly better objective value in at least one of the Q objective functions and does not have a worse value in any of the other objective functions, i.e., $\phi_{q'}(\hat{\mathbf{x}}) < \phi_{q'}(\mathbf{x})$ for some $q' \in [Q]$, and $\phi_q(\hat{\mathbf{x}}) \leq \phi_q(\mathbf{x})$ for all $q \in [Q]$.

Further, $\hat{\mathbf{x}}$ is said to be *Pareto-optimal* if $\hat{\mathbf{x}}$ is not dominated by any other solution in \mathcal{X} . Hence, a Pareto-optimal solution is one that cannot be improved with respect to one objective function without sacrificing the value of another objective function. The goal of solving a multi-objective optimization problem is to find the set of all Pareto-optimal solutions. Let \mathcal{P} denote the set of all Pareto-optimal solutions, or the Pareto set. The *Pareto-frontier* is the set of objective values $\{\{\phi_1(\hat{\mathbf{x}}), \phi_2(\hat{\mathbf{x}}), \dots, \phi_Q(\hat{\mathbf{x}})\}\}_{\hat{\mathbf{x}} \in \mathcal{P}}$ in the objective space corresponding to Pareto-optimal solutions. The set of values in the Pareto-frontier provides a trade-off of the extent to which one objective function can be made better by sacrificing another objective function.

The ϵ -constraint method is a popular method to compute the Pareto set in multi-objective problems (Chiandussi et al. 2012). This method first selects one of the objectives as a *primary* objective, while the other *secondary* objectives are introduced as constraints of the form $\phi_q(\mathbf{x}) \leq \epsilon_q \forall \mathbf{x} \in \mathcal{X}$, where ϵ_q is a user-imposed upper bound on a secondary objective $q \in [Q]$. This is an iterative procedure where the values of ϵ_q are gradually reduced, starting from a large value. The amount by which the values of ϵ_q are reduced in every iteration should be small enough that no non-dominated solution is skipped over. For each value of ϵ_q , a single-objective problem is solved with the primary objective function as the main objective. The optimal solution found in each iteration is stored, and is a candidate to be a Pareto-optimal solution. The method terminates when the value of ϵ_q is so small that the optimization solver either returns an infeasibility certificate or does not find a feasible solution within the time limit. If the primary objective function is such that several solutions could have the same value, the ϵ -constraint method may produce dominated solutions. When the dominated solutions are removed from the set of candidates, the rest of the solutions form the Pareto set.

For a bi-objective PFDP like those that will be solved in Section 5, the compactness objective (ϕ_{comp}) as defined in (10) is chosen as the primary objective since the moment of inertia function (i.e., sum of the weighted squared distances from the district centers) is structurally similar to the p -median problem, which is known to possess good linear programming lower bounds (Salazar-Aguilar et al. 2011). Hence, solving a single-objective problem in each iteration of the ϵ -constraint method is relatively more computationally efficient with compactness as the primary objective. This was experimentally verified by setting each of the other fairness objectives as the primary objective, which resulted in a lack of convergence to optimality. Additionally, since compactness is universally regarded as an essential requirement in political districting, it is considered a primary criterion (National Conference of State Legislatures 2020). Hence, the solution approach is to solve Compact-DP supplemented by constraints related to the other three fairness metrics (as formulated in their single-objective problems) and the ϵ -constraints. The corresponding ϵ -constraints for

ϕ_{EG} , ϕ_{cmpttv} and ϕ_{PA} are given by constraints (35), (36) and (37), where ϵ_{EG} , ϵ_{PA} and ϵ_{cmpttv} are upper bounds on the efficiency gap, partisan asymmetry, and competitiveness, respectively.

$$-\epsilon_{EG} \sum_{i \in V} (p_i^A + p_i^B) \leq \sum_{i \in V} a_i \leq \epsilon_{EG} \sum_{i \in V} (p_i^A + p_i^B), \quad (35)$$

$$\sum_{k=1}^K |\hat{\omega}_k - (1 - \hat{\omega}_{K-k+1})| \leq K\epsilon_{PA}, \quad (36)$$

$$h \leq \epsilon_{cmpttv}. \quad (37)$$

As described in Section 3, a_i is the difference in wasted votes between parties A and B in a district with center $i \in V$, $\hat{\omega}_k$ is the minimum average vote-share required for party A to win $k \in [K]$ districts, and h is the maximum margin across all the districts. Note that the absolute value in constraint (36) can be linearized by standard linearization techniques. In particular, for each $k \in [K]$, the absolute value term in constraint (36) is replaced by an additional continuous variable $\Omega_k \geq 0$, along with additional constraints $-\Omega_k \leq (\hat{\omega}_k - (1 - \hat{\omega}_{K-k+1})) \leq \Omega_k$. Then, the constraint $\sum_{k=1}^K \Omega_k \leq K\epsilon_{PA}$ implies constraint (36).

In every iteration of the ϵ -constraint method, it is important to decrement the value of ϵ such that no non-dominated solutions is skipped between successive iterations. In this paper, ϵ is decremented by 10^{-7} from the corresponding fairness value of the solution found in the previous iteration. This value is obtained by first expressing each of the three fairness objectives (ϕ_{EG} , ϕ_{cmpttv} and ϕ_{PA}) as fractions, and noting that for the Wisconsin instance solved in this paper, 10^7 is larger than the denominator term in each of these fractions. Hence, when ϵ is decremented by 10^{-7} , no non-dominated solutions is skipped between successive iterations.

While this section describes a method to obtain a Pareto-frontier, the case study presented in this paper finds an *approximate Pareto-frontier*. The approximation arises from three reasons. First, in a particular iteration for a given ϵ , if the optimization solver terminates due to time limit, it may return a non-optimal solution which is not Pareto-optimal. Second, in this scenario, it is possible that non-dominated solutions in future iterations may be skipped over. Third, in the last iteration, when the solver terminates due to time limit without a feasible solution, it is still theoretically possible that a non-dominated solution exists whose fairness value is smaller than the (smallest) ϵ . To minimize these issues in the case study, the time limit for each iteration of the ϵ -constraint method is generously set to 24 hours.

4.3. Initial Solution Heuristic

In the exact optimization stage to solve an MIP, the computational performance of an optimization solver can be improved when the solver is given an initial feasible solution. This section presents

a multistart local search heuristic used to generate a feasible solution for a PFDP. A feasible solution is one that satisfies contiguity, population balance, and an ϵ -constraint from (35)–(37). The heuristic proceeds in two phases. The first phase creates a feasible solution, and the second phase improves the compactness objective using local search.

To generate a feasible solution, the heuristic proceeds as in Bozkaya et al. (2003), Vickrey (1961), and Thoreson and Liittschwager (1967), as follows. An arbitrary unit is first picked uniformly at random from V to seed a district, and this district is extended by randomly assigning neighboring units to it. This district is iteratively grown by adding one unassigned neighboring unit at a time. This process continues until either the population in the district exceeds \bar{P} or there are no more unassigned units in any of the neighborhoods of the units in the district. Next, a new unassigned unit is randomly chosen to seed and grow a new district. This proceeds continues until all the units in the graph are assigned to districts. At the end of this procedure, it is possible to have either more than, equal to, or fewer than K districts. If there are more than K districts, pairs of neighboring districts are iteratively merged until the number of districts equals K . Hence, this method finds K contiguous districts. If there are fewer than K districts, or if the solution does not satisfy the population balance or fairness-based ϵ -constraints, it is discarded, and the method is repeated until a feasible solution is found. Based on experiments not included in this paper, it is observed that the computational time to find an initial solution increases as the feasible region to the districting problem is made smaller, either by tightening the ϵ -constraint or by reducing the population deviation threshold τ .

In addition to finding a feasible solution, it is preferable to find one that is highly compact (i.e., the primary objective). In a local search method, starting with a feasible solution, the neighborhood of the solution is explored for any improvement in the objective (Ricca et al. 2013). The neighborhood considered in this paper is obtained by re-assigning a single unit from one district to another district to reduce the compactness objective ϕ_{comp} . Note that such a re-assignment is referred to as a “flip” operation in algorithms that generate random district plans using Markov chains, e.g., DeFord et al. (2019). After every such re-assignment, the feasibility of the solution is assessed. The solution in the neighborhood with the smallest compactness value, ϕ_{comp} , is chosen as the starting point for the next iteration, and the local search terminates either when the number of iterations reaches a limit, or when there is no further improvement. The quality of the best solution from the local search method is typically influenced by the initial solution. To encourage a more complete exploration of the solution space, a multistart local search procedure is used, where multiple feasible solutions are generated as starting points for local search. The initial solution for the MIP is the solution with the smallest ϕ_{comp} produced by the multistart local search.

4.4. Un-coarsening

The final stage of the multilevel algorithm (lines 3-5 in Algorithm 1) iteratively maps the optimal solution from level L back to level 0. Note that the optimal solution at level L may not be optimal at lower levels, and hence this solution could potentially be improved using a heuristic. Let $z_{l+1} : V_{l+1} \rightarrow [K]$ be the district solution at level $l+1$. First, an initial district solution at level l is created, i.e., $z_l(i) = z_{l+1}(\mathcal{C}(i))$ for all $i \in V_l$, where \mathcal{C} is the mapping function defined in the coarsening procedure. Then, using the local search method described in Section 4.3, this solution is improved further with respect to ϕ_{comp} while enforcing contiguity, population balance and an ϵ -constraint. After L levels of un-coarsening, the algorithm returns the locally-optimized solution z_L .

5. Case Study in Wisconsin

This section applies the solution approach described in Section 4 to approximate the trade-offs between compactness and each of the political fairness metrics (i.e., efficiency gap, partisan asymmetry, and competitiveness) for congressional districting in the state of Wisconsin. The population data and adjacency information are obtained from the [U.S. Census Bureau \(2010\)](#), and are processed using QGIS. Wisconsin has $|V| = 1,409$ census tracts (units) and $K = 8$ congressional districts with ideal district population $\bar{P} = 710,873$. Although census blocks are more granular than census tracts, this paper uses census tracts as the basic units since fewer than 5% of the census tracts are split between multiple districts according to Wisconsin's 2010 district plan. Voter information is compiled as the average of the 2012 and 2016 presidential election results ([McGovern 2017](#), [The Guardian 2018](#)). The Democratic (D) and Republican (R) parties have an approximately equal overall vote-share (51.6% and 48.4% respectively) across the state. This information is available at the precinct level, which is aggregated to the county level, and then distributed to the census tracts proportionally based on their populations, after which the number of voters in each census tract is rounded down to the nearest integer. While more sophisticated aggregation techniques exist, where precinct-level data is de-aggregated to census blocks and is then aggregated to census tracts ([MGGG 2021](#)), the coarse approximation of voting patterns yielded by disaggregating from counties still provides an example instance on which to test the multilevel algorithm for the PFDP proposed in this study.

Population balance requires that the relative difference between every district's population and the ideal district population lies within a given deviation threshold $\tau \geq 0$ ([Kalsics and Ríos-Mercado 2019](#)). Alternatively, it is also common to require that the relative difference between the most and least populous districts be within a tolerable range ([Grofman and Cervas 2020](#)). Even though the Supreme Court does not set an exact legal limit for the value of τ , they require that district populations be “as nearly [equal] as is practicable” ([Wesberry et al. v. Sanders, Governor](#)

of Georgia, et al. 1964). In the case study presented in this paper, τ is set to 2%, which allows for a maximum deviation of up to 14,217 people in each district from the ideal district population. To achieve the goal of analyzing the trade-offs between the fairness objectives, a 2% deviation is preferred since it allows the algorithm to examine the tradeoffs within a larger pool of candidate district plans compared to using a smaller deviation threshold. Moreover, even though a district plan with a 2% deviation may be considered unconstitutional, such a plan can be refined at the census block level to achieve a plan with a tighter population deviation. DeFord et al. (2020) note that a plan with a 1% deviation is “easily tuneable to 1-person deviation by refinement at the [census] block level”, and the same refinement technique can be extended to a plan with a 2% deviation. Further, the 2% deviation has precedence in prior work in Congressional districting, such as in DeFord and Duchin (2019). As a proof of concept, Section 5.2 presents district plans with τ as low as 0.25% to showcase the ability of the multilevel algorithm to produce practical district plans even without census block level refinements.

