S^{4}W: Globally Optimized Design of Wireless Communication Systems

A. Verstak, J. He, L. T. Watson,
N. Ramakrishnan, and C. A. Shaffer
Department of Computer Science
Virginia Polytechnic Institute
and State University
Blacksburg, VA 24061-0106

T. S. Rappaport, C. R. Anderson,
K. Bae, J. Jiang, and W. H. Tranter
Bradley Department of Electrical and Computer
Engineering, Virginia Polytechnic Institute
and State University
Blacksburg, VA 24061-0111

Abstract

In this paper, a global optimization technique is applied to solve the optimal transmitter placement problem for indoor wireless systems. An efficient pattern search algorithm—DIRECT (DIviding RECTangles) of Jones, Perttunen, and Stuckman (1993)—has been connected to a parallel 3D radio propagation ray tracing modeler running on a 200-node Beowulf cluster of Linux workstations. Surrogate functions for a parallel WCDMA (wideband code division multiple access) simulator were used to estimate the system performance for the global optimization algorithm. Power coverage and BER (bit error rate) are considered as two different criteria for optimizing locations of a specified number of transmitters across the feasible region of the design space. This paper briefly describes the underlying radio propagation and WCDMA simulations and focuses on the design issues of the optimization loop.

1. Introduction

Optimal transmitter placement provides high spectral efficiency and system capacity while reducing network costs, which are the key criteria for wireless network planning. As the complexity and popularity of modern wireless networks increases, automatic transmitter placement provides cost savings when compared to the traditional human process of site planning. Automatic design tools are being developed to offer efficient and optimal planning solutions. Besides [3], [6], and [13], S^4W (Site-Specific System Simulator for Wireless system design) is among the few known wireless system tools for in-building network design. It is being developed jointly by the Mobile & Portable Radio Research Group (MPRG) and the Problem Solving Environment (PSE) research group at Virginia Polytechnic Institute & State University. An optimization loop in S^4W is proposed to maximize the efficiency of simulated channel models and surrogate functions are proposed to reduce the cost of simulations. Transmitter placement optimization is one

Figure 1.1. Durham Hall, fourth floor.

specific problem that can be solved by S^4W. An example of an S^4W model consisting of a propagation model, a wireless system model, and an optimizer is given in [17].

In general, transmitter placement optimization is aimed at ensuring an acceptable level of wireless system performance within a geographical area of interest (Figure 1.1 shows an indoor environment for the present study) at a minimum cost. [3] considers the major performance factor to be the power coverage, defined as the ratio of the number of receiver locations with received power above the threshold to the total number of receiver locations. This nonsmooth function leads to the rank based methods used by [3]. In [6] and [14], the objective function is based on several weighted factors, such as covered area, interference area, and mean signal path loss. In the present work, two performance metrics form objective functions for optimal transmitter placement. The metrics are continuous penalty functions defined in terms of power levels (i.e., power coverage) and bit error rates of receiver locations within the covered region. Both objective functions are devised to minimize the average shortfall of the estimated performance metric with respect to the corresponding threshold. 3D ray tracing is used as a deterministic propagation model to estimate power coverage levels and impulse responses within the region of interest for transmitter locations sampled by the optimization algorithm [15][16]. Surrogates for the Monte Carlo WCDMA simulation are used to estimate the BERs (bit error rates) for the second optimization criterion. Both the surrogates and the WCDMA

simulation utilize the impulse responses estimated by the ray tracing model. Since 3D ray tracing and WCDMA simulation are computationally expensive, MPI-based parallel implementations are used in the present work.

The underlying optimization algorithm is known as DIRECT (DIviding RECTangles), a direct search algorithm proposed by Jones et al. [9]. It was proposed as an effective approach to solve global optimization problems subject to simple constraints. Jones et al. [9] named the algorithm after one of its key steps—dividing rectangles. DIRECT is a pattern search method that is categorized as a direct search technique by Lewis et al. [10]. Generally speaking, "pattern search methods are characterized by a series of exploratory moves that consider the behavior of the objective function at a pattern of points" [10], which are chosen as the centers of rectangles in the DIRECT algorithm. This center-sampling strategy reduces the computational complexity, especially for higher dimensional problems. Moreover, DIRECT adopts a strategy of balancing local and global search by selecting potentially optimal rectangles to be further explored. This strategy gives rise to fast convergence with reasonably broad space coverage. These features have motivated its successful application in modern large-scale multidisciplinary engineering problems [18]. The present work is the second known application of DIRECT to wireless communication systems design other than the previous work in [7].

This paper is organized as follows. Section 2 presents the parallel 3D ray tracing model. Section 3 describes the parallel WCDMA simulator and the surrogate functions. An overview of the DIRECT algorithm is given in Section 4, followed by a description of dynamic data structures. In Section 5, optimization results are presented and analyzed. Finally, Section 6 summarizes some key contributions of the present work and suggests directions for future research.

