Optimizing Revenue in CDMA Networks Under Demand Uncertainty

Eli V. Olinick • Jay M. Rosenberger

Department of Engineering Management, Information, and Systems, School of Engineering, Southern Methodist University, P.O. Box 750123, Dallas, Texas 75275-0123, USA

Department of Industrial and Manufacturing Systems Engineering, University of Texas at Arlington, P.O. Box 19017 Arlington, TX 76019-0017, USA

olinick@engr.smu.edu • *jrosenbe@uta.edu*

Code division multiple access (CDMA) networks allow subscribers to share frequencies, so optimizing revenue involves reserving channel equivalents to each market subject to signal-to-interference constraints at the radio towers. We present a yield management model inspired by those from the airline industry to optimize revenue under uncertainty. We describe the optimality conditions and develop a supergradient algorithm. We provide computational results that show the effects of the distribution and variance of demand. Finally, we discuss areas of future research, including a method to optimize the locations of the towers.

(*Keywords: OR in Telecommunications; Nonlinear Programming; Stochastic Programming*)

---

1. Introduction

Optimization models for network design, routing, and capacity expansion in the telecommunications industry are typically driven by demand forecasts. As Laguna [11] observes, however, it is particularly difficult to predict the demand pattern that will be placed on a telecommunications network. Building or expanding these networks often entails considerable investment in infrastructure. The investment could be lost if the network is designed based upon an overly optimistic forecast. Moreover, if the forecast is too conservative, the resulting network will not have enough capacity to meet all of the demand. Consequently, the service provider risks losing market share while it scrambles to install additional capacity to meet unanticipated demand. Thus, the problem of network capacity planning under demand uncertainty has received considerable attention in the literature.

Sen et al. [19] present a stochastic programming model for capacity expansion in which the demand is realized after additional lines are added to an existing network. The recourse problem is to find optimal routes for the traffic given the link capacities determined in the first stage. Other examples of stochastic programming approaches to capacity planning in fiber-optic networks include Lisser et al. [12], Riis and Lodalhl [17], Riis and Andersen [16], and Smith et al. [21]. Another popular approach to incorporating demand uncertainty in optimization models is the robust optimization methodology described by Mulvey et al. [14]. Some examples of applying this technique to capacity planning and traffic routing problems in fiber-optic networks are Laguna [11], Gryseels et al. [7], Birkan et al. [3] and Kennington et al. [10]. One objective in capacity planning is to optimize revenue, and in this paper we maximize expected revenue in cellular radio networks using *code division multiple access* (CDMA) technology under demand uncertainty.

Mobile telephone networks are based on the concept of cellular communications in which the service area is divided into subsections called *cells* that are each served by a radio tower. Thus, *tower location* is an important subproblem in cellular network design. The other factor determining the capacity and coverage area of a cellular network is how users share bandwidth (radio frequencies). There are two basic strategies for sharing bandwidth in commercial cellular systems. Systems using *frequency division multiple access* (FDMA), divide the bandwidth available to the service provider into distinct frequency channels. A central problem in maximizing capacity in a FDMA system is that of assigning frequencies to towers in such a way that users in nearby cells do not cause interference with each other by using the same frequencies. Thus, optimization models for FDMA design often include a graph-coloring component. Murphey et al. [15] give an extensive survey of frequency assignment problems.

CDMA networks use a considerably different approach to bandwidth sharing than FDMA. By using orthogonal codes and strict control of transmission power, CDMA systems enable a service provider to use all the frequencies it has licensed in each cell [9]. Thus, frequency assignment is not an explicit issue in CDMA planning problems. Instead, the primary issue is to keep the signal-to-interference ratio of each subscriber at an acceptable level. In a CDMA system, each call interferes with every other call. If the total interference at a particular tower reaches a certain threshold, then the next subscriber in that cell who attempts to make a call will be denied access to the system. As we describe in Section 2, the amount of interference caused by a particular subscriber's handset depends on the tower to which the subscriber is assigned.

CDMA is a relatively new technology, and so there has been little research reported on optimizing the joint tower location and subscriber assignment problems for CDMA networks. Galota et al. [6] present a polynomial-time approximation scheme for a profit maximization model for radio tower location and subscriber assignment in CDMA networks. Amaldi et al. [1] minimize the system cost, measured in terms of tower construction and overall system interference, subject to providing enough capacity to allow all subscribers in the area simultaneous access to the system. Mathar and Schmeink [13] maximize the system capacity subject to a budget constraint. Kalvenes et al. [9] propose a profit maximization model that considers the trade-off between the revenue generated per customer served and the cost of constructing towers. Their basic model is a large integer program that requires a special algorithm to find practical solutions.

Optimization models for cellular networks typically model demand by specifying a set of discrete points in the service area where a given amount of traffic is requested. These points represent "hot spots" where demand tends to be concentrated and the demand is usually measured in Erlangs or *channel equivalents* (i.e., the number of simultaneous calls in progress at any given time) [1]. Likewise, we define a *market* as a demand concentration point in the service provider's coverage area and pose the *channel allocation* problem as determining how to reserve channel equivalents to markets. Although the demand at a market may fluctuate throughout a given day and/or vary from one day to the next, the cellular network design models described above are deterministic. This is consistent with how planning is typically done in the industry today. Providers design their systems to be able to accommodate a small amount more than 100% of the anticipated average peak demand. Our model addresses the uncertainty about the values for the average peak demands that are used in the deterministic models for cellular network design.

m_2

ℓ_1 ℓ_2

m_1 m_3

Figure 1: Channel Allocation Example

Regardless of the bandwidth sharing scheme, the allocation of channel equivalents to a market can dramatically affect revenue. Consider the two towers, ℓ_1 and ℓ_2 and three markets m_1, m_2, and m_3 as depicted in Figure 1. Although the subscribers' handsets in each market interfere with both towers, markets m_1 and m_3 interfere insignificantly with towers ℓ_2 and ℓ_1, respectively. Allowing too many subscriptions in m_2 interferes with both towers ℓ_1 and ℓ_2 and crowds out potential subscribers in markets m_1 and m_3. Consequently, we can capture more revenue by reserving many channel equivalents in markets m_1 and m_3, but limiting those of market m_2. Kalvenes et al. [9] implicitly solve a deterministic channel allocation problem as a subproblem to their integer program, as demonstrated in Section 2.

In this paper, we present a yield management model (YM) inspired by those in the airline industry to optimize revenue under demand uncertainty. In Section 2, we describe the deterministic models in Amaldi et al. [1] and Kalvenes et al. [9]). Section 3 develops our YM model for channel allocation. We present a supergradient algorithm to solve the YM model in Section 4. We demonstrate the effectiveness of the algorithm on CDMA network problems from the literature and show the effects of distribution and the variance of demand in our computational results in Section 5. Finally, we discuss conclusions and future research, particularly how we can use this model to optimize the locations of towers, in Section 6.

