A Bisection Protocol for Political Redistricting

Ian G. Ludden
Department of Computer Science, University of Illinois Urbana-Champaign, iludden2@illinois.edu

Rahul Swamy
Department of Industrial and Enterprise Systems Engineering, University of Illinois Urbana-Champaign, rahulswa@illinois.edu

Douglas M. King
Department of Industrial and Enterprise Systems Engineering, University of Illinois Urbana-Champaign, dmking@illinois.edu

Sheldon H. Jacobson
Department of Computer Science, University of Illinois Urbana-Champaign, shj@illinois.edu

Political redistricting in the United States is the process of drawing congressional and state legislative district boundaries. This work introduces the *bisection protocol*, a two-player zero-sum extensive-form game motivated by political redistricting in which two players alternately divide pieces of a region in half (up to rounding) to obtain a district plan. A subgame perfect Nash equilibrium is presented for the protocol in a relaxed continuous nongeometric (CN) setting, and a recurrence is given for the optimal strategies. The bisection protocol is compared to the recently proposed I-cut-you-freeze protocol across a variety of standard fairness metrics. A hardness result is presented for the bisection protocol in the more realistic discrete geometric setting, along with exact equilibrium strategies for small grid graphs. Since equilibrium computation is intractable for practical instances, heuristics are developed for both protocols which model players' drawing strategies as mixed-integer linear programs. When the heuristics are applied to congressional redistricting in Iowa with counties preserved, both protocols produce district plans that are fairer (according to three popular metrics) than Iowa's 115th congressional districts. Finally, the bisection heuristic is used to generate congressional district plans from census tracts for 18 states, demonstrating its potential for practical use.

Key words: political redistricting, game theory, integer programming

---

1. Introduction

Legislative bodies in the United States must draw electoral district boundaries for state and federal elections every ten years in response to new census data, a process referred to as *redistricting*. *Gerrymandering* is the manipulation of district boundaries for political advantage, a term coined in 1812 when Massachusetts governor Elbridge Gerry approved a district plan featuring a salamander-shaped district (Davis 2017). Gerrymandering typically takes one or both of two forms: *packing*, the creation of some districts in which the opposing party wins by a large margin (and hence wastes surplus votes), or *cracking*, in which the opposing party's voters are divided across many districts to be wasted in lost elections. Several gerrymandering court cases, including *Gill v. Whitford*, *Rucho v. Common Cause*, and *Lamone v. Benisek*, highlight the potential value for mathematical insights into the redistricting process. Judges may use these insights to decide fairly whether district plans are gerrymandered, potentially leading states to consider changing their redistricting procedures.

Since the 1960's, many algorithmic approaches to political redistricting have been proposed, many of which attempt to optimize nonpartisan objectives such as population balance, compactness, and contiguity (Ricca et al. 2013). However, political geography may lead to partisan bias even when objectives are not explicitly political (Chen and Rodden 2013); for example, maximizing compactness tends to disadvantage Democrats in states where their voters are heavily concentrated in cities. Recent approaches have therefore considered objectives that explicitly incorporate the partisan leanings of districts. These recent attempts to combat gerrymandering generally fall into three main categories: the design of redistricting games, or protocols, with theoretical guarantees for certain fairness metrics (e.g., Landau et al. 2009, Pegden et al. 2017, Tucker-Foltz 2019); the optimization of district plans with objectives such as compactness, population balance, and partisan fairness measures (Liu et al. 2016, Cohen-Addad et al. 2018, Levin and Friedler 2019, Gurnee and Shmoys 2021, Validi et al. 2021, Validi and Buchanan 2021, Swamy et al. 2022); and the pursuit of justiciable criteria for identifying gerrymandering (e.g., Grofman and King 2007, Stephanopoulos and McGhee 2015, Duchin 2018, Ramachandran and Gold 2018). This paper contributes to the first.

Three redistricting protocols with a game-theoretic fair division approach have been proposed. Landau et al. (2009) proposes a redistricting protocol in which an independent third party (i.e., an arbiter with no political motives) constructs a nested set of possible splits (i.e., partitions of the state into two contiguous pieces) and, for each split, solicits preferences of which side each party would rather draw. Assuming an additive utility function (e.g., number of districts controlled), this protocol guarantees each party an average or better than average choice among the splits offered, but its reliance on an impartial arbiter is a significant drawback in practice. Pegden et al. (2017) introduces the *I-cut-you-freeze* (ICYF) redistricting protocol inspired by classical results for cake-cutting problems. Players alternate between drawing the remaining district boundaries and choosing one district to *freeze*. While this protocol avoids the need for an independent third party, outcomes are not always symmetric, giving the starting player a substantial advantage in some cases. Moreover, for an n-district state, the protocol requires $n-1$ rounds of drawing district plans and freezing one district, a considerable increase in communication when compared to one party drawing the final plan. Tucker-Foltz (2019) proposes a *cut-and-choose* mechanism in which one party draws a district plan and the other chooses the vote-share margin needed to win a district; districts in which neither party wins are allocated by fair coin tosses. While the theoretical equilibria give each party about the same proportion of seats as they have votes, legislatures and the general public may balk at the suggestion of a party potentially losing a district in which they have a simple majority.

This paper introduces the *bisection protocol*, in which two players (e.g., the Republican and Democratic parties in the U.S.) alternately divide pieces of an n-district state in half (up to rounding) until each piece has a $1/n$ fraction of the total population. Throughout this paper, a *piece* is simply a nonempty (but not necessarily contiguous) subset of the state. The bisection protocol avoids the need for an impartial arbiter (present in the Landau et al. (2009) protocol), requires only $O(\log n)$ rounds (rather than $O(n)$ in the ICYF protocol), and assigns seats using the simple plurality rule (unlike the Tucker-Foltz (2019) protocol). The bisection and ICYF protocols are compared using three fairness metrics: *proportionality*, the notion that seat-shares should match vote-shares; *efficiency gap*, the difference between the number of wasted votes for each party; and *partisan symmetry*, which loosely requires that party outcomes are symmetric with respect to vote-share. Since optimal play of both protocols sometimes involves extreme packing, we constrain the maximum district vote-share to 75% and show equivocal effects on fairness measures.

The relative merit of these and other fairness metrics has been debated (e.g., Lowenstein and Steinberg 1985, Grofman 1985, Niemi 1985, Cain 1985, McDonald 2007, Chambers et al. 2017, Cho 2017, Stephanopoulos and McGhee 2018), but these three are considered here based on past discussions in U.S. Supreme Court cases. In *Davis v. Bandemer*, the Court ruled that partisan gerrymandering claims are justiciable. However, the Court also upheld the Indiana district plan in question, stating that "the mere lack of proportional representation will not be sufficient to prove unconstitutional discrimination" (Davis 1986). In *LULAC v. Perry*, several Supreme Court justices made comments suggesting partisan symmetry could possibly be one criterion in a general test for resolving gerrymandering claims (Grofman and King 2007). In 2016, a U.S. District Court used the efficiency gap measure introduced by Stephanopoulos and McGhee (2015) in ruling the 2011 Wisconsin district plan an unconstitutional partisan gerrymander, the first condemnation of a district plan by a court in over 30 years (Bernstein and Duchin 2017). In June 2019, the Supreme Court decided not to overturn district plans in North Carolina and Maryland, vacating lower court decisions and remanding the cases gerrymandering question back to the state level. The majority opinion held that "Partisan gerrymandering claims present political questions beyond the reach of the federal courts" (*Rucho v. Common Cause* 2019). While the Supreme Court declines to establish specific gerrymandering tests, individual states which do so may incorporate at least one of these three fairness metrics.

The paper is organized as follows. Section 2 defines the bisection protocol and, in a relaxed continuous nongeometric setting, finds a subgame perfect Nash equilibrium and player utilities in the limit as the number of districts goes to infinity. Section 3 analyzes the more realistic discrete geometric setting, establishing a hardness result, studying optimal strategies on small grid instances, and discussing feasibility issues. Section 4 implements heuristic drawing strategies for the bisection and ICYF protocols using mixed integer program formulations and applies them to congressional redistricting in 18 states, including a detailed case study for Iowa ($n = 4$ districts). Section 5 addresses limitations of the computational experiments and the bisection protocol more broadly. Section 6 provides concluding remarks and directions for future work. In the e-companion, Appendix A expounds on the impact of the packing constraint described in Section 2.3. Appendix B motivates the use of a cut-edge compactness objective in the bisection heuristic. Appendix C provides details of the protocol heuristics applied to real-world redistricting instances, including figures of all final district plans produced. Appendix D collects all proofs.

2. The Bisection Protocol

The bisection protocol is a two-player redistricting game similar to ICYF. For a state with n districts, player 1 first partitions the state into two pieces with sufficient populations to comprise $\lfloor n/2 \rfloor$ and $\lceil n/2 \rceil$ districts. Player 2 then divides each of these in half (up to rounding). Player 1 then divides each of the four pieces created by player 2 in half (up to rounding), and so on. A piece with an appropriate population for a single district is a finalized district, and the process terminates when all pieces are finalized districts. Figure 1 demonstrates how the eight congressional districts of Wisconsin for 2013–2022 could be obtained via the bisection protocol. A similar recursive bisection approach has been applied to graph partitioning before but with a single agent making each bisection (Bichot and Siarry 2013, §1.9). Our bisection protocol may be viewed as a game-theoretic variation of this canonical recursive bisection algorithm with redistricting-focused objectives. Levin and Friedler (2019) propose a non-game-theoretic redistricting algorithm with a similar divide-and-conquer structure. Algorithm 1 describes the bisection protocol at a high level and is similar to Algorithm 1.2 of Bichot and Siarry (2013).

In a real-world redistricting instance, the input graph G has a vertex for each indivisible geographic unit (e.g., census block) and includes an edge between two units if they have a (nontrivial) shared border. Vertices are weighted with the populations of the corresponding units, so the player-controlled bisection step in Line 9 of Algorithm 1 must produce pieces with populations approximately proportional to the district counts (i.e., V_1 comprises approximately an a/m fraction of the total population in H). Rather than directly analyzing the bisection protocol for the real-world setting, the theoretical analysis in this section assumes the same *continuous nongeometric* (CN) setting considered in Pegden et al. (2017) for the ICYF protocol, in which all citizens are voters, districts need not be contiguous, and the population is continuous rather than discrete. Sections 3

Figure 1 **The bisection protocol applied to Wisconsin.**

(a) WI counties (b) Player 1 bisects (c) Player 2 bisects (d) Player 1 bisects

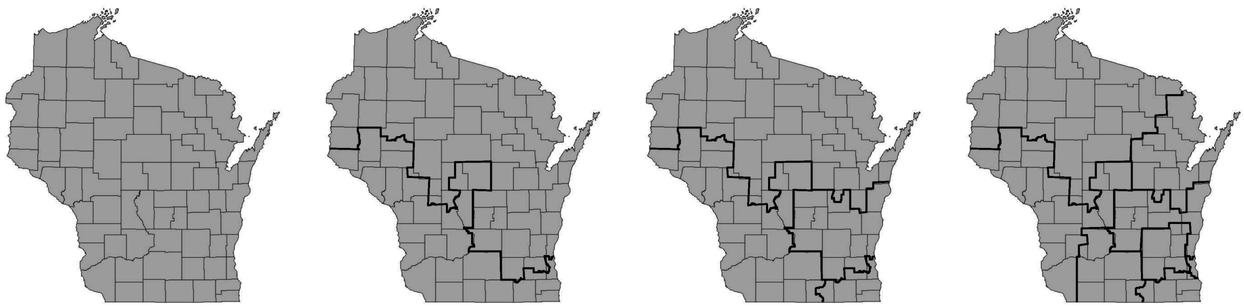

Note. The four panels illustrate how the eight-district plan of Wisconsin for 2013–2022 could be obtained via the bisection protocol. Panel (a) shows the counties all in one piece at the beginning of the protocol. Panel (b) shows the state partitioned into two pieces by player 1. Panel (c) shows each of the two pieces cut into two pieces by player 2. Panel (d) shows each of the four pieces cut into two pieces again by player 1, which comprise the final district plan.

Algorithm 1: Bisection protocol for partitioning a graph into a set $\mathcal{D}$ of n districts

Algorithm BisectionProtocol (graph $G = (V, E)$, number of districts $n \in \mathbb{Z}^+$)
  $\mathcal{D} \leftarrow$ RecursiveBisect (G, n, 1)
  return $\mathcal{D}$

Procedure RecursiveBisect (subgraph H, number of districts m, player $d \in \{1, 2\}$)
  if $m = 1$ **then**
    return $\{V(H)\}$
  end
  $a \leftarrow \lfloor m/2 \rfloor$
  $(V_1, V_2) \leftarrow$ Player d bisects H into connected components of weights a and $m - a$
  $d \leftarrow 3 - d$ `// Switch players`
  $\mathcal{D}_1 \leftarrow$ RecursiveBisect ($H[V_1]$, a, d)
  $\mathcal{D}_2 \leftarrow$ RecursiveBisect ($H[V_2]$, $m - a$, d)
  return $\mathcal{D}_1 \cup \mathcal{D}_2$

and 4 revisit the *discrete geometric* (DG), or graph-based, setting. Define the *vote-share* of player i as the number of statewide votes for player i divided by the average number of voters per district. For a state with n districts, the vote-shares s_1 and s_2 for player 1 and player 2, respectively, satisfy $s_1 + s_2 = n$. It will sometimes be convenient to refer to the *normalized vote-share* of player i, which is the vote-share of player i divided by the number of districts, so that the sum of the players' normalized vote-shares is 1 rather than n.

In the CN setting, the district plan is fully characterized by a list $D_n = [s_{1,1}, s_{1,2}, \ldots, s_{1,n}]$, where $s_{1,i} \in [0,1]$ is the vote-share of player 1 in district i, and $\sum_{i=1}^{n} s_{1,i} = s_1$. Player 1 wins district i if and only if $s_{1,i} \geq 1/2$, (i.e., player 1 achieves a simple majority or the players are tied). Algorithm 2 describes the bisection protocol in the CN setting.