2. Ray Tracing Propagation Model

Received impulse responses are approximated with a 3D ray tracing propagation model that is based on geometrical optics. Electromagnetic waves are modeled as rays that are traced through reflections and transmissions through the walls. Beams [2] are shot from geodesic domes drawn around transmitters. Each beam is a triangular pyramid formed by the point location of the transmitter and one of the triangles on the surface of the dome. Essentially, the spherical wavefront is triangulated and the 3D sphere is split into pyramidal beams. Following the argument in [16], all such beams are disjoint and have nearly the same shape and angular separation. Only the central ray of each beam is traced to identify reflection locations. However,
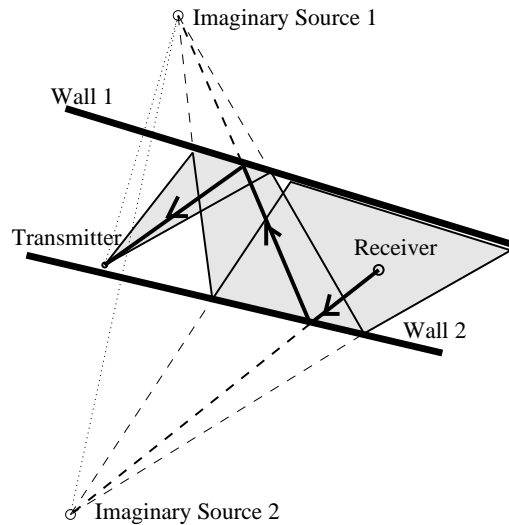

Figure 2.1. 2D beam tracing: a beam (shadowed region) is traced from the transmitter location to the receiver location through two reflections, and then a ray (bold line) is traced back.

the whole beam is used for ray-receiver intersection tests. Once an intersection with a receiver location is detected, a ray will be traced back from the receiver to the transmitter through the sequence of reflections and transmissions (penetrations) encountered by the beam. The illustration of this process in 2D is given in Figure 2.1. Figure 2.2 depicts a fast intersection test of a beam with a grid of receiver locations. Neither diffraction nor scattering are modeled for computational complexity reasons, although these phenomena play an important role in propagation [12]. Octree space partitioning [5] and image parallelism with dynamic scheduling [4] are used to reduce simulation run time.

Although material parameters and incidence angles affect losses in a wireless channel, a constant 6 dB reflection loss (same as in [15]) and a constant 4.6 dB transmission (penetration) loss (the loss for plaster board in [1]) are assumed. The power contribution of each ray, in dBW, is calculated according to the model developed in [16]:

$$P_j = P(d_0) - 20\log_{10}(d/\lambda) - nL_r - mL_t, \quad (2.1)$$

where P_j is the power of the j-th ray, d is the total distance traveled by the ray, $P(d_0)$ is the transmitter power at a reference distance d_0 from the transmitter, n and m are the numbers of reflections and transmissions, $L_r = 6$ dB and $L_t = 4.6$ dB are reflection and transmission losses, and λ is the wavelength.

The ray tracer has been validated and calibrated with a series of measurements in the corridor of the fourth floor of Durham Hall, Virginia Tech. An ultrawideband sliding
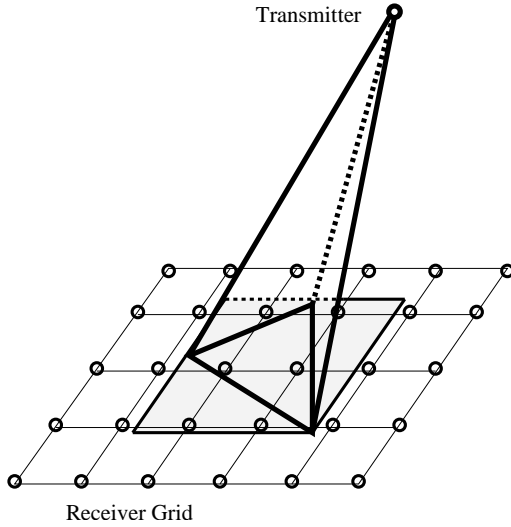
Figure 2.2. Beam intersection with a receiver grid: only the locations inside of the bounding box of the projection of the beam onto the grid (shadowed region) are tested for intersection with the beam pyramid.

correlator channel sounder [12] operating at 2.5 GHz and outfitted with omnidirectional antennas was used to record impulse responses at six separate locations. The sliding correlator utilized an 11-bit, 400 MHz pseudo-noise spreading code for a time domain multipath resolution of 2.5 nanoseconds and a dynamic range of 30 dB. Simulated power delay profiles were post-processed and compared to the measured ones location by location.