2. Deterministic Models for Tower Selection and Channel Assignment

Before describing our new yield management model for CDMA systems, we first review the framework for the deterministic problems treated by Amaldi et al. [1] and Kalvenes et al. [9]. Let L denote the set of potential tower locations, and a_ℓ denote the amortized cost of building and operating a tower at location $\ell \in L$. The set of markets is denoted by M, and d_m denotes the expected demand for service, measured in number of channel equivalents, at market $m \in M$. The revenue received per channel equivalent served is denoted by r. There are two types of decision variables in this model: the binary variable y_ℓ indicates whether or not a tower is built at location $\ell \in L$, and the integer variable $x_{m\ell}$ denotes the number of channel equivalents at tower $\ell \in L$ that are allocated for subscribers in market $m \in M$. As we describe below, the number channel equivalents may also be thought of as the number of simultaneous calls supported by a tower.

The parameter P_{target} specifies the target power level for signals received at the towers. When a subscriber in market m transmits to tower ℓ the signal strength is weakened by a given *attenuation factor* $g_{m\ell}$. Thus, the subscriber's handset transmits with power level $\frac{P_{\text{target}}}{g_{m\ell}}$ so that the strength of the signal received at tower ℓ is $g_{m\ell} \times \frac{P_{\text{target}}}{g_{m\ell}} = P_{\text{target}}$. Each tower in a CDMA system receives signals from many subscriber handsets in the surrounding neighborhood. In order for the signals to be processed with a reasonable error rate, the signal-to-interference ratio for any call must be at least a given threshold value, $SIR_{\min}$. To illustrate how this requirement is captured in optimization models, consider the number of channel equivalents allocated for market $m \in M$ assigned to tower $\ell \in L$, $x_{m\ell}$. Since each of these subscribers' handsets transmits at power level $\frac{P_{\text{target}}}{g_{m\ell}}$, the received power from each of these handsets at some other tower $j \in L \setminus \{\ell\}$ is $g_{mj} \times \frac{P_{\text{target}}}{g_{m\ell}} = \frac{g_{mj}}{g_{m\ell}} P_{\text{target}}$. In general, the total received power at tower location ℓ from all markets is given by

$$P_\ell^{\text{TOT}} = P_{\text{target}} \sum_{m \in M} \sum_{j \in L} \frac{g_{m\ell}}{g_{mj}} x_{mj}. \tag{1}$$

Consider a subscriber assigned to tower ℓ. When that subscriber makes a call, P_{target} represents the strength of signal for the session, and all other power received at tower ℓ, $P_\ell^{\text{TOT}} - P_{\text{target}}$, is interference ([1]). Thus, the channel allocation must satisfy the requirement

$$\frac{P_{\text{target}}}{P_\ell^{\text{TOT}} - P_{\text{target}}} \geq SIR_{\min}, \tag{2}$$

which is equivalent to imposing the following quality-of-service (QoS) constraint:

$$\sum_{m \in M} \sum_{j \in L} \frac{g_{m\ell}}{g_{mj}} x_{mj} \leq 1 + \frac{1}{SIR_{\min}}. \tag{3}$$

Observe that for any tower $\ell \in L$, (3) implies that

$$\sum_{m \in M} \sum_{j = \ell} \frac{g_{m\ell}}{g_{mj}} x_{m\ell} = \sum_{m \in M} x_{m\ell} \leq 1 + \frac{1}{SIR_{\min}}. \tag{4}$$

Thus, even if we ignore the attenuation factors there is an implicit limit on the number of simultaneous calls a tower can support.

The profit-maximization model described by Kalvenes et al. (2002 [9]) is the following integer program (KKOIP)

$$\max \sum_{m\in M}\sum_{\ell\in L} r x_{m\ell} - \sum_{\ell\in L} a_\ell y_\ell \tag{5}$$

$$\sum_{m\in M}\sum_{j\in L} \frac{g_{m\ell}}{g_{mj}} x_{mj} \le s + (1-y_\ell)\beta_\ell \qquad \forall \ell \in L, \tag{6}$$

$$x_{m\ell} \le d_m y_\ell \qquad \forall m \in M, \ell \in L, \tag{7}$$

$$\sum_{\ell\in L} x_{m\ell} \le d_m \qquad \forall m \in M, \tag{8}$$

$$x_{m\ell} \in \mathcal{Z}^+ \qquad \forall m \in M, \ell \in L, \tag{9}$$

$$y_\ell \in \{0,1\} \qquad \forall \ell \in L, \tag{10}$$

where $s = 1 + 1/SIR_{\min}$, and β_ℓ is a sufficiently large constant so that the QoS constraint (6) for tower ℓ is only binding if tower ℓ is built. Constraint set (7) links the decision variables for tower location and market allocation variables, and the domains for the variables are provided by constraint sets (8)-(10). The cost-minimization model presented by Amaldi et al. [1] is similiar to KKOIP except that it is a pure binary program requiring that all demand is served in every market and that all subscribers in a given market are assigned to the same base station.

Our yield management model relies on a result which states that revenue is maximized by assigning subscribers at market m to the available tower with the largest attenuation factor [1, 9]. More precisely, suppose that towers are built at locations $i \in L$ and $j \in L$ where $g_{mi} < g_{mj}$, then there exists and optimal solution in which $x_{mi} = 0$. Since an important component influencing the attenuation factor between two locations is distance, an intuitive statement of this result is that there is always an optimal solution in which all subscribers in a given market m are assigned to the nearest constructed tower. In our model, we assume that the tower-selection problem has already been solved and find the maximum expected revenue for the set of towers selected. This would likely be used in a planning tool where the network designer would evaluate the maximum expected revenue for a variety of tower configurations. Furthermore, we intend to use it in future research as a subproblem in a more complex model that addresses tower-selection and subscriber assignment simultaneously.

3. Yield Management Model

Most major domestic airline carriers use yield management (YM) techniques to optimize their revenue. American Airlines, for example, reported a \$1.4 billion revenue increase over a three-year period due to improved YM methods [20]. The airlines use a complex fare structure, which identifies several types of passengers, such as business and leisure. Passengers fly a sequence of flights called an *itinerary*, and a *fare class* includes a type of passenger and an itinerary. For each fare class, airline planners estimate the demand and assign a fare. They maximize revenue by allocating a number of seats for each fare class subject to the capacities of the aircraft assigned to the flights. When allocating seats, the airlines determine a break-even price, or *bid price*, for each flight. Every fare class that is allocated at least one seat has a fare greater than or equal to the sum of the bid prices of the flights in the associated itinerary. Airlines execute these models several times before departure, so the allocations and bid prices change dynamically. D'Slyvia [4] develops a stochastic model for airline yield management, and Günther [8] and Chen et al. [22] discuss a variety of airline yield management models.

