

This is the accepted manuscript made available via CHORUS. The article has been published as:

Empirical sampling of connected graph partitions for redistricting

Elle Najt, Daryl DeFord, and Justin Solomon

Phys. Rev. E **104**, 064130 — Published 20 December 2021

DOI: [10.1103/PhysRevE.104.064130](https://doi.org/10.1103/PhysRevE.104.064130)

Empirical Sampling of Connected Graph Partitions for Redistricting

Lorenzo Najt*

Department of Mathematics, University of Wisconsin

Daryl DeFord†

Department of Mathematics and Statistics, Washington State University

Justin Solomon‡

Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology

(Dated: August 16, 2021)

The space of connected graph partitions underlies statistical models used as evidence in court cases and reform efforts that analyze political districting plans. In response to the demands of redistricting applications, researchers have developed sampling methods that traverse this space, building on techniques developed for statistical physics. In this paper, we study connections between redistricting and statistical physics, and in particular with self-avoiding walks. We exploit knowledge of phase transitions and asymptotic behavior in self avoiding walks to analyze two questions of crucial importance for Markov Chain Monte Carlo analysis of districting plans. First, we examine mixing times of a popular Glauber dynamics based Markov chain and show how the self-avoiding walk phase transitions interact with mixing time. We examine factors new to the redistricting context that complicate the picture, notably the population balance requirements, connectivity requirements, and the irregular graphs used. Second, we analyze the robustness of the qualitative properties of typical districting plans with respect to score functions and a certain lattice-like graph, called the state-dual graph, that is used as a discretization of geographic regions in most districting analysis. This helps us better understand the complex relationship between typical properties of districting plans and the score functions designed by political districting analysts. We conclude with directions for research at the interface of statistical physics, Markov chains, and political districting.

I. Introduction

Several recent court cases and reform efforts have created demand for quantitative techniques to analyze *political districting plans*, which are divisions of a state or municipality into regions that each elect a candidate to office. Reflecting political practice, districting plans are considered as connected partitions of geographic subunits, e.g., voting precincts or census blocks. Since, **outside of small scale validation studies** [32], the space of all possible such partitions is too large for complete enumeration, most analyses rely on random sampling to generate an ensemble of plans for a given piece of geography; the properties of this ensemble are used to understand the achievable qualities of potential plans and inevitable trade-offs between the many objectives for plan design.

An increasingly popular methodology for sampling districting plans employs Markov Chain Monte Carlo (MCMC) techniques. A representative subset of examples includes [9, 15, 20, 22, 23, 29, 30, 34, 35, 49]; see [5] for broader discussion of methods in computational redistricting. Although districting plan samplers have been applied convincingly in several court cases [12, 34, 35, 57], these arguments relied on sampling methods manually tuned to the geographies and legal constraints of specific

states, suggesting the need for further scientific research into this potentially useful tool.

To enhance understanding of these MCMC methods, in this paper we draw connections between political districting and statistical physics. We will see that statistical physics helps discover and explain the behavior of sampling methods used in redistricting. In the reverse direction, redistricting motivates new variants of commonly-studied statistical physics models, including new questions that are especially natural in the districting context.

One of the popular Markov chains for districting explores the space of districting plans through a variant of Glauber dynamics. However, a distinguishing feature between redistricting and traditional spin glass models is that districting plans are required to be *contiguous*. That is, the nodes assigned to each label must induce connected subgraphs. This constraint connects the behavior of redistricting methods to known results about self-avoiding walks, but makes it more difficult to use intuitions from spin glass models. In particular, fast mixing properties of cluster-based spin glass models over portions of the parameter space [25, 36, 37, 56] do not appear to hold generically for the samplers we consider. For instance, [54] demonstrates general obstructions to efficient uniform sampling from the space of connected districting plans, along with explicit families of graphs where the Glauber dynamics based Markov chains on districting plans mix slowly. However, the behavior of these Markov chains on regular lattices has so far resisted anal-

* lnajt@math.wisc.edu

† daryl.deford@wsu.edu

‡ jsolomon@mit.edu

ysis; in this paper, we provide some empirical evidence regarding the behavior of these chains on lattices.

Moving beyond contiguity, other factors such as population balance and irregularity of the underlying graph present significant departures from physical models. These complications affect mixing times and qualitative properties of stationary distributions; however, many facts about phase transitions and lattice dependencies carry over to the redistricting context.

We also encounter some questions not explicitly studied in statistical physics, since they have to do with the specific way that samplers are used in redistricting. In the redistricting analysis pipeline, a finite graph representing the state is constructed, called the state-dual graph (see §II), and the aforementioned MCMC samplers are used to produce partitions of that graph that correspond to legal districting plans. We demonstrate, using knowledge of phase transitions gleaned from statistical physics, that the way that the graph is constructed can influence the analysis, even as the parameters of a sampling algorithm are held constant. Such examples raise important questions about the practical interpretation of statistics about random redistricting plans.

A theory for resolving the problem of *which* distributions over partitions are appropriate for redistricting analysis has thus far escaped analysis, or even a precise formulation of the basic questions. We hope that progress on this question can be made by strengthening connections with statistical physics.

II. Background on redistricting

In a typical redistricting process, states are partitioned into small units such as census blocks or voting precincts for data recording, which are then aggregated into larger districts in which political representatives are elected. The units can be represented by the nodes of a graph, where units that share common boundaries are adjacent, as in Figure 1. This graph is called the *state dual graph*. Assuming that voting districts must be contiguous, a districting plan with k districts is modeled by a connected k -partition of the state dual graph. Each state has its own rules for determining which partitions are permissible, although these statutes rarely provide sufficient detail to provide a formal mathematical description of the space of politically valid plans. For more details on the problem setup, see [22, §3].

For our purposes, two of these constraints are particularly relevant. The first is population balance, which requires that the population within each of the k districts be near the the state’s total population divided by k . The second is the notion of geographic compactness, which refers to the ‘niceness’ of district boundaries. Compactness is evaluated using a variety of geographic metrics that are designed to compare the shape of a given district to some idealized geometric property. One of the most common is called the Polsby–Popper score [58], which




Figure 1: a) Kansas with county units [8], along with a connected 2-partition. b) The corresponding state dual graph overlaid. Much more granular subdivisions of a state such as the census units shown in the bottom row are often used instead of counties.

measures the isoperimetric ratio of perimeter to area of each district. Other measures compute ratios of areas to regularized bounding shapes or measure the length of the perimeter directly. In this paper we use a family of scores that measure the compactness of a districting plan as a function of the number of edges in the state dual graph going between different districts. This score is compelling because of its mathematical convenience and relationship to the statistical physics literature, and it has also received some attention in the redistricting literature [19, 22].

It was quickly observed [55] that by a clever choice of districting plan one could engineer aspects of electoral outcomes, a practice that became known as “gerrymandering.” **Despite the focus on using geographic compactness as key indicator of district quality, compactness alone is not enough to prevent gerrymandering, as shown in [2]. In other efforts to counteract gerrymandering,** there have been many proposals to design districting plans algorithmically, a process which often involves grappling with computationally intractable problems [3, 41]. The reality of redistricting, however, is that the power to draw the graph partition is in the hands of a legislature, dedicated committee, or hired expert—

rather than a piece of software. For this reason, rather than using an algorithm to draw plans in the first place, some have suggested to analyze already drawn plans for compliance with civil rights law and other redistricting principles.

One approach to this analysis tries to understand a plan in the context of what alternative plans are possible or typical. For instance, an assertion that it is impossible to provide a certain number of majority-minority districts can sometimes be falsified by a computer search for such plans, as in [13]. However, since the space of districting plans is in general too large to be completely enumerated it is generally computationally intractable to prove that certain kinds of plans are impossible. There are exceptions, however: one can refer to [32] for enumerations of medium size state dual graphs and [13] for situations where impossibility arguments based on non-geographic constraints are possible. Additionally, in many cases the relevant questions have to do with how unusual a plan is along a certain dimension, rather than the impossibility or possibility of a certain kind of plan.

As such, appeals to notions of a *typical* plan and statistical outliers have developed into popular lines of argument. For instance, an argument that a plan was drawn with the intent to discriminate might calculate that the proposed plan has more discriminatory properties than the vast majority of plans from a randomly generated collection of comparable plans, as was done in an amicus brief that was presented to the Supreme Court of the United States in *Rucho v. Common Cause* [11]. The arguments given in that amicus brief relied on a choice of distribution over the space of plans and attempted to tailor that distribution to sample from a *diverse ensemble* of plans that are each compliant with the redistricting principles laid out by the governing body. More generally, a variety of algorithms have been proposed to sample ensembles of graph partitions for similar purposes, from genetic algorithms [17, 44] to random walks [15, 34]; recent expert reports in redistricting cases have used these tools to generate quantitative assessments of proposed plans [12, 34, 35, 57].

In this paper, we focus on algorithms based on Markov chain Monte Carlo, due to a connection with the statistical physics of random spanning trees, self-avoiding walks, and phase transitions in Gibbs energies. Methods making use of these tools have already found wide application in court cases and academic analyses [11, 34, 35, 57]. We will concern ourselves primarily with questions about the mixing time of these Markov chains, and the dependence of the properties of typical plans on phase transitions in the stationary distributions. For instance, we will find that the phase transitions are useful for understanding how a redistricting analysis depends on the discretization of the geographic region into a state dual graph, which is an important point for developing “objective” (in the sense of [33]) standards for restricting analysis. Phase transitions in a Potts model related to these redistricting problems are considered from a statistical physics per-

spective in [24].

