Proceedings of the 12th Annual
International Conference on Industrial Engineering
Theory, Applications and Practice
Cancun, Mexico
November 4-7, 2007

EMPIRICAL EVALUATION OF A METAHEURISTIC FOR COMMERCIAL TERRITORY DESIGN WITH JOINT ASSIGNMENT CONSTRAINTS

Saúl I. Caballero-Hernández[1*], Roger Z. Ríos-Mercado[1], Fabián López[2], and Satu Elisa Schaeffer[1]

[1]Graduate Program in Systems Engineering
Universidad Autónoma de Nuevo León
San Nicolás de los Garza, Nuevo León, Mexico
Corresponding author's e-mail: saul@yalma.fime.uanl.mx

[2]Grupo ARCA
Monterrey, México

Abstract: In industry, territory design is motivated by changes in the number or properties of customers served by a given company. The goal of territory design is to group the customers into manageable-sized territories. It is often required to balance the demand among the territories in order to delegate responsibility fairly. In this paper, we present a metaheuristic solution approach based on GRASP (Greedy Randomized Adaptive Search Procedure) to a particular commercial territory design problem motivated by a real-world application in a beverage distribution company in the city of Monterrey, Mexico. Our empirical work includes an evaluation of the overall algorithmic performance. In addition, we study the effect of the weight parameter of the GRASP greedy function (which is a convex combination of the original dispersion-based objective function and the relative violation of the balance constraints) on the quality of the final solution. The experiments were carried out over a set of randomly generated instances based on real-world data from the industrial partner.

1. INTRODUCTION

The *territory design problem* is concerned with grouping small geographic areas into larger geographic clusters called territories in such a way that the latter fulfill certain planning criteria. This problem belongs to the family of districting problems that have a broad range of applications like political districting and the design of sales and services territories. The work of Kalcsics, Nickel, and Schröder (2005) is an extensive survey on approaches proposed for territory design that presents an up to date state of the art on the topic.

The problem addressed in this paper is motivated by a real-world application in the city of Monterrey, Mexico. The firm wishes to partition the city area into disjoint territories that are suitable for their commercial purposes. In particular, the firm wishes to design territories that be *balanced* (similar in size) with respect to each of two different *activity measures* (number of customers and sales volume). Additional planning criteria include: (i) *contiguity* of each territory, so that each *basic unit* (BU) can reach each other by traveling within the territory; (ii) territory *compactness*, so that customers within a territory are relatively close to each other; (iii) *joint assignment* of BUs, so that specified pairs of BUs must be assigned to the same territory in the design, and (iv) a fixed number of territories.

Although several commercial territory design approaches have appeared in the literature, the specific features present in this concrete problem make it very unique, and it has not been addressed before to the best of our knowledge. Vargas-Suárez, Ríos-Mercado, and López (2005) studied a similar problem without compactness, contiguity and joint assignment constraints and where the objective was to minimize the unbalance among territories. Ríos-Mercado and Fernández (2006) studied the problem considering compactness and contiguity but without joint assignment constrains.

In this work we propose and develop a solution approach based on GRASP (Feo and Resende, 1995) to construct high-quality solutions. Our empirical work includes an evaluation of the overall algorithmic performance of the proposed approach. Our computational experiments show that this method yields good results.

2. PROBLEM DESCRIPTION

Let $G = (V, E)$ be a graph where each BU is represented by a node $i \in V$ and an arc connecting nodes i and j exists in E if units i and j are located in adjacent blocks. Each node has three properties: geographical coordinates (c_i^x, c_i^y), and two measurable activities. Let ω_i^a be the value of activity a at node i, where the first activity, $a = 1$, represents the number of customers within that BU and the second activity, $a = 2$, represents the total sales volume of the BU. A *territory* is a subset of nodes $V_k \subset V$. The number of territories p is fixed and given as a parameter. It is required that each node is assigned to only one territory.

One of the desired properties in a solution is that the territories be balanced with respect to each of the activity measures. Let us define the size of territory V_k with respect to activity a as: $\omega^a(V_k) = \sum_{i \in V_k} \omega_i^a, a = 1,2$. 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 overcome this difficulty we measure the degree of balance by computing the *relative deviation* of each territory from its average value μ^a of activity a that is, $\mu^a = \omega^a(V)/p, a = 1,2$. We are also given a collection H of pairs of nodes such that $[i, j] \in H$ implies that node i and node j must be assigned to the same territory. In addition, the industrial partner requires that in each of the territories, the units must be geographically located relatively close to each other. One way to achieve this is to assign one node in each territory to serve as a territory center, denoting the center of territory k by $c(k)$ and then to define a distance measure such as

$$f(V_1, V_2, \ldots, V_p) = \max_{\substack{k=1,\ldots,p \\ j \in V_k}} \{d_{c(k)j}\}, \tag{1}$$