Similar to the airlines, telecommunications providers segregate their demand into markets and allocate channel equivalents to them. Although we assume the revenue per subscriber is the same for different markets, using a variety of prices is one topic of future research. The number of allocated channel equivalents in a market in a CDMA network is limited by the QoS constraints (6) at the towers. In optimizing channel allocation, we determine an *interference price* at each tower. A channel equivalent in a market will be allocated only when the revenue is greater than or equal to the sum of the interference prices weighted by the attenuation factor from the market to the towers.

3.1 Model

In this section, we provide a deterministic and a stochastic market allocation model. In both models, we relax the allocation to be continuous.

For each market $m \in M$, let $\widetilde{d}_m$ be a continuous random variable for the channel equivalent demand in market m, and let $f_{\widetilde{d}_m}(x)$, $F_{\widetilde{d}_m}(x)$, and $E[\widetilde{d}_m]$ be the probability density function, the cumulative distribution function, and the expected value of the demand $\widetilde{d}_m$, respectively. The result in Amaldi et al. [1] and Kalvenes et al. [9] states that markets are assigned to the tower with the largest attenuation factor, so for each market $m \in M$, let ℓ_m be the tower assigned to market m. To simplify notation, let the channel allocation x_m be equivalent to $x_{m\ell_m}$ and the attenuation factor g_m be equivalent to $g_{m\ell_m}$.

Let $\widehat{L}$ be the set of constructed towers. The deterministic channel allocation problem

(DCA) is given by

$$\max Det(x) = \sum_{m \in M} r x_m \tag{11}$$

$$\sum_{m \in M} \frac{g_{m\ell}}{g_m} x_m \leq s \qquad \forall \ell \in \widehat{L} \tag{12}$$

$$E[\widetilde{d}_m] \geq x_m \geq 0 \qquad \forall m \in M \tag{13}$$

Observe that given a set of constructed towers, (DCA) is the linear programming relaxation of the integer program (KKOIP). Constraint set (12) includes the QoS constraints (6), and constraints (5)-(9) reduce to constraint set (13).

The deterministic model (DCA) assumes that subscribers will consume all of the allocated channel equivalents, so the objective of (DCA) overestimates the expected revenue. The expected number of subscribers in terms of channel equivalents is given by,

$$E\left[\min\left(\widetilde{d}_m, x_m\right)\right] = \int_{-\infty}^{x_m} t f_{\widetilde{d}_m}(t)dt + x_m\left(1 - F_{\widetilde{d}_m}(x_m)\right).$$

$E\left[\min\left(\widetilde{d}_m, x_m\right)\right]$ is concave since the second partial derivative of the expected number of subscribers is always nonpositive. That is,

$$\frac{d^2 E\left[\min\left(\widetilde{d}_m, x_m\right)\right]}{dx_m^2} = -f_{\widetilde{d}_m}(x_m) \leq 0 \text{ for all } x_m \in \Re.$$

The stochastic channel allocation problem (SCA) is given by

$$\max Stoch(x) = \sum_{m \in M} r E\left[\min\left(\widetilde{d}_m, x_m\right)\right] \tag{14}$$

$$\sum_{m \in M} \frac{g_{m\ell}}{g_m} x_m \leq s \qquad \forall \ell \in \widehat{L} \tag{15}$$

$$x_m \geq 0 \qquad \forall m \in M \tag{16}$$

The airline yield management problem described in D'Sylvia [4] is similar to (SCA), but the attenuation ratios are either one, if the flight is in the fare class, or zero, if it is not. The stochastic and deterministic airline yield management models in Chen et al. [22] are analogous to (DCA) and (SCA), respectively. Neither D'Sylvia [4] nor Chen et al. [22] discuss algorithms to solve the analogous airline YM problems.

For each tower $\ell \in \widehat{L}$, denote δ_ℓ as the *violation* of the interference constraint (15) at tower ℓ; that is,

$$\delta_\ell = \sum_{m \in M} \frac{g_{m\ell}}{g_m} x_m - s.$$

For each tower $\ell \in \widehat{L}$, let λ_ℓ be the interference price associated with the constraint in set (15). For each market $m \in M$, let μ_m be the *market price* given by

$$\mu_m = \sum_{\ell \in \widehat{L}} \frac{g_{m\ell}}{g_m} \lambda_\ell.$$

Let $\mathcal{L}(x, \lambda)$ be the Lagrangian of (SCA), and so the optimality conditions in addition to (15)-(16) are given by

$$\frac{\partial \mathcal{L}(x,\lambda)}{\partial x_m} = r\left(1 - F_{\tilde{d}_m}(x_m)\right) - \mu_m \leq 0 \qquad \forall m \in M, \tag{17}$$

$$x_m \frac{\partial \mathcal{L}(x,\lambda)}{\partial x_m} = x_m\left(r\left(1 - F_{\tilde{d}_m}(x_m)\right) - \mu_m\right) = 0 \qquad \forall m \in M, \tag{18}$$

$$\lambda_\ell \frac{\partial \mathcal{L}(x,\lambda)}{\partial \lambda_\ell} = \lambda_\ell \delta_\ell = 0 \qquad \forall \ell \in \widehat{L}, \tag{19}$$

$$\lambda_\ell \geq 0 \qquad \forall \ell \in \widehat{L}. \tag{20}$$

Constraint set (17) includes the *dual feasibility constraints*, and the *primal complimentary slackness* constraints are given by (18). Constraint sets (19) and (20) are the *dual complimentary slackness* and *dual nonnegativity constraints*, respectively. Observe that from (18) the channel allocation x_m is positive only when $r - \mu_m = rF_{\tilde{d}_m}(x_m) \geq 0$, and so the revenue r must be greater than or equal to the market price μ_m.

4. Algorithm

Consider the supergradient algorithm *SCAOPT*, as depicted as algorithm 1, to solve (SCA). On each iteration, SCAOPT finds the optimal channel allocations based upon the interference prices because constraint sets (15) - (18) and (20) are satisfied. Consequently, the violation vector δ is a supergradient. SCAOPT updates the interference prices in the directions of the supergradient and terminates if a stopping criteria, which is discussed in Section 4.1, is satisfied.

SCAOPT will converge as long as the step size α^k is such that

$$0 < \alpha^k < \frac{2\left(\mathcal{L}(x^k, \lambda^k) - \mathcal{L}(x^*, \lambda^*)\right)}{||\delta||^2},$$

where (x^*, λ^*) is an optimal primal-dual solution to (SCA) [2].

Algorithm 1 SCAOPT.