A. Two key questions for MCMC samplers in redistricting

Two critical challenges arise when using statistics of random districting plans to analyze political redistricting. First, when the random plans are produced by a Markov chain, we would like to estimate the computation time required to produce a statistically representative sample from the stationary distribution of that chain. Second, to make the analysis more *transparent* and *impartial* (in the sense of [33]), we would like to characterize how the statistics used in the analysis are influenced by features of the sampling algorithm.

This second challenge is especially subtle in the context of the data analysis pipeline for political redistricting, since an outlier classification of a districting plan could be affected by ad-hoc seeming features of the discretization of the state into a state dual graph, as we explore in §V. Additionally, although explicit score functions can be designed for use with the Metropolis-Hastings algorithm, such as those in [10, 30, 34, 48], there is often a significant difference between the apparent preferences of a score function and the typical properties of a sample from the corresponding stationary distribution. Concepts from statistical physics, such as phase transitions around lattice-dependent constants, are key for understanding the relationship between designed score functions and the qualitative properties of typical samples.

In this paper, we will use those physical concepts to build evidence for conjectures about mixing times and to investigate the robustness of redistricting analysis with respect to the state dual graph. We believe that our investigation provides strong evidence that insights from statistical physics are valuable for developing best practices in redistricting analysis.

III. Metropolis-Hastings based samplers

We recall that a general paradigm for Metropolis-Hastings based sampling requires that we build a reversible Markov chain M over the state space Ω of interest and define an energy function $E : \Omega \rightarrow \mathbb{R}$ proportional to the desired stationary distribution. The Metropolis-Hastings algorithm uses the energy function to modify the proposal of M in the following way: if the proposal chain suggests moving to a state of lower energy, the Metropolis-Hastings chain always accepts. On the other hand, if the proposal chain suggests moving to a state of higher energy, the Metropolis-Hastings chain only accepts with probability $\exp(E(x_{\text{current}}) - E(x_{\text{proposal}}))$.

In redistricting, a popular approach has been to define an energy function that incorporates many of the relevant criteria, such as geographic compactness (see §II), population balance, county preservation, and VRA com-

pliance, among others [10, 21, 22, 30, 34, 48]. Analysis depending on such Markov chains have appeared in expert reports in federal court cases, as such as those discussed in [11, 48, 57]. We explain these choices in further detail in §III B. As is typical of many such systems, there is a tug of war between combinatorial explosion and the energy minimization resulting in phase transitions in the tuning of these energy functions; these phase transitions will be a central part of the story we tell in §IV.

In general, one can use the Metropolis–Hastings scheme over a larger state space, say Ω' , than the one, Ω , that one is explicitly interested in, and then rejection sample to obtain samples in Ω . This approach is useful if enlarging the state space makes the Markov chain mix more quickly and the probability that the stationary distribution assigns to Ω is not negligible. This approach can be useful for certain features of redistricting plans, such as population balance. However, in general those interested in analyzing redistricting plans have found it convenient to preserve the connectedness requirement in the proposal function; thus, we will be focusing on proposals that return a connected partition at each step of the chain.

There are two predominant proposal types used in redistricting, one is a Glauber dynamics like proposal based on node flips and the other is a family of proposals based on using spanning trees as a certificate of connectedness. We discuss these in more detail in the next section.

A. Two proposal methods for connected partitions

In this section, we define connected partitions and describe two types of Markov chains on connected partitions.

Definition III.1 (Connected Partitions) *Let G be a graph. A connected partition of G is a partition of the vertices such that each block induces a connected subgraph of G . We let $P_k(G)$ denote the set of connected, ordered k -partitions; equivalently, $P_k(G)$ is the set of k colorings of G where each coloring induces a connected subgraph. We call such a coloring a connected k -coloring.*

The first Markov chain we define is called the flip walk, which is a variant on Glauber dynamics:

Definition III.2 (Flip Walk) *Let $G = (V, E)$ be a graph. We define a Markov chain on the state space $P_k(G)$ by defining a step from a connected coloring P by picking a $v \in V$ uniformly at random, and a $c \in [k]$ uniformly at random, and proposing a move to the coloring P' with $P'(w) = P(w)$ for $w \neq v$, and $P'(v) = c$. If P' is a connected coloring, this move is accepted, otherwise it is rejected.*

Additional discussion of flip variants can be found in [22, §4.2].

Remark III.3 *In §IV we describe the connection between 2-connected partitions of the grid and self-avoiding walks. The flip walk on connected 2-partitions can thus be translated over to self-avoiding walks, and results in a chain that is essentially identical to the BFACF walk [59]. See [54] for a literature review.*

Another MCMC algorithm for redistricting is based on spanning trees:

Definition III.4 (Recombination) *Let $G = (V, E)$ be a graph. We define a Markov chain on the state space $P_k(G)$ by defining a step from a connected coloring P by picking two adjacent colors at random and drawing a random spanning tree T for the subgraph of G induced on the union of the nodes of those two colors. We then randomly select an edge of T to cut, leaving behind two connected components, which are randomly assigned one of the two originally selected colors.*

The recombination algorithm was introduced to the redistricting literature in [22], which contains a discussion of several variants of the method depending on the algorithms used in the randomized parts of the procedure, such as selecting the initial colors, the spanning tree, or the edge to cut. However, the idea of using spanning trees as a certificate for connectedness is much older and pervasive throughout combinatorics and statistics [45, 50, 61, 62]. Recently, this approach has been generalized [10] to operate on the extended state space of spanning forests.

For the most part, our empirical investigation in this paper will focus on the flip walk. **However, we will return to the 2-partition version of recombination in §V. In the 2-partition case, recombination simplifies to sampling a partition by removing an edge from a random spanning tree.**

B. Gibbs energy

Given one of these proposal mechanisms, the Metropolis–Hastings algorithm can be used to construct a Markov chain whose stationary distribution is proportional to a Gibbs energy:

$$\mathbb{P}(P) = \frac{e^{-\beta(\sum_i \alpha_i \sigma_i(P))}}{Z} \quad (\text{III.1})$$

where P is a partition of the given graph, β is the inverse temperature, the σ_i are functions measuring legislatively relevant quantities of P with corresponding weights α_i , and Z is the corresponding partition function:

$$Z = \sum_{P \in \mathcal{G}} e^{-\beta(\sum_i \alpha_i \sigma_i(P))}. \quad (\text{III.2})$$

Common measurements represented by the σ_i include population balance, geographic compactness (see §II),

preservation of county boundaries, and compliance with the Voting Rights Act of 1965 among others. As the relevant legislation is rarely written in sufficient detail to specify a single σ much less an entire distribution, there is a significant amount of modeling freedom in the construction of such an energy.

As an example, [34] uses simulated annealing with the following energy function to sample Congressional districting plans in North Carolina:

$$e^{-\beta(3000 \sigma_{\text{pop}} + 2.5 \sigma_{\text{compact}} + 0.4 \sigma_{\text{county}} + 800 \sigma_{\text{VRA}})}, \quad (\text{III.3})$$

where, in addition to terms measuring the population balance and compactness, the energy function incorporates an entropy-based measure of county splits and a term incentivizing districts with Black Voting Age Population percentages at least as large as those permitted by previous legal challenges [21].

IV. Phase transitions

In this section, we consider a specific score function for Metropolis–Hastings on the flip walk; namely, for a fixed $\lambda > 0$, we define $\beta_\lambda(P) = \lambda^{\text{cut}(P)}$, where $\text{cut}(P)$ is the number of edges connecting the different connected components of P . We explain a connection to theory of self-avoiding walks and how the phase transition behavior for self-avoiding walks carries over to the case of connected partitions and the flip walk. Following the self-avoiding walk literature, we will refer to the term λ in $\beta_\lambda(P)$ as the *fugacity*. Later on, we will use our knowledge about phase transitions and other self-avoiding walk theory to guide an empirical investigation of bottlenecks in the flip walk chain. To begin with, we introduce a toy model of a state dual graph that will serve as a running example throughout.

Definition IV.1 (Gridlandia) *Gridlandia, G_n , is an $(n+1) \times (n+1)$ grid graph where each node has population 1. In other words, it is an arrangement of $n \times n$ squares as in Figure 3. We think of population as a function of N -valued weights on the vertices of a state dual graph.*

Definition IV.2 (Population Balance) *We consider α -balanced connected partitions of G_n , $P_{2,\alpha}(G_n)$, which are the connected 2-partitions of G_n where the total population of each district is within 100α percent of exactly half the total population.*

Note that increasing α amounts to loosening the population constraint; if $\alpha \leq \alpha'$, then $P_{2,\alpha} \subseteq P_{2,\alpha'}$.

Remark IV.3 *The definition of population balance used in Definition IV.1 is just one of many reasonable definitions. The choice of function measuring the population balance does not appear to have a significant effect on the empirical behavior. It is possible that another definition would be analytically preferable.*




Figure 2: The relationship between connected partitions and self-avoiding walks.

