Interfaces
Vol. 43, No. 3, May–June 2013, pp. 209–220
ISSN 0092-2102 (print) | ISSN 1526-551X (online)
http://dx.doi.org/10.1287/inte.1120.0675

Embotelladoras ARCA Uses Operations Research to Improve Territory Design Plans

J. Fabián López-Pérez, Roger Z. Ríos-Mercado
Universidad Autónoma de Nuevo León (UANL), San Nicolás de los Garza, NL 66450, Mexico
{jesus.lopezpz@uanl.edu.mx, roger.rios@uanl.edu.mx}

Embotelladoras ARCA is a Mexican company dedicated to the production, distribution, and sales of soft drink brands owned by the Coca-Cola Company, ARCA, and third parties. One of the most relevant problems the company faces is how to segment its customers into territories to accommodate the most efficient handling of marketing and distribution decisions. This territory design must also satisfy several planning requirements, such as territory compactness, connectivity, and balancing, and similarity with the existing design. Before 2009, the company planners defined these units based on their experiences, without using quantitative tools, giving more weight to the territory compactness criterion. This method of definition led to a number of undesirable issues, including highly unbalanced territories; that is, the final plans had large disparities in size with respect to both number of customers and total product demand. In this paper, we apply operations research methods to determine configurations of the territorial units that ensure territory balance with respect to both number of customers and total product demand, while complying with other important planning requirements and maximizing territory compactness. Using this methodology has resulted in significant enhancements with respect to the territory imbalance issue while maintaining the level of territory compactness. We also highlight additional benefits of this approach. Embotelladoras ARCA has adopted this proposed tool to make its design decisions.

Key words: bottled beverage distribution; operations research; commercial districting; basic units; mixed-integer programming model.
History: This paper was refereed. Published online in *Articles in Advance* March 28, 2013.

Embotelladoras ARCA was formed in 2001 by integrating three of the oldest bottlers in Mexico, and it became the second-largest bottler of Coca-Cola products in Latin America. The company distributes its products in the northern region of Mexico and, since 2008, in the north of Argentina and Ecuador. It also produces and distributes branded salty snacks, Bokados. Embotelladoras ARCA merged with Grupo Continental in January 2011 in a $2.3 billion exchange of stock; the resulting conglomerate is called Embotelladoras ARCA-Continental (or Arca-Contal) (http://www.arcacontal.com/). Headquartered in Monterrey, Mexico, Embotelladoras ARCA is dedicated to the production, distribution, and sales of snacks and soft drink brands. The company has soft drink sales of more than 1.2 billion unit cases and ranks as the third-largest Coca-Cola bottler in the world. It sells Ciel bottled water, tea, energy drinks, and snacks, and the usual Coke brands. The market's large size and the relevance of several problems that it faces motivated Embotelladoras ARCA to use and apply operations research (OR) models and techniques.

One of the company's most relevant problems is how to segment or partition its customers into clusters or territories to accommodate the efficient handling of marketing and distribution decisions. In essence, this is a commercial territory design problem (TDP). A commercial TDP may be viewed as the problem of grouping basic units (i.e., city blocks, zip codes, or individual customers), which we call BUs, into subsets according to specific planning criteria. These subsets are known as territories or districts. The geographic definition of the problem also includes some other spatial constraints.

In the Embotelladoras ARCA application, each territory is served by a single resource; therefore, using planning criteria to balance the quantity of customers and the product demand required by the dispatchers or truck drivers to cover each territory seems sensible. Moreover, balancing demand among the territories to delegate responsibility fairly is often necessary. To this end, Arca-Contal wishes to partition the city area into disjoint territories suitable for its commercial purposes. In particular, given a set of city blocks or BUs, it wants to create a specific number of territories according to planning criteria, such as: (1) compactness, whereby customers are as close to each other as possible; (2) balance, with respect to each of two activity measures (number of customers and product demand); (3) territory connectivity, such that a truck assigned to a territory can deliver goods without leaving the territory; (4) disjoint BU assignment, which avoids assigning two given BUs to the same territory; and (5) similarity to an existing plan for a subset of BUs. That is, the main objective of solving the TDP is to group the customers into manageably sized territories to guarantee that BUs assigned to a territory are relatively close to each other and meet the aforementioned planning criteria.

We note that one of the main motivations driving this work was the need for more balanced territories with respect to both the number of customers and product demand. Before we undertook this work, the company was facing imbalances of up to 30 percent off target. The practical implication of these imbalances was that some workers (i.e., salesmen or truck drivers) were finishing their shifts very early and others were finishing very late, creating dissatisfaction among the workers. The model we developed effectively addresses this imbalance.

To address this TDP, we derive a mixed-integer linear programming (MILP) model and develop a solution framework based on the iterative solution of an associated MILP model with a cut-generation strategy. We implemented this with an off-the-shelf modeling and optimization suite, X-PRESS from FICO™ (Fair Isaac, formerly known as Dash Optimization). The model and solution method are integrated into an interactive and user-friendly geographic information system (GIS) application, MAPINFO© (http://www.pbinsight.com/welcome/mapinfo/). This paper describes and illustrates the potential of the proposed approach as a user-friendly decision tool in the context of a case study for Embotelladoras ARCA. In the following sections, we (1) describe in detail the problem being addressed, (2) highlight an overview of other approaches to commercial districting problems, (3) present a general description of the solution framework, (4) illustrate the value of the approach using a case study, and (5) discuss the practical benefits that resulted from this work.

Overview of Related Work

Depending on the context of the problem, territory design (TD) may be used in the same manner as districting, which is a multidisciplinary research field that includes several areas, such as geography, political science, public administration, and OR. We can generalize that TD is common to all applications that operate with a group of resources that must be assigned in an optimal way to subdivide the work area into balanced regions of responsibility. Some of these applications are pick-up and delivery applications, including waste collection, political districting, school districting, and sales workforce TD, that are related to geopolitical concerns. Most public services, including hospitals and schools, are managed according to territorial boundaries. We can mention either economic or demographic issues that we might consider in establishing a well-balanced territory.

