

A geometric framework for resource allocation problems

John Gunnar Carlsson and Yinyu Ye
Stanford University

Sharon Arroyo, Roman Fresnedo, and Raju Mattikalli
Boeing Company

Abstract: In a wide variety of industrial problems, a geographic region must be segmented into smaller regions to allocate resources. This is often done using combinatorial devices that discard the underlying geometry of the problem. In this paper we describe a geometric framework, *map segmentation*, for solving such problems and give several practical examples demonstrating the efficacy of such approaches.

1. Introduction: In the *map segmentation problem*, we have a planar geographic region C that we must partition into a collection of k smaller sub-regions C_i , while optimizing some criteria and satisfying a regularity condition. “Regular” might mean convex, simply connected, connected, or merely measurable. The criteria to optimize might be the area of a sub-region, the number of elements of a given point set \mathcal{X} in a sub-region, or the integral $\int_{C_i} f(x, y) dA$ of a specified density $f(\cdot)$. Figure 1 gives an abstract instance of such a problem.


In a variety of practical settings, map segmentation problems can be used to allocate tasks to vehicles or other agents in surveying or providing service to a region. In this paper we demonstrate several problems for which map segmentation can serve as a useful tool.

2. Background: A number of geometric results form the theoretical framework for our problems, motivated for the most part by the *ham sandwich theorem* [12]:


Theorem 1. *The volumes of any n solids of dimension n can always be bisected by a $(n - 1)$ -dimensional hyperplane.*

Corollary 2. *Any two planar regions can be simultaneously bisected by a single line.*

Analogous results hold for probability densities instead of solids. Stojmenović [13] gives a linear-time algorithm that finds the bisecting line of corollary 2 when the two regions are disjoint convex polygons. More recently, Bepamyatnikh et al. [3] and Ito et al. [8] have proven the following discrete result:



(a)



(b)

Figure 1: An example of a map segmentation problem; in (a) we have a probability density defined on a region, and in (b) we partition the region into $k = 8$ convex sub-regions, each having the same area and containing the same amount of the density.

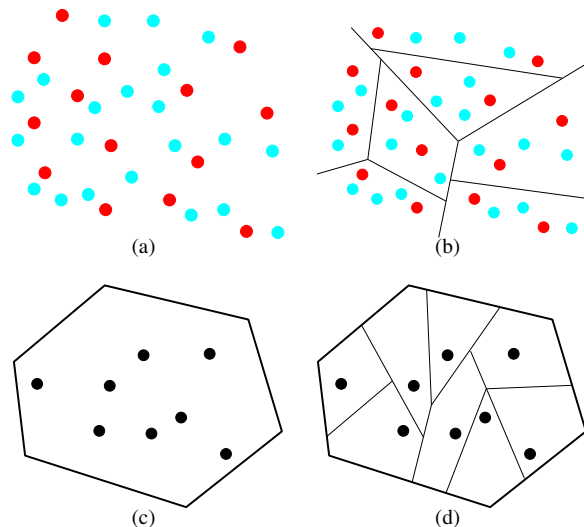


Figure 2: (a) and (b) show an equitable partition of a red and blue point set; each convex cell contains 3 blue points and 2 red points. (c) and (d) show an equitable partition of a point set and a polygon, so that each convex cell has the same area and contains exactly one point.

Theorem 3. *Given gn red points and gm blue points in the plane in general position, there exists a subdivision of the plane into g disjoint convex polygons, each of which contains n red points and m blue points.*

In addition, [3] gives a polynomial-time algorithm for finding such a partition. Carlsson et al. [4] prove the following result, motivated by a vehicle routing heuristic:

Theorem 4. *Given a convex polygon C with m vertices and area 1 containing a point set \mathcal{X} with $|\mathcal{X}| = n$, there exists a partition of C into n disjoint convex cells, each containing 1 point and having area $1/n$. Furthermore, one can find such a partition in running time $\mathcal{O}(Nn \log N)$, where $N = m + n$.*

This algorithm is readily extended to an arbitrary probability density on C , so that we obtain a collection of convex sub-regions each having the same probability mass.