over the p territories, $d_{c(k)j}$ represents the Euclidean distance from node j to center of territory k. All parameters of the model are assumed to be known with certainty. The problem can be thus described as finding a p-partition of V that satisfies the specified planning criteria of balancing:

$$(1-\tau^a)\mu^a \le \omega^a(V_k) \le (1+\tau^a)\mu^a \quad \forall k, \tag{2}$$

contiguity, and joint assignment, that minimizes the above distance measure (1). This problem can be modeled in terms of integer programming as a p-center problem with additional constraints on capacity, contiguity and joint assignment. Given the complexity of graph partitioning problems and the nature of the additional constraints, such as contiguity, we resort to a metaheuristic that handles some of the difficulty already in the construction phase, and further assuring in each step that the contiguity is not broken.

3. PROPOSED ALGORITHM

GRASP (Feo and Resende, 1995), a well-known metaheuristic that captures good features of both pure greedy algorithms and random construction procedures, has been widely used for solving many combinatorial optimization problems. GRASP is a two-phase iterative process. In GRASP iteration, we first perform a *construction phase* for building a feasible solution and then a *post-processing phase* that aims to improve the solution. The post-processing typically consists of a local search within suitable neighborhoods.

The proposed algorithm starts with an infeasible territory assignment in which each BU is assigned to a singleton territory, $V^o = (V_1, \ldots, V_n)$ with $V_i = \{i\}$. Before continuing with the construction, we force the joint assignment of the specified pairs of BUs (v, w) in H: we find a path P in G from v to w and merge into a single territory all the territories that contain one or more nodes included in P. In each subsequent step of the construction phase, we merge two adjacent territories until we reach the desired number of territories. By construction, all the territories will be contiguous. However, this solution might not be feasible with respect to the balance constraints (2). The post-processing phase attempts to reduce if not completely overcome the balance-constraint infeasibility. The post-processing phase also aims to reduce the maximum distances within the territories. The relative importance of these two tasks is controlled by a weight parameter. We now describe in detail each of the components of the algorithm.

3.1 Pre-Processing Phase

To address the joint assignment of basic units, we need to take into account also the contiguity requirement – the final solution must fulfill both requirements. One way to do this is to find, for each (v, w) in H, a path P in G between v and w

that must be assigned to the same territory and then by merging all territories that contain one or more nodes in the path *P*. The newly obtained territory then contains both of the desired basic units and is necessarily contiguous if all the merged territories were contiguous. The initial territory assignment consists of singleton territories and the merging is done with respect to paths in *G*, and hence all created territories are contiguous as well. For additional diversity of solutions, in the spirit of GRASP, instead of using always the shortest path, we compute *k* shortest paths (Martins, Pascoal, and Santos, 1999) between the BUs and randomly select one of these paths.

3.2 Construction Phase

After the pre-processing phase, the number of territories is different from *p*. For the majority of natural instances, the number of territories is significantly larger than *p*; our algorithm does not handle the rare cases where the number of territories is smaller than *p*. We decrease the number of territories by iteratively merging adjacent territories two at a time until *p* territories remain. The two territories to merge are chosen using a greedy function that weighs both a *distance-based dispersion measure* $f(V_k) = \max_{i,j \in V_k} \{d_{ij}\}$ and the relative violation of the balance constraints (2). Let us define $\omega^a(V_k) = \sum_{i \in V_k} \omega_i^a$ as the *size* of the territory V_k with respect to the activity *a*. Let V_i and V_j be two adjacent territories (that is, territories connected by at least one edge in *G*), let us define their greedy function as

$$\phi(V_i, V_j) = \lambda\, f(V_i \cup V_j) + (1-\lambda) G(V_i \cup V_j), \tag{3}$$

where $V_i \cup V_j$ is the new territory that is the union of the two smaller territories, λ is a parameter, and

$$G(V_i \cup V_j) = \sum_{a \in A} g^a(V_i \cup V_j), \tag{4}$$

with $g^a(V_k) = (1/\mu^a) \max\{\omega^a(V_k) - (1+\tau^a)\mu^a, 0\}$ as the sum of the relative infeasibilities with respect to the upper bound of the balance constraint of activity *a*.

In the spirit of GRASP, the list of all possible candidate moves are sorted by non-decreasing value of its greedy function, and a restricted candidate list (RCL) is constructed by selecting all those moves which are within α (or 100 α %) of the best move. A move from the RCL is chosen randomly, the greedy function of the remaining moves is updated and we proceed iteratively until the number of territories equals *p*.