For excellent reviews on TD and districting models, algorithms, and applications, refer to recent works by Kalcsics et al. (2005), Duque et al. (2007), Zoltners and Sinha (2005), and Ricca et al. (2011). The first two papers survey the most important and relevant applications of districting problems. The last two focus on sales and political districting, respectively.

Although somewhat related, sales districting must locate or deploy a sales workforce to minimize costs, such as an office lease and travel. Commercial districting addresses finding compact territories subject to different planning requirements. In this section, we highlight the most relevant work on commercial TD, which most closely relates to the problem of interest. The first work on commercial TD, Vargas-Suárez et al. (2005), addresses a problem with a variable number of territories; the authors aim to optimize territory balance with respect to three activity measures (i.e., the number of customers, product demand, and workload). They did not consider the compactness criterion; they developed a greedy randomized adaptive search procedure (GRASP), tested it in a few instances, and obtained reasonably good results.

Ríos-Mercado and Fernández (2009) studied a version of the commercial TDP that considers compactness and connectivity, but with no joint or disjoint assignment constraints. Two of the most common methods of measuring territory dispersion result from well-known location problems: the p-center problem (pCP) and the p-median problem (pMP). These location problems require locating a set of p facilities and then assigning customers to facilities to minimize a distance-based dispersion measure. In the pCP, this dispersion function penalizes the largest distance of a unit to its assigned center; in the pMP, the dispersion function considers the sum of the distance of each unit to its assigned center. Ríos-Mercado and Fernández (2009) used the objective function of the pCP for modeling territory dispersion. They proposed and developed a reactive GRASP algorithm for handling relatively large instances of the TDP, evaluated their algorithm on 500- and 1,000-node instances, and obtained good results. More recently, Salazar-Aguilar et al. (2011a) developed an exact optimization scheme for solving the TDP with double-balancing and connectivity constraints. They used their framework to solve problems with both types of dispersion functions: those based on the pCP and those based on the pMP. They observed that models with a pMP objective function had a tighter linear programming (LP) relaxation; therefore, these models were solved faster than the ones using a pCP objective. Furthermore, they also observed that solutions obtained from the relaxation of the pMP-based models had a very high degree of connectivity. Nevertheless, the largest instance they were able to solve for the pMP-based models was approximately 150 BUs. In our approach, we use a similar framework to the one they used in their work; however, we focus on the allocation phase and aim at significantly larger instances. Ríos-Mercado and Salazar-Acosta (2011) present a model for a commercial TDP under a routing budget constraint. More recently, several approaches have been developed for multiobjective versions of the commercial TDP, including both exact optimization approaches (Salazar-Aguilar et al. 2011b) and metaheuristic methods (Salazar-Aguilar et al. 2012, 2013).

Our model has features that extend previous models and that have not yet been addressed in the commercial territory design context, such as the disjoint assignment constraints and a similarity with existing plans. Furthermore, the mathematical structure of our problem differs from that of earlier models, making previous approaches inapplicable for our specific model. For example, location-allocation methods (Kalcsics et al. 2005) have been popular in some districting problems. In particular, these methods seem to give reasonably good performance when the pMP objective function is used as a dispersion measure and single balancing constraints are present. This method relies on relaxing the integrality requirements of the binary variables and then exploiting the substructure of an embedded transportation subproblem. There is an important theoretical result that bounds by $p-1$ the number of fractional variables found in the subproblem. Finally, these fractional variables (corresponding to multiple-center assignments) are resolved by means of a split resolution subproblem. However, in our case, the presence of multiple balancing constraints and connectivity constraints results in the loss of the location-allocation structure and properties.

We note that in our problem, we assume a fixed set of territory centers; that is, we focus on the allocation phase of the process. This set of centers is an input to the model, remaining fixed. By making this assumption, the problem becomes more tractable because the number of binary variables is considerably lower than that of the model with no fixed centers. To the best of our knowledge, no other work in the TDP literature has studied this variation. In the context of the application, the solution to the reduced problem is still clearly a p-partition; however, it depends on the fixed set of centers chosen. Because the chosen set of centers remains fixed, a bad choice of this set (i.e., one in which centers are relatively close to each other) will yield a bad design, that is, a design that might be very unbalanced or have very high dispersion function values. To mitigate this effect, this set of centers must be chosen such that the centers are relatively far apart from each other. We address this issue in our solution methodology.

Territory Design Application at Embotelladoras ARCA

TDP models and solutions are application specific; each has its own constraints and objectives, making the creation of a general-purpose algorithm that can be applied to all types of situations practically impossible. When reviewing the literature, one can observe that only a few papers consider TD problems independently from a concrete practical background. Hence, we do not observe the tendency in OR to separate the model from the application and to establish the model itself as a self-contained topic of research. Therefore, we introduce the TDP in the context of Embotelladoras ARCA and present a solution framework tailored for this particular application.

Our TD problem can be defined as the process of grouping small geographic areas (i.e., BUs) into clusters or territories. Each BU must be contained in exactly one territory. Moreover, we require compactness and connectivity for the territories constructed. A territory is considered to be connected if we can travel between each pair of BUs by following a path contained in the territory. We know the number of territories to be constructed in advance. Each BU has two measurable attributes or activities associated with it: (1) the number of customers and (2) the product demand. For each BU, we know the following information with certainty: location coordinates (from the firm's GIS), number of customers, and product demand as measured by the number of 12-bottle boxes. We compute the activity measure of a given territory by adding up the values of each individual BU that belong to the territory.

To manage its operations more effectively, the company operates by assigning a team consisting of a salesperson and truck drivers to each territory. Each territory is the responsibility of a salesperson who takes the orders from the customers associated with this territory. Similarly, each territory has a group of truck drivers who operate a fleet of vehicles. Therefore, the workload associated with a salesperson is proportional to the number of individual orders he or she must handle, and the workload associated with the truck drivers is proportional to the product demand. As we stated above, each territory must be balanced (or similar in size) with respect to each activity. For example, balancing the number of customers implies a fair work distribution among the salespersons who handle the individual orders. Balancing product demand implies a fair distribution among the truck drivers. We model these balancing constraints by restricting the deviation of the size of each territory from the mean, which we call the ideal target, to within a specified tolerance level. It is noteworthy that only a few authors simultaneously consider more than one criterion for designing balanced territories (Deckro 1977, Zoltners 1979, Zoltners and Sinha 1983).