We will refer to a non-self intersecting path between any two boundary points of G_n as a *chordal self-avoiding walk*. Using bond-cycle duality for plane graphs, we can interpret connected partitions of G_n as simple cycles on the dual graph. For connected and simply connected partitions, these are effectively the same thing as chordal self-avoiding walks in the dual graph minus the node corresponding to the unbounded face, see Figure 2; the only minor discrepancy being that each of the four corner faces has two edges connecting it to the supernode. Given a connected and simply connected 2-partition P , we refer to the boundary self-avoiding walk as the *partition boundary* of P .

An important family of distributions on chordal self-avoiding walks in a plane graph is the following: For any $\lambda > 0$, define the probability measure $\beta_\lambda(\omega) \propto \lambda^{|\omega|}$, where $|\omega|$ is the number of edges used in the self-avoiding walk. On the grid graph, the limiting behavior of these distributions experience a phase transition at $1/\mu_S$ [27], where $\mu_S = 2.635\dots$ is the *connective constant* of the grid lattice, which measures the rate of growth of the number of length n self-avoiding walks starting at the origin in $\mathbb{Z}^2$. This phase transition leads to three regimes of behavior: $\lambda < 1/\mu_S$ (subcritical), $\lambda = 1/\mu_S$ (critical) and $\lambda > 1/\mu_S$ (supercritical). Between each regime, the behavior of a typical self-avoiding walk is very different [27], although within regimes the behavior has a qualitative uniformity. See §IV D for a discussion the properties of typical partitions sampled in these three regimes.

Connective constants for triangular and hexagonal lattices, denoted μ_T and μ_H , are defined similarly to μ_S , and stationary distributions for self-avoiding walks exhibit identical phase transitions for λ at $1/\mu_L$, where L is the lattice. It is known that $\mu_T = 4.15\dots$ and $\mu_H = \sqrt{2 + \sqrt{2}} = 1.85\dots$ [28, 39]; in particular, $\mu_T < \mu_S < \mu_H$, and therefore the phase transitions occur at different parameter settings on different lattices. To discuss this precisely, we consider two additional graphs:

Definition IV.4 (Hexlandia) *Hexlandia, H_n , is a graph of $n \times n$ hexagonal tiles arranged as in Figure 3. Each node has population 1.*

Definition IV.5 (Trilandia) *Trilandia, T_n , is a graph of $n \times n$ triangular tiles arranged as in Figure 3. Each node has population 1.*




Figure 3: From top to bottom, the 4×4 versions of Squarelandia, Trilandia and Hexlandia in black, displayed with their supernode-deleted duals in blue.

Connected 2-partitions on Hexlandia correspond to self-avoiding walks on a dual triangular lattice graph, and connected 2-partitions on Trilandia correspond to self-avoiding walks on a dual hexagonal lattice graph. In particular, $1/\mu_H$ is the phase transition for β_λ in Trilandia, and $1/\mu_T$ is the phase transition for β_λ in Hexlandia. In general, the behavior of Markov chains on connected partitions can depend on which side of the phase transition one is on, and therefore the behavior of the same algorithm can vary widely based on whether the state is discretized into triangles, squares or hexagons. We return to the role that topology the state dual graph plays in §V, but for the remainder of this section we focus on gridlandia.

Summary of experiments. The remainder of this section consists of experiments investigating the mixing time and qualitative properties of the flip walk Markov

chain in each of three fugacity regimes. Specifically:

- In §IV D, we explore the mixing time of the Markov chain by measuring statistics about the slope of the partition boundary.
- In §IV E, we follow up on the apparent success of the critical regime flip walk chain by testing it against known statistics of Schramm-Loewner evolution.
- In §IV F, we examine the distribution of population balance in the three fugacity regimes.

Each experiment computes statistics from an MCMC run through connected partitions on gridlandia. In §IV A, §IV B, and §IV C, we summarize precisely how we measure the length of the Markov chain experiments, the summary statistics we use and the precise specifications of the Markov chains, respectively. In §IV G, we conclude with a conjecture summarizing the experimental data and related questions.

Implementation. Code for reproducing the experiments in this paper is available at: https://github.com/LorenzoNajt/Code_For_Empirical-Sampling-of-Connected-Graph-Partitions-for-Redistricting

A. Experimental details: Measuring run time

We use the Metropolis-Hastings algorithm on a reversible Markov chain M that can be represented as an undirected, edge-weighted graph G_M . In our case, G_M has many self loops, and instead of simulating those self loops during implementation we can simulate the random number of self-loops taken before moving to a different state. To add to the confusion, we do not simulate every self loop this way, but only a large fraction of them. In this section we include definitions for *compressed steps*, *steps* and *accepted proposals* that are intended to clear up the potential ambiguity that can occur when reporting on Markov chains that simulate self-loops. We also specify the computing resources used.

Counting the number of steps. Recall that the flip walk on $P_2(G)$ picks a node $v \in V(G)$ uniformly at random and a color $c \in \{0, 1\}$ uniformly at random, and proposes changing the color of v to c . The flip walk accepts this proposal if the coloring remains connected and if the number of colors remains equal to 2. At a partition P , the only nodes that can change color are those adjacent to an edge of $\text{cut}(P)$; call that set $\text{bdnodes}(P)$. This means that instead of uniformly sampling from vertices v until we obtain a sample from $\text{bdnodes}(P)$, we can sample the number of times we expect to draw a sample outside of $\text{bdnodes}(P)$ and then uniformly sample from $\text{bdnodes}(p)$. This is achieved by the following observation:

Observation IV.6 *Let $B \subseteq N$ be two finite sets. Let $X_1, X_2, \dots$ be iid samples from N . Let T be the first*

time that $X_i \in B$. Then T is geometrically distributed according with success probability $|B|/|N|$, and X_T is a uniform sample from B .

Thus, instead of picking a uniform $v \in V$, one can pick v uniformly from $\text{bdnodes}(P)$ and use a geometric random variable with success probability $|\text{bdnodes}(P)|/|V|$ to count the number of self loops taken because of drawing $v \notin \text{bdnodes}(P)$.

We let *compressed steps* refer to the number of times we draw from $\text{bdnodes}(P)$, and *steps* refer to the total sum of the aforementioned geometric random variables. Finally, if we use an additional score function/Metropolis–Hastings step on top of the flip walk proposal function, we let *accepted proposals* refer to the number of proposals accepted by the Metropolis–Hastings algorithm, and *proposals* refer to the number of proposals made.

Remark IV.7 Only certain nodes in $\text{bdnodes}(P)$ result in legal flips, so that in fact even more of the self-loops of the underlying directed graph could be ignored and wrapped into a geometric random variable. However, we did not do so in the implementation.

Computation time. Computational time refers to the number of seconds the 0.2.12 release of the open-source GerryChain software ran for. All experiments about mixing times and phase transitions on gridlandia ran on a Xeon(R) CPU E5-2683 v4 at 2.10Ghz. The annealing experiments were run on an Ubuntu 16.04 machine with 64 GB memory and an Intel Xeon Gold 6136 CPU (3.00GHz) processor.

B. Experimental details: Statistics measured

This subsection will summarize the statistics we use to empirically analyze the behavior of these Markov chains.

Boundary slope. A natural qualitative measurement of the shape of a simply connected, connected 2-partition of gridlandia is whether it is like a vertical, diagonal or horizontal partition. The line connecting the endpoints of the boundary of a simply connected partition gives us a simple way to turn that intuition into a quantitative measurement whose long-term behavior can be visualized and measured. This measurement will give us insight into bottlenecks, or metastable regions, in the flip walk.

Definition IV.8 (Boundary line segment) Let P be a connected partition of gridlandia G_n . Then P induces a connected 2-partition Q of the boundary of G_n , which is a cycle with $4n-4$ nodes. If $\text{cut}(Q)$ is non-empty, then it has 2 edges, e_1 and e_2 , and the boundary line segment is the line segment connecting the midpoints of e_1 and e_2 . If Q is the trivial partition, and $\text{cut}(Q) = \emptyset$, then we define the boundary line segment to be $\perp$.




Figure 4: (Left) The boundary line segment. (Right) An example boundary slope plot.

We summarize a time series of boundary line segments via their slopes, which we plot in polar coordinates with the radial axis denoting time. That is,

Definition IV.9 (Boundary Slope Plot) Let

$P_1, \dots, P_n$ be a sequence in $P_2(G_n)$, and let $l_1, \dots, l_n$ be the corresponding boundary slopes. For each $l_i \neq \perp$, there are two associated unit vector directions $\pm v_i$. We choose signs so that v_i is closer to v_{i-1} than $-v_{i-1}$ and draw an intervals connecting $(i-1)v_{i-1}, iv_i$; in the limit, this amounts to picking a continuous path of unit vectors that are aligned with the boundary slope.

Our implementation produces the sequence $(P_1, \dots, P_n)$ by taking every 10,000 steps (which include self loops, as discussed in §IV A); this choice generally captures every compressed step. For aesthetics, we interpolate between points (θ, r) and (θ', r') in polar coordinates with the curved line $(\lambda\theta + (1-\lambda)\theta', \lambda r + (1-\lambda)r')$. The radial direction measures time, and the stretching that comes with being further from the origin is a distortion that has nothing to do with the underlying path of partitions; only the angle represents a feature of the corresponding partition. Finally, we draw circles at regular intervals to give a sense of time segments. An example boundary line segment and boundary slope plot is given in Figure 4.

