OPTIMIZING REDISTRICTING PLAN SELECTION FOR U.S. CONGRESSIONAL DISTRICTS WITH VARIOUS DEFINITIONS

JOHN ASPLUND

ABSTRACT. Redrawing lines for redistricting plans that represent U.S. congressional districts is a tricky business. There are many laws that dictate how lines can and cannot be drawn such as contiguity. In fact, building all redistricting plans for a single U.S. state is an intractable problem. Researchers have turned to heuristics in order to analyze current redistricting plans. Many of these heuristics (e.g. local search heuristics and Markov chain Monte Carlo algorithms) used by researchers form new congressional districts by switching the smaller pieces (e.g. precincts or census blocks) that make up congressional districts from one congressional district to another. In this paper, we discuss the various natural definitions involved in satisfying rules for contiguity and simply connectedness of precincts or census blocks and how these relate to contiguity and simply connectedness of congressional districts. We also propose and analyze several constructions to alleviate violations of contiguity and simply connectedness in precincts and census blocks. Finally, we develop efficient algorithms that allow practitioners to assess redistricting plans using local search heuristics or Markov chain Monte Carlo algorithms efficiently.

Keywords: redistricting, contiguity, simply connected, planarity

Mathematics Subject Classification (2010) 05C85, Graph algorithms · 05C40, Connectivity · 05C10, Planar graphs · 05A18, Partitions of sets

1. INTRODUCTION

Every ten years, the United States redraws lines representing congressional districts for each state in response to data provided by the national census. The geometry representing all of the congressional districts of a state is called a *redistricting plan*. The rules and laws governing how to redraw congressional district lines varies from state to state. In the U.S., one iron-clad rules must be followed by all states: one-person one-vote. That means the population in each congressional district in a state must

Date: March 19, 2019.

have roughly the same population as any other congressional district in that state.

People who redraw the district lines are typically the same people who have a personal stake in retaining power or disadvantaging certain groups such as minorities. Consequently, the past is riddled with cases of implicit or overt favoritism towards one's political party or socioeconomic group. Take for instance the case of Elbridge Gerry, who was the governor of Massachusetts in the early nineteenth century. The legislature at the time created congressional districts that highly favored one party and Gerry signed the legislation for the oddly shaped redistricting plan. One cartoonist likened one of the districts to that of a mythical salamander. A portmanteau of Elbridge's last name and the salamander from the cartoon formed the term, gerrymandering. Specifically, if a particular group manipulates the boundaries of any region for political advantage, then it is called *gerrymandering*.

Assessing whether a congressional district has been gerrymandered has proven difficult for many reasons. First and foremost, the geometry of state borders is restricted by the geography of the land and water, (e.g. Hawaii). Another limitation in determining whether a state has been gerrymandered is a lack of consistent rules among states for redrawing lines for congressional districts. There is also the complication of determining a "fair" metric that would assess whether a redistricting plan has been gerrymandered. Even if such a metric could be developed, one would still have to check all redistricting plans with this metric. However, such an undertaking would be an astronomically difficult task, as explained in Section 2. Instead of determining all possible redistricting plans, researchers have used various heuristics to explore a sampling of all the possible redistricting plans as a proxy for finding all redistricting plans. One such heuristic is called a local search heuristic.

Local searches are heuristics used to solve difficult optimization problems by moving in a step-by-step manner "locally" in the solution space. The aim is to find a new solution by improving on the current solution. One modern type of algorithm that also moves in the same step-by-step manner as local searches and is used to assess redistricting plans is called a *Markov chain Monte Carlo* (MCMC) algorithm. An MCMC algorithm is a sampling technique used on complicated distributions or on distributions we know nothing about. It explores the solution space in a step-by-step manner by randomly picking the next state based solely on the current state, just like local searches.

Recently, MCMC algorithms have been built to assess redistricting plans of congressional districts. This is done by dividing each congressional district into smaller pieces, typically census blocks or voting tabulation districts (precincts). For the sake of argument, suppose the congressional districts are divided into census blocks. Then one redistricting plan can move to another redistricting plan by randomly selecting one census block on the border between two congressional districts, and selecting a neighboring congressional district. Finally, if the chosen census block and congressional district pass predetermined criteria, the chosen census block is moved into the chosen congressional district. Since MCMC algorithms and local searches can be defined to move between redistricting plans in the same manner, the same criteria can be used for local searches with the addition of an objective function. Two criteria we will use to assess new redistricting plans in this paper are contiguity and simply connectedness; we will discuss these criteria further in Section 2. In Section 2, we also describe various definitions for adjacency between census blocks that will be used throughout this paper.

In Section 3, we describe several constructions that turns discontiguous census blocks into contiguous census blocks and that turns non-simply connected census blocks into simply connected census blocks. We also discuss how these constructions and the choice of definitions from Section 2 can impact data typically associated by the census blocks when using MCMC algorithms or local searches.

In Section 4, we examine and discuss several algorithms that can be used to check whether the census blocks or congressional districts on a redistricting plan are contiguous and simply connected. Additionally, we focus on determining whether the congressional districts of a redistricting plan are contiguous and simply connected after each iteration of a local search or MCMC algorithm has been applied to said redistricting plan. We describe this iteration process in Construction 11.

Since we can form a graph that is similar to the dual graph on the census blocks (call it G^* for the sake of simplicity), we can leverage the structure on G^* to identify when census blocks or congressional districts are contiguous or simply connected. In [15], when using a particular definition for adjacency, contiguity, and simply connectedness, the authors showed that all non-simply connected congressional districts could be identified in $\mathcal{O}(k^2)$ time on a redistricting plan where all k congressional districts are contiguous and all census blocks are contiguous and simply connected. For a census block v, let $R(v)$ represent the set of census blocks that are adjacent to v at at least one point. in [15], the authors also showed that discontiguous congressional districts can be identified in $\mathcal{O}(R(v))$ time for some census block v after each iteration of a local search or MCMC algorithm. In Section 4.1, we explain how to exploit articulation points (cut-vertices) on G^* to find all non-simply connected census blocks in $\mathcal{O}(|G^*|)$ time if G^* is planar. Similarly, if G^* is planar, we explain how to find all the non-simply connected congressional districts in $\mathcal{O}(k)$ time where k is the number of zones on a redistricting plan. This method can also be applied after each iteration of a local search or MCMC algorithm with the same efficacy as described in Section 4.1. Note that using the articulation points on G^* does not require G^* to be planar, and so we explain what this means when G^* is non-planar. See Section 4.1 for more details.

In Section 4.2, given a redistricting plan where all census blocks and congressional districts are both simply connected and contiguous, we provide algorithms that can be applied to the redistricting plan formed after each iteration of a local search or MCMC algorithm to determine whether the congressional districts in this new redistricting plan are *both* contiguous and simply connected in $\mathcal{O}(|R(v)|)$ time where v is a unit chosen by the local search or MCMC algorithm. We prove these claims in Section 4.3.

2. Background & Definitions

To assess whether a redistricting plan has been gerrymandered, one technique would be to examine the distribution of voting results for all redistricting plans in comparison to the current redistricting plan. Building the entire space of all possible redistricting plans on a set of census blocks or precincts is astronomically large. Let us examine Hawaii, which is the state with the smallest number of census blocks with at least two congressional districts as of the 2010 national census. With 25,016 census blocks, Hawaii has $2^{25,016} \approx 10^{7530}$ possible redistricting plans if no additional requirements are placed on the redistricting plans. Even if we used precincts instead of census blocks, Hawaii has 248 precincts, and so there would be $2^{248} \approx 10^{74}$ redistricting plans, which is slightly less than the number of atoms in the universe (about 10^{80}), but more than the number of stars in the universe (about 10^{24} according to the European Space Agency in 2015). Now what?