Additionally, assigning some BUs to the same or different territories might be necessary. For example, we must assign some BUs to a specific territory, which we call joint assignment constraints, and other BUs to different territories, which we call disjoint assignment constraints. These constraints align with either political or strategic decisions. This modeling feature can also be applied to consider geographical obstacles (e.g., rivers and mountains). Finally, the new plan must not differ significantly from an existing plan. Therefore, we consider keeping a subset of BUs in each territory of an existing design in the new design.

Note that we assume deterministic product demand to make the problem more tractable. In the real world, this demand is subject to some variation, which can be seasonal; for example, demand is higher in summer than in winter. However, the demand remains somewhat stable within a given season. Given that TD planning changes every three months, the deterministic assumption for the demand seems reasonable. We developed an MILP model (see the appendix) to solve this problem.

Solution Framework Overview

In this section, we present a solution strategy for solving the allocation model (AM) given in the appendix. One major difficulty is the exponential number of connectivity constraints, as constraint (5) shows, implying that writing them out explicitly is practically impossible. Therefore, we consider instead the AM relaxation (AMR), which consists of relaxing these connectivity constraints. Our method's basic concept is to solve the AMR and check if the solutions obtained satisfy the connectivity constraints. To determine the violated connectivity constraints, we solve a relatively easy separation problem and add these cuts to the AMR. This procedure iterates until no additional connectivity constraints are found; therefore, an optimal solution to the AM is obtained. The solution's optimality is guaranteed because the separation

```
function method()
Input: A problem instance.
Output: A feasible solution X.
1 Solve model AMR and obtain solution X.
2 Identify a set C of violated constraints of model
    AM for solution X.
3 If |C| > 0, add these constraints to model AMR
    and go to Step 1.
4 Return X.
end method
```

Figure 1: This figure shows the steps in the pseudocode of the solution procedure.

problem for identifying violated cuts is solved exactly. Figure 1 depicts a general overview of the method.

In Step 1, we use a branch-and-bound method (because we are not relaxing the integrality requirements of the binary variables). This approach is motivated because the AMR can be solved optimally by current branch-and-bound methods relatively quickly for relatively large instances. For example, 2,000-node instances can be solved in a few seconds of CPU time using a personal computer. In addition, identifying and generating the violated cuts in Step 2 can also be performed in polynomial time, such that the overall procedure may be suitable, provided that the number of iterations needed to reach optimality is not too large. The algorithm gives an optimal solution for the AM.

That we are assuming a fixed set of centers is further exploited to develop several algorithmic strategies for speeding up convergence. Some strategies implemented include (1) variable fixing at preprocessing that allows for identifying and fixing BUs that are far away (close to) territory centers at 0 (1), and (2) strengthening connectivity constraints by adding to the relaxed model (a polynomial number of) some connectivity constraints that prevent forming unconnected territories of size 1 (which are the most commonly found in a disconnected solution from the relaxed model). An accompanying paper (Ríos-Mercado and López-Pérez 2013) fully describes and widely tests all these strategies. One interesting observation is that depending on how these variables are fixed in advance at 0 or 1, a natural trade-off exists between algorithmic efficiency and solution quality. In a preliminary study (Ríos-Mercado and López-Pérez 2013) we empirically show how a significant number of these binary variables can be fixed in advance, thus allowing the method to find near-optimal solutions, and in many cases, global optimal solutions.

Case Study

We implement the solution framework using the X-PRESS MILP solver and language capabilities from FICO. We executed the solution using a personal computer with two Intel core processors at 1.4 GHz and a Windows XP operating system. To assess the proposed method and illustrate its usefulness, we use some real-world instances of 5,000 BUs and 50 territories.

According to our solution procedure, some data are required as inputs to the solution method. Therefore, we performed the following tasks.

- We created a geodatabase layer with the set of points representing the BUs to be clustered into territories. To develop this database, locating all customers was necessary. Salespersons with handheld equipment performed the data collection. They visited all customers in the city of Monterrey (approximately 65,000) and used a GPS device that received latitude and longitude coordinates to mark points representing each customer.
- Using a simple GIS application, we aggregated all these customers into a number of 5,000 BUs. Each BU corresponds to a physical block in Monterrey.
- We developed an infodatabase layer containing the two activity measures (attributes) for each BU. As previously mentioned, these are (1) the number of customers and (2) product demand.
- We set the number of territories the end user requires to construct this study to $p = 50$.

Prior to the development of this tool, the significant amount of territory imbalance with respect to both the number of customers and total product demand was a key issue. The planners constructed the designs by hand, based on experience. Basically, the heuristic solution generated by the company involved (1) making a guess regarding p city blocks to act as territory seeds, and (2) assigning the remaining BUs to its closest seed (i.e., making the compactness criterion play a major role in this operation). Then, by looking at the GIS representing the solution, an operator manually adjusted some bordering BUs by moving them to an adjacent territory if she or he observed an improvement in terms of the balance of the number of customers and demand. Nevertheless, this was a cumbersome process that usually resulted in designs with a very high degree of imbalance among territories (20–25 percent) with respect to both the number of customers and product demand. Of course, this caused dissatisfaction among the teams assigned to different territories. Recall that the ideal target for each activity is the sum of the individual activity values over all BUs divided by the number of territories. With the explicit modeling of the balancing requirement, we can now achieve feasible designs with deviations from the ideal target of within 10 percent (in some cases, within 5 percent), keeping approximately the same level of the territory compactness measure. This model resolved one of the major issues that the company faced.

To illustrate how we improved the existing design, we compare the original plan with the new plan, which the tool we developed generated, in a 5,000-BU, 50-territory instance. Figure 2 shows two charts. The upper chart plots the distribution of the number of customers; the lower chart plots the distribution of product demand for both the old and new designs. The dashed lines depict the 5 percent tolerance limits. The figure shows that the disparity in size among the territories under the previous design was very large with respect to both activities. This disparity contrasts significantly with the new plan's efficiency. Under the new plan, this deviation falls to within 5 percent for all territories.