Flips and edge cuts. Here, we summarize two closely related statistics that can be used to see metastable regions in the Markov chain. Unlike measuring the boundary slope, these store an aggregate measure of variation, rather than a detailed history. This can make it harder to see metastability after sufficiently long runs, but these statistics are especially useful because they generalize easily to k -partitions. Additionally, patterns that also have to do with the internal behavior of partition boundaries, such as those we highlight in §V, are easier to see from these plots since boundary slope plots only record the relative behavior of the endpoints of the partition boundary.

Definition IV.10 (Flips) We define a function $\text{flips} : V \rightarrow \mathbb{N}$ by setting $\text{flips}(v)$ to be the number of times the district containing v changed: $\text{flips}(v) = |\{i \in \{1, \dots, n -$

$1\} : f_i(v) \neq f_{i+1}(v)\} |$. Usually we will plot $\text{lflips} := \log(\text{flips} + 1)$ instead.

We plot lflips using a heat map where red means a node is flipped more often, and blue means it is flipped less often.

Definition IV.11 (Edge Cuts) We define a function $c : E \rightarrow \mathbb{N}$ by setting $c(e)$ to be the number of times that e was a cut edge of a partition P_i : $c(e) := |\{i \in \{1, \dots, n\} : e \in \text{cut}(P_i)\}|$.

We visualize cut edges as heat maps on the state dual graph; lighter green/yellow edges are those that are cut more than darker green/purple edges. Edge cuts are a useful statistic for examining proposals that do not flip individual vertices, such as the tree based proposals. See Figure 20 for an example.

C. Experimental details: Proposal & score function

Here we provide further details of a flip walk that is modified to reject proposals leading to heavily population unbalanced states, and that is further modified by using the Metropolis–Hastings algorithm with the score function λ^{cut} .

Hard population constraints, rejection sampling population constraints. We define the population balance score PopBal of a partition (A, B) of $n \times n$ gridlandia as $\text{PopBal}(A, B) = \max(|A|/\frac{n}{2}, |B|/\frac{n}{2})$. We set a threshold $\alpha > 0$, and define the flip walk with population constraint α , Flip_α to be the flip walk modified to reject all proposals that would produce a partition with $\text{PopBal}(A, B) > 1 + \alpha$.

Metropolis–Hastings for λ^{cut} compactness constraint. Fix some $\lambda > 0$. We use the Metropolis–Hastings algorithm on top of Flip_α with score function $f((A, B)) = \lambda^{\text{cut}(A, B)}$, where $\text{cut}(A, B) = \{e \in E(G) : e = \{v, w\}, (v \in A \wedge w \in B) \vee (v \in B \wedge w \in A)\}$. We denote the resulting Markov chain by $\text{Flip}_{\alpha, \lambda}$.

D. Experiment: Metastable Regions in the flip walk chain

In this section, we will see that the boundary slope plots give insights into bottlenecks in $\text{Flip}_{\alpha, \lambda}$. We break our analysis down into the three phase transition regimes of the λ^{cut} score function.

The subcritical case. Subcritical self-avoiding walks in the grid are known to converge to geodesics [27]. In the context of connected partitions, this means that the score function β_λ strongly prefers minimizing the boundary length of a partition. Unlike the self-avoiding walk case, however, in redistricting we have to contend with population balance. It turns out that the behavior of the




Figure 5: (left) Boundary slope plot and (right) the end state, from a subcritical ($\lambda = \frac{1}{2\mu} = .189\dots$) flip walk constrained to population balance 10 percent (bottom) and 50 percent (top). The top plot was run for 21,000,000 compressed steps, which amounted to 703,328,732,936 proposals, over roughly one month of computing time. The bottom plot was run for 23,000,000 compressed steps, 712,097,053,677 proposals, over the same time period.

subcritical cases depends strongly on whether or not the two districts are required to have balanced population.

If the population balance is unrestricted, the chain appears to mix thoroughly in its state space, which consists of partitions where one block occupies most of the region and the other block is a small bubble. As the population constraint is tightened, the ability of the Markov chain to move around the state space becomes more tightly constrained. We illustrate the evolution of 10% and 50% population balances in Figure 5: with strong population constraints, the chain gets stuck around one of the partitions satisfying the population constraint with the shortest possible boundary. We note that the results of these experiments change when λ is closer to the critical fugacity, and will discuss that case in more detail after reviewing the critical and supercritical regimes.

In §IV F, we revisit the influence of population balance on subcritical sampling through rejection sampling.




Figure 6: $\log(\text{flips} + 1)$, after 254 hours (equivalently, 64,054,464,437,342 proposals) on a 40×40 grid graph. pop bounds 5 percent. This is stuck in a bottleneck set; it has not reached the equally probable region around the vertical partition. Loosening either the population constraint or bringing the fugacity closer to the critical value can help mixing.




Figure 7: A typical supercritical sample.

The supercritical case. The supercritical case is generally unsuitable for redistricting, since the boundaries are almost-surely space filling [27]; see Figure 7 for a typical example. Experiments show that the flip walk Markov chain also has bottlenecks that are revealed by the boundary slope statistics, see Figure 8. On the other




Figure 8: Boundary slope plots (left) for a super critical ($\lambda = 1$) run with population constraint 100%. The apparent rapid rotation in the beginning of the run was due to the boundary end points ending up near the two corners of one edge. This was run for 11,000,000 compressed steps, or 20,302,926,563 proposals, over the course of one month of computing time.




Figure 9: Boundary Slope Plots for $\lambda = 1/\mu$ (critical regime), with population constraints 10%, along with the final state on a 40×40 grid graph. Run for 25,000,000 compressed steps, amounting to 354,337,192,915 proposals, over the course of approximately one month of computing time. The metastable region that was apparent in the subcritical case is not an obstruction here.

hand, the supercritical case appears to mix in space, in the sense (see [51]) that the configuration of the partition boundary in the interior of gridlandia seems to be independent of the partition restricted to the boundary of gridlandia.

The critical fugacity case. Unlike the previous two cases, this setting of the score function does not cause the chain to get trapped in a metastable region corresponding to the slope of the boundary. This fact can be seen in the boundary slope plot Figure 9. Additionally, the chain appears to continue to mix rapidly even with strong population constraints. Still, whether or not the flip walk is rapidly mixing at this setting of λ remains unknown, since there could be more subtle bottlenecks.

To explore that possibility, more intricate tests of mixing are needed; in the next section we show that samples from the flip walk at $\lambda = 1/\mu$ produce empirical statistics that agree with predictions from Schramm–Loewner evolution. This is further, although still inconclusive, evidence of rapid mixing of the flip walk on gridlandia at critical fugacity.

When does the phase transition occur? The phase transition results in [27] are asymptotic, and for any given grid graph the distinctions between the three regimes are fuzzy. For instance, in Figure 10 we take λ at three values closer to $1/\mu$ than $1/2\mu$, and see that doing so produces behavior that is more the critical regime than the subcritical regime. In particular, the bottlenecks we saw in the subcritical regime become gradually less pronounced as λ approaches $1/\mu$. Exactly how the asymptotic phase transitions kick in seems like a challenging question, and one with implications for sampling maps for redistricting.




Figure 10: (Left) Setting $\lambda = 1/\mu - .3 = .35\dots$, population bound at 50% and using a 40×40 grid graph results in a run that looks more like the critical case than the subcritical case. (Middle) The same but $\lambda = .32$, (Right) same but $\lambda = .29$. Each of these was run for about a week of computing time, which amounted to the order of 250,000,000,000 proposals each.

E. Experiment: SLE based test of mixing at criticality

Although the boundary slope plots suggest bottlenecks in both the subcritical and supercritical regimes, they do not suggest the same bottlenecks in the critical regime. In this section, we provide another test for mixing at the critical regime, and show that the Markov chain passes this test.

Background on Schramm–Loewner Evolution.

The theory of Schramm–Loewner evolution (SLE) can be used to make a prediction about the expectation of a statistic on the Markov chain of connected 2-partitions in the critical regime. We thus obtain a test for mixing by empirically calculating this expectation and comparing it to the theoretically predicted value. Our experiments in this section are indebted to the investigation in [40], which used the same statistic to verify conjectures about scaling limits of self-avoiding walks.

Schramm–Loewner evolution concerns a family of probability distributions, SLE_κ with $\kappa \in [0, \infty)$, on compact subsets of the upper half plane $\mathbb{H}$. For a precise definition of Schramm–Loewner evolution, see [7]. Many discrete models from statistical physics converge to SLE_κ for appropriate values of κ . In particular, $SLE_{8/3}$ has close, albeit conjectural, relationships with self-avoiding walks.

$SLE_{8/3}$ is a probability distribution over simple paths from 0 to ∞ in $\mathbb{H}$. This distribution is invariant under scaling, so if γ is a sample from $SLE_{8/3}$, so is $\alpha\gamma$, for any $\alpha > 0$. In particular, if $a, b \in \partial\mathbb{D}$, the scale invariance of $SLE_{8/3}$ means that the distribution over simple paths in $\mathbb{D}$ from a to b obtained by conformally mapping $\mathbb{D}$ to $\mathbb{H}$ and sending a to 0 and b to ∞ is well-defined; this distribution over paths in $\mathbb{D}$ is called chordal SLE from a to b .