The computations reported in this paper are performed using a 2.7 GHz Intel Core i5 machine with 8 GB memory. The data sets, code, and the results can be found at Swamy (2022). The algorithms are written in Python 3.7. The computational times are reported using both clock time (seconds) and CPU time (ticks). For each of the PFDPs, the initial solution from the multistart local search is first computed with the termination criteria being either 100 iterations of the multistart local search, or a time limit of 60 minutes. To solve the MIP formulations, IBM’s CPLEX Optimizer 12.6 is used with its default settings which include a host of cuts and heuristics. The MIPGap parameter is set to zero in order to find the optimal solutions to the MIPs. The initial solution found using the multi-start heuristic is passed into the root of the branch-and-cut tree. A time limit of 24 hours is set for each run of solving the MIPs. If the optimality gap is non-zero, the best feasible solution is returned as an approximate optimal solution. If no feasible solution is found, the bi-objective multilevel algorithm terminates, and the current Pareto frontier is returned.

The rest of this section is organized as follows. Section 5.1 presents results from the two matching-based coarsening procedures used to coarsen Wisconsin’s census tract level instance. Section 5.2 presents district plans produced from solving the Compact Districting Problem with the population deviation threshold (τ) value ranging from 0.25% to 2%. Sections 5.3, 5.4 and 5.5 present district plans obtained from solving three bi-objective problems, each with compactness as the primary objective and a respective political fairness metric — efficiency gap, partisan asymmetry and competitiveness — as the secondary objective. The approximate Pareto-frontiers produced highlight the trade-offs between compactness and the three political fairness metrics.




Figure 3 Adjacency graphs of Wisconsin for three levels of coarsening. Each level is obtained by merging matched units in a minmax maximal matching in the previous level, as described in Section 4.1.2.

5.1. Coarsening

Starting with the adjacency graph at the finest level (census tracts), the matching-based coarsening procedure proceeds as described in Section 4.1. Two coarse-level instances are produced when coarsening using (a) a minmax maximum matching (MM), and (b) a minmax maximal matching (ML). For both methods, the number of levels L is set to three, since it was observed that the optimization solver performs well for instance sizes around 200 units, and three levels of near-perfect matchings would attain instances near this size.

Coarsening using MM reduces the graph size to $|V_3| = 177$ nodes, whereas ML reduces the graph size to $|V_3| = 224$ nodes. Figure 3 depicts the adjacency graphs of the three levels of graphs produced when coarsening using ML. As expected, coarsening using MM produces a smaller graph instance than when using ML, and hence could lead to faster computation of Pareto-optimal solutions at the coarsest level. However, in the coarsest graph G_3 , the population of the most-populous unit when using MM is 52,207, whereas when using ML is 47,932. The subsequent Sections 5.3 — 5.5 present results for PFDPs using both coarsening methods. Note that the definition of compactness depends on the structure of a graph instance. Since the two coarsening methods produce two different coarse graph instances, the compactness values cannot be directly compared between the two methods. Hence, in Figures 5, 7 and 9, the vertical axes representing compactness have different scales between the two coarsening methods.

5.2. Compact Districts

The multilevel algorithm first solves the traditional Compact-DP formulated by (2)-(10). Coarsening using ML produces the graph instance $G_3 = (V_3, E_3)$ after three levels of coarsening. Since this section serves as a proof-of-concept that the multilevel approach can produce compact districts meeting tighter population balance tolerances than $\tau = 2\%$, coarsening using MM is not presented here as results using ML are sufficient to demonstrate this capability. In the definition of compactness in (10), the distance d_{ij} between every pair of units $i, j \in V_3$ is measured by the number of

edges in the shortest path between i and j in G_3 . Note that Hess et al. (1965)'s original model defines d_{ij} to be a general distance function. Mehrotra et al. (1998) conclude that the shortest path-based definition used in this paper produces less geographical bias than using Euclidean distances since it is invariant of the geographical sizes of the units. This definition of compactness is used in the rest of Section 5.

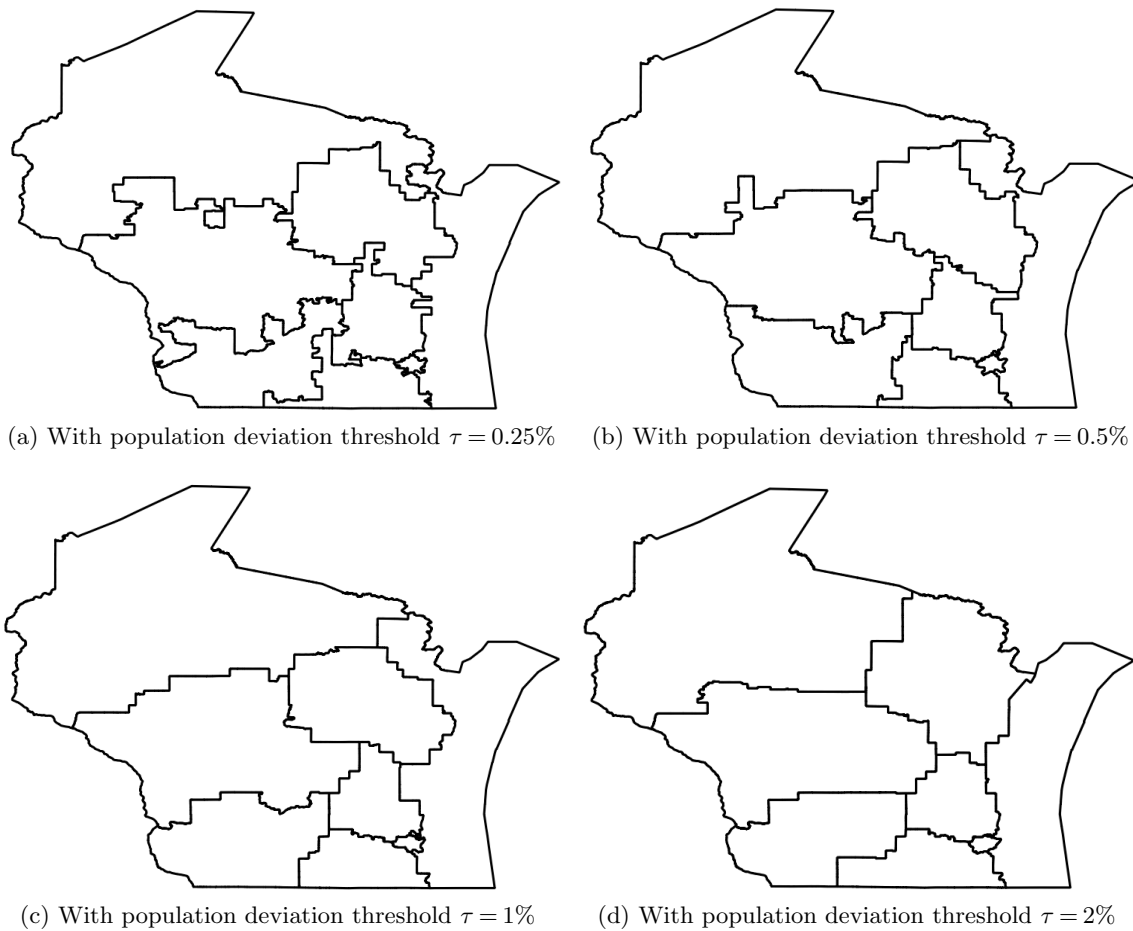


Figure 4 Compact district plans for varying tightness of the population balance constraints.

To observe the impact of the population balance threshold (τ) on the compactness of the district plan produced by the algorithm, four values for τ are chosen from $\{0.25\%, 0.5\%, 1\%, 2\%\}$. Figure 4 depicts the four district plans produced for these values of τ . The inner-level exact optimization stage solves the coarsened instance with 224 units within 1 hour for each of the four settings of τ .

Further, the plan obtained with $\tau = 0.25\%$ is testament to the ability of the multilevel algorithm to produce district plans that satisfy practical population balance requirements. It can be seen that as τ is increased from 0.25% to 2%, the corresponding districts produced by the algorithm are visually more compact. Comparing the four plans, the majority of the units in the districts remain

in the same district across the plans. The differences arise in the district boundaries, where the compact shapes are perturbed to satisfy a tighter population balance constraint for smaller values of τ . These plans illustrate the trade-off between the compactness objective and the tightness of the population balance constraint. To allow more flexibility when approximating the trade-offs between compactness and the three political fairness objectives, a population deviation threshold of 2% is set in the rest of this section.

5.3. Efficient Districts

To approximate the trade-off between compactness and efficiency gap (EG), the approximate Pareto-optimal solutions are obtained by solving the PFDP with objective set $\{\phi_{comp}, \phi_{EG}\}$, formulated by (2)-(10), (14)-(20) and (35) using the coarsest levels obtained from the two coarsening methods. The value of ϵ_{EG} is iteratively reduced from 0.5 (i.e., the theoretical maximum possible value for EG), and the MIP is solved either to optimality or until the time limit is reached. Figure 5 depicts the Pareto-frontiers highlighting the trade-off between the compactness and EG objectives. Coarsening using MM and ML produce 14 and 11 district plans respectively, with similar median EG values of 0.023 and 0.027, respectively. Every iteration terminated to optimality within the 24 hour time limit. The computational results are presented in Section EC.4.1 of the e-companion. All the solutions terminated to optimality.

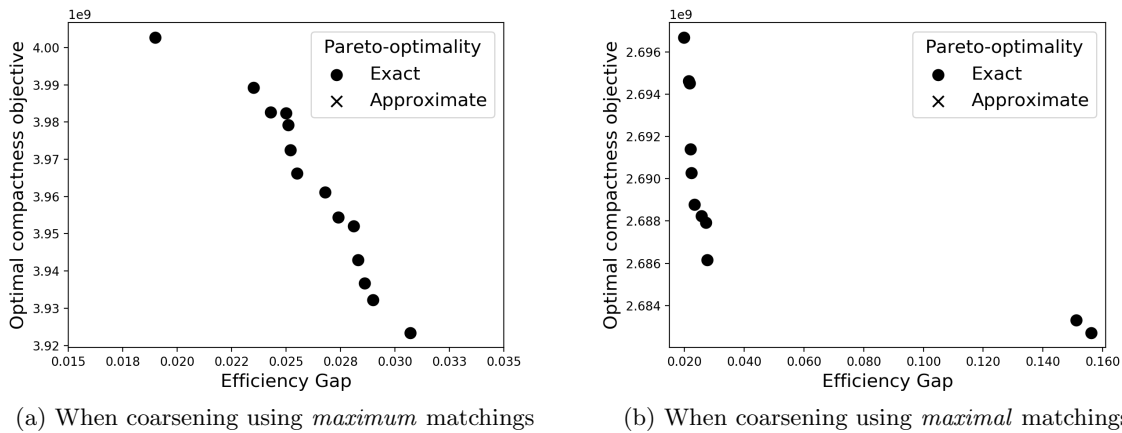



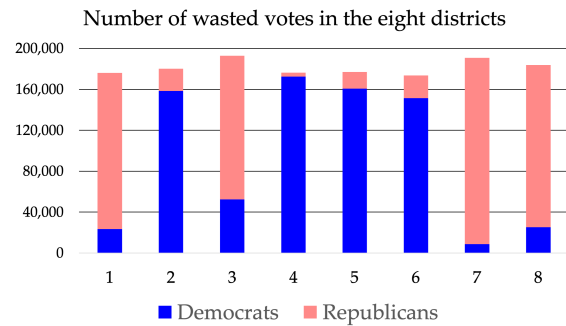
Figure 5 Approximate Pareto-frontier using the coarsest graph instance G_3 highlighting the trade-off between the efficiency gap and compactness objectives. Each solution represents either an exact or an approximate solution found in the coarsest graph instance G_3 .