Some researchers have turned to heuristics as a method to avoid finding all possible redistricting plans. Kalcsics et al. [14] provided a detailed description of several heuristic methods that went into these redistricting problems. More recently, several groups have been utilizing heuristics based on local search methods and computational geometry [15, 16, 17, 21] to explore the space of all redistricting plans given a set of census blocks. Other groups have been using MCMC algorithms with the Potts Model [2, 5, 11, 13]. We will focus solely on local searches and MCMC algorithms as was done in [2, 5, 9, 11, 13, 15, 16, 17].

2.1. Definitions.

Depending on the discipline, a simple closed curve can have several definitions. For our purposes, a *simple curve* is a curve that does not intersect itself. A *simple closed curve* is a simple curve that begins and ends at the same point. A *unit* is the union of open sets in $\mathbb{R}^2$ with finite area. Essentially, each unit represents a geographic area. Let $\mathcal{U}$ be the set of all units in $\mathbb{R}^2$. We call each disconnected open set of a unit v

a *piece*. A more rigorous definition for a piece of a unit will be given after we give various definitions for contiguity.

There are several possible candidates that can be used to represent a unit. For example, one can use counties to represent units, although these are sometimes partitioned by congressional districts due to population restrictions placed on redistricting plans. Another candidate that can be used to represent units is a *voting tabulation district* (VTD). These are also sometimes called precincts, wards, voting precincts, municipalities, or municipal wards. VTDs are typically smaller than a county, but are sometimes partitioned further by congressional districts due to the constraints placed on the U.S. states by laws (e.g., one-person one-vote). The smallest unit used for the entire U.S. that has accurate population data is the census block. Census blocks are units built and maintained by the Census Bureau every 10 years for the national census. We will primarily consider census blocks as our unit even though we continue using the term "unit."

We will use units to form a partition of congressional districts. As such, we need a term for each part of our partition that will represent congressional districts. A *zone* z' is the union of a set of non-overlapping (the interiors of two units do not overlap) units. To make some explanations that refer to both zones and units non-redundant, we define a *region* to be either a zone or a unit; whenever we discuss the interactions between several regions, either all regions are assumed to be zones or all regions are assumed to be units. A *map* on a set of units $\mathcal{U}$ is defined as the set of all units in $\mathcal{U}$ where each unit is assigned to exactly one zone. A redistricting plan is an example of a map. An *admissible* map (*admissible* redistricting plan) is a map that follows a set of conditions (e.g. contiguity or simply connectedness) by which to judge the validity of a map. U.S. states are an example of a map, congressional districts are an example of a zone, and census blocks are an example of a unit. An assumption we will make throughout this paper is that given a map M, there is a fixed integer $k \geq 2$ representing the number of zones for which units in $\mathcal{U}$ are assigned. Note that even though a zone is defined using units, it, too, is composed of one or more simple closed curves. Note that it is rare to find U.S. state redistricting legislation that defines these terms so carefully.

Unfortunately, there are few criteria that all U.S. states have agreed upon when dictating rules to follow when building redistricting plans. Some criteria that have been considered by both academics and governments include the following: contiguity, compactness, simply connectedness, communities of interest, and political subdivisions. We will discuss only contiguity and simply connectedness in this paper, although it should be noted that compactness is in the legislation for more than a third of U.S. states. We do not discuss compactness here since a thorough investigation into the complexity and methods for calculating compactness of a congressional district was conducted by Barnes and Solomon [3].

The most widely adopted rule in legislation regarding redistricting plans, not including one-person one-vote, is the contiguity criterion. One way some people define contiguity is as follows: a region is contiguous if one can travel from any point in a region to any other point in a region while staying in said region. Unfortunately, the manner in which contiguity is defined varies from U.S. state to U.S. state. For example, consider the region labeled v_1 in Figure 3(a). The laws of some U.S. states—such as South Carolina [8]—would say that the left piece of v_1 in Figure 3(a) is connected to the right piece of v_1, while the laws of other states—such as Minnesota [19]—would disagree. For the purposes of being as specific as possible, we give three definitions of contiguity for a unit based on either definitions used by other researchers, based on legislation, or based on definitions that fit naturally with the other definitions we have already stated. We also give three definitions for the adjacency between units and zones, and for when a unit or zone is simply connected. All definitions are stated in terms of units on $\mathbb{R}^2$.

The following three competing definitions for adjacency between two units can be attributed to GIS software and the Dimensionally Extended Nine-Intersection Model (DE-9IM) [6, 7]. For each $v \in \mathcal{U}$, let $\partial(v)$ be the set of points representing the boundary of v.

(A1) Let $\mathcal{U}$ be a set of units. Two units are *rook-adjacent* if and only if they share a boundary of some positive length in common. That is, two units u, v are *rook-adjacent* if and only if the length of $\partial(u) \cap \partial(v)$ is positive (the length of boundary shared by u and v is positive).

(A2) Let $\mathcal{U}$ be a set of units. Let $X = \{X_v : v \in \mathcal{U}\}$ where X_v is a finite set of points on $\partial(v)$ and for each pair of units, $v_1, v_2 \in \mathcal{U}$, let $X_{v_1} \cap X_{v_2} = \varnothing$. Let $u, v \in \mathcal{U}$. Then u and v are *semi-queen-adjacent* if and only if either the length of $\partial(u) \cap \partial(v)$ is positive or u and v share at least one point from $X_v \cup X_u$.

(A3) Let $\mathcal{U}$ be a set of units. Two units $u, v \in \mathcal{U}$ are *queen-adjacent* if and only if $|\partial(u) \cap \partial(v)| > 0$.

Definitions (A1) and (A3) are used in several geospatial software such as the R package spdep [22].

To make defining contiguity easier, we denote the interior of a unit v as $\text{INT}(v)$.

(B1) Let $\mathcal{U}$ be a set of units. Then $v \in \mathcal{U}$ is *rook-contiguous* if and only if for each pair of points in $\text{INT}(v)$, there exists a simple curve entirely contained in $\text{INT}(v)$.

(B2) Let $\mathcal{U}$ be a set of units. Let $Y = \{Y_v : v \in \mathcal{U}\}$ where Y_v is a finite set of points in $\partial(v)$ and for each pair of units v_1 and v_2, $Y_{v_1} \cap Y_{v_2} = \varnothing$. Then $v \in \mathcal{U}$ is *semi-queen-contiguous* if and only if for each pair of points in $\text{INT}(v) \cup Y_v$, there exists a simple curve that only contains points in $\text{INT}(v) \cup Y_v$.

(B3) Let $\mathcal{U}$ be a set of units. Then $v \in \mathcal{U}$ is *queen-contiguous* if and only if for each pair of points in $\text{INT}(v)$, there exists a simple curve entirely contained in $\text{INT}(v) \cup \partial(v)$.

Rook-contiguous and queen-contiguous can be found in [24]. Definition (B2) is based on the description for contiguity provided by the state of South Carolina for building redistricting plans [8].

Now that we have defined contiguous, we can properly define a piece of a unit. Let contiguity be defined by either Definition (B1), (B2), or (B3). A *piece* P of a unit v is a contiguous (either rook-, semi-queen-, or queen-contiguous) union of open sets from v where for each piece P' of v disjoint from P, $P \cup P'$ is discontiguous.

