[14] F. H. Walters, L. R. Parker, S. L. Morgan, and S. N. Deming (1991). *Sequential Simplex Optimization*. CRC Press, Boca Raton, USA.

[15] M. H. Wright (1996). Direct search methods: once scorned, now respectable. in *Numerical Analysis 1995: Proceedings of the 1995 Dundee Biennial Conference in Numerical Analysis*, D. F. Griffiths and G. A. Watson (eds.), 191–208, Addison Wesley Longman, Harlow, UK.

method [1], [3]. If we wish to minimize $f(\mathbf{x})$ subject to the equality constraints $c_i(\mathbf{x}) = 0$, $i = 1, \ldots, m$, the associated ℓ_1 (absolute-value) penalty function is

$$P(\mathbf{x}, \tau) = f(\mathbf{x}) + \tau \sum_{i=1}^{m} |c_i(\mathbf{x})|, \tag{5}$$

where $\tau > 0$ is the *penalty parameter*. If minimizing $f(\mathbf{x})$ subject to the inequality constraints $c_i(\mathbf{x}) \geq 0$, $i = 1, \ldots, m$, the corresponding ℓ_1 penalty function is

$$P(\mathbf{x}, \tau) = f(x) + \tau \sum_{i=1}^{m} \max(0, -c_i(\mathbf{x})). \tag{6}$$

The idea in either case is that a new composite objective function is created by adding a positive multiple of the constraint violations to the original objective function. It can be shown (see, e.g., [1], [3]) that, for a sufficiently large finite value of τ, an unconstrained minimizer of the penalty function is a solution of the constrained problem.

Although nondifferentiable penalty functions have certain theoretical advantages such as the finiteness of the required penalty parameter, they are sometimes considered undesirable because of their lack of smoothness. However, this is not an issue with direct search methods, which, as observed earlier, do not require smoothness. According to folklore, direct search methods are effective in minimizing penalty functions, and our favorable experience with (5) and (6) in the wireless application confirms this property. Thus both constrained and unconstrained problems can be handled with the same basic approach.

Direct search methods cannot be guaranteed to converge to the global (i.e., best overall) optimum except for certain highly restricted problem classes. In general, even very simple forms of the base station location problem have multiple local optima. Standard approaches (see, e.g., [3], [14]) are to apply the Nelder-Mead algorithm with multiple starting points, or to restart the method with a large initial simplex near a previously found local optimum. Given the generic complexity of global optimization, further research is needed to develop specialized techniques for finding global optima in the wireless problem.

VI. SUMMARY

The base placement problem in Wise [2] is only one of a wide variety of optimization problems that may arise in wireless systems. For example, base station power levels may be treated as variables to be optimized, or a system design may be sought that minimizes cost while satisfying performance constraints. Our experience during four years of extensive numerical testing strongly suggests that the Nelder-Mead method, or some other direct search method, will be highly effective for many of these formulations, particularly as reliable problem-specific initialization heuristics are developed.

REFERENCES

[1] R. Fletcher (1987). *Practical Methods of Optimization* (second edition), John Wiley and Sons, Chichester.

[2] S. J. Fortune, D. M. Gay, B. W. Kernighan, O. Landron, R. A. Valenzuela, and M. H. Wright (1995). WISE design of indoor wireless systems: practical computation and optimization, *IEEE Computational Science and Engineering* 2, 58–68.

[3] P. E. Gill, W. Murray, and M. H. Wright (1981), *Practical Optimization*, Academic Press, London and New York.

[4] C. T. Kelley (1997). Detection and remediation of stagnation in the Nelder-Mead algorithm using a sufficient decrease condition, Technical Report, Department of Mathematics, North Carolina State University, Raleigh, North Carolina, USA.

[5] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright (1997). Convergence properties of the Nelder-Mead simplex algorithm in low dimensions, Technical Report 96-4-07, Computing Sciences Research Center, Bell Laboratories, Murray Hill. New Jersey 07974. To appear, *SIAM Journal on Optimization.*

[6] J. C. Lagarias and M. H. Wright (1998). Explaining and improving the performance of the Nelder-Mead method for unconstrained optimization, to appear.

[7] R. M. Lewis and V. Torczon (1996). Pattern search algorithms for bound-constrained minimization, Technical report 96–20, ICASE, Hampton, Virginia.

[8] R. M. Lewis and V. Torczon (1997). Pattern search algorithms for linearly constrained minimization, Technical report 97–15, ICASE, Hampton, Virginia.

[9] K. I. M. McKinnon (1996). Convergence of the Nelder-Mead simplex method to a non-stationary point, Technical report, Department of Mathematics and Computer Science, University of Edinburgh, Edinburgh, UK. To appear, *SIAM Journal on Optimization.*

[10] J. A. Nelder and R. Mead (1965). A simplex method for function minimization, *Computer Journal* 7, 308–313.