When the solutions reported in Table EC.4.1 of the e-companion are un-coarsened with local search improvements to compactness, the smallest value of EG attained is 0.0188 (0.0198) when the

coarsening procedure uses MM (ML). The district plan corresponding to an EG of 0.0188 is depicted in Figure 6a. Both the parties win four districts each, and hence the seat-share is proportional to the overall vote-share. Figure 6b depicts the number of votes wasted by the two parties in the eight districts. This district plan is an ideal candidate to a districting process that seeks to minimize the difference in wasted votes between the two parties. Even though this district plan minimizes EG, this plan is neither competitive nor symmetric. It can be seen all eight districts are highly packed and cracked by both parties to an almost equal degree, leading to poor competitiveness as measured by a maximum margin of 27.19%. Moreover, this plan also has a high partisan asymmetry value of 0.0218, which is much larger than the best asymmetry value of 0.0002 achieved in Section 5.4, due to a large difference in the two parties' vote-seat curves.



(a) Party majorities in the eight districts classified based on their margins of victory.



(b) Number of votes wasted by both parties.

Figure 6 An efficient district plan for Wisconsin with an efficiency gap of 0.0188.

5.4. Symmetric Districts

To approximate the trade-off between compactness and partisan asymmetry (PA), the approximate Pareto-optimal solutions are obtained by solving the PFDP with objective set $\{\phi_{comp}, \phi_{PA}\}$, formulated by (2)-(10), (36) and the linearized constraints (EC.2)-(EC.14), (EC.17)-(EC.22) given in the proof of Proposition 1 in Section EC.2 of the e-companion. A reduced time limit of 6 hours is used for each run of solving the MIP since the 24 hour time limit produced large branch-and-bound trees whose sizes exceeded the computer's memory. This is due to the relative complexity of the MIP that models PA. To produce the district plans, the value of ϵ_{PA} is iteratively reduced from 1 (i.e., the theoretical maximum possible value for PA). The solver terminates with an approximate solution in every iteration owing to the complexity of solving the MIP. Coarsening using MM and ML produce 4 district plans each, with median PA values of 0.0021 and 0.0073, respectively. Figure


(a) When coarsening using *maximum* matchings.(b) When coarsening using *maximal* matchings.

Figure 7 Approximate Pareto-frontier using the coarsest graph instance G_3 highlighting the trade-off between the partisan asymmetry and compactness objectives. Each solution represents either an exact or an approximate solution found in the coarsest graph instance G_3 .

7 depicts the approximate Pareto-frontiers highlighting the trade-off between the compactness and PA objectives. The computational results for the eight district plans produced are presented in Section EC.4.2 in the e-companion.

When the solutions reported in Table EC.4.2 are un-coarsened with local search improvements to compactness, the smallest value of PA attained is 0.0002 (0.0016) when the coarsening procedure uses MM (ML). Further, note that among the six district plans with PA values less than 0.01, five of them have R winning 4 districts out of 8 and the sixth one has R winning 5 districts out of 8. This suggests that the majority of the symmetric district plans align with vote-seat proportionality. Figure 8a depicts a district plan with a PA value of 0.0002 found when coarsening using MM. In this plan, both the parties win four districts each. Figure 8b shows the corresponding vote-seat curves for the two parties. It can be seen that the two curves are close to each other with an area of 0.0002 between them, signifying that both parties gain or lose seats symmetrically when voting levels fluctuate from their actual values. Hence, this district plan is an ideal candidate for a districting process that seeks to minimize PA. Even though this plan has been optimized for PA, it is not competitive. The high concentration of D voters in district 4 gives this plan a competitiveness value (maximum margin) of 14.39%. However, since both parties win an equal number of districts, the number of seats won by each party is proportional to its overall vote-share, giving rise to a small EG value of 0.0216 where D wastes more voters than R.

5.5. Competitive Districts

To approximate the trade-off between compactness and competitiveness, the approximate Pareto-optimal solutions are obtained by solving the PFDP with objective set $\{\phi_{comp}, \phi_{cmpttv}\}$, formulated

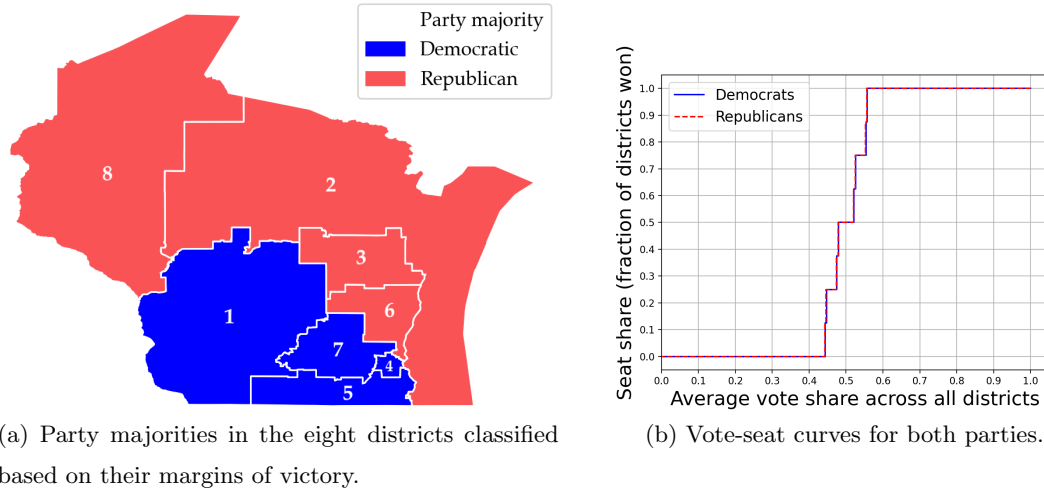


Figure 8 A symmetric district plan for Wisconsin with a partisan asymmetry of 0.0002.

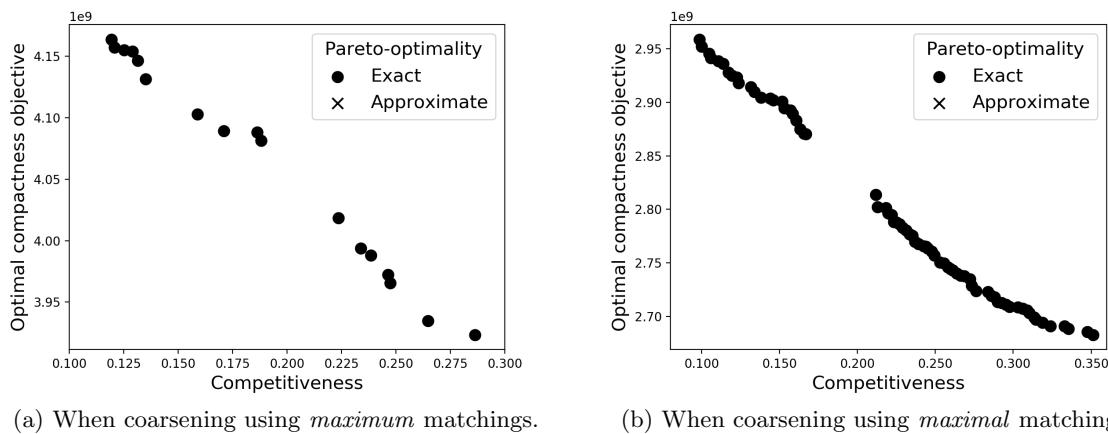


Figure 9 Approximate Pareto-frontier using the coarsest graph instance G_3 highlighting the trade-off between the competitiveness and compactness objectives. Each solution represents either an exact or an approximate solution found in the coarsest graph instance G_3 .

by (2)-(10), (29)-(34), and (37). The value of ϵ_{cmpttv} is iteratively reduced from 1 (i.e., the theoretical maximum possible value for the maximum margin), and the MIP is solved either to optimality or until the time limit is reached. Figure 9 depicts the Pareto-frontiers highlighting the trade-off between the compactness and competitiveness objectives. Coarsening using MM and ML produce 17 and 72 district plans respectively, with respective median competitiveness values of 18.6% and 23.8%. Every iteration terminated to optimality within the 24 hour time limit. The instances produced by MM are solved slightly slower (with an average of 64.2 minutes) than the instances produced by ML (with an average of 59.1 minutes).

When the solutions reported in Table EC.4.3 are un-coarsened with local search improvements to compactness, the smallest value of competitiveness value attained is 11.94% (9.78%) when the

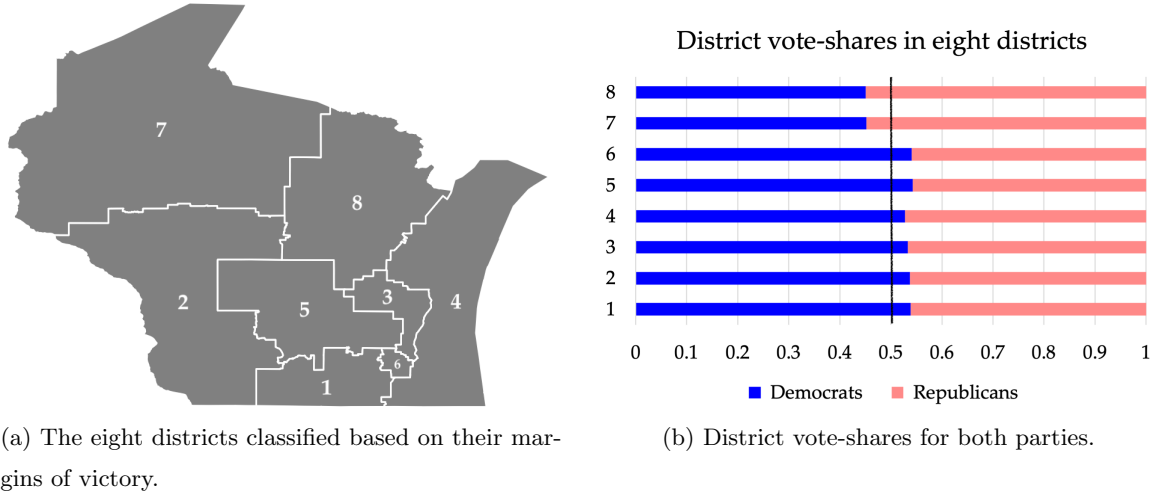


Figure 10 A competitive district plan for Wisconsin with a maximum margin of 9.78%.

coarsening procedure uses MM (ML). Figure 10a depicts a district plan with all eight districts have their margins less than 9.78% found when coarsening using ML, and Figure 10b depicts the corresponding district vote-shares. Party D disproportionately wins six out of the eight districts, and hence this plan has a high EG value of 0.2202. This plan also has a high PA value of 0.0374 (compared to the district plan in Section 5.4 with a PA of 0.0002) due to the asymmetry in the seat gains between the two parties. The computational details are presented in Section EC.4.3 of the e-companion.

6. Conclusions

Political districting is a problem of national interest with consequences to electoral representation. This paper addresses gerrymandering by providing a practical multi-objective framework that can be adopted by a districting process that explicitly focuses on political fairness. The mathematical models presented in this paper optimize a set of fairness metrics that cater to different facets of fairness. The computational experiments demonstrate an algorithm for creating district plans using two objectives at a time. The case study for congressional districting in Wisconsin highlights that solutions that are approximately optimal with respect to one objective may be poor with respect to another. These results reiterate the need for a holistic multi-objective perspective in the design of political districts to best serve the interests of multiple stakeholders.