Let D_n denote the nearest neighbor graph on the intersection of $\mathbb{D} = \{z \in \mathbb{R}^2 : |z| \leq 1\}$ with $\frac{1}{n}\mathbb{Z} \times \frac{1}{n}\mathbb{Z}$. We recall the conjecture discussed in [27, 43] that a $\beta_{1/\mu}$ -sample from a self-avoiding walk from -1 to 1 in D_n ,

considered as the trace of a path in the plane, will converge to $SLE_{8/3}$ in $\mathbb{D}$ from -1 to 1 as $n \rightarrow \infty$; a similar conjecture holds for other pairs of points in the boundary. Thus, if a sample from the flip walk closely approximates a $\beta_{1/\mu}$ -sample from a self-avoiding walk, then a sample from the flip walk will closely approximate a sample from $SLE_{8/3}$.

Explicit formulas for the probabilities that the upper half-plane $SLE_{8/3}$ hits certain sets are known [40, 42], and we use those to design tests of mixing. In particular, we use the same statistics as those used by the investigation in [40], which relied on the calculation in [42] of the probability that an $SLE_{8/3}$ path in the upper half plane $\mathbb{H}$ intersects $B_r(x)$, the ball of radius r centered around $x \in \mathbb{R} \subseteq \mathbb{H}$. We will denote this probability by $\phi(r, x)$. What follows is the description of an experiment using this to test mixing of the flip walk at fugacity $1/\mu$:

1. We are given a sequence of chordal self-avoiding walks $\gamma_1, \dots, \gamma_N$ as the boundary paths of a sequence of connected 2-partitions of D_n obtained from a length M run of the flip walk Markov chain at $\lambda = 1/\mu$.
2. For each i , map γ_i to $\mathbb{H}$ via a conformal map from $\mathbb{D}$ to $\mathbb{H}$ sending one endpoint of γ to 0 and the other to ∞ . Call the resulting path $\tilde{\gamma}_i$.
3. Set $\hat{\phi}(r, x) = \frac{1}{N} \sum_{i=1}^N 1_{B_r(x) \cap \tilde{\gamma}_i \neq \emptyset}$.

If the chain is rapidly mixing, and ignoring the significant issues of discretization, $\hat{\phi}(x, r)$ should approximate with $\phi(x, r)$.

SLE experiments. We present two experiments. First, we fix the endpoints at $-1 + 0i$ and $1 + 0i$, meaning that we start with a connected partition that has a boundary path from -1 to 1 , and reject all proposals of the flip walk that would change the endpoints; results are displayed in Figure 11. Since we do not appear to need many steps to see agreement between the theoretical and empirical statistics, we take this as evidence that the flip walk mixes rapidly at the critical fugacity.

For the second experiment, we drop the fixed-endpoint requirement. In this case, we do not see strong agreement between $\phi(x, r)$ and $\hat{\phi}(x, r)$, as can be seen in Figure 12. This may have as much to do with the discretization of the disk into D_n as the mixing time of the random walk.

F. Experiment: Population balance.

The regions of apparent rapid and slow mixing, which depend on population balance and temperature, suggest that parallel tempering could be a useful way to sample from these stationary distributions. Parallel tempering between phase transitions probably would not work well, given that the corresponding distributions have negligible overlap in their support. On the other hand, if one fixes λ to be subcritical or critical and varies the population threshold α , then there appears to be considerable overlap between the distributions; this is demonstrated




Figure 11: Taking $n = 20$ (Top) and $n = 30$ (Bottom), with γ_i obtained after 100,000 steps each with $N = 500$. Comparing $\hat{\phi}(r, x)$ with $\phi(r, x)$ for $x = \pm 1$ and $r \in \{i/101 : 0 \leq i \leq 99\}$. The solid red line is $\phi(r, 1) = \phi(r, -1)$, and the other two lines are $\hat{\phi}(r, 1)$ and $\hat{\phi}(r, -1)$. **The x -axis is r , and the y -axis is the value of the function.** Apparently 100,000 is sufficient for mixing when $n = 20$, at least for this statistic, but not when $n = 30$.

in Figure 13. Such overlap could be exploited by MCMC approaches using parallel tempering, as in [30].

G. Discussion

Mixing times. Summarizing the data in §IV D and §IV E, where we saw evidence about the interaction of the mixing time of flip walk with population constraints and fugacity regimes, we make the following conjectures. We recall that in the theoretical computer science literature, “rapid” mixing means polynomial mixing-time in the parameter, and “torpid” generally means super-




Figure 12: With variable endpoints on a 30×30 (top) and 40×40 (bottom), with $N = 501$ independent runs starting from the horizontal partition for 1,000,000 steps each. Notation as in fig. 11.

polynomial or exponential time in the parameter. Here the parameter is n , the size of the grid graph.

Conjecture IV.12 Let $\mathcal{G} = (G_n)_{n \geq 0}$ denote the family of $n \times n$ grid graphs. Then, there are functions $\epsilon_i : \mathbb{N} \rightarrow (0, 1)$, $\epsilon_1 < \epsilon_2$ and $\lim_{n \rightarrow \infty} \epsilon_2(n) = 0$ so that the following holds. For $\lambda > 1/\mu + \epsilon_2(n)$ the flip walk on $\mathcal{G}$ is torpidly mixing, regardless of population constraints. For $\lambda < 1/\mu - \epsilon_2(n)$, if there is any constant percentage population balance constraint then the flip walk on $\mathcal{G}$ is torpidly mixing, but with no population constraints it is rapidly mixing. For $|\lambda - 1/\mu| < \epsilon_1(n)$, the flip walk is rapidly mixing, regardless of of population constraints.

Although in [54] we calculated worst case lower bounds on the mixing time of the flip walk, the graphs that provided the examples were unlike the graphs used in redistricting analysis. On the other hand, the grid graph seems to capture some of the difficulties involved in sampling connected partitions of the state dual graphs used in political districting. Thus resolving the above conjecture would bring us closer to understanding the reliability of flip-based Markov chains in redistricting. As previously mentioned, the flip walk appeared in statistical physics as the BFACF chain [59], but known lower




Figure 13: The population balance histogram after running with balance constraints for steps starting from a vertical partition of a 40×40 grid graph. Here, the x -axis represents the balancedness of a block, measured in terms of $|A|/h$ where $|A|$ is the size of a block, and h is exactly half the total population. The y -axis counts the number of times a partition containing a block with that balance was encountered during the run. After normalization, this provides an approximation of the marginal distribution of population balance of the stationary distribution. Since we start with a perfectly balanced partition, we do not include the first 1,000,000 partitions in the histogram, and these figures were obtained after a month of computation. Left is subcritical ($\lambda = 1/2\mu$), middle is critical ($\lambda = 1/\mu$), right is supercritical ($\lambda = 1$). The balance constraints for the subcritical and supercritical runs are 50 percent, which explains the sharp cutoffs, and the balance constraint for the supercritical run is 100 percent. For the subcritical run, most of the mass is on partitions which are nearly as unbalanced as allowed by the constraints; experiments evaluating softer thresholds on population balance or using simulated tempering could provide additional insight in future work.

bounds on mixing time [60] use techniques that we were not able to modify for this conjecture. In the remainder of this section we examine the applicability of annealing through different fugacities, the impact of fixing endpoints, present a schematic capturing some of the above conjecture and of the effects of annealing, and close by mentioning k -partitions briefly.

Annealing. A common technique for escaping bottlenecks in MCMC is to use simulated annealing, where the inverse temperature parameter is varied to promote efficient movement throughout the state space. Although this technique has been used in court cases [34, 35], the rigidity introduced by the contiguity and population balance terms for the redistricting problem mean that there are additional tradeoffs that must be evaluated in practice.

A key issue in many of the examples in Figure 14 is that while there is a significant amount of movement within the center of the grid, there is little change in the partition of the boundary component, which indicates that this simulated annealing strategy does not help us move out of the metastable region corresponding to the vertical partition. More explicitly, whether the induced partition of the boundary of the grid (a cycle graph) is




Figure 14: Samples from distributions with energy function $e^{-(\gamma_{\text{pop dev}} + \delta \text{boundary length})}$ with $\delta = 1$ as γ varies from .11 to .81 (top two rows, read left to right) and $\gamma = 1$ as δ varies from .11 to .81 (bottom two rows, read left to right). The images in each experiment are snapshots of a single run of the Markov chain.

vertical or horizontal decomposes the configuration space into two equal sized pieces, and Figure 14 shows that annealing does not essentially change the partition of that cycle graph—in fact, throughout the annealing process the boundary slope remains roughly vertical. A similar observation was made on real world data in [22]. Additionally, one source of computational inefficiency of these chains operating in the super-critical regime is that the space filling boundaries make it more likely that any individual proposal will be rejected since changing the assignment of a node is likely to disconnect the partition.

Fixing endpoints. As with the population balance constraint, there appears to be an appreciable difference between the mixing of connected partitions with fixed partition boundary endpoints, and where the partition boundary endpoints are allowed to move. Our empirical evidence does not exclude the possibility that the bottlenecks revealed by the boundary slope plots are the only bottlenecks in the super-critical and subcritical case; that would mean that sampling from the stationary distribution could be possible by picking endpoints according to a partition function, and then running the flip walk on self-avoiding walks with fixed endpoints. Of course, such a strategy would become more difficult for k -partitions, where there is a combinatorial explosion in the different ways the endpoints of the boundary segments can be arranged.