Comparing ray tracer output with a physical channel requires accounting for antennas and resampling the signal to match the sampling rate of the measurement system. The same conversion sequence was used for both validation against measurements and interfacing with the WCDMA simulation. The received E-field envelope of ray j (in V/m) that arrived at time t_j is $E_j = \sqrt{\eta 10^{0.1 P_j}}$, where P_j is the output of the ray tracer (in dBW) and $\eta = 120\pi$ Ω is the impedance of free space [12]. To account for antenna directivity, an omnidirectional antenna pattern must be applied to all E_js. The electric field that would be registered at time t_j by a hypothetical measurement system with infinite bandwidth resolution is

$$E'_j = E_j G_t G_r \cos \Theta_t \cos \Theta_r, \qquad (2.2)$$

where Θ_t and Θ_r are ray transmission and reception elevation angles relative to the horizon, and G_t and G_r are maximum transmitter and receiver antenna gains, respectively. Further, the discrete impulse response must be convolved with a Gaussian filter and sampled at uniform time intervals of width δ. The measurement system output samples with $\delta = 1$ ns while the WCDMA
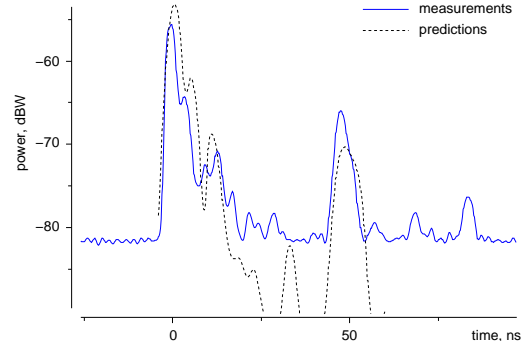

Figure 2.3. Measurement vs. prediction of channel impulse response.

simulation used chip time $\delta \approx 260$ ns. The measured electric field E_k^m of bin k centered at time $k\delta$ is

$$E_k^m = C \sum_{j=1}^{Q} E'_j e^{i\phi_j} \int_{t_j - k\delta - \delta/2}^{t_j - k\delta + \delta/2} e^{-\tau^2/(2\sigma^2)} \mathrm{d}\tau, \quad (2.3)$$

where Q is the number of rays, σ is the half-width of the Gaussian pulse (1.25 ns for measurements), and C is a scale factor that fits this generic equation to a particular system. Since most of the energy in the Gaussian pulse should fall into one time interval of width δ, assume that

$$C \int_{-\delta/2}^{\delta/2} e^{-\tau^2/(2\sigma^2)} \, d\tau = 1. \qquad (2.4)$$

The complex factor $e^{i\phi_j}$ accounts for ray interference. Phase angles ϕ_j were determined from transmitter wavelength λ, total ray path length d_j, and number of reflections n (a 180 degree phase shift per reflection was assumed). Another interpretation of (2.3) is that every time bin registers a weighted average of the energies of all predicted rays, where the weight decreases exponentially as the time difference of the ray and the bin increases. Finally, $P_k^m = |E_k^m|^2/\eta$ gives the measured power of bin k, in watts.

Figure 2.3 shows measurements and predictions for one location with relatively strong multipath. As can be seen from the graph, the predictions are within 3–5 dB of the measurements, which is similar to the results achieved by earlier research [16]. The difference can be explained by device positioning errors (devices were positioned with ± 3 cm precision, which is crude given that the wavelength was 12 cm) and imprecise modeling of reflections. Additionally, small multipath components were missed by the ray tracer. These components are probably due to scattering and diffraction, which were not simulated. Geodesic tessellation frequency was 700 (9.8×10^6 beams) for calibration because the simulation results for frequencies above 700 were indistinguishable.
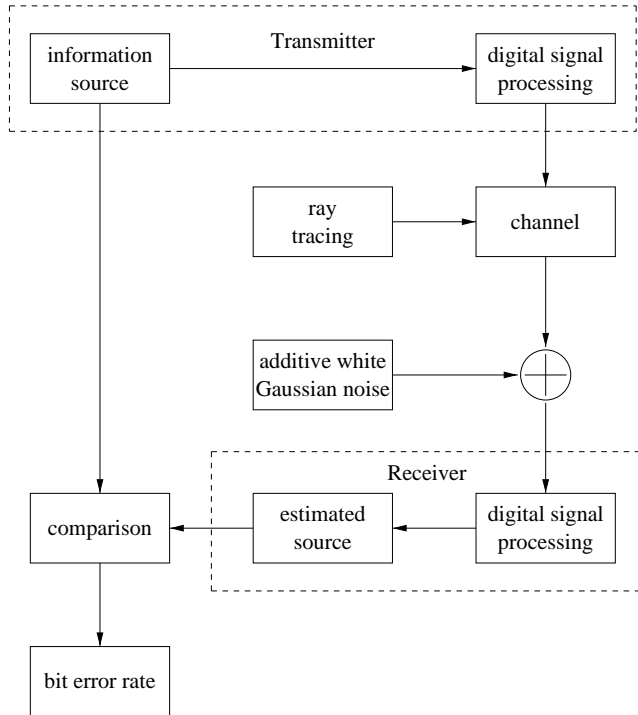
Figure 3.1. Block diagram of the WCDMA simulator.