$k \leftarrow 0$
$STOP \leftarrow FALSE.$
Initialize $\lambda^0 > 0$.
while $STOP = FALSE$ **do**
  for all $m \in M$ **do**
    $$x_m^k \leftarrow \begin{cases} \left[F_{\tilde{d}_m}^{-1}\left(1 - \frac{\mu_m^k}{r}\right)\right]^+ & \text{if } (r > \mu_m^k), \\ 0 & \text{otherwise.} \end{cases}$$
  end for
  if Stopping Criteria Satisfied **then**
    $STOP \leftarrow TRUE.$
  else
    $\lambda^{k+1} \leftarrow \left[\lambda^k + \alpha^k \delta^k\right]^+.$
    $k \leftarrow k + 1$
  end if
end while

4.1 Feasible Projection

On iteration k of SCAOPT, the channel allocation x^k typically violates some of the QoS constraints in set (15), and a projection onto the feasible region can provide a lower bound on the optimal Lagrangian value $\mathcal{L}(x^*, \lambda^*)$. For each market $m \in M$, let $\overline{x}_m^k$ be a feasible projection of the channel allocation x_m^k, and let $\gamma_m^k = \overline{x}_m^k - x_m^k$. By weak duality,

$$Stoch(\overline{x}^k) \leq \mathcal{L}(x^k, \lambda^k). \tag{21}$$

Although there are several methods to find a feasible projection, one method is by solving the following projection problem (SCAPROJ)

$$\min \sum_{m \in M} |\gamma_m^k| \tag{22}$$

$$\sum_{m \in M} \frac{g_{m\ell}}{g_m} \gamma_m^k \leq -\delta_\ell^k \qquad \forall \ell \in \widehat{L}, \tag{23}$$

$$\gamma_m^k \geq -x_m^k \qquad \forall m \in M. \tag{24}$$

Using weak duality, we define the gap of iteration k to be

$$gap^k = \frac{\mathcal{L}(x^k, \lambda^k) - Stoch(\overline{x}^k)}{Stoch(\overline{x}^k)} \times 100\%.$$

For an optimal solution (x^*, λ^*), x^* is feasible, and so $x^* = \overline{x}^*$. By strong duality, we know that $Stoch(\overline{x}^*) = Stoch(x^*) = \mathcal{L}(x^*, \lambda^*)$, so $gap^* = 0$. Thus, we use the stopping criteria $\left|gap^k\right| < \varepsilon$.

4.2 Step Size

We use three step-size approaches within SCAOPT. Approaches A and B are from [2] and are given by

$$\alpha^k \leftarrow \frac{\rho^k \left(\mathcal{L}(x^k, \lambda^k) - \overline{\mathcal{L}}\right)}{||\delta^k||^2}, \tag{25}$$

where $\overline{\mathcal{L}}$ is an estimation of the optimal value of the Lagrangian $\mathcal{L}(x^*, \lambda^*)$. Approach A sets $\rho^k = 1$ for every iteration and calculates $\overline{\mathcal{L}}$ using

$$\overline{\mathcal{L}} \leftarrow (1 - \phi_k)\widehat{\mathcal{L}}, \tag{26}$$

where $\widehat{\mathcal{L}} = \max_{0 \le i < k} \mathcal{L}(x^i, \lambda^i)$ and ϕ_k is a scalar between zero and one. The scalar ϕ_k is increased when the Lagrangian is less than all of its previous values— $\mathcal{L}(x^k, \lambda^k) < \widehat{\mathcal{L}}$, and ϕ_k is decreased otherwise.

Approach B uses the optimal value of projected channel allocation to estimate the optimal value of the Lagrangian; that is,

$$\overline{\mathcal{L}} \leftarrow \mathit{Stoch}(\overline{x}_m^k). \tag{27}$$

The coefficient ρ^k is given by

$$\rho^k \leftarrow \frac{1 + \omega}{k + \omega}, \tag{28}$$

where ω is a constant.

Approach C considers the distance of the complimentary slackness vector $||\lambda \cdot \delta||$. If $||\lambda^k \cdot \delta^k|| \ge ||\lambda^{k-1} \cdot \delta^{k-1}||$, then α^k is decreased slightly by a constant factor, ψ; otherwise, $\alpha^k \leftarrow \alpha^{k-1}$.

4.3 Newton Method

In addition to the supergradient approach, we implemented a Newton Method with a barrier penalty that maximized the following function

$$\mathit{Stoch}^+(x) = \mathit{Stoch}(x) + \varepsilon \left[\sum_{m \in M} \ln(x_m) + \sum_{\ell \in \widehat{L}} \ln(-\delta_\ell) \right].$$

The function $\mathit{Stoch}^+(x)$ is concave with a defined value, gradient vector, and Hessian matrix for each point in the interior of the feasible region of x. The algorithm starts with an interior point and finds a direction using the Newton Method. It then uses a line search to find a new interior point. One major drawback to this approach is that the Hessian matrix is $|M| \times |M|$, so calculating its inverse can be time consuming.

4.3.1 Dual Projection

Similar to the dual supergradient, primal Newton Methods may stop when solutions are near optimality. On occasion, our Newton Method finds a set of Lagrange multipliers using a linear programming projection. Let $(\overline{x}, \overline{\delta})$ be a primal feasible solution. Consider the following projection onto the Lagrange multipliers λ.

$$\min \sum_{m \in M} \overline{x}_m b_m - \varepsilon \sum_{\ell \in \widehat{L}} \overline{\delta}_\ell \lambda_\ell \tag{29}$$

$$\sum_{\ell \in \widehat{L}} \frac{g_{m\ell}}{g_m} \lambda_\ell - b_m = r\left(1 - F_{\tilde{d}_m}(\overline{x}_m)\right) \qquad \forall m \in M \tag{30}$$

$$\lambda_\ell \geq 0 \qquad \forall \ell \in \widehat{L} \tag{31}$$

$$b_m \geq 0 \qquad \forall m \in M. \tag{32}$$

The objective is to minimize the complementary slackness values associated with (18) and use the values associated with (19) to break ties. If $b = 0$ in a projected solution, then complementary slackness conditions (18) are held, as they are also held in SCAOPT. Because $g_{m\ell} > 0$, $\overline{x}_m \geq 0$, $\overline{\delta}_\ell \leq 0$, and $1 \geq F_{\tilde{d}_m}(\overline{x}_m) \geq 0$, the Lagrange multiplier projection linear programming problem is always feasible and bounded. Using a solution to the multiplier projection LP and a single iteration of SCAOPT, we can find an upper bound on a primal problem and a primal-dual gap for the Newton Method approach.

