

Assessing a Metaheuristic for Large Scale Commercial Districting

Roger Z. Ríos-Mercado

Graduate Program in Systems Engineering

Universidad Autónoma de Nuevo León (UANL), Mexico

roger.rios@uanl.edu.mx

August 2015

Accepted April 2016

Abstract

In this paper, we present an empirical evaluation of a metaheuristic approach to a commercial districting problem. The problem consists of partitioning a given set of basic units into p districts so as minimize a measure of territory dispersion. Additional constraints include territory connectivity and balancing with respect to several criteria. To obtain feasible solutions to this NP-hard problem, a reactive greedy randomized adaptive search metaheuristic procedure is used. Previous work addressed medium-scale instances. In this paper, we report our computational experience when addressing larger instances resembling more closely the size of real-world instances. The empirical work includes full assessment of the algorithmic parameters and its local search phase, and a sensitivity analysis of the balance tolerance parameter in terms of solution quality and feasibility. The empirical evidence shows the effectiveness of the proposed approach, and how this approach is significantly better than the method used by the industrial partner. The complexity of the planning constraints make the current practice method struggling on getting feasible designs. Even for the larger cases, the proposed procedure successfully solved instances with balance tolerance parameter values of as low as 3%, something impossible to achieve by the company current standars.

Keywords: Combinatorial optimization; Territory design; Multiple balancing requirements; Metaheuristics; Reactive GRASP.

1 Introduction

In this paper, a commercial territory design problem (TDP) arising from a beverage distribution firm is addressed. A territory design problem consists of grouping small geographic or basic units (BUs) into larger geographic clusters, called territories, in a way that the territories are acceptable (or optimal) according to relevant planning requirements. This problem belongs to the family of districting problems that have a broad range of applications such as political districting and the design of sales and services territories. For survey papers in districting, or some its important applications such as political and sales districting, the reader is referred to the works by Kalcsics, Nickel, and Schröder [19] , Zoltners and Sinha [35] , Duque, Ramos, and Suriñach [9] , and Ricca, Scozzari, and Simeone [24].

Given a set of city blocks, where three different activities are present in each block (number of customers, product demand, and workload), the firm wants to partition the area of the city into disjoint territories according to several criteria such as:

- *Balanced territories*: Territories must be balanced, i.e, similar in size, with respect to each of the three node activity measures.
- *Contiguity*: For each formed territory, BUs can reach each other by traveling within the territory.
- *Compactness*: BUs assigned to a territory are relatively close to each other.
- *Number of territories*: A fixed number of territories must be sought.

The problem addressed in this paper, motivated by a real-world application, was introduced by Ríos-Mercado and Fernández [26]. They modeled the problem as a p -center problem (with additional side constraints) where the focus was on locating p centers, one for each territory. It is clear that for modeling this problem in principle it is not needed to associate a center with each territory. However, this provides a simple tool for defining a compactness measure and for formulating the contiguity requirements. The specification that the territories be simultaneously balanced with respect to all the three measures has been modeled by requiring that each territory be within a threshold of a target value for each activity measure. This is motivated by the fact that for a given instance a solution where all the territories are simultaneously balanced with respect to all three measures may not exist. In that work, the authors proposed and developed a reactive GRASP algorithm with excellent results when compared to current industry practice in medium-size instances, that is, instances with about 500 basic units.

The goal of this paper is to extend that work by reporting further results following some of the directions for future work pointed out by the authors. In particular, it is of interest to manipulate larger instances, so computational work with 1000- and 2000-basic areas instances is presented. In

addition, the greedy function of the GRASP construction phase is a weighted combination of the original objective function and the violation of the balancing constraints. In that work, the weight parameter (named λ) was fixed, so here a study of the sensitivity of the solutions when variation on this parameter is permitted is presented. Given the algorithm is very successful in finding feasible solutions (i.e., feasible with respect to the balancing constraints) using a tolerance level of 0.05 (referred to as τ), the algorithm is now applied to more tightly constrained instances. Finally a comparison with current industry standard is performed.

The obtained results are very satisfactory. The local search has a very good performance as it was able to improve the phase 1 solutions by over 55% on average. Furthermore, the algorithm consistently found feasible solutions to all instances, which is something very hard to achieve by current firm methods in practice. In addition, evidence showed that larger values of the weight parameter (that is, more weight given to the dispersion function) deliver better solutions. In addition, it was also observed that the algorithm was able to find feasible solutions for tighter values of the tolerance parameter. Finally, the empirical work showed how the Reactive GRASP consistently outperformed the method previously used by the firm.

The paper is structured as follows. A description of the problem is presented in Section 2. An overview of the most significant work on territory design is given in Section 3. Section 4 describes the Reactive GRASP. The empirical work is presented in Section 5. The paper is wrapped-up with some final remarks in Section 6.