This paper re-frames the districting process and serves as a starting point to investigate how political fairness can be integrated into optimization-based approaches for district design. There are several aspects of this framework that can be improved with future work. First, optimizing over one pair of fairness objectives could unintentionally produce solutions with poor values for the other objectives, as was observed in the Wisconsin case study reported in Section 5. Hence, extending the ϵ -constraint method to optimize more than two objectives at a time could yield a more

holistic approach that considers simultaneous trade-offs between multiple objectives. Second, the algorithm relies on the use of historic population and voting data to draw districts for the future. Incorporating potential variations in voting trends into the mathematical models by using robust optimization could produce district plans that are more relevant to future needs when populations and voter preferences are unknown. Third, the matching-based coarsening procedure used in this paper weights the graph edges to encourage a homogeneous population distribution in the coarse level graph. However, a population-based coarsening could have unintended consequences to the demographic composition of the coarse level units. For example, merging a high-population unit in an urban area with a low-population unit in a surrounding rural area could create coarse units that dilute the voice of the voters in either unit, which could have unintended consequences for racial, socioeconomic, or other representation. Hence, an analysis of the impact of coarsening on the demographic composition of the coarse level graph could yield a more informed coarsening procedure that best fits the goals of the districting process. Fourth, improving the efficiency of exactly solving the PFDPs using methods that utilize the structure of the MIPs (e.g., using decomposition techniques) could yield a more exhaustive Pareto-frontier within a shorter computational time.

7. Acknowledgements

The third author's research (Jacobson) has been supported in part by the Air Force Office of Scientific Research (FA9550-19-1-0106). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the United States Government, or the Air Force Office of Scientific Research. The authors would like to thank the referees and the Editorial team at *Operations Research* for their extensive feedback in improving this manuscript.

Author Biographies

Rahul Swamy

Rahul Swamy is a Ph.D. Candidate in Industrial Engineering (Operations Research) at the University of Illinois at Urbana-Champaign. His research interests include discrete optimization and algorithmic game theory. His work on political redistricting has been recognized with the INFORMS Service Science Best Paper Award 2019 (First place), INFORMS Poster Competition 2019 (First place), and the INFORMS Public Sector Operations Research (PSOR) Best Paper Award 2018 (Finalist). He has served in student leadership roles in the OR/MS community, such as Lead Editor for INFORMS OR/MS Tomorrow (2017-2019) and as President of the INFORMS Student Chapter at UIUC (2016-2018).

Douglas M. King, Ph.D.

Douglas M. King is a Teaching Assistant Professor in the Department of Industrial and Enterprise Systems Engineering at the University of Illinois. His research interests include discrete optimization and mathematical modeling in applications related to public policy. His work on redistricting has been recognized with the IISE Pritsker Doctoral Dissertation Award (Second Place) and in the INFORMS Section of Public Programs, Services, and Needs Best Paper Competition (Finalist), among others.

Sheldon H. Jacobson, Ph.D.

Sheldon H. Jacobson is a Founder Professor of Computer Science at the University of Illinois. He has a B.Sc. and M.Sc. (both in Mathematics) from McGill University, and a Ph.D. (in Operations Research) from Cornell University. His research interests span theory and practice, covering decision-making under uncertainty and optimization-based artificial intelligence, with applications in aviation security, public policy, public health, and sports. He has been recognized by numerous awards including a Guggenheim Fellowship and the INFORMS Saul Gass Expository Writing Award. He is a fellow of AAAS, INFORMS, and IISE.

References

- Abramowitz AI, Alexander B, Gunning M (2006) Incumbency, redistricting, and the decline of competition in US House elections. *The Journal of Politics* 68(1):75–88.
- Adams WP, Sherali HD (1990) Linearization strategies for a class of zero-one mixed integer programming problems. *Operations Research* 38(2):217–226.
- Bacao F, Lobo V, Painho M (2005) Applying genetic algorithms to zone design. *Soft Computing* 9(5):341–348.
- Ballotpedia (2020) Missouri Senate Joint Resolution No. 38. URL <https://legiscan.com/MO/text/SJR38/2020>, [Accessed: 28-Jan-2021].
- Benadè G, Procaccia AD, Tucker-Foltz J (2021) You can have your cake and redistrict it too.
- Bernstein M, Duchin M (2017) A formula goes to court: Partisan gerrymandering and the efficiency gap. *Notices of the AMS* 64(9):1020–1024.
- Bozkaya B, Erkut E, Laporte G (2003) A tabu search heuristic and adaptive memory procedure for political districting. *European Journal of Operational Research* 144(1):12–26.
- Browdy MH (1990) Simulated annealing: An improved computer model for political redistricting. *Yale Law & Policy Review* 8(1):163–179.
- Buluç A, Meyerhenke H, Safo I, Sanders P, Schulz C (2016) Recent advances in graph partitioning. *Algorithm Engineering*, 117–158 (Springer).

- Chambers CP, Miller AD, Sobel J (2017) Flaws in the efficiency gap. *Journal of Law & Politics* 33:1.
- Chatterjee T, DasGupta B, Palmieri L, Al-Qurashi Z, Sidiropoulos A (2020) On theoretical and empirical algorithmic analysis of the efficiency gap measure in partisan gerrymandering. *Journal of Combinatorial Optimization* 40:512–546.
- Chen J, Rodden J (2013) Unintentional gerrymandering: Political geography and electoral bias in legislatures. *Quarterly Journal of Political Science* 8(3):239–269.
- Chiandussi G, Codegone M, Ferrero S, Varesio FE (2012) Comparison of multi-objective optimization methodologies for engineering applications. *Computers & Mathematics with Applications* 63(5):912–942.
- D’Amico SJ, Wang SJ, Batta R, Rump CM (2002) A simulated annealing approach to police district design. *Computers & Operations Research* 29(6):667–684.
- De Silva J, Gales B, Kagy B, Offner D (2018) An analysis of a fair division protocol for drawing legislative districts. URL <http://dx.doi.org/10.48550/ARXIV.1811.05705>.
- DeFord D, Dhamankar N, Duchin M, Gupta V, McPike M, Schoenbach G, Sim KW (2020) Implementing partisan symmetry: Problems and paradoxes. URL <http://dx.doi.org/10.48550/ARXIV.2008.06930>.
- DeFord D, Duchin M (2019) Redistricting reform in Virginia: Districting criteria in context. *Virginia Policy Review* 12(2):120–146.
- DeFord D, Duchin M, Solomon J (2019) Recombination: A family of markov chains for redistricting. URL <http://dx.doi.org/10.48550/ARXIV.1911.05725>.
- Drexel A, Haase K (1999) Fast approximation methods for sales force deployment. *Management Science* 45(10):1307–1323.
- Duchin M (2018) Gerrymandering metrics: How to measure? What’s the baseline? URL <http://dx.doi.org/10.48550/ARXIV.1801.02064>.
- Edmonds J (1965) Paths, trees, and flowers. *Canadian Journal of Mathematics* 17:449–467.
- Eppstein D (2003) Find maximum cardinality matching in general undirected graph. URL <https://www.ics.uci.edu/~eppstein/PADS/CardinalityMatching.py>, [Accessed: 5-May-2020].
- ESRI (2021) Efficiency gap (measure of partisan bias in redistricting) by state, 2016. URL <https://www.arcgis.com/home/item.html?id=563f61cd84a344deb707eed3e0b258bd>, [Accessed: 27-Oct-2021].
- Fifield B, Higgins M, Imai K, Tarr A (2015) A new automated redistricting simulator using Markov chain Monte Carlo. *Working Paper, Princeton University, Princeton, NJ* URL <http://imai.princeton.edu/research/files/redist.pdf>, [Accessed: 25-Dec-2020].
- Friedman JN, Holden RT (2009) The rising incumbent reelection rate: What’s gerrymandering got to do with it? *The Journal of Politics* 71(2):593–611.

- Gabow HN, Tarjan RE (1988) Algorithms for two bottleneck optimization problems. *Journal of Algorithms* 9(3):411–417.
- Garey MR, Johnson DS (1978) “Strong” NP-completeness results: Motivation, examples, and implications. *Journal of the ACM (JACM)* 25(3):499–508.
- Garfinkel RS, Nemhauser GL (1970) Optimal political districting by implicit enumeration techniques. *Management Science* 16(8):B–495.
- Gentry S, Chow E, Massie A, Segev D (2015) Gerrymandering for justice: Redistricting U.S. liver allocation. *Interfaces* 45(5):462–480.
- Gill v Whitford (2018) *Gill v. Whitford* (Case 16-1161, U.S. Supreme Court).
- Grimstad B, Knudsen BR (2020) Mathematical programming formulations for piecewise polynomial functions. *Journal of Global Optimization* 77(3):455–486.
- Grofman B (1983) Measures of bias and proportionality in seats-votes relationships. *Political Methodology* 295–327.
- Grofman B, Cervas JR (2020) The Terminology of Districting. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3540444.
- Grofman B, King G (2007) The future of partisan symmetry as a judicial test for partisan gerrymandering after LULAC v. Perry. *Election Law Journal* 6(1):2–35.
- Haase K, Müller S (2014) Upper and lower bounds for the sales force deployment problem with explicit contiguity constraints. *European Journal of Operational Research* 237(2):677–689.
- Hendrickson B, Leland RW (1995) A multi-level algorithm for partitioning graphs. *Proceedings of the 1995 ACM/IEEE Conference on Supercomputing* 95(28).
- Hess S, Weaver J, Siegfeldt H, Whelan J, Zitlau P (1965) Nonpartisan political redistricting by computer. *Operations Research* 13(6):998–1006.
- Jackman S (1994) Measuring electoral bias: Australia, 1949–93. *British Journal of Political Science* 24(3):319–357.
- Jones PE (2013) The effect of political competition on democratic accountability. *Political Behavior* 35(3):481–515.
- Joshi D, Soh J, Samal A (2012) Redistricting using constrained polygonal clustering. *IEEE Transactions on Knowledge and Data Engineering* 24(11):2065–2079.
- Kalcsics J, Ríos-Mercado RZ (2019) Districting problems. *Location Science* 705–743.
- Karypis G, Kumar V (1998) A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on Scientific Computing* 20(1):359–392.
- Katz JN, King G, Rosenblatt E (2020) Theoretical foundations and empirical evaluations of partisan fairness in district-based democracies. *American Political Science Review* 114(1):164–178.