Algorithm 2: Bisection protocol for dividing statewide vote-shares into a list D of n district vote-shares for the continuous nongeometric setting

1 **Algorithm** BisectionProtocolCN (player 1 vote-share s_1, number of districts $n \in \mathbb{Z}^+$)
2     $D \leftarrow$ RecursiveBisectCN (s_1, n, 1)
3   **return** D

4 **Procedure** RecursiveBisectCN (s_1, m, player d)
5   **if** $m = 1$ **then**
6   **return** $[s_1]$
7   **end**
8     $a \leftarrow \lfloor m/2 \rfloor$
9     $(s_1^1, s_1^2) \leftarrow$ player d partitions s_1 into $0 \leq s_1^1 \leq a$, $0 \leq s_1^2 \leq m - a$
10     $D_1 \leftarrow$ RecursiveBisectCN (s_1^1, a, $3 - d$)
11     $D_2 \leftarrow$ RecursiveBisectCN (s_1^2, $m - a$, $3 - d$)
12   **return** `concatenate`(D_1, D_2)

2.1. Optimal Strategies in the CN Setting

This section presents optimal strategies for players 1 and 2 in the bisection protocol for every possible vote-share in the CN setting, assuming each player tries to maximize their seat-share. These optimal strategies are applied in Line 9 of Algorithm 2. The following definitions of thresholds and the protocol utility curve allow for the optimal strategies to be computed and evaluated.

Definition 1. Consider the bisection protocol in the CN setting on n districts. Suppose player 1 goes first and wins any ties. For $1 \leq j \leq n$, define the *threshold* $t_{n,j}$ as the minimum vote-share needed by player 1 to win at least j districts under optimal play by both players. That is, player 1 wins at least j districts if and only if $s_1 \geq t_{n,j}$. Define $t_{n,0}$ as 0 and $t_{n,n+1}$ as n.

Definition 2. For $n \in \mathbb{N}$, the *protocol utility curve* for the bisection protocol in the continuous nongeometric setting is $f_n\colon [0,1] \to [0,1]$, where $f_n(v) = \frac{1}{n} \cdot \max\{j \in \mathbb{N}\colon 0 \leq j \leq n, t_{n,j}/n \leq v\}$ is the normalized seat-share won by player 1 given normalized vote-share v.

In Definition 1, if ties are not broken in favor of player 1, then the inequality is strict ($s_1 > t_{n,j}$). For simplicity, ties are always broken in favor of player 1, and player 1 always has the first turn. Awarding ties to player 2 does not substantially change the results, since a small perturbation of the vote-shares could ensure s_1 does not equal any threshold $t_{n,j}$. Note that $t_{n,n+1}$ is defined as n only to simplify statements of lemmas and proofs; it is, of course, impossible to win $n+1$ districts in a state with n districts. In Definition 2, $t_{n,j}/n$ represents the normalized vote-share threshold needed to win the normalized seat-share j/n.

The protocol utility curve should be understood in relation to the standard definition of a *seats-votes curve* for a district plan, which is an estimate of normalized seat-shares under hypothetical elections as normalized vote-share varies across all districts from a given election (DeFord et al. 2021a). One way to obtain a seats-votes curve from a district plan is the *uniform partisan swing* strategy. The observed seats and votes point (i.e., a particular election) is plotted first, and the rest of the curve is estimated by uniformly adjusting the normalized vote-share across all districts and observing the resulting normalized seat-share (Grofman and King 2007). Critically, the district plan remains the same throughout this counterfactual analysis. In contrast, any two distinct points on the protocol utility curve correspond to different executions of the protocol and may therefore represent vastly different district plans.

Definitions 1 and 2 are first used in Lemma 1 to characterize the conditions under which player 1 can win all ($s_1 \geq t_{n,n}$) or none ($s_1 < t_{n,1}$) of the districts when n is a power of two.

Lemma 1. *Consider the bisection protocol in the CN setting on n districts. When $n = 2^r$ for some $r \in \mathbb{N}$, the thresholds $t_{n,1}$ and $t_{n,n}$ satisfy the system of recurrences*

$$t_{n,1} = t_{n,n} = 1/2 \text{ for } n = 1,$$
$$t_{n,n} = n - 2t_{n/2,1}, \text{ and}$$
$$t_{n,1} = n/2 - t_{n/2,n/2}.$$

Furthermore, these recurrences have closed-form solutions

$$t_{n,1} = 2^{\lceil (r-1)/2 \rceil - 1} = \begin{cases} \sqrt{n}/2 & \text{if } r \text{ is even} \\ \sqrt{2n}/4 & \text{otherwise} \end{cases}$$

and

$$t_{n,n} = n - 2^{\lceil r/2 \rceil - 1} = \begin{cases} n - \sqrt{n}/2 & \text{if } r \text{ is even} \\ n - \sqrt{2n}/2 & \text{otherwise.} \end{cases}$$

The closed-form solution for $t_{n,1}$ in Lemma 1 captures the intuition that if player 1 wants to win exactly one district in a state with $n = 2^r$ districts, then player 1 should keep as much of their share in one piece as possible each time they bisect a piece. Throughout r rounds, player 2 has $\lceil (r-1)/2 \rceil$ turns and can divide each of player 1's separate vote-shares in half in each of those rounds. For player 1 to win one of the final districts, player 1 must have at least a vote-share of 1/2 intact after the $\lceil (r-1)/2 \rceil$ divisions by 2 caused by player 2's splits, and hence $t_{n,1} = 2^{\lceil (r-1)/2 \rceil - 1}$.

While closed-form expressions for $t_{n,j}$ are not known, Lemma 2 presents recurrences for the thresholds.

LEMMA 2. *Consider the bisection protocol in the CN setting on n districts. Let $j \in \mathbb{N}$, $0 \leq j \leq n$. If $j = 0$, then $t_{n,j} = 0$. When $n = j = 1$, $t_{n,j} = 1/2$. For $n \geq 2$ and $j \geq 1$,*

$$\begin{aligned} t_{n,j} &= \min_{k \in K(j)} \left\{ (a - t_{a,a-k+1}) + \left(b - t_{b,b-(j-k)+1}\right) \right\} \\ &= \min_{k \in K(j)} \left\{ n - t_{a,a-k+1} - t_{b,b-(j-k)+1} \right\}, \end{aligned}$$

where $a = \lfloor n/2 \rfloor$ is the size of the smaller piece after the cut (side A), $b = \lceil n/2 \rceil$ is the size of the larger piece (side B), and $K(j) = \{k \in \mathbb{N} \colon 0 \leq k \leq j, 0 \leq k \leq a, 0 \leq j - k \leq b\}$ is the set of permissible seat-shares from the part of size a.

Lemma 2 allows for the protocol utility curve f_n, which is a step function with equal increments of $1/n$ at the thresholds $t_{n,j}$, to be computed via dynamic programming for any integer n. Figure 2 presents some examples of the protocol utility curves f_n with the perfectly proportional curve $f(v) = v$ included for comparison. As an example of how to interpret the protocol utility curves, consider f_3 in Figure 2a. For $n = 3$, player 1 starts by cutting the state into a one-district piece and a two-district piece. If player 1 starts with a normalized vote-share between 0 and $t_{3,1}/3 = (1/2)/3 = 1/6$, then player 1 cannot win a single district. However, if player 1 has a normalized vote-share between 1/6 and 1/2, then player 1 can win the district in the one-district piece by allocating a vote-share of at least 0.5 to that piece. By similar analysis, $t_{3,2} = 1/2$ and $t_{3,3} = 2/3$.

In Figure 2, the protocol utility curves appear to approach a straight line (i.e., perfect proportionality) as n increases. Theorem 1 formalizes this convergence when n is a power of two.

THEOREM 1. *Let $f \colon [0,1] \to [0,1]$ be given by $f(v) = v$. Then $(f_{2^r})_{r=1}^{\infty}$ converges uniformly to f.*

The proof of Theorem 1 relies on the following lemma, which presents a closed-form solution and optimal strategy for the recurrence of Lemma 2 for the special case when n is a power of two.

LEMMA 3. *Consider the bisection protocol in the CN setting on $n = 2^r$ districts, $r \in \mathbb{N}$. Then $t_{n,0} = 0$, and for all integers $1 \leq j \leq n$,*

$$t_{n,j} = t_{n,1} + \sum_{i=0}^{\infty} \left(2^{i-2} \left\lfloor \frac{j-1}{2^{2i-1}} \right\rfloor \right). \tag{1}$$

Furthermore, an optimal choice of the number of districts to win from one of the two pieces is

$$\arg\min_{k \in K(j)} \{ n - t_{n/2, n/2-k+1} - t_{n/2, n/2-(j-k)+1} \}, \tag{2}$$

where $K(j)$ is the set of permissible seat-shares from side A as defined in Lemma 2. If $1 \leq j \leq n/2$, then j is (one of) the minimizer(s) in (2), and if $n/2 < j \leq n$, then $n/2$ is the minimizer. These choices form a subgame perfect Nash equilibrium.

Informally, the optimal strategy presented in Lemma 3 is: **"When you divide a piece in half, win all seats on one side, if possible, with as little vote-share as possible. If you cannot win all seats on one side, then put all of your vote-share in one side."** However, this strategy is sometimes suboptimal when n is not a power of two. Consider a player trying to win $j=4$ of $n=6$ districts with a vote-share of $t_{6,4}=3$. The region will be split into pieces of sizes 3 and 3. With a vote-share of $3-t_{3,2}=1.5$, the player can win two districts in a piece of size 3, so splitting a vote-share of 3 evenly wins four districts total. If the player instead follows the informal version of the optimal strategy from Lemma 3, they win all three districts from one piece of size 3, using $3-t_{3,1}=2.5$ of their vote-share, but then they need an additional $3-t_{3,3}=1$ vote-share to win one district from the other piece of size 3, which is more than their remaining 0.5 vote-share.

While Theorem 1 proves convergence of the seats-votes curve to perfect proportionality only for the special case where n must be a power of two, it is conjectured that the convergence holds even if non-powers of two are included. Even if the convergence does not hold for non-powers of two, the subgame perfect Nash equilibrium of Lemma 3 is interesting in its own right, as the optimal strategies lead players to pack *themselves* in one side. This self-packing keeps their vote-shares consolidated to avoid falling victim to cracking by their opponent in later rounds. Furthermore, by self-packing in one side, the player implicitly packs their *opponent* in the other side.

Figure 2 **Examples of protocol utility curves for the bisection protocol.**

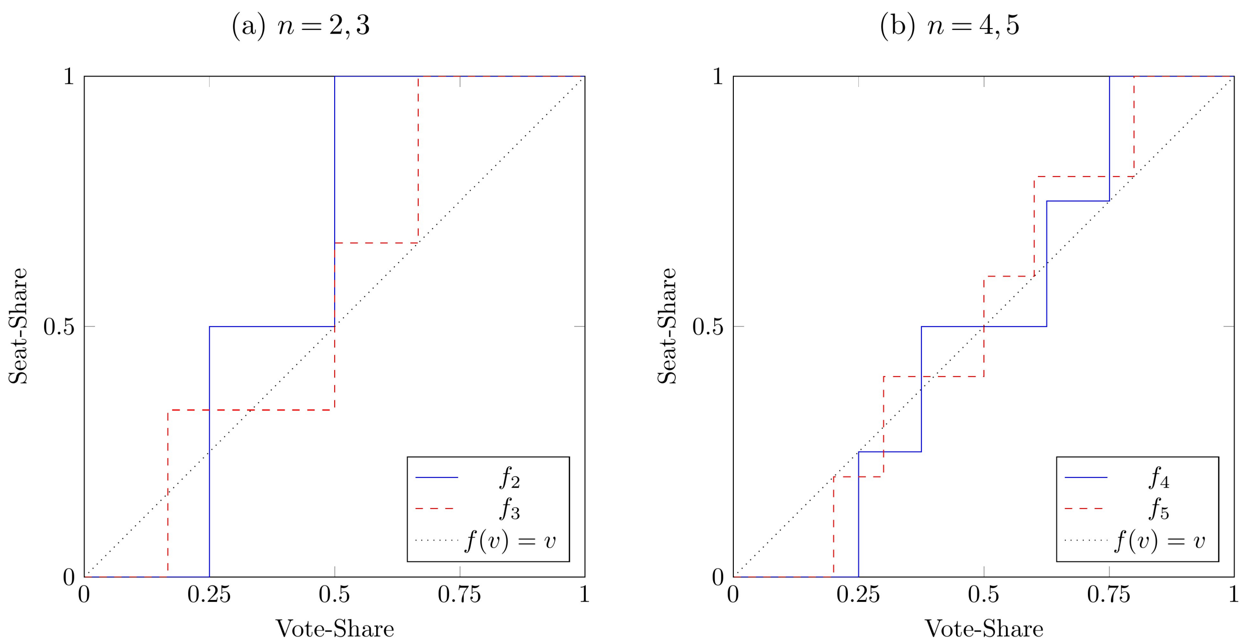

Note. Panels (a) and (b) display protocol utility curves for the bisection protocol in the CN setting for (a) $n=2,3$ and (b) $n=4,5$ districts. Both panels include the line $f(v)=v$ (perfect proportionality) for reference. The vote-share and seat-share axes are normalized to $[0,1]$.

2.2. Protocol Fairness Comparisons

This section compares the bisection and ICYF protocols using the fairness metrics of proportionality, efficiency gap, and partisan symmetry. Pegden et al. (2017) comments on the visual symmetry of ICYF but does not quantify the partisan asymmetry as defined in political science (see Definition 5), likely because the authors emphasize the qualitative improvement in the symmetry of the protocol utility curve compared to that of the one-player-draws protocol. However, a direct comparison between the bisection and ICYF protocols requires quantitative assessments, which all three selected fairness metrics provide. We first formally define the fairness metrics in the context of a two-party election, beginning with the Sainte-Laguë index of disproportionality.