[11] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vettering (1988). *Numerical Recipes in C*, Cambridge University Press, Cambridge, UK.

[12] V. Torczon (1989). *Multi-directional Search: A Direct Search Algorithm for Parallel Machines*. Ph.D. thesis, Rice University, Houston, Texas, USA.

[13] V. Torczon (1997). On the convergence of pattern search algorithms, *SIAM J. Optimization* 7, 1–25.

where $\bar{\mathbf{x}} = \sum_{i=1}^{n} \mathbf{x}_i / n$ is the centroid of the n best points (all vertices except for $\mathbf{x}_{n+1}$). Evaluate $f_r = f(\mathbf{x}_r)$.

If $f_1 \le f_r < f_n$, accept the reflected point $\mathbf{x}_r$ and terminate the iteration.

3. **Expand.**

If $f_r < f_1$, calculate the *expansion point* $\mathbf{x}_e$:

$$\mathbf{x}_e = \bar{\mathbf{x}} + \rho\chi(\bar{\mathbf{x}} - \mathbf{x}_{n+1}), \tag{2}$$

and evaluate $f_e = f(\mathbf{x}_e)$. If $f_e < f_r$, accept $\mathbf{x}_e$ and terminate the iteration; otherwise (if $f_e \ge f_r$), accept $\mathbf{x}_r$ and terminate the iteration.

4. **Contract.**

If $f_r \ge f_n$, perform a *contraction* between $\bar{\mathbf{x}}$ and the better of $\mathbf{x}_{n+1}$ and $\mathbf{x}_r$.

a. Outside. If $f_n \le f_r < f_{n+1}$ (i.e., $\mathbf{x}_r$ is strictly better than $\mathbf{x}_{n+1}$), perform an *outside contraction*: calculate

$$\mathbf{x}_c = \bar{\mathbf{x}} + \gamma\rho(\bar{\mathbf{x}} - \mathbf{x}_{n+1}), \tag{3}$$

and evaluate $f_c = f(\mathbf{x}_c)$. If $f_c \le f_r$, accept $\mathbf{x}_c$ and terminate the iteration; otherwise, go to Step 5 (perform a shrink).

b. Inside. If $f_r \ge f_{n+1}$, perform an *inside contraction*: calculate

$$\mathbf{x}_{cc} = \bar{\mathbf{x}} - \gamma(\bar{\mathbf{x}} - \mathbf{x}_{n+1}), \tag{4}$$

and evaluate $f_{cc} = f(\mathbf{x}_{cc})$. If $f_{cc} < f_{n+1}$, accept $\mathbf{x}_{cc}$ and terminate the iteration; otherwise, go to Step 5 (perform a shrink).

5. **Perform a shrink step.**

Evaluate f at the n points $\mathbf{v}_i = \mathbf{x}_1 + \sigma(\mathbf{x}_i - \mathbf{x}_1)$, $i = 2, \ldots, n+1$. The (unordered) vertices of the simplex at the next iteration consist of $\mathbf{x}_1, \mathbf{v}_2, \ldots, \mathbf{v}_{n+1}$.

Shrink steps are extremely rare in practice [12], which means that the NM method typically requires only one or two objective function values per iteration. In contrast, an iteration of other direct search methods requires n (the number of variables) or more evaluations of the objective function.

Typical termination criteria for the NM method involve a sufficiently small simplex and/or a sufficiently small variation in the function values at the simplex vertices; see [12] and [6] for further comments about how to terminate direct search methods. There is a trove of folklore, practical experience, and rules of thumb about issues such as choosing the initial simplex shape and size; see, e.g., [14].

Relatively strong convergence results have been obtained [13] for *pattern search* methods (a subset of direct search methods), which require, in the worst case, $n+1$ function evaluations per iteration. Surprisingly, no published theoretical analysis treated the *original* NM algorithm in the 30 years following its publication. The first such results for Nelder-Mead appeared in 1996 [5], [9], but are limited to dimensions one and two. However, additional convergence theory is likely to be developed in the near future because of rapidly growing interest in the topic of non-derivative optimization.

Beyond the lack of general convergence theory, the NM method is known to suffer in practice from serious, even fatal, stagnation and failure [12], [4]. Despite these known inefficiencies and failures, the method remains extraordinarily, and rightly, popular for two reasons that apply directly in the present context of wireless design. First, the Nelder-Mead algorithm typically produces significant improvement in the objective function during the first few iterations [6]; thus it is efficient in applications where one wants to find parameter values that improve (rather than optimize) a performance measure. Second, the relative "best-case efficiency" of the NM method is appealing when each function evaluation is enormously expensive or time-consuming.