Figures 3 and 4 display the graphical solution of the previous design and the new design, respectively. Note that the previous design is infeasible with respect to the tolerance level of 5 percent because it has some territories with deviations from the ideal target of up to 25 percent, whereas the design that our solution generated is always feasible under a tolerance level of 5 percent. The legend next to the graph indicates the number of BUs in each territory. We solved the problem in a few minutes using the tool we developed.

Our tool is also useful to assess the natural trade-off between balance and territory compactness. To illustrate this, we solved the 5,000-BU, 50-territory

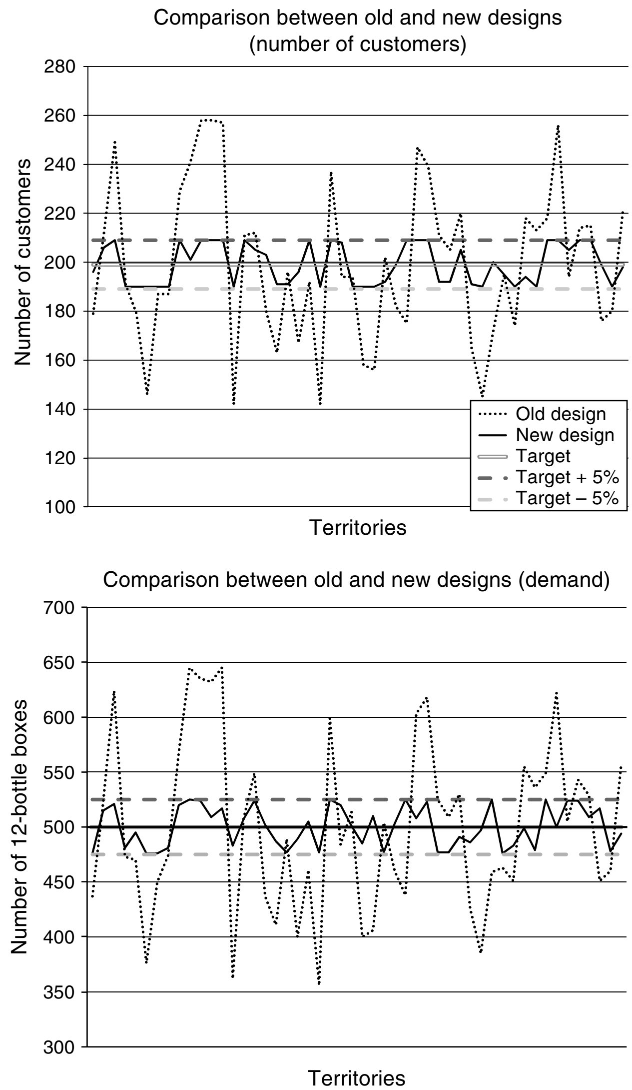

Figure 2: This figure shows a comparison between old and new design in terms of number of customers (top) and product demand (bottom). The charts plot the distribution of the corresponding activity measure for each of the 50 territories. The tolerance limits are depicted as dotted lines.

problem instance for the following values of the user-defined tolerance: 5, 6, 7, 8, 9, and 10 percent (see Figure 5). We can observe that these measures conflict: as a tighter (smaller) value for the balance-constraint allowed tolerance is attempted, the dispersion measure value increases, and vice versa. Therefore, our method becomes a valuable tool for evaluating different solutions in terms of these two factors. When the tolerance changes from 10 percent to 5 percent,