- Kennedy L, Corriher B, Root D (2016) Redistricting and Representation. URL <https://www.americanprogress.org/issues/democracy/reports/2016/12/05/294272/redistricting-and-representation/>, [Accessed: 27-Oct-2021].
- Kernighan BW, Lin S (1970) An efficient heuristic procedure for partitioning graphs. *The Bell System Technical Journal* 49(2):291–307.
- King DM, Jacobson SH, Sewell EC (2015) Efficient geo-graph contiguity and hole algorithms for geographic zoning and dynamic plane graph partitioning. *Mathematical Programming* 149(1-2):425.
- King G, Grofman B, Gelman A, Katz JN (2006) Brief of Amici Curiae Professors Gary King, Bernard Grofman, Andrew Gelman, and Jonathan N. Katz, in Support of Neither Party.
- Landau Z, Su FE (2014) Fair division and redistricting. *The Mathematics of Decisions, Elections, and Games. American Mathematical Society* 17–36.
- Lavrov M (2019) Is there a polynomial time algorithm for min max maximal matching problem? Mathematics Stack Exchange, URL <https://math.stackexchange.com/q/3456687>, [Accessed: 30-Nov-2019].
- Liu YY, Cho WKT, Wang S (2016) PEAR: A massively parallel evolutionary computation approach for political redistricting optimization and analysis. *Swarm and Evolutionary Computation* 30:78–92.
- Ludden IG, Swamy R, King DM, Jacobson SH (2022) A bisection protocol for political redistricting. Technical report, University of Illinois at Urbana-Champaign.
- LULAC v Perry (2006) *LULAC v. Perry* (548 U.S. 399, U.S. Supreme Court).
- Magleby DB, Mosesson DB (2018) A new approach for developing neutral redistricting plans. *Political Analysis* 26(2):147–167.
- Mann TE (2005) Redistricting reform. *The National Voter* URL <https://www.brookings.edu/articles/redistricting-reform/>, [Accessed: 25-Dec-2020].
- McCarty N, Poole KT, Rosenthal H (2009) Does gerrymandering cause polarization? *American Journal of Political Science* 53(3):666–680.
- McDonald MP (2006) Drawing the line on district competition. *PS: Political Science & Politics* 39(1):91–94.
- McGovern T (2017) United States General Election Presidential Results by County from 2008 to 2016. URL https://github.com/tonmcg/US_County_Level_Election_Results_08-16, [Accessed: 25-Dec-2020].
- Megiddo N, Supowit KJ (1984) On the complexity of some common geometric location problems. *SIAM Journal on Computing* 13(1):182–196.
- Mehrotra A, Johnson E, Nemhauser G (1998) An optimization based heuristic for political districting. *Management Science* 44(8):1100–1114.
- MGGG (2021) The geospatial toolkit for redistricting data. URL <https://github.com/mggg/maup>, [Accessed: 24-Oct-2021].

- Micali S, Vazirani VV (1980) An $O(\sqrt{|V|}|E|)$ algorithm for finding maximum matching in general graphs. *21st Annual Symposium on Foundations of Computer Science (SFCS 1980)*, 17–27 (IEEE).
- National Conference of State Legislatures (2020) Redistricting criteria. URL <https://www.ncsl.org/research/redistricting/redistricting-criteria.aspx>, [Accessed: 25-Dec-2020].
- Oehrlein J, Haunert JH (2017) A cutting-plane method for contiguity-constrained spatial aggregation. *Journal of Spatial Information Science* 2017(15):89–120.
- Office of the Governor (2019) *Executive Order 66: Relating to creating the people’s maps commission* (The State of Wisconsin).
- Pegden W, Procaccia AD, Yu D (2017) A partisan districting protocol with provably nonpartisan outcomes. URL <http://dx.doi.org/10.48550/ARXIV.1710.08781>.
- Ricca F, Scozzari A, Simeone B (2008) Weighted Voronoi region algorithms for political districting. *Mathematical and Computer Modelling* 48(9-10):1468–1477.
- Ricca F, Scozzari A, Simeone B (2013) Political districting: From classical models to recent approaches. *Annals of Operations Research* 204(1):271.
- Ricca F, Simeone B (2008) Local search algorithms for political districting. *European Journal of Operational Research* 189(3):1409–1426.
- Royden L, Li M (2017) *Extreme maps* (Brennan Center for Justice at New York University School of Law), URL https://www.brennancenter.org/sites/default/files/publications/Extreme%20Maps%205.16_0.pdf, [Accessed: 2-May-2018].
- Rubin PA (2015) Optimizing part of the objective function II. URL <https://orinanobworld.blogspot.com/2015/08/optimizingpartoftheobjectivefunction-ii.html>, [Accessed: 27-Oct-2021].
- Salazar-Aguilar MA, Ríos-Mercado RZ, González-Velarde JL (2011) A bi-objective programming model for designing compact and balanced territories in commercial districting. *Transportation Research Part C: Emerging Technologies* 19(5):885–895.
- Shirabe T (2009) Districting modeling with exact contiguity constraints. *Environment and Planning B: Planning and Design* 36(6):1053–1066.
- Silverberg K (1995) Illegitimacy of the incumbent gerrymander. *Texas Law Review* 74:913.
- Stephanopoulos NO, McGhee EM (2015) Partisan gerrymandering and the efficiency gap. *The University of Chicago Law Review* 831–900.
- Stephanopoulos NO, McGhee EM (2018) The measure of a metric the debate over quantifying partisan gerrymandering. *Stanford Law Review* 70(5).
- Swamy R (2022) Code Companion for Multi-objective Optimization for Politically Fair Districting: A Scalable Multilevel Approach. URL <https://github.com/rahul-swamy/Multi-Objective-Politically-Fair-Districting>.

- Tapp K (2018) Measuring political gerrymandering. URL <http://dx.doi.org/10.48550/ARXIV.1801.02541>.
- The Guardian (2018) U.S. 2012 election data. URL <https://www.theguardian.com/news/datablog/2012/nov/07/us-2012-election-county-results-download>, [Accessed: 5-May-2018].
- Thoreson JD, Liittschwager JM (1967) Computers in behavioral science. Legislative districting by computer simulation. *Behavioral Science* 12(3):237–247.
- US Census Bureau (2010) Census block tallies by state or state equivalent. URL <https://www.census.gov/data.html>, [Accessed: 27-Dec-2017].
- Validi H, Buchanan A, Lykhovyd E (2021) Imposing contiguity constraints in political districting models. *Operations Research* 70(2):867–892.
- Veomett E (2018) Efficiency gap, voter turnout, and the efficiency principle. *Election Law Journal: Rules, Politics, and Policy* 17(4):249–263.
- Vickrey W (1961) On the prevention of gerrymandering. *Political Science Quarterly* 76(1):105–110.
- Vielma JP, Ahmed S, Nemhauser G (2010) Mixed-integer models for nonseparable piecewise-linear optimization: Unifying framework and extensions. *Operations Research* 58(2):303–315.
- Walshaw C (2000) A multilevel algorithm for force-directed graph drawing. *International Symposium on Graph Drawing*, 171–182 (Springer).
- Walshaw C (2002) A multilevel approach to the travelling salesman problem. *Operations Research* 50(5):862–877.
- Warrington GS (2018) Quantifying gerrymandering using the vote distribution. *Election Law Journal* 17(1):39–57.
- Wesberry et al v Sanders, Governor of Georgia, et al (1964) *Wesberry et al. v. Sanders, Governor of Georgia, et al.* (Case 376 U.S. 1., No. 22., U.S. Supreme Court).
- Young HP (1988) Measuring the compactness of legislative districts. *Legislative Studies Quarterly* 105–115.

E-Companion

The following sections are part of the e-companion.

EC.1. Proof for Theorem 1

THEOREM 1 *For a general graph $G = (V, E)$, unit populations $\{p_i, p_i^A, p_i^B\}_{i \in V}$ and number of districts $K \geq 2$, the Aggregate Competitive Districting Problem (ACDP) with or without the population balance constraints is NP-complete if $K = 2$ and is strongly NP-complete if $K \geq 3$.*

Proof of Theorem 1. The decision version of ACDP is defined as follows. The inputs are a graph $G = (V, E)$, unit populations $\{p_i, p_i^A, p_i^B\}_{i \in V}$, a balance threshold $\tau \geq 0$, number of districts $K \geq 2$, and a target $W \geq 0$. Then, the decision version of ACDP without population balance constraints asks if there exists a feasible district plan $z : V \rightarrow [K]$ such that

$$\max_{k \in [K]} |P_k^A(z) - P_k^B(z)| \leq W, \quad (\text{EC.1})$$

where for each district $k \in [K]$ and party $l \in \{A, B\}$, $P_k^l(z) = \sum_{i \in V: z(i)=k} p_i^l$. Here, a feasible district plan is one that satisfies the MIP constraints (3)-(9). For ACDP with population balance constraints, a feasible district plan also satisfies constraints (2) for a given $\tau \geq 0$. Given a district plan, since feasibility for ACDP with or without population balance constraints can be verified in polynomial time using the constraints (2)-(9) and (EC.1), it is clear that ACDP with or without population balance constraints is in NP.

It is now shown that ACDP is NP-complete for $K = 2$ districts, and is strongly NP-complete for $K \geq 3$ districts. A polynomial-time reduction is shown from the NP-complete *2-partitioning problem* for $K = 2$ and from the strongly NP-complete *3-partitioning problem* for $K \geq 3$ (Garey and Johnson 1978). The two reductions are first shown for ACDP without the population balance constraints. Then, it is shown how a τ value can be chosen so that the reductions hold when the population balance constraints are added.

For $K = 2$ without population balance:

The 2-partitioning problem (or 2-PART) is defined as follows. For a given set of integers $V' = \{a_1, a_2, \dots, a_n\}$ and a target $W' \geq 0$, 2-PART divides V' into two partitions $\{S_1, S_2\}$ such that $\max_{k=1,2} |\sum_{i \in S(k)} a_i| \leq W'$. For a given instance of 2-PART (V', W') , an instance of ACDP can be constructed such that solving the ACDP is equivalent to solving 2-PART. The set of units V is chosen to be the indices of the set of integers $[n]$, the number of districts K is equal to 2, and the adjacency graph G is constructed to be a complete graph. For every unit $i \in V$, the unit populations (p_i, p_i^A, p_i^B) are chosen to be $(a_i, a_i, 0)$. The target W is chosen to be W' . To show that the two

problems are equivalent, it shown that the constructed instance of ACDP is a “yes” instance if and only if the given 2-PART instance is a “yes” instance.

First, let the instance of 2-PART be a “yes” instance, i.e., there exists a partition $\{S_1, S_2\}$ of V' such that $\max_{k=1,2} |\sum_{i \in S_k} a_i| \leq W'$. Consider a district plan z constructed by assigning units in V to the two districts according to the corresponding assignment of integers in the two sets S_1 and S_2 . Then, for every district $k = 1, 2$,

$$|P_k^A(z) - P_k^B(z)| = \left| \sum_{i \in V: z(i)=k} (p_i^A - p_i^B) \right| = \left| \sum_{i \in V'_k} (a_i - 0) \right| \leq W'.$$

Since $W = W'$, the district plan z satisfies constraint (EC.1). Note that since G is a complete graph, z satisfies the contiguity constraints. Hence, the constructed instance to ACDP is a “yes” instance.

Second, let the constructed instance of ACDP be a “yes” instance, i.e., there exists a district plan z that satisfies constraints (3)-(9) and (EC.1). Consider a partition of V' into 2 partitions $\{S_1, S_2\}$ constructed according to the assignment of the units in $V = V'$ to the $K = 2$ districts. For every $k = 1, 2$, as determined by z , $|\sum_{i \in S(k)} a_i| = |\sum_{i \in V: z(i)=k} (p_i^A - p_i^B)| \leq W$. Since $W = W'$, the partition $\{S_1, S_2\}$ satisfies $\max_{k=1,2} |\sum_{i \in S(k)} a_i| \leq W'$. Hence, the instance of 2-PART is a “yes” instance.

For $K \geq 3$ without population balance:

The 3-partitioning problem (or 3-PART) is defined as follows. For a given set of integers $V' = \{a_1, a_2, \dots, a_n\}$ and a target $W' \geq 0$, 3-PART divides V' into three partitions $\{S_1, S_2, S_3\}$ such that $\max_{k=1,2,3} |\sum_{i \in S(k)} a_i| \leq W'$. For a given instance of 3-PART (V', W') , an instance of ACDP can be constructed such that solving the ACDP is equivalent to solving 3-PART. The number of districts is chosen to be $K \geq 3$. The set of units V is chosen to be the indices of the set of integers in V' padded with $K - 3$ additional units, i.e., $V = [n] \cup \{n+1, n+2, \dots, n+K-3\}$. The adjacency graph G is constructed to be a complete graph. For every unit $i \in V$, the unit populations (p_i, p_i^A, p_i^B) are chosen to be $(a_i, a_i, 0)$ if $i \in [n]$, or $(W', W', 0)$ if $n+1 \leq i \leq n+K-3$. The target W is chosen to be W' . To show that the two problems are equivalent, it is to be shown that an instance of ACDP is a “yes” instance if and only if the constructed 3-PART instance is a “yes” instance.