5. Computational Results

To demonstrate the effectiveness and benefits of SCAOPT and the Newton Method on CDMA network problems from the literature, we implemented them in the C programming language on a Dell Precision Workstation with dual Intel Xeon Processors and 2 Gb of memory, and we tested them on a series of 36 problem instances. The instances collectively referred to as Scenario 1 are based on a problem proposed by Amaldi et al. (2001a, 2001b). This problem has 22 potential tower locations and 95 markets randomly placed in a 400 square meter grid. Each market has a demand for one channel equivalent and a minimum signal-to-interference ratio of $SIR_{\min} = 0.03125$. To generate the data for Scenario 1, we solved this problem with (KKOIP) to determine the set of constructed towers $\hat{L}$. The solution of the tower location problem has 4 towers and 95 markets from which we derived 18 stochastic problems by letting $\tilde{d}_m$ have a Normal or Gamma distribution with a mean of 1 and varying the standard deviation over a set of nine values.

To generate the data for Scenario 2, we started with a set of 40 towers that cover a large area in the northern plains of the United States (see F [5]). Following the procedure in Kalvenes et al. (2002), we drew a sample of 250 randomly placed markets in this service area each with a demand d_m drawn randomly from the range $[1, 16]$ and used $SIR_{\min} = 0.009789$. Solving this deterministic problem gave us a stochastic instance with 24 towers and 250 markets from which we generated 18 problems by letting $\tilde{d}_m$ have a Normal or Gamma distribution with a mean d_m and varying the standard deviation as we did with Scenario 1. For both scenarios, the revenue per channel was $r = \$42, 820$.

In all of our computational experiments we stopped the supergradient algorithm after $k = 300, 000$ iterations or when we had an optimality gap $gap^k < 0.01\%$. Before solving (SCA), we used CPLEX 8.1 to solve (DCA). None of the QoS constraints in the (DCA) solution for Scenario 1 were binding, so we initialized the interference prices to be the average deterministic revenue per tower $\frac{Det(x)}{|\widehat{L}|}$. For Scenario 2, we initialized the tower interference prices to be those of the (DCA) solution.

We solved the projection problem (SCAPROJ) with CPLEX 8.1, and we used functions from the GNU Scientific Library to evaluate *Stoch* and F^{-1} for both the distributions. For the Normal distribution, we also used GNU's quadrature adaptive integration function, `gsl_integration_qagil`, with 10,000 double precision intervals to calculate the integral component of *Stoch*. The integration function terminated when the error was within 0.1% of the expected demand $E\left[\tilde{d}_m\right]$. For a Gamma random variable with distribution shape parameter n, scale parameter l, and probability density function $f_{n,l}(x)$, we used cumulative distribution function $F_{n+1,l}$ and equation (33) to evaluate *Stoch*;

$$\int_0^x s f_{n,l}(s) ds = nlF_{n+1,l}(x). \tag{33}$$

For step size A, we started with the scalar $\phi = 0.0001$. As described in Section 4.2, when the Lagrangian was less than all of its previous values— $\mathcal{L}(x^k, \lambda^k) < \widehat{\mathcal{L}}$, we increased ϕ so that $\phi^{k+1} \leftarrow \min(1, 1.01\phi^k)$. Otherwise, we decreased the value so that $\phi^{k+1} \leftarrow 0.99\phi^k$. With step size B, we used a value of 299 for the parameter ω. For the implementation of step size C, we started with $\alpha = 10.0$. This value was decreased by a factor of $\psi = 0.9995$ whenever the distance of the complimentary slackness vector increased from one iteration to the next (i.e., $||\lambda^k \cdot \delta^k|| > ||\lambda^{k-1} \cdot \delta^{k-1}||$).

Tables 1 to 4 give the results of our experiments. In each table, the first column lists the *coefficient of variation* (CV) for the distribution. The CV gives a measure of dispersion of the distribution and is defined as the ratio of the standard deviation to the mean. For each of the nine CV values, we ran the algorithm three times using step sizes A, B, and C. For each run, we report the number of iterations k^*, the quality of the projected channel allocation as measured by expected revenue $Stoch(\overline{x}^{k^*})$, the distance of the complimentary slackness vector at termination $||\lambda^{k^*} \cdot \delta^{k^*}||$, and the optimality gap gap^{k^*}.

Table 1 about here.

Table 1 gives the results for Scenario 1 when the demand in each market has a Gamma distribution. In all but one of these 27 experiments, SCAOPT found a solution with a gap $gap^{k^*} < 0.01\%$. For the problem with CV = 0.01, SCAOPT had a gap of 0.02% after 300,000 iterations using step size B. For CV > 0.01, step size B used an average of 1,215 iterations. For step sizes A and C, SCAOPT required an average of 2,443 and 4,269 iterations to find solutions within 0.01% of optimality.

Table 2 about here.

The results for Scenario 2 when the demand in each market has a Gamma distribution are shown in Table 2. SCAOPT found a solution within 0.01% of optimality for all experiments using step sizes B and C, and the average number iterations were 1,449 and 572, respectively. For each experiment in which CV $\leq$ 0.10, SCAOPT did not find a solution within 0.01% of optimality using step size A. However, the geometric mean of the optimality gap was 0.02% overall and 0.04% for those experiments that terminated after 300,000 iterations. When CV > 0.10, the average number of iterations for step size A was 508.

Table 3 about here.

Table 3 gives the results for Scenario 1 when the demand in each market has a Normal distribution. Results using the Normal distribution are very similar to those using the Gamma distribution for CV $\leq$ 0.10. However, overall step sizes A and B used markedly fewer iterations for the experiments that used Normal distribution than for those of the Gamma distribution. For CV > 0.01, step size B used an average of 766 iterations when demand was assumed to be a Normal random variable, versus 1,215 when demand was a Gamma random variable. For step size A, SCAOPT required an average of 917 iterations assuming the demand is from a Normal distribution, which is less than half of the average

2,443 iterations when demand was a Gamma random variable. For step size C, the results from Tables 1 and 3 are very similar, as SCAOPT required an average of 4,354 iterations to find solutions within 0.01% of optimality in Table 3.

$\boxed{\text{Table 4 about here.}}$

Table 4 gives the results for Scenario 2 when the demand in each market has a Normal distribution. In all but two of these experiments (CV = 0.05 and 0.01), SCAOPT found solutions within 0.01% of optimality using step size A. Excluding this case, SCAOPT terminated after an average number of 501 iterations for these runs. With step sizes B and C, the algorithm found solutions for these problems that were all within 0.01% of optimality and averaged 916 and 538 iterations, respectively.

Table 5: Summary of SCAOPT Results

Step Size	No. Opt.	Ave. Iter.
A	31	42,689
B	34	17,698
C	36	2,434

Table 5 summarizes our results by step size. The second column gives the number of problems (out of 36) solved to within 0.01% of optimality, and third column shows the average number of iterations iterations for SCAOPT with each step size. Based on these results, step size C is clearly superior to step sizes A and B, even though it rarely outperformed both step sizes A and B. To get even better performance using three processors, one could execute SCAOPT with each step size on separate processors simultaneously and terminate it when the first one finished. Using parallel processing, no instance would require much more than 1000 iterations, which takes less than 40 seconds of CPU time. Thus, SCAOPT is a practical algorithm for these problems.