3.3 Post-Processing Phase

After the construction phase has reached the desired number of territories, the post-processing phase consisting of a local search is performed. The goal of this phase is both to reduce the (possible) infeasibility with respect to the balance constraints and to improve the value of the objective function. The local search uses a *merit function* that weighs both the infeasibility of the balance constraints and the objective function. For a partition $S = \{V_1, \ldots, V_p\}$, the merit function is given by

$$\psi(S) = \lambda F(S) + (1-\lambda) I(S) \tag{5}$$

where *F*(*S*) is the dispersion measure given by

$$F(S) = \max_{k=1,\ldots,p} \left\{ \max_{j \in V_k} d_{c(k)j} \right\} \tag{6}$$

and *I*(*S*) is the sum of the relative infeasibilities of the balance constraints,

$$I(S) = \sum_{k=1}^{p} \sum_{a \in A} i^a(V_k), \tag{7}$$

with $i^a(V_k) = (1/\mu^a) \max\{\omega^a(V_k) - (1+\tau^a)\mu^a, (1-\tau^a)\mu^a - \omega^a(V_k), 0\}$.

We define a neighborhood $N(S)$ of a partition S to consists of all partitions S' reachable from S by moving a basic unit i from its current territory $t(i)$ to the territory of another BU j such that the edge $(i, j) \in E$ and $t(j) \neq t(i)$. By only considering territories directly adjacent to the BU i, we can assure that also the partition S' has only contiguous territories.

4. EXPERIMENTAL RESULTS

In this section we present the experimental results obtained with a C++ implementation of the proposed algorithm, compiled with the Sun C++ compiler workshop 8.0 under the Solaris 9 operating system and executed on a SunFire V440 server. For the experiments, we generated randomly problem instances based on real-world data provided by the industrial partner. Each instance consists of node coordinates generated uniformly at random in a [0; 500] x [0; 500] plane and a set of edges placed among the nodes to make a planar graph. The node activities were generated as the sum of d random numbers with uniform distribution, where d depends on the number of nodes, contracting together nodes from the real-world data. We experimented with 20 instances of size $n = 500$ and $p = 10$. The number of joint assignment constraints considered in the problem instances is about 5 % of n, this was suggested by the industrial partner. For all instances, the allowable deviation for the balancing constraints was set at 5 % ($\tau^a = 0.05$). The local search phase is executed until we reach a local optimum; that is, when we cannot find an improvement movement in the neighborhood.

4.1 Experiment A: GRASP Iteration Count

As GRASP is an iterative procedure, we need to determine for how many iterations we need to run the algorithm to reach a good solution in a reasonable time of computation. In order to determine the iteration count to use for the rest of the experiments, we first executed the algorithm five times over three instances for 1,000 GRASP iterations. Figure 1 shows how the quality of the best solution found improves over the iterations performed for two of these instances (each of the five repetitions shown as a separate line); the third instance behaved very similarly.

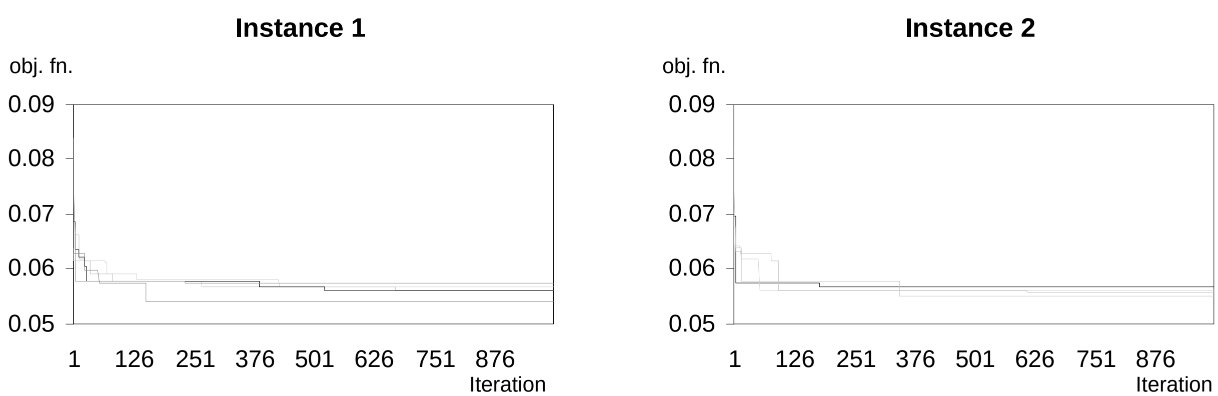

Figure 1. GRASP Convergence

Figure 1 shows that all major improvements in the value of the objective function tend to occur during the first 400 – 500 iterations, for which we find it safe to conclude that running the algorithm up to 500 iterations suffices for finding a good solution.

