Modeling UpLink Power Control with Outage Probabilities

Kenneth L. Clarkson[†], K. Georg Hampel[*], and John D. Hobby[*]

* Bell Laboratories, Alcatel-Lucent

[†] IBM Research (work was done at Bell Laboratories)

Abstract— **We investigate models for uplink interference in wireless systems. Our models account for the effects of outage probabilities. Such an accounting requires a nonlinear, even nonconvex model, since increasing interference at the receiving base station increases both mobile transmit power *and* outage probability, and this results in a complex interaction. Our system model always has at least one solution, a fixed point, and it is provably unique under certain reasonable conditions. Our main purpose is to model real wireless systems as accurately as possible, and so we test our models on realistic scenarios using data from a sophisticated simulator. Our algorithm for finding a fixed point works very well on such scenarios, and is guaranteed to find the fixed point when we can prove it is unique. A slightly simplified model reduces the main data structure for a K-sector market to $16K^2$ bytes of memory.**

I. Introduction

While many factors can cause problems for a cell phone call, and many design goals must be balanced in designing a cell phone system, the power used by the phone (the "mobile") is particularly important: this power is limited, and the less it is used, the longer the mobile battery will last and the smaller it can be. Moreover, in a spread spectrum system, the signal from each mobile can interfere with the signal from every other mobile. This motivates the use of sophisticated power control methods: by a variety of means, the system determines how much power is needed by a mobile to carry its call, and the mobile transmits using just that much power. Part of this determination is done at the base stations interacting with the mobile. The result is a complex dynamical system, as mobiles move, signal losses vary, and calls begin and end.

We describe here a computational model of spread-spectrum uplink power control. The model is used within Alcatel-Lucent's Ocelot software for wireless optimization, which handles for example CDMA2000 and UMTS voice and circuit data services. With this application in mind, the model has several properties:

- It models existing wireless systems: it is not a proposal for a new power control system;
- It can be evaluated with reasonable speed;
- It is a differentiable function of relevant parameters, and the derivatives can be evaluated with reasonable speed.

The first property implies that we cannot, for example, simply insist of the wireless system that no calls are dropped; such a requirement corresponds to the inclusion of an upper bound on mobile power as a constraint in an optimization problem [1].

We focus here mainly on voice, not data, but some parts of our modeling apply to certain kinds of circuit data services.

Our basic setting is as follows (see for example Lee and Miller [2]). In a given region, there are K base-station antennas (hereafter "sectors"), and a given sector k, with $1 \le k \le K$, receives total radio power x_k, from mobiles in the region and from thermal noise and external radio interference sources. Based on the frame error rate, the sector determines a target SIR (signal-to-interference ratio) ϕ_k such that if the signal power received from a mobile is at least $\phi_k x_k$, then the error rate for the mobile will be acceptably low. For each mobile m in soft handoff with the sector, the sector determines the received power s_m, and sends a *power control bit* to the mobile, whose value depends on whether $s_m > \phi_k x_k$, and tells the mobile to increment or decrement its transmit power accordingly. The mobile looks at all such power control bits, and decrements its power if any of the bits suggest it. This protocol keeps the mobile power near the smallest possible such that some sector will receive it with adequate SIR.
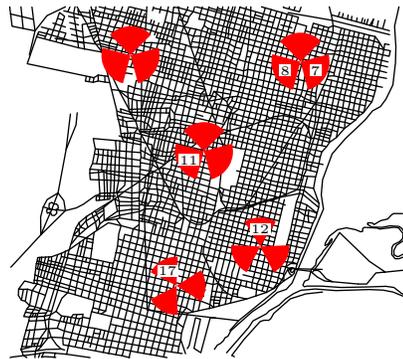

Fig. 1. A typical test scenario with some of the sectors labeled by index k.

Suppose each mobile m is transmitting with power s_m, and is received at sector k with power s_m/L_{km}, where L_{km} is the pathloss factor for the signal traveling from m to k. Let $k(m)$ denote the sector that is currently determining the power transmitted by m, and S_k denotes the set of mobiles m such that $k = k(m)$. Then the power transmitted by $m \in S_k$ is

$$s_m := \phi_k x_k L_{km}, \qquad (1)$$

and so for each k', $x_{k'} = \eta_{k'} + \sum_k \sum_{m \in S_k} L_{km} \phi_k x_k / L_{k'm}$, where $\eta_{k'}$ is the power of the noise plus external interference received by sector k'. This equation is correct only under some

approximations and assumptions, but it suggests something of the nature of the model that must be evaluated. Put another way, the vector of total radio powers x satisfies the fixed point condition $x = \eta + Ax$, where

$$A_{k'k} := \sum_{m \in S_k} \frac{L_{km}\phi_k}{L_{k'm}}. \tag{2}$$

Note that A is a nonnegative matrix, that is, all its entries are nonnegative. If $\eta = 0$, then $x = Ax$, and x is an eigenvector of the nonnegative matrix A.