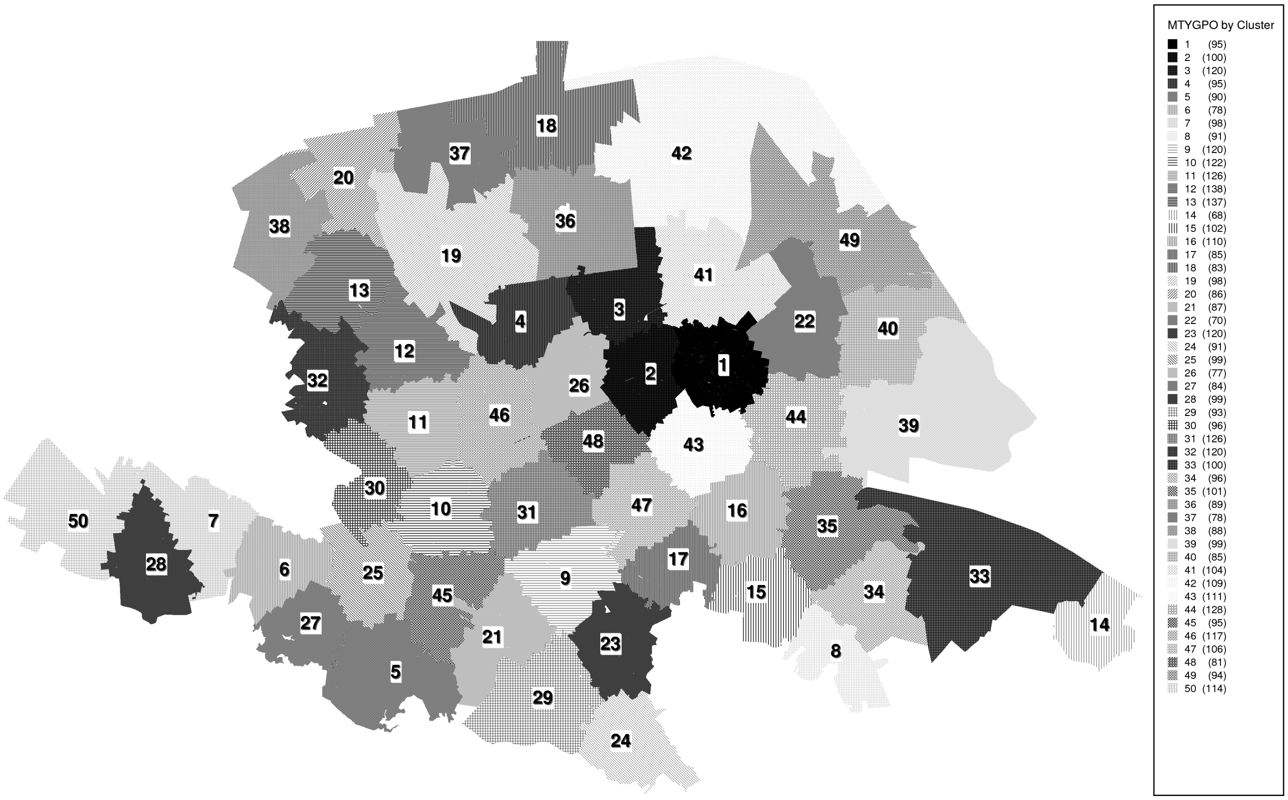

Figure 3: This figure displays the previous territory plan used by Embotelladoras ARCA in Monterrey.

the dispersion increases by approximately 2 percent. In real-world operations, this 2 percent gap means that each worker requires more travel distance (i.e., longer working times).

Implementation Challenges

One of the major challenges was the geographic codification of the customer database. The database maintenance required the addition of specific management and codification modules. At the beginning of the project, we had to assign this task to the presales employees, given the scarce amount of human resources. We also had to involve the human resources department and make an arrangement with the presales employees to add labor compensation. This was a political challenge because of the time and cost required, but we resolved it satisfactorily. The limited functionality of MAPINFO, the company's GIS system, was also a technical issue. This tool was useful for viewing some planning results; however, it did not offer any type of functionality in terms of scenario generation or problem optimization. Thus, we integrated the model into MAPINFO to achieve practical functionality for the end users.

Several technical issues arose during the implementation of the new designs. Initially, the major challenge was to reduce the many territory imbalances. We achieved this by incorporating the balancing constraints into the model and defining an appropriate value for the tolerance parameter. Planners were reluctant to approve major changes because they felt comfortable with the manual designs they had been using for several years, although these designs had 25 percent deviations from the target. However, when the designs that our proposed tool generated showed deviations of 5 percent from the target, we could make a strong case to upper management to use this tool.

The last problem was the issue of realignment. An earlier version of the model did not consider this

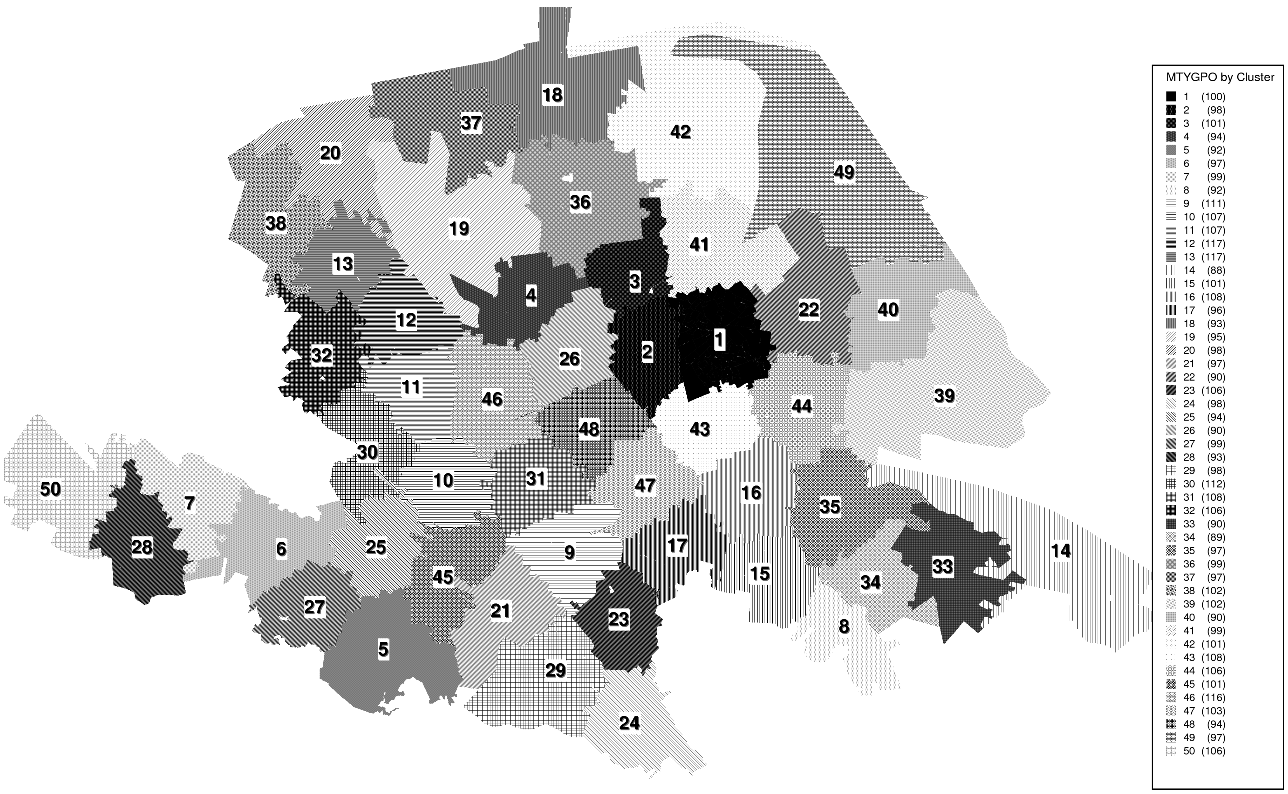

Figure 4: This figure shows the new territory plan in Monterrey (tolerance = 5 percent) after the solution of our developed model.

issue; thus, some of the optimal designs obtained differed greatly from the existing design. From a practical standpoint, drastically changing the original design might have a negative side effect, because