Space of graph configurations. We give a schematic of the space of configurations in Figure 15. As in our discussion of annealing, walks that attempt to travel vertically in our schematic by increasing or decreasing the fugacity do not appear to generate diverse collections of plans, since the boundary slopes remain essentially fixed. Traveling horizontally is also constrained outside of the critical region, as we saw earlier. In general, it appears

that moving to other qualitatively distinct “compact” partitions requires traversing those with poorer compactness scores, such as those that are typical for the critical fugacity. On the other hand, allowing the compactness score to become too loose incurs a penalty in the form of the rigidity that comes from interlocking tendrils of supercritical fugacity states. Without population balance constraint, conditions are known for when these configurations spaces are connected [1].




Figure 15: A schematic showing the local geometry of connected graph partitions. The y-axis represents ‘compactness’ (boundary length) as a measure of niceness of the partition.

***k*-partitions.** Most political districting plans involve more than 2 districts, and it is natural to ask which of our observations extend to the *k*-partition setting. An issue that arises when dealing with *k*-partitions instead of 2-partitions is that compactness scores defined for individual districts have to be aggregated into compactness scores for an entire plan [22, 34, 57]. There are various ways to perform this aggregation, and to the best of our knowledge phase transitions in choices of the aggregation have not been studied.

V. The role of the state dual graph

The experiments in §IV give an indication of the important role that the underlying graph plays in the way any given algorithm samples from the set of districting plans, thought of as partitions of the underlying geographic region. In this section, we highlight some subtle features of this dependence by showing that the behavior of a fixed sampling algorithm can be controlled by subdividing faces of the underlying state dual graph. We connect this to phase transitions in the self-avoiding walk model and discuss some implications for redistricting analysis.




Figure 16: The graph G' .




Figure 17: The shaded squares is the chosen set F , which are subdivided to form the graph on the right.

A. Concentration of boundary location

In this section we work with variations on the gridlandia graph, wherein we allow certain faces to be subdivided into triangles.

Definition V.1 (Subdivided Gridlandia) Given gridlandia G_n , and a set $F \subseteq \text{Faces}(G_n)$, where Faces denotes the set of bounded faces of G_n , define G_n^F by adding edges diagonally from the lower left to upper right corner of each face in F . See Figure 17.

Remark V.2 G_n is the dual graph of a map, that is, a partition of the unit square into regions wherein each region becomes a node, and two nodes are adjacent if the regions share a 1 dimensional boundary. In particular, G_n is the dual graph of the map of the unit square obtained by subdividing it into regular squares, and G_n^F is the dual graph obtained by modifying the partition of the square as in Figure 18.




Figure 18: Minor modifications of the partition of a state into geographic units can change the lattice structure of the state dual graph.




Figure 19: Even if we start with the vertical partition, we rotate into the horizontal partition and then remain there. 1000000 steps. Logflips on the left and edge cuts heat map on the right.

Our aim here is to describe how carefully choosing F can cause the boundary of “random partition” to concentrate around a horizontal line dividing gridlandia in half. We recall from §IV that the phase transition for connected partitions on the triangular lattice occurs at a lower fugacity than that of the square lattice. Thus, if parameters are set so the flip walk chain is subcritical on the triangular part, but critical on the square part, the chain will push partition boundaries off of the triangular portion. If the triangular portion surrounds the boundary of a particular partition P , it is therefore intuitive that that a stationary sample will be likely to thread its boundary along that of P .

This intuition is born out in the experiments we performed. Suppose we start with gridlandia G_n , but we triangulate the faces away from the horizontal partition boundary, as in Figure 16. We will call the resulting graph G'_n . We run the flip walk at fugacity λ where $1/\mu_T < \lambda < 1/\mu_S$, and find a preference for partitions near to the horizontal partition. The result of this experiment can be seen in Figure 20, where we can see that the boundary is indeed more likely to run vertically. This is a different feature from the subcritical metastable regions discussed in §IV D, since a vertical partition of G'_n will quickly rotate into being horizontal, see Figure 19.

The typical behavior of partitions of G'_n behaves similarly for the various popular sampling algorithms that partition by using components of random spanning forests [10, 22] (see §III A). One such algorithm samples a uniform spanning tree and deletes a random edge such that the two components of the resulting forest are roughly balanced in size. When we sample connected




Figure 20: (Left) Log flips plot of running the flip walk with $\lambda = 1/\mu_{\text{square}}$ on G' for 10,000,000 steps. (Right) Edge cut heat map of splitting 1000 uniform spanning trees. In both cases a preference for the horizontal divide is apparent.

partitions in this way and display the edge cut frequencies on the right side of Figure 20, there is a clear bias toward the horizontal partition. These biases can have significant effects on the number of seats won in a “typical map,” representing a challenge for ensemble based methods for curtailing extreme gerrymanders. Nearly identical effects are seen for producing a random spanning tree by drawing random edge weights and calculating the minimum spanning tree. See [54] for further discussion, including an explicit example of the influence on the number of seats won.

Remark V.3 We explain a potential intuition for these examples. First, since we are partitioning the vertices of these graphs, we can represent the boundary by a simple cycle in the plane dual. A dual path passing through the triangulated part is in general twice as long as a corresponding boundary was before, since for every square face it passed through it now passes through two triangular faces. Even though exponentially many potential boundary paths are added after this triangulation, the compactness constraint is such that each is down-weighted exponentially. The balancing act between having many paths and each path having negligible probability determines which orientation of boundary paths are more likely. While the phase transitions give a conceptual, although speculative, explanation of how this tug-of-war works out for the β_λ distribution on the flip walk, it remains a challenge to explain the why G' influences the uniform spanning tree based distribution in the way that it does.

Remark V.4 It is plausible that one could engineer a compactness score that depends on the metric length of the boundary, to avoid dependency on the discretization. However, defining such a compactness score raises several other important issues such as the fractal nature of coastlines, dependencies on map projections, and others [4, 26]. Investigating the relationship between state dual graph topology and compactness represents an important




Figure 21: Heat map of boundary edge frequency for the states of New Hampshire and Kansas as computed over four different levels of resolution derived from census data. For each state, we draw 10,000 districting plans, with a 2% population deviation tolerance and compute how often each edge in the dual graph is cut in a partition.

challenge for redistricting.

B. Implications for Redistricting

We discuss the implications of the above experiments for redistricting analysis.

Interpretation of the distribution. Recall that redistricting is modeled by producing a state dual graph out of small units. Common choices for these units include census geographies like blocks and tracts as well as voting precincts, whose boundaries are frequently determined locally. In Figure 21 we show the edge cut frequencies of spanning tree partitions for various choices of units. While the differences are not as extreme as the synthetic examples in Figure 20, they highlight the need for careful consideration of the underlying graph as a part of experimental design. There are many points in the analysis pipelines where choices about the state dual graph are made, such as handling state dual graph adjacencies over mountain ranges or lakes and prorating data between different geographic units. Thus, to rigorously compare plans to a typical plan it is important to understand to what extent the properties of typical plans depends on choices made earlier in the analysis, and to make such choices in a transparent and impartial way [33].

Tuning score functions. Investigating a local notion of the connective constant could lead to insights about a score function for which the flip walk mixes rapidly on a broad class of graphs. Following the definition of the connective constant in [46], given a graph G and $x \in V(G)$, we define $\mu_{G,n}(x) := \sqrt[n]{S_n(G, x)}$, where $S_n(G, x)$ is the number of simple paths in G starting at x of length n . Then, for any self-avoiding walk γ in G , defining $f_n(\gamma) := \prod_{x \in \gamma} \frac{1}{\mu_{G,n}(x)}$ and taking n sufficiently large would give one possible analog to the $\beta_{1/\mu}(\gamma) = (1/\mu)^{|\gamma|}$ distribution, for which we gave evidence of rapid mixing in §IV D and §IV E. A challenge for computing with this is that estimating $S_n(G, x)$ is not straightforward for

large n and is already challenging for highly structured graphs like $G = \mathbb{Z}^2$ [18].

Metamandering. One could hypothetically exploit phase transitions and lattice-dependent properties to cause a sampling algorithm to have its sampling distribution concentrate around some preferred options, by an appropriate modification or subdivision of a state dual graph. For instance, if an analyst prefers vertical partitions of a square state, for instance because of the arrangement of voters, and had some freedom over the state dual graph, they could use the subdivided state dual graph to make a “typical” partition vertical. Whereas gerrymandering refers to the practice of drawing district boundaries to create desirable outcomes, we refer to practices that modify the input to a sampling algorithm to shift the seats histogram as *metamandering*. Unlike gerrymandering, metamandering is more of academic interest than a practical concern, since scenarios allowing metamandering require far fetched and sophisticated conspiracies. Nevertheless, it raises a point of scrutiny in the analysis of districting plans by Markov chains.

Implications for the $\sqrt{\epsilon}$ -test approach to redistricting analysis. The influence of the state-dual graph on the output of sampling algorithms also matters for the interpretation of methods like [14, 16, 57]. Their work provides a hypothesis test, called the $\sqrt{\epsilon}$ -test, for whether a plan was drawn from a distribution near to the stationary distribution of a Markov chain on districting plans. Crucially, and unlike outlier analysis based methods, the $\sqrt{\epsilon}$ -test does not rely on guarantees about convergence of the Markov chain. However, since our work shows that tunings of the score function and discretizations of the state can influence the stationary distribution in subtle ways, the results of the $\sqrt{\epsilon}$ -test should also be examined with respect to the influence of discretization and parameter tuning. We will briefly summarize this test, then relate our experiments to it.