First, let the instance of 3-PART be a “yes” instance, i.e., there exists a partition $\{S_1, S_2, S_3\}$ of V' such that $\max_{k=1,2,3} |\sum_{i \in S(k)} a_i| \leq W'$. Consider a district plan z constructed by assigning units in $[n]$ to districts 1, 2 and 3 according to the corresponding assignment of integers in the three sets S_1, S_2, S_3 . Each unit in $\{n+1, n+2, \dots, n+K-3\}$ is assigned to occupy a single district among the remaining districts in $[K] \setminus [3]$. Then, for every district $k = 1, 2, 3$,

$$|P_k^A(z) - P_k^B(z)| = \left| \sum_{i \in V: z(i)=k} (p_i^A - p_i^B) \right| = \left| \sum_{i \in V'_k} (a_i - 0) \right| \leq W'.$$

For every district $k \in K \setminus [3]$,

$$|P_k^A(z) - P_k^B(z)| = \left| \sum_{i \in V: z(i)=k} (p_i^A - p_i^B) \right| = |(W' - 0)| = W'.$$

Since $W = W'$, the district plan z satisfies constraint (EC.1). Note that since G is a complete graph, z satisfies the contiguity constraints. Hence, the constructed instance to ACDP is a “yes” instance.

Second, let the instance of ACDP be a “yes” instance, i.e., there exists a district plan z that satisfies constraints (3)-(9) and (EC.1). Note that each of the $K - 3$ padded units will have to be assigned to its own district in order to satisfy constraint (EC.1). These units will occupy $K - 3$ of the K districts. A “yes” to 3-PART can be constructed using the units in $[n]$ to the remaining three districts; suppose, without loss of generality, that the three remaining districts correspond to districts $k = 1, 2, 3$. Consider a partition of V' into 3 partitions, $\{S_1, S_2, S_3\}$, constructed according to the assignment of the units in $[n]$ to the three districts. For every $k = 1, 2, 3$, $|\sum_{i \in S(k)} a_i| = |\sum_{i \in V: z(i)=k} (p_i^A - p_i^B)| \leq W$. Since $W = W'$, the partition $\{S_1, S_2, S_3\}$ satisfies $\max_{k=1,2,3} |\sum_{i \in S(k)} a_i| \leq W'$. Hence, the instance of 3-PART is a “yes” instance.

Inclusion of population balance constraints:

The population balance constraints (2) require that the population in every district is in the range $[\bar{P}(1 - \tau), \bar{P}(1 + \tau)]$, where $\bar{P} := (\sum_{i \in V} p_i)/K$. When these constraints are included in ACDP, it is needed to show that an instance of ACDP can be constructed for some $\tau \geq 0$ such that the two reductions in this proof hold true. Let $\tau = K - 1$. Then, the range of allowable district populations is $[\bar{P}(2 - K), \bar{P}K]$ which is equivalent to $[0, \sum_{i \in V} p_i]$ for $K \geq 2$ since unit populations are non-negative. Then, every district plan will satisfy the population balance constraints, and the two reductions in this proof will continue to hold true. \square

EC.2. Proof for Proposition 1

PROPOSITION 1 *Symmetric-DP can be formulated as a Mixed Integer Linear Program (MILP).*

Proof of Proposition 1. To construct an MILP formulation for Symmetric-DP, each of the three constraints (23) - (25) are linearized one at a time in the following three parts.

- (a) First, constraints (23) are linearized for given values of $\{x_{ij}\}_{i,j \in V}$, the decision variables that assign units to district centers. For each unit $i \in V$, the following set of constraints define ω_i to be the vote-share for party A in a district with center i ; 0 if i is not a district center. Constraints (23) are given by,

$$\omega_i = \begin{cases} \frac{\sum_{j \in V} p_j^A x_{ij}}{\sum_{j \in V} (p_j^A + p_j^B) x_{ij}}, & \text{if } x_{ii} = 1 \\ 0, & \text{otherwise} \end{cases} \quad \forall i \in V,$$

The following additional variables are defined.

$$\beta_{ij} = \begin{cases} \omega_i, & \text{if } i \in V \text{ is a district center, and } j \in V \text{ is assigned to that district.} \\ 0, & \text{otherwise.} \end{cases}$$

The first claim is that the following set of constraints linearize constraints (23).

$$\omega_i \leq x_{ii} \quad \forall i \in V, \quad (\text{EC.2})$$

$$\beta_{ij} \leq \omega_i \quad \forall i, j \in V, \quad (\text{EC.3})$$

$$\omega_i - 1 + x_{ij} \leq \beta_{ij} \quad \forall i, j \in V, \quad (\text{EC.4})$$

$$\beta_{ij} \leq x_{ij} \quad \forall i, j \in V, \quad (\text{EC.5})$$

$$\sum_{j \in V} (p_j^A + p_j^B) \beta_{ij} = \sum_{j \in V} p_j^A x_{ij} \quad \forall i \in V, \quad (\text{EC.6})$$

$$\beta_{ij}, \omega_i \geq 0 \quad \forall i, j \in V. \quad (\text{EC.7})$$

For every $i \in V$, constraint (EC.2) ensures that if i is not a district center, then $\omega_i = 0$. For every $i, j \in V$, constraints (EC.3)- (EC.5) together establish the relationship $\beta_{ij} = \omega_i \cdot x_{ij}$. When $x_{ij} = 0$, constraint (EC.5) ensures that $\beta_{ij} = 0$, and when $x_{ij} = 1$, constraints (EC.3) and (EC.4) ensure that $\beta_{ij} = \omega_i$. Given that $\beta_{ij} = \omega_i \cdot x_{ij}$, it can be seen that constraint (EC.6) is equivalent to the non-linear constraint (23). Constraints (EC.7) establish the non-negativity of the variables.

- (b) Second, constraints (24) are linearized. Here, the values for $\{\omega_i\}_{i \in V}$ defined by constraints (EC.2)-(EC.7) are used to define α_k , the k -th largest district vote-share for $k \in [K]$ among the set of all district vote-shares. Constraints (24) are given by,

$$\alpha_k = k \max_{i \in V} \{\omega_i\} \quad \forall k \in [K],$$

where the k max function returns the k -th largest value of a given set. To linearize this constraint, for each $i \in V$ and $k \in [K]$, additional variables $\delta_{i,k}$ and $\gamma_{i,k}$ are now defined. Some additional notation is first introduced for each $k \in [K]$. Let $\omega_{(k)}$ denote the value of the k -th largest element in the set $\{\omega_i\}_{i \in V}$. For each $k \in [K]$, let $S^+(k), S(k), S(k)^- \subseteq V$ be subsets of units that partition V which are defined as $S^+(k) := \{i \in V : \omega_i > \omega_{(k)}\}$, $S(k) := \{i \in V : \omega_i = \omega_{(k)}\}$, and $S(k)^- := \{i \in V : \omega_i < \omega_{(k)}\}$. It can be observed that $|S^+(k)| < k$ and $|S^-(k)| < |V| - k + 1$ since $|S(k)| \geq 1$ by definition. Among the units in $S(k) \cup S^+(k)$, choose an arbitrary subset of units $S'(k) \subseteq S(k) \cup S^+(k)$ such that $|S'(k)| = k$. Such an $S'(k)$ is possible since $|S(k) \cup S^+(k)| = |V| - |S^-(k)| > |V| - (|V| - k + 1) = k - 1$, and hence $|S(k) \cup S^+(k)| \geq k$ since the size of the set is an integer. Further, among the units in $S(k) \cup S^-(k)$, choose an arbitrary subset of units $S''(k) \subseteq S(k) \cup S^-(k)$ such that $|S''(k)| = |V| - k + 1$. Such an $S''(k)$

is possible since $|S(k) \cup S^-(k)| = |V| - |S^+(k)| > |V| - k$, and hence $|S(k) \cup S^-(k)| \geq |V| - k + 1$ since the size of the set is an integer. Based on the constructed sets $S'(k)$ and $S''(k)$, the $\delta_{i,k}$ and $\gamma_{i,k}$ values are defined as follows.

$$\delta_{i,k} = \begin{cases} 1, & \text{if } i \in S'(k), \\ 0, & \text{otherwise.} \end{cases} \quad (\text{EC.8})$$

$$\gamma_{i,k} = \begin{cases} 1, & \text{if } i \in S''(k), \\ 0, & \text{otherwise.} \end{cases} \quad (\text{EC.9})$$

For each $k \in [K]$, ensuring that α_k is the k -th largest element in the set $\{\omega_i\}_{i \in V}$ is equivalent to ensuring that α_k is less than or equal to at least k elements in the set, and that α_k is greater than or equal to at least $|V| - k + 1$ elements in the set. The second claim in this proof is that the following set of constraints linearize constraints (24).

$$\alpha_k \leq \omega_i + 1 - \delta_{i,k} \quad \forall i \in V, k \in [K] \quad (\text{EC.10})$$

$$\sum_{i \in V} \delta_{i,k} = k \quad \forall k \in [K] \quad (\text{EC.11})$$

$$\omega_i - 1 + \gamma_{i,k} \leq \alpha_k \quad \forall i \in V, k \in [K] \quad (\text{EC.12})$$

$$\sum_{i \in V} \gamma_{i,k} = |V| - k + 1 \quad \forall k \in [K] \quad (\text{EC.13})$$

$$\gamma_{i,k}, \delta_{i,k} \in \{0, 1\} \quad \forall i \in V, k \in [K]. \quad (\text{EC.14})$$

For each $k \in [K]$, constraints (EC.10) ensure that for every $i \in V$, $\delta_{i,k} = 1 \Rightarrow \alpha_k \leq \omega_i$. Constraints (EC.11) ensure that $\delta_{i,k} = 1$ for exactly k elements in $\{\omega_i\}_{i \in V}$. Constraints (EC.12) ensure that for $i \in V$, $\gamma_{i,k} = 1 \Rightarrow \alpha_k \geq \omega_i$. Constraints (EC.13) ensure that $\gamma_{i,k} = 1$ for exactly $|V| - k + 1$ elements in $\{\omega_i\}_{i \in V}$. Constraints (EC.14) define the binary nature of the variables. This linearization is similar to the constraints provided by Rubin (2015), where the goal is to minimize the sum of the k largest elements in a vector.

The correctness of constraints (EC.10)-(EC.14) is now shown. It is required to show that a solution satisfies these constraints if and only if it defines α_k for all $k \in [K]$ to be the k -th largest element in the set $\{\omega_i\}_{i \in V}$.

For the sufficiency condition, consider a solution $(\alpha_k, \delta_{i,k}, \gamma_{i,k})_{i \in V, k \in [K]}$ that satisfies constraints (EC.10)-(EC.14). It is required to show that $\alpha_k = \omega_{(k)}$ for all $k \in [K]$. Suppose for the purpose of contradiction that $\alpha_{k'} \neq \omega_{(k')}$ for some $k' \in [K]$. Then, one of the following two cases holds for that k' :

- *Case 1:* $\alpha_{k'} < \omega_{(k')}$. Any $i \in V$ with $\gamma_{i,k'} = 1$ must have $i \in S_{k'}^-$, since it must have $\omega_i \leq \alpha_{k'} < \omega_{(k')}$. Hence, $\{i \in V \mid \gamma_{i,k'} = 1\} \subseteq S_{k'}^-$ and it can be shown that $\sum_{i \in V} \gamma_{i,k'} \leq \sum_{i \in S_{k'}^-} 1 = |S_{k'}^-| < |V| - k' + 1$, which provides a contradiction with constraint (EC.13).