Definition 3. Consider an n-district plan in which player 1 has normalized vote-share $v_i \in [0,1]$ and normalized seat-share $s_i \in [0,1]$. The *Sainte-Laguë (SL) index of disproportionality* is

$$SL(v_1, s_1) = \frac{(s_1 - v_1)^2}{v_1} + \frac{(1 - s_1 - (1 - v_1))^2}{1 - v_1} = \frac{(s_1 - v_1)^2}{v_1} + \frac{(v_1 - s_1)^2}{1 - v_1}, \tag{3}$$

assuming $0 < v_1 < 1$. By definition, $SL(0,0)$ and $SL(1,1)$ are zero. Otherwise ($v_1 = 0$ or 1 but $s_1 \neq v_1$), $SL(v_1, s_1) \equiv \infty$; this scenario is only hypothetical, and cannot occur in a practical districting scenario.

Definition 3 is adapted from Gallagher (1991), which describes the SL index as "the soundest of all the [disproportionality] measures." Further studies confirm its strength and robustness (Pennisi 1998, Goldenberg and Fisher 2019). The SL index bears a resemblance to Pearson's χ^2 statistic and can be viewed as a goodness of fit measure against the baseline of perfect proportionality (Goldenberg and Fisher 2019).

The second fairness metric is the efficiency gap, which is "the difference between the parties' respective wasted votes, divided by the total number of votes cast in the election" (Stephanopoulos and McGhee 2015). A party's wasted votes include all votes in a lost district and any votes beyond 50% in a won district. The efficiency gap is positive (negative) when player 1 (2) wastes more votes, and its magnitude never exceeds 0.5 (i.e., 50% of votes).

Definition 4. Consider an n-district plan $D_n = [s_{1,1}, s_{1,2}, \ldots, s_{1,n}]$, where $s_{1,i} \in [0,1]$ is the vote-share of player 1 in district i. For $A \in \{1,2\}$ and $i \in [n]$, define $w_{A,i}$ to be the total wasted vote-share of player A,

$$w_{1,i} \equiv \begin{cases} s_{1,i} - 1/2 & \text{if } s_{1,i} \geq 1/2 \\ s_{1,i} & \text{otherwise,} \end{cases}$$

with $w_{2,i}$ computed analogously. The *efficiency gap* of the district plan is then

$$EG(D_n) \equiv \frac{1}{n} \left(\sum_{i=1}^{n} (w_{1,i}) - \sum_{i=1}^{n} (w_{2,i}) \right). \tag{4}$$

The third, and most involved, fairness metric is partisan asymmetry, which measures the degree to which seat-share outcomes for the two parties are asymmetric with respect to vote-share. The following definition is adapted from the partisan Gini score of Grofman (1983). We refer the reader to Section 2.1 and DeFord et al. (2021a) for details regarding the seats-votes curve and uniform partisan swing.

DEFINITION 5. Suppose $D_n = [s_{1,1}, s_{1,2}, \ldots, s_{1,n}]$ is the sorted list of player 1 vote-shares in each district resulting from some protocol. The total player 1 vote-share is $s_1 = \sum_i s_{1,i}$, and the normalized vote-share of player 1 is s_1/n. Let $g\colon [0,1] \to [0,1]$ be the (uniform partisan swing) seats-votes curve of player 1 constructed from D_n. The seats-votes curve for player 2 is then $1 - g(1-v)$.

The *partisan asymmetry* of the district plan is the area between the seats-votes curves for the two players,

$$PA(g) = \int_0^1 |g(v) - (1 - g(1-v))| dv. \tag{5}$$

The definition of partisan asymmetry in (5) is agnostic to which player is labeled player 1 and summarizes the asymmetry of the district plan across all possible vote-shares. When n is finite, g is a step function, reducing the integral in (5) to a straightforward sum of rectangular areas.

Note that computing partisan asymmetry requires counterfactual analysis (i.e., election outcomes following a hypothetical uniform partisan swing) to produce the seats-votes curve, whereas the SL index and efficiency gap follow directly from the list of district vote-shares. Hence the SL index and efficiency gap are computed from a fixed district plan produced by a protocol, but partisan asymmetry is computed from a seats-votes curve representing a hypothetical variety of plans derived from the protocol-produced district plan.

To compare the fairness of the bisection and ICYF protocols, we compute all three metrics for the district plans produced by each protocol, varying $v_1 \in \{40\%, 41\%, \ldots, 60\%\}$ and $n \in \{1, 2, \ldots, 150\}$. This range for n covers not only all congressional apportionments (up to 52 districts) but also most state legislative district counts.

The bisection protocol generally achieves a fairer SL index at equilibrium than does ICYF in the CN setting. Figure 3 presents heat maps of the SL indices for each vote-share and district count considered. Both protocols tend to produce more proportional district plans as the number of districts increases. On one hand, the SL index for bisection appears more volatile than that of ICYF, especially near odd powers of two (2, 8, 32, 128, etc.) for which player 1 gets a full extra turn in the bisection protocol. On the other hand, the SL index for ICYF stratifies based on normalized vote-share, with greater deviations from 50% producing less proportional plans, whereas the SL index for bisection is often below 1% for all normalized vote-shares considered. Furthermore, the limiting protocol utility curve of the bisection protocol stated in Theorem 1 for n being a power of two achieves an SL index of zero for all normalized vote-shares. By contrast, the limiting protocol utility curve of ICYF is not linear but piecewise-quadratic (see Theorem 2.3 of Pegden et al. 2017), resulting in higher SL indices further from 50% normalized vote-share.

Figure 3 The Sainte-Laguë index of disproportionality for the bisection and I-cut-you-freeze protocols.

(a) Bisection (b) I-cut-you-freeze

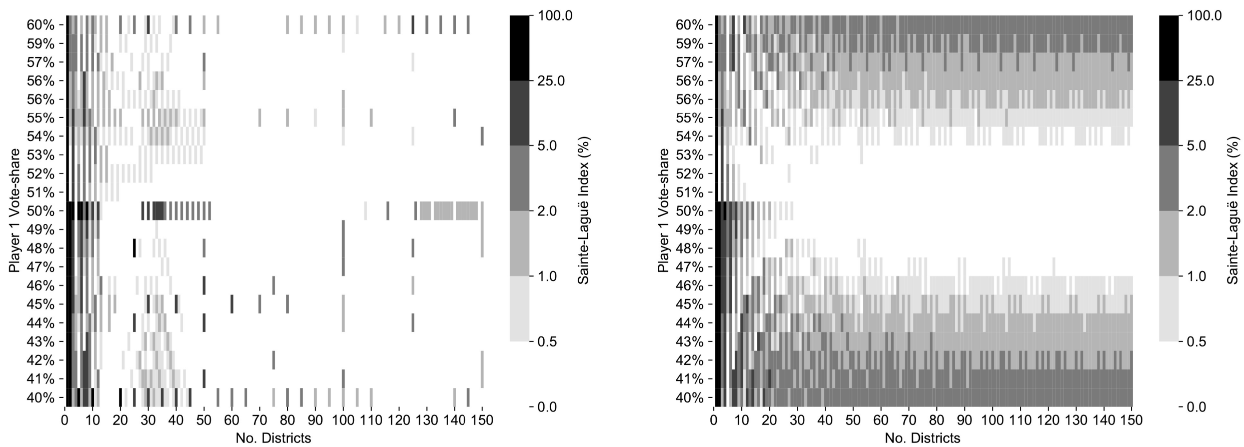

Note. Panels (a) and (b) plot the Sainte-Laguë index for 1–150 districts and normalized vote-shares from 40% to 60% resulting from the (a) bisection and (b) I-cut-you-freeze protocols in the CN setting.

Figure 4 gives heat maps of efficiency gap as a function of normalized vote-share and number of districts for the bisection and ICYF protocols. For both protocols, results for $n \leq 10$ districts align with previous studies suggesting the efficiency gap is volatile and counterintuitive when the number of districts is small (Cho 2017). Based on a historical analysis of U.S. congressional and state house district plans, Stephanopoulos and McGhee (2015) recommend using $\pm 8\%$ as the threshold for claiming a district plan contains excessive partisan bias. Hence Figure 4 distinguishes $\pm 4\%$ and $\pm 8\%$ from more extreme efficiency gap values.

In the CN setting considered, (4) generally yields fairer (smaller magnitude) efficiency gap values for ICYF than for the bisection protocol. Above 15 districts, ICYF achieves efficiency gaps between $\pm 8\%$ for all normalized vote-shares considered. In contrast, bisection generally stays within the desired efficiency gap range for 45%–55% normalized vote-share, but often produces high-magnitude efficiency gaps for 40% and 60% normalized vote-share. For both protocols, efficiency gap correlates with normalized vote-share, which means the stronger player tends to waste more votes.

Figure 5 offers partisan asymmetry heat maps for the bisection and ICYF protocols. ICYF achieves lower partisan asymmetry, as formulated in Definition 5, than the bisection protocol for

Figure 4 The efficiency gap of the bisection and I-cut-you-freeze protocols.

(a) Bisection (b) I-cut-you-freeze

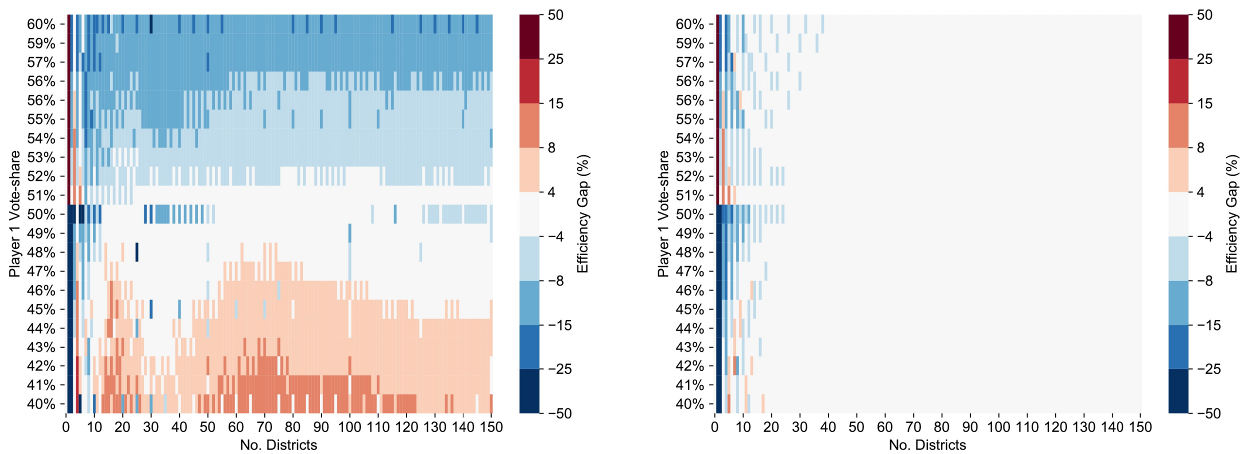

Note. Panels (a) and (b) present the efficiency gap plots for 1–150 districts and 40%–60% normalized vote-shares resulting from the (a) bisection and (b) I-cut-you-freeze protocols in the CN setting.

most district counts of interest in the United States (1–53 congressional districts, 30+ state house districts). Both protocols tend to produce more symmetric district plans when the initial normalized vote-share is closer to 50%. For a given normalized vote-share deviation from 50%, bisection tends to produce more symmetric plans when player 1 is the minority (Fig. 5a). In this regard, ICYF seems less sensitive to which player goes first; for sufficiently many districts, partisan asymmetry scores correlate with the absolute value of the normalized vote-share deviation from 50% (Fig. 5b).

2.3. Impact of a Packing Constraint

Optimal play in both protocols sometimes requires a player to completely pack a district with their opponent's voters. This is likely impossible in practice due to the political geography of the state. One straightforward way to approximate the effect of political geography is to impose a packing constraint on the district plan in the form of an upper bound δ on the maximum vote-share margin in a district. Each $s_{1,i}$ in the list D_n must then lie in the interval $[0.5-\delta, 0.5+\delta]$. Call δ the *packing parameter*. The packing parameter seeks to capture both implicit packing limits imposed by political geography and potential explicit constraints imposed by legislatures. The ICYF protocol utility curve with no packing constraint, given by Theorem 2.3 of Pegden et al. (2017), is

$$f^{\text{ICYF}}(v) = \begin{cases} 2v^2 & \text{if } 0 \le v \le \frac{1}{2}, \\ 1-2(1-v)^2 & \text{if } \frac{1}{2} < v \le 1. \end{cases} \tag{6}$$

Theorem 2 shows how the unrestricted protocol utility curves for both bisection and ICYF change in an environment where packing is limited based on δ.

Figure 5 The partisan asymmetry of the bisection and I-cut-you-freeze protocols.

(a) Bisection (b) I-cut-you-freeze

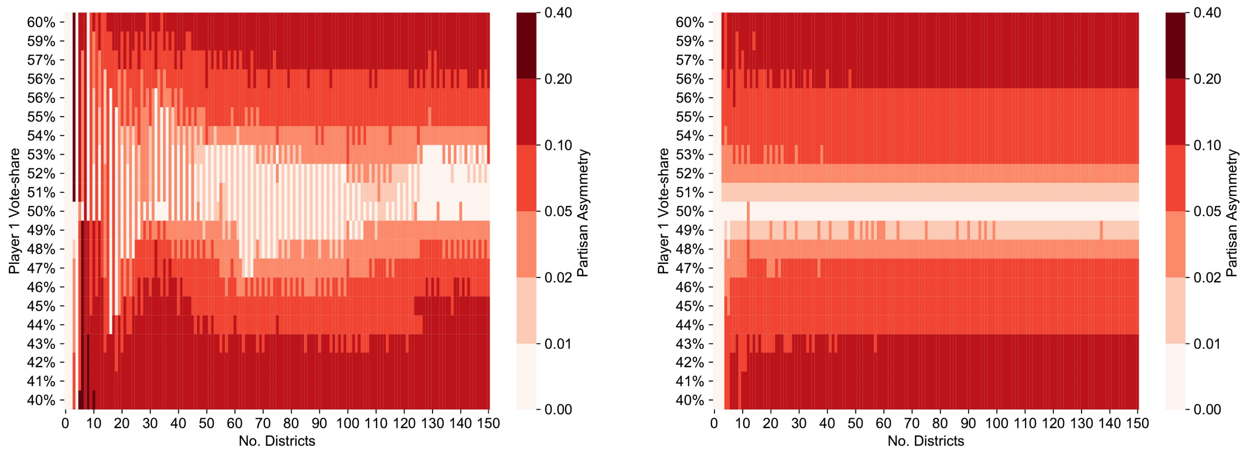

Note. Panels (a) and (b) plot the partisan asymmetry for district plans produced by the (a) bisection and (b) I-cut-you-freeze protocols in the CN setting for 1–150 districts and normalized vote-shares from 40% to 60%.