2 Problem Description

The problem is modeled by a graph $G = (V, E)$, where a city block or BU i is associated with a node, and an arc connecting nodes i and j exists in E if blocks i and j are adjacent to each other. Now each node $i \in V$ has several associated parameters such as geographical coordinates (c_i^x, c_i^y) , and three measurable activities. Let w_i^a be the value of activity a at node i , where $a = 1$ (number of customers), $a = 2$ (product demand), and $a = 3$ (workload). A territory is a subset of nodes $V_k \subset V$. The number of territories is given by the parameter p , so the set of territories is given by $K = \{1, \dots, p\}$. It is required that each node is assigned to only one territory. Thus, the territories define a partition of V . One of the properties sought in a solution is that the territories are balanced with respect to each of the activity measures. So, let us define the size of territory V_k with respect to activity a as: $w^a(V_k) = \sum_{i \in V_k} w_i^a$, $a \in A = \{1, 2, 3\}$. Due to the discrete structure of the problem and to the unique assignment constraint, it is practically impossible to have perfectly balanced territories with respect to each activity measure. To account for this, we measure the balance degree by computing the relative deviation of each territory from its average size μ^a , given by $\mu^a = w^a(V)/p$, $a = 1, 2, 3$. Another important feature is that all of the nodes assigned to each territory are connected by a path contained totally within the territory. In other

words, each of the territories V_k must induce a connected subgraph of G . In addition, industry demands that in each of the territories, blocks must be relatively close to each other. One way to achieve this is for each territory to select an appropriate node to be its center, and then to define a distance measure such as $D = \max_{k \in K} \max_{j \in V_k} d_{c(k),j}$, where $c(k)$ denotes the index of the center of territory k so $d_{c(k),j}$ represents the Euclidean distance from node j to the center of territory k . So maximizing compactness is equivalent to minimizing this dispersion function D . All parameters are assumed to be known with certainty.

The combinatorial optimization problem can be described as finding a p -partition $X = (V_1, \dots, V_p)$ that minimizes the distance-based dispersion measure given above.

$$\text{Minimize} \quad f(X) = \max_{k \in K, j \in V_k} \{d_{c(k),j}\} \quad (1)$$

$$\text{subject to:} \quad \bigcup_{k \in K} V_k = V \quad (2)$$

$$V_{k_1} \cap V_{k_2} = \emptyset \quad k_1, k_2 \in K \quad (3)$$

$$\frac{|w^a(V_k) - \mu^a|}{\mu^a} \leq \tau^a \quad k \in K, a \in A \quad (4)$$

$$G_k = (V_k, E(V_k)) \text{ is connected} \quad k \in K \quad (5)$$

Objective (1) measures territory dispersion. Constraints (2)-(3) define the p -partition. Constraints (4) represent the territory balance with respect to each activity measure as it establishes that the size of each territory must lie within a range (measured by tolerance parameter τ^a) around its average size. Constraints (5) guarantee the connectivity of the territories. These constraints are similar to the constraints used in routing problems to guarantee the connectivity of the routes. Note that, as usual, there is an exponential number of such constraints. This particular TDP is \mathcal{NP} -hard [26]. A Mixed-Integer Linear Programming (MILP) formulation can be found in [26].

3 Related Work

Territory design or districting research includes work in many different areas including design of sales territories [15, 12, 8, 19], design of political districts [16, 14, 17, 21, 3, 1], turfing in telecommunications [33], police districting [6], districting for salt spreading operations [22], home-care districting [2], school districting [5], and recollection of waste equipment [11], to name a few. For survey papers in districting, or some its important applications such as political and sales districting, the reader is referred to the works by Kalcsics et al. [19], Zoltners and Sinha [35], Duque et al. [9], Ricca et al. [24], and Kalcsics [18]. Most of these applications are basically node-

based partitioning problems. A discussion of literature of edge/arc based districting problems can be found in García-Ayala et al. [13]. In this section we focus on reviewing the most relevant work in node-based commercial territory design.

Vargas-Suárez et al. [34] addressed a related commercial TDP with a variable number of territories p , using as an objective a weighted function of the activity deviations from a given goal. No compactness was considered. A basic GRASP was developed and tested in a few instances obtaining relatively good results. The main limitation was that no geographical information was considered, therefore resulting territories were rather disperse.

Later, Ríos-Mercado and Fernández [26] introduced a model considering geographical information aiming at finding compact territories. In that problem a dispersion function based on Euclidean distances between units was considered as a measure of dispersity. In addition, balancing with respect to three activity measures, and territory contiguity was included. In that work, a Reactive GRASP was proposed and tested in a variety of problem instances. Particularly, the value of the reactivity, that is, the self-adjustment of the GRASP quality parameter (α) was empirically proved. The experimental set consisted of data sets of size 500 nodes, and tolerance levels of 30, 20, 10, and 5%. The present is a follow-up of that work.

Caballero-Hernández et al. [4] extended the work in [26] by considering a problem with additional joint assignment constraints, i.e, when some units are required to belong to the same territory. Given the nature of this work, the approach developed in [26] no longer applied, thus the authors proposed a pre-processing phase based on the k -Shortest Path Problem, which finds pieces of territories that satisfy the joint assignment constraints. This is followed by a GRASP-based phase aiming at merging the isolated components until p territories are formed. The reported results were relatively good, and significantly better than those reported by industry.

Ríos-Mercado and Salazar-Acosta [28] studied a commercial TDP with both design and routing decisions simultaneously. They developed a three-phase heuristic consisting of a construction phase and a local search phase for obtaining feasible designs, and then a routing cost computation phase where optimal TSPs were obtained for each formed territory by branch and cut. These phases were iteratively executed until a stopping criteria was satisfied. The results indicated that their method was able to find good designs with relatively low routing costs.

Ríos-Mercado and López-Pérez [27] considered a different commercial TDP. First, as a measure of dispersion, they used a p -median objective function (rather than the p -center objective). They introduced additional requirements such as joint and disjoint assignment requirements, and similarity with existing plan. The problem was in fact posed as a re-districting problem, i.e., find a new districting plan under the new requirements but not too different from the existing design. They proposed a heuristic approach based on a surrogate mixed-integer linear programming (MILP) model, with good results. The practical success of this approach is further documented in [20].