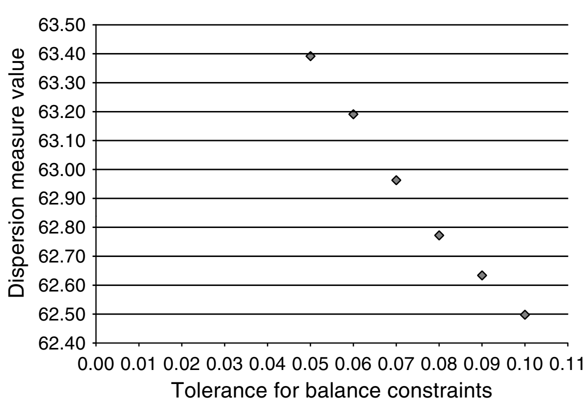

Figure 5: This figure plots the trade-off between the balance constraint tolerance parameter and the dispersion measure value.

many customers have established long-term relationships with the team (salespersons and truck drivers) responsible for serving them. We accommodated this requirement by incorporating a penalty term into the objective to penalize for differences from the existing plan. However, a small number of customers near the territory borders could be switched from one territory to another. To keep these customers happy, the company might offer them some special promotions or discounts on certain products.

Benefits

We integrated our model into an advanced interactive tool based on the MAPINFO application, achieving a practical functionality for end users. This GIS environment can be used in different contexts. At the operational level, it represents a valuable tool to quickly produce and deploy different solutions. At the tactical level, it can be used to simulate alternative scenarios and evaluate the impact of changes in territories. Identifying the interest of the end user with respect to how our model can easily take into account existing territories is important. In particular, we designed the model to consider any prescribed and forbidden assignment of BUs. This means that one can impose some fixed territory centers or BU allocations, which must be taken into account, to territory centers. All these features can be extended for any case in which some territory information is present at the beginning of the planning process. Because the issue of territory realignment is crucial for customer satisfaction, it is an important feature for the company. Thus, the company recognizes how our model efficiently accommodates system changes, such as customer additions or dropouts, and attempts not to significantly disrupt the previous design.

From a business standpoint, we developed and implemented the TDP application at Embotelladoras ARCA to optimize the distribution operation to the end customers. This is the first OR application that it has implemented. The company notes that the overall results have been very positive, and its senior management recognizes the value of the features included in the implemented OR model. The project was a major challenge, requiring a great deal of thought and effort. The first plans for territory design suggested by the optimization model were implemented in mid-2010. Throughout the ramp-up and launch of the project, we analyzed the plans for distribution operation. Subsequently, the project has resulted in a significant increase in productivity and direct savings for the firm. Next, we list some of the benefits that were achieved as a result of this project.

- Identification of a rational set of activity measures to target and balance each truck resource. This results in an optimal fleet of trucks, drivers, and salespeople.
- An increase in efficiency and effectiveness in the planning process is required to set up territory and route designs. The typical fully manual planning process time was reduced from two weeks to less than an hour using the new OR application, permitting the company to dynamically refine its capacity each season. As a result, it has achieved an optimal capacity to satisfy demand in each territory with an optimization of 30 delivery routes in the Monterrey metropolitan area. This represents a 15 percent reduction from the original number of routes.
- Streamlined truck capacity aligned to a new end-customer distribution strategy. The added throughput allows the firm to defer trucks and equipment investments that were originally allocated to the product routing. The investment savings for trucks was approximately 8 percent of the entire fleet.
- Identification and implementation of an optimal cost of service, depending on each route model type. This allowed the firm to set an optimal frequency for customer delivery operations; the benefits were less travel time between customers and a 5 percent increase in product volume delivered per route per day.
- Elimination of territory overlapping. As a result of the connectivity-constrained featured model, we have eliminated territory overlaps, and the territories are now more appropriately defined in geographic terms. Determining which salesperson should be responsible for a new customer (or handle a dropout) is now easy. The company is now able to better define areas of responsibility and loading.
- Improved territory compactness. Because the territories are more compact, the total travel time has decreased, improving the productivity of the distributors. In addition, improving the compactness measurement has allowed managers to reduce the number of trucks available to the distributors.
- Improved balance with respect to the two activity measures. Our model features a small territory tolerance of approximately 5 percent in the lower and upper bounds for the two activity measures. The after-alignment structure is more balanced than the former structure. The standard deviation of the number of customers per territory or the level of workload for each salesperson decreased on average by 24 percent. This alignment allows for an increase in the level of service to the end customers. The company estimates a 3 percent sales increase as a direct benefit of the new territory alignment.
- Value at the strategic level. The approach discussed in this paper focuses mostly on the model's operational aspects; however, the model might also be useful at the strategic level. For example, in this work, we assume a fixed value for the number of territories; however, a sensitivity analysis might be conducted as

a function of the number of territories to assess the effect on the final design in terms of planning criteria, such as compactness and territory imbalance. The results of this analysis might provide a more appropriate value for the number of territories.

In addition to these business benefits, the new OR model will allow the company to speed up some other route-to-market initiatives of special interest among Coca-Cola bottlers worldwide. The proposed model approach can extend the basic model to address other specific business rules or additional planning criteria. Some of these issues can be easily modeled as activity measures on the BUs. Overall, we have provided a valuable tool for a more efficient territory design plan based on the company's business requirements. Our model can manage very large problem instances—instances even larger than 10,000 BUs. Currently, the firm is using our model to obtain a business solution with significant benefits.

One important lesson we learned was that in addition to finding a way to improve existing designs, the implementation of this project helped achieve a closer working relationship among different company departments that were unaccustomed to working together. This benefit, of course, represents an aggregated value for the company.

Final Remarks

In this paper, we have addressed a territory design problem as a critical component of the operational planning process in sales and services companies. Many logistics problems that are found in the service industry can be modeled as a TDP. TDPs are multidisciplinary and have been widely studied in the OR literature. However, because a real-world TDP includes logic and many business rules beyond those addressed in mathematical models in the literature, solving it poses a significant challenge for both researchers and practitioners. In particular, some business rules, such as territory connectivity, are fairly complex to address. We particularly emphasize a business application case at Embotelladoras ARCA. Using a real-world application from the service industry, we present a richly featured TDP model. We include some extensions that are common to some problems encountered in the industry. Because of the characteristics of a TDP, obtaining solutions that are based on concrete business requirements within a reasonable computational time is also challenging. Furthermore, field employees who will deploy the TDP's solution may have to pay more attention to the solution's feasibility in practice than they would to a purely optimal solution in terms of mathematics.