The eigenvalues and eigenvectors of nonnegative matrices are well-studied, as the *Perron-Frobenius* theory, and that theory has been applied to the understanding of power control [3]. However, the presence of noise and external interference, implying $\eta > 0$, means that such theory does not directly give the most detailed understanding of power control.

A further complication is reverse-link outage: a call may be dropped if the mobile cannot transmit the target power $\phi_k x_k L_{km}$ as in (1); thus the power s_m is not a linear function of x_k, but instead a sawtooth: at a certain x_k threshold, it goes to zero.

Another complication is *noise rise limitation*: a sector may block calls if the total radio power it receives is above a pre-set threshold. Such a limitation is discussed in Section V.

A dynamic model of power control might maintain a collection of active mobiles, adding some as calls arrive and dropping others either as normal call termination, or as outages, the result of reverse-link failure. Simulation over time would then yield outage probabilities, average values for the x_k, and so on. However, such a scheme would be too slow for our optimization application, and also, not smooth enough. We use instead a static framework: a (large) discrete set of locations is fixed, each of which has an estimated probability of being the location of a transmitting mobile. The locations and probabilities are determined elsewhere, and are based on input by the Ocelot user, from a variety of sources, and also on estimates of forward-link coverage probability, and other considerations. The locations could just be points on a regular grid, but Ocelot provides many other options.

We thereby model a dynamic set of mobiles simply as the expectation of the mobile power generated at each location. With some abuse of notation, we index the locations with m, and have corresponding losses L_{km}, power levels s_m, and so on. Since we are modeling probabilities and expectations and not specific mobiles, the power s_m need not be a discontinuous sawtooth function of x_k, but instead can drop off smoothly, as an ensemble average. We might base this dropoff on a log-normal probability distribution for shadow fading. Such a model is discussed in Section II. However, for efficiency reasons, we use spline-based approximations to the mobile response, as discussed in Section III. This leads to a function $F : \mathcal{R}^K \to \mathcal{R}^K$, which disregarding call dropping would be $F(x) = \eta + Ax$, but instead we have $F(x)_{k'} = \eta_{k'} + \sum_k \tilde{A}_{k'k}(x_k)$, where $\tilde{A}_{k'k}(x_k)$ is a spline function of the sector k interference x_k. This smooth replacement for the sawtooth is not only more plausible as an estimate, but is convenient computationally: with discontinuities, there may not be a solution to the fixed point problem $x = F(x)$; with the smooth version, we are able to show that under some reasonable conditions, a fixed point solution exists. The solution of the fixed point problem is discussed in Section V.

Here is the outline for the rest of the paper: we begin in Section II by explaining the smoothed sawtooth functions and their relationship to log-normal fading. Then Section III gives a spline-based approximation that allows contributions for various locations m to be combined and manipulated efficiently, and Section IV presents a resource-saving refinement. Next, Section V presents robust algorithms for finding fixed points and gives appropriate theorems. The results in Section VI include discussions of the various models, and tests on realistic scenarios (not just hexagonal grids). Finally, Section VII presents conclusions.

II. SMOOTHED SAWTOOTH FUNCTIONS

Consider a single term from the sum (2) as changed to account for the dropping of calls due to reverse-link limitations. If $\hat{s}$ is the maximum mobile uplink transmit power, multiplying the term by x_k, if $s_m < \hat{s}$, and 0 otherwise, where again $s_m := x_k L_{km} \phi_k$ for $m \in S_k$, gives the sawtooth function that we need to smooth by considering the ensemble average of a dynamic set of mobiles and finding the contribution to interference due to location m. The sawtooth-based interference contribution to sector k' from location m transmitting to location k could thus be expressed as

$$\frac{\hat{s}}{L_{k'm}}\gamma_m Q^-(\gamma_m), \quad \text{where} \quad Q^-(t) := \begin{cases} 1 & \text{if } t < 1 \\ 0 & \text{otherwise,} \end{cases}$$