THEOREM 2. *Let $\delta \in (0, 0.5)$, $\gamma = 0.5 - \delta$, and $\alpha \in [\gamma, 1-\gamma]$. In the bisection protocol with packing parameter δ and n a power of two, the seat-share won by player 1 with a normalized vote-share of α obeys the limiting distribution*

$$\lim_{r\to\infty} f_{2^r}(\alpha) = \frac{\alpha - \gamma}{1 - 2\gamma}.$$

In ICYF with packing parameter δ and arbitrary n, the seat-share won by player 1 with a normalized vote-share of α obeys the limiting distribution

$$\lim_{n\to\infty} \frac{\sigma(n, \alpha n)}{n} = \begin{cases} 2\left(\frac{\alpha-\gamma}{1-2\gamma}\right)^2 & \text{for } \gamma \le \alpha \le \frac{1}{2} \\ 1 - 2\left(1 - \left(\frac{\alpha-\gamma}{1-2\gamma}\right)\right)^2 & \text{for } \frac{1}{2} < \alpha \le 1-\gamma, \end{cases}$$

where $\sigma(n, s)$ denotes the number of seats won by player 1 under optimum play of ICYF with n districts and player 1 vote-share s.

Theorem 2 demonstrates that the packing constraint essentially compresses the protocol utility curve horizontally, reducing its domain from $[0, 1]$ to $[0.5 - \delta, 0.5 + \delta]$. This compression also impacts all three fairness measures, and this impact can differ between the bisection and ICYF protocols. For example, adding the packing constraint with parameter $\delta = 0.25$ (chosen based on real-world congressional district margins, as discussed in Appendix A) has the following impacts on the fairness measures:

- worsens the proportionality of both protocols by horizontally compressing the seats-votes curve, maintaining an advantage for the bisection protocol;

- improves (worsens) the efficiency gap of the bisection (ICYF) protocol, changing the advantage from ICYF to the bisection protocol; and
- worsens the partisan asymmetry of both protocols, maintaining an advantage for ICYF.

Appendix A illustrates these effects with heat maps comparable to Figures 3-5.

3. Bisection in the Discrete Geometric Setting

The protocol fairness comparisons in Section 2.2 ignore the real-world constraints of contiguous districts, discrete voter populations, and political geography. This section considers the *discrete geometric* (DG) setting, in which districts are connected components of a graph whose vertices represent indivisible units with discrete populations and fixed vote counts for each party. Section 3.1 formalizes the DG setting. Section 3.2 partially characterizes the subgame perfect Nash equilibria. Section 3.3 considers the bisection protocol on the restricted class of grid graphs. Section 3.4 discusses feasibility concerns. Section 3.5 summarizes key insights from the DG setting and motivates the MIP-based heuristics of Section 4.

3.1. Definitions and Notation

Recall from Algorithm 1 that the bisection protocol in the DG setting takes as input a positive integer n of required districts and an undirected, simple graph $G = (V, E)$, where V represents the set of indivisible geographic units, and $ij \in E$ if and only if units i and j share at least one boundary of nonzero length. In other words, the input graph G is the dual graph of the region's unit boundaries. To achieve (approximate) population balance, each $i \in V$ is weighted with its unit's population p_i. We also associate with each unit the *vote counts* $v_{1,i}$ for player 1 and $v_{2,i}$ for player 2, which one may estimate from past elections. A player's normalized vote-share is simply the sum of their vote counts in every unit divided by the total number of votes. For each unit i, $v_{1,i} + v_{2,i} \leq p_i$, and equality holds in the rare case that all residents are eligible, active voters. Note p_i, $v_{1,i}$, and $v_{2,i}$ are nonnegative and can be nonintegral when derived from statistical estimates or (dis)aggregation. For brevity, let $p(S) = \sum_{i \in S} p_i$ denote the combined population for any $S \subseteq V$. A district plan with n districts is a partition $\mathcal{D} = \{D_k\}_{k=1}^{n}$ of V, where each D_k induces a (nonempty) connected component in G. A district plan is *balanced* if and only if each district's total population is within the closed interval $\left[\overline{P}(1-\tau), \overline{P}(1+\tau)\right]$, where $\overline{P} = (\sum_{i \in V} p_i)/n$ is the mean district population and $\tau \geq 0$ is the maximum allowable deviation from the ideal. For example, $\tau = 0.03$ for Montana's state house and senate districts, requiring each district's population to fall within 3% of the mean (Montana Districting and Apportionment Commission 2010). To enforce population balance with a threshold of τ in the final districts, the bisection returned by player d in Line 9 of Algorithm 1 must satisfy $\left|p(V_1) - a\overline{P}\right| \leq a\tau\overline{P}$ and $\left|p(V_2) - (m-a)\overline{P}\right| \leq (m-a)\tau\overline{P}$. This requirement is necessary but not sufficient, because imbalanced unit populations or the structure of the piece could prevent a balanced bisection in the next round. Section 3.4 further discusses these and other feasibility concerns.

3.2. In Search of Equilibria

The bisection protocol applied to the DG setting is a perfect information finite extensive-form game, so it is guaranteed to have a subgame perfect Nash equilibrium by backward induction (Osborne 2004). At each decision point (node) in the game tree, however, the current player may have exponentially many actions (feasible partitions) from which to choose. Naïve backward induction on the game tree is therefore intractable for even moderately-sized instances (tens of units, four districts). But could the optimal strategies be specified in a general manner that avoids enumerating best responses to every feasible bisection in the previous round? Unlikely. As Theorem 3 states, it is NP-complete to determine whether the vote counts can be evenly split with a bisection, even for a restricted class of graphs. Such a balanced bisection would be an optimal strategy when $n = 2$ and the player with greater vote-share is cutting, since having the same vote-share in each of the two districts would give them both seats. Not only is it NP-complete to determine whether such a basic strategy can be implemented, but also Chataigner et al. (2007) establish inapproximability results for maximizing balance when partitioning a 2-connected graph into two connected parts. Hence there is little hope for computing a subgame perfect Nash equilibrium in polynomial time.

THEOREM 3. *Let $G = (V, E)$ be an undirected graph with vertex weights $w : V \to \mathbb{R}_{\geq 0}$ and vertex populations $p : V \to \mathbb{R}_{\geq 0}$. For a vertex subset S, define $w(S) = \sum_{v \in S} w(v)$ and $p(S) = \sum_{v \in S} p(v)$. Let* BALANCEDBISECTION *be the problem of deciding whether G admits a bisection (connected 2-partition) into components D_1, D_2 such that $w(D_1) = w(D_2)$ and $p(D_1) = p(D_2)$. The problem* BALANCEDBISECTION *for $3 \times M$ grid graphs is NP-complete.*

The proof of Theorem 3 provided in the e-companion is nearly identical to the proof of Theorem 1.1 in Becker et al. (1998) for a similar problem.

3.3. Optimal Bisection for Small Grids

Since computing a subgame perfect Nash equilibrium directly from the full game tree is generally intractable, we find optimal bisection strategies for small grid graphs and compare the results to those predicted by the CN setting. How closely will the optimal strategies for these DG instances resemble those of the corresponding CN instance? The answer depends on the spatial distribution of voters, as the following examples illustrate.

Figure 6 presents 4×6 and 6×6 grid graphs with three hypothetical spatial voter distributions: one large metropolitan area on the western border (`lm`), a symmetric north/south split (`ns`), and two medium cities in opposite corners (`oc`). Each unit has the same population and is fully for

Figure 6 **Examples of two-party spatial distributions of voters for small grid graphs.**

(a) `4x6_lm` (b) `4x6_ns` (c) `4x6_oc` (d) `6x6_lm` (e) `6x6_ns` (f) `6x6_oc`

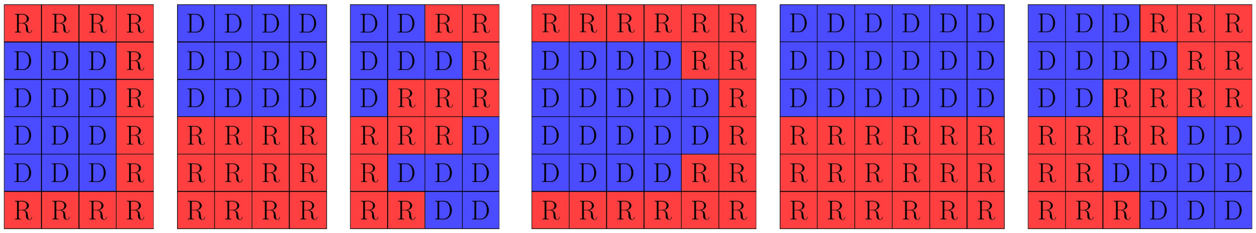

one party or the other. All six instances have 50% normalized vote-share. We consider partitioning these instances into $n=4$ or $n=6$ districts with perfect population balance.

For such small grid graphs, enumeration of all district plans is tractable. We generate all possible partitions of the 4×6 and 6×6 grid graphs into $n=4$ or $n=6$ districts using the `enumerator` tool (Schutzman 2019). The 4×6 grid has 2,369 four-district plans and 2,003 six-district plans. There are far more partitions of the 6×6 grid: 442,791 four-district plans and 451,206 six-district plans.

For each party and each instance, we find a sequence of optimal bisections by a recursive branch-and-bound algorithm with memoization of maximum utilities for particular subgraphs. The search for optimal bisections with $n=4$ or $n=6$ takes 1–11 seconds for the 4×6 instances and 49–8,379 seconds (over two hours) for the 6×6 instances, running on a personal machine at 1.6 GHz with an Intel Core i5-8250U CPU, using a single core.

Table 1 presents the outcomes under optimal play for these six instances in the DG setting, as well as their CN counterparts computed using Lemma 2. A table entry of "3D, 1T, 2R" indicates three wins for party D, one tie, and two wins for party R. (Note the 6×6 grid with $n=4$ precludes ties, since each district has 9 units.) For $n=4$ districts, the DG setting gives an advantage to the first player that was not present in the CN setting. This phenomenon is symmetric for all instances but `4x6_oc`, in which D can win three districts when given the first turn, but R has to settle for two wins and one tie.

Table 1 **Equilibrium utilities for the bisection protocol on small grid instances.**

Setting	Instance	$n=4$ Districts		$n=6$ Districts	
		D first	R first	D first	R first
CN	—	2D, 2R	2R, 2D	4D, 2R	4R, 2D
DG	`4x6_lm`	3D, 1R	3R, 1D	3D, 1T, 2R	3R, 1T, 2D
	`4x6_ns`	3D, 1R	3R, 1D	3D, 1T, 2R	3R, 1T, 2D
	`4x6_oc`	3D, 1R	2R, 1T, 1D	3D, 1T, 2R	4R, 2D
	`6x6_lm`	3D, 1R	3R, 1D	4D, 2R	4R, 2D
	`6x6_ns`	3D, 1R	3R, 1D	4D, 2R	4R, 2D
	`6x6_oc`	3D, 1R	3R, 1D	4D, 2R	4R, 2D

For $n = 6$ districts, all three 6×6 instances align with the CN setting outcome: four districts for the first player, two for the second. (Recall ties go to the first player.) The 4×6 instances mostly agree on three first-player wins, one tie, and two second-player wins. The only exception is, again, the `4x6_oc` instance, for which when R goes first, they can win a fourth district. Paradoxically, this voter distribution gives D a greater first-player advantage when $n = 4$ but gives R a greater first-player advantage when $n = 6$. These differences stem from the extent to which R can pack D voters, concentrated in opposite corners, into forced districts.

Differences in equilibrium utilities between the CN and DG settings flow from differences in optimal bisection strategies. Table 2 presents the optimal first-round bisection vote-shares for each instance. Recall vote-shares are scaled so that each district population is 1. The core tactics from the CN setting, such as winning as many districts as possible from one piece, carry over to the DG setting, but the details vary. For instance, when $n = 4$ in the CN setting with 50% normalized vote-shares, the first player must allocate 1.5 vote-share to one piece to guarantee winning both districts. In the corresponding DG instances, the structure of the grid graph allows the first player to spend less vote-share in one piece but still win both districts; the excess vote-share translates to an extra win or tie in the other piece. Similarly, when $n = 6$, the CN setting calls for a balanced first-round bisection, which guarantees two wins from each side since ties are awarded to the first player. In the corresponding DG instances, the optimal first-round bisection at times follows the balanced CN strategy but at other times packs the opponent into one piece while still carving out a win or tie. The precise vote-share split varies depending on the spatial voter distribution.