Although it is not as widely discussed as contiguity, holes in congressional districts are sometimes defined or discussed using the term *simply connected* [1, 23]. There are some mediums for which ensuring that a zone is simply connected is a desirable trait [12, 20, 25]. Like the various definitions of contiguity, we provide three definitions for simply connected. Typically, a unit v is defined as simply connected if any simple closed curve in v can be shrunk to a point continuously. The variations in the definitions below for a simply connected unit stem from whether a simple closed curve is allowed to pass through the boundary of the unit. We present three options for the definition of a simply connected unit based on this premise.

(C1) Let $\mathcal{U}$ be a set of units. Then $v \in \mathcal{U}$ is *simply interior-connected* (simply *i*-connected) if and only if any closed simple curve in $\text{INT}(v)$ can be shrunk to a point continuously.

(C2) Let $\mathcal{U}$ be a set of units. Let $W = \{W_v : v \in \mathcal{U}\}$ where W_v is a finite set of points from $\partial(v)$ and for each pair of units v_1 and v_2, $W_{v_1} \cap W_{v_2} = \varnothing$. Then $v \in \mathcal{U}$ is *simply semi-boundary-connected* (or simply *sb*-connected) if and only if any simple closed curve in $\text{INT}(v) \cup W_v$ can be shrunk to a point continuously.

(C3) Let $\mathcal{U}$ be a set of units. Then $v \in \mathcal{U}$ is *simply boundary-connected* (or simply *b*-connected) if and only if any simple closed curve in $\text{INT}(v) \cup \partial(v)$ can be shrunk to a point continuously.

Note that when we say two units are adjacent, or a unit is contiguous, or a unit is simply connected, then it is assumed we could be using any of the above respective definitions. If v is not simply connected, we say *v has a hole* and if there exists a simple closed curve in v that entirely contains a unit w, then we say *w is a hole in v* or *w is a hole of v*. Since zones are a partition of $\mathcal{U}$, two zones z_1, z_2 are adjacent if a unit in z_1 is adjacent to z_2. Similarly, a zone z_1 is contiguous if the union of all units in z_1 is contiguous. Finally, a zone z_1 is simply connected if the union of all units in z_1 is simply connected.

There are situations when a unit can have a hole regardless of the definition used for simply connectedness. For example, no matter which simply connected definition is chosen, v_1 is always a hole and v_2 always has a hole in Figure 1. That is not the case with the map in Figure 2. In Figure 2, v_2 has a hole by Definition (C3), sometimes has a hole by Definition (C2), and has no hole by Definition (C1). In Figure 2, v_1 is called a *degenerate hole.*

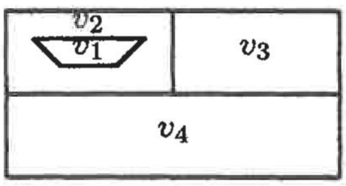

FIGURE 1. Example of a unit with a non-degenerate hole

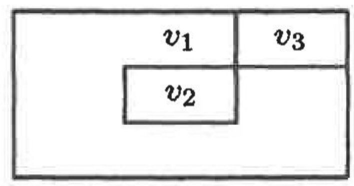

FIGURE 2. Example of a unit that may or may not be a hole

A serious complication can arise from conflicts between the definitions we defined above for adjacency, contiguity, and simply connectedness. For example, if M is the map in Figure 2, then using Definitions (A3) and (C3) would mean v_2 is a hole in v_1, and v_2 is adjacent to v_3, yet v_1 is not simply connected. We propose a reasonable assumption to resolve this discrepancy (supported by topological definitions of simply connectedness) that would need to hold for any combination of definitions described above for adjacency, contiguity, and simply connectedness. Recall that a region is either a unit or a zone.

(2.1) *If a region u or piece of u is in a hole of region v, then all of the area enclosed by u must also be in some hole in v, and any region adjacent to u must also be in a hole in v.*

In other words, if a region u or piece of u is in a hole of region v, then u cannot be adjacent to any unit that is not in a hole of v besides v itself. Below, we give the 3-set combinations of definitions for adjacency, contiguity, and simply connectedness that satisfy Assumption (2.1) above. Let $\mathcal{D}$ be the following set of 3-sets of definitions:

$$\{(A1),(B1),(C3)\},\{(A1),(B1),(C2)\},\{(A1),(B1),(C1)\},$$
$$\{(A1),(B2),(C1)\},\{(A1),(B3),(C1)\},\{(A2),(B1),(C1)\},$$
$$\{(A2),(B2),(C1)\},\{(A2),(B3),(C1)\},\{(A3),(B1),(C1)\},$$
$$\{(A3),(B2),(C1)\},\{(A3),(B3),(C1)\}.$$

$\mathcal{D}$ represents the set of allowable combinations of definitions for adjacency, contiguity, and simply connectedness by Assumption (2.1). To make notation easier, for example, instead of saying that a map M is using Definitions (A1), (B1), and (C3), we will say that M is using definition

$$\{(A1),(B1),(C3)\}.$$

Additionally, if we are using definition $D \in \mathcal{D}$ and Definition (A1) is in D, we will simply say $(A1) \in D$. As a justification for why a 3-set $D \notin \mathcal{D}$ violates Assumption (2.1), see either Figure 4(a) (when $(A3) \in D$) or Figure 7 (when $(A1) \in D$).

A *dual graph* G^* of a planar graph G is a graph where each vertex in G^* represents a face of G (including the outside face), and for each edge e in G, join two vertices $u, v \in V(G^*)$ if the two faces in G that represent u and v both contain the edge e on the boundary of their respective faces. Note that a dual graph is allowed to have multiedges or loops, but both multiedges and loops are meaningless in the context of redistricting plans. (Perhaps one day someone will find meaning behind multiedges and loops.) Dual graphs can also be formed from a set of units by using the boundaries of units just like edges in a graph. To account for the multiedges, loops, and the multiple definitions for adjacency, we will need to build a new definition for a dual graph. Let $\mathcal{U}_0 \supset \mathcal{U}$ be the set of units in $\mathbb{R}^2$ that includes a dummy unit u_0 that corresponds to the single infinite area in $\mathbb{R}^2 - \mathcal{U}$. We also assign this dummy unit to zone 0, a dummy zone. An *augmented dual graph* G^* of a set of units $\mathcal{U}_0$ is a graph where each unit (including the dummy unit) represents a vertex in G^* and two vertices in G^* are adjacent if the respective units are adjacent. We follow the example set by [15, 16, 17] and build a structure from this augmented dual graph called a geo-graph. A *geo-graph*, $\mathcal{G}(M) = (V, E, \partial, z_0, k, \mathcal{U}_0, X, Y, Z, D)$, is an augmented dual graph of a map M with k zones on the units in $\mathcal{U}_0$ using definition $D \in \mathcal{D}$ where the following hold:

- V and E are the vertex and edge sets of the augmented dual graph respectively,
- ∂ is a boundary function that maps each unit $v \in \mathcal{U}_0$ to all of the simply closed curves that represent v denoted as $\partial(v)$,
- $z : \mathcal{U} \to \{1, \dots, k\}$ is a zoning function that maps each unit $v \in \mathcal{U}$ to exactly one zone according to M, and $z_0 : \mathcal{U} \to \{0, \dots, k\}$ is a zoning function such that $z_0(v) = z(v)$ for all $v \in \mathcal{U}$ and $z_0(u_0) = 0$, and
- X, Y, Z are each a set of sets of points as defined in Definitions (A2), (B2), (C2) respectively (but X, Y, Z may be empty if we are not using the respective definition in $\{(A2), (B2), (C2)\}$).