Figure 2 gives some examples of these problems. The algorithms described above provide us with polynomial time approximation schemes for the general problems described in the following section.

3. Applications:

3.1. Street Scanning: The *Chinese Postman Problem* (CPP) is a well-studied optimization problem first described in [9]. The objective of CPP is to find a tour of a weighted graph $G = (V, E)$ of minimal total length that traverses each edge or arc at least once. As the name suggests, CPP is a natural situation encountered

by delivery services. Recently, large-scale CPP and variants have become of significant practical importance for internet-based map companies in scanning roads to generate street-level views of a location in an urban environment [7].

In such applications, the goal is to traverse every street in a city with a survey vehicle, which typically has a camera attached. However, merely solving CPP gives only a route for a single vehicle, so we do not have any idea of how to distribute the workload between a fleet of vehicles. Hence, treating the problem as an instance of CPP is a poor solution technique. Although approximation algorithms for the “ k -postman problem” exist (see [1]), they are strictly combinatorial and do not take advantage of the fact that our road map is a planar graph, and consequently vehicle tours may not be geographically separate. In a practical setting it is desirable to clearly separate one vehicle’s route from another in an obvious geographic way.

Suppose that we have k vehicles, located at different given starting points (x_i, y_i) on a map of a (convex) city. We can generate vehicle tours using map segmentation by partitioning the map into k convex sub-regions, each containing one point (x_i, y_i) , and each containing the same amount of total road length.

We then solve an instance of CPP at each sub-region while using combinatorial tools to minimize the amount of left turns and U-turns that each vehicle takes. In [6] we solve a problem instance provided by an industrial affiliate with 900,000 nodes while introducing only 4% overhead cost and maintaining a load-balancing ratio of 1.34, so that the longest vehicle tour is at most 1.34 times the shortest vehicle tour. Figure 3 shows the input and outputs of this algorithm.

3.2. The Min-Max Multiple Depot Vehicle Routing Problem: The *Vehicle Routing Problem* (VRP) has been a key problem in discrete optimization for almost 50 years. The problem is to route a set of service vehicles to visit all clients in a geographical region while minimizing cost. In the *multi-depot vehicle routing problem* (MDVRP) variant, multiple vehicles start from multiple depots and return to their original depots at the end of their assigned tours. We have studied the min-max MDVRP where the objective is to minimize the longest travel time of all the vehicles. This objective is rarely studied in the literature, although it is a common objective in CPU and job scheduling, where the time to completion of all jobs is called the “makespan”. This objective has substantial practical use, for example to minimize the amount of overtime for the drivers of the vehicles. This objective also balances the loads of all vehicles more than the traditional objective, which minimizes the total tour length of all vehicles.