Table 2 **Optimal first-round bisections, expressed as first-player vote-shares in each piece.**

Setting	Spatial Voter Distribution	$n = 4$ Districts		$n = 6$ Districts	
		D first	R first	D first	R first
CN	—	1.5, 0.5	1.5, 0.5	1.5, 1.5	1.5, 1.5
DG	`4x6_lm`	1.33, 0.67	1.33, 0.67	2.25, 0.75	1.5, 1.5
	`4x6_ns`	1.33, 0.67	1.33, 0.67	2.5, 0.5	2.25, 0.75
	`4x6_oc`	1.33, 0.67	1.17, 0.83	1.5, 1.5	1.5, 1.5
	`6x6_lm`	1.22, 0.78	1.33, 0.67	2.17, 0.83	1.67, 1.33
	`6x6_ns`	1.33, 0.67	1.44, 0.56	2.33, 0.67	1.67, 1.33
	`6x6_oc`	1.33, 0.67	1.11, 0.89	2, 1	1.5, 1.5

In summary, the small grid results show the DG setting can increase or maintain bisection's first-player advantage, depending on the spatial voter distribution and number of districts.

3.4. Feasibility Issues

In the CN setting, the bisection protocol has two desirable properties: (1) every possible district plan can be reached, and (2) for any sequence of strategies, the current player has at least one available strategy. In the DG setting, however, the bisection protocol retains neither property. These issues persist even for grid graphs with equipopulous units. Figure 7 demonstrates both an unreachable district plan (Fig. 7a) and a first-player bisection in an eight-district instance that makes the second player's task impossible (Fig. 7b). The four-district plan in Figure 7a is unreachable because no matter which district is in the same piece as district 3 after the first bisection, the other two districts will be disconnected. The first-round bisection in Figure 7b produces a claw-shaped piece, side 1, that player 2 cannot evenly bisect. There is, however, a district plan (thick lines) compatible with the initial bisection, so a compatible full district plan does not suffice to prove feasibility.

Figure 7 **Examples of bisection feasibility issues.**

(a) Unreadable district plan

(b) No available strategies in second round

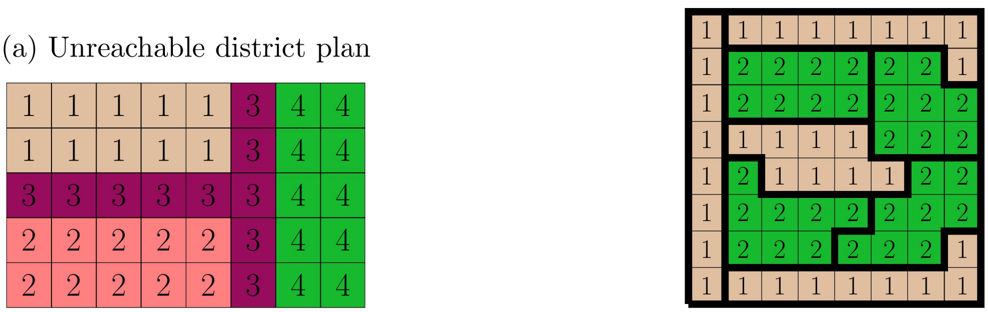

Heterogeneous unit populations further complicate feasibility, as mentioned at the end of Section 3.1. For example, if a two-district piece induces a path with relative unit populations 1–3–7–5–3–1, and the ideal district population is 10, then the best possible population deviation is $\pm 10\%$. In practice, units with large populations may be split to resolve such issues, but this increases the problem size and may divide municipalities or other communities of interest.

While Figure 7a shows not every balanced district plan is reachable via the bisection protocol, all four-district and six-district plans for the 4×6 and 6×6 grids are reachable. The problem of determining whether a plan is reachable reduces to a property of the plan's *district adjacency graph*, in which nodes represent districts and edges indicate shared district borders. A district adjacency graph H with $n = 2^r$ nodes, for some $r \in \mathbb{N}$, can be reached via bisection if and only if H can be reduced to a single node by repeatedly finding a perfect matching and contracting its edges. Intuitively, this process reconstructs the pieces from each round of the bisection protocol; an edge in a matching joins two sides of a bisection. A thorough characterization of the graphs which admit such a recursive perfect matching is an open problem beyond the scope of this paper.

3.5. Key Insights for the Discrete Geometric Setting

Moving to the DG setting from the CN setting limits the ways in which units can be mixed and matched into pieces, complicating the search for equilibria. Finding optimal bisections in the DG setting is related to the balanced connected partitioning problem, for which there are many known hardness and inapproximability results. Naïve backward induction on the game tree is impractical, but an enumerative branch-and-bound algorithm can solve the bisection protocol on 4×6 and 6×6 grid graphs with $n = 4, 6$ districts. Results for such instances with three hypothetical voter distributions highlight the impact of political geography on first-player advantage.

Although enumeration is tractable for the small grid instances considered in Section 3.3, the number of possible partitions explodes as the grid size increases. For example, there are 187,497,290,034 partitions of the 8×8 grid into 8 parts (OEIS Foundation Inc. 2022). We therefore turn to heuristics in Section 4 to study the bisection and ICYF protocols on U.S. congressional redistricting instances with hundreds to thousands of units. The feasibility concerns raised in Section 3.4 must also be addressed for these real-world instances.

4. Toward Practical Implementations

For small DG instances such as the grids studied in Section 3, one may efficiently compute equilibrium bisection strategies by implicit enumeration of the game tree. But when the input graph is too large for exact equilibrium computation, how should the players bisect on their turns? This section provides heuristic strategies for both the bisection (Section 4.1) and ICYF (Section 4.3) protocols. Section 4.2 compares the strategies and outcomes of the bisection heuristic applied to the grid instances of Section 3.3 against the exact equilibria. Section 4.4 applies the bisection and ICYF heuristics to a congressional redistricting instance in Iowa with a detailed case study on the heuristic strategies. Finally, Section 4.5 demonstrates the scalability of the bisection heuristic by applying it to larger congressional redistricting instances in 18 states.

4.1. Discrete Geometric Bisection Heuristic

The bisection heuristic for the DG setting is motivated by the game tree of the bisection protocol in the CN setting. If a player had unlimited time and computational resources, they could build the entire DG game tree, then pick the bisection that gives them the most districts in each of their turns, assuming their opponent's play is optimal. As shown in Section 3.3, implicit enumeration of the game tree can take over two hours for a small 6×6 grid. Consequently, even the smallest practical instances are unlikely to be tractable using an implicit enumeration approach.

Rather than attempting to look ahead at future rounds, the bisection heuristic replicates the vote-share allocation strategy in the CN setting, which essentially packs the opponents voters' into one piece, and cracks their voters in the other piece. The choice to imitate the CN vote-share allocation is motivated by the similarity between the CN and DG optimal first-round bisections in the small grid instances (as summarized in Table 2).

Furthermore, preliminary investigations (see Appendix B in the e-companion) revealed that when the physical shapes of the bisected pieces are not *compact*, or round-shaped, there may not be feasible bisections in future rounds, as previously depicted in Figure 7b. To prevent such infeasibility, the bisection heuristic finds the most compact bisection that achieves the target vote-shares from the CN setting in the two pieces. There are many ways to quantify compactness (Duchin and Tenner 2020); we choose to minimize the number of *cut-edges*, that is, edges whose end points are in different pieces. The cut-edge compactness measure is a popular discrete compactness score that agrees well with visual intuition (Validi and Buchanan 2021, DeFord et al. 2021b). Note the compactness objective might unintentionally bias the pieces in favor of a party whose spatial voter distribution benefits from compact pieces (Chen and Rodden 2013).

The remainder of this section explains how the heuristic determines the target vote-shares and finds a bisection. In a given call to RECURSIVEBISECT in Algorithm 1, player $d \in \{1,2\}$ receives a subgraph $H \subseteq G$ to be divided into two parts totaling $m \in [n]$ districts (e.g., in the first round, $d=1$, $H=G$, and $m=n$.) Recall from Section 3.1 that each unit $i \in V(H)$ has population p_i and player vote counts $v_{1,i}$, $v_{2,i}$. Let $\phi = 3-d$ denote the other player's index, and let $v_d(H) = m \cdot \left(\sum_{i\in V(H)} v_{d,i}\right) / \left(\sum_{i\in V(H)}(v_{d,i}+v_{\phi,i})\right)$ denote the total vote-share in H for player d. By Definition 1, the maximum number of districts that player d can win in the CN setting with vote-share $v_d(H)$ is given by

$$j^* = \underset{j\in[m]}{\arg\max}\{t_{m,j} : t_{m,j} \le v_d(H)\}. \tag{7}$$

Player d wins j^* districts in the CN setting by winning k^* districts from the piece of size a and j^*-k^* districts from the piece of size b, where k^* is the argmin of the recursive threshold definition in Lemma 2. Hence, when applying the CN strategy in the DG setting, player d targets vote-shares of $v_d^a = a - t_{a,a-k^*+1}$ and $v_d^b = b - t_{b,b-(j^*-k^*)+1}$ in the pieces of sizes a and b, respectively, expecting to win j^* districts total. It is possible that the structure of the graph and the distribution of vote counts in its units allow player d to win more than j^* districts (e.g., if $v_{d,i} > v_{\phi,i}$ in every unit i, then player d will win every district), but in this heuristic, player d pessimistically assumes that j^* is an upper bound on the number of districts they can hope to win in the DG setting. This sacrifice is necessary to avoid the computationally expensive task of looking ahead in the DG game tree.

If player d begins with a total vote-share of at least $v_d^a + v_d^b$, then splitting the vote-share into v_d^a and v_d^b is always feasible in the CN setting and yields j^* districts. In the DG setting, however, such a split may not be feasible. To address this, a feasibility mixed-integer program (MIP) named Bisection is designed to check whether there exists a feasible bisection with at least v_d^a in the side of size a and v_d^b in the side of size b. If the MIP is infeasible, then j^* is iteratively decremented by 1 until the MIP yields a feasible solution. This assumes that player d will, upon realizing they cannot win j^* districts, attempt to win one fewer. Note that ties go to the drawing player throughout this section. In contrast, the CN setting awards ties to the first player. Algorithm 3 describes this bisection heuristic to be used in Line 9 of Algorithm 1.

Algorithm 3: Bisection protocol for the discrete geometric setting implementing Line 9 of Algorithm 1

Algorithm BISECT (H, m, d)
    Compute j^*, k^* using CN thresholds (Lemma 2) and Equation (7)
    while Bisection(H, m, j^*, k^*, d) is infeasible **do**
        $j^* \leftarrow j^* - 1$
        Recompute k^* using Lemma 2
    end
    return feasible solution of Bisection(H, m, j^*, k^*, d) as vertex sets V_1, V_2

We now define the Bisection(H, m, j^*, k^*, d) feasibility MIP. Let $\overline{P}_a := (a/m) \cdot (\sum_{i \in V(H)} p_i)$ be the ideal total population of the piece of size a (i.e., $p(V_1)$). To support flow-based contiguity constraints adapted from Shirabe (2009), let $\overrightarrow{H} = (V(H), \overrightarrow{E}(H))$ be a bidirected version of H (i.e., $\overrightarrow{E} = \bigcup_{uv \in E} \{\vec{uv}, \vec{vu}\}$), and let $N_H(i) := \{j \in V(H) : ij \in E(H)\}$ be the neighbors of unit i in H.

The variables for the MIP are

- the assignment indicators $x_i = 1$ if and only if $i \in V(H)$ is assigned to V_1 (otherwise, i is in V_2),
- the cut-edge indicators $y_{ij} = 1$ for $ij \in E(H)$ if and only if exactly one endpoint is assigned to V_1,
- the flows $g_{i,j}^k$ = amount of directional flow on edge $ij \in \overrightarrow{E}(H)$ for piece $k \in \{1, 2\}$ (adapted from Shirabe (2009)) and
- the sink indicators $\beta_{i,k} = 1$ if and only if $i \in V(H)$ is the sink for piece $k \in \{1, 2\}$ (adapted from Shirabe (2009)).

The x variables fully determine $V_1 = \{i \in V(H) : x_i = 1\}$ and $V_2 = \{i \in V(H) : x_i = 0\}$, which are returned at the end of Algorithm 3 if a feasible solution is found. The y variables are used in formulating the compactness objective which minimizes the number of cut-edges. The flow and sink variables enforce contiguity by directing at least one unit of flow from each unit in a district to its district's sink through other units in the district; see Shirabe (2009) for details. The program Bisection(H, m, j^*, k^*, d) is then

$$\text{minimize} \quad \sum_{ij \in E(H)} y_{i,j} \tag{8a}$$

$$
\begin{aligned}
\text{subject to}\quad & (1-\tau)\overline{P}_a \le \sum_{i\in V(H)} p_i x_i \le (1+\tau)\overline{P}_a, && && \text{(8b)}\\
& \sum_{i\in V(H)} v_{d,i}x_i \ge \frac{v_d^a}{a}\sum_{i\in V(H)}(v_{d,i}+v_{\phi,i})x_i, && && \text{(8c)}\\
& \sum_{i\in V(H)} v_{d,i}(1-x_i) \ge \frac{v_d^b}{b}\sum_{i\in V(H)}(v_{d,i}+v_{\phi,i})(1-x_i), && && \text{(8d)}\\
& \sum_{i\in V(H)} \beta_{i,k} = 1 && \forall k\in\{1,2\}, && \text{(8e)}\\
& \beta_{i,1} \le x_i && \forall i\in V(H), && \text{(8f)}\\
& \beta_{i,2} \le 1-x_i && \forall i\in V(H), && \text{(8g)}\\
& \sum_{j\in N_H(i)} (g^1_{i,j}-g^1_{j,i}) \ge x_i - |V(H)|\,\beta_{i,1} && \forall i\in V(H), && \text{(8h)}\\
& \sum_{j\in N_H(i)} (g^2_{i,j}-g^2_{j,i}) \ge (1-x_i) - |V(H)|\,\beta_{i,2} && \forall i\in V(H), && \text{(8i)}\\
& \sum_{j\in N_H(i)} g^1_{i,j} \le |V(H)|\,x_i && \forall i\in V(H), && \text{(8j)}\\
& \sum_{j\in N_H(i)} g^2_{i,j} \le |V(H)|\,(1-x_i) && \forall i\in V(H), && \text{(8k)}\\
& y_{ij} \le x_i + x_j && \forall ij\in E(H), && \text{(8l)}\\
& y_{ij} \le 2 - x_i - x_j && \forall ij\in E(H), && \text{(8m)}\\
& y_{ij} \ge x_i - x_j && \forall ij\in E(H), && \text{(8n)}\\
& y_{ij} \ge -x_i + x_j && \forall ij\in E(H), && \text{(8o)}\\
& x_i, y_{ij}, \beta_{i,k} \in \{0,1\} && \forall i\in V(H), k\in\{1,2\}, j\in N_H(i), && \text{(8p)}\\
& g^k_{i,j} \ge 0 && \forall i\in V(H), k\in\{1,2\}, j\in N_H(i). && \text{(8q)}
\end{aligned}
$$

Note that the indices j and k used in defining the MIP are not related to the variables j^* and k^* introduced above to define the target seat-shares. The objective in (8a) minimizes the number of cut-edges. Constraints (8b) ensure population balance in the smaller (or equal size) side, V_1. Note that constraints (8b) generalize the classical partition problem, which is NP-complete (Garey and Johnson 1979). Furthermore, constraints (8b) also ensure population balance in the larger (or equal size) side, V_2, because the absolute population deviations in V_1 and V_2 are equal in magnitude but are divided by a and $b \ge a$ to obtain the relative deviations. Constraints (8c) and (8d) ensure that player d has at least v_d^a vote-share in V_1 and at least v_d^b vote-share in V_2. Constraints (8e)-(8k) are flow-balance constraints that ensure that each piece is contiguous, adapted from Shirabe (2009). Constraints (8e) ensure that each piece has a single unit serving as a sink for flow, and constraints (8f)-(8g) ensure that a unit i can serve as the sink of piece k only if i is assigned to piece k.