Even though each geo-graph contains a large list of specific variables, we will assume that k is fixed; the zoning assignment is based on M; M is defined based on the structure of $\mathcal{U}$; X, Y, Z are all defined based on D; and D was already set when M was defined. In this way, we will treat $\mathcal{G}(M)$ just like another graph $G = (V, E)$ almost exclusively. To help identify the units in a particular zone, we let $\mathcal{U}_M(j) \subseteq \mathcal{U}$ denote the set of units in zone j on map M.

Note that the definitions for adjacency, contiguity, and simply connectedness used in [15, 16, 17] are equivalent to definition $\{(A1), (B1), (C1)\}$.

For all maps M on $\mathcal{U}$, it is assumed that the boundary of each unit is a collection of simple closed curves. We make the additional assumption that the intersection of the interior of any two regions is empty (i.e. there is no area shared between any pair of regions). We assume that all sets of units $\mathcal{U}$ are connected, that is, for each pair of points p_1, p_2 in units in $\mathcal{U}$, there exists a simple curve through the units in $\mathcal{U}$ connecting p_1 and p_2.

3. Reducing Computational Difficulties Through Geometry

In this section, we propose several constructions to address when units that are discontiguous or not simply connected by building a new set of units. We also address the flow of data typically associated with units. For each unit, we will assume the following additional information has been associated with each unit: perimeter, area, number of votes for each party candidate, the total population of the unit, the voting age population, and total population for each demographic. An entire paper could be devoted solely to a discussion on the manner and difficulty of collecting all of this data, but we will forgo this experience.

There are two broad categories of violations on the boundaries of units that we consider in this paper: contiguity and simply connectedness. Note that these violations on the boundaries of the units were called a piece violation and a hole violation respectively in [16]. Depending on the definition used from $\mathcal{D}$, there will be distinctive differences in the manner in which on the boundaries of the units occur.

Note that an analysis of various on the boundaries of the units was conducted in [16] along with several proposed ways in which to resolve each geometry, using only definition $D = \{(A1), (B1), (C1)\}$.

3.1. Contiguity.

We still assume that we are operating under a definition in $\mathcal{D}$ and Assumption (2.1), but we are aiming to fix contiguity violation, so we allow an initial set of units $\mathcal{U}$ to violate Assumption (2.1) in this section. Since we are entirely concerned with contiguity in this section, we will disregard concerns of simply connectedness.

For the purpose of making contiguous units from discontiguous units, we build constructions in the spirit of certain methods described in [16]. Visualizations of these constructions are shown in Figure 3.

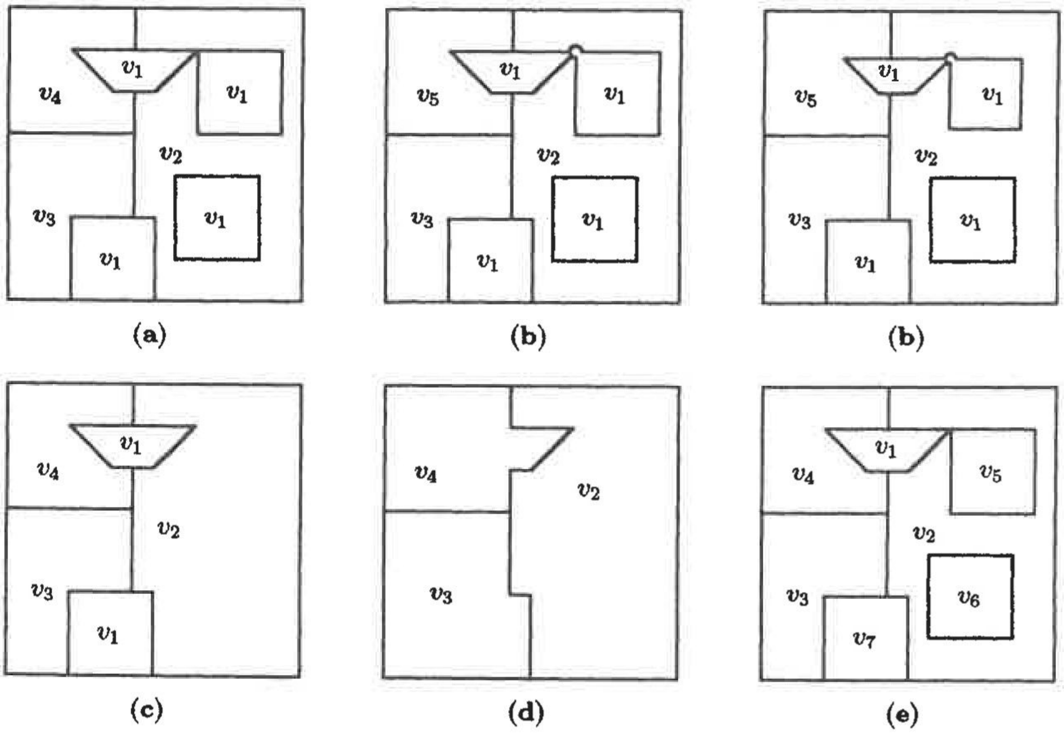

FIGURE 3. Five methods to resolve discontiguous units: (a) original map, (b) violation resolved with Construction 1, (c) violation resolved with Construction 2, (d) violation resolved with Construction 3, (e) violation resolved with Construction 4, and (f) violation resolved with Construction 5.

Construction 1. Let $\mathcal{U}$ be the set of units on $\mathbb{R}^2$ and let $v \in \mathcal{U}$ be discontiguous where v' and v'' are two pieces of v that share a point x. Let v_1 be formed from v by joining v' and v'' into one continuous open set defined by a single simple closed curve by creating an ε radius ball B around x, and "widening" the point x between all pieces of v with non-empty intersection with B. Let $U \subseteq \mathcal{U}$ be a set of units such that $\widehat{v} \in U$ if and only if the intersection of $\widehat{v}$ and v_1 has non-empty area. Let $U' = \{\widehat{v} \cap \overline{v_1} : \widehat{v} \in U\}$ where $\overline{v_1}$ is the complement of v_1 in $\mathbb{R}^2$. Form a new set of units $\mathcal{U}'$ by replacing $\{v\} \cup U$ in $\mathcal{U}$ with $\{v_1\} \cup U'$. This is called *opening*. ■

In the next construction, we use a term from Euclidean geometry to describe how to change units called uniform scaling. *Uniform scaling* is a linear transformation that enlarges or shrinks an object by a scale factor that is the same in all directions.