Our extensive experience with the propagation model in Wise [2] shows that Nelder-Mead method is effective and reliable in finding acceptable approximations to local optima within a small number of function evaluations. As a rule of thumb, with one base station (three variables), only 20–25 function evaluations are typically needed to satisfy rather stringent termination criteria (the maximum simplex edge must be less than 0.1 meter, or the maximum difference between function values at the vertices must be less than 0.001). In numerous trials with four base stations for which the z coordinates were fixed (eight variables), an acceptably accurate optimum was almost always located in 40–45 function evaluations.

As discussed in [14], the size of the initial simplex has a substantial effect on efficiency. For almost all practical problems, the advice in [14] is to begin the NM method with a relatively large initial simplex. A too-small initial simplex can lead in the worst case to premature termination, and, even when the method continues to iterate, progress can be extremely slow.

V. EXTENSIONS

Most direct search methods are designed for unconstrained optimization of a nonlinear function f. Recently, [7], [8] have shown how to adapt pattern search methods to handle both simple bound constraints (i.e., constraints of the form $\mathbf{x}_i \ge \ell_i$ and/or $\mathbf{x}_i \le u_i$) and general linear constraints ($A\mathbf{x} = b$ or $Ax \ge b$, where A is a matrix and b a vector of appropriate dimension). The approaches given in [7] and [8] rely completely on the *a priori* known structure of the constraints.

A generic technique for transforming a constrained problem into an unconstrained problem is to use an *exact penalty function*, of which there are several varieties. In the wireless application, we have used a nondifferentiable exact penalty specified level. The treatment of constraints will be discussed briefly in Section V.

III. PROPERTIES OF THE OBJECTIVE FUNCTION

For indoor Wise (see [2]), the objective function is the percentage of the building (or target subarea of the building) within which received power reaches or exceeds a specified threshold. This value is obtained by computing received power at "virtual receivers" on a grid of points (typically thousands) in the building. The objective function value, "coverage", is the ratio of grid points covered to the total number of grid points.

This function has several properties that are unattractive from the viewpoint of numerical optimization. First, since it is defined as the ratio of two (not very large) integers, it can assume only a relatively small finite set of values, which means that its behavior is quite different from that of a finite-precision approximation to a smooth function. Second, it is unpredictably nondifferentiable (in some cases, even discontinuous) because the coverage can change dramatically following a tiny change in a base station location—for example, when a base moves from just inside to just outside a room with non-transmissive walls. Third, the form of some propagation models introduces discontinuities; most commonly, this happens when physical quantities that would ideally be continuous are approximated by table lookup. Another form of discontinuity arises in situations when the value of one parameter depends on the range in which a second parameter lies. Even if present at a very low level, specifications of the form

$$\text{if } 0 \le \beta \le 2 \text{ then } \alpha \leftarrow 1 \quad \text{else } \alpha \leftarrow 2.5$$

inherently introduce discontinuities, albeit small, that can significantly affect the efficiency of optimization methods based on smoothness.

IV. THE NELDER-MEAD METHOD

Because of the properties just described, there is no hope of obtaining the gradient of the objective function in the base station placement problem, or even a finite-difference approximation to its gradient: only function values are available. Thus the workhorses of optimization like Newton and quasi-Newton methods (see, e.g., [1], [3]), which require first or second derivatives, cannot be applied. Furthermore, computation of a single function value can be extremely costly for large, complex buildings and campus environments, even with the fastest available computational geometry techniques (such as those in Wise [2]). Thus it is essential for the optimization method to produce an acceptable answer in a relatively small number of function evaluations. (Asymptotic properties, such as ultimate superlinear convergence, are irrelevant in this context.)

The most plausible solution technique is a *direct search method*, one of a family of optimization methods that have been popular with practitioners since the 1960s [11], [15], particularly in engineering. Although direct search methods were in large part ignored by the mainstream optimization community until recently, they are now undergoing a renaissance of intense interest, with active development during the past two years of the underlying theory and algorithmic improvements (see, for example, [13], [15], [4], [6]).

The direct search method used in Wise [2] is a version of the Nelder-Mead (NM) "simplex" method [10], first published in 1965. The NM method is without question the most popular direct search method, used in solving tens of thousands of problems, especially in chemistry, chemical engineering, and medicine. The book [14] is devoted entirely to the NM method and variations. The NM method appears in the best-selling handbook *Numerical Recipes* [11], where it is called the "amoeba algorithm", and in the Matlab™ interactive matrix software package.

The Nelder-Mead method is conceptually simple, although care is needed to implement it properly. The NM method attempts to minimize a scalar-valued nonlinear function of n real variables using only function values, without any derivative information (explicit or implicit). Many of the most well known direct search methods, including the NM method, maintain at each step a nondegenerate *simplex*, a geometric figure defined by $n+1$ vertices (real n-vectors). (For example, a simplex in two dimensions is a triangle.)