More recently, Ríos-Mercado and Escalante [25] studied the commercial districting problem

but using a diameter-based objective function as dispersion measure. They proposed a GRASP that incorporates a novel construction procedure where territories were formed simultaneously in two main stages using different criteria. This differed from previous literature where GRASP was used to build one territory at a time. The GRASP was further enhanced with two variants of forward-backward path relinking, namely static and dynamic. Path relinking is a sophisticated and very successful search mechanism. Experimental results revealed that the construction mechanism produced feasible solutions of acceptable quality, which were improved by an effective local search procedure. In addition, empirical evidence indicated that the two path relinking strategies had a significant impact on solution quality when incorporated within the GRASP framework.

From the exact optimization perspective Salazar-Aguilar et al. [29] presented a computational study of existing and new MILP formulations for two versions of the commercial TDP. One using the p -center problem measure of dispersion as its objective function and the other using the p -median problem objective. They developed an exact optimization algorithm based on an iterative relaxation of the connectivity constraints. Being an exact approach, they based their study in relatively small and medium size instances.

There are also multi-objective optimization approaches to commercial TDP, where two or more functions are considered as optimization objectives. Salazar-Aguilar et al. [30, 32, 31] considered a family of problems where both territory dispersion and balancing criteria are optimized. In [30], they addressed small size instances from an exact optimization perspective by means of an ϵ -constraint method. In [32, 31], they developed heuristic methods for large scale instances combining GRASP and scatter search approaches.

In terms of developing lower bounds for commercial territorial design problems, the only work we are aware of is that of Elizondo-Amaya et al. [30] who derived a lower bounding scheme for commercial territory design under a p -center based function dispersion objective subject to multiple balance constraints with no connectivity constraints. Lower bounds were obtained using a binary search over a range of coverage distances. For each coverage distance a Lagrangian relaxation of a maximal covering model was used effectively. Empirical evidence showed that the bounding scheme provided tighter lower bounds than those obtained by the linear programming relaxation.

4 The Reactive GRASP for TDP

GRASP [10], a fairly well-known metaheuristic that captures good features of both pure greedy algorithms and random construction procedures, has been widely used for successfully solving many combinatorial optimization problems. In our work, we make use of the Reactive GRASP for this TDP proposed in [26]. Thus in this section we present a general description of the main components of the method. The details can be found in that work.

A GRASP is an iterative process in which each major iteration consists typically of two phases:

construction and post-processing. The construction phase attempts to build a feasible solution S , and the post-processing phase attempts to improve it. When a feasible solution is successfully found in phase one, phase two is typically a local search within suitable neighborhoods with the aim of improving the objective function value. In this particular case, the construction phase does not necessarily terminate with a feasible solution, since the solution found may not be a p -partition or may violate constraints (4). Thus, both an adjustment phase, that modifies the current solution so as to assure a p -partition, and a post-processing phase, that attempts to improve the solution quality and to reduce the total relative infeasibility with respect to (4), are developed. The algorithm takes as an input an instance of the TDP, the maximum number of GRASP iterations, the restricted candidate list (RCL) quality parameter α , and the number of territories, and returns a solution S^{best} .

The motivation for GRASP in this particular application stems from the fact that it seems more appealing than current state-of-the-art approaches based on two-stage location-allocation algorithms for handling the connectivity constraints (5). By handling these constraints within a construction heuristic such as GRASP, the connectivity is always kept so it remains to appropriately address the balancing constraints (4).

Construction Phase At a given iteration a partial territory is considered and an attempt is made to either allocate an unassigned node to it or to “close” the current territory and “start” a new one. For favoring contiguity, when a new territory is started, the first node is an unassigned one with the smallest degree. When assigning a node that is not the first one in a territory, a greedy function that weighs both a distance-based dispersion measure and the relative violation of the balance constraints (4) is used. Let V_k be the current territory being built. Let $f(V_k) = \max_{i,j \in V_k} d_{ij}$ denote its corresponding dispersion measure. Note that this is an approximation to the original objective function (where the dispersion is taken with respect to a center node) which is less expensive to update and compute during the construction phase. In the local search, the original objective function is used. Recall that $w_a(V_k) = \sum_{i \in V_k} w_i^a$ is referred to as the size of V_k with respect to activity a , $a \in A$.

For a candidate node v , its greedy function is defined as

$$\phi(v) = \lambda F_k(v) + (1 - \lambda) G_k(v), \quad (6)$$

where

$$\begin{aligned} F_k(v) &= \left(\frac{1}{d_{\max}} \right) f(V_k \cup \{v\}) \\ &= \left(\frac{1}{d_{\max}} \right) \max \left\{ f(V_k), \max_{j \in V_k} d_{vj} \right\}, \end{aligned}$$

accounts for the original objective function, and

$$G_k(v) = \sum_{a \in A} g_k^a(v),$$

with

$$g_k^a(v) = (1/\mu^a) \max\{w^a(V_k \cup \{v\}) - (1 + \tau^a)\mu^a, 0\},$$

accounts for the sum of relative infeasibilities for the balancing constraints. Here, $d_{\max} = \max_{i,j \in V} \{d_{ij}\}$ is used for normalizing the objective function. Note that $g_k^a(v)$ represents the infeasibility with respect to the upper bound of the balance constraint for activity a , and these two factors are weighted by a parameter λ in function (6).