In addition to SCAOPT experiments, we also tested each instance with each CV value using the Newton Method. However, the efficiency of the Newton Method was poor for two reasons. One, computing the inverse of the Hessian matrix using LU-decomposition was significantly time consuming. Two, conducting a line search to find a new interior point required evaluating constraints (15) and (16) for several points. Consequently, an average iteration of the Newton Method required 13 times more CPU time than an average iteration of SCAPOPT. We also estimated the efficiency for which each algorithm iteration reduced the gap using the following metric

$$\text{gap efficiency} = \frac{100 - gap}{k}.$$

Although the gap efficiency for the Newton Method was more than three times that of SCAOPT, it never compensated for the increased CPU time. Consequently, SCAOPT outperformed the Newton Method.

Observe that the total expected revenue decreases as the variance increases. For moderate CV values of 0.15 and 0.25, the expected revenue is between 3% and 7% less than the deterministic optimal revenue. For larger CV values of 0.50, the revenue is as much as 18% less than the deterministic revenue. Consequently, the deterministic optimal solution overstates expected revenue.

To determine the effects of this overstatement, we compared the value of the optimal solution of the deterministic problem, denoted as $Stoch(\overline{x}^{DET})$, to the best value from SCAOPT, $Stoch(\overline{x}^{STOCH})$, for each scenario, CV value, and demand distribution in Tables 6 and 7. The column labeled "% improvement" refers the difference between the stochastic solution value $Stoch(\overline{x}^{STOCH})$ and the deterministic solution value $Stoch(\overline{x}^{DET})$ as a percentage of the deterministic solution value.

Table 6 about here.

As CV increases, the percentage improvement increases. Even for modest uncertainty, CV = 0.10, the improvement is more than 2% for Scenario 1, while for larger uncertainty, CV = 0.50, the improvement is about 6%.

Table 7 about here.

For Scenario 2, the improvement is less dramatic, but even for moderate uncertainty, CV = 0.25, the improvement is over 3%, and so network designers should consider variance in the demand forecasts.

6. Conclusions and Future Research

Using mathematical programming for telecommunications network design is prevalent in the literature, but very little research has been reported on stochastic models for cellular networks. This paper develops a stochastic programming model that allocates channel equivalents to optimize revenue given a set of constructed radio towers in a CDMA network. The model is based upon a popular yield management model from the airline industry, and the objective function uses a continuous random variable to model the demand for channels. We provide optimality conditions for our model, and we present an algorithm (SCAOPT) to solve it. Finally, we provide computational results for the Normal and the Gamma distributions and for several variances. SCAOPT solved every problem scenario for both Normal and Gamma distributions and nine values for the variance.

SCAOPT could be used as a planning tool to evaluate the revenue potential of a set of radio towers, and unlike previous literature, our model considers the demand to be random. Ideally, however, it could be incorporated in a solution procedure that jointly optimizes the tower location and subscriber assignment decisions like those described earlier for the deterministic problem (e.g., [1] and [9]). Thus, one area of continuing research is integrating this model with an integer programming model to determine a robust set of locations for the towers. In [18] we propose a Benders' decomposition algorithm to determine the tower locations in the master problem, and use SCAOPT to calculate the optimal revenue in a subproblem. Using the optimal revenue and the dual interference prices from the solution, SCAOPT provides a Benders' optimality cut to the master problem. The algorithm continues until the master problem finds a set of locations that satisfy every optimality cut generated from SCAOPT.

Another interesting direction for future research would be to extend our stochastic model to the *SIR-based power control model* described by Amaldi et al. [1]. With SIR-based power control, the strength of the received signal required for users in market m to connect to their assigned tower is a variable, p_m, rather than the fixed value of P_{target}. In this model, the total power received at tower ℓ is

$$P_\ell^{TOT} = \sum_{m \in M} p_m x_m g_{m\ell}.$$

If we consider a single user in market m assigned to tower ℓ_m, then the signal strength for that user is p_m and the interference is $(p_m(x_m - 1)g_m) + \sum_{i \in M \setminus \{m\}} (p_i x_i g_{i\ell_m})$. Thus, the QoS constraint becomes

$$\frac{p_m}{(p_m(x_m - 1)g_m) + \sum_{i \in M \setminus \{m\}} (p_i x_i g_{i\ell_m})} \geq SIR_{min}. \tag{34}$$

The advantage of this model over using a preset value of P_{target} is that by adjusting the handset transmission power according to the actual assignment of markets to towers, it is possible to reduce the total interference in the system. Amaldi et al. [1] compared the two models and found that the SIR-based power control model produced solutions with fewer towers than those found using the fixed P_{target} value. The disadvantage of the model is that constraint (34) is nonlinear. In the context of our model, the QoS constraints resulting from (34) would be

$$\frac{p_m}{(p_m(x_m-1)g_m)+\sum_{i\in M\setminus\{m\}}(p_i x_i g_{i\ell_m})} \geq z_m SIR_{min} \qquad \forall m \in M \tag{35}$$

where z_m is a binary variable that is equal to one whenever x_m is greater than zero, and zero otherwise. Our analysis of these constraints indicates that the feasible region they define may not be convex, and so a local optimum is not necessarily a global optimum. Thus, an exact solution procedure appears to be beyond the capabilities of the current state-of-the art of mathematical programming techniques.

Traditional models and our stochastic model assume that customers in every market pay the same price. Another area of future research would be using different prices for different markets. An analogy in the airline industry is how business and leisure travellers pay different fares for the same flight. Moreover, using a variety of prices allows us to model several discrete points on the demand curve for each market. For example, suppose a demand curve for a market m has one point on the demand curve with price r_1 and mean demand d_1 and another higher priced point with price r_2 and mean demand d_2. We could split market m into two markets m_1 and m_2 with the same attenuation factors as market m, and market m_1 has price r_1 and mean demand d_1, while m_2 has price r_2 and mean demand d_2.

Acknowledgements

This research was supported in part by the Office of Naval Research Award Number N00014-96-1-0315.

References

[1] E. Amaldi, A. Capone, and F. Malucelli. Planning UMTS base station location: Optimization models with power control algorithms. *IEEE Transactions on Wireless Com-*

munications, 2:939–952, 2003.

[2] D. P. Bertsekas. *Nonlinear Programming.* Athena Scientific, Belmont, Massachusetts, 1995.

[3] G. Birkan, J. Kennington, E. Olinick, A. Ortynski, and G. Spiride. Making a case for using integer programming to design DWDM networks. *Optical Networks Magazine*, 4:107–120, 2003.

[4] E. D'Sylvia. O and D seat assignment to maximize expected revenue. Technical report, Boeing Commercial Airplane Company, 1982.