and the ratio $\gamma_m := s_m/\hat{s}$. Another way to describe this is as

$$\frac{\hat{s}}{L_{k'm}}G(\gamma_m),$$

where $G(\gamma)$ is the sawtooth function

$$G(\gamma) := \gamma Q^-(\gamma). \tag{3}$$

Replace pathlosses L_{km} and $L_{k'm}$ by $L_{km}\exp(R_{km})$ and $L_{k'm}\exp(R_{k'm})$, where random variables R_{km} and $R_{k'm}$ are $N(0,\sigma)$, that is, normally distributed with zero mean and standard deviation σ. This implies replacing γ_m by $\gamma_m \exp(R_{km})$ as well. Then the expected interference contribution for location m could be estimated as

$$\mathbf{E}\left[\frac{\hat{s}}{L_{k'm}\exp(R_{k'm})}\gamma_m \exp(R_{km})Q^-(\gamma_m \exp(R_{km}))\right]$$
$$= \frac{\hat{s}\gamma_m}{L_{k'm}}\mathbf{E}\left[\exp(R_{km} - R_{k'm})Q^-(\gamma_m \exp(R_{km}))\right], \quad (4)$$

although the non-outage condition $Q^-(\gamma_m \exp(R_{km})) = 1$ might be better modeled as a non-outage probability $Q()$.

It may be that R_{km} and $R_{k'm}$ are partially correlated, so we assume that there is some $\beta \in [0,1]$ and $N(0,\sigma)$-distributed random variable $\hat{R}_{k'm}$, independent of R_{km}, so that $R_{k'm} = (1-\beta)R_{km} + \beta\hat{R}_{k'm}$, and so

$$\exp(R_{km} - R_{k'm}) = \exp(-\beta\hat{R}_{k'm})\exp(\beta R_{km}).$$

If we use this in (4), use the independence of R_{km} and $\hat{R}_{k'm}$, and observe that $\mathbf{E}[\exp(-\beta\hat{R}_{k'm})] = \mathbf{E}[\exp(\beta\hat{R}_{k'm})]$, the expected interference contribution becomes

$$\frac{\hat{s}\gamma_m}{L_{k'm}}\mathbf{E}\left[\exp(R_{km}-R_{k'm})Q^-(\gamma_m\exp(R_{km}))\right]$$
$$= \frac{\hat{s}\gamma_m}{L_{k'm}}\mathbf{E}[-\beta R_{k'm}]\mathbf{E}\left[\exp(\beta R_{km})Q^-(\gamma_m\exp(R_{km}))\right]. \quad (5)$$

It is not hard to show that this is

$$\frac{\hat{s}\exp(\beta^2\sigma^2)}{L_{k'm}}\gamma_m\Phi[-\ln(\gamma_m)/\sigma-\beta\sigma]. \quad (6)$$

Here $\Phi(x)$ is the normal cumulative distribution function at x, the probability that a $N(0,1)$ random variable is less than x. Note that $\gamma_m\Phi[-\ln(\gamma_m)/\sigma-\beta\sigma]$ is essentially a smoothed version of (3). In the following discussion, we assume $\beta = 0$ for brevity, but it is not hard to generalize to $\beta > 0$, which our algorithms handle.

III. Spline Approximations

The normal CDF needed for (6) is easily evaluated via the error function erf(), since $2\Phi(x) = 1 + \text{erf}(x/\sqrt{2})$. However, it would be awkward to handle sums of such interference contributions for many different locations $m \in S_k$. Viewing (6) as a function of x_k rather than γ_m requires a substitution $\gamma_m = s_m/\hat{s} = \phi_k x_k L_{km}/\hat{s}$ that shifts the x_k values at which (6) begins to fall to zero. Even if we were to scale them so that their initial slopes match, the functions (6) for different locations m would be as in Figure 2a.
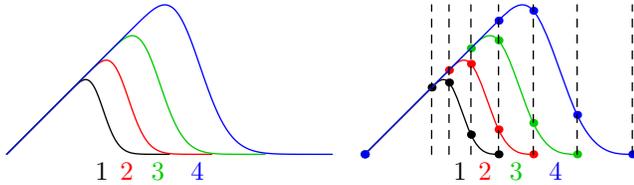

Fig. 2. (a) Smoothed sawtooth functions for various locations m where γ_m/x_k differs by factors of $\alpha = 1.26$; (b) the corresponding spline approximations with dashed lines at the ends of spline segments.