The GRASP construction phase works as follows. In a given iteration, for each possible candidate move v its greedy function (6) is computed. Then, a restricted candidate list (RCL) by objective function quality value α is built. That is, the RCL contains all possible candidate moves whose greedy function value is within α % of the best move. An element is chosen randomly from the RCL. Then criterion for closing the current territory is checked. If met, that is, a balance constraint upper bound has been violated, the current territory is “closed” and a new one is “started”. Note that in fact, this threshold is adjusted by a parameter $\rho > 0$ which allows for further flexibility in succeeding stages. A value of $\rho < 1$ allows, for instance, to close a territory having a relatively small size. This could allow, however, a violation of the upper bound of constraints (4) when merging territories in the adjustment phase. The construction procedure is referred to as `BuildGreedyRandomized(α)` which takes as input the quality parameter α (in addition to the TDP instance naturally).

Adjustment Phase Procedure `BuildGreedyRandomized()` does not necessarily return a feasible solution. In particular, a solutions may not be a p -partition and (4) may not be satisfied. To address this issue, a two-step post-processing phase is applied (`Adjustment()` and `LocalSearch()`). First, the number of territories q found in the construction phase is different from p , the procedure `Adjustment()` either merges territories (when $q > p$) or splits territories (when $q < p$). The merging operation consists of iteratively considering a territory of smallest size and merging it with its smallest neighboring territory. By smallest size it is meant the relative territory size with respect to the sum of the three node attributes. This reduces the number of connected territories by one at each iteration. This is iteratively repeated until $q = p$. The splitting operation consists of taking a territory of largest size, and splitting it into two connected territories. By noting that this problem is in a fact another territory design problem with $p = 2$, it can be solved by recursively applying the same GRASP to the subgraph induced by this territory to be splitted taking $p = 2$ as input. This increases the number of territories by one at each iteration, so the procedure is performed

iteratively until $q = p$. Note that the merging operation can be done very efficiently, while the splitting operation is itself another TDP problem. However, the nature of the construction phase makes merging more likely to be applied than splitting. In fact, in the empirical evaluation of the procedure, it has been found that the splitting operation is required in less than 0.4% of the cases.

Local Search After this adjustment step, a post-processing phase consisting of a local search is performed. Procedure `LocalSearch()` attempts both to recover feasibility of constraints (4) and to improve the objective function value. In this local search, a merit function that weighs both infeasibility with respect to (4) and the objective function value is used. In fact, this function is similar to the greedy function used in the construction phase with the exception that now the sum of relative infeasibilities takes into consideration both lower and upper bound violation of the balancing constraints. Specifically, for a given partition $S = \{V_1, \dots, V_p\}$, its merit function $\psi(S)$ is given by

$$\psi(S) = \lambda F(S) + (1 - \lambda)G(S)$$

where

$$F(S) = \left(\frac{1}{d_{\max}} \right) \max_{k=1, \dots, p} \left\{ \max_{i, j \in V_k} d_{ij} \right\},$$

and

$$G(S) = \sum_{k=1}^p \sum_{a \in A} g^a(V_k),$$

with $g^a(V_k) = (1/\mu^a) \max\{w^a(V_k) - (1 + \tau^a)\mu^a, (1 - \tau^a)\mu^a - w^a(V_k), 0\}$, being the sum of the relative infeasibilities of the balancing constraints.

A neighborhood $N(S)$ made up of all solutions reachable from S by moving a basic unit i from its current territory $t(i)$ to a neighbor district $t(j)$, where j is the corresponding basic unit in territory $t(j)$ adjacent to i , without creating a non-contiguous solution is used. Such a move is denoted by $move(i, j)$. Note that $move(i, j)$ is allowed only if $V_{t(j)} \cup \{i\}$ is connected (which is always the case if arc (i, j) exists), and $V_{t(i)} \setminus \{i\}$ remains connected. In practice an additional stopping criteria, such as $limit_moves$, is added to avoid performing the search for a relatively large amount of time. So the procedure stops as soon a local optima is found or the number of moves exceeds $limit_moves$. A first improving rule is used, that is, each potential move is examined at a time and a move is actually made as soon as an improving move is found.

Reactive GRASP The RCL quality parameter α is basically the only parameter to be calibrated in a practical implementation of a GRASP. Feo and Resende [10] have discussed the effect of the choice of the value of α in terms of solution quality and diversity during the construction phase and how it impacts the outcome of a GRASP. In a Reactive GRASP approach [7, 23] this α is self-adjusted according to the quality of the solutions previously found.

Instead of using a fixed value for the parameter α , which determines what elements are placed in the RCL at each iteration of the construction phase, the procedure randomly selects this value α from a discrete set $A = \{\alpha_1, \dots, \alpha_m\}$ containing m predetermined acceptable values. Using different values of α at different iterations allows for building different RCLs, possibly leading to the construction of different solutions which would never be built if a single, fixed value of α was used. Let p_i denote the probability associated with the choice of α_i , for $i = 1, \dots, m$. Initially, $p_i = 1/m, i = 1, \dots, m$, corresponding to a uniform distribution. Then these probabilities are periodically updated using information collected during the search. Different strategies for this update can be explored.

At any GRASP iteration, let A_i be the average value of the solutions obtained with $\alpha = \alpha_i$ in the construction phase. The probability distribution is periodically updated every *update_period* iterations (a value of *update_period* = 200 is used in the implementation) as follows. Compute first $q_i = (1/A_i)^\delta$ for $i = 1, \dots, m$, and then update the new values of the probabilities by normalization of the q_i as $p_i = q_i / (\sum_j q_j)$. Note that the smaller the A_i , the higher the corresponding p_i . Consequently, in the next block of iterations, the values of α that lead to better solutions have higher probabilities and are more frequently used in the construction phase. The exponent δ may be used and explored to differently attenuate the updated values of the probabilities. In our case, a value of $\delta = 8$ is used.