3. WCDMA Simulation

The ray tracing propagation model predicts a measured impulse response $P_1^m, P_2^m, \ldots, P_n^m$ of a wireless channel (see Section 2). This propagation model does not directly predict the performance of any particular wireless system that operates in this channel. A meaningful performance metric is the bit error rate (BER) defined as the ratio of the number of incorrectly received bits to the total number of bits sent. The BER of a narrowband system (designed for $n = 1$) correlates with P_1^m, so the power level at the receiver maps directly to the BER of a narrowband system. However, estimating the BER of a wideband system (designed for $n > 1$) in a mobile wireless environment usually involves analytically non-tractable problems. This work uses simple least square fit models to the results of a Monte Carlo simulation of a WCDMA system. The WCDMA simulation models channel variation due to changes in the environment as a random process [8]. Notice that channel variation due to receiver movement is modeled in both the ray tracing and the WCDMA simulations, but other kinds of variation are modeled only in the WCDMA simulation. This section outlines the WCDMA simulation and describes the surrogate functions used for optimization.

Figure 3.1 briefly describes the computational steps of the WCDMA simulator. The source module of the transmitter generates information bits to be sent through a wireless channel. The generated information is processed with a series of digital signal processing techniques to reduce the potential channel errors. The wireless channel is modeled as a linear time varying filter in the present work. The channel is characterized by the impulse response predicted by the ray tracer. Before being sent to the receiver, the channel output is combined with Gaussian noise from the electronic system. Similarly, the received distorted signal is processed with a series of digital signal processing techniques by the receiver, which thereafter estimates the information bits to be compared with the original information bits for the BER.

The WCDMA simulation is computationally intensive since a satisfactory BER value ranges from 10^{-3} to 10^{-6}. The parallelized WCDMA simulator significantly speeds up the simulation process, but its run time is still far from practical for optimization problems. The BER depends on small-scale propagation effects that exhibit large variation with respect to receiver location. Practical coverage optimization problems involve wavelengths of less than a foot and areas of thousands of square feet. Four samples per wavelength should be taken to obtain meaningful aggregate results. Therefore, the BER results of the WCDMA simulation were approximated by simple models.

Consider a distribution of impulse responses in the environment shown in Figure 1.1, as measured by the WCDMA system with the standard chip time $\delta \approx 260$ ns and a dynamic range of 12 dB. Empirically, 49% of the impulse responses have only one component ($n = 1$), 42% have two components where the first one is dominant ($n = 2$, $P_1^m \geq P_2^m$), 7% have two components where the second one is dominant ($n = 2$, $P_1^m < P_2^m$), and the remaining 2% have three components ($n = 3$). It turns out that the BERs at the majority of the receiver locations can be approximated by simple functions. This work considers the first two cases that account for 91% of the data.

Given a measured impulse response $P_1^m, P_2^m, \ldots, P_n^m$, define the relative strength of the first component

$$p_1 = P_1^m / \sum_{1 \leq i \leq n} P_i^m$$

and the signal-to-noise ratio (SNR)

$$S = \max_{1 \leq i \leq n} \left\{ 10 \log_{10}(P_i^m / N_0) \right\}$$

(in dB), where N_0 is the power at the noise level (in watts).

The BER b_1 of a WCDMA system in the first case ($n = 1$, $p_1 = 1$) was approximated by

$$\log_e(b_1) = -0.251 S - 2.258, \qquad (3.1)$$

obtained by a linear least squares fit of the simulated BERs for $S = 0, 2, \ldots, 30$ in steps of 2 dB (16 points). In other words, the BER of a WCDMA system with no

multipath is a simple monotonically decreasing function of the SNR. This observation justifies the use of power levels to predict system performance when there is no multipath. However, using the strongest component to predict the BER does not work when $n > 1$.

The second case ($n = 2$, $p_1 \geq 0.5$) was approximated by the slightly more complicated model

$$\log_e(b_2) = -0.467\,S + 0.930\,e^{p_1}$$
$$+\ 0.037\,Se^{p_1} - 5.272, \quad (3.2)$$

obtained by a linear least squares fit, where b_2 is the BER. The data consisted of 65 points for a cross product of $S = 0$, 2, $\ldots$, 20 and $p_1 = 0.5$, 0.6, $\ldots$, 0.9, 0.925, excluding the points with $b_2 < 3 \cdot 10^{-5}$ that required an enormous computation time for accurate results. Equation (3.2) implies that the logarithm of BER is a bilinear function of SNR and e^{p_1}, with a relatively weak cross-term. A further examination of the fitted constants in (3.2) reveals that the BER approaches zero as the SNR increases and that stronger multipath greatly improves performance for a fixed SNR. The latter needs some explanation because multipath is often thought of as an obstacle that impairs system performance. In this work, the SNR is defined in terms of the strongest component of the impulse response. When the SNRs of two channels that meet the criteria for this case are the same, the channel with a stronger second component transfers more total power than the channel with a weaker second component. In this case, the benefits of more power outweigh the disadvantages of multipath.

Both surrogate models were validated with the simulated BER results. In the first case, the approximate values were within 9.7% on average (0.9% minimum, 19.4% maximum) of the simulation output at $S = 1$, 3, $\ldots$, 29. The validation set for the second model consisted of 87 points with a cross product $S = 0$, 1, $\ldots$, 30 and $p_1 = 0.55$, 0.65, 0.75, 0.85 pruned according to the same criterion as the least squares data. The average, minimum, maximum error of the least squares fit was 13.1%, 0.5%, 54%, respectively.