[5] A. Farmehr. Design strategies for both survivable and non-survivable personal communication networks. Praxis Report, School of Engineering and Applied Science, Southern Methodist University, Dallas, TX, 2000.

[6] M. Galota, C. Glasser, S. Reith, and H. Vollmer. A polynomial-time approximation scheme for base station positioning in UMTS networks. In *Proceedings of the 5th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, pages 52–59. ACM, 2001.

[7] M. Gryseels, L. Sorbello, and P. Demeester. Network planning in uncertain dynamic environments. In *Networks 2000, 9th International Telecommunication Network Planning Symposium*, 2000. Published on CD ROM.

[8] D. Günther. *Airline Yield Management.* PhD thesis, Georgia Institute of Technology, School of Industrial and Systems Engineering, 1998.

[9] J. Kalvenes, J. Kennington, and E. Olinick. Base station location and service assignments in W-CDMA networks. *INFORMS Journal on Computing (to appear).*

[10] J. Kennington, K. Lewis, E. Olinick, A. Ortynski, and G. Spiride. Robust solutions for the WDM routing and provisioning problem: models and algorithms. *Optical Networks Magazine*, 4:74–84, 2003.

[11] M. Laguna. Applying robust optimization to capacity expansion of one location in telecommunications with demand uncertainty. *Management Science*, 44:5101–5110, November 1998.

[12] A. Lisser, A. Ouorou, J.-Ph. Vial, and J. Gondzio. Capacity planning under uncertain demand in telecommunications networks. Available on line at http://citeseer.nj.nec.com/lisser99capacity.html.

[13] R. Mathar and M. Schmeink. Optimal base station positioning and channel assignment for 3g mobile networks by integer programming. *Annals of Operations Research*, 107:225–236, 2001.

[14] J. M. Mulvey, R. J. Vanderbei, and S. A. Zenios. Robust optimization of large-scale systems. *Operations Research*, 43:264–281, March-April 1995.

[15] R. Murphey, P. Pardalos, and G. Resende. Frequency assignment problems. In D. Zu and P. Pardalos, editors, *Handbook of Combinatorial Optimization*, pages 1–55. Kluwer Academic Publishers, 1999.

[16] M. Riis and K. A. Andersen. Capacitated network design with uncertain demands. *INFORMS Journal on Computing*, 14:247–260, 2002.

[17] M. Riis and J. Lodahl. A bicriteria stochastic programming model for capacity expansion in telecommunications. *Mathematical Methods of Operations Research*, 56:83–100, 2002.

[18] Jay Rosenberger and Eli Olinick. Robust tower location for CDMA networks. Technical Report 04-EMIS-11, Southern Methodist University, Dallas, TX, 2004.

[19] S. Sen, R. Doverspike, and S. Cosares. Network planning with random demand. *Telecommunication Systems*, 3:11–30, 1994.

[20] B. Smith, J. Leimkuhler, and R. Darrow. Yield management at American Airlines. *Interfaces*, 22:8–31, 1992.

[21] J. C. Smith, A. J. Schaefer, and J. Yen. A stochastic intra-ring synchronous optical network design problem. *Networks*, 44:12–26, 2004.

[22] E. L. Johnson V. C. P. Chen, D. Günther. A markov decision problem based approach to the airline ym problem. Technical Report TLI/LEC-98-12, Georgia Institute of Technology, Atlanta, GA, 1998.

Table 1: Scenario 1, Gamma distribution

CV	Step Size	k^*	$Stoch(\overline{x}^{k^*})$ (\$)	$\|\|\lambda^{k^*} \cdot \delta^{k^*}\|\|$ (\$)	gap^{k^*}
0.01	A	594	4,062,436	2,296	<0.01%
0.01	B	300,000	4,061,684	835	0.02%
0.01	C	4,576	4,062,060	403	<0.01%
0.05	A	594	4,034,394	1,600	<0.01%
0.05	B	843	4,034,175	1,027	<0.01%
0.05	C	6,437	4,034,178	908	<0.01%
0.10	A	14,218	3,988,431	204	<0.01%
0.10	B	12	3,988,512	652	<0.01%
0.10	C	5,538	3,988,432	466	<0.01%
0.15	A	1,165	3,936,098	1,600	<0.01%
0.15	B	884	3,936,244	663	<0.01%
0.15	C	4,639	3,936,038	384	<0.01%
0.25	A	1,181	3,813,610	562	<0.01%
0.25	B	715	3,813,376	2,184	<0.01%
0.25	C	3,791	3,813,266	352	<0.01%
0.50	A	966	3,460,042	1,992	<0.01%
0.50	B	1,319	3,460,318	854	<0.01%
0.50	C	3,616	3,459,997	346	<0.01%
1.00	A	990	2,769,217	2,035	<0.01%
1.00	B	1,788	2,769,180	331	<0.01%
1.00	C	3,369	2,768,986	227	<0.01%
1.50	A	1,148	2,199,119	343	<0.01%
1.50	B	1,999	2,198,961	522	<0.01%
1.50	C	3,259	2,198,929	163	<0.01%
2.00	A	1,127	1,762,067	410	<0.01%
2.00	B	2,161	1,762,064	181	<0.01%
2.00	C	3,200	1,761,944	127	<0.01%

Table 2: Scenario 2, Gamma distribution

CV	Step Size	k^*	$Stoch(\overline{x}^{k^*})$ (\$)	$\|\|\lambda^{k^*} \cdot \delta^{k^*}\|\|$ (\$)	gap^{k^*}
0.01	A	300,000	65,082,245	473,734	0.03%
0.01	B	4,003	65,098,852	192,956	<0.01%
0.01	C	806	65,099,451	672,031	<0.01%
0.05	A	300,000	64,522,454	259,059	0.09%
0.05	B	355	64,570,558	540,035	<0.01%
0.05	C	1,450	64,569,576	424,117	<0.01%
0.10	A	300,000	63,750,216	105,521	0.01%
0.10	B	104	63,752,784	271,524	<0.01%
0.10	C	777	63,751,535	260,034	<0.01%
0.15	A	427	62,832,919	146,171	<0.01%
0.15	B	166	62,836,138	85,611	<0.01%
0.15	C	597	62,833,024	83,781	<0.01%
0.25	A	297	60,862,149	65,482	<0.01%
0.25	B	66	60,864,368	71,528	<0.01%
0.25	C	744	60,862,176	56,219	<0.01%
0.50	A	289	55,647,783	27,808	<0.01%
0.50	B	291	55,644,776	2,959	<0.01%
0.50	C	321	55,644,392	1,531	<0.01%
1.00	A	559	45,146,500	14,684	<0.01%
1.00	B	2,052	45,147,380	2,401	<0.01%
1.00	C	200	45,146,460	1,034	<0.01%
1.50	A	722	36,010,548	3,699	<0.01%
1.50	B	2,564	36,009,442	3,202	<0.01%
1.50	C	144	36,008,987	620	<0.01%
2.00	A	752	28,898,441	13,743	<0.01%
2.00	B	3,437	28,897,423	4,006	<0.01%
2.00	C	109	28,897,269	483	<0.01%