5 Empirical Work

The procedure was compiled with the Sun C++ compiler workshop 8.0 under the Solaris 9 operating system and run on a SunFire V440 with 4 UltraSPARC IIIi processors at 1062 MHz. For the experiments, randomly generated problems based on real-world data provided by the industrial partner were generated.

Data set DU was randomly generated as follows. Each instance topology was randomly generated as a planar graph in the $[0, 500] \times [0, 500]$ plane. Then, each of the three node activities were generated from a uniform distribution in the following way. Each node in a particular instance represents the aggregated information of $34000/n$ blocks in the original graph, where $n = |V|$. This is so because the original problem contains 34000 city blocks. Thus, each node is the sum of $34000/n$ independent uniformly distributed random variables. Thus for the 1000-node instances, each node contains the sum of 34 uniform random variables, and for the 2000-node instances this is the sum of 17 uniform random variables. The number of customers, in each block, is randomly generated in the $[0, 3]$ range. Product demand and workload are generated in the $[1, 12]$ range. A data set named DU05, DU03, and DU01 was generated according to the parameter τ_a equal to 0.05, 0.03, and 0.01, respectively. Recall that this parameter sets the allowable deviation from the target of the balancing constraints. For each of these set types, 20 different instances of size $n = 1000$,

DU05 1000-node Instances

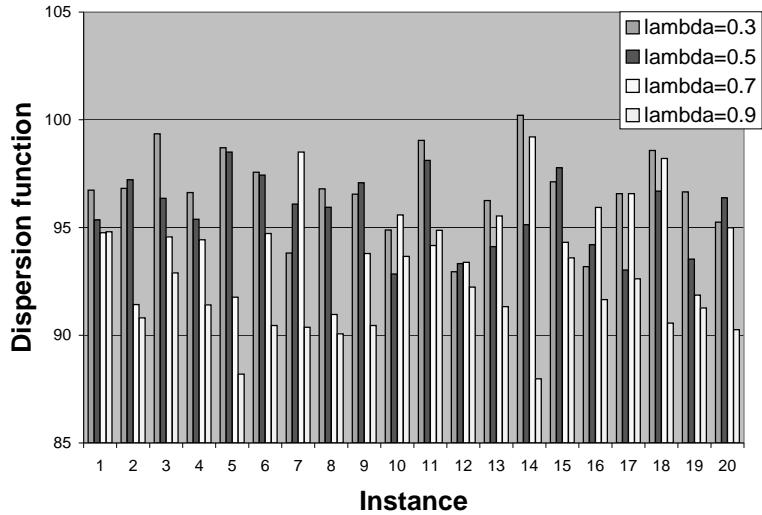


Figure 1: Effect of the weight parameter λ in the reactive GRASP.

$n = 2000$ and $p = 20$ were generated. The closing criteria parameter ρ , which is used within the GRASP construction phase, for deciding when to close (stop allocating new nodes to it) a currently active territory and start a new one, is set to 1.0.

The sensitivity of the algorithm with respect to the choice of the parameter λ , which is used within the GRASP construction phase, as a weight parameter in the greedy function, is investigated. See Section 4. To this effect, the GRASP with the local search phase for different values of λ (0.3, 0.5, 0.7, 0.9, 1.0) is executed, and the quality of the weighted objective function (ψ), the distance-based measure (F), and the degree of infeasibility (G) is measured. In this study, the number of GRASP iterations was set to 500.

Results over twenty 1000-node instances both for DU05 sets are shown in Figures 1 and 2. The former shows the value of the dispersion function for each of the 20 instances for the different choices of λ different from 1. The latter shows the same comparison, but displaying the relative gap from the best known solution. The choice $\lambda = 1.0$ is not displayed in the figures because its associated results are extremely bad. As it can be seen, the results when $\lambda = 0.9$ dominate the other three.

Summary of the results using a finer choice for λ (0.70, 0.75, 0.80, 0.85, 0.90, 0.95, 1.00) is shown in Table 1. The quality of the solutions are better on average when $\lambda = 0.95$. In addition, more best solutions were found under this same setting. As can be seen, there is a tendency that results get better as λ gets large. Naturally, the extreme case $\lambda = 1.0$ is the worst since this means

DU05 1000-node Instances

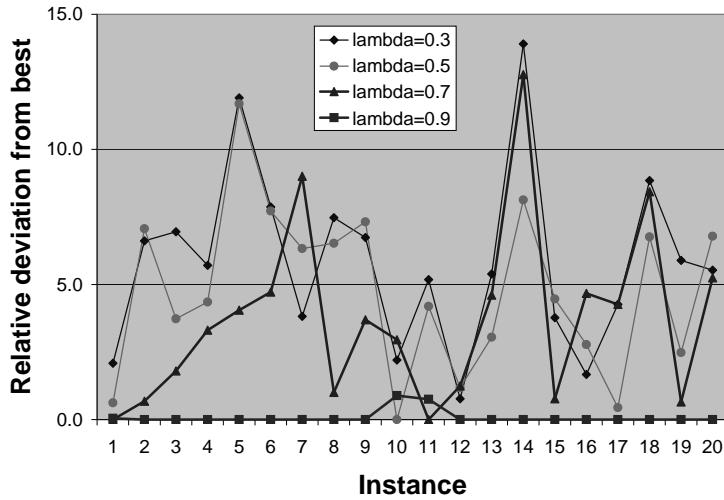


Figure 2: Effect of the weight parameter λ in the reactive GRASP.

that no regard to the violation of the balance constraints is considered whatsoever. Table 2 displays the results for the 2000-node instances. As it can be seen, the results also show a tendency towards a large value for λ . Note that in any event, the method was able to find feasible solutions in all instances tested except for the extreme case $\lambda = 1.0$. So as a conclusion, the evidence indicates that better results are obtained when more weight is given to the dispersion function violation than the one given to the violation of the balance constraints. For the 1000-node instances, CPU times range in 306-376 sec. For the 2000-node instances, CPU times were in the 2137-2286 sec. range. Another observation is that in general, the local search brought a very high benefit to the quality of the solutions found in the construction phase. These phase 1 solutions were improved by over 75% on average by the local search. Incidentally, about the running times, it was observed that about 40% of the effort was spent in the construction phase, 15% in the adjustment phase, and 45% in the local search. Another important observation is that the algorithm delivered 100% of feasible solutions for all cases, except for the case $\lambda = 1.0$, where again, this was to be expected.