Construction 2. Let $\mathcal{U}$ be the set of units on $\mathbb{R}^2$ and let $v \in \mathcal{U}$ be discontiguous where v' and v'' are two pieces of v that share a point x. Let v_1 be formed from v by joining v' and v'' into one continuous open set defined by a single simple closed curve by creating an ε radius ball B around x, and "widening" the point x between all pieces of v with non-empty intersection with B while shrinking the area of the joined pieces v' and v'' with uniform scaling. That is, after widening the point at x and uniformly scaling v' and v'', the area in v_1 will be the same as v, and any unit in $\mathcal{U} \setminus \{v\}$ adjacent to v will fill in the area that was occupied by v by not by v_1. Let $U \subseteq \mathcal{U} \setminus \{v\}$ be a set of units that had their boundaries changed in this process. Let U' be the new units formed from units in U. Form a new set of units $\mathcal{U}'$ by replacing $\{v\} \cup U$ in $\mathcal{U}$ with $\{v_1\} \cup U'$. This is called a *proportional opening*. ■

Construction 3. Let $\mathcal{U}$ be the set of units on $\mathbb{R}^2$ and let $v \in \mathcal{U}$ be discontiguous. Let $v_1, v_2, \ldots, v_t$ be units in $\mathcal{U}$ such that for each v_i, there is a piece of v that is in a hole of v_i. For each $i \in \{1, \ldots, t\}$, form a unit v_i' from v_i by removing all pieces from v that are in a hole in v_i, and taking the union of v_i and all of these pieces of v; let v' be the unit formed from v by removing all of the pieces that were in holes in some v_i. Then we form a new set of units $\mathcal{U}'$ by replacing $v, v_1, v_2, \ldots, v_t$ in $\mathcal{U}$ with $v', v_1', v_2', \ldots, v_t'$ respectively. This is called *annexing*. ■

Construction 4. Let $\mathcal{U}$ be the set of units on $\mathbb{R}^2$ and let $v \in \mathcal{U}$ be discontiguous. For each piece $\widehat{v_i}$ in v, designate an adjacent unit $v_i \in \mathcal{U} \setminus \{v\}$. For each $i \in \{1, \ldots, t\}$, form a unit v_i' from v_i by removing $\widehat{v_i}$ from v, and taking the union of v_i and $\widehat{v_i}$; let v' be the unit formed from v by removing each $\widehat{v_i}$. Then we form a new set of units $\mathcal{U}'$ by replacing $v, v_1, v_2, \ldots, v_t$ in $\mathcal{U}$ with $v', v_1', v_2', \ldots, v_t'$ respectively. This is called *merging*. ■

Construction 5. Let $\mathcal{U}$ be the set of units on $\mathbb{R}^2$ and let $v \in \mathcal{U}$ be discontiguous. Let $v_1', v_2', \ldots, v_t'$ be each piece of v. Then we form a new set of units $\mathcal{U}'$ by replacing v with $v_1', v_2', \ldots, v_t'$. This is called *splitting*. ■

3.1.1. *Opening.* Note that the proof to Theorem 2 in [16] gives an explanation for another way to perform Construction 1 without a ball of radius ε. Let $\mathcal{U}$ be a set of units under definition $D \in \mathcal{D}$. We discuss both Constructions 1 and 2 in this section. We skip an analysis of how Construction 1 affects the solution space of maps when using $D = \{(A1), (B1), (C1)\}$ as this was thoroughly done in [16]. Correspondingly, we need not discuss how Construction 2 impacts the redistricting plans under the same definition. Since we are interested in contiguity violations, we are assuming two pieces of some unit v are non-adjacent, and since Definition (B3) implies that all pieces can be adjacent at a point, we will ignore the case where $(B3) \in D$. Similarly, since the case when $(B2) \in D$ will be identical to either the case when $(B1) \in D$ or $(B3) \in D$, we will ignore this case as well. For the remainder of this section, we will assume that $(B1) \in D$ unless stated otherwise.

Since both area and perimeter are used for some compactness calculations, it is worth noting that Construction 1 will change the area and

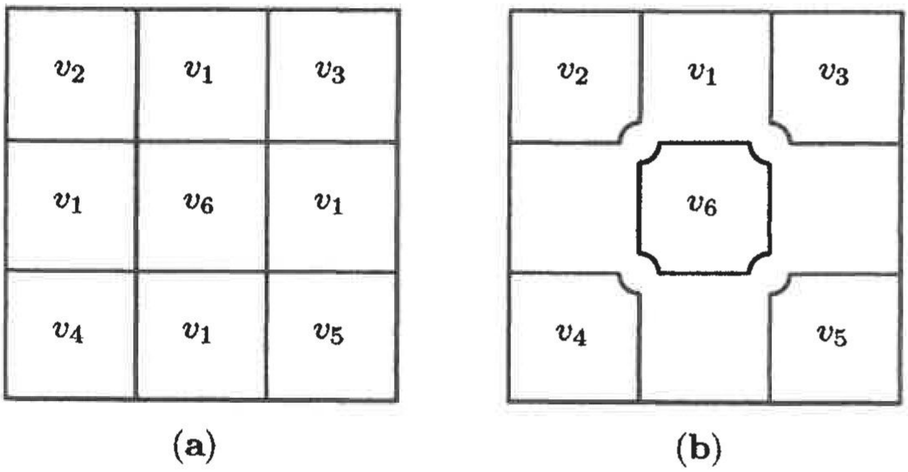

FIGURE 4. Example of a hole being created after multiple applications of Construction 1

perimeter of several units and Construction 2 will change the area of the neighbors of a particular unit. Since the widening of the gap between the two pieces of some unit v is accomplished by using a ball of radius ε in either construction, there is a small change to the area or perimeter of the affected units. Additional concern should be paid to the effect Construction 1 or 2 has on the various population and voting data stored. If none of these values change after using Construction 1 or 2, this would imply that the population is zero in the area that changed from one unit to another unit. This is not an entirely unreasonable assumption, as the amount of area that changes from one unit to another unit is dictated by the ball of radius ε. The advantage that Construction 2 has over Construction 1 is that the area of one unit is preserved. Unfortunately, Construction 2 may be difficult to execute in practice.

There are two additional concerns about the use of Construction 1 or 2 in any map. The first of these occurs when applying Construction 1 or 2 multiple times. If Construction 1 or 2 is used to fix all of the discontiguity problems in unit v_1 in Figure 4, we have created a contiguous region v_1, but, at the same time, we have also created a hole under any simply connected definition. The second concern occurs when using Definition (A3). In Figure 4(a), v_2 is adjacent to v_6 under Definition (A3), but after using Construction 1 or 2, that adjacency is deleted from the augmented dual graph of $\mathcal{U}$. Under certain definitions, edges in the geo-graph can be deleted, which would decrease the number of redistricting plans in the solution space.

A useful consequence of Construction 1 is that it provides the architecture necessary to show that $\mathcal{G}(M)$ is planar under certain definitions in $\mathcal{D}$. Note that the next lemma is only true for certain definitions in $\mathcal{D}$. For this reason we define a subset of $\mathcal{D}$. Let $\mathcal{D}' \subset \mathcal{D}$ be the set containing

$$\{(A1),(B1),(C3)\},\{(A1),(B1),(C2)\},\{(A1),(B1),(C1)\},$$
$$\{(A1),(B2),(C1)\},\{(A2),(B1),(C1)\}.$$

Lemma 6. *Let $\mathcal{U}$ be a set of contiguous units under definition $D \in \mathcal{D}'$. Let G be the augmented dual graph of $\mathcal{U}$. Then G is planar. Furthermore, if $D \in \mathcal{D} \setminus \mathcal{D}'$, then G is non-planar.*

Proof. Since simply connectedness plays no role in whether G is planar, it does not matter if $(C1) \in D$, $(C2) \in D$, or $(C3) \in D$. It is clear that G is planar for $\{(A1),(B1),(C1)\}$, $\{(A1),(B1),(C2)\}$, and $\{(A1),(B1),(C3)\}$ since a dual graph is planar and defined using Definitions (A1) and (B1).

Suppose that $D = \{(A1),(B2),(C1)\}$. For any unit w, if two pieces of w are adjacent at a point x, then there are no other units with two or more pieces adjacent only at x. Thus, for each unit w, we can apply Construction 1 to $\mathcal{U}$ at each point in $Y_w \in Y$ to form $\mathcal{U}'$ using a ball of radius $\varepsilon > 0$. Let G' be the augmented dual graph of $\mathcal{U}'$. Since all units are contiguous and all adjacencies are defined by Definition (A1), G' must be planar for the same reason that a dual graph of planar graph is planar.