A solution to the min-max MDVRP can be found in

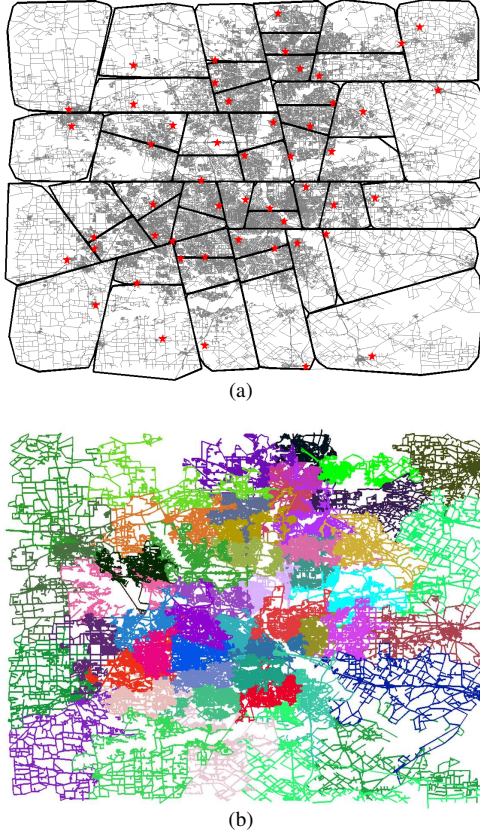


Figure 3: In (a) we segment the region into convex sub-regions, each containing the same amount of road length and containing one vehicle's starting point (indicated by the red stars). These sub-regions are part of the input to a local search algorithm that generates CPP tours (b) that balance the loads between vehicles and minimize the number of left turns and U-turns.

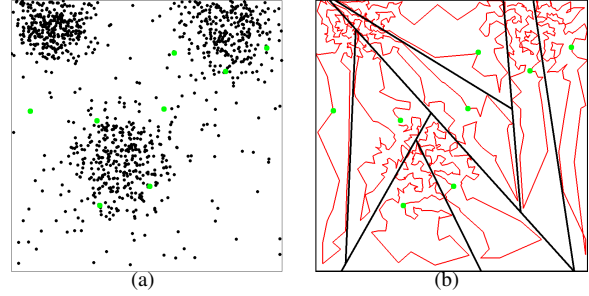


Figure 4: Segmenting a map to generate vehicle tours. Each sub-region contains exactly one depot, indicated by the green points, and is guaranteed by theorem 6 to asymptotically induce equal tour lengths.

two steps. In the first step we assign the clients to individual vehicles located at their respective depots and in the second we solve a *traveling salesman problem* (TSP) for each vehicle in order to route it from its depot to the clients assigned to it and back to the depot with minimal travel time. Our heuristic for min-max MDVRP first solves a map segmentation problem where we divide the service region into n pieces, one for each vehicle, such that each piece contains the depot for that vehicle, is convex, and contains an equal area. The heuristic then solves a TSP for each vehicle to determine a route starting at the vehicle's depot, visiting all the clients located in that piece of the service region before returning to the depot. This heuristic is shown to be asymptotically optimal in [5] when the clients are uniformly distributed by virtue of the *BHH theorem* [2]:

Theorem 5. Suppose that (X_1, \dots, X_k) are random points i.i.d. uniformly in a compact region R . Then the length TSP (X_1, \dots, X_n) of the optimal travelling salesman tour traversing points X_1, \dots, X_k satisfies

$$k^{-1/2} \text{TSP}(X_1, \dots, X_k) \rightarrow \alpha \sqrt{\text{Area}(R)}$$

We therefore find that sub-regions with equal area have the same load, asymptotically speaking, up to $o(\sqrt{k})$. A similar result, also asymptotically optimal, exists when demand is non-uniformly distributed:

Theorem 6. Suppose that (X_1, \dots, X_k) are random points i.i.d. according to a probability density function $f(\cdot)$ supported on a compact region R . Then the length TSP (X_1, \dots, X_n) of the optimal travelling salesman tour traversing points X_1, \dots, X_k satisfies

$$k^{-1/2} \text{TSP}(X_1, \dots, X_k) \rightarrow \alpha \int_R \sqrt{f(x,y)} dA \quad (1)$$

A problem instance is illustrated in figure 4.

The requirement that each sub-region be convex might have some further benefits. We expect that odd-shaped regions take longer to service than more compact




Figure 5: Based on voting data for the state of Pennsylvania from the 2008 election, the districts shown above contain the same number of Republicans and Democrats (and hence the same total population).

regions with the same area. Convex regions also ensure that a vehicle's traveling-salesman tour remains in its service region and may be a robust way to handle some small uncertainty in the customers' exact locations.