A second experiment consists of evaluating of both the algorithmic performance and the effect both in the dispersion function and the feasibility with respect to the balancing constraints when tighter tolerance levels for τ^a are used. So, we run the reactive GRASP (fixing λ at 0.95) under the same instances, but using tolerance levels of 0.05, 0.04, 0.03, and 0.02. Table 3 display a summary of the results for the 1000- and 2000-node instances.

It can be seen how as the tolerance gets tighter the dispersion function grows. An interesting

Table 1: Evaluation of the effect of the weight parameter λ in the reactive GRASP on DU05 1000-node instances.

Statistic	λ						
	0.70	0.75	0.80	0.85	0.90	0.95	1.0
RGB (average)	5.8	5.3	4.0	2.8	2.2	0.9	32.2
RGB (worst)	12.8	11.8	9.2	7.4	7.5	3.8	67.1
RLSI	90.8	89.8	87.3	83.3	75.4	62.5	7.7
NIS	0	0	0	0	0	0	20
NBS	0	0	1	3	8	8	0

RGB: Relative gap from best
 RLSI: Relative local search improvement
 NIS: Number of infeasible solutions
 NBS: Number of best solutions

Table 2: Evaluation of the effect of the weight parameter λ in the reactive GRASP on DU05 2000-node instances.

Statistic	λ						
	0.70	0.75	0.80	0.85	0.90	0.95	1.0
RGB (average)	9.2	5.3	3.4	3.2	8.1	0.7	39.9
RGB (worst)	43.3	17.5	7.7	11.3	30.6	3.7	69.6
RLSI	69.0	83.2	80.5	76.2	55.5	56.2	6.0
NIS	0	0	0	0	0	0	20
NBS	0	1	2	3	2	12	0

observation is that the algorithm still finds feasible solutions to all instances with tolerance factors as low as 0.03, but when going down to the DU02 instances ($\tau^a = 0.02$) the algorithm can no longer find feasible solutions to all instances. For the 1000-node instances it found 4 feasible solutions (20%). For the 2000-node instances, it found 16 feasible solutions (80%).

Finally, a comparison with current industry standards is performed. Results over twenty 1000-node and 2000-node instances both for DU03 sets are shown in Figures 3 and 4, respectively. We first note that, the solution found by the firm alone has extremely large deviations from feasibility with respect to the balancing constraints. In the figures, our local search scheme was applied to the solution found by the firm method, so this is indicated by the term “Firm+LS”. It can be observed that the Reactive GRASP find consistently significantly better solutions than those found by the firm method. In fact, the quality of the results is even greater since our proposed method found 100% of feasible solutions while the firm method found only 2.5% of feasible solutions (with respect to the tolerance level of 0.03 used here). So this magnifies the value of the proposed approach and is consistent with the results reported in [26].

Table 3: Evaluation of the effect of the tolerance parameter τ^a in the reactive GRASP on 1000- and 2000-node instances.

1000 nodes		τ^a			
Statistic		0.05	0.04	0.03	0.02
RGB (average)	0.9	1.1	4.2	26.5	
RGB (worst)	6.0	3.7	12.4	47.9	
NIS	0	0	0	16	
AIS	0.0	0.0	0.0	0.003	
NBS	10	9	1	0	
2000 nodes		τ^a			
Statistic		0.05	0.04	0.03	0.02
RGB (average)	1.0	3.1	3.6	22.5	
RGB (worst)	5.0	10.9	10.1	84.6	
NIS	0	0	0	3	
AIS	0.0	0.0	0.0	$< 10^{-3}$	
NBS	11	3	6	0	

AIS: Average of infeasible solutions

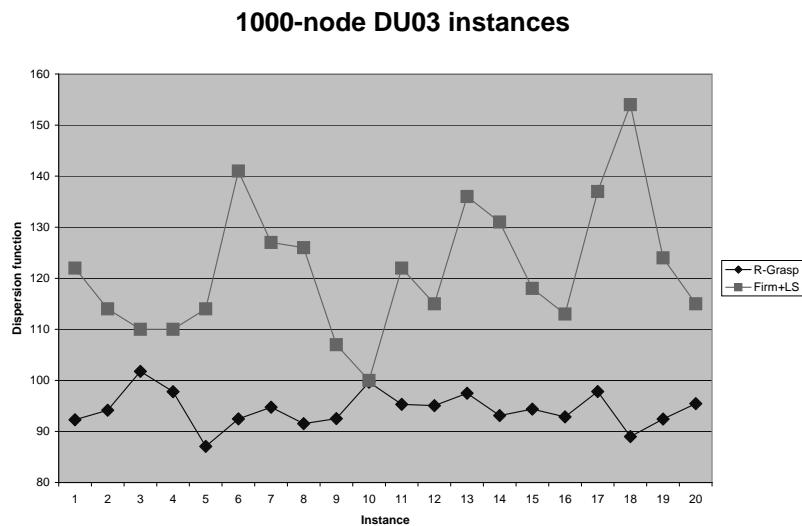


Figure 3: Reactive GRASP vs. Firm method on 1000-node DU03.