Spline approximations such as those shown in Figure 2b are much more convenient when taking smoothed sawtooth functions for various locations m, and adding them up in a manner analogous to (2). The splines are piecewise-cubic polynomial functions chosen to have second-order continuity at the *knots* where one cubic polynomial segment joins the next one. Placing the knots at powers of a parameter α ensures that any linear combination of these spline functions will be a piecewise cubic spline with the same knot spacing.

We have found a spline with four knots, leading to a linear system of ten equations in ten unknowns, to be an effective approximation. Call such a spline approximation $G_\alpha(x)$; it is normalized to have an initial slope of one.

The complete set of normalized smoothed sawtooth splines is

$$\left\{\frac{G_\alpha(b\alpha^i x)}{b\alpha^i}\ \middle|\ i \in \mathbf{Z}\right\},$$

where b is a bias parameter to be chosen along with α. We must choose α and b so that

$$\frac{G_\alpha(b\gamma)}{b} \approx \gamma\Phi[-\ln(\gamma)/\sigma-\beta\sigma] \quad (7)$$

as functions of γ. (Again, we assume for brevity that $\beta = 0$.)

We can quantify the difference between the two functions by evaluating each side of (7) at γ_m values $\sigma^{-3.00}$, $\sigma^{-2.97}$, $\sigma^{-2.94}, \ldots, \sigma^{3.00}$ and taking the RMS mean of the differences. For any given σ, it is easy to choose α and b so as to minimize this. For example, exhaustively trying multiples for 0.0001 for α and b gives the results in Table I. We have

TABLE I
Choosing α and b so as to satisfy (7)

σ	$\ln\alpha$	b	RMS error
0.09212	0.1799	1.14574	0.00379052
0.2303	0.457	1.76343	0.00498369
0.4606	0.9679	4.09981	0.00748236
0.6909	1.5904	12.3288	0.00915151
0.9212	2.3385	50.1016	0.0108936
1.1515	3.1548	239.085	0.0173641
1.3818	3.9875	1140.56	0.030776
1.6121	4.8184	5061.15	0.0540373
1.8424	5.668	21703.5	0.0943915
2.0727	6.6114	141023	0.165292

also defined and used an empirical formula giving α and b as functions of σ and β.

These smoothed sawtooth splines need to be added up for all locations m served by sector k so as to obtain a function $\bar{A}_{k'k}$ for interference at sector k' due to mobiles owned by sector k as a function of x_k, the interference at sector k. In other words, we need a spline-based generalization of (2). For each m, we must choose a smoothed sawtooth spline $\frac{G_\alpha(b\alpha^{i_m}x_k)}{b\alpha^{i_m}}$ so that $\alpha^{i_m}x_k \approx \gamma_m$. This gives

$$\bar{A}_{k'k}(x_k) = \sum_{m\in S_k}\frac{\hat{s}}{L_{k'm}}\frac{G_\alpha(b\alpha^{i_m}x_k)}{b\alpha^{i_m}}.$$

IV. Using Similarity to Save Resources

Since the path loss L_{km} can vary by more than a factor of 1000 as the location m ranges over S_k, the $\bar{A}_{k'k}$ function will typically have dozens of spline segments. This seems like a lot of information to store and manipulate for each pair of sectors $k'k$, especially since $\bar{A}_{k'k}$ functions for a common k but differing k' tend to be related, as shown in Figure 3.

Since the knots are all aligned, it is easy to add up the the spline functions $\bar{A}_{k'k}$ to obtain a master spline function $A_k(x)$ that can be thought of as an average $\bar{A}_{k'k}$, normalized to have unit initial slope.

The resource-saving idea is to store K master spline functions and K^2 simple transformations instead of K^2 spline functions. We use

$$\tau_{k'k}x\left(\frac{\bar{A}_k(x)}{x}\right)^{\bar{\tau}_{k'k}} \quad (8)$$

in place of $\bar{A}_{k'k}(x)$, where $\tau_{k'k}$ and $\bar{\tau}_{k'k}$ are chosen based on $\bar{A}'_{k'k}(0)$ and $\bar{A}_{k'k}(\bar{x})$ for some fixed $\bar{x}$ that can be thought of as
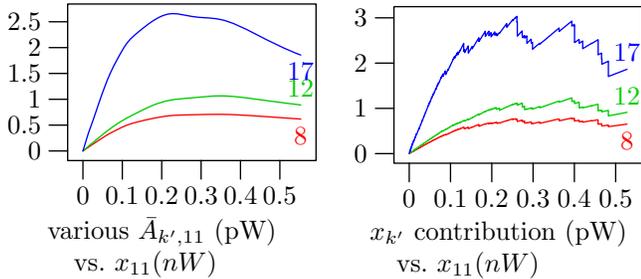
Fig. 3. (a) $\bar{A}_{k'k}$ functions for $k = 11$ and various k'; (b) the corresponding functions based on raw unsmoothed sawteeth. The test data are from the scenario shown in Figure 1.