Now we need to show that G is a subgraph of G' since this shows G is planar. The vertex sets of G and G' are the same since Construction 1 does not delete any units. It is clear that all edges between vertices in G representing units that did not intersect the ball of radius ε will still be edges in G'. By our choice in definitions, if there is a ball of radius ε at a point x, we can choose ε to be small enough that the only adjacencies lost at x would be between units that were adjacent at x. However, our choice of definition D prevents this, so no adjacencies were lost in the formation of G' . Thus $G = G'$ and so G is planar.

Suppose that $D = \{(A2),(B1),(C1)\}$. By the definition of X in Definition (A2), if a point $x \in X_v$ for some unit v, then for each unit $u \in \mathcal{U} \setminus \{v\}$, $x \notin X_u$. This means that G must be planar as we could place a ball of radius ε at each point in each $X_v \in X$ just as we did in the argument above.

When D is any other set in $\mathcal{D}$, we can form a map whose geo-graph is a K_5, clearly non-planar. Figure 5(a) produces a K_5 for definitions $\{(A3),(B1),(C1)\}$, $\{(A3),(B2),(C1)\}$, and $\{(A3),(B3),(C1)\}$ when zone 1 is $\{v_1, v_2\}$, zone 2 is $\{v_3, v_4\}$, zone 3 is $\{v_5, v_6\}$, zone 4 is $\{v_7\}$, and zone 5 is $\{v_8\}$. Figure 5(b) produces a K_5 for definition $\{(A2),(B2),(C1)\}$ when zone 1 is $\{v_1\}$, zone 2 is $\{v_2, v_3\}$, zone 3 is $\{v_4, v_5\}$, zone 4 is $\{v_6\}$, and zone 5 is $\{v_7\}$. Figure 5(c) produces a K_5 for definitions $\{(A1),(B3),(C1)\}$ and $\{(A2),(B3),(C1)\}$ when zone 1 is $\{v_1\}$, zone 2 is $\{v_2\}$, zone 3 is $\{v_3, v_4\}$, zone 4 is $\{v_5\}$, and zone 5 is $\{v_6\}$. □

3.1.2. *Annexing and Merging.* Since Constructions 3 and 4 are quite similar, we shall discuss them at the same time.

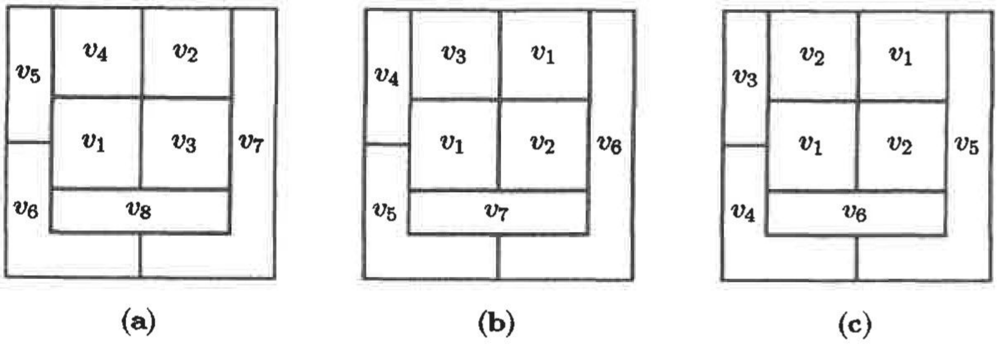

FIGURE 5. Examples that would produce a K_5 in the geo-graph for various definitions in $\mathcal{D}$

A similar procedure to Construction 3 was described in [16], but there were extra conditions that needed to be met. In [16], to annex pieces of a unit, all of the pieces of some unit v must be in holes of the same unit. Additionally, there must be only one piece of v that is not in a hole, and this one piece needs to be adjacent to the same unit as all of the other pieces of v. It is quite easy to fail any one of these requirements as seen in Figure 3(a). The benefits of performing Construction 3 under these restrictions is that no adjacencies are gained or lost from the augmented dual graph of $\mathcal{U}$ to the augmented dual graph of $\mathcal{U}'$. However, ti is possible that Construction 3 may force the removal of some maps from the set of all contiguous maps with k zones if the additional restrictions in [16] are not followed.

Construction 4 has similar concerns as Construction 3. Also, as with Construction 3, the authors of [16] have similar procedures with additional requirements for a process similar to that stated in Construction 4. The authors of [16] intended to merge units only when all pieces of a discontiguous unit are adjacent to the same unit to prevent loss of edges in the geo-graph. Unfortunately, it is possible that even while using the restrictions outlined in [16], some redistricting plans may be removed from the solution set after using Construction 4. For example, if Construction 4 is applied to Figure 6(a) to form Figure 6(b), then we have lost the map where v_1 and v_3 are in one zone while v_2 is in a different zone. There is also the possibility of choice for Construction 4, which adds to the ambiguity of this construction. We can merge as shown in Figure 6(b) or as shown in Figure 6(c). There is no discernible reason why one would choose Figure 6(b) over Figure 6(c) or vice versa without more information. In fact, deciding how to merge may have to be manually executed without creating a random assignment, making this construction difficult to execute in practice. We forgo discussing the impact of the various definitions in $\mathcal{D}$, as there is already too much variability to give an appropriate assessment of the impact from Constructions 3 and 4.

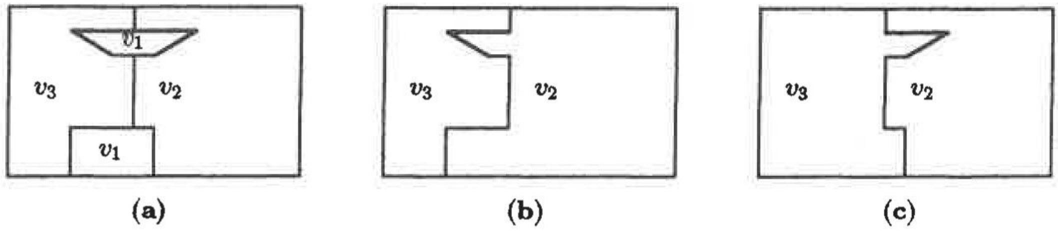

FIGURE 6. Various ways to use Construction 4 on a map: (a) original map, (b) merged v_1 with v_2, (c) merged v_1 with v_3

The movement of the data associated with the units complicates the use of Constructions 3 and 4 quite a bit more, especially the voting and population data. Consider the use of Construction 3 on the map in Figure 3(a) to build the map in Figure 3(d). When two of the pieces of unit v_1 in Figure 3(a) are joined with v_2, there are two natural ways to resolve the voting results and various population data. The first is to do nothing. That is, if at least one piece of the unit v involved in Construction 3 still exists, have v retain all of its population numbers and voting results. As suggested in [16], this assumes that the pieces that were annexed into other units were unpopulated. In most cases, this seems reasonable as most instances of discontiguous units have one large piece with all other pieces being relatively small. But there are still a non-trivial number of units like in Figure 3(d) where each piece has less than half of the area of the whole unit. For example, in the state of Georgia, there are 291,086 census blocks, and of those there are 333 discontiguous census blocks. Of these 333 discontiguous census blocks, there are 49 census blocks where the area of each piece of a unit is less than half of the area of the whole unit.