2000-node DU03 instances

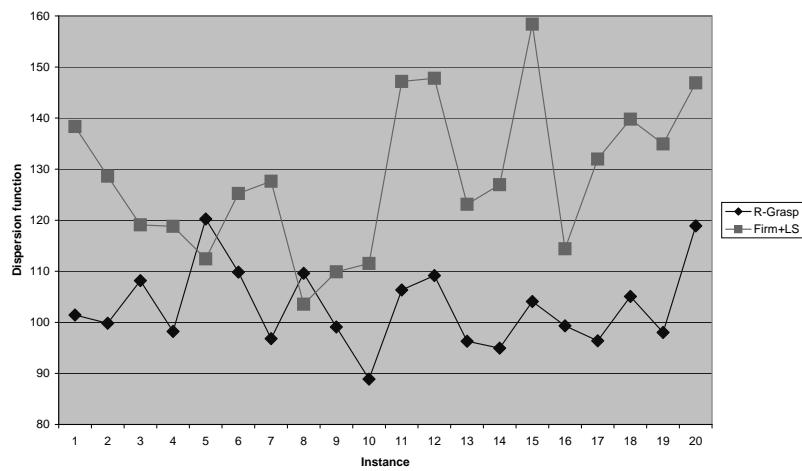


Figure 4: Reactive GRASP vs. Firm method on 2000-node DU03.

6 Conclusions

In this paper, computational experience with a reactive GRASP heuristic for a commercial territory design problem was presented. The heuristic was evaluated on 1000- and 2000-node instances. In particular, an evaluation of the sensitivity of a greedy function parameter that weights the contribution of the original objective function and the violation of the balance constraints was assessed. The results indicate that using a greedy function with larger values of this weight parameter is preferred, that is, more weight given to the dispersion function is preferred. In addition, the results show the benefit of the local search since the solution found in the construction phase improved over 55% on average in all cases where λ is different from 1.0. This extends the results reported in [26] for 500-node instances. Another important result from the practical perspective is that the algorithm was able to find feasible solutions consistently, even for tolerance levels as low as 3%. Nowadays, the current methodology used by the firm to try to obtain solutions has tremendous difficulties on finding feasible designs not only to DS05 type of instances, but even for DS10 instances, that is instances where a deviation of $\tau_a = 0.10$ is permitted in the balancing constraints. In the present work, we showed a comparison that shows the Reactive GRASP approach outperforms both in dispersion function quality and feasibility violation. Overall, we have provided a very valuable tool for a more efficient territory design planning according to the company planning requirements.

There are still several areas of opportunity for further work on this problem. From the heuristic perspective, the local search phase may be improved by developing other neighborhoods, such as node swapping for instance. A natural extension could be the development of other metaheuristics such as tabu search or scatter search. From the practical standpoint, the issue of territory realignment is an important area of opportunity. This problem consists of, given a current design, how to efficiently accommodate for system changes such as customers additions or dropouts trying not to disrupt the previous design considerably. Another line of research is addressing parameters such as product demand from a stochastic programming perspective. This will lead to an integer stochastic programming model with stochastic parameters in the input coefficient matrix, which is certainly a very challenging problem.

Acknowledgments: I am grateful to the anonymous referees for reviewing this paper. This work was supported by the Mexican National Council for Science and Technology under its Basic Research Program (grants SEP-CONACYT CB05-1-48499-Y and CB11-1-166397), and Universidad Autónoma de Nuevo León under its Scientific and Technological Research Support Program (grant UANL-PAICYT CA1478-07).

References

- [1] F. Bação, V. Lobo, and M. Painho. Applying genetic algorithms to zone design. *Soft Computing*, 9(5):341–348, 2005.
- [2] M. Blais, S. D. Lapierre, and G. Laporte. Solving a home-care districting problem in an urban setting. *Journal of the Operational Research Society*, 54(11):1141–1147, 2003.
- [3] B. Bozkaya, E. Erkut, and G. Laporte. A tabu search heuristic and adaptive memory procedure for political districting. *European Journal of Operational Research*, 144(1):12–26, 2003.
- [4] S. I. Caballero-Hernández, R. Z. Ríos-Mercado, F. López, and S. E. Schaeffer. Empirical evaluation of a metaheuristic for commercial territory design with joint assignment constraints. In J. E. Fernandez, S. Noriega, A. Mital, S. E. Butt, and T. K. Fredericks, editors, *Proceedings of the 12th Annual International Conference on Industrial Engineering Theory, Applications, and Practice (IJIE)*, pages 422–427, Cancun, Mexico, November 2007. ISBN: 978-0-9654506-3-8.
- [5] F. Caro, T. Shirabe, M. Guignard, and A. Weintraub. School redistricting: Embedding GIS tools with integer programming. *Journal of the Operational Research Society*, 55(8):836–849, 2004.
- [6] S. J. D'Amico, S.-J. Wang, R. Batta, and C. M. Rump. A simulated annealing approach to police district design. *Computers & Operations Research*, 29(6):667–684, 2002.
- [7] H. Delmaire, J. A. Díaz, E. Fernández, and M. Ortega. Reactive GRASP and tabu search based heuristics for the single source capacitated plant location problem. *INFOR: Information Systems and Operational Research*, 37(3):194–225, 1999.
- [8] A. Drexel and K. Haase. Fast approximation methods for sales force deployment. *Management Science*, 45(10):1307–1323, 1999.
- [9] J. C. Duque, R. Ramos, and J. Suriñach. Supervised regionalization methods: A survey. *International Regional Science Review*, 30(3):195–220, 2007.
- [10] T. A. Feo and M. G. C. Resende. Greedy randomized adaptive search procedures. *Journal of Global Optimization*, 6(2):109–133, 1995.
- [11] E. Fernández, J. Kalcsics, S. Nickel, and R. Z. Ríos-Mercado. A novel maximum dispersion territory design model arising in the implementation of the WEEE-directive. *Journal of the Operational Research Society*, 61(3):503–514, 2010.