- *Case 2:* $\alpha_{k'} > \omega_{(k')}$. Any $i \in V$ with $\delta_{i,k'} = 1$ must have $i \in S_{k'}^+$, since it must have $\omega_i \geq \alpha_{k'} > \omega_{(k')}$. Hence, $\{i \in V \mid \delta_{i,k'} = 1\} \subseteq S_{k'}^+$ and it can be shown that $\sum_{i \in V} \delta_{i,k'} \leq \sum_{i \in S_{k'}^+} 1 = |S_{k'}^+| < k'$, which provides a contradiction with constraint (EC.11).

Hence, either case contradicts the assumption that $\alpha_{k'} \neq \omega_{(k')}$ for some $k' \in [K]$. Therefore, it can be concluded that $\alpha_k = \omega_{(k)}$ for all $k \in [K]$.

For the necessary condition, for each $k \in [K]$, it is required to show that a solution $(\alpha_k, \delta_{i,k}, \gamma_{i,k})_{i \in V}$ that defines α_k to be the k -th largest element in $\{\omega_i\}_{i \in V}$ will satisfy constraints (EC.10)-(EC.14). In such a solution, $\alpha_k = \omega_{(k)}$ by definition. The $\delta_{i,k}$ and $\gamma_{i,k}$ values for each $i \in V$ are derived from (EC.8) and (EC.9). First, consider constraint (EC.10). For all $i \in S'(k)$, since $\delta_{i,k} = 1$ and $\alpha_k \leq \omega_i$ by the definitions of $S(k)$ and $S^+(k)$ (whose union is a superset of $S'(k)$), this constraint is satisfied. For all $i \in V \setminus S'(k)$, since $\delta_{i,k} = 0$ and $\alpha_k \leq 1$ by constraint (27), this constraint is satisfied. Hence, this solution satisfies constraint (EC.10). Further, since

$$\sum_{i \in V} \delta_{i,k} = \sum_{i \in S'(k)} \delta_{i,k} + \sum_{i \in V \setminus S'(k)} \delta_{i,k} \quad (\text{EC.15})$$

$$= k + 0 = k, \quad (\text{EC.16})$$

this solution satisfies constraint (EC.11). Similarly, using the $\{\gamma_{i,k}\}_{i \in V}$ values from (EC.9), it can be shown that the solution satisfies constraints (EC.12) and (EC.13). The proof proceeds similar to the proof for the solution satisfying constraints (EC.10) and (EC.11), and is hence omitted. Finally, the $\{\delta_{i,k}, \gamma_{i,k}\}_{i \in V}$ values satisfy constraint (EC.14) by definition.




Figure EC.1 District votes-shares in K districts sorted in a non-increasing order, $\{\alpha_k\}_{k \in [K]}$.

- (c) Third, constraints (25) are linearized. Here, for each $k \in [K]$, the value of α_k defined by constraints (EC.10)-(EC.14) is used to define $\mu_{k,m}$ to be the minimum average vote-share for A to win exactly k districts. To achieve this, the vote-share α_m in each district $m \in [K]$ is decreased by a value of $\alpha_k - \frac{1}{2}$, which is the fraction of voters needed to just flip district k ;

note that this will increase the vote-share if $\alpha_k < 1/2$. The intuition behind this is illustrated in Figure EC.1. Constraints (25) are given by,

$$\mu_{km} = \begin{cases} 0, & \text{if } \alpha_m + \frac{1}{2} - \alpha_k \leq 0 \\ \alpha_m + \frac{1}{2} - \alpha_k, & \text{if } \alpha_m + \frac{1}{2} - \alpha_k \in (0, 1) \\ 1, & \text{if } \alpha_m + \frac{1}{2} - \alpha_k \geq 1 \end{cases} \quad \forall k, m \in [K].$$

It can be seen that for each $k, m \in [K]$, μ_{km} is a non-convex piecewise linear function of $(\alpha_m + 1/2 - \alpha_k)$ with 2 breakpoints to ensure that μ_{km} does not leave the $[0, 1]$ interval. Linearizing a piecewise linear function in an MIP has been well studied (Grimstad and Knudsen 2020), and this paper presents a concise version of the standard big-M constraints. The concision comes from using only 2 binary variables for each $k, m \in [K]$ instead of the standard 3 variables needed when the piecewise linear function has 2 breakpoints. See Vielma et al. (2010) for convex combination models to linearize non-convex piecewise linear functions.

The linearization is now described. For each $k, m \in [K]$, additional binary variables Ω_{km} and Ω'_{km} are defined as follows

$$\begin{aligned} \Omega_{km} &= \begin{cases} 1, & \text{if } \alpha_m + \frac{1}{2} - \alpha_k > 0. \\ 0, & \text{if } \alpha_m + \frac{1}{2} - \alpha_k < 0. \end{cases} \\ \Omega'_{km} &= \begin{cases} 1, & \text{if } \alpha_m + \frac{1}{2} - \alpha_k < 1. \\ 0, & \text{if } \alpha_m + \frac{1}{2} - \alpha_k > 1. \end{cases} \end{aligned}$$

Here, Ω_{km} (Ω'_{km}) arbitrarily takes the value of either 0 or 1 if $\alpha_m + \frac{1}{2} - \alpha_k = 0$ (1). The first idea in the linearization is to use big-M constraints to ensure that Ω_{km} and Ω'_{km} are defined with respect to the value of $(\alpha_m + \frac{1}{2} - \alpha_k)$. The second idea is to define μ_{km} using Ω_{km} , Ω'_{km} and $(\alpha_m + \frac{1}{2} - \alpha_k)$. Note that for every $k, m \in [K]$, the expression $(\alpha_m + \frac{1}{2} - \alpha_k)$ ranges between $-\frac{1}{2}$ and $\frac{3}{2}$. The third claim is that the following set of constraints linearize constraints (25). The linearized constraints are given by,

$$\alpha_m + \frac{1}{2} - \alpha_k \leq \frac{3}{2}\Omega_{km} \quad \forall k, m \in [K], \quad (\text{EC.17})$$

$$\alpha_m + \frac{1}{2} - \alpha_k \geq \mu_{km} + \frac{3}{2}(\Omega_{km} - 1) \quad \forall k, m \in [K], \quad (\text{EC.18})$$

$$\alpha_m + \frac{1}{2} - \alpha_k \geq -\frac{3}{2}\Omega'_{km} + 1 \quad \forall k, m \in [K], \quad (\text{EC.19})$$

$$\alpha_m + \frac{1}{2} - \alpha_k \leq \frac{3}{2}(1 - \Omega'_{km}) + \mu_{km} \quad \forall k, m \in [K], \quad (\text{EC.20})$$

$$-\Omega'_{km} + 1 \leq \mu_{km} \leq \Omega_{km} \quad \forall k, m \in [K], \quad (\text{EC.21})$$

$$\Omega_{km}, \Omega'_{km} \in \{0, 1\} \quad k, m \in [K]. \quad (\text{EC.22})$$

For each $k, m \in [K]$, constraint (EC.22) establishes the binary domain of both Ω_{km} and Ω'_{km} . To prove that constraints (EC.17) - (EC.22) linearize constraint (25), there are two parts to the proof. First, it is shown that for a given value of $(\alpha_m + \frac{1}{2} - \alpha_k)$, $\mu_{k,m}$ is correctly defined through the variables $\Omega_{k,m}$ and $\Omega'_{k,m}$. Consider the following five cases for $(\alpha_m + \frac{1}{2} - \alpha_k)$:

- *Case 1:* when $\alpha_m + \frac{1}{2} - \alpha_k < 0$, then $\Omega_{km} = 0$ by combining constraints (EC.18) and (27), and $\Omega'_{km} = 1$ by constraint (EC.19). Substituting these values of Ω_{km} and Ω'_{km} into constraint (EC.21) gives $\mu_{km} = 0$. These values also satisfy constraints (EC.17) and (EC.20).
- *Case 2:* when $\alpha_m + \frac{1}{2} - \alpha_k = 0$, then $\Omega'_{km} = 1$ by constraint (EC.19). There are two subcases based on the value of Ω_{km} :
 - (a) if $\Omega_{km} = 1$, then $\mu_{km} = 0$ by combining constraints (EC.18) and (27). Constraints (EC.17), (EC.20), and (EC.21) are also satisfied.
 - (b) if $\Omega_{km} = 0$, then $\mu_{km} = 0$ by constraint (EC.21). Constraints (EC.17), (EC.18), and (EC.20) are also satisfied.
- *Case 3:* when $\alpha_m + \frac{1}{2} - \alpha_k \in (0, 1)$, then $\Omega_{km} = 1$ and $\Omega'_{km} = 1$ by constraints (EC.17) and (EC.19), respectively. Substituting these values of Ω_{km} and Ω'_{km} into constraints (EC.18) and (EC.20) gives $\mu_{km} = \alpha_m + \frac{1}{2} - \alpha_k$. These values also satisfy constraint (EC.21).
- *Case 4:* when $\alpha_m + \frac{1}{2} - \alpha_k = 1$, then $\Omega_{km} = 1$ by constraint (EC.17). There are two subcases based on the value of Ω'_{km} :
 - (a) if $\Omega'_{km} = 1$, then $\mu_{km} = 1$ by combining constraints (EC.20) and (27). Constraints (EC.18), (EC.19), and (EC.21) are also satisfied.
 - (b) if $\Omega'_{km} = 0$, then $\mu_{km} = 1$ by constraint (EC.21). Constraints (EC.18), (EC.19), and (EC.20) are also satisfied.
- *Case 5:* when $\alpha_m + \frac{1}{2} - \alpha_k \geq 1$, then $\Omega_{km} = 1$ by constraint (EC.17) and $\Omega'_{km} = 0$ by combining constraints (EC.20) and (27). Substituting these values of Ω_{km} and Ω'_{km} into constraint (EC.21) gives $\mu_{km} = 1$. These values also satisfy constraints (EC.18) and (EC.19).

Therefore, regardless of the value of $\alpha_m + \frac{1}{2} - \alpha_k$, the value of μ_{km} is determined as defined in constraint (25).

Second, it is shown that for given values of Ω_{km} and Ω'_{km} , the values of μ_{km} and $(\alpha_m + \frac{1}{2} - \alpha_k)$ are correctly determined. There are 4 cases:

- *Case 1:* $\Omega_{km} = \Omega'_{km} = 0$ is not possible by constraint (EC.21).
- *Case 2:* when $\Omega_{km} = 0$ and $\Omega'_{km} = 1$, then $\mu_{km} = 0$ by (EC.21), and $(\alpha_m + \frac{1}{2} - \alpha_k) \leq 0$ by constraint (EC.17) (Other constraints are also satisfied by noting that $\alpha_m + \frac{1}{2} - \alpha_k \geq -1/2$).
- *Case 3:* when $\Omega_{km} = 1$ and $\Omega'_{km} = 0$, then $\mu_{km} = 1$ by (EC.21), and $(\alpha_m + \frac{1}{2} - \alpha_k) \geq 1$ by constraint (EC.19) (Other constraints are also satisfied by noting that $\alpha_m + \frac{1}{2} - \alpha_k \leq 3/2$).
- *Case 4:* when $\Omega_{km} = \Omega'_{km} = 1$, then $\mu_{km} = \alpha_m + \frac{1}{2} - \alpha_k$ by combining constraints (EC.18) and (EC.20). Other constraints are also satisfied. \square

EC.3. Approximating the Minmax Maximal Matching Problem (MLP)

This section presents an approximation result for the *minmax maximal matching problem* (MLP). Given a connected undirected graph $G = (V, E)$ and edge weights u_e for all $e \in E$, MLP finds a maximal matching M in G that minimizes the maximum edge weight $\max_{e \in M} u_e$. MLP is NP-complete via a reduction from the dominating set problem (Lavrov 2019). This paper uses a greedy algorithm to find a feasible solution to MLP; this matching can then be used to merge in the coarsening procedure in Section 4.1.2. For this greedy algorithm, the edge weights for the population-based coarsening is given by $u_{(i,j)} = p_i + p_j$ for every edge $(i, j) \in E$, where $p_i \geq 0$ is the population in every unit $i \in V$.