Finally, observe that the models for the two cases are not asymptotically matched. The simulated WCDMA receiver had two antennas, one of which was turned on or off depending on whether or not the second component met the relative power threshold. Discontinuity can pose problems for the DIRECT optimization algorithm, which assumes Lipschitz continuity.

To summarize, this work considers two surrogate models for the BER of a WCDMA system. Both models were obtained by a linear least squares fit of the logarithm of the BER to a combination of channel characteristics. Empirically, both models cover 91% of the data with about 13% average relative error. However, no confident claims can be made because the distribution of the fitted

data is unknown. In particular, these models do not apply for $n > 2$. On the other hand, the models predict sensible responses outside of the range of the fitted data. The latter is crucial for their application to solve the optimization problem described in this paper. The formulation of the optimization problem does not allow directly limiting the model variables. Thus, a more accurate model that produces unreasonable values outside of the fitted data is less desirable than a less accurate model that produces reasonable values outside of the fitted data.

4. DIRECT

The multivariate DIRECT algorithm can be described by the following six steps [9].

Given an objective function f and the design space $D = \{x \in E^n \mid \ell \leq x \leq u\}$:

Step 1. Normalize the design space D to be the unit hypercube. Sample the center point c_i of this hypercube and evaluate $f(c_i)$. Initialize $f_{\min} = f(c_i)$, evaluation counter $m = 1$, and iteration counter $t = 0$.

Step 2. Identify the set S of potentially optimal boxes.

Step 3. Select any box $j \in S$.

Step 4. Divide the box j as follows:

   (1) Identify the set I of dimensions with the maximum side length. Let δ equal one-third of this maximum side length.

   (2) Sample the function at the points $c \pm \delta e_i$ for all $i \in I$, where c is the center of the box and e_i is the ith unit vector.

   (3) Divide the box j containing c into thirds along the dimensions in I, starting with the dimension with the lowest value of $w_i = \min\{f(c + \delta e_i), f(c - \delta e_i)\}$, and continuing to the dimension with the highest w_i. Update $f_{\min}$ and m.

Step 5. Set $S = S - \{j\}$. If $S \neq \emptyset$ go to Step 3.

Step 6. Set $t = t + 1$. If iteration limit or evaluation limit has been reached, stop. Otherwise, go to Step 2.

Steps 2 to 6 form a processing loop controlled by two stopping criteria—limits on iterations and function evaluations. Starting from the center of the initial hypercube, DIRECT makes exploratory moves across the design space by probing potentially optimal subsets. 'Potentially optimal' is an important concept defined next [9].

Definition 4.1. Suppose that the unit hypercube has been partitioned into m (hyper) boxes. Let c_i denote the center point of the ith box, and let d_i denote the distance from the center point to the vertices. Let $\epsilon > 0$ be a positive
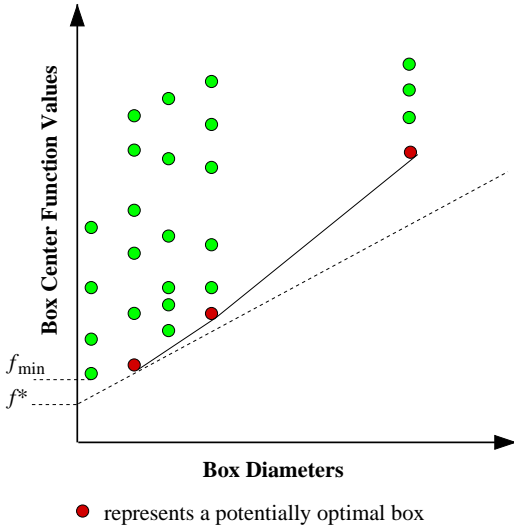
Figure 4.1. Illustration of potentially optimal boxes on convex hull with ϵ test. Note that $f^* = f_{\min} - \epsilon|f_{\min}|$. Potentially optimal boxes are on the lower-right convex hull.

constant. A box j is said to be *potentially optimal* if there exists some $\tilde{K} > 0$ such that for all $i = 1, \ldots, m$,

$$f(c_j) - \tilde{K}d_j \le f(c_i) - \tilde{K}d_i, \qquad (4.1)$$

$$f(c_j) - \tilde{K}d_j \le f_{\min} - \epsilon|f_{\min}|. \qquad (4.2)$$

Figure 4.1 represents the set of boxes as points in a plane. The first inequality (4.1) screens out the boxes that are not on the lower right of the convex hull of the plotted points, as shown in Figure 4.1. Note that $\tilde{K}$ plays the role of the (unknown) Lipschitz constant. The second inequality (4.2) prevents the search from becoming too local and ensures that a nontrivial improvement will (potentially) be found based on the current best solution. In Figure 4.1, $f_{\min}$ is the current best solution, but its associated box is screened out of the potentially optimal box set due to the second inequality (4.2). An example illustrating the behavior of DIRECT on a simple 2D function is given in [18].