An alternative to doing nothing with the data would be to move the population and voting results proportionally using some metric. The choice of metric can depend on the availability of data. We propose three metrics on which to base this movement of the various population measures and voting data. For each discontiguous unit v affected by Construction 3, one can split the population and voting results to each piece of v based on the following:

(a) the proportion of area of a piece to the overall area of the unit,
(b) the proportion of registered voters of a piece to the overall number of registered voters in the unit, or
(c) the proportion of voting-age population of a piece to the overall voting-age population for the unit.

There are several advantages and disadvantages to each of these approaches depending on the unit. If a unit is a VTD, then there are no values collected for the metrics (b) or (c), so one would need to aggregate the data from the census blocks to the VTDs. But this will not work for metric (b) as the registered voter data has not been collected at the census block level. In fact, the smallest refinement for registered voter data is census tracts (a much larger unit than census blocks) which are sometimes smaller or larger than VTDs, and typically overlap VTDs in unhelpful ways.

Even though metric (b) would likely be the most accurate, metric (c) is the next best alternative. If the unit is a census block, then metric (c) is not viable, as we would need to know the voting-age population for each piece of a census block, and census blocks are the most refined measurement used by the Census Bureau. Luckily, metric (a) will always work as the shapefile provided by the Census Bureau can be used to find the area of the pieces of census blocks. Unfortunately, metric (a) supposes the population is evenly spread throughout each census block, which may add bias to any study conducted in this area. However, if we use census blocks as units, this is our best alternative without more refined geographic units than census blocks or more accurate data at the census block level.

3.1.3. *Splitting.* The last construction we will examine that will heal contiguity violations is called splitting (see Figure 3(f)). This method has a clear advantage over all other constructions presented thus far, as it can be used on any set of units with discontiguous units. Unlike Construction 4, there is no ambiguity with how to carry out the construction. As mentioned in Section 3.1.2, we would need to either choose to do nothing, or use one of the three metrics ((a), (b), or (c)) to separate the population and voting results data when each discontiguous unit is split.

It is possible that splitting discontiguous units will add more admissible maps, but Construction 5 will never remove maps from the set of all maps on $\mathcal{U}$. It is easy to see that we do not lose any maps in the set of all redistricting maps on $\mathcal{U}$ by splitting a discontiguous unit v since we can choose all pieces of v to be in the same zone. In the same way, by not choosing all pieces of a split unit to be in the same zone, we will produce new maps.

3.2. **Simply Connectedness.** Suppose that all units in $\mathcal{U}$ are contiguous and that the units in $\mathcal{U}$ satisfy Assumption (2.1).

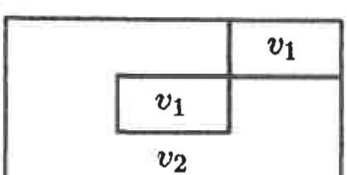

FIGURE 7. Example of a unit that has a degenerate hole

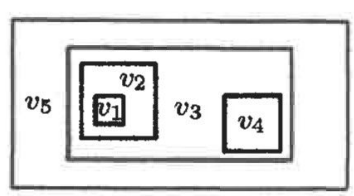

FIGURE 8. Unit that has a hole in a hole

In [16], the authors proposed a method to eliminate holes by merging hole(s) with the units that have holes; the areas, population counts, votes, etc. were combined as in Section 3.1.2. We handle the data associated with each unit by summing together all interval metric data, similar to what we did in Section 3.1.2. In this way, we avoid losing any data and retain all relevant records.

Recall that Assumption (2.1) ensures that if a piece of a unit is in a hole of some unit v, then the entire unit is in holes of v. When writing code to merge the holes with the violating unit, it should be noted that since there is a possibility of having a hole inside of a hole, the unit that is in a hole may change during this process (see Figure 8).

Construction 7 merges holes in a similar manner to that in Construction 3. The main difference is that we have assumed that all units are already contiguous.

Construction 7. Let $\mathcal{U}$ be a set of units using definition $D \in \mathcal{D}$. Let $v, v_1 \in \mathcal{U}$ where v_1 is in a hole of v. Form a new set of units $\mathcal{U}'$ from $\mathcal{U}$ by replacing v_1 and v with the union of v and v_1. This is called *merging holes.* ■

Note that using Construction 7 may not fill a hole in a traditional sense. For example in Figure 8, if we aim to use Construction 7 to fill a hole in v_3 with v_1, we will form a discontiguous unit. However, by applying Construction 7 enough times to "fill all holes," the resulting map will have all units that are contiguous and simply connected.

Lemma 8. *Let $\mathcal{U}$ be a set of contiguous units under definition $D \in \mathcal{D}$ and Assumption (2.1). Let $\mathcal{U}'$ be constructed from $\mathcal{U}$ using multiple applications of Construction 7 until all units are simply connected. Then all units in $\mathcal{U}'$ are also contiguous.*

Proof. Let $\mathcal{U} = \mathcal{U}_1, \mathcal{U}_2, \ldots, \mathcal{U}_t = \mathcal{U}'$ be a sequence of sets of units where $\mathcal{U}_i$ is the set of units formed by applying Construction 7 to $\mathcal{U}_{i-1}$ for each $i \in \{2, \ldots, t\}$. Let $\mathcal{U}_j$ contain a discontiguous unit v. Let $v_1, v_2, \ldots, v_p$ be all of the pieces of v. By definition D and Assumption (2.1), all but one piece of v is surrounded by some other piece of v, say v_1. That is, without loss of generality there exists a simple closed curve in v_1 that encloses each piece in $\{v_2, \ldots, v_p\}$. For each simple closed curve in v_1 that encloses some $v_i \in \{v_2, \ldots, v_p\}$, this simple closed curve also encloses another unit or piece of a unit w that is in a hole of v_1. Since w is in a hole of v, Construction 7 will eventually join w with v, thus either reducing the number of pieces of v or reducing the area represented by holes in v. Either way, when Construction 7 cannot be used anymore, we will have a set of units $\mathcal{U}'$ that are all contiguous. □

A benefit of using Construction 7 is that no redistricting plans were created or destroyed. Recall our discussion in Section 2.1 about how zones are adjacent, contiguous, or simply connected.

Lemma 9. *Let $\mathcal{U}_1$ be a set of contiguous units under definition $D \in \mathcal{D}$ and Assumption (2.1). Let $\mathcal{U}_2$ be the set of units constructed from $\mathcal{U}_1$ using multiple applications of Construction 7 until all units are contiguous. Let Ω_1 be the collection of all maps on $\mathcal{U}_1$ with k contiguous and simply connected zones and Ω_2 be the collection of all maps on $\mathcal{U}_2$ with k contiguous and simply connected zones. Then $\Omega_1 = \Omega_2$.*

Proof. By Lemma 8, the units in $\mathcal{U}_2$ are contiguous and simply connected. Suppose that $\Omega_1 \neq \Omega_2$. Then either there exists a map $M_1 \in \Omega_1$ and $M_1 \notin \Omega_2$ or there exists a map $M_2 \in \Omega_2$ and $M_2 \notin \Omega_1$. Suppose that there exists a map $M_1 \in \Omega_1$ and $M_1 \notin \Omega_2$. Since Construction 7 only merges holes with units that have holes, there must be a zone z_1 in M_1 that contains a unit v that is in a hole of u but u is in a zone $z_2 \neq z_1$. When the two units v, u are merged in M_2, then they will be in the same zone. Since $\mathcal{U}_1$ has only contiguous units, Assumption (2.1) is satisfied, and all zones in M_1 are contiguous, it follows that either contiguity of units is violated, contiguity of zones is violated, or Assumption (2.1) is violated, ultimately leading to a contradiction.