The $\sqrt{\epsilon}$ -test takes as input a reversible and ergodic Markov chain M on a finite state space Ω , an initial point $X_0 \in \Omega$ and a label function $\omega : \Omega \rightarrow \mathbb{R}$, and is based on to what extent $\omega(X_0)$ is an “outlier” in $\text{Values} = (\omega(X_0), \omega(X_1), \dots, \omega(X_n))$, where X_i a sample path of M started at X_0 . The notion of outlier used here involves the index of $\omega(X_0)$ after sorting Values ; specifically, the test bounds the probability of $\omega(X_0)$ having index in the sorted list below $\lfloor \epsilon(n+1) \rfloor$, under the null hypothesis that X_0 is sampled from the stationary distribution of M . For applications to redistricting, the label ω can be some measurement of how much a plan favors one party.

One intuition for this test is discussed in [14, 57]: at least for small enough values of n a plan being rejected by this test can be interpreted as evidence that it is “carefully crafted,” as nearby plans have generally less extreme values of ω . It should be emphasized that the notion of nearby here refers to the typical plans arrived at after evolving the Markov chain, which is not necessarily an intuitive notion of locality. Indeed, as we explain next,




Figure 22: Take ω to be the number of seats won by party one. Arrange the voters so that every voter with x -coordinate $\leq .6$ is party one (yellow), and every other voter is party two (purple). In this set up, every roughly vertical partition gives one block to each party and every roughly horizontal partition gives both blocks to party one. The horizontal and vertical partitions are represented by dashed lines.

our examples from §IV D provide challenges for this intuition when n is sufficiently large. See question 8 in the conclusion below for some additional discussion.

We will use fig. 19 and fig. 20, which we introduced as examples where the stationary distribution of a given Markov chain algorithm is strongly influenced by the discretization, to show how the output of the $\sqrt{\epsilon}$ test can be influenced by the discretization. First, arrange the voters on the nodes in a grid graph discretization of $[0, 1]^2$ so that the horizontal and vertical partitions are starkly different under ω ; see fig. 22 for an example. If we let G be the graph constructed in fig. 19 and G' its 90 degree rotation, then using G or G' as the underlying graph to evaluate the horizontal partition under the $\sqrt{\epsilon}$ -test can result in very different outcomes. This is because, as fig. 19 and fig. 20(left) show, when using the flip walk and the tuning as the in those example, the partitions of G will be horizontal with very high probability, and partitions of G' will be vertical with high probability, regardless of the starting position.

Limitations. Algorithmic feasibility and compactness are but two of the many considerations taken into account when choosing a distribution for analyzing districting plans. While the mixing time of the critical fugacity chain appears promising, we do not explicitly recommend its use in redistricting analysis.

In this paper, we examined a single family of score functions on districting plans, using a specific notion of population balance and a notion of compactness largely driven by connections to statistical physics and SLE. However, in redistricting, score functions are constructed out of a variety of different political compactness measurements, notions of population balance, and other considerations, such as adherence to voting rights laws. These other scores come with their own complications. For instance, the Polsby-Popper compactness score is famously susceptible to the coastline paradox [26], while other common compactness scores interact with map projections used in practice to calculate them in ways that

can flip the preference ordering [4]. We view our study not as a comprehensive evaluation of potential distributions over partitions for use in redistricting analysis, but rather as an illustration of subtle challenges that arise when using these methods in practice—many of which are well-known in related statistical physics applications.

VI. Conclusion

As methods from statistical physics continue to be incorporated into political redistricting analysis, domain knowledge of statistical physics becomes increasingly relevant to developing best practices. To support collaboration, we outline some challenging open questions at the interplay of redistricting and statistical physics.

Statistical physics motivates many questions in redistricting analysis, including the following:

1. Are there phase transitions in the Gibbs energy scores mentioned in §III B? For instance, what about when other compactness scores are used? Do phase transition in other compactness scores, if they exist, interact with the effects of map projections studied in [4]?
2. Besides the β_λ distributions, are there other analytically approachable distributions on the space of connected partitions of the grid graph?
3. Are certain natural distributions over connected partitions of the grid graph, such as max entropy distributions, efficiently sampleable?
4. We know from statistical physics that there are certain features of models that are independent of the lattice being used. Although it would be surprising, perhaps there are similarly universal properties of outlier analysis?
5. Can we construct distributions over districting plans with fewer phase transitions or dependencies on discretization?
6. How does the interaction between phase transitions, connective constants, and population balance influence mixing time of algorithms used to sample districting plans?
7. In physics, Glauber dynamics-style chains are used with simulated tempering. Some redistricting studies, such as [29, 31], use this method. Can tempering weights on the population balance criterion, for instance, be chosen so as to make the bottlenecks we identified in the flip walk on the grid graph disappear? We note that in general the problem of uniformly sampling balanced connected partitions is NP-hard [54], but there are no known theoretical results forbidding uniform sampling from connected partitions of grid graphs.
8. Recall that n denotes the number of steps used for the $\sqrt{\epsilon}$ -test. Given that results show that large n can change the output of the $\sqrt{\epsilon}$ -test, what is a principled way to pick n in practice? For a given choice of n and some real world data, is it possible to beat the test by finding a gerrymandered map surrounded, in

the sense of the evolution of the Markov chain, by gerrymandered maps?

Redistricting also motivates statistical physics style questions, such as:

1. What is the influence of population balance on phase transitions in subcritical connected partitions of the grid?
2. Is it possible to efficiently sample from self-avoiding walks in the $n \times n$ grid graph according to the β_λ distribution?
3. When the endpoints are held fixed, does the flip walk mix rapidly at critical fugacity?
4. What is the correct way to generalize the connective constant to irregular lattice-like graphs?
5. How much of the population balance distribution in the critical fugacity case of fig. 13 can be calculated asymptotically via Schramm-Loewner evolution? See [52] for some discussion.
6. Is a typical connected partition of the grid approximately balanced? See [53] for some discussion.
7. Is there a scaling limit of the distribution of connected k -partitions of the $n \times n$ grid graph, with score function $\beta_\lambda(P) = \lambda^{\text{cut}(P)}$ and $\lambda = 1/\mu$, where μ is the connective

constant of the grid? What about just the 2-partition case, with or without population balance? If so, can discretizations of it be sampled directly, as is the case with Schramm-Loewner evolution [6]?

VII. Acknowledgements

We want to thank Assaf Bar-Natan, Sarah Cannon, Jowei Chen, Moon Duchin, Jordan Ellenberg, Gregory Herschlag, Matt Karmann, Ian Gordon Mark, Mathoverflow user Kostaya.I, Jonathan Mattingly, Marshall Mueller, Zohar Nussinov, John O'Neill, Wes Pegden, Dana Randall, Zach Schutzmann, Hannah Silverman, Annie Yun, Paul Zhang, and Seanna Zhang for helpful discussions and/or feedback during the preparation of this manuscript.

L. Najt, D. DeFord and J. Solomon acknowledge the support of the Prof. Amar G. Bose Research Grant. J. Solomon acknowledges the support of Army Research Office grants W911NF1710068 and W911NF2010168, of Air Force Office of Scientific Research award FA9550-19-1-031, and of National Science Foundation grant IIS-1838071. L. Najt acknowledges the support of NSF TRIPODS Award #1740707.

-
- [1] Hugo A Akitaya, Matthew D Jones, Matias Korman, Christopher Meierfrankenfeld, Michael J Munje, Diane L Souvaine, Michael Thramann, and Csaba D Tóth. Reconfiguration of Connected Graph Partitions. *arXiv preprint arXiv:1902.10765*, 2019.
 - [2] Boris Alexeev and Dustin G. Mixon. Partisan gerrymandering with geographically compact districts. *Journal of Applied Probability*, 55(4):1046–1059, 2018.
 - [3] Micah Altman. Is automation the answer: The computational complexity of automated redistricting. *Rutgers Computer and Law Technology Journal*, 23, 1997.
 - [4] Assaf Bar-Natan, Lorenzo Najt, and Zachary Schutzman. The gerrymandering jumble: map projections permute districts' compactness scores. *Cartography and Geographic Information Science*, pages 1–15, 2020.
 - [5] Amariah Becker and Justin Solomon. *Redistricting Algorithms*, chapter 11. Birkhäuser, 2020.
 - [6] Vincent Beffara. Coding SLEs (Schramm-Loewner Evolution) eg. SLE(6). MathOverflow. URL: <https://mathoverflow.net/q/244227> (version: 2016-07-13).
 - [7] Nathaniel Berestycki and James Norris. Lectures on Schramm-Loewner Evolution. *Lecture notes, available on the webpages of the authors*, 2014.
 - [8] US Census Bureau. Tiger/line shapefiles, 2010.
 - [9] Sophia Caldera, Daryl DeFord, Moon Duchin, Samuel C Gutekunst, and Cara Nix. Mathematics of Nested Districts: The Case of Alaska. *Statistics and Public Policy*, 7:39–51, 2020.
 - [10] Daniel Carter, Gregory Herschlag, Zach Hunter, and Jonathan Mattingly. A Merge-Split Proposal for Reversible Monte Carlo Markov Chain Sampling of Redistricting Plans. *arXiv:1911.01503*, 2019.
 - [11] Guy-Uriel E. Charles et al. Amicus brief of mathematicians, law professors, and students in support of the appellees and affirmance.
 - [12] Jowei Chen. Expert report of Jowei Chen, Ph.D. *Raleigh Wake Citizen's Association et al. vs. The Wake County Board of Elections*, November 2017.
 - [13] Jowei Chen. Expert report of Jowei Chen, Ph.D (Plaintiffs' brief regarding the voting rights act). *Common Cause et al. v. David Lewis, in his official capacity as senior chairman of the house select committee on redistricting, et al.*, September 2019.
 - [14] Maria Chikina, Alan Frieze, Jonathan Mattingly, and Wesley Pegden. Practical tests for significance in markov chains. *arXiv preprint arXiv:1904.04052*, 2019.
 - [15] Maria Chikina, Alan Frieze, and Wesley Pegden. Assessing significance in a Markov chain without mixing. *Proceedings of the National Academy of Sciences*, 114(11):2860–2864, March 2017.
 - [16] Maria Chikina, Alan Frieze, and Wesley Pegden. Assessing significance in a markov chain without mixing. *Proceedings of the National Academy of Sciences*, 114(11):2860–2864, 2017.
 - [17] Wendy Cho and Yan Liu. Toward a Talismanic Redistricting Tool: A Computational Method for Identifying Extreme Redistricting Plans. *Election Law Journal: Rules, Politics, and Policy*, 15, November 2016.
 - [18] Nathan Clisby. Calculation of the connective constant for self-avoiding walks via the pivot algorithm. *Journal of Physics A: Mathematical and Theoretical*, 46(24):245001, 2013.
 - [19] Vincent Cohen-Addad, Philip N Klein, and Dániel Marx.