3.3. Voter redistricting: In the USA, a state's congressional and legislative districts are redrawn every ten years following a census. A Supreme Court ruling requires that the districts in each state be contiguous and contain an equal number of voters. Partitioning a state equitably with respect to these constraints is straightforward, provided the geometry of the state to be partitioned is not extremely abnormal. The single constraint of containing equal numbers of voters leaves much flexibility; for example, suppose we desired convex districts (for the appearance of simplicity) and that the residences of n favorite politicians be in separate districts (a congressman is required to reside in the district he represents). Figure 5 shows a map of the state of Pennsylvania divided into smaller districts while maintaining equal numbers of Republicans and Democrats in each. The algorithm of [4] describes an approximation scheme for solving such a problem.

A more effective "gerrymandering" scheme to enhance the control of the majority party would seek to spread its advantage equally among the districts in the hope of winning all of them. In recent years there has been a push for fair redistricting by independent panels. These might have one criteria to ensure equal area and another to select for simple district shapes (e.g., the ratio of the perimeter to the square root of the area). The distributed algorithm of [10] describes a means for partitioning a region equitably with respect to a density using *power diagrams*, with the intention of producing simpler shapes than [4].

3.4. Natural resource allocation: Suppose we have a geographical region known to contain some natural resource, such as water, petroleum, or soil. We assume that this resource is distributed in the region according to a known probability density function. A naturally decentralized problem we may encounter is to divide the geographical region into smaller regions while preserving an equitable distribution of the resource, subject to additional constraints.

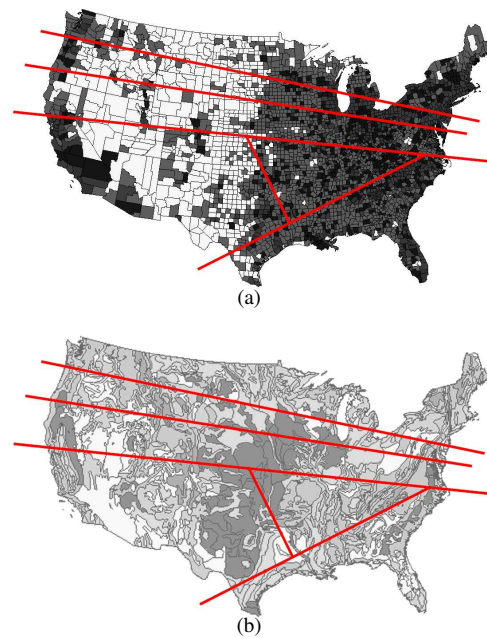


Figure 6: In (a), a map of the population density of the USA, and in (b), a map of the soil quality. The same partition, shown in red, equally divides both population and soil. (Pictures taken from Wisconsin Business School http://www.bus.wisc.edu/realestate/images/resources/us_density.gif and US Geological Survey http://biology.usgs.gov/luhna/images/fig3_2.gif)

For example, Figure 6 displays an equitable partition of the United States with respect to population density and soil quality; each sub-region has the same total population and the same average soil quality. Farmers could theoretically use such a partition to best determine the most useful regions to farm. In this case, we impose the constraint that the partition divide the plane \mathbb{R}^2 into a collection of convex cells; while this does not guarantee that *sub-regions* themselves will be convex, it clearly generates regions whose boundaries are somewhat regular.

4. Further Work:

4.1. Theoretical work: In the examples given above, the criteria that we use to define the sub-regions is always *additive*; for example, for any two disjoint regions C_1 and C_2 , we have $\text{Area}(C_1 \cup C_2) = \text{Area}(C_1) + \text{Area}(C_2)$ and $|(C_1 \cup C_2) \cap \mathcal{X}| = |C_1 \cap \mathcal{X}| + |C_2 \cap \mathcal{X}|$, where \mathcal{X} denotes a specified point set. Other criteria of practical interest which are not additive include the diameter or perimeter of a sub-region. We hope to obtain some bounds that can be placed on problems in similar settings, two of which are described below.