Now suppose that there exists a map $M_2 \in \Omega_2$ and $M_2 \notin \Omega_1$. Since Construction 7 only merges units, $M_2 \in \Omega_1$ and thus we have a contradiction. □

3.3. **Corrections on the Initial Redistricting.** Weighing all of the advantages and disadvantages of the various constructions, to ensure we get a map where all units and zones are contiguous and simply connected, we use Constructions 5 and 7. The greatest advantage these two constructions possess is that they can always be applied regardless of the map. Additionally, as we show in the next corollary, we do not lose any feasible solutions by filling holes and splitting pieces. Corollary 10 follows directly from Lemma 8 and the discussion provided in Section 3.1.3.

Corollary 10. *Let $\mathcal{U}_1$ be a set of units under definition $D \in \mathcal{D}$ and Assumption (2.1). Let $\mathcal{U}_2$ be a set of contiguous and simply connected units formed from $\mathcal{U}_1$ using Construction 5 and 7 multiple times. Let Ω_1 be the collection of all maps on $\mathcal{U}_1$ with k contiguous and simply connected zones and Ω_2 be the collection of all maps on $\mathcal{U}_2$ with k contiguous and simply connected zones. Then $\Omega_1 \subseteq \Omega_2$.*

4. Algorithms

4.1. **Articulation Points.** Let M be a map on the units in $\mathcal{U}$. From the perspective of graph theory, ensuring that local search and MCMC algorithms will not form a discontiguous zone when a single unit v is removed from a zone j is equivalent to identifying whether vertex v is an articulation point in the induced subgraph of $\mathcal{G}(M)$ on $\mathcal{U}_M(j)$ (set of units in zone j on map M). This was mentioned in both [16, 21], but the authors were using definition $D = \{(A1), (B1), (C1)\}$. We can also take what they said one step further to identify holes.

Assumption (2.1) allows us to use the geo-graph to find a hole using the articulation points. In terms of the geo-graph, Assumption (2.1) ensures that if u is a hole in v, then all paths (walks on a graph that do not repeat any vertices) with u as an endpoint must contain v if the other endpoint of the path is any vertex representing a unit in $\mathcal{U} \setminus \{v\}$ that is not a hole in v. In other words, v represents an articulation point (a cut-vertex) in the geo-graph. Since all definitions in $\mathcal{D}$ satisfy Assumption (2.1), a unit u that is a hole in some unit v is adjacent only to v itself or other units that are holes in v. Thus, if v is an articulation point and not the dummy unit (unit representing the unbounded area), then all disconnected components in the graph $\mathcal{G}(M) - v$ (remove vertex v and all incident edges) not containing the dummy unit represent holes in v. So as long as we can identify all articulation points, we can identify all units or zones that have holes. As stated in [10, 18], the set of units with holes (articulation points) can be identified in $\mathcal{O}(|\mathcal{U}|)$ time for a geo-graph. Further, it is said that holes in any unit of $\mathcal{U}$ can be identified while finding the articulation points. Thus, all units with holes and in holes can be identified in $\mathcal{O}(|\mathcal{U}|)$ time. If $\mathcal{G}(M)$ is non-planar, then by the work in [10, 18], we can find all of the articulation points in $\mathcal{O}(|\mathcal{U}| + |E(\mathcal{G}(M))|)$ time where $E(\mathcal{G}(M))$ is the set of edges in $\mathcal{G}(M)$. Hence, as long as Assumption (2.1) is satisfied, we can can find all units that have holes and all units in holes in $\mathcal{G}(M)$ in $\mathcal{O}(|\mathcal{U}| + |E(\mathcal{G}(M))|)$ time under any definition in $\mathcal{D}$.

We can also perform a similar action to find all zones with holes and zones in holes. We first need to build the graph $G_z(M)$, the *zone dual graph*. Let $G_z(M)$ be the dual graph of the zones in M where each vertex is represented by a zone in M and two vertices z_1, z_2 are adjacent if there is a unit in z_1 adjacent to a unit in z_2. Recall that due to the adjacency rule in D, we can attribute the same adjacency definitions to the zones as we do for units. Similarly, we can attribute the contiguity and simply connectedness definitions in D to each zone as well. Thus, by our above discussion on identifying holes in units, if we are using definition $D \in \mathcal{D}'$, then we can find all articulation points in $G_z(M)$ in $\mathcal{O}(|V(G_z(M))|) = \mathcal{O}(k)$ time, and we can also find all zones with holes and zones in holes in $\mathcal{O}(k)$ time with the Hopcroft-Tarjan algorithm in [10, 18]. If $G_z(M)$ is non-planar, then we can find all zones with holes and all zones in holes in $\mathcal{O}(k + |E(G_z(M))|)$ times for the same reason as explained in the previous paragraph.

A similar argument can be made for assessing contiguity since articulation points signify when a zone can become discontiguous after using Construction 11. Unfortunately, it will take $\mathcal{O}(|\mathcal{U}_M(z(v))|)$ time and $\mathcal{U}_M(z(v))$ could be considerably large as compared to the number of zones depending on the number of units in a zone. In the context of congressional districts and census blocks, the average number of units per zone will be large, meaning $|\mathcal{U}_M(z(v))|$ will be large in comparison to the number of zones in this context.

4.2. **Main Algorithm.** Before we can talk about the main algorithm, we need to make some assumptions about the state of the map, which are the same assumptions made by most scholars when performing local search heuristics or MCMC algorithms [2, 5, 9, 11, 13, 15, 16, 17]. All algorithms described in this section assume that all units and zones are initially contiguous and simply connected.

The following construction is applied to a geo-graph as part of a local search heuristic (such as steepest descent [16]) or an MCMC algorithm (see [2, 5, 11, 13]) to move one unit locally from one zone into an adjacent zone.

Construction 11. Let M be a map on a set of units $\mathcal{U}$ where each unit in $\mathcal{U}$ and zone in M is contiguous and simply connected. Let X be the set of pairs (v, z') where v is a unit on the boundary of at least one zone adjacent to another zone (not the dummy zone) and z' be one zone (not the dummy zone) to which v shares a boundary. Then take one pair (v, z') uniformly at random from X. Form a new map M' where $z_M(u) = z_{M'}(u)$ for all $u \in \mathcal{U} \setminus \{v\}$ and $z_{M'}(v) = z'$. ■

As this construction is stated, it is possible that applying Construction 11 could make the zones in M' discontiguous and/or not simply connected, even though M contains only contiguous and simply connected zones. We utilize the algorithms below to check that all zones are contiguous and simply connected after applying Construction 11.

Let $D \in \mathcal{D}$. Let M be a map as defined in Construction 11. Then, by Lemma 6, $\mathcal{G}(M)$ is planar. The beginning of our main algorithm depends on another algorithm for finding planar embeddings. Boyer and Myrvold [4] showed the following result regarding planar embeddings.

Theorem 12. [4] *Given a planar graph G with n vertices, algorithm* PLANARITY *in* [4] *produces a planar embedding of G in $\mathcal{O}(n)$ time.*

Since $\mathcal{G}(M)$ is planar, we can find a planar embedding of $\mathcal{G}(M)$ in $\mathcal{O}(|\mathcal{U}|)$ time.

Since $\mathcal{G}(M)$ contains information about the boundary of v, we can form a clockwise ordered list $R(v)$ of the queen-adjacencies of v using the planar