Some modifications with respect to the stopping rules and box selection rules are proposed in the present implementation to offer more choices. Two new stopping criteria are (1) minimum diameter (terminate when the best potentially optimal box's diameter is less than this minimum diameter) and (2) objective function convergence tolerance (exit when the objective function does not decrease sufficiently between iterations). The minimum diameter of a hyperbox represents the degree of space partition, and therefore is a reasonable criterion for optimization problems requiring only some depth of design space exploration. The objective function convergence tolerance was inspired by some
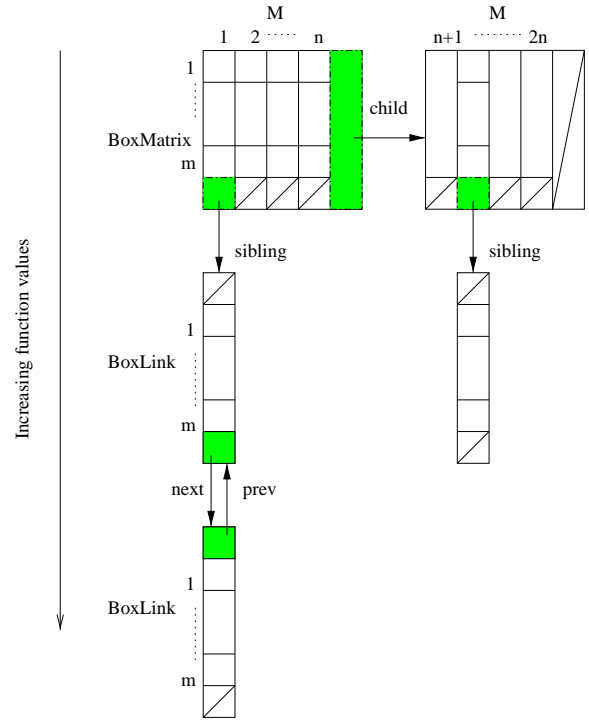

Figure 4.3. Box structures comprised of `HyperBox`es.

experimental observations in the later stages of running the DIRECT algorithm, when the objective function convergence tolerance test avoids wasting a great number of expensive function evaluations in pursuit of very small improvements. This is a reasonable stopping criterion for large-scale engineering design problems.

The present implementation of the DIRECT algorithm addresses an efficiency issue involved in an unpredictable storage requirement in the phase of space partitioning. To reduce the execution overhead and adapt to varying memory requirements, a set of dynamic data structures is proposed. They are extensible and flexible in dealing with information generated by the space partitioning process in high dimensions.

Two groups of dynamic structures have been implemented in Fortran 90: box structures and linked list structures illustrated by Figure 4.3. The box structures (`BoxMatrix`, `BoxLink`, and `HyperBox`) are responsible for holding boxes. The linked lists are built out of linked real and integer vectors, and manage the allocated memory for the box structures.

In [9], Graham's scan method is recommended for finding the convex hull of a set of m arbitrary points in time $O(m \log_2 m)$. Here, a different approach is taken to shrink the initial set with m points to a much smaller set of vertices exclusively around the low edge of the convex hull. With all the hyperboxes linked logically in the scatter plot pattern, Jarvis's march (or

gift wrapping) method is applied starting from the box sequence with the biggest size, and eventually identifies all the potentially optimal boxes to be further subdivided for the next iteration.

The linked list data structures play an important role in maintaining the logical scatter plot pattern and recycling memory cells. They are doubly linked lists constructed with two derived data types. In some sense, `M` (the two-dimensional array defined in `BoxMatrix`) acts as a memory pool of recyclable cells for holding boxes. When cells are used up, a new `BoxMatrix` is allocated and connected as the `child` link at the end of the chain of `BoxMatrices`, so that the memory pool can be filled up again using new cells from `M` in the newly allocated `BoxMatrix`. For faster execution, sorting is not involved in the strategy for maintaining a logical scatter plot pattern of hyperboxes. Instead, binary search is used in locating the insertion positions in sorted sequences, in both the cases of inserting boxes and box sizes. Some shifting operations are needed for inserting/deleting boxes in a particular column of boxes in `M` and its box links, if any, while shifting boxes among columns is avoided by keeping column indices sorted (by decreasing box sizes).

5. Optimization Results

Ray tracing was performed on a 200-node Athlon 650 Beowulf cluster of Linux workstations. Two sets of simulations for optimizing transmitter placement were executed with respect to the two performance criteria—coverage and BER—discussed in Section 1. The ray tracer's tessellation frequency was 100 for coverage and 700 for BER. The former was sufficient to match the peak powers against measurements, while the latter was required to match the whole impulse responses. The optimizer and the user interface ran on a Sun workstation outside the cluster. Tcl/Tk scripts glued the pieces together and provided a graphical user interface. Similar to [11], users could select regions for transmitter placement (to be optimized) and regions to be covered.