4.2. Further applications:

4.2.1. Organ transplant regions: The USA is currently divided into 11 regions that govern the assignment of donor organs to recipients, based on travel time and proximity to organ transplant facilities. [11] gives an integer program to maximize a bicriteria objective function combining the number of intra-regional transplants and intra-regional transplant rate. We believe that the partitioning algorithm of [4] can be used to model this problem using a continuous formulation, equitably partitioning some measure (population density or average demand for donor organs, for example) while generating regions containing specified organ transplant facilities, or the regions for organ transplantation in which one looks first for possible matches.

4.2.2. Online vehicle routing: The *online vehicle routing problem* (OLVRP) is a variant of the stochastic vehicle routing problem in which new demand points arrive while vehicles are completing their tours. The result (1) is no longer guaranteed to hold, and we are not aware of a comparable law of large numbers for this situation. If a similar result does exist, we can obtain a new partitioning criterion specifically for OLVRP.

One might also analyze the long-term performance of specific routing policies; for example, if all vehicles use a nearest-neighbor rule to compute their routes, the system becomes an $M/G/k$ queue (provided arrivals follow a Poisson process). If we partition the service region, the resulting system is a collection of independent $M/G/1$ queues in which the service times are dependent on the shape of the sub-regions. Our goal would be to find a partitioning criterion that gives us useful information about

this distribution of service times.

5. Acknowledgements: The authors gratefully acknowledge the support of NSF award #0800151. In addition, Y.Y. and J.G.C. gratefully acknowledge the support of the Boeing Company.

6. References:

- [1] D. Ahr and G. Reinelt, "New heuristics and lower bounds for the Min-Max k-Chinese Postman Problem," Proceedings of the 10th Annual European Symposium of Algorithms (ESA), September 2002.
- [2] J. Beardwood, J. H. Halton, and J. M. Hammersley, "The shortest path through many points," Proc. Cambridge Philos. Soc 55:299-327, 1959.
- [3] S. Bessamyatnikh, D. Kirkpatrick, and J. Snoeyink, "Generalizing ham sandwich cuts to equitable subdivisions," Discrete Comput. Geom., 24:605-622, 2000.
- [4] J. G. Carlsson, B. Armbruster, and Y. Ye, "Finding equitable convex partitions of points in a polygon efficiently," ACM Transactions on Algorithms. To appear.
- [5] J. G. Carlsson, D. Ge, A. Subramaniam, and Y. Ye, "Solving the Min-Max Multi-Depot Vehicle Routing Problem," Proceedings of the FIELDS Workshop on Global Optimization, April 2007.
- [6] J. G. Carlsson and Y. Ye, "Practical Distributed vehicle routing for street-level map scanning."
- [7] M. Helft, "Google zooms in too close for some," *New York Times*, June 1 2007.
- [8] H. Ito, H. Uehara, and M. Yokoyama, "2-dimension ham sandwich theorem for partitioning into three convex pieces," In Proc. Japan. Conf. Discrete Comput. Geom. (JCDCG'98), volume 1763 of Lecture Notes in Computer Science, pages 129-157. Springer-Verlag, 1998.
- [9] M. K. Kwan, "Graphic programming using odd or even points," Chinese Math., 1:273-277, 1962.
- [10] M. Pavone, A. Arsie, E. Frazzoli, and F. Bullo, "Equitable partitioning policies for mobile robotic networks," Under review. See <http://mit.edu/pavone/www/papers/Pavone.Arsie.ea.TAC08.pdf>.

- [11] J. E. Stahl, N. Kong, A. Schaefer, and M. Roberts, "A methodological framework for optimally reorganizing liver transplant regions," *Medical Decision Making*, 25(35), 2005.
- [12] H. Steinhaus, "A note on the ham sandwich theorem," *Mathesis Polska* 9, 26–28, 1938.
- [13] I. Stojmenović, "Bisections and ham-sandwich cuts of convex polygons and polyhedra," *Information Processing Letters* 38.1:15-21, 1991.