[12] B. Fleischmann and J. N. Paraschis. Solving a large scale districting problem: A case report. *Computers & Operations Research*, 15(6):521–533, 1988.

[13] M. G. García-Ayala, J. L. González-Velarde, R. Z. Ríos-Mercado, and E. Fernández. A novel model for arc territory design: Promoting Eulerian districts. *International Transactions of Operational Research*, 23(3):433–458, 2016.

[14] R. S. Garfinkel and G. L. Nemhauser. Optimal political districting by implicit enumeration techniques. *Management Science*, 16(8):B495–B508, 1970.

[15] S. W. Hess and S. A. Samuels. Experiences with a sales districting model: Criteria and implementation. *Management Science*, 18(4):P41–P54, 1971.

[16] S. W. Hess, J. B. Weaver, H. J. Siegfeldt, J. N. Whelan, and P. A. Zitzlau. Nonpartisan political redistricting by computer. *Operations Research*, 13(6):998–1006, 1965.

[17] M. Hojati. Optimal political districting. *Computers & Operations Research*, 23(12):1147–1161, 1996.

[18] J. Kalcsics. Districting problems. In G. Laporte, S. Nickel, and F. Saldanha da Gama, editors, *Location Science*, chapter 23, pages 595–622. Springer, Cham, Switzerland, 2015. ISBN: 978-3-319-13110-8.

[19] J. Kalcsics, S. Nickel, and M. Schröder. Towards a unified territorial design approach: Applications, algorithms, and GIS integration. *TOP*, 13(1):1–56, 2005.

[20] J. F. López-Pérez and R. Z. Ríos-Mercado. Embotelladoras ARCA uses operations research to improve territory design plans. *Interfaces*, 43(3):209–220, 2013.

[21] A. Mehrotra, E. L. Johnson, and G. L. Nemhauser. An optimization based heuristic for political districting. *Management Science*, 44(8):1100–1114, 1998.

[22] L. Muyldermans, D. Cattrysse, D. Van Oudheusden, and T. Lotan. Districting for salt spreading operations. *European Journal of Operational Research*, 139(3):521–532, 2002.

[23] M. Prais and C. C. Ribeiro. Reactive GRASP: An application to a matrix decomposition problem in TDMA traffic assignment. *Journal on Computing*, 12(3):164–176, 2000.

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

[25] R. Z. Ríos-Mercado and H. J. Escalante. Grasp with path relinking for commercial districting. *Expert Systems with Applications*, 44:102–113, 2016.

[26] R. Z. Ríos-Mercado and E. A. Fernández. A reactive GRASP for a commercial territory design problem with multiple balancing requirements. *Computers & Operations Research*, 36(3):755–776, 2009.

[27] R. Z. Ríos-Mercado and J. F. López-Pérez. Commercial territory design planning with realignment and disjoint assignment requirements. *Omega*, 41(3):525–535, 2013.

[28] R. Z. Ríos-Mercado and J. C. Salazar-Acosta. A GRASP with strategic oscillation for a commercial territory design problem with a routing budget constraint. In I. Batyrshin and G. Sidorov, editors, *Advances in Soft Computing*, volume 7095 of *Lecture Notes in Artificial Intelligence*, pages 307–318. Springer, Heidelberg, Germany, 2011.

[29] M. A. Salazar-Aguilar, R. Z. Ríos-Mercado, and M. Cabrera-Ríos. New models for commercial territory design. *Networks & Spatial Economics*, 11(3):487–507, 2011.

[30] M. A. Salazar-Aguilar, R. Z. Ríos-Mercado, and J. L. González-Velarde. A bi-objective programming model for designing compact and balanced territories in commercial districting. *Transportation Research Part C: Emerging Technologies*, 19(5):885–895, 2011.

[31] M. A. Salazar-Aguilar, R. Z. Ríos-Mercado, and J. L. González-Velarde. GRASP strategies for a bi-objective commercial territory design problem. *Journal of Heuristics*, 19(2):179–200, 2013.

[32] M. A. Salazar-Aguilar, R. Z. Ríos-Mercado, J. L. González-Velarde, and J. Molina. Multiobjective scatter search for a commercial territory design problem. *Annals of Operations Research*, 199(1):343–360, 2012.

[33] M. Segal and D. B. Weinberger. Surfing. *Operations Research*, 25(3):367–386, 1977.

[34] L. Vargas-Suárez, R. Z. Ríos-Mercado, and F. López. Usando GRASP para resolver un problema de definición de territorios de atención comercial. In M. G. Arenas, F. Herrera, M. Lozano, J. J. Merelo, G. Romero, and A. M. Sánchez, editors, *Proceedings of the IV Spanish Conference on Metaheuristics, Evolutionary and Bioinspired Algorithms (MAEB)*, pages 609–617, Granada, Spain, September 2005. In Spanish.

[35] A. A. Zoltners and P. Sinha. Sales territory design: Thirty years of modeling and implementation. *Marketing Science*, 24(3):313–331, 2005.