Constraints (8h)-(8i) ensure that for every unit i assigned to piece k, the net piece-k-flow leaving i is at least 1 if i is not chosen as sink; the net piece-k-flow leaving i is unbounded if i is chosen as a sink. Constraints (8j)-(8k) ensure a unit i sends out flow in piece k only if i is assigned to k. Constraints (8l)-(8o) define the cut-edge indicator $y_{ij} = x_i \oplus x_j$ for each $ij \in E(H)$ with respect to the assignment variables. Constraints (8p)-(8q) characterize the binary and continuous variables.

4.2. Comparing Heuristic and Optimal Bisection on Small Grids

The bisection heuristic in Section 4.1 offers a practical, but not necessarily optimal, method for each player to bisect in their round. This section explores the extent to which the heuristic differs from optimal play on the small grid instances of Section 3.3. Table 3 reports the utilities from each combination of voter distribution, first player, and number of districts ($n = 4$ or $n = 6$). A table entry of "3D, 1T, 2R" indicates three wins for party D, one tie, and two wins for party R. The instances with $n = 4$ and $n = 6$ districts require two and three rounds of bisection, respectively.

The heuristic utilities align with the optimal utilities for the DG setting (in Table 1) when a 6×6 grid is partitioned into $n = 6$ districts, assuming that ties are awarded to the first player. Among the 6×6 grids, when $n = 6$, the heuristic utilities align with the optimal DG utilities in all three grid instances; when $n = 4$, the optimal strategy awards the first player one extra district. Furthermore, the heuristic utilities align with the utilities from the CN setting irrespective of n. Among the 4×6 grids, the results are mixed; compared to the optimal DG utilities, the heuristic yields one more district to the first player in 5 instances, one less district in 2 instances, and the same seat-share in 5 instances.

The outcomes and optimal strategies differ substantially between the heuristic and optimal bisection strategies. Hence real-world players would likely apply more sophisticated strategies, incorporating political geography and feasibility issues, that achieve better results than the CN threshold-based heuristic. Such strategies would require additional computational resources and algorithmic development, and are therefore beyond the scope of this paper, but Section 6 addresses possible directions.

Table 3 Utilities from applying the bisection heuristic on small grid instances.

Spatial Voter Distribution	$n = 4$ Districts		$n = 6$ Districts	
	D first	R first	D first	R first
CN setting	2D, 2R	2R, 2D	4D, 2R	4R, 2D
4x6_lm	1D, 2T, 1R	1R, 2T, 1D	1D, 4T, 1R^M	2R, 2T, 2D
4x6_ns	2D, 2R^L	2R, 2D^L	2D, 2T, 2R	2R, 2T, 2D
4x6_oc	4T^M	4T^M	1D, 4T, 1R^M	1R, 4T, 1D^M
6x6_lm	2D, 2R^L	2R, 2D^L	4D, 2R	2R, 2T, 2D
6x6_ns	2D, 2R^L	2R, 2D^L	2D, 2T, 2R	2R, 2T, 2D
6x6_oc	2D, 2R^L	2R, 2D^L	2D, 2T, 2R	2R, 2T, 2D

Superscripts ^{M} and ^{L} indicate one *more* or one *less* district is awarded to the first player compared to the optimal DG utilities in Table 1, assuming that ties are awarded to the first player.

Optimal and heuristic bisection strategies also differ in the computational time required to run the protocol. The implicit enumeration technique takes 1–11 seconds for each of the 4×6 instances and 49–8,379 seconds for each of the 6×6 instances. The bisection heuristic, on the other hand, takes less than 1 second to solve each of the 24 small grid instances. As will be shown in Sections 4.4 and 4.5, the bisection heuristic can therefore complete the protocol for larger instances.

4.3. Discrete Geometric I-Cut-You-Freeze Heuristics

The heuristic strategies for ICYF in the DG setting are based on the optimal strategies presented in Pegden et al. (2017) for ICYF in the CN setting. Algorithm 4 presents ICYF in the DG setting with only the drawing step of Line 9 and the freezing step of Line 10 left undetermined. Note that a heuristic implementation of ICYF requires both a freezing heuristic and a drawing heuristic, whereas the heuristic implementation of the bisection protocol needs only a drawing heuristic. In the CN setting, the optimal freezing strategy is to choose a district the freezing player wins by the least margin; if no such district exists, the freezing player should choose a district in which their opponent wins by the greatest margin (Pegden et al. 2017). Since the freezing strategy only depends on victory margins, it can be directly applied in the DG setting.

Algorithm 4: I-cut-you-freeze protocol for partitioning a graph into a set $\mathcal{D}$ of n districts, adapted from Pegden et al. (2017)

1 **Algorithm** IcyfProtocol (graph $G=(V,E)$, number of districts $n \in \mathbb{Z}^+$)
2     $\mathcal{D} \leftarrow$ RecursiveIcyf (G, n, 1)
3   **return** $\mathcal{D}$

4 **Procedure** RecursiveIcyf (subgraph H, number of districts m, player $d \in \{1,2\}$)
5   **if** $m=1$ **then**
6   **return** $\{V(H)\}$
7   **end**
8     $\phi \leftarrow 3-d$ `// Other player index`
9     $\{V_i\}_{i=1}^m \leftarrow$ player d partitions H into m population-balanced districts
10     $i^* \leftarrow$ player ϕ chooses index of district to freeze
11     $d \leftarrow 3-d$ `// Players switch roles for next round`
12     $\mathcal{D} \leftarrow$ RecursiveIcyf ($H[V(H) \setminus V_{i^*}]$, $m-1$, d)
13   **return** $\mathcal{D} \cup \{V_{i^*}\}$

Similarly, the ICYF drawing strategy may be adapted from the CN setting. In the CN setting, the drawing player should maximize the number of districts they win while maximally packing their opponent's voters in the districts they lose (Pegden et al. 2017). The DG setting inherently limits packing, but one may pursue the same objectives by formulating and solving a MIP. Applying this drawing strategy for ICYF is a heuristic rather than an exact algorithm because it does not attempt to look ahead to future rounds. The structure of the subgraph in later rounds, which is partially determined by the drawing and freezing of the current round, may prevent partitions that mirror the optimal strategies for ICYF in the CN setting. But among the class of restricted heuristics that do not explore further down the game tree, this drawing heuristic for ICYF is a reasonable approach that captures the essential packing and cracking tactics from the CN setting.

The ICYF MIP for the heuristic drawing strategy used by player d in Line 9 of Algorithm 4 is constructed from the parameters H, m, and d of RecursiveIcyf. Define ϕ, $N_H(i)$, and $\overrightarrow{H}$ as in Section 4.1. The variables for the ICYF MIP are

- the assignment indicators $x_{i,k} = 1$ if and only if $i \in V(H)$ is assigned to district $k \in [m]$,
- the district winner indicators $z_k = 1$ if and only if player d wins district $k \in [m]$,
- the opponent's unit winner indicators $u_{i,k} = 1$ if and only if $i \in V(H)$ is assigned to district $k \in [m]$ and player ϕ wins district k,
- the opponent's winning vote-shares divided by m,

$$q_k = \begin{cases} \frac{1}{m} \cdot (\text{player } \phi \text{ vote-share in district } k) & \text{if player } \phi \text{ wins district } k \in [m] \\ 1 & \text{otherwise,} \end{cases}$$

- the minimum of all the q_k values, $q = \min\{q_k : k \in [m]\}$ (which is 1 if player d wins every district),
- the flows $g^k_{i,j}$ = amount of directional flow on edge $(i,j) \in \overrightarrow{E}(H)$ for district $k \in [m]$ (adapted from Shirabe (2009)), and
- the sink indicators $\beta_{i,k} = 1$ if and only if $i \in V(H)$ is the sink for district $k \in [m]$ (adapted from Shirabe (2009)).

In defining the winner indicators z, a tie is awarded to the drawing player. The flow and sink variables enforce contiguity by directing at least one unit of flow from each unit in a district to its district's sink through other units in the district; see Shirabe (2009) for details. The MIP ICYF(H, m, d) is then

$$\text{maximize} \quad \left(\sum_{k \in [m]} z_k \right) + q \tag{9a}$$

$$\text{subject to} \quad \sum_{k \in [m]} x_{i,k} = 1 \qquad \forall i \in V(H), \tag{9b}$$

$$(1-\tau)\overline{P} \leq \sum_{i \in V(H)} p_i x_{i,k} \leq (1+\tau)\overline{P} \quad \forall k \in [m], \tag{9c}$$

$$-M(1-z_k) \leq \sum_{i \in V(H)} (v_{d,i} - v_{\phi,i})\, x_{i,k} \qquad \forall k \in [m], \tag{9d}$$

$$\sum_{i\in V(H)} (v_{d,i} - v_{\phi,i})\, x_{i,k} \le M z_k \qquad \forall k \in [m], \tag{9e}$$

$$u_{i,k} \le x_{i,k} \qquad \forall i \in V(H), k \in [m], \tag{9f}$$

$$u_{i,k} \le 1 - z_k \qquad \forall i \in V(H), k \in [m], \tag{9g}$$

$$u_{i,k} \ge x_{i,k} - z_k \qquad \forall i \in V(H), k \in [m], \tag{9h}$$

$$q_k = z_k + \frac{1}{m\overline{P}} \sum_{i\in V(H)} v_{\phi,i} u_{i,k} \qquad \forall k \in [m], \tag{9i}$$

$$q \le q_k \qquad \forall k \in [m], \tag{9j}$$

$$\sum_{i\in V(H)} \beta_{i,k} = 1 \qquad \forall k \in [m], \tag{9k}$$

$$\beta_{i,k} \le x_{i,k} \qquad \forall i \in V(H), k \in [m], \tag{9l}$$

$$\sum_{j\in N_H(i)} (g^k_{i,j} - g^k_{j,i}) \ge x_{i,k} - |V(H)|\, \beta_{i,k} \qquad \forall i \in V(H), k \in [m], \tag{9m}$$

$$\sum_{j\in N_H(i)} g^k_{i,j} \le |V(H)|\, x_{i,k} \qquad \forall i \in V(H), k \in [m], \tag{9n}$$

$$x_{i,k}, z_k, u_{i,k}, \beta_{i,k} \in \{0,1\} \qquad \forall i \in V(H), k \in [m], \tag{9o}$$

$$q, q_k \in [0,1] \qquad \forall k \in [m], \tag{9p}$$

$$g^k_{i,j} \ge 0 \qquad \forall i,j \in V(H), k \in [m]. \tag{9q}$$

The objective in (9a) maximizes the number of districts won by player d with a secondary objective of maximizing the minimum margin among the districts won by player ϕ. The minimum margin is scaled down by m to ensure q is fractional, and therefore a secondary objective to the integral sum of z_k indicators, whenever player ϕ wins at least one district. These two objectives align with those of the drawing player in the CN setting. Constraints (9b) assign each unit to exactly one district. Constraints (9c) ensure that each district population lies within the permissible range. Constraints (9d)-(9e) relate the binary variables x and z for some $M > \max\{0, \sum_{i\in V}(v_{d,i} - v_{\phi,i})\}$. Constraints (9f)-(9h) linearize the quadratic constraints given by $u_{i,k} = x_{i,k}(1 - z_k)$. When there is a tie in district k, constraints (9f)-(9h) allow z_k to be either 0 or 1. In this case, maximizing the objective function in (9a) sets z_k to be 1, awarding the district to the drawing player. Constraints (9i) use the vote counts and the $u_{i,k}$ indicator variables to define the victory margin for each district that player ϕ wins, dividing by the total district population $m\overline{P}$ rather than the total number of voters given by $\sum_{i\in V(H)}(v_{d,i} + v_{\phi,i})x_{i,k}$ in each district $k \in [m]$ to keep the constraints linear. Dividing by $m\overline{P}$ rather than simply $\overline{P}$ prevents q_k from exceeding 1 in the case when district k has population greater than $\overline{P}$. Scaling victory margins down by a factor of m does not change their ordering, so maximizing the minimum scaled-down victory margin is equivalent to maximizing the minimum original victory margin. Constraints (9j) define the (scaled-down) minimum margin among the districts won by player ϕ. Constraints (9k)-(9n) are flow-balance constraints that ensure each district is contiguous, adapted from Shirabe (2009). Constraints (9k) ensure each district has a single unit serving as a sink for flow, and constraints (9l) ensure that unit i can serve as the sink of district k only if i is assigned to district k. Constraints (9m) ensure that for every unit i assigned to district k, the net district-k-flow leaving i is at least 1 if i is not chosen as sink; the net district-k-flow leaving i is unbounded if i is chosen as a sink. Constraints (9n) ensure that unit i sends out district-k-flow only if i is assigned to district k. Hence, if district k is noncontiguous, these constraints ensure that the flow from every unit assigned to k will not reach the sink of district k. Constraints (9o)-(9q) enforce the binary or continuous bounds on the variables.