as an a priori guess at a typical interference level. This way, there are just two values to keep track of for each $k'k$ pair while we consider various locations m. After finding these $2K^2$ values and the K master spline functions $\bar{A}_k$, we can set

$$\tau_{k'k} = \bar{A}'_{k'k}(0) \quad \text{and} \quad \bar{\tau}_{k'k} = \frac{\log(\bar{A}_{k'k}(\bar{x})/(\tau_{k'k}\bar{x}))}{\log(\bar{A}_k(\bar{x})/\bar{x})}$$

for each pair $k'k$. (In practice, one must impose also a positive lower bound on $\bar{\tau}_{k'k}$ to avoid 0^0 in (8).)

V. FINDING A FIXED POINT

The last three sections have described different versions of a function that gives the (expected) interference received at sector k' from mobiles whose primary sector is k. In whatever way each such function $A_{k'k}(x_k)$ is defined, the result is an estimate $\eta_{k'} + \sum_k A_{k'k}(x_k)$ of the interference received at sector k'. This could also be written as a vector $\eta + \bar{F}(x_k)\mathbf{1}$, where $\bar{F}(x_k)$ is a $K \times K$ matrix with $\bar{F}(x_k)_{k'k} = A_{k'k}(x_k)$, and $\mathbf{1}$ is the K-vector of all ones. Defining the function $F : \mathcal{R}^K \to \mathcal{R}^K$ by $F(x) = \eta + \bar{F}(x_k)\mathbf{1}$, the interference vector thus obeys the condition $x = F(x)$. In other words, it is a fixed point of the mapping F.

When the sector-to-sector interference function is linear, a sum of sawtooth, or smoothed via log-normal fading, the corresponding fixed-point problem is easy, unsolvable, or impractically slow. It remains to consider finding fixed points for the last two versions, with spline-based functions, and with the more compact spline-based scheme of Section IV.

Before discussing methods of solution of such general fixed point problems, we consider the implementation of a model of *noise-rise limits* for power control. Here sector k tries to keep $x_k \leq \hat{\rho}\eta_k$, for some $\hat{\rho} > 1$, by blocking calls if necessary. The simplest way to model this is just to compose F with a function that limits the kth component to at most $\hat{\rho}\eta_k$, but it would be more accurate to replace $\mathbf{1}$ in $\eta + \bar{F}(x)\mathbf{1}$ with a vector $h(x)$ whose kth component is a probability that sector k decides not to block a call due to noise rise concerns.

By any of these definitions, F is a smooth function that maps the positive orthant into a rectilinear region

$$U = \{x \in \mathcal{R}^K \mid \eta_k \leq x_k \leq \mu_k\},$$

where the upper bound μ_k is readily derived. Thus we have points μ and η that are opposite corners of U.

Since U is homeomorphic to a closed ball and the continuous function F maps the whole positive orthant (a superset of U) into U, the Brouwer fixed point theorem guarantees that there is at least one fixed point $x = F(x)$.

It would also be desirable to guarantee a unique fixed point and provide an algorithm that finds it efficiently. A popular approach used by Yates [4] and others is to let the algorithm be Picard iteration, where repeatedly $x \leftarrow F(x)$, and also to give conditions under which Picard iteration provably converges to a unique fixed point. For example, if we can exhibit a real number $\kappa < 1$ and a vector norm $\|\cdot\|_*$ under which

$$\|F(x) - F(y)\|_* \leq \kappa \|x - y\|_* \quad \text{for all } x, y \in U, \quad (9)$$

then the fixed point is unique and Picard iteration converges from any starting point in U. Nuzman [5] has shown that this type of argument can be applied to a class of non-monotonic functions that unfortunately does not contain F.

In order to have an algorithm that is as reliable as possible, we certainly need guaranteed convergence if (9) holds. We can do this by producing a sequence of iterates $x^{(1)}, x^{(2)}, x^{(3)}, \ldots$, where each

$$\left\|x^{(i+1)} - F(x^{(i+1)})\right\| \leq \kappa \left\|x^{(i)} - F(x^{(i)})\right\| \quad (10)$$

for the standard Euclidean norm. Picard iteration under (9) achieves this if a few iterations, each of which reduces $\|x - F(x)\|_*$ by the factor κ suffice to reduce $\|x - F(x)\|$ by that factor.