Our TDP instance is motivated by a real-world application in the soft drink industry. In particular, addressing very large-scale problem instances is of interest. We identify and discuss several different objectives and constraints in the TD process. To address these simultaneous and conflicting objectives, we develop an MILP-based solution framework to accommodate the particular business requirements. The proposed framework incorporates some algorithmic strategies that permit solving the problem more efficiently. Our implementation is based on a cut-generation strategy that solves a relaxed model (relaxing the exponential number of connectivity constraints) and then iteratively identifies and adds violated cuts by solving an easy separation problem.

The proposed model addresses both the difficulties embedded in the typical TDP problem and also some practical concerns about predefined and (or) forbidden joint assignments of BUs. Preassigned or forbidden requirements arise from business issues, such as territory realignment. From a practical standpoint, the issue of territory realignment focuses on how the model might efficiently accommodate changes, such as customer additions or dropouts, without significant disruption of the previous design. With respect to our industrial experience and end-user perspectives at Embotelladoras ARCA, we believe that our model can be applied in quite different settings, such as sales territories, locations of new stores in a chain, and delivery areas for distribution. In summary, our model and approach are capable of solving very large-scale real-world TDP instances. Embotelladoras ARCA has successfully used it and has received many benefits as a result.

Appendix. The Model for the ARCA Territory Design Problem

Let

- V be the set of BUs, $|V| = n$;
- E be set of edges representing adjacency between BUs;

- V_c be the set of territory centers, $|V_c| = p$;
- $A = \{1, 2\}$ be the set of BU attributes corresponding to the number of customers ($a = 1$) and product demand ($a = 2$);
- $c(k)$ denote the index of the territory center k;
- d_{ij} be Euclidean distance between BUs i and j;
- w_i^a be the value of attribute $a \in A$ in BU $i \in V$;
- $w^a(V_k) = \sum_{i \in V_k} w_i^a$, the size of territory $V_k \subset V$ with respect to activity $a \in A$;
- $\mu^a = w^a(V)/p$ be the average size (and therefore target value) of activity $a \in A$;
- $N^i = \{j \in V\colon (i, j) \in E\}$ be the set of BUs adjacent to BU $i \in V$;
- H be the set that contains all pairs of BUs that must be assigned to different territories;
- F^i be the prespecified subset of BUs associated with center i from an existing plan;
- q_{ij} be the penalty term for not assigning BU j to center $i \in V_c$, equal to $0.5d_{ij}$ for $j \in F^i$;
- τ^a be the user-specified tolerance parameter for activity $a \in A$;
- α be the user-specified parameter for setting a minimum threshold on the number of BUs that must be assigned from the existing plan.

The decision variables are defined as $x_{ij} = 1$ if the BU j is assigned to a territory with a center in i, and 0 otherwise; $i \in V_c$, $j \in V$. Note that $x_{ii} = 1$ implies that BU i is a territory center.

Allocation Model (AM)

$$\min \left\{ \sum_{\substack{i \in V_c \\ j \in V}} d_{ij} x_{ij} + \sum_{\substack{i \in V_c \\ j \in F^i}} q_{ij}(1 - x_{ij}) = f(x) \right\} \tag{1}$$

$$\text{subject to} \quad \sum_{i \in V_c} x_{ij} = 1 \quad j \in V, \tag{2}$$

$$\sum_{j \in V} w_j^a x_{ij} \le (1 + \tau^a)\mu^a \quad i \in V_c,\ a \in A, \tag{3}$$

$$\sum_{j \in V} w_j^a x_{ij} \ge (1 - \tau^a)\mu^a \quad i \in V_c,\ a \in A, \tag{4}$$

$$\sum_{j \in \cup_{v \in S} N^v \setminus S} x_{ij} - \sum_{j \in S} x_{ij} \ge 1 - |S| \quad i \in V_c, \quad S \subset V \setminus (N^i \cup \{i\}), \tag{5}$$

$$x_{ij} + x_{ih} \le 1 \quad i \in V_c,\ (j, h) \in H, \tag{6}$$

$$\sum_{i \in V_c} \sum_{j \in F^i} x_{ij} \ge \alpha |\cup_i F^i|, \tag{7}$$

$$x_{ii} = 1 \quad i \in V_c, \tag{8}$$

$$x_{ij} \in \{0, 1\} \quad i \in V_c,\ j \in V. \tag{9}$$

Objective (1) incorporates a term that measures territory dispersion and a term that favors the assignment of a subset of BUs from the existing plan. Constraints (2) guarantee that each BU j is assigned to a territory. Constraints (3)–(4) represent the territory balance with respect to each activity measure, because they establish that the size of each territory must lie within a range (measured by tolerance parameter τ^a) of its average size. Constraints (5) guarantee the connectivity of the territories. These constraints, taken from Ríos-Mercado and Fernández (2009), can be explained as follows. Let $S \subset V \setminus (N^i \cup \{i\})$, that is, a subset whose BUs are not adjacent to center i and its adjacent neighbors N^i. Note that if at least one BU j in S, which is not assigned to territory center i, exists, then the second term of the left side becomes strictly less than $|S|$, making the constraint redundant. When all BUs in S are assigned to center i, then the first term on the left side must be greater than 1, that is, at least one BU k surrounding S that must be also be assigned to center i exists. Recursively applying the same rationale to the set $S \cup \{k\}$ results in a territory connected to BU i. Note that an exponential number of such constraints exists. The disjoint assignment is represented by constraints (6). Constraints (7) ensure that at least a minimum number of BUs from the existing plan are assigned, where α is usually set between 0.10 and 0.20 in practice. Constraints (8) fix the territory centers. The choice of the values of parameters q_{ij} and α is not entirely arbitrary. Here we have used values that have shown better results in preliminary work and that are somewhat reasonable in a practical setting.