4.4. Computational Results: Iowa Case Study

This section presents computational results for the application of the discrete protocol implementations from Sections 4.1 and 4.3 to the state of Iowa. Iowa has only four congressional districts and requires that none of its 99 counties are split across districts, making it an ideal state for a small-scale case study on a practical districting scenario. The vote counts for the two major parties, Democratic (D) and Republican (R), are derived by averaging the 2016 and 2020 presidential election results aggregated from voting precincts to each county (MIT Election Data and Science Lab 2020). The county-level vote counts are then disaggregated to census tracts proportional to their populations. Based on this dataset, the normalized vote-share for R in Iowa is 55%. Figure 8 categorizes the vote counts in each county, giving an overview of the political geography of Iowa. The population balance threshold τ is set to 1%. All computations were performed on a 3.1 GHz Intel Core i5-2400 CPU machine using a single core. The mixed integer programs were solved using the IBM CPLEX Optimizer 12.10.

Figure 8 County-level party support based on 2016 and 2020 presidential elections in Iowa

(a) County-level party support for D and R

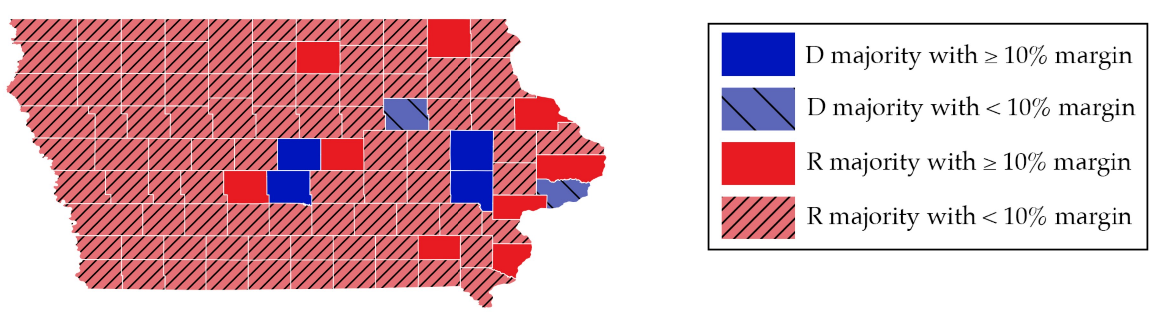

The two redistricting protocols are each applied to Iowa twice: once for each player (D or R) taking the first turn. Further, the two protocols are compared with respect to fairness metrics and computational times. Note the results presented in this section are obtained using the particular implementation of the heuristics from Sections 4.1 and 4.3 for both players, and may not reflect the (unknown) optimal play for the two protocols. See Appendix C for a detailed, round-by-round breakdown of heuristic play in Iowa, including figures depicting the final district plans produced for the four instances (each protocol with each player drawing first).

The geographical distribution of each party's voters impacts the ability of each player to pack and crack the opponent's voters. In particular, Iowan D voters are generally clustered in a few urban counties, and Iowan R voters are generally spread across rural counties. As evidence of this clustering, D has a majority vote count in only 6 of the 99 counties in Iowa despite having 45% normalized vote-share.

When both players use the heuristics of Sections 4.1 and 4.3, neither protocol exhibits a first-player advantage, as evidenced by the seat-share results in Table 4. In fact, the bisection heuristic yields a second-player advantage when D draws first, giving D a seat-share of only 0.25, half the seat-share D receives when R draws first. ICYF yields the same seat-share regardless of who draws first. The asymmetric political geography may partially explain the difference in outcomes between the two protocols. Note that these observations are anecdotal to the Iowa instance and might not generalize to districting instances in other states.

Table 4 District plans obtained from the bisection protocol, I-Cut-You-Freeze protocol and Iowa's 115th congressional districts compared across three fairness metrics.

Fairness metric	Bisection		I-Cut-You-Freeze		115th CD
	D first	R first	D first	R first	
Seat-share for D	0.25	0.5	0.5	0.5	0.25
SL Index	16.70%	0.88%	0.87%	0.87%	7.92%
Efficiency Gap	15.49%	9.57%	8.39%	9.24%	14.68%
Partisan Asymmetry	2.84%	2.27%	2.16%	0.20%	7.95%
Computational Time (s)	3	3	6,183	935	-

Table 4 also reports the three fairness metrics defined in Section 2.2—Sainte Laguë (SL) index, efficiency gap, and partisan asymmetry—for the district plan produced from each protocol with each player taking the first turn. The metrics for Iowa's 115th congressional districts (115th CD), used for 2017-2019, are also included. Recall that the SL index is a measure of disproportionality, efficiency gap captures the difference in wasted votes between the two players, and (partisan) asymmetry is the total area between the seats-votes curves of the two players. Hence, smaller values of all three metrics indicate fairer districts. For the bisection protocol, the district plan obtained when R goes first produces fairer districts (in all three metrics) than the plan when D goes first.

For the ICYF protocol, the district plan obtained when D goes first has a slightly smaller efficiency gap, although the district plan obtained when R goes first has a smaller partisan asymmetry.

The execution of the heuristic bisection protocol is drastically faster than that of the heuristic ICYF protocol. The observed CPU time for dividing a region into two pieces, as in the bisection protocol, is significantly lower than for dividing it into three or more pieces, as in the ICYF protocol. Additionally, bisection involves $\lceil \log_2 n \rceil$ rounds, whereas ICYF requires $n-1$ rounds. These computational advantages suggest the bisection heuristic may be more likely than the ICYF heuristic to be tractable for larger redistricting instances.

4.5. Computational Results: Larger Instances

The Iowa case study demonstrates the protocol heuristics can be efficiently applied to a practical instance and produce fairer district plans than Iowa's 115th congressional districts. We now examine the robustness and scalability of the bisection heuristic by applying it to 18 congressional redistricting instances in U.S. states at the census-tract level. The smallest instance is Utah with 716 census tracts and 4 districts, and the largest is Missouri with $1,654$ census tracts and 8 districts. Among these states, six are D-leaning, that is, have a D voter majority. The most D-leaning state is Maryland (65.6% D voters) and the most R-leaning state is Oklahoma (68.1% R voters).

The MIP-based heuristic bisection protocol is applied to each instance with $\tau = 1\%$ and a one-hour MIP solver time limit as in Section 4.4. The four-district states have two rounds of bisection, and the others have three.

The bisection heuristic produces feasible district plans for all 18 states. Table 5 summarizes key information for each instance, including the fairness metrics (SL index, efficiency gap, and partisan asymmetry) of the final district plans obtained when D bisects first. Table 6 provides the same information for when R bisects first. See Appendix C.3 for figures depicting the final district plans resulting from all 36 bisection instances. In 12 states, the seat-share does not depend on who bisects first. It is interesting to note that the bisection heuristic produces the same Oklahoma district plan regardless of the first player. Furthermore, four states (Arkansas, Mississippi, Iowa and South Carolina) give R a greater seat-share when D goes first, and two states (Alabama and Wisconsin) give R a lower seat-share when R goes first. This suggests the bisection heuristic gives a slight seat-share advantage to the second player by not allowing the first player to draw convoluted districts that restrict future bisection strategies.

Political geography impacts the utilities that the heuristic awards to the players. For example, all five districts in Connecticut (41.2% R voters) are won by D, regardless of who bisects first, whereas all five districts of Oklahoma (68.2% R voters) are awarded to R. These differences are due to the spatial distribution of voters across the states.

The combined effects of the DG setting and the heuristic bisection strategies can be inferred by comparing against the CN setting results. For each instance, the seat-share from the CN setting with the same normalized vote-share and number of districts is reported in Tables 5 and 6 (column 10). Comparing R's seat-share from the DG heuristic (column 6) to R's seat-share from the CN setting, it is clear that the DG setting gives R a greater advantage for these instances when the players use the heuristic strategies presented in this study. When D draws first, compared to the CN setting, the DG setting gives R a greater seat-share in 13 states, a lesser seat-share in four states, and the same seat-share in one state. However, the results are mixed when R draws first; compared to the CN setting, the DG setting gives R a greater seat-share in six states, a lesser seat-share in five states, and the same seat-share in seven states. Overall, R benefits from the DG setting, especially as the second player.

Tables 5 and 6 also report the total time spent executing Algorithm 1 for each instance. States with more census tracts and districts tend to reach the 1 hour time limit more often than those with fewer. Curiously, the CPU time is generally much lower when R draws first than when D draws first. While the exact cause of this gap is unclear, the typical pattern of urban clusters of Democrats and rural/suburban concentrations of Republicans may render the first round bisection MIP easier when R draws first due to increased opportunities for packing and cracking.

Table 5 Bisection heuristic results when D draws the first round.

State	Number of census tracts	Number of districts	Normalized R vote-share	R seat-share	CN R seat-share	CPU Time (s)	SL Index	EG*	PA
Utah	716	4	61.30%	75.00%	50.00%	615	7.90%	1.90%	4.80%
Nevada	778	4	48.70%	25.00%	50.00%	71	22.60%	-18.80%	4.40%
Arkansas	823	4	64.30%	100.00%	75.00%	188	55.60%	21.50%	1.80%
Kansas	829	4	59.10%	75.00%	50.00%	288	10.40%	4.40%	3.00%
Mississippi	873	4	58.70%	100.00%	50.00%	$3{,}716^{T}$	70.30%	32.60%	2.10%
Connecticut	883	5	41.20%	0.00%	40.00%	104	70.10%	-32.50%	0.70%
Iowa	895	4	54.70%	75.00%	50.00%	90	16.70%	15.60%	2.50%
Oregon	1,000	5	42.60%	20.00%	40.00%	257	20.90%	-15.10%	3.40%
Oklahoma	1,205	5	68.10%	100.00%	60.00%	562	46.90%	13.90%	2.90%
Kentucky	1,305	6	64.40%	83.30%	66.70%	501	15.70%	3.50%	3.20%
South Carolina	1,321	7	56.60%	85.70%	42.90%	$7{,}591^{TT}$	34.50%	23.00%	1.30%
Louisiana	1,384	6	59.80%	83.30%	66.70%	4,041	23.00%	14.30%	7.60%
Alabama	1,436	7	63.60%	71.40%	57.10%	$10{,}146^{TT}$	2.70%	-5.30%	3.90%
Colorado	1,436	7	45.00%	42.90%	42.90%	$10{,}921^{TTT}$	0.20%	3.50%	4.50%
Maryland	1,474	8	34.40%	25.00%	25.00%	1,067	3.90%	7.60%	8.80%
Minnesota	1,505	8	47.60%	50.00%	37.50%	$7{,}238^{TT}$	0.20%	3.70%	1.70%
Wisconsin	1,541	8	50.00%	50.00%	25.00%	$7{,}610^{TT}$	0.00%	1.20%	5.60%
Missouri	1,654	8	60.60%	75.00%	50.00%	$9{,}583^{TT}$	8.60%	1.10%	8.50%

The bisection protocol with the MIP-based heuristic (Section 4.1) is applied to U.S. states with 4 to 8 congressional districts at the census tract level. Abbreviations include SL for Sainte-Laguë, EG for efficiency gap, PA for partisan asymmetry, and CN for the continuous non-geometric setting. *Positive efficiency gap is in favor of R; negative is in favor of D. Superscripts ^{T}, ^{TT}, and ^{TTT} indicate one, two, or three iterations of the MIP solver reached the one-hour time limit.

Table 6 Bisection heuristic results when R draws the first round.

State	Number of census tracts	Number of districts	Normalized R vote-share	R seat-share	CN R seat-share	CPU Time (s)	SL Index	EG*	PA
Utah	716	4	61.30%	75.00%	50.00%	72	7.90%	2.10%	11.40%
Nevada	778	4	48.70%	25.00%	50.00%	51	22.60%	-18.80%	4.30%
Arkansas	823	4	64.30%	75.00%	75.00%	335	5.00%	-5.00%	6.00%
Kansas	829	4	59.10%	75.00%	50.00%	76	10.40%	2.50%	5.90%
Mississippi	873	4	58.70%	50.00%	50.00%	639	3.10%	-19.10%	2.20%
Connecticut	883	5	41.20%	0.00%	40.00%	87	70.10%	-32.50%	0.70%
Iowa	895	4	54.70%	50.00%	50.00%	81	0.90%	-9.10%	3.40%
Oregon	1,000	5	42.60%	20.00%	40.00%	136	20.90%	-15.10%	3.40%
Oklahoma	1,205	5	68.10%	100.00%	80.00%	350	46.90%	13.90%	2.90%
Kentucky	1,305	6	64.40%	83.30%	66.70%	333	15.70%	3.40%	3.00%
South Carolina	1,321	7	56.60%	57.10%	57.10%	460	0.00%	-6.50%	2.30%
Louisiana	1,384	6	59.80%	83.30%	66.70%	398	23.00%	14.30%	7.50%
Alabama	1,436	7	63.60%	85.70%	71.40%	$9{,}730^{TT}$	21.10%	8.90%	2.70%
Colorado	1,436	7	45.00%	42.90%	57.10%	868	0.20%	1.50%	2.90%
Maryland	1,474	8	34.40%	25.00%	37.50%	1,267	3.90%	7.80%	7.20%
Minnesota	1,505	8	47.60%	50.00%	50.00%	$3{,}843^{T}$	0.20%	3.90%	2.10%
Wisconsin	1,541	8	50.00%	75.00%	75.00%	671	25.00%	25.00%	9.70%
Missouri	1,654	8	60.60%	75.00%	75.00%	$8{,}090^{TT}$	8.60%	1.00%	6.70%

The bisection protocol with the MIP-based heuristic (Section 4.1) is applied to U.S. states with 4 to 8 congressional districts at the census tract level. Abbreviations include SL for Sainte-Laguë, EG for efficiency gap, PA for partisan asymmetry, and CN for the continuous non-geometric setting. *Positive efficiency gap is in favor of R; negative is in favor of D. Superscripts ^{T}, ^{TT}, and ^{TTT} indicate one, two, or three iterations of the MIP solver reached the one-hour time limit.