Another way to find a fixed point is to use Newton iteration to look for a zero of $F(x) - x$. Such an iteration does not require (9) and is known to converge quadratically if the initial x is sufficiently close to a solution of $F(x) - x = 0$. This suggests a hybrid algorithm that uses an intelligent starting point, does Newton iterations, but switches to Picard iterations if necessary to obey (10).

1) Use binary search to find a point x on the line between η and μ where the number of negative components in $x - F(x)$ is between $\frac{1}{3}K - \frac{1}{2}$ and $\frac{2}{3}K + \frac{1}{2}$.
2) Compute a Newton step $\Delta x = (J_F(x) - I)^{-1}(F(x) - x)$ and find the maximum $\bar{\lambda}$ such that $x - \bar{\lambda}\Delta x \in U$. Here $J_F(x)$ is the Jacobian of F.
3) Let $y = x$, $e_0 = \|x - F(x)\|$ and exit if e_0 is tiny. Then if $\bar{\lambda} < 0.8$, do Picard iterations and go to Step 2 as soon as $\|x - F(x)\| \leq \kappa e_0$.
4) Let $\hat{\lambda} = \max(\bar{\lambda}, 1 + 10^{-6})$ and $\lambda = \min(1, \bar{\lambda})$. While $\|x - F(x)\| > e_0$ for $x = y - \lambda\Delta x$, iterate $\lambda \leftarrow \max(\lambda/2, 2\lambda - \hat{\lambda})$.
5) If $\|x - F(x)\| > \kappa e_0$, do Picard iterations until $\|x - F(x)\| \leq \kappa e_0$. Then go to Step 2

For each i, the ith iteration of Steps 2–5 advances from $x^{(i)}$ to $x^{(i+1)}$ while trying to ensure that the condition of (10) is satisfied. The testing of λ values less than one, if one Newton step is not feasible or does not help, is a *line search* along the Newton step direction.

If a Newton step is significantly out of bounds or makes so little progress that $e_0 \geq \|x - F(x)\| > \kappa e_0$, the algorithm

resorts to Picard iterations. An excessive number of such iterations in Steps 3 or 5 should be treated as a failure indication.

VI. Results

A version of the algorithm has been tested extensively as part of Alcatel-Lucent's Ocelot software, but it is not practical to rerun it on more than a few test scenarios as shown in Table II. We will not consider defining the function

TABLE II
THE TEST SCENARIOS

Label	K	Explanation
H57	57	Hex pattern with 3 sectors per base station
M39	39	Metro area with population ≈ 500 thousand
M39h	39	M39 with 14 times the traffic near Sector 7
M60	60	Metro area with population ≈ 700 thousand
M115	115	Metro area with population ≈ 2 million

$F : \mathcal{R}^K \to \mathcal{R}^K$ based on unsmoothed sawtooth functions, because fixed point iteration did not converge for any of the scenarios in Table II. (It continually jumps back and forth across a particular discontinuity.) This leaves $F(x) = \eta + Ax$, and also the spline based functions, with and without resource-saving transformations, which we shall refer to as $F\S IV$ and $F\S III$. Except as stated below, these F functions do not include the noise rise limitation.

TABLE III
SAMPLE RESULTS WITH POWERS IN PICOWATTS, ALL $\eta_k = 17.3pW$, AND "$5+1P$" MEANING A PICARD ITERATION WAS NEEDED.

	average x_k			max x_k			Newton steps
	$\eta+Ax$	$F\S III$	$F\S IV$	$\eta+Ax$	$F\S III$	$F\S IV$	
H57	19.8	19.4	19.4	20.2	19.7	19.7	4
M39	35.5	23.7	23.5	99.6	31.0	30.6	4
M39h	—	25.1	25.1	—	141	131	$5+1P$
M60	23.1	21.6	21.6	60.4	35.6	36.1	4
M115	—	21.3	21.2	—	35.2	33.9	5 or 6

Table III summarizes the results for the three versions of F and the four scenarios. As the table suggests, it is usually easy to solve for $x = F(x)$ under typical scenarios, and very few Newton iterations are needed to achieve full accuracy in 64-bit floating point. The only case where a Newton step failed and Picard iteration was needed was the contrived scenario M37h.

Note that the Brouwer fixed point theorem does not apply to $x = \eta + Ax$ since there is no vector of upper bounds μ in that case. Of course, singularity of the $A - I$ matrix is not a problem in practice, but there were two cases where