This section shows that any arbitrary maximal matching (and by extension, the solution found by the greedy algorithm in Section 4.1.2) has an approximation factor that depends on the relative populations of neighboring units in G . To formalize this, let $\rho := \max_{(i,j) \in E} \max\{p_i/p_j, p_j/p_i\}$ be the largest ratio of populations among neighboring units in G . Clearly, $\rho \geq 1$. In defining ρ , a key assumption is that unit populations are strictly positive. Theorem 2 shows that the maximum edge weight (u') in any arbitrary maximal matching is at the most ρ times the maximum edge weight (u^*) in an optimal maximal matching to MLP. In the special case when all the unit populations are equal, this result is trivially true since all the edge weights are equal, where $\rho = 1$. In this case, every maximal matching is optimal to MLP. Theorem 2 derives an approximation ratio for the general case for any $\rho \geq 1$.

THEOREM 2 *Consider a connected graph $G = (V, E)$, positive unit populations $p_i > 0$ for every unit $i \in V$, and edge weights $u_{ij} = p_i + p_j$ for every edge $(i, j) \in E$. Let M^* be an optimal solution to the minmax maximal matching problem with inputs $(G, \{u_{ij}\}_{(i,j) \in E})$, and let $u^* := \max_{e \in M^*} u_e$ be its maximum edge weight. Let M be an arbitrary maximal matching with maximum edge weight $u' := \max_{e \in M} u_e$. Then, $u'/\rho \leq u^* \leq u'$, where $\rho := \max_{(i,j) \in E} \max\{p_i/p_j, p_j/p_i\}$.*

Proof of Theorem 2. Since M is a maximal matching (i.e., M is a feasible solution to the minmax maximal matching problem), $u^* \leq u'$. The proof for $u^* \geq u'/\rho$ proceeds by contradiction. Assume the contrary, that $u^* < u'/\rho$. Let $e' = (i', j')$ be an edge with maximum weight in M , i.e., $e' := \arg \max_{e \in M} u_e$ and $u_{e'} = u'$. There are two cases:

1. *Case 1:* edge e' is a part of M^* . In this case, $u' \leq u^*$ since u^* is the maximum edge weight in M^* . This is contrary to the assumption that $u^* < u'/\rho$ since $\rho \geq 1$.
2. *Case 2:* edge e' is not a part of M^* . Since M^* is a maximal matching, at least one of i' or j' is a part of a matched edge in M^* that is not e' . Without loss of generality, select i' as a unit

that is part of a matched edge in M^* . Let that edge be (i', j'') for some $j'' \in V$, and its weight is given by,

$$u_{(i', j'')} = p_{i'} + p_{j''} \geq p_{i'} + \frac{p_{i'}}{\rho} = p_{i'} \frac{(1 + \rho)}{\rho}. \quad (\text{EC.23})$$

Since (i', j'') is a part of M^* , and u^* is the maximum edge weight in M^* ,

$$u_{(i', j'')} \leq u^*. \quad (\text{EC.24})$$

From (EC.23) and (EC.24),

$$p_{i'} \frac{(1 + \rho)}{\rho} \leq u^*. \quad (\text{EC.25})$$

On the other hand, the weight of edge (i', j') is given by,

$$u' = p_{i'} + p_{j'} \leq p_{i'} + \rho p_{i'} = (1 + \rho)p_{i'}. \quad (\text{EC.26})$$

From (EC.25) and (EC.26), it is evident that $u' \leq \rho u^*$. This is contrary to the assumption that $u^* < u'/\rho$. \square




Figure EC.2 An example instance of MLP illustrating that Theorem 2 provides a tight approximation; a graph with four units with respective unit populations $1/P$, $1/P$, 1 and P for some $P > 1$. The unit populations and edge weights are depicted at the units and edges, respectively.

It can be shown that the greedy algorithm in Section 4.1.2 provides a tight approximation, i.e, there exists an instance of MLP such that $u'/u^* = \rho$, where u' and u^* are the greedy and optimal solutions, respectively.

EXAMPLE EC.1. Consider an instance of MLP with four units and four edges as depicted in Figure EC.2. The respective unit populations are $1/P$, $1/P$, 1 and P for some $P > 1$. The maximum population ratio among neighboring units is $\rho = P$.

For the instance in Example 1, the optimal maximal matching to MLP selects exactly one edge whose edge weight is $u^* = (P + 1)/P$. On the other hand, the greedy algorithm in Section 4.1.2 first sorts the edges in the non-decreasing order: $\{2/P, (P + 1)/P, (P + 1)/P, P + 1\}$. The maximal matching returned by the greedy algorithm selects the first and the fourth edges with respective weights $2/P$ and $P + 1$. The maximum edge weight for the greedy solution is $u' = P + 1$. Hence, the approximation ratio is $u'/u^* = (P + 1)/((P + 1)/P) = P$, which is equal to ρ . Note that if the unit populations are restricted to integers, setting P to be an integer and multiplying all the unit populations in this example by P would yield an instance with integral unit populations, and the tightness result still holds.

EC.4. Computational Results for Pareto-Frontier

EC.4.1. Efficiency Gap and Compactness

Coarsening	ϕ_{EG}	$\phi_{comp} (\times 10^9)$	Clock time (s)	CPU time (ticks)	# B&C nodes	Opt. gap ($\frac{UB-LB}{UB}$)
ML	0.1561	2.683	416	180,309	112	-
	0.1511	2.683	784	290,876	199	-
	0.0277	2.686	1,161	467,881	987	-
	0.0272	2.688	28,927	9,402,030	7,091	-
	0.0257	2.688	8,480	3,010,334	3,187	-
	0.0234	2.689	1,695	638,511	1,756	-
	0.0224	2.690	1,625	599,558	2,347	-
	0.0220	2.691	12,849	4,160,289	4,860	-
	0.0217	2.695	81,616	24,191,404	16,204	-
	0.0215	2.695	52,968	15,764,469	12,876	-
	0.0198	2.697	13,370	3,801,803	10,131	-
MM	0.0307	3.923	295	134,708	123	-
	0.0290	3.932	7,584	2,452,971	4,200	-
	0.0286	3.937	10,582	3,517,502	5,491	-
	0.0283	3.943	9,689	2,610,137	3,907	-
	0.0281	3.952	23,973	7,019,182	14,997	-
	0.0274	3.954	15,944	5,100,907	7,839	-
	0.0268	3.961	14,622	4,349,295	8,056	-
	0.0255	3.966	3,289	1,093,726	6,296	-
	0.0252	3.973	45,793	14,142,057	23,586	-
	0.0251	3.979	32,342	9,499,810	16,455	-
	0.0250	3.982	5,678	1,828,732	11,261	-
	0.0243	3.983	48,467	15,287,890	32,957	-
	0.0235	3.989	34,041	10,235,375	22,377	-
	0.0190	4.003	39,687	12,684,359	44,313	-

Table EC.1 Computational results from the ϵ -constraint method to obtain the Pareto-frontier between the efficiency gap and compactness, when coarsening using maximum matchings (MM) and maximal matchings (ML). An optimality gap of “-” indicates that the solution is optimal, i.e., the upper bound (UB) and lower bound (LB) on the compactness objective found in the branch-and-cut algorithm are equal.

EC.4.2. Partisan Asymmetry and Compactness

Coarsening	ϕ_{PA}	$\phi_{comp}(\times 10^9)$	Clock time (s)	CPU time (ticks)	# B&C nodes	Opt. gap ($\frac{UB-LB}{UB}$)
ML	0.0633	3.166	21,601*	5,457,399	4,285	0.154
	0.0127	3.277	21,600	6,187,473	1,554	0.184
	0.0020	3.535	21,600	6,141,634	6,796	0.243
	0.0016	3.929	21,600	5,934,803	3,484	0.319
MM	0.0034	4.605	21,608*	6,102,360	14,573	0.148
	0.0024	5.516	21,600	6,034,105	14,503	0.289
	0.0018	5.800	21,600	6,045,607	8,643	0.324
	0.0002	6.729	21,600	5,861,623	4,098	0.420

Table EC.2 Computational results from the ϵ -constraint method to obtain the Pareto-frontier between the partisan asymmetry and compactness objectives, when coarsening using maximum matchings (MM) and maximal matchings (ML). *The solver takes extra time beyond the 6 hour MIP time limit to completely terminate.

EC.4.3. Competitiveness and Compactness

Coarsening	ϕ_{cmpptv}	$\phi_{comp}(\times 10^9)$	Clock time (s)	CPU time (ticks)	# B&C nodes	Opt. gap ($\frac{UB-LB}{UB}$)
ML	0.333	2.691	897	443,614	642	-
	0.313	2.699	590	251,040	1,208	-
	0.303	2.709	651	296,992	1,211	-
	0.290	2.714	475	220,096	426	-
	0.276	2.724	574	278,495	1,608	-
	0.266	2.738	529	243,615	807	-
	0.258	2.746	520	239,000	604	-
	0.248	2.761	693	317,177	2,105	-
	0.239	2.768	755	349,210	2,952	-
	0.231	2.781	825	383,338	2,746	-
	0.223	2.788	972	496,707	548	-
	0.213	2.802	1,639	869,083	1,916	-
	0.163	2.875	1,755	809,647	4,702	-
	0.153	2.895	11,345	3,633,036	7,867	-
	0.138	2.905	1,825	710,101	1,506	-
	0.123	2.923	11,290	3,416,250	2,539	-
	0.111	2.939	16,952	5,092,444	3,451	-
	0.099	2.959	33,652	11,747,083	5,430	-
MM	0.286	3.923	1,094	503,219	918	-
	0.265	3.935	321	156,208	182	-
	0.247	3.965	2,866	1,029,225	3,429	-
	0.246	3.972	11,033	3,425,895	5,827	-
	0.238	3.988	995	417,346	1,199	-
	0.234	3.994	444	203,912	414	-
	0.224	4.018	1,025	463,780	2,001	-
	0.188	4.081	4,212	1,368,974	2,730	-
	0.186	4.088	6,886	2,225,135	4,722	-
	0.171	4.089	3,115	1,095,628	3,175	-
	0.159	4.103	556	253,614	616	-
	0.135	4.131	16,198	5,681,976	9,785	-
	0.131	4.146	2,928	1,161,383	3,175	-
	0.129	4.154	5,835	2,082,055	6,933	-
	0.125	4.155	3,121	1,083,149	5,181	-
	0.121	4.157	2,255	780,248	4,666	-
	0.119	4.164	2,952	1,000,392	4,791	-

Table EC.3 Computational results from the ϵ -constraint method to obtain the Pareto-frontier between competitiveness and compactness objectives, when coarsening using maximum matchings (MM) and maximal matchings (ML). From the 72 solutions obtained when coarsening using ML, a subset of 18 solutions are showcased here for simplicity. This subset is selected by picking every fourth solution in the sorting of the 72 solutions according to their respective competitiveness (ϕ_{cmpptv}) values. An optimality gap of “-” indicates that the solution is optimal, i.e., the upper bound (UB) and lower bound (LB) on the compactness objective found in the branch-and-cut algorithm are equal.