- On the computational tractability of a geographic clustering problem arising in redistricting. *arXiv preprint arXiv:2009.00188*, 2020.
- [20] Daryl DeFord and Moon Duchin. Redistricting Reform in Virginia: Districting Criteria in Context. *Virginia Policy Review*, 12(2):120–146, 2019.
- [21] Daryl DeFord and Moon Duchin. *Random Walks and the Universe of Districting Plans*, chapter 12. Birkhäuser, 2020.
- [22] Daryl DeFord, Moon Duchin, and Justin Solomon. Recombination: A family of Markov chains for redistricting. *arXiv preprint arXiv:1911.05725*, 2019.
- [23] Daryl DeFord, Moon Duchin, and Justin Solomon. A Computational Approach to Measuring Vote Elasticity and Competitiveness. *Statistics and Public Policy*, 7:39–51, 2020.
- [24] Daryl DeFord and Zohar Nussinov. Universal phase transitions in segmentation or districting problems. *in preparation*, 2020.
- [25] Jian Ding and Yuval Peres. Mixing time for the Ising model: a uniform lower bound for all graphs. In *Annales de l’IHP Probabilités et statistiques*, volume 47, pages 1020–1028, 2011.
- [26] Moon Duchin and Bridget Eileen Tenner. Discrete geometry for electoral geography. *arXiv preprint arXiv:1808.05860*, 2018.
- [27] Hugo Duminil-Copin, Gady Kozma, and Ariel Yadin. Supercritical self-avoiding walks are space-filling. In *Annales de l’IHP Probabilités et statistiques*, volume 50, pages 315–326, 2014.
- [28] Hugo Duminil-Copin and Stanislav Smirnov. The connective constant of the honeycomb lattice equals $\sqrt{2} + \sqrt{2}$. *Annals of Mathematics*, 175(3):1653–1665, 2012.
- [29] Benjamin Fifield, Michael Higgins, Kosuke Imai, and Alexander Tarr. A new automated redistricting simulator using markov chain monte carlo. *Work. Pap., Princeton Univ., Princeton, NJ*, 2015.
- [30] Benjamin Fifield, Michael Higgins, Kosuke Imai, and Alexander Tarr. Automated redistricting simulation using markov chain monte carlo. *Journal of Computational and Graphical Statistics*, 0(0):1–14, May 2020.
- [31] Benjamin Fifield, Michael Higgins, Kosuke Imai, and Alexander Tarr. Automated redistricting simulation using markov chain monte carlo. *Journal of Computational and Graphical Statistics*, 29(4):715–728, 2020.
- [32] Benjamin Fifield, Kosuke Imai, Jun Kawahara, and Christopher T Kenny. The essential role of empirical validation in legislative redistricting simulation. *Statistics and Public Policy*, 7(1):52–68, 2020.
- [33] Andrew Gelman and Christian Hennig. Beyond subjective and objective in statistics. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 180(4):967–1033, 2017.
- [34] Gregory Herschlag, Han Sung Kang, Justin Luo, Christy Vaughn Graves, Sachet Bangia, Robert Ravier, and Jonathan C. Mattingly. Quantifying Gerrymandering in North Carolina. 2018.
- [35] Gregory Herschlag, Robert Ravier, and Jonathan C. Mattingly. Evaluating Partisan Gerrymandering in Wisconsin. *arXiv:1709.01596 [physics, stat]*, September 2017.
- [36] M. Huber. Efficient exact sampling from the Ising model using Swendsen-Wang. *Soda ’99*, pages 921–922, 1999.
- [37] M. Huber. A bounding chain for Swendsen-Wang. *Random Structures and Algorithms*, 22:43–59, 2002.
- [38] Thomas Rogers Hunter. The First Gerrymander? Patrick Henry, James Madison, James Monroe, and Virginia’s 1788 Congressional Districting. *Early American Studies*, pages 781–820, 2011.
- [39] Iwan Jensen. Improved lower bounds on the connective constants for two-dimensional self-avoiding walks. *Journal of Physics A: Mathematical and General*, 37(48):11521, 2004.
- [40] Tom Kennedy. Monte Carlo tests of stochastic Loewner evolution predictions for the 2D self-avoiding walk. *Physical review letters*, 88(13):130601, 2002.
- [41] Richard Kueng, Dustin G Mixon, and Soledad Villar. Fair redistricting is hard. *Theoretical Computer Science*, 2019.
- [42] Gregory Lawler, Oded Schramm, and Wendelin Werner. Conformal restriction: the chordal case. *Journal of the American Mathematical Society*, 16(4):917–955, 2003.
- [43] Gregory F Lawler, Oded Schramm, and Wendelin Werner. On the scaling limit of planar self-avoiding walk. *arXiv preprint math/0204277*, 2002.
- [44] Yan Y. Liu, Wendy K. Tam Cho, and Shaowen Wang. PEAR: a massively parallel evolutionary computation approach for political redistricting optimization and analysis. *Swarm and Evolutionary Computation*, 30:78–92, October 2016.
- [45] László Lovasz. A homology theory for spanning tress of a graph. *Acta Mathematica Hungarica*, 30(3-4):241–251, 1977.
- [46] Neal Madras and Gordon Slade. *The self-avoiding walk*. Springer Science & Business Media, 1996.
- [47] Kenneth C Martis. The original gerrymander. *Political Geography*, 8(27):833–839, 2008.
- [48] Jonathan Mattingly. Declaration of Jonathan Mattingly. *Common Cause vs. Rucho*, March 2017.
- [49] Jonathan C Mattingly and Christy Vaughn. Redistricting and the will of the people. *arXiv preprint arXiv:1410.8796*, 2014.
- [50] Marina Meila. How to sample connected K -partitions of a graph. *arXiv:1808.00050*, 2018.
- [51] Cristopher Moore and Stephan Mertens. *The nature of computation*. OUP Oxford, 2011.
- [52] Lorenzo Najt. The distribution of the area of a region cut out by chordal sle. *MathOverflow*. <https://mathoverflow.net/q/333336>.
- [53] Lorenzo Najt. Left passage probability of $SLEs$. *MathOverflow*. <https://mathoverflow.net/q/335793>.
- [54] Lorenzo Najt, Daryl Deford, and Justin Solomon. Complexity and geometry of sampling connected graph partitions. *arXiv preprint arXiv:1908.08881*, 2019.
- [55] See [38, 47] for historical accounts.
- [56] Sejun Park, Yunhun Jang, Andreas Galanis, Jinwoo Shin, Daniel Stefankovic, and Eric Vigoda. Rapid Mixing Swendsen-Wang Sampler for Stochastic Partitioned Attractive Models. *arXiv preprint arXiv:1704.02232*, 2017.
- [57] Wesley Pegden. Pennsylvania’s congressional districting is an outlier: Expert report. *League of Women Voters vs. Pennsylvania General Assembly*, November 2017.
- [58] Daniel D Polsby and Robert D Popper. The third criterion: Compactness as a procedural safeguard against partisan gerrymandering. *Yale Law & Policy Review*, 9(2):301–353, 1991.

- [59] Alan D Sokal. Monte Carlo methods for the self-avoiding walk. *arXiv preprint hep-lat/9405016*, 1994.
- [60] Alan D Sokal and Lawrence E Thomas. Exponential convergence to equilibrium for a class of random-walk models. *Journal of Statistical Physics*, 54(3-4):797–828, 1989.
- [61] Hamidreza Validi, Austin Buchanan, and Eugene Lykhovyd. Imposing contiguity constraints in political districting models. 2020.
- [62] Tien Vo, Vamsi Ithapu, Vikas Singh, and Michael A Newton. Dimension constraints improve hypothesis testing for large-scale, graph-associated, brain-image data. *arXiv preprint arXiv:1908.07176*, 2019.