Consider the placement of n transmitters in an indoor environment located on the fourth floor of Durham Hall at Virginia Tech (see Figure 1.1). Suppose, the objective is to maximize the average performance over m receiver locations. The variables are the transmitter coordinates

$$X = (x_1, y_1, z_1, x_2, y_2, z_2, \ldots, x_n, y_n, z_n),$$

where all $z_j = z_0$ are fixed, a reasonable assumption for indoor environments. Let transmitter (k, i), located at $(x_k, y_k, z_0), 1 \leq k \leq n$, generate the highest power level $P_{ki}(x_k, y_k, z_0) \geq P_{ji}(x_j, y_j, z_0), 1 \leq j \leq n$, at the receiver location $i, 1 \leq i \leq m$. The objective function is
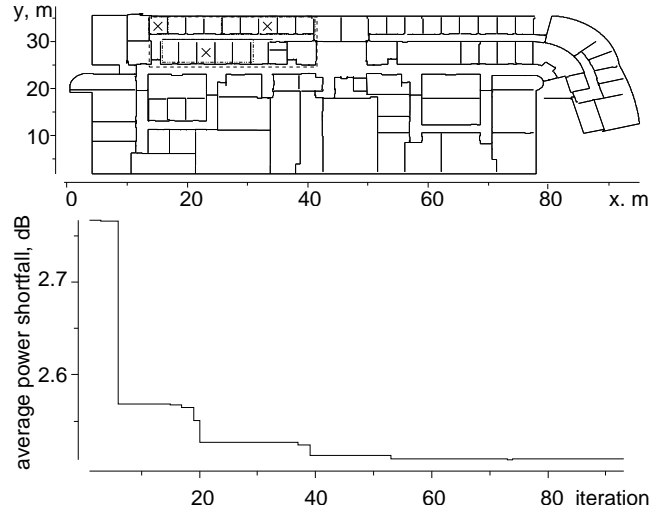

Figure 5.1. Power coverage optimization results for three transmitters. Bounds on transmitter placement are drawn with dotted lines and their initial positions are marked with crosses. The dashed line delimits the region to be covered.

the average shortfall of the estimated performance metric from the given threshold T, given by

$$f(X) = \begin{cases} \dfrac{1}{m} \sum_{i=1}^{m} (T - p_{ki})_+, & \text{coverage}, \\[2ex] \dfrac{1}{m} \sum_{i=1}^{m} (p_{ki} - T)_+, & \text{BER}, \end{cases} \quad (5.1)$$

where p_{ki} is the performance metric of transmitter (k, i) evaluated at the ith receiver location. For power coverage optimization, p_{ki} is $P_{ki}(x_k, y_k, z_0)$ and $(T - p_{ki})_+$ is the penalty for a low power level. For BER optimization, p_{ki} is $\log_{10}(\text{BER}_{ki})$ and $(p_{ki} - T)_+$ is the penalty for a high bit error rate.

Figure 5.1 illustrates power coverage optimization of the locations of three transmitters to cover eighteen rooms and a corridor bounded by the box in the upper-left corner. 93 function evaluations reduced the objective from 2.77 dB to 2.51 dB, or by 9.4%, in 38 minutes on 40 machines. Figure 5.2 depicts BER optimization of the locations of two transmitters to cover half of the former region. 51 iterations reduced the objective function from 0.091 to 0.021 in 7 hours and 26 minutes on 44 machines. The BER threshold was $\log_{10}(\text{BER}) = -3$ (voice quality), so this improvement corresponds to a 17% reduction in the average BER. In both cases, the optimization loop stops with the minimum diameter required by the problem. System performance was significantly improved by DIRECT with a reasonable number of evaluations.
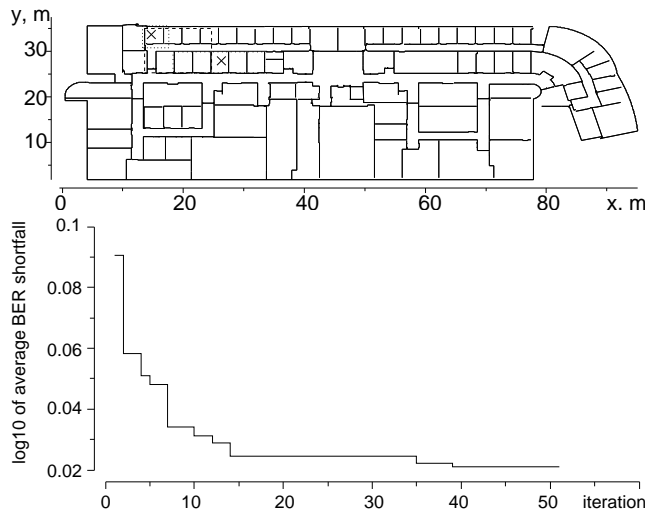
Figure 5.2. BER optimization results for two transmitters. The region to be covered is half of that in Figure 5.1.

6. Conclusion

The main contribution of the present work is the design of an optimization loop that takes feedback from a sophisticated wireless system model. The model involves parallel 3D ray tracing and two surrogate functions that estimate the BER of a WCDMA system. DIRECT has demonstrated its effectiveness in solving the global optimization problem of transmitter placement in wireless communication systems design.