Suppose that we are minimizing the function $f(\mathbf{x})$, where $\mathbf{x}$ denotes a real n-vector (a point in n-space). The NM method includes four possible operations on the current simplex, each associated with a coefficient: reflection (ρ), expansion (χ), contraction (γ), and shrinkage (σ). The nearly universal choices for these coefficients in the standard NM algorithm are

$$\rho = 1, \quad \chi = 2, \quad \gamma = \tfrac{1}{2}, \quad \text{and} \quad \sigma = \tfrac{1}{2}.$$

The result of each NM iteration is either: (1) a single new vertex—the *accepted point*—which replaces the current worst vertex (the vertex with the largest function value) in the set of vertices for the next iteration; or (2) if a shrink is performed, a set of n new points that, together with the previous best vertex, form the simplex at the next iteration.

A single iteration of the NM method is defined according to the following steps.

1. **Order.**

   Order the $n+1$ vertices to satisfy

   $$f(\mathbf{x}_1) \le f(\mathbf{x}_2) \le \cdots \le f(\mathbf{x}_{n+1}),$$

   using a consistent tie-breaking rule (see [5]).

2. **Reflect.**

   Compute the *reflection point* $\mathbf{x}_r$ from

   $$\mathbf{x}_r = \bar{\mathbf{x}} + \rho(\bar{\mathbf{x}} - \mathbf{x}_{n+1}), \tag{1}$$

OPTIMIZATION METHODS FOR BASE STATION PLACEMENT IN WIRELESS APPLICATIONS

Margaret H. Wright

Bell Laboratories, Lucent Technologies
Murray Hill, New Jersey 07974, USA
email: mhw@research.bell-labs.com

Abstract — Current wireless design tools contain one or more propagation models whose purpose is to predict quantitative measures of system performance as a function of system parameters. A logical next step is to *optimize* predicted overall performance by choosing the best values for parameters that can be controlled by the provider and/or maintainer of the system. However, because the function to be optimized is often highly nonlinear and nasty (for example, unpredictably discontinuous, nondifferentiable, or noisy), only a limited number of optimization techniques can be applied.

Direct search methods are well suited to finding the optimal placement of base stations, since they require only the value of the function to be optimized. In extensive numerical testing, we have found a customized variant of the Nelder-Mead "simplex" method to be reasonably efficient and reliable at finding local optima. Our experience with producing improved or optimal base station placements suggests that optimization can enhance understanding of propagation models as well as improve wireless system designs.

I. THE BASE STATION PLACEMENT PROBLEM

Propagation models of widely varying degrees of sophistication, ranging from simple statistical models to full three-dimensional ray and/or beam tracing, provide detailed predictions of quantities that characterize performance (for example, the received power at a given location) within wireless systems in both indoor and outdoor environments. These models typically depend on many parameters—some dictated in advance, others that can be varied by the provider or maintainer of the system.

Given a propagation model with acceptable predictive capability and an initial design, it is desirable to produce an improved (ideally, the optimal) design within the framework of that model by systematically adjusting free system parameters to achieve the best performance. We consider the specific problem of optimizing the *base locations* in a fixed-base system; these locations are the most obvious, perhaps the most important, free parameters. Other variables, such as power, antenna tilt, etc., seem likely to be amenable to the same general approach described in this paper.

II. FORMULATION OF THE PROBLEM

A crucial question relevant to optimization but not considered here is how to define appropriate numerical measures of the desirable qualities in a wireless system. Standard optimization methods require a *single number* (the objective function) that represents the "goodness" of the system's performance. For a measure such as the cost of equipment and maintenance, we seek to reduce its value; for others, such as the size of the coverage region, larger values are better. The performance of a realistic system involves numerous quantities, sometimes incommensurate, which means that the problem of characterizing good performance may be just as complicated as finding its optimal value. Fortunately, the optimization methods considered here can be applied to essentially any objective function, and furthermore their performance should not depend heavily on the choice of parameters.

Implicit in the concept of optimization is that the chosen performance measure depends on a set of *variables*, i.e., quantities that can be chosen or modified in the system design. In the WiSE model [2], the usual variables are the x, y, and z coordinates of the base stations. In some instances only a subset of these coordinates may be allowed to vary, for example when all antennas will be installed at a fixed height. Other possible variables include features such as the tilt angle or power level of each individual antenna.

In addition to selecting the objective function variables, wireless system design may involve *constraints* arising from physical, contractual, or legal requirements. A common example of such a constraint is the existence of site-specific "forbidden regions" within the building (such as the chief executive's office or a hospital operating room) where base stations cannot be placed. Another form of constraint specifies that power received within a certain zone outside the building perimeter must not exceed a