5. Limitations

Limitations of this study include both computational hindrances in the experiments and philosophical drawbacks of fair division-inspired redistricting protocols.

One technical limitation is the computational effort required to apply the MIP-based bisection protocol heuristic to real-world redistricting instances. In particular, when CPLEX solves the formulations, the weak relaxation bounds result in long computation times for large problem instances. Hence the experiments of Section 4.5 only cover states with 4 to 8 districts (i.e., 2 to 3 rounds of bisection), preventing this heuristic from being applied to—and drawing insights for—larger instances at the scale typically found in congressional and state legislative redistricting.

Another limitation is the observed gap between heuristic and optimal play in the DG setting, not only in strategy but also in outcome. Applying CN-based strategies in the DG setting overlooks the crucial effects of political geography and piece shapes. While the CN-based strategies hinder the opponent by packing and cracking their votes, optimal DG setting strategies likely involve trapping the opponent into drawing unfavorable districts by severely restricting the feasible strategy space in later rounds. Due to the computational difficulty of incorporating future-round effects into DG strategy selection, we leave refinement of DG heuristics as future work, with some possible directions mentioned in Section 6.

On the philosophical side, the bisection protocol and similar fair division-inspired redistricting methods may perpetuate politicians picking their voters (since the parties are the players in each protocol), albeit in a more bipartisan manner, instead of voters picking their politicians. Further work is needed to explore possible negative side effects, such as incentivizing incumbent gerrymandering. Another is the difficulty of enforcing compliance with the Voting Rights Act (VRA): neither bisection nor ICYF guarantees the formation of majority-minority districts. Any required majority-minority districts could be drawn in advance before applying either protocol to the rest of the state, but doing so reintroduces reliance on a neutral third party.

6. Conclusions

This paper presents a bisection protocol for political redistricting in which two players alternately bisect pieces of the state until each piece is a district. This paper also extends the analysis of the I-cut-you-freeze (ICYF) protocol proposed by Pegden et al. (2017), evaluating its fairness with respect to common fairness measures from the political science literature. The simplicity of the bisection and ICYF protocols renders them accessible to legislators, judges, and citizens alike. Allowing both major parties to contribute to the final district plan may reduce accusations of excessive partisan gerrymandering.

Assessing the protocols in both a relaxed, continuous nongeometric (CN) setting and the more realistic discrete geometric (DG) setting, we find both bisection and ICYF achieve reasonable values for proportionality, efficiency gap, and partisan asymmetry. Furthermore, comparing the protocols' utility curves against that of the one-player-draws protocol, the current status quo in 33 states, reveals both bisection and ICYF greatly reduce the first-player advantage. A case study of Iowa congressional redistricting further supports this conclusion. Bisection tends to produce more proportional district plans, and ICYF leads to lower-magnitude efficiency gaps.

Implicit enumeration of the bisection protocol game tree for small grid instances reveals some voter distributions give one party a greater first-player advantage. For census tract-level redistricting in states with 4 to 8 districts, we apply a mixed-integer programming heuristic bisection strategy. The heuristic generates feasible district plans, and results suggest political geography greatly influences protocol outcomes.

Future work could refine the bisection heuristic for the DG setting (e.g., by strengthening the MIP formulation with valid inequalities) and compare against randomized play as a baseline. Comparing the outcome distribution from randomized play to the heuristic outcomes could help quantify the quality of the heuristic. The details of randomized play would require thoughtful development, especially to handle cases when early-round bisections lead to infeasible subproblems. But an approach based on balanced cuts of random spanning trees, currently successful in the

Recombination Markov chains for redistricting (DeFord et al. 2021b), offers a reasonable starting point for randomized bisection.

This study also generates several open problems. The bisection and ICYF analyses assume players consider all wins equal; how would protocol outcomes shift, in both the CN and DG settings, if players valued 60%–40% landslide victories more than 51%–49% races? How robust are protocol outcomes to small perturbations in unit vote-shares? Answering such questions will help determine whether bisection and ICYF align with real-world goals and preferences for redistricting reform.

Acknowledgments

The authors thank the editor-in-chief, associate editor, and two anonymous referees for their helpful comments and suggestions, which significantly improved the manuscript.

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE – 1746047. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

The fourth author has been supported in part by the Air Force Office of Scientific Research under Grant No. 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.

7. Author Biographies

Ian G. Ludden is a Ph.D. student in Computer Science at the University of Illinois Urbana-Champaign. He has a B.S. in Computer Engineering and Mathematics from Rose-Hulman Institute of Technology. His research interests include algorithmic game theory, combinatorial optimization, and graph algorithms. He has served as a Graduate Fellow for the Computing Research Association Education Committee (CRA-E) since 2020.

Rahul Swamy is a Ph.D. student in Industrial Engineering (Operations Research) at the University of Illinois Urbana-Champaign. His research interests lie in discrete optimization and algorithmic game theory with application focus in domains such as political redistricting, logistical planning, and network games. His work on political redistricting has been recognized by First place in the INFORMS Service Science Best Paper Award 2019, First place in the INFORMS Poster Competition 2019, and as a Finalist at the INFORMS Public Sector Operations Research (PSOR) Best Paper Award 2018. 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 the president of the INFORMS Student Chapter at UIUC (2016-2018).

Douglas M. King, Ph.D, 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 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.

8. Research Story

The authors conceived of the bisection protocol during a research meeting discussing recent political redistricting literature, in particular, the I-cut-you-freeze protocol pre-print. After establishing the theoretical results for the continuous nongeometric setting, they discussed ways to implement both protocols on real-world data, culminating in the Iowa case study and computational experiments with 17 other states.

References

Becker R, Lari I, Lucertini M, Simeone B (1998) Max-min partitioning of grid graphs into connected components. *Networks* 32(2):115–125.

Bernstein M, Duchin M (2017) A formula goes to court: Partisan gerrymandering and the efficiency gap. *Notices of the AMS* 64(9):1020–1024.

Bichot CE, Siarry P, eds. (2013) *Graph Partitioning* (John Wiley & Sons, Inc.).

Cain BE (1985) Simple v. complex criteria for partisan gerrymandering: A comment on Niemi and Grofman. *UCLA Law Review* 33:213.

Chambers CP, Miller AD, Sobel J (2017) Flaws in the efficiency gap. *Journal of Law & Politics* 33:1.

Chataigner F, Salgado LRB, Wakabayashi Y (2007) Approximation and inapproximability results on balanced connected partitions of graphs. *Discrete Mathematics & Theoretical Computer Science* Vol. 9 no. 1(Graph and Algorithms), URL http://dx.doi.org/10.46298/dmtcs.384.

Chen J, Rodden J (2013) Unintentional gerrymandering: Political geography and electoral bias in legislatures. *Quarterly Journal of Political Science* 8:239–269.

Cho WKT (2017) Measuring partisan fairness: How well does the efficiency gap guard against sophisticated as well as simple-minded modes of partisan discrimination? *University of Pennsylvania Law Review Online* 166:17.

Cohen-Addad V, Klein PN, Young NE (2018) Balanced centroidal power diagrams for redistricting. *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems* (ACM), URL http://dx.doi.org/10.1145/3274895.3274979.

Davis J (2017) Elbridge Gerry and the monstrous gerrymander. https://blogs.loc.gov/law/2017/02/elbridge-gerry-and-the-monstrous-gerrymander/, accessed on 04 June 2019.

DeFord D, Dhamankar N, Duchin M, Gupta V, McPike M, Schoenbach G, Sim KW (2021a) Implementing partisan symmetry: Problems and paradoxes. *Political Analysis* 1–20, URL http://dx.doi.org/10.1017/pan.2021.49.

DeFord D, Duchin M, Solomon J (2021b) Recombination: A family of markov chains for redistricting. *Issue 3.1, Winter 2021* URL http://dx.doi.org/10.1162/99608f92.eb30390f.

Duchin M (2018) Gerrymandering metrics: How to measure? What's the baseline? *arXiv preprint* arXiv:1710.08781.

Duchin M, Tenner BE (2020) Discrete geometry for electoral geography. *arXiv pre-print* .

Gallagher M (1991) Proportionality, disproportionality and electoral systems. *Electoral Studies* 10(1):33 – 51.

Garey MR, Johnson DS (1979) *Computers and Intractability: A Guide to the Theory of NP-Completeness* (New York, NY, USA: W. H. Freeman & Co.), ISBN 0716710447.

Goldenberg J, Fisher SD (2019) The Sainte-Laguë index of disproportionality and Dalton's principle of transfers. *Party Politics* 25(2):203–207.

Grofman B (1983) Measures of bias and proportionality in seats-votes relationships. *Political Methodology* 9(3):295–327, ISSN 01622021, URL http://www.jstor.org/stable/25791195.

Grofman B (1985) Criteria for districting: A social science perspective. *UCLA Law Review* 33:77.

Grofman B, King G (2007) The future of partisan symmetry as a judicial test for partisan gerrymandering after LULAC v. Perry. *Election Law Journal: Rules, Politics, and Policy* 6(1):2–35.

Gurnee W, Shmoys DB (2021) Fairmandering: A column generation heuristic for fairness-optimized political districting. *SIAM Conference on Applied and Computational Discrete Algorithms (ACDA21)*, 88–99 (Society for Industrial and Applied Mathematics), URL http://dx.doi.org/10.1137/1.9781611976830.9.

Landau Z, Reid O, Yershov I (2009) A fair division solution to the problem of redistricting. *Social Choice and Welfare* 32(3):479–492.

Levin HA, Friedler SA (2019) Automated congressional redistricting. *ACM Journal of Experimental Algorithmics* 24:1–24, URL http://dx.doi.org/10.1145/3316513.

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, URL http://dx.doi.org/10.1016/j.swevo.2016.04.004.

Lowenstein DH, Steinberg J (1985) The quest for legislative districting in the public interest: Elusive or illusory. *UCLA Law Review* 33:1.

McDonald MP (2007) Regulating redistricting. *PS: Political Science & Politics* 40(4):675–679.

MIT Election Data and Science Lab (2020) County Presidential Election Returns 2000-2020. URL http://dx.doi.org/10.7910/DVN/VOQCHQ, accessed on 04 February 2022.

Montana Districting and Apportionment Commission (2010) Congressional and legislative redistricting criteria. https://leg.mt.gov/content/Committees/Interim/2011-2012/Districting/Other-Documents/1124RWFA-corrected-criteria-updated-2011.pdf, accessed on 26 June 2019.

Niemi RG (1985) Relationship between votes and seats: The ultimate question in poltical gerrymandering. *UCLA Law Review* 33:185.

OEIS Foundation Inc (2022) Entry A172477 in The On-Line Encyclopedia of Integer Sequences. *On-Line Encyclopedia of Integer Sequences* URL https://oeis.org/A172477.

Osborne M (2004) *An Introduction to Game Theory* (New York: Oxford University Press), ISBN 0195128958.

Pegden W, Procaccia AD, Yu D (2017) A partisan districting protocol with provably nonpartisan outcomes. *arXiv preprint* arXiv:1710.08781.

Pennisi A (1998) Disproportionality indexes and robustness of proportional allocation methods. *Electoral Studies* 17(1):3 – 19.

Ramachandran G, Gold D (2018) Using outlier analysis to detect partisan gerrymanders: A survey of current approaches and future directions. *Election Law Journal: Rules, Politics, and Policy* 17(4):286–301.

Ricca F, Scozzari A, Simeone B (2013) Political districting: From classical models to recent approaches. *Annals of Operations Research* 204(1):271–299.

Schutzman Z (2019) zschutzman/enumerator: v0.1.5. URL http://dx.doi.org/10.5281/zenodo.3467675.

Shirabe T (2009) Districting modeling with exact contiguity constraints. *Environment and Planning B: Planning and Design* 36(6):1053–1066, URL http://dx.doi.org/10.1068/b34104.

Stephanopoulos NO, McGhee EM (2015) Partisan gerrymandering and the efficiency gap. *The University of Chicago Law Review* 82(2):831–900.

Stephanopoulos NO, McGhee EM (2018) The measure of a metric: The debate over quantifying partisan gerrymandering. *Stanford Law Review* 70:1503.

Supreme Court of the United States (1986) *Davis v. Bandemer*. 478 U.S. 109 (1986); 106 S. Ct. 2797; 92 L. Ed. 2d 85.

Supreme Court of the United States (2019) *Rucho v. Common Cause*. 588 U.S. ___ (2019).

Swamy R, King DM, Jacobson SH (2022) Multiobjective optimization for politically fair districting: A scalable multilevel approach. *Operations Research* URL http://dx.doi.org/10.1287/opre.2022.2311.

Tucker-Foltz J (2019) A cut-and-choose mechanism to prevent gerrymandering. *arXiv preprint* arXiv:1802.08351v3.

Validi H, Buchanan A (2021) Political districting to minimize cut edges. *Optimization Online pre-print* URL http://www.optimization-online.org/DB_FILE/2021/04/8349.pdf.

Validi H, Buchanan A, Lykhovyd E (2021) Imposing contiguity constraints in political districting models. *Operations Research* URL http://dx.doi.org/10.1287/opre.2021.2141.