Table 3: Scenario 1, Normal distribution

CV	Step Size	k^*	$Stoch(\overline{x}^{k^*})$ (\$)	$\|\|\lambda^{k^*} \cdot \delta^{k^*}\|\|$ (\$)	gap^{k^*}
0.01	A	594	4,062,457	2,322	<0.01%
0.01	B	300,000	4,061,687	866	0.02%
0.01	C	4,576	4,062,081	390	<0.01%
0.05	A	594	4,034,513	1,976	<0.01%
0.05	B	958	4,034,421	1,048	<0.01%
0.05	C	6,512	4,034,425	933	<0.01%
0.10	A	768	3,989,522	1,673	<0.01%
0.10	B	48	3,989,434	879	<0.01%
0.10	C	5,495	3,989,428	566	<0.01%
0.15	A	1,223	3,938,437	449	<0.01%
0.15	B	837	3,938,564	617	<0.01%
0.15	C	4,692	3,938,426	417	<0.01%
0.25	A	942	3,820,963	5,605	<0.01%
0.25	B	589	3,820,933	2,305	<0.01%
0.25	C	3,901	3,820,888	816	<0.01%
0.50	A	928	3,476,417	1,223	<0.01%
0.50	B	615	3,476,258	534	<0.01%
0.50	C	3,796	3,476,122	368	<0.01%
1.00	A	1,056	2,734,512	926	<0.01%
1.00	B	843	2,734,303	2,001	<0.01%
1.00	C	3,491	2,734,278	300	<0.01%
1.50	A	1,058	1,973,115	977	<0.01%
1.50	B	842	1,973,238	1,019	<0.01%
1.50	C	3,385	1,973,070	183	<0.01%
2.00	A	1,093	1,198,675	490	<0.01%
2.00	B	1,398	1,198,670	1,315	<0.01%
2.00	C	3,341	1,198,663	92	<0.01%

Table 4: Scenario 2, Normal distribution

CV	Step Size	k^*	$Stoch(\overline{x}^{k^*})$ (\$)	$\|\|\lambda^{k^*} \cdot \delta^{k^*}\|\|$ (\$)	gap^{k^*}
0.01	A	300,000	65,082,245	199,190	0.03%
0.01	B	466	65,097,739	1,363,504	<0.01%
0.01	C	718	65,099,474	417,385	<0.01%
0.05	A	300,000	64,525,139	213,870	0.08%
0.05	B	291	64,569,225	244,722	<0.01%
0.05	C	1,480	64,568,994	143,980	<0.01%
0.10	A	705	63,744,564	177,816	<0.01%
0.10	B	142	63,745,332	278,668	<0.01%
0.10	C	769	63,744,731	277,863	<0.01%
0.15	A	413	62,813,772	129,939	<0.01%
0.15	B	312	62,813,621	40,350	<0.01%
0.15	C	609	62,813,721	130,150	<0.01%
0.25	A	302	60,788,742	49,247	<0.01%
0.25	B	838	60,788,411	5,134	<0.01%
0.25	C	450	60,788,410	4,344	<0.01%
0.50	A	253	55,047,140	31,655	<0.01%
0.50	B	100	55,047,019	97,418	<0.01%
0.50	C	292	55,045,274	23,761	<0.01%
1.00	A	469	40,596,887	80,132	<0.01%
1.00	B	1,655	40,592,883	75,422	<0.01%
1.00	C	205	40,593,692	987	<0.01%
1.50	A	651	23,867,799	38,914	<0.01%
1.50	B	2,368	23,866,419	69,353	<0.01%
1.50	C	178	23,866,530	647	<0.01%
2.00	A	717	6,126,113	40,626	<0.01%
2.00	B	2,074	6,124,402	49,194	<0.01%
2.00	C	145	6,123,414	25,385	<0.01%

Table 6: Scenario 1, Stochastic vs. Deterministic Solution Quality

CV	Distribution	$Stoch(\overline{x}^{STOCH})$ (\$)	$Stoch(\overline{x}^{DET})$ (\$)	% improvement
0.01	Gamma	4,062,436	4,051,672	0.3%
0.01	Normal	4,062,457	4,051,671	0.3%
0.05	Gamma	4,034,394	3,986,774	1.2%
0.05	Normal	4,034,513	3,986,757	1.2%
0.10	Gamma	3,988,512	3,905,749	2.1%
0.10	Normal	3,989,522	3,905,614	2.1%
0.15	Gamma	3,936,244	3,824,927	2.9%
0.15	Normal	3,938,564	3,824,471	3.0%
0.25	Gamma	3,813,610	3,664,293	4.1%
0.25	Normal	3,820,963	3,662,186	4.3%
0.50	Gamma	3,460,318	3,273,167	5.7%
0.50	Normal	3,476,417	3,256,471	6.8%
1.00	Gamma	2,769,217	2,571,403	7.7%
1.00	Normal	2,734,512	2,445,043	11.8%
1.50	Gamma	2,199,119	2,014,265	9.2%
1.50	Normal	1,973,238	1,633,614	20.8%
2.00	Gamma	1,762,067	1,596,400	10.4%
2.00	Normal	1,198,675	822,185	45.8%

Table 7: Scenario 2, Stochastic vs. Deterministic Solution Quality

CV	Distribution	$Stoch(\overline{x}^{STOCH})$ (\$)	$Stoch(\overline{x}^{DET})$ (\$)	% improvement
0.01	Gamma	65,099,451	64,954,851	0.2%
0.01	Normal	65,099,474	64,954,849	0.2%
0.05	Gamma	64,570,558	63,970,120	0.9%
0.05	Normal	64,569,225	63,969,657	0.9%
0.10	Gamma	63,752,784	62,726,626	1.6%
0.10	Normal	63,745,332	62,722,857	1.6%
0.15	Gamma	62,836,138	61,472,594	2.2%
0.15	Normal	62,813,772	61,460,710	2.2%
0.25	Gamma	60,864,368	58,953,431	3.2%
0.25	Normal	60,788,742	58,905,667	3.2%
0.50	Gamma	55,647,783	52,745,931	5.5%
0.50	Normal	55,047,140	52,272,470	5.3%
1.00	Gamma	45,147,380	41,475,553	8.9%
1.00	Normal	40,596,887	36,522,292	11.2%
1.50	Gamma	36,010,548	32,484,640	10.9%
1.50	Normal	23,867,799	18,621,424	28.2%
2.00	Gamma	28,898,441	25,736,797	12.3%
2.00	Normal	6,126,113	-125,792	NA