Allocation Model Relaxation (AMR)

Given the exponential number of connectivity constraints, constraints (5), for our solution procedure, we consider the relaxation of these constraints and call this relaxed model AMR. Note that the integrality constraints are kept.

Acknowledgments

We thank the anonymous referees for the valuable suggestions that improved this work. This research has been supported by the Mexican National Council for Science and Technology [Grants SEP-CONACYT CB05-01-48499-Y and CB11-01-166397] and by Universidad Autónoma de Nuevo León through its Scientific and Technological Research Support Program [Grants UANL-PAICYT CE012-09, CS470-10, and IT511-10].

References

Deckro RF (1977) Multiple objective districting: A general heuristic approach using multiple criteria. *Oper. Res. Quart.* 28(4): 953–961.

Duque JC, Ramos R, Suriñach J (2007) Supervised regionalization methods: A survey. *Internat. Regional Sci. Rev.* 30(3):195–220.

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

Ricca F, Scozzari A, Simeone B (2011) Political districting: From classical models to recent approaches. *4OR* 9(3):223–254.

Ríos-Mercado RZ, Fernández EA (2009) A reactive GRASP for a commercial territory design problem with multiple balancing requirements. *Comput. Oper. Res.* 36(3):755–776.

Ríos-Mercado RZ, López-Pérez JF (2013) Commercial territory design planning with realignment and disjoint assignment requirements. *Omega* 41(3):525–535.

Ríos-Mercado RZ, Salazar-Acosta JC (2011) A GRASP with strategic oscillation for a commercial territory design problem with a routing budget constraint. Batyrshin I, Sidorov G, eds. *Advances in Soft Computing* (Springer, Heidelberg, Germany), 307–318.

Salazar-Aguilar MA, Ríos-Mercado RZ, Cabrera-Ríos M (2011a) New models for commercial territory design. *Networks Spatial Econom.* 11(3):487–507.

Salazar-Aguilar MA, Ríos-Mercado RZ, González-Velarde JL (2011b) A bi-objective programming model for designing compact and balanced territories in commercial districting. *Transportation Res. Part C: Emerging Tech.* 19(5):885–895.

Salazar-Aguilar MA, Ríos-Mercado RZ, González-Velarde JL (2013) GRASP strategies for a bi-objective commercial territory design problem. *J. Heuristics*. Forthcoming.

Salazar-Aguilar MA, Ríos-Mercado RZ, González-Velarde JL, Molina J (2012) Multiobjective scatter search for a commercial territory design problem. *Ann. Oper. Res.* 199(1):343–360.

Vargas-Suárez L, Ríos-Mercado RZ, López F (2005) Usando GRASP para resolver un problema de definición de territorios de atención comercial. Arenas MG, Herrera F, Lozano M, Merelo JJ, Romero G, Sánchez AM, eds. *Proc. IV Spanish Conf. Metaheuristics, Evolutionary Bioinspired Algorithms (MAEB)* (Sociedad Científica Informática de España, Granada, Spain), 609–617.

Zoltners AA (1979) A unified approach to sales territory alignment. Bagozzi R, ed. *Sales Management: New Developments from Behavioral and Decision Model Research* (Marketing Science Institute, Cambridge, MA), 360–376.

Zoltners AA, Sinha P (1983) Sales territory alignment: A review and model. *Management Sci.* 29(11):1237–1256.

Zoltners AA, Sinha P (2005) Sales territory design: Thirty years of modeling and implementation. *Marketing Sci.* 24(3):313–331.

Verification Letter

Jose Luis González García, Marketing Vice-President, ARCA Corporative S. A. de C.V., Monterrey, Mexico, writes:

"I am pleased to write this letter to verify that Embotelladoras Arca, a leading Coca-Cola bottler in Latin America, has been using an Operations Research Model developed in house by a team led by Fábian López, Ph.D., to optimize its product distribution to end customers, with outstanding results. The territory routes designed with the OR model, in use now for over eight months, have yielded substantial benefits:

• "Higher productivity in planning routes. The time involved in the planning process was reduced from two weeks to a single day. This permitted Arca to fine tune its capacity by season, on a dynamic basis.

• "Identify and implement an optimal cost of service depending on route to market model type. This allowed ARCA to set an optimal customer delivery frequency, yielding less travel time between customers and a 5 percent increase in volume delivered per route per day.

• "Streamline truck capacity to align it to our new distribution strategy. The added throughput allowed us to defer investments on trucks and other equipment.

• "Helped speed up other route-to-market initiatives in ARCA, a key need among bottlers.

"Accordingly, I would like to recommend this OR contribution as an excellent case for publication in your journal. With this letter I confirm that Arca has been using the model continuously and profitably since last year [2009]. Feel free to contact us if you require additional information."

J. Fabián López-Pérez is a professor of management science in the Public Accounting and Management School at Universidad Autonoma de Nuevo Leon (UANL), Mexico. He holds a doctorate degree in management science from UANL. His main research interest are logistics, supply chain, vehicle routing, and transportation. He is a member of the Mexican System of Research Scientists. Prior to joining UANL, he worked in the industry for more than 25 years as a supply chain optimization and business analyst, where he conducted applied research and development for consumer products, third party logistics, financial services, retailing, pharmaceutical, petro-chemical, and computer and electronic companies.

Roger Z. Ríos-Mercado is a professor of operations research in the Graduate Program in Systems Engineering at Universidad Autonoma de Nuevo Leon, Mexico. He holds a PhD in operations research from the University of Texas at Austin. He has held visiting scholar positions at the Graduate Program in OR/IE (University of Texas at Austin), Department of Operations Research (Universitat Politecnica de Catalunya, Spain), Leeds School of Business (University of Colorado), and High Performance Computing Center (University of Houston). His research interests are mainly in designing and developing efficient solution methods to hard combinatorial optimization problems. In particular, he has recently addressed applied decision-making problems on territory design systems, forestry management, optimization of natural gas transportation systems, and scheduling in manufacturing systems. His research has been published in leading journals in the field. He is a member of the Mexican Academy of Sciences, and the Mexican System of Research Scientists. More about his work can be found at http://yalma.fime .uanl.mx/~roger/.