Several extensions to the present work are envisioned. First, the surrogate functions for the BER can be extended to channels with relatively strong multipath and interference. Second, wireless systems with data quality bit error rates (10^{-6}) can be considered. Third, different ways of combining the deterministic propagation model with the stochastic wireless system model can be explored.

7. Acknowledgments

References

[1] G. D. Durgin, T. S. Rappaport, and Hao Xu, "Measurements and models for radio path loss and penetration loss in and around homes and trees at 5.85 GHz," *IEEE Transactions on Communications*, vol. 46(11), pp. 1484–1496, 1998.

[2] S. J. Fortune, "A beam-tracing algorithm for prediction of indoor radio propagation," in *WACG: 1st Workshop on Applied Computational Geometry: Towards Geometric Engineering, Lecture Notes in Computer Science*, vol. 1148, pp. 157–166, 1996.

[3] S. J. Fortune, D. M. Gay, B. W. Kernighan, O. Landron, R. A. Valenzuela, and M. H. Wright (AT&T Bell Laboratories), "WISE design of indoor wireless systems: practical computation and optimization," *IEEE Computational Science & Engineering*, vol. 2(1), pp. 58–68, Spring, 1995.

[4] B. Freisleben and D. Hartmann and T. Kielmann, "Parallel raytracing: a case study on partitioning and scheduling on workstation clusters," in *Proc. Thirtieth International Conference on System Sciences, Hawaii*, vol. 1, pp. 596–605, 1997.

[5] A. Glassner, "Space subdivision for fast ray tracing," *IEEE Computer Graphics and Applications*, vol. 4(10), pp. 15–22, October, 1984.

[6] X. Huang, U. Behr, and W. Wiesbeck, "Automatic base station placement and dimensioning for mobile network planning," in *Proc. of Vehicular Technology Conference*, IEEE VTS Fall VTC 2000. 52nd, vol. 4, pp. 1544–1549, 2000.

[7] J. He, A. Verstak, L. T. Watson, T. S. Rappaport, C. R. Anderson, N. Ramakrishnan, C. A. Shaffer, W. H. Tranter, K. Bae, and J. Jiang, "Global optimization of transmitter placement in wireless communication systems," to appear in *Proc. High Performance Computing Symposium 2002*, A. Tentner (ed.), Soc. for Modeling and Simulation International, San Diego, CA, 2002.

[8] M. C. Jeruchim, Philip Balaban, and K. S. Shanmugan, *Simulation of Communication Systems*, Plenum Press, New York, 1992.

[9] D. R. Jones, C. D. Perttunen, and B. E. Stuckman, "Lipschitzian optimization without the Lipschitz constant," *Journal of Optimization Theory and Applications*, vol. 79(1), pp. 157–181, 1993.

[10] R. M. Lewis, V. Torczon, and M. W. Trosset, "Direct search methods: then and now," *Journal of Computational and Applied Mathematics*, vol. 124, pp. 191–207, 2000.

[11] M. A. Panjwani, A. L. Abbott, and T. S. Rappaport, "Interactive computation of coverage regions for wireless communication in multifloored indoor environments," *IEEE Journal on Selected Areas in Communications*, vol. 14(3), pp. 420–430, 1996.

[12] T. S. Rappaport, *Wireless Communications: Principles and Practice*, Prentice Hall, New Jersey, 1996.

[13] T. S. Rappaport and R. R. Skidmore, Wireless Valley Communications, Inc., *Method and System for Automated Optimization of Antenna Positioning in 3-D*, US Patent 6,317,599, November, 2001.

[14] H. D. Sherali, C. M. Pendynala, and T. S. Rappaport, "Optimal location of transmitters for micro-cellular radio communication system design," *IEEE Journal on Selected Areas in Communications*, vol. 14(4), pp. 662–673, 1996.

[15] K. R. Schaubach, N. J. Davis IV, and T. Rappaport, "A ray tracing method for predicting path loss and delay spread in microcellular environments," in *Proc. IEEE Vehicular Technology Conference*, vol. 2, pp. 932–935, 1992.

[16] S. Y. Seidel and T. S. Rappaport, "Site-specific propagation prediction for wireless in-building personal communication system design," *IEEE Transactions on Vehicular Technology*, vol. 43(4), pp. 879–891, 1994.

[17] A. Verstak, M. Vass, N. Ramakrishnan, C. Shaffer, L. T. Watson, K. K. Bae, J. Jiang, W. H. Tranter, and T. S. Rappaport, "Lightweight data management for compositional modeling in problem solving environments," in *Proc. High Performance Computing Symposium 2001*, A. Tentner (ed.), Soc. for Modeling and Simulation Internat., San Diego, CA, pp. 148–153, 2001.

[18] L. T. Watson and C. A. Baker, "A fully-distributed parallel global search algorithm," *Engineering Computations*, vol. 18(1/2), pp. 155–169, 2001.

IEEE
COMPUTER
SOCIETY