4.2 Experiment B: The effect of the weight parameter (λ)

We then studied the effect of the weight parameter λ (5) in the quality of the solutions found in the local search phase. We examined four different values of the parameter: 0.0, 1/3, 2/3, and 1.0, in order to assess this effect. Table 1 shows for each value of the parameter studied, the relative mean deviation from the best solution found, the number of times this value of the parameter lead to the best solution, and the number of infeasible solutions found. We exclude from the table the case λ = 1 due to not having found any feasible solutions with this value.

Table 1. Calibration of Parameter λ

Value of the parameter λ	0	1/3	2/3
Relative deviation from best solution (average)	0.05767	0.02335	0.00238
Frequency of being the best solution	1	5	15
Number of infeasible solutions	4	0	1

As we can learn from the table, larger values of the parameter had a better effect. This means that giving more weight to the objective function rather than to the violation of the constraints lead to better solutions.

4.3 Experiment C: Evaluation of the local search

In this experiment we studied the behavior of the solution quality during the local search procedure. We computed the average relative improvement on the quality of the solutions obtained by the local search at the end of each GRASP iteration. The results are shown in Figure 2. Each curve shows the path followed by the local search (in terms of the best solution found until that point) for five specific iterations (shown in the right).

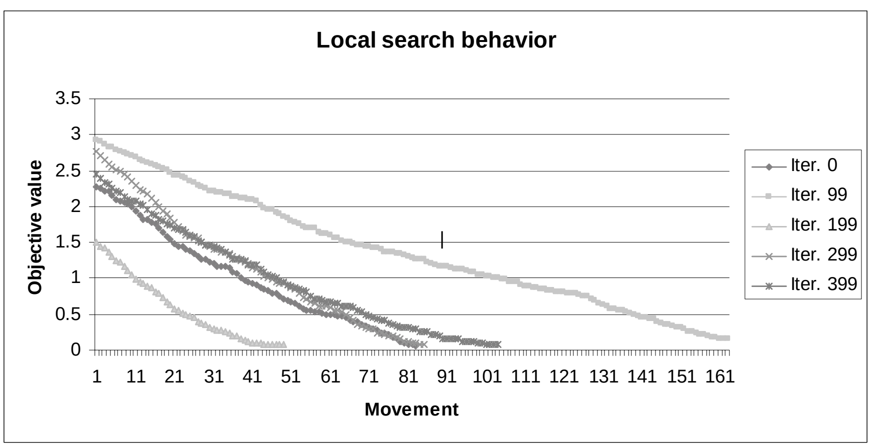

Figure 2. Behavior of the Local Search

The numerical results of this experiment show that the local search produces an improvement of about 90 percent in the quality of the solution. Moreover, as we can see, there are cases where convergence is relatively slow and other s where a local optimum is found relatively fast.

5. CONCLUSIONS

In this work we have presented a GRASP algorithm for this territory design problem. One of the main contributions of this work is the incorporation of joint-assignment constraints that to the best of our knowledge have not been treated before.

Experimental results show the effectiveness of the method in finding good-quality solutions for problems of size n = 500 in reasonably short computation times (about 170 seconds) using 500 GRASP iterations. In addition, we observed an excellent performance of the local-search phase. Our experiments indicate that a value of the weight parameter (λ) of 2/3 leads to the best solutions for the instances of the problem out of the values studied.

As future work we plan to execute a more extensive experimentation considering instances of size n = 1,000 and up, and study the behavior of another local-search procedure based on swap moves.

Acknowledgements*:* This research has been supported by the Mexican National Council for Science and Technology (grant SEP-CONACYT 48499-Y). In addition, the first author is supported by a fellowship for graduate studies from CONACYT.

6. REFERENCES

1. T. A. Feo and M. G. C. Resende (1995). Greedy randomized adaptive search procedures. Journal of Global Optimization, 6(2):109-133.

2. J. Kalcsics, S. Nickel, and M. Schröder (2005). Toward a unified territorial design approach: Applications, algorithms, and GIS integration. Top, 13(1):1-74.

3. E. Q. V. Martins, M. M. B. Pascoal, and J. L. E. Santos (1999). Deviation algorithms for ranking shortest paths. International Journal of Foundations of Computer Science, 10(3):247-263.

4. R. Z. Ríos-Mercado and E. Fernández (2006). A reactive GRASP for a sales territory design problem with multiple balancing requirements. Technical report PISIS-2006-12, Graduate Program in Systems Engineering, Universidad Autónoma de Nuevo León, San Nicolás de los Garza, Mexico, September.

5. L. Vargas-Suárez, R. Z. Ríos-Mercado, and F. López (2005). 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, Memorias del IV Congreso Español sobre Metaheurísticas, Algoritmos Evolutivos y Bioinspirados (MAEB), pp. 609-617, Granada, Spain, September. In Spanish.