CHAPTER 2

DESIGNING ELECTRICAL POWER DISTICTS IN THE REPUBLIC OF GHANA

INTRODUCTION

Deregulation of electrical power industries is actively taking place in the United States and in a number of regions throughout the world. The motivation for this transformation is the generally accepted belief that a monopolistic industry fails to achieve economic efficiency for consumers over the long run. New cost effective technologies allow electricity generation to be "unbundled" from the functions of transmission and distribution and are a driving force behind emerging competitive electricity markets.

Deregulation of the existing electrical power monopolies means that industries are restructuring from a centralized system to a decentralized structure with independently operated business units. This necessitates boundaries of ownership to be established within a common distribution network. Over contractual periods, these boundaries may change or be reorganized to adapt to a developing economy or to incorporate new technologies.

For any electricity service provider in a geographic region, the quality of service (and hence customer satisfaction) is highly dependent upon the ability of the service provider to effectively manage its operation. Thus, an interesting problem that has emerged directly from deregulation of the electrical power industry is the problem of allocating customers located within a geographic region to distinct groups that represent profitable investment opportunities. Recall that we have termed this problem the electrical power districting problem (EPDP).

DISTRICTING PROBLEMS

The research literature contains a variety of mathematical characterizations for the generalized districting problem as well as several suitable application areas. The application areas include, but are not limited to, political redistricting, sales territory alignment, school redistricting, and turf allocation models for the forestry and telecommunications industries. We posit that the EPDP has similar constructs to the above applications and represent a huge open opportunity for researchers interested in districting problems to create substantial social and economic benefit. This section provides a review of the various applications of districting problems covered in the research literature, the manner in which the problem has been characterized, and the various solution techniques applied.

The most prevalent redistricting applications in the management science literature are political redistricting and sales territory alignment. The primary goals of political redistricting problems are to provide geographical districts that are compact and contiguous, yet respectful of existing political units to the maximum extent possible. In addition, the districts must have populations with approximately equal voting potential for the subsequent period. The motivation for using a computerized solution method is to reduce the effects of "Gerrymandering," which may occur when political incumbents bias the redistricting solution to accommodate their political agenda.

Automated political redistricting has been of great interest to politicians and researchers for the last four decades (Vickery 1961, Hess et al. 1965, Nagel 1965, Garfinkel and Nemhauser, 1970, Torricelli and Porter 1979, Backstrom 1982, Browdy 1990, Altman 1995). The first mathematical characterization of the political districting problem was proposed by (Vickery 1961). The redistricting problem has been characterized in the recent literature as a set partitioning problem (Altman 1995), a graph partitioning problem (Browdy 1990, El-Farzi and Mitra 1992), a polygon dissection problem (Hershberger 1991), and an integer programming problem for redrawing congressional districts in the state of South Carolina (Mehrotra et al. 1998).

Designing sales territories is another application area that can be viewed as a districting problem. The sales territory alignment problem is concerned with grouping a number of smaller geographic regions into clusters forming non-overlapping sales territories that span a larger geographic region. Sales territories may need to be realigned whenever changing market conditions warrant, such as the introduction of a new product or variation in sales force size. Sales managers are motivated to carefully design an equitable districting plan, otherwise they risk low morale, poor performance, high turnover rate, and ultimately low productivity within the sales force. Furthermore, a balanced territorial design provides a more consistent and effective means of evaluating individual performance.

There are a number of criteria that can be used to assign geographic regions to districts or sales people to customers. Some single alignment criteria methods seek to balance income or revenue potential (Lodish 1975), while other methods attempt to balance workload or effort (Easingwood 1973). Another common objective is to maximize overall expected profitability. The consideration of multiple objectives was first proposed by (Zoltners 1979).

Sales territory redistricting has a considerable history in the research literature. Similar to the political redistricting problem, a variety of solution techniques have been applied to this problem. A set-partitioning approach was investigated by (Shanker et al. 1975). Various assignment methods have been employed by (Hess and Samuels 1971, Segal and Weinbergey 1977, Zoltners 1979, Richardson 1979). For a thorough review of integer programming approaches, see (Zoltners and Sinha 1980). Heuristic approaches that utilize incremental improvements were proposed by (Easingwood 1973, Lodish 1974, Heschel 1977).

The relationship between a political district and a sales territory is fairly obvious. The political district is driven by the "one man - one vote" principle, which seeks to balance legislative power among districts. When assigning geographic regions to sales territories or salespeople, most equitable solutions seek to balance income potential or effort. Both scenarios can be treated as a bin-packing problem, which seeks to minimize the total deviation of some criteria in each bin (district) from the ideal (global) bin mean. Typically, there are other non-trivial considerations distinguishing a bin-packing problem from a districting problem, such as the compactness of a districting plan or contiguity among the geographic regions allocated to a particular district.

THE ELECTRICAL POWER DISTRICTING PROBLEM

The design of a competitive electricity market is driven by two mitigating factors. First, energy flows along a physical network according to the laws of physics, which requires a coordination of effort in order to achieve balance, reliability, and frequency control. Second, there is presently no economical way to store electricity, which means that it must be delivered in real-time on demand. Some of the major design considerations that result from the above requirements were proposed by (Rassenti and Smith 1998):

1. Coordination of the dispatch and delivery of electricity for a centralized network with decentralized suppliers and buyers.
2. Financial instruments (futures and spot markets with bilateral contracts) which yield appropriate market signals for trading and making long term investment decisions. For details on this issue see (Jamison 2000).
3. Defining divisible property rights for an indivisible common transmission network.
4. Establishing pricing policies such as Zonal or Nodal pricing. For details on this issue see (Schweppe et al.1998).
5. Facilitating competition at the local distribution level.

The EPDP has risen directly out of issue number 3 above. Between the electricity generator (supplier) and the customer (demand) stretches a system of transmission lines that are interconnected to form a physical network. Once electricity is generated, it is transmitted with high voltage over large geographic regions to distribution nodes referred to as bulk supply points (BSPs). From the BSPs, the voltage is reduced and the electricity is distributed to local user groups such as residential neighborhoods, industrial users, or commercial developments. Due to the large costs of equipment and maintenance, transmission and distribution is likely to remain more economical if a large proportion of the power grid is maintained by a single entity in a given geographic area. Thus, the available customer base, which can be represented by BSPs, must be grouped into districts that can be effectively managed in a competitive market.

Similar to political redistricting and sales territory alignment problems, the EPDP is primarily concerned with creating groups of approximately equal potential (profitability). The motivation for this objective is to foster competition by attracting private investment. A good districting plan must also consider the compactness and contiguity of BSPs that comprise a district. Districts that are compact over a geographic region rather than disbursed will be more economical to maintain and thus more profitable. The requirement for contiguity is rather intuitive since product delivery takes place over physical wires connecting the BSPs. From a practical point of view, it is desirable to deliver electricity between BSPs within a common district without paying rents to another transmission enterprise. However, while contiguity appears to be a simple and intuitive requirement, there are a number of hidden complexities that arise when actual power flow is considered.

Perhaps the most difficult and sensitive issue regarding competitive design of the electricity market is establishing a pricing policy. There are two pricing configurations presently being explored in the research literature and implemented in emerging markets throughout the world: nodal pricing and zonal pricing (Chao and Peck 1996, Schweppe et al. 1998). Nodal pricing schemes establish a unit charge for electricity demanded at each BSP. Under this scenario, unit prices of electricity may be based upon historical consumption and delivery costs. Another form of nodal pricing is locational marginal pricing (LMP). LMP is not based upon historical data but driven by the concepts of economic efficiency. The curious reader is referred to (Schweppe et al. 1998, Jamison 2000) for a detailed description of LMP methods. The zonal approach aggregates a number of BSPs into larger zones under the assumption that this would reduce the complexity of the pricing issue. In this case, a number of BSPs share a common unit price.

In our research, we have implemented a nodal pricing scheme based upon historical data provide by the World Bank. The research literature suggests that nodal pricing accounts for "loop-flow" phenomenon more accurately than zonal pricing (Hogan 1997,1998). Furthermore, because the nodal pricing that we implement is based upon historical information, it is less ambiguous and therefore more appealing to the decision makers (DMs) at the World Bank.

A GENERAL MATHEMATICAL FRAMEWORK

A districting plan is a partitioning of units (populations, sales regions, BSPs) into non-overlapping districts (groups) that the are contiguous (adjacent) and geographically compact. A desirable districting plan conforms to one or more objectives such as balanced legislative power or revenue potential. By associating each BSP with a node, and associating each long distance transmission line between BSPs with an edge, the EPDP can be modeled as a graph partitioning problem. This method was applied to the political districting problem by (Altman 1995, Mehrotra et al. 1998) where the weight on the node was equal to the corresponding population size. For the EPDP, we begin by letting the weight on the node correspond to an aggregate measure of revenue potential for each BSP. To obtain *k* transmission districts of approximately equal earning potential we can minimize the total deviation of the revenue in each district from a target value. The model is as follows:

Minimize: $$\sum_{j=1}^{k} \left| R_j - \bar{r} \right| \tag{2.1}$$

Where: $$\bar{r} = \frac{1}{k} \sum_{i=1}^{n} r_i \tag{2.2}$$

$$R_j = \sum_{i=1}^{n} r_i X_{ij} \quad , \; j=1\ldots k \tag{2.3}$$

$$\left| R_j - \bar{r} \right| \le \delta \bar{r} \quad , \; j=1\ldots k \tag{2.4}$$

$$X_{ij} = \begin{cases} 1, x_i = j \\ 0, o.w. \end{cases} \quad , \; i=1\ldots n \quad , \; j=1\ldots k \tag{2.5}$$

$$m_j = \sum_{i=1}^{n} X_{ij} \quad , \; j=1\ldots k \qquad (2.6)$$

$x = (x_1 , x_2 ,\ldots, x_n)$
x_i= district assignment for node i
R_j = sum total of revenue potential in district j
k = number of districts (partitions)
r_i = amount of revenue potential contained in node i
n = number of nodes in the network
$100\delta, (0 \le \delta \le 1)$, *is the maximum allowable percentage deviation of the actual revenue in a district from the target*
m_j = number of nodes in district j

Another criterion that can be applied to locate an equitable solution is the mini-max objective function. This criteria finds the district that is maximally separated from the others, with regard to the target revenue potential, and minimizes it. The mini-max objective is sometimes referred to as the "bottle neck" objective function, which is described in detail in (Edmonds and Fulkerson 1968).

Minimize:

$$\underset{j=1}{\overset{k}{Max}} \left| R_j - \bar{r} \right| \qquad (2.7)$$

Let us define *G (N, E)* to be a graph *G* with nodes *N* defined to be the set of BSPs that comprise *G*, and edges *E* defined to be a pair-wise connection matrix corresponding to the set of long distance transmission lines that connect *N*. A district is a node induced subgraph *G' (N', E')* of *G (N, E)* that is contiguous and conforms to the measure of economic potential defined in (2.4). A district is also referred to as a partition or a group belonging to *G (N, E)* in later sections. See Figure 2.1.

FIGURE 2.1
A Graph Representation of an Electrical Transmission and Distribution Network

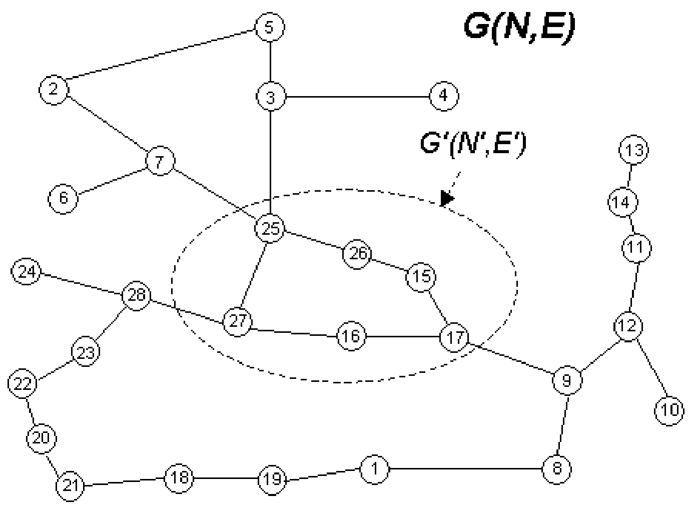

A district G' is contiguous if all nodes N' assigned to the district are connected by edges E'. Contiguity also implies that it is possible to reach any node in a particular district from any other node assigned to the same district without leaving the district.

As an alternative (or additional) objective, we can seek districts that are compact. To measure the compactness of a district j, we use the total euclidean distance of each node assigned to a district from the district centroid $C_x(j)$ and $C_y(j)$. Thus, for each node i in G' there exists coordinates x and y defining its location $N_x(i)$ and $N_y(i)$. The compactness of a district j, K_j is provided in (2.9) and the objective would be to minimize the function provided in (2.8) to yield an overall districting plan that is compact.

Minimize:
$$K = \sum_{j=1}^{k} K_j \tag{2.8}$$

Where:
$$K_j = \sum_{i=1}^{n}\left(\left(C_x(j)-N_x(i)\right)^2+\left(C_y(j)-N_y(i)\right)^2\right)^{1/2} X_{ij} \quad , \; j=1\ldots k \tag{2.9}$$

$$C_x(j) = \frac{1}{m_j}\sum_{i=1}^{n} N_x(i) X_{ij} \quad , \; j=1\ldots k \tag{2.10}$$

$$C_y(j) = \frac{1}{m_j}\sum_{i=1}^{n} N_y(i) X_{ij} \quad , \; j=1\ldots k \tag{2.11}$$

$$X_{ij} = \begin{cases} 1, x_i = j \\ 0, o.w. \end{cases} \quad , \; i=1\ldots n \quad , \; j=1\ldots k \tag{2.12}$$

$$m_j = \sum_{i=1}^{n} X_{ij} \quad , \; j=1\ldots k \tag{2.13}$$

K_j =*Compactness for district j*
k = number of districts (partitions)
$C_x(j)$= *x-Coordinate for the Centroid of district j*
$C_y(j)$= *y-Coordinate for the Centroid of district j*
N_x *(i)= x-Coordinate for node i*
N_y *(i)= y-Coordinate for node i*
n = number of nodes in the network
m_j = *number of nodes in district j*

We could also measure compactness as the "farthest neighbor" in a district. The farthest neighbor is the maximum pairwise euclidean distance between nodes assigned to a district. This is a similar extension of the "bottle neck" objective function described above.

SOLVING THE DISTRICTING PROBLEM

There are two broad classes of solution techniques that have been applied to the districting problem, exact methods and approximate methods. Exact methods search (implicitly or explicitly) every point in the solution space and thus guarantee optimality of the best solution. Due to their computational intensity, their usefulness is restricted to solving only small districting problems.

The most common solution method applied to this class of problems involves approximate methods, which often use heuristics. Heuristics can be broadly defined as "rules of thumb" or "educated guesses." Quite often, heuristics can be designed that consistently solve intractable problems to within an acceptable neighborhood of the optimal solution. In exchange for their efficiency, heuristics sacrifice the guarantee that the best solution found is in fact an optimal solution.

While exact techniques have been applied to districting problems in the past, we believe that most effective solution methods for the EPDP will make use of heuristics. For a comprehensive review of exact methods that have been applied to districting problem, the reader is referred to (Garfinkel and Nemhauser 1970, Zoltners and Sinha 1980). The heuristic solution methods that have been investigated for solving districting problems include simulated annealing (Browdy 1990), branch-and-price heuristic strategies (Mehrotra et al. 1998, Barnhart et al. 2000), hill climbing (Moshman and Kokiko 1973, Nagel 1965) and a divide-and-conquer heuristic (Hershberger 1991).

Simulated Annealing

In this section, we provide a brief overview of the Simulated Annealing (SA) algorithm. The SA algorithm is a general purpose optimization procedure based upon the thermodynamic process of annealing metals by slow cooling (Kirkpatrick et al. 1983, Kirkpatrick 1984). At high temperatures molecules in metal move rapidly with respect to one another. If the metal is cooled sufficiently slow then thermal mobility is lost. The resulting configuration of atoms are aligned to form a pure crystal that is completely ordered. This ordered state occurs when the system has achieved minimum energy.

To achieve low energy configurations, the annealing process must be cooled sufficiently slow to reach thermal equilibrium. The compelling reason for this is that the energy of a system (in thermal equilibrium) is probabilistically distributed according to the Boltzmann distribution. At high temperatures, the algorithm visits a very large neighborhood of the current state. As cooling takes place, the system accepts high and low energy states, but the lower the temperature, the lower the likelihood of accepting an uphill move (higher energy states). Therefore, the transitions to higher energy states become less frequent, and the solution stabilizes.

The SA algorithm starts at some high temperature T_c and in some state s. A sequence of points is then generated using a neighborhood operator to find some other configuration t. The probability of moving from state s to a neighboring state t is a function of the level of energy, $E(s)$ and $E(t)$, of each state. The state transition can be characterized by two conditions: First, if $E(t) - E(s) < 0$ then the new configuration is accepted deterministically because it represents a desirable reduction in energy level. Second, if $E(t) - E(s) \geq 0$, indicating an uphill move to a higher energy state, the probability that state t is accepted is provided in (2.14).

$$\text{P(accept state } t\text{)} = e^{-\frac{E(t)-E(s)}{T_c}} \tag{2.14}$$

Thus, a cooling schedule for T_c is also necessary. The cooling schedule consists of the sequence of temperature changes and the amount of time spent at each temperature. If the annealing process reaches thermal equilibrium at each temperature, it can be proven that it produces the global minimum energy at absolute zero T_z (Metropolis et al. 1953). An appropriate cooling schedule will conform to expressions (2.15) and (2.16).

Cooling Schedule: $T_1 > T_2 > \ldots T_c \ldots > T_z$ (2.15)

Where: $\lim_{c \to z} T_z = 0$. (2.16)

The process is again repeated until equilibrium is reached for each temperature increment or no more useful improvements can be expected (the analogue to thermal equilibrium). Neighboring states are continually visited until thermal equilibrium is approached for each temperature using the following generalized procedure:

SA Procedure:

Select cooling schedule $T_1 > T_2 > \ldots T_c \ldots > T_z$
Create initial solution in state s and evaluate
c = 1 (counter)

Do

Do

Create neighbor in state t and evaluate

$\Delta E = E(t) - E(s)$

If $\Delta E < 0$ *then*

Set s = t

Else if $e^{-\frac{\Delta E}{Tc}}$ *> Random [0,1] Then*

Set s = t

End if

Loop Until thermal equilibrium achieved at T_c

c = c+1

Loop Until $T_c \leq T_z$

End Procedure

An inherent difficulty of implementing the SA algorithm lies in determining an appropriate cooling schedule such that we can determine when thermal equilibrium has been reached. A pragmatic way of dealing with this is to introduce control parameters such as the acceptance frequency, acceptance probability limit p_a and cooling factor *d*. By monitoring the rate at which new neighbors are accepted or rejected for any given set of iterations at temperature T_c, we can determine the probability of finding an improving solution on the next step. When further improvements appear to be unlikely, say the probability drops below 1/100 (thus, we use a moving average with a window *W* of size equal to 100 iterations), the annealing process is deemed to be in equilibrium for temperature T_c. The cooling factor d $(0 < d < 1)$, is a multiplier used to reduce the annealing temperature to T_{c+1} for the next iteration.

$$T_{c+1} = d\,T_c \tag{2.17}$$

The current solution in thermal equilibrium at T_c is slightly out of equilibrium at T_{c+1}, so the annealing process must be repeated again for T_{c+1}. This method of developing a generic, yet flexible cooling schedule is popularly supported in the research literature (Kirkpatrick et al. 1983, Johnson et al.1989). A similar configuration of control parameters was implemented in the SA algorithm on graph partitioning problems by (Johnson et al.1989).

In order to implement the SA algorithm, it is also necessary to develop an analogue of the concepts in the physical system (state and energy) to the properties of the problem at hand. For the EPDP, which we have characterized as a graph partitioning problem, consider the follow the analogies:

state: a feasible partition s *= Union of* k *sub-graphs* $G' (N', E') \in G (N, E)$

energy: the cost of the solution $E(s)= f(x)$

We can now see that the SA algorithm is simulating the thermal history of an analogous material that is being subjected to a cooling schedule. Instead of the energy of a real material, SA uses the objective function of an optimization problem. The simulated thermal perturbations are changes in the decision variables rather than crystalline structures. If the process achieves thermal equilibrium at each temperature, then the objective function reaches its global minimum when the simulated temperature reaches approximately zero. By using the analogies and the control parameters described above, we develop a SA procedure for the EPDP. For a detailed description of our SA procedure and the operators that we have tailored to the EPDP, the reader is referred to Appendix A.

Evolutionary Algorithms

One of the most interesting subclasses of heuristic search techniques involves evolutionary algorithms (EAs). EAs are a mathematical modeling paradigm inspired by Darwin's theory of evolution. An EA adapts during the search process, using the information discovered to break the curse of dimensionality that makes nonrandom and exhaustive search methods computationally intractable.

Over the past three decades, three primary evolutionary-based problem solving techniques have emerged: the Genetic Algorithm (GA), the Evolution Strategy (ES) and Genetic Programming (GP). For the most part, the techniques remain isolated in their methodology and preferred domain representations. The GA and GP have targeted mostly combinatorial optimization problems while the ES has emphasized optimization in the continuous domain. The ES was developed in Germany (Rechenberg 1973) while the GA and GP were developed in the United States (Holland 1975, Koza 1990). For an excellent comparison of the three methods the reader is referred to (Back and Schwefel 1993). The remainder of this section develops the framework and motivation for applying the GA to the EPDP.

The Genetic Algorithm

The GA was first introduced at the University of Michigan by (Bagley 1967). However, it was not until Holland published the original schemata theorem, that the GA technique began to take hold in the research community (Holland 1975). Holland's GA contained three operators: crossover, mutation, and selection. His GA also represented the problem domain as a binary string of numbers called a chromosome, where each binary digit within the chromosome represented a gene. The GA model consisted of a collection of chromosomes that formed a gene pool (population) of potential solutions to a problem that were initially scattered throughout the search space. Using a process modeled after the biological theory of evolution, Holland's GA compelled the gene pool to remain diverse yet continuously improve, increasing the likelihood that the population would "converge" upon a global rather than a local optimal solution.

Holland's genetic operators formed the foundation for subsequent GAs. Although many variants of the original operators exist, their original purpose remains intact for most implementations, and can be described as follows:

1. Crossover. Crossover is the process that mimics reproduction in biological organisms. Crossover is implemented in a GA by exchanging chromosome segments between two or more parent chromosomes to form a child chromosome. The crossover operator typically serves a dual purpose. First, to effectively reduce the search space to regions of greater promise. Second, to provide a mechanism for allowing offspring to inherit the properties of their parents. The crossover process is also commonly referred to as "Recombination."
2. Mutation. Mutation is the process that mimics the unpredictable and unexpected developments that occur in biological reproduction. In a GA, mutation is a random perturbation to one or more genes that occurs infrequently during the evolutionary process. The purpose of the mutation operator in a GA is to provide a mechanism to help stave off misconvergence due to a homogeneous population.
3. Selection. Selection is the process that mimics the "survival of the fittest" principle in the biological theory of evolution. Selection is implemented in a GA as a policy for determining which chromosomes in the population will survive and be carried over into the next generation. Selection is typically the only operator that considers the fitness of the population members in a GA.

Given the extensive applications of the GA in the field of combinatorial optimization, it is surprising that our literature review uncovered only a single application of GAs to districting problems (Ding et al. 1992). However, there have been several applications where GAs have proven to be effective on various partitioning problems (Jones and Beltramo 1991, Hohn and Reeves, 1996). An extensive study of hybrid GAs for closely related bin-packing and grouping problems has also been performed by (Falkenauer and Delechambre 1992, Falkenauer 1994).

In Holland's original GA, the genetic operators were effective on the binary representation of the problem in the chromosome string. The crossover operation had a "constructive" effect upon the production of offspring, because the transfer of partial schema from parents to children preserved the information discovered about the problem at hand.   For example:

P1: 10010010101110101010

C1: 10010010110000010111

P2: 11101011110000010111

The recombination of Parents 1 & 2 resulted in a Child with schema representative of both parents. This is the underlying concept behind the "Building Block Hypothesis" of the schemata theorem (Holland 1975). The building blocks refer to sequential groups of binary values that are exchanged in small segments to propagate the traits carried by a particular gene. As recombination takes place new genes emerge that are tested against their environment (ie. the fitness function). The first attempts at using GAs to solve problems involving grouping or partitioning implemented the classical operators described above (Jones and Beltramo 1991, Van Driessche and Piessens 1992, Ding et al. 1992). For partitioning problems, a direct translation of the classic genetic operators has a "disruptive" effect and requires special considerations in designing grouping (or partitioning) GAs (Falkenauer 1994). For an extensive review of GA design concepts and a detailed description of the genetic operators that we have tailored to the EPDP, the reader is referred to Appendix B.

CREATING POWER DISTRICTS IN THE REPUBLIC OF GHANA

The Republic of Ghana, situated on the Gold Coast of Africa, is one of the most developed African countries south of the Sahara. Like all existing electricity infrastructures in the world today, the Ghana power sector began as a non-competitive national monopoly. With the objective of developing an economically stable, customer oriented industrial culture in Ghana, the Government of Ghana (GoG) is undertaking an industry reform program for their emerging competitive electricity market. The GoG has enlisted the services of the World Bank in hopes of bringing their objectives to fruition. To this end, the authors are engaged in a joint research effort with the World Bank to develop a solution methodology and a Decision Support System (DSS) to assist with restructuring the electric power sector in Ghana, with the intentions of applying this methodology to other countries. Specifically, the focus of this research is to develop a restructuring plan comprised of districts (groupings) such that independent distribution enterprises may operate as reliable and economically viable electricity service providers.

Under the present system, there are two monopoly organizations installed as the national service providers of electricity in the Republic of Ghana, the Electricity Corporation of Ghana (ECG) and the Northern Electricity Department (NED). Together, these organizations share the ten geographic regions that span Ghana. The ECG is the bigger and older of the two with six regions, while the remaining four are cared for by the NED.

FIGURE 2.2

The National Electricity Power Grid of Ghana

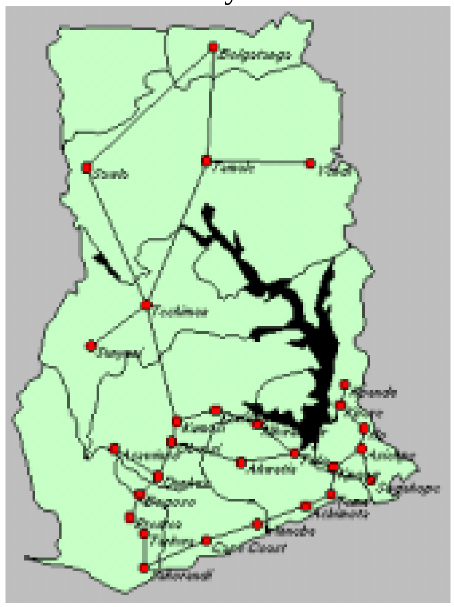

The national electricity transmission grid of Ghana is provided in Figure 2.2. The grid consists of high voltage transmission connections and 28 BSPs. Recall that BSPs represent bulk distribution junctions points where high voltage electricity is reduced for local distribution services. Table 2.1 provides a legend of the BSPs included in this study. The information consists of the index number and name corresponding to each BSP shown in Figure 2.2. Also, the expected annual revenue potential is provided for each BSP in millions of dollars along with the standardized physical coordinates for a two dimensional map.

TABLE 2.1

BSP Information

Index	BSP	Exp Rev($M)	X	Y
1	Achimota	$ 30.518	21	8
2	Sawla	$ 0.006	6	34.5
3	Tamale	$ 1.943	17	35
4	Yendi	$ 0.007	23	35
5	Bolgatanga	$ 0.873	17	45
6	Sunyani	$ 1.978	6	20.5
7	Techiman	$ 1.127	8.5	23
8	Tema	$ 12.369	23	9
9	Kpong	$ 2.511	23	12
10	Sogakope	$ 0.294	27	11.5
11	Ho	$ 0.403	26	15
12	Asiekpe	$ 0.010	26.5	13
13	Kpandu	$ 0.616	25	18
14	Kpeve	$ 0.208	25	16
15	Nkwankaw	$ 1.585	7.5	20
16	Akwatia	$ 0.571	17	11

17	Tafo	$	0.985	19.3	13
18	Cape Coast	$	1.654	14	5
19	Winneba	$	1.053	18	6.5
20	Tarkwa	$	3.104	9	6
21	Takorandi	$	5.112	10	3
22	Prestea	$	1.610	8	7
23	Bogoso	$	1.479	8.5	8
24	Asawinso	$	0.623	6.5	13
25	Kumasi	$	11.284	11	16
26	Konongo	$	0.434	14	15
27	Obuasi	$	1.111	11	12.5
28	Dunkwa	$	0.163	10	11.5

Recall that *G (N, E)* is a graph *G* with nodes *N* defined to be the set of BSPs that comprise *G,* and edges *E* defined to be a pair-wise connection matrix corresponding to the set of long distance transmission lines that connect *N*. The graph representation of Ghana's electricity network in Figure 2.3 is simply a logical representation of the physical transmission system.

FIGURE 2.3
Graph of Ghana's Electricity Grid

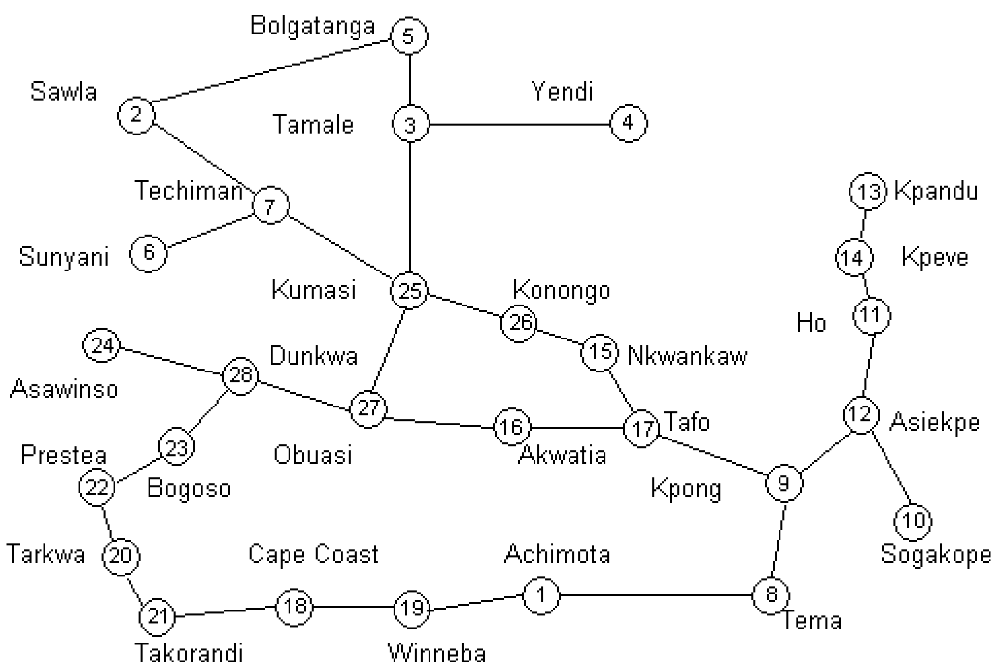

A district is a node induced subgraph *G' (N', E')* of *G (N, E)* that is contiguous and conforms to a specified range of earning potential (2.4). To obtain transmission districts of approximately equal revenue potential, we minimize the total absolute deviation of the revenue in each district from the ideal district revenue. Thus, the weight on each node is calculated by finding the product of the unit price and the total yearly demand for each BSP. The district revenue is the sum of expected revenues across BSPs assigned to the district. Figure 2.4 indicates a feasible partitioning of *G (N, E).*

FIGURE 2.4
Feasible solution with 5 electrical power districts

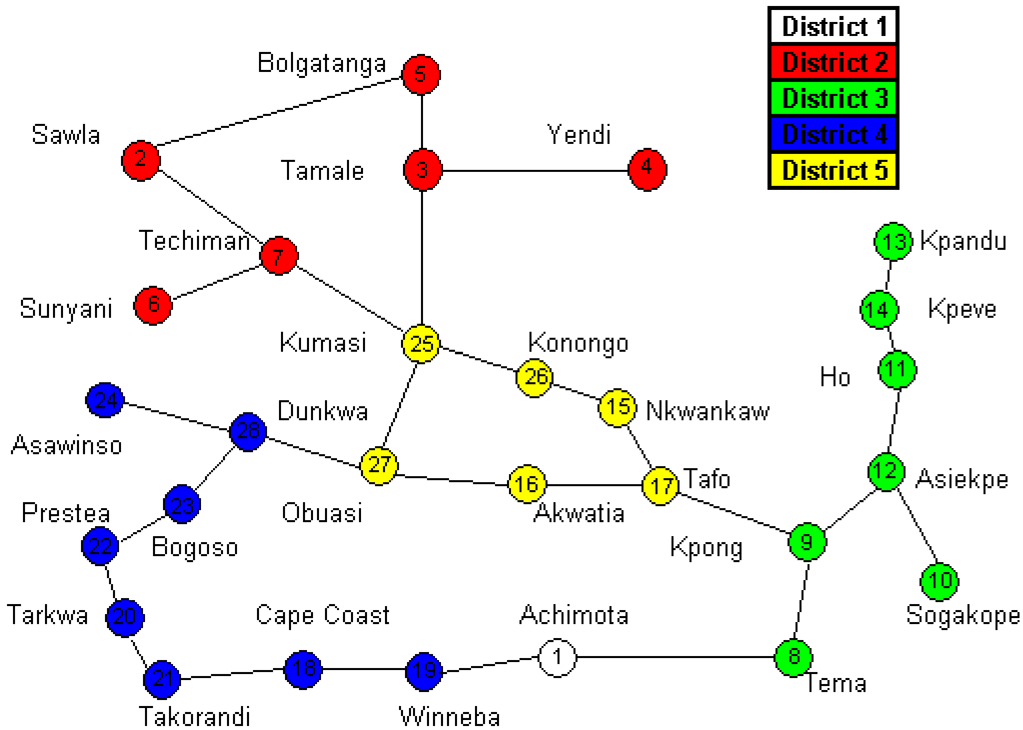

Table 2.2 indicates the summary information for each district corresponding to the partition provided in Figure 2.4.

TABLE 2.2
Summary Information by District

Districts	1	2	3	4	5	Average	Total
Revenue	30.52	5.93	16.41	14.80	15.97	16.73	83.63
Deviation	13.79	10.79	0.31	1.93	0.76	5.52	27.58

EMPIRICAL STUDY – SINGLE CRITERIA

The purpose of this empirical study is to examine and compare the behavior of the GA and SA for solving the EPDP using Ghana's current transmission network. The empirical study in this chapter is limited to an analysis of the single criteria version of the EPDP for the total absolute revenue deviation function (f_1) and total Euclidean distance function (f_2). While a more general study that is independent of Ghana is undertaken in the subsequent chapter.

In the design of this study, it is fundamental to consider the range of k that is expected for solving an EPDP. The EPDP is manifested when national electric utilities make the initial transition from a monopoly, where k=1 to an oligopoly where k>1. No explicit number of firms is required for oligopoly conditions, but the number is usually somewhere between 2 and 10 (Baye and Beil 1994), depending upon the specific industry. Based upon our conversations with DMs at the World Bank, we believe that a reasonable range for k for the electric utility industry is somewhere between $4 \leq k \leq 6$. It is crucial to keep this information in mind as we make decisions for investigating the performance of solution algorithms on f_1 and f_2.

A noteworthy characteristic that we observed about f_1 is that it has a propensity to produce multiple district configurations (alternate solutions) having the same measure of overall absolute revenue deviation. It is also significant to note that this tendency is exacerbated as the number of desired districts in the districting plan is increased. In order to understand this anomaly, consider any arbitrary solution with n nodes assigned to k districts. If a node is removed from a district where the total district revenue is above (or below) the ideal target value and reassigned to another district that is also above (or below) the ideal target value, then the net effective change to f_1 is zero. As the number of districts in the plan is increased one can see that the number of scenarios that will result in multiple alternate solutions increases because more districts are present. Thus f_1 has "flat spots" in the response surface. This function characteristic can pose considerable difficulty for solution techniques that rely upon a gradient to determine search directions. Algorithms that invoke a "point-based" search as opposed to a "population-based" search are likely to be adversely affected by this anomaly.

In contrast, f_2 exhibits no such characteristic. While it is possible for multiple alternate solutions to exist for this measure, it is highly unlikely since the distance measure is continuous. Thus, it is our expectation that a point based solution technique will not be hindered by f_2 as it would be for f_1. Similarly, we expect a population based search to perform better on f_1.

We believe that understanding the final outcome of each procedure is important to our study. However, we also believe that it is equally important (perhaps even more important) to understand and compare the behavior of each method during the initial phase of an optimization run. Our rationale for this comparison is based upon the fact that we wish to determine if one algorithm is better suited than the other for our DSS. Since we anticipate that the DSS will be used in the context of a negotiation involving multiple DMs, the final solution will likely depend upon the collective utility of the parties involved. It is very likely that the final (unique) solution preferred by the parties will be sub-optimal with regard to the explicit cost functions in the model.

The rationale for this is that the agreed upon solution will implicitly incorporate the utility functions of the DMs as well as other criteria, such as the "clarity of demarcation." The term, "clarity of demarcation" can be interpreted as the extent to which a districting plan is respectful of existing political boundaries and natural geographic barriers. This criteria is extremely difficult to model explicitly without extensive utility assessment. Thus, it is the goal of our DSS to be able to produce "very good" solutions quickly and efficiently for the specified criteria in order to maximize the value added to the DMs. The DSS is designed to support an interactive search to accomplish this difficult task.

The empirical study is organized into two sections. The first section addresses the dynamic behavior of the GA and SA over the initial 1,000 function evaluations for 30 independent optimization runs. The analysis is dynamic in the sense that it is based upon the temporal performance of each method. The second section is a static analysis of the same set of optimization runs. The analysis is static because it is based upon the final outcomes of the optimization runs. We conclude with some justification for our approach.

Parameter Settings

The parameter settings used to obtain the results in this study were determined via preliminary experiments. We readily admit that the parameter settings of both algorithms can have a profound affect on the results of the study. However, the settings that we implement were found by expending a reasonable amount of computational effort and a practical amount of time.

For the GA, the initial population size was set to 25 chromosomes and remained constant throughout performance testing. Every member in the population was recombined in each generation to produce a single offspring. Following recombination, each child chromosome was mutated with a neighborhood search as described in Appendix B. The GA optimization runs were terminated if no improvements in the population were found in 1,000 function evaluations or a maximum of 20,000 function evaluations – whichever came first.

For the SA procedure, the initial temperature of system was set to 5.0^o and the final cut-off temperature was set to 0.1^o. The cooling factor used to lower the system temperature was 0.95 and invoked when no improvements were found in the prior 100 function evaluations. In similar fashion to the GA, the run was terminated at a maximum of 20,000 function evaluations or whenever the system cut-off temperature was reached – whichever came first.

Dynamic Analysis

In this section, we analyze the difference between the average performance of the GA and the average performance of SA on f_1 and f_2 with respect to the number of function evaluations. We test for statistically significant differences between the population means of the GA and SA in increments of 100 function evaluations from the start of the optimization run and extending to 1,000 function evaluations. An identical test is performed for each setting of $4 \leq k \leq 6$ that we determined was relevant to our investigation. It is important to point out that the method used to produce initial starting solutions was identical for both techniques. Thus, for each case examined herein, we note that difference in starting solution values for the GA and SA, for any setting of k districts, was not statistically significant.

The "No Free Lunch Theorem" asserts that all algorithms that search for the extremum of a cost function perform exactly the same, according to any performance measure, when averaged over all possible cost functions (Wolpert and Macready 1995). In search and optimization, the fundamental question is, "How do I find good solutions for a given cost function?" To answer this question, we must first determine which solution method matches our cost function best under the conditions for which it is intended to be used.

Revenue Deviation

Figure 2.5 shows the initial convergence behavior of the GA and SA for the case of k=4 districts. The curves shown in Figure 2.5 correspond to the data provided in Table 2.3. The data in Table 2.3 indicate that there is a statistically significant difference between the average performance of the GA and the average performance of SA at the α=.005 level across the entire range of the initial 1,000 function evaluations.

FIGURE 2.5

Revenue Deviation – 4 Districts

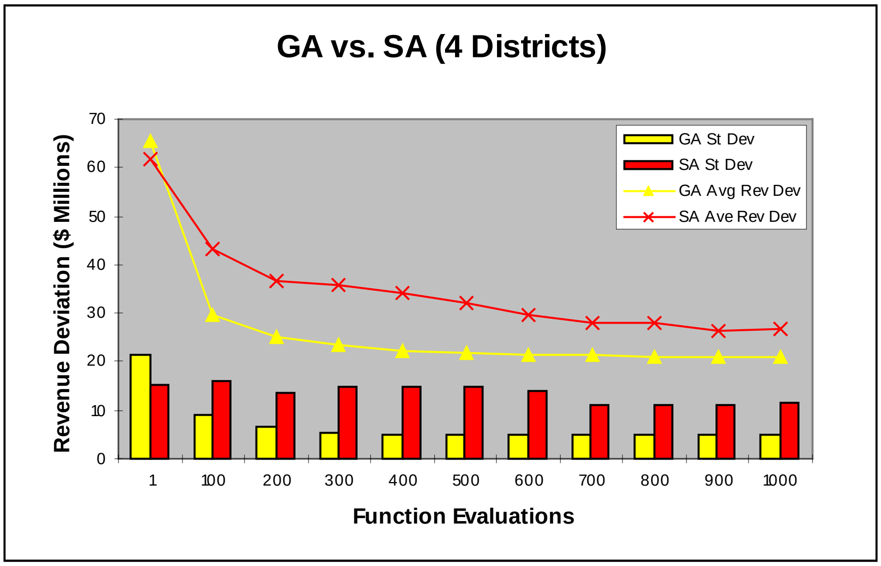

TABLE 2.3

Dynamic Analysis of Results for Revenue Deviation – 4 Districts

Function Evaluations	Genetic Algorithm Average	Genetic Algorithm Standard Deviation	Simulated Annealing Average	Simulated Annealing Standard Deviation	Difference in Means	Standard Error	t-values
1	65.515	21.254	61.572	15.169	3.943	4.767	0.827
100	29.796	8.858	43.143	16.261	-13.347	3.381	-3.948*
200	25.295	6.662	36.639	13.654	-11.343	2.774	-4.090*
300	23.313	5.300	35.979	14.713	-12.666	2.855	-4.436*
400	22.140	4.927	34.260	15.011	-12.121	2.884	-4.202*
500	21.723	4.822	32.044	14.618	-10.321	2.810	-3.672*
600	21.548	4.832	29.725	14.048	-8.177	2.712	-3.015*
700	21.348	4.816	28.113	11.316	-6.766	2.245	-3.013*
800	21.120	4.781	28.062	11.283	-6.942	2.237	-3.103*
900	21.120	4.781	26.519	10.917	-5.399	2.176	-2.481*
1000	21.035	4.798	26.905	11.369	-5.870	2.253	-2.605*

* t-values significant at α=.005

Figure 2.6 shows the initial convergence behavior of the GA and SA for the case of k=5 districts. The curves shown in Figure 2.6 correspond to the data provided in Table 2.4. The data indicate that there is a statistically significant difference between the average performance of the

GA and the average performance of SA at the α=.005 level for every increment up to and including 900 function evaluations. When the algorithms reach 1,000 function evaluations, the performance difference is no longer statistically significant.

FIGURE 2.6
Revenue Deviation – 5 Districts

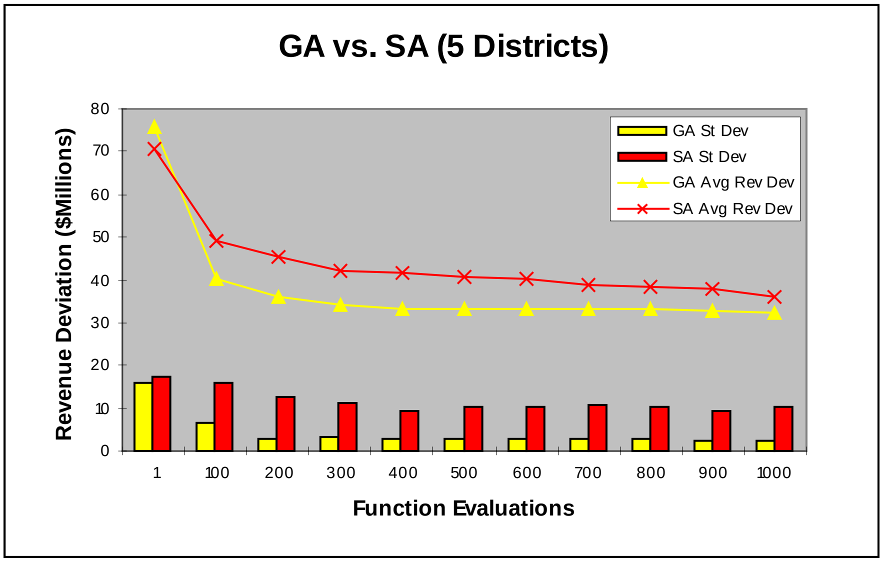

TABLE 2.4
Dynamic Analysis of Results for Revenue Deviation – 5 Districts

Function Evaluations	Genetic Algorithm		Simulated Annealing		Difference in Means	Standard Error	t-values
	Average	Standard Deviation	Average	Standard Deviation			
1	75.818	15.683	70.635	17.297	5.184	4.263	1.216
100	40.342	6.748	49.281	15.878	-8.940	3.150	-2.838*
200	35.799	2.959	45.304	12.536	-9.506	2.352	-4.042*
300	34.370	3.234	42.095	11.087	-7.725	2.108	-3.664*
400	33.443	2.799	41.513	9.194	-8.070	1.755	-4.599*
500	33.443	2.799	40.836	10.289	-7.393	1.947	-3.798*
600	33.265	2.777	40.264	10.320	-6.999	1.951	-3.587*
700	33.189	2.834	39.015	10.917	-5.826	2.059	-2.829*
800	33.132	2.850	38.566	10.403	-5.433	1.969	-2.759*
900	32.555	2.509	37.762	9.346	-5.207	1.767	-2.947*
1000	32.360	2.516	36.096	10.409	-3.736	1.955	-1.911

* t-values significant at α=.005

Figure 2.7 shows the initial convergence behavior of the GA and SA for the case of *k*=6 districts. The curves shown in Figure 2.7 correspond to the data provided in Table 2.5. The data indicate that there is a statistically significant difference between the average performance of the GA and the average performance of SA at the α=.005 level for every increment up to and including 800 function evaluations. When the optimization runs reach 900 function evaluations the difference is no longer statistically significant at the α=.005 level.

FIGURE 2.7
Revenue Deviation – 6 Districts

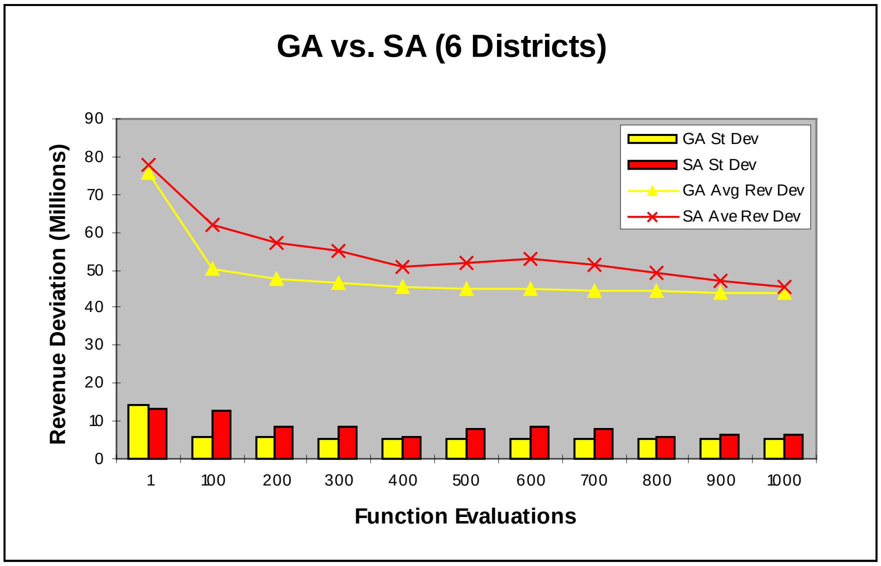

TABLE 2.5
Dynamic Analysis of Results for Revenue Deviation – 6 Districts

Function Evaluations	Genetic Algorithm		Simulated Annealing		Difference in Means	Standard Error	t-values
	Average	Standard Deviation	Average	Standard Deviation			
1	75.924	14.228	77.654	13.279	-1.731	3.553	-0.487
100	50.112	6.039	61.961	12.857	-11.849	2.594	-4.569*
200	47.691	5.621	57.304	8.425	-9.613	1.849	-5.199*
300	46.601	5.265	55.200	8.421	-8.599	1.813	-4.742*
400	45.481	5.231	50.735	5.912	-5.254	1.441	-3.646*
500	45.039	5.250	51.759	7.855	-6.721	1.725	-3.896*
600	44.857	5.146	53.066	8.208	-8.208	1.769	-4.641*
700	44.352	5.483	51.286	7.909	-6.933	1.757	-3.946*
800	44.213	5.417	49.361	5.594	-5.149	1.422	-3.621*
900	43.989	5.433	47.097	6.509	-3.108	1.548	-2.008
1000	43.989	5.433	45.644	6.454	-1.654	1.540	-1.074

* t-values significant at α=.005

Compactness

Figure 2.8 shows the initial convergence behavior of the GA and SA for the case of k=4 districts. The curves shown in Figure 2.8 correspond to the data provided in Table 2.6. The data indicate that there is not a statistically significant difference between the average performance of the GA and the average performance of SA at the α=.005 level over the first 1,000 function evaluations.

FIGURE 2.8
Compactness – 4 Districts

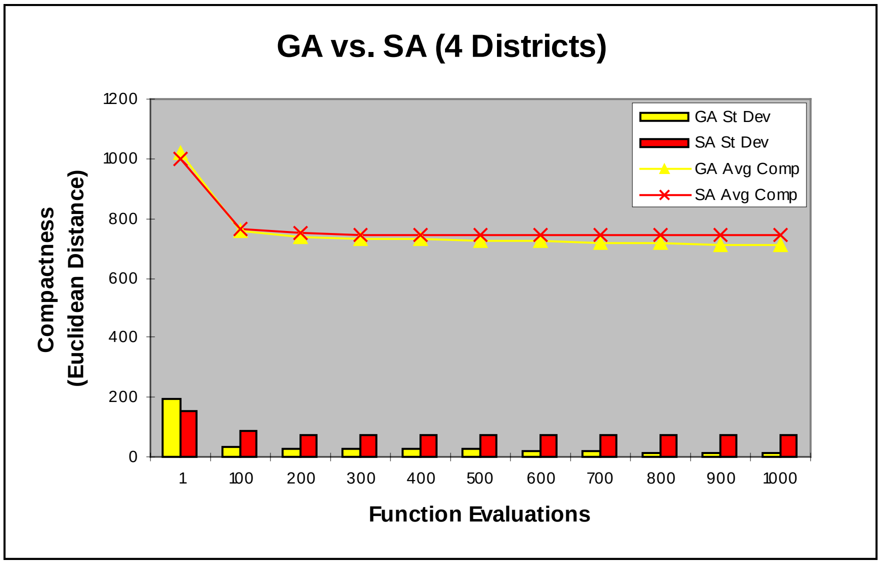

TABLE 2.6
Dynamic Analysis of Results for Compactness – 4 Districts

Function Evaluations	Genetic Algorithm		Simulated Annealing		Difference in Means	Standard Error	t-values
	Average	Standard Deviation	Average	Standard Deviation			
1	1018.488	194.956	1001.933	152.338	16.555	45.172	0.366
100	757.545	35.070	766.063	87.802	-8.518	17.262	-0.493
200	738.609	27.297	748.603	75.823	-9.994	14.713	-0.679
300	729.844	25.699	744.746	75.263	-14.902	14.520	-1.026
400	728.302	25.161	742.433	75.180	-14.131	14.474	-0.976
500	724.279	24.047	742.106	74.863	-17.827	14.356	-1.242
600	722.540	23.112	742.106	74.863	-19.566	14.305	-1.368
700	717.387	19.009	742.106	74.863	-24.719	14.102	-1.753
800	714.358	16.574	742.106	74.863	-27.748	13.999	-1.982
900	713.620	14.799	742.106	74.863	-28.486	13.932	-2.045
1000	712.381	14.937	742.106	74.863	-29.725	13.937	-2.133

Zero t-values significant at α=.005

Figure 2.9 shows the initial convergence behavior of the GA and SA for the case of k=5 districts. The curves shown in Figure 2.9 correspond to the data provided in Table 2.7. The data indicate that there is not a statistically significant difference between the average performance of the GA and the average performance of SA at the α=.005 level over the first 1,000 function evaluations.

FIGURE 2.9
Compactness – 5 Districts

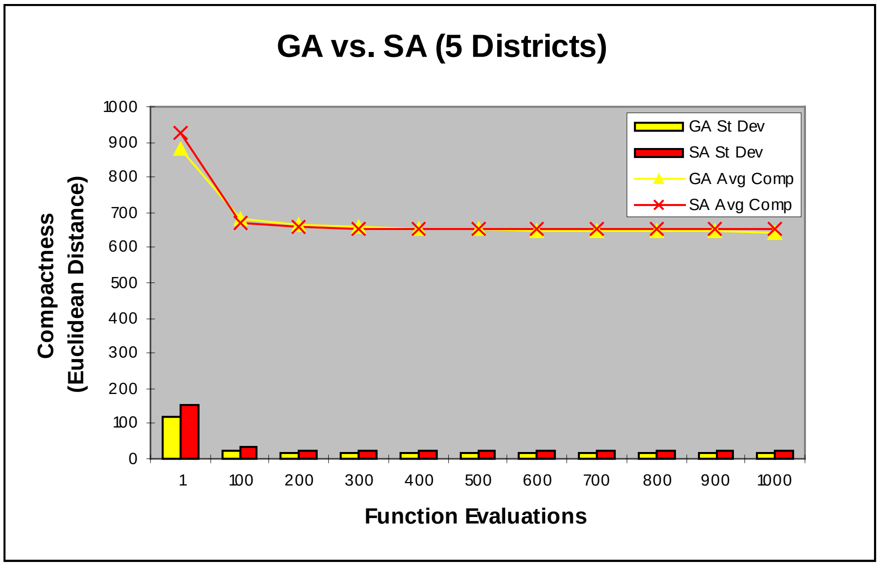

TABLE 2.7
Dynamic Analysis of Results for Compactness – 5 Districts

Function Evaluations	Genetic Algorithm Average	Genetic Algorithm Standard Deviation	Simulated Annealing Average	Simulated Annealing Standard Deviation	Difference in Means	Standard Error	t-values
1	878.186	120.105	925.495	152.084	-47.309	35.381	-1.337
100	679.327	24.695	671.430	32.128	7.896	7.398	1.067
200	663.722	14.541	657.648	24.284	6.074	5.168	1.175
300	656.665	16.878	655.989	23.816	0.676	5.329	0.127
400	654.838	16.535	654.943	23.715	-0.105	5.278	-0.020
500	650.739	17.757	655.063	23.665	-4.324	5.402	-0.800
600	649.939	16.892	654.906	23.766	-4.967	5.323	-0.933
700	647.043	16.497	654.984	23.820	-7.942	5.290	-1.501
800	645.325	15.524	654.906	23.766	-9.581	5.183	-1.849
900	645.073	15.447	654.906	23.766	-9.833	5.175	-1.900
1000	644.438	15.145	654.441	24.047	-10.004	5.189	-1.928

Zero t-values significant at α=.005

Figure 2.10 shows the initial convergence behavior of the GA and SA for the case of *k*=6 districts. The curves shown in Figure 2.10 correspond to the data provided in Table 2.8. The data indicate that there is not a statistically significant difference between the average performance of SA and the average performance of the GA at the α=.005 level for the initial 1,000 function evaluations.

FIGURE 2.10
Compactness – 6 Districts

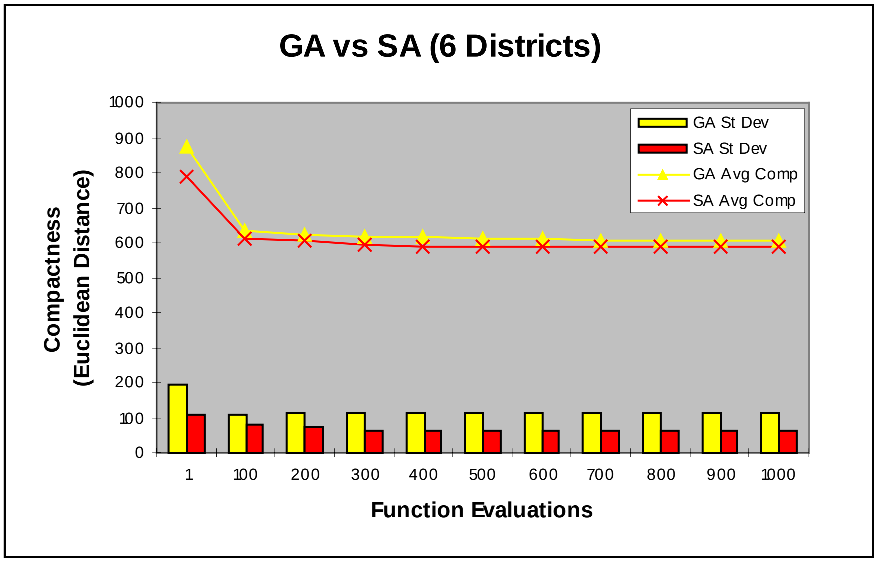

TABLE 2.8
Dynamic Analysis of Results for Compactness – 6 Districts

Function Evaluations	Genetic Algorithm Average	Genetic Algorithm Standard Deviation	Simulated Annealing Average	Simulated Annealing Standard Deviation	Difference in Means	Standard Error	t-values
1	873.963	194.428	786.058	108.777	-87.905	40.675	-2.161
100	634.035	110.810	609.824	77.214	-24.212	24.658	-0.982
200	625.271	112.185	603.104	75.977	-22.168	24.737	-0.896
300	617.726	112.981	592.684	63.377	-25.042	23.651	-1.059
400	617.421	113.021	590.422	62.883	-26.999	23.614	-1.143
500	612.702	113.772	590.420	63.014	-22.282	23.745	-0.938
600	611.233	113.805	590.075	62.743	-21.158	23.726	-0.892
700	608.237	114.321	590.049	62.790	-18.188	23.813	-0.764
800	606.944	114.551	590.034	62.720	-16.911	23.844	-0.709
900	605.002	114.804	590.034	62.720	-14.968	23.884	-0.627
1000	604.002	115.021	590.034	62.720	-13.969	23.919	-0.584

* Zero t-values significant at α=.005

Static Analysis

In this section, we perform a static analysis of the final solutions produced using each solution technique on Ghana's EPDP for each objective function independently. We report the statistical significance of the average performance difference between the GA and SA. We also report the quality of the best case solution and worst case solution for each procedure across all optimization runs. In addition, we are interested in the reliability of each method. To this end, we report the proportion of times each method repeatedly locates the best case and worst case solutions found during its respective performance tests.

Revenue Deviation

Table 2.9 shows that the average performance of the GA and SA on the total absolute revenue deviation function with respect to the quality of the final solution values for $4 \leq k \leq 6$ districts. The data indicate that there is a statistically significant difference between the average performance of the GA and the average performance of SA at the α=.005 level for the case of k=4 and k=6 districts. For the case of k=5 districts the performance difference is not statistically significant.

TABLE 2.9

Static Analysis of the Average Case Solution for Revenue Deviation

Districts	Average Case Solution		Difference in Means	Standard Error	t-value
	SA	GA			
4	24.226	19.220	-5.006	1.861	-2.691*
5	29.268	28.367	-0.901	1.125	-0.801
6	38.419	35.145	-3.275	0.731	-4.482*

* t-values significant at α=.005

Table 2.10 shows that both the GA and SA have approximately the same "best case" performance result. Note that the best solution found for k=4 and k=5 districts are identical for the GA and SA, while the best solution found by the GA for k=6 is slightly better than the best solution found by SA. However, we point out that the reliability of the GA is better than SA overall. For the case of k=4 districts, the GA consistently produced the best solution 30/30=100% of the time while SA finds the same solution 23/30=76.67% of the time. For k=5 districts, the reliability of each method is identical at 25/30=83.88%. For k=6 districts, the reliability of the GA is exactly twice that of SA, 10/30=33.33% versus 5/30=16.67% respectively. It is also important to point out that for k=6 the SA procedure consistently converged to a higher solution value. This is an indication of its inability to escape local minima once the system temperature has dropped sufficiently low.

TABLE 2.10

Static Analysis of the Best Case Solution for Revenue Deviation

Districts	Best Case Solution		Frequency of Best Case	
	SA	GA	SA	GA
4	19.220	19.220	23	30
5	27.583	27.583	25	25
6	33.305	33.158	5	10

Table 2.11 shows the "worst case" performance results for both the GA and SA. The worst case performance measure is an indication of the downside risk of choosing one solution technique over the other. Note that the worst case solution value of the GA is consistently better than the worst case solution value of SA for all k.

TABLE 2.11

Static Analysis of the Worst Case Solution for Revenue Deviation

Districts	Worst Case Solution		Frequency of Worst Case	
	SA	GA	SA	GA
4	45.156	19.220	3	30*

5	59.427	29.436	1	1
6	45.839	39.719	2	3

*Worst case solution value is also the best case solution value

Compactness

Table 2.12 shows that the average performance of the GA and SA on the total Euclidean distance function with respect to the quality of the final solution values for $4 \leq k \leq 6$ districts. The data indicate that there is a statistically significant difference between the average performance of the GA and the average performance of SA at the α=.005 level for the case of *k*=5 districts. For the case of *k*=4 and *k*=6 districts the performance difference is not statistically significant.

TABLE 2.12

Static Analysis of the Average Case Solution for Compactness

Districts	Average Case Solution		Difference in Means	Standard Error	t-value
	SA	GA			
4	742.106	709.922	-32.136	18.196	-1.769
5	654.906	619.426	-35.48	7.711	-4.601*
6	590.034	569.756	-20.278	15.472	-1.312

* t-value is significant at α=.005

Table 2.13 shows that both the GA and SA have similar "best case" performance results. Neither procedure performs best during the entire optimization test. Rather, we find that the better performing algorithm is a function of the number of desired districts in the plan. Note that for *k*=4 districts, the best solution found by both techniques is identical. Again the GA is more consistent than SA in locating this solution 9/30=30% versus 2/30=6.67% respectively. For *k*=5 districts the best solution found by the SA is worse than the best solution produced by the GA. Thus, the best case performance of SA involved a convergence to a local minimum for *k*=5 districts. In contrast, the GA is able to locate its best solution 6/30=20% of the time. When *k*=6 districts, the SA procedure is better than the GA with regard to the quality of the best solution found. Thus, the best case performance of the GA involved a convergence to a local minimum for *k*=6 districts. However, both solution techniques were unable to repeat their best case performance.

TABLE 2.13

Static Analysis of the Best Case Solution for Compactness

Districts	Best Case Solution		Frequency of Best Case	
	SA	GA	SA	GA
4	691.903	691.903	2	9
5	608.605	600.561	2	6
6	535.330	539.465	1	1

Table 2.14 shows the "worst case" performance results for both the GA and SA. Note that the worst case solution value of SA is consistently better than the worst case solution value of the GA for all *k*.

TABLE 2.14

Static Analysis of the Worst Case Solution for Compactness

Districts	Worst Case Solution		Frequency of Worst Case	
	SA	GA	SA	GA
4	995.056	1056.949	2	1
5	699.695	798.775	1	1
6	860.471	865.381	1	1

Summary

The central story of this empirical analysis was to match the relevant performance characteristics of the GA and SA procedures to the needs of our DSS. The story began with a dynamic analysis that is based upon the performance of each technique during the initial phase of repeated optimization runs. We believe that using the number of function evaluations as a basis of comparison for this test is appropriate because it provides a common measure of computational effort for comparing a population based technique to a point based search method.

In this study, we address a new class of districting problems – the EPDP. We are unaware of any test suite appropriate for our cause, and thus, we attempt to study an “open problem.” The static analysis that we perform in the second half of this study can, at best, provide us with a distribution of the final solution vector set. We can offer no claim of global optimality of our best results. The value that can be gleaned from the results of the static analysis is that the GA appears to have a slight performance advantage over SA for solving Ghana’s EPDP. We believe that this is attributable to the population based approach that is inherent to the GA. We have also gained an increased level of confidence in the relevance of our dynamic analysis. By this we mean that having expended considerably more computational effort, we find that the corresponding improvements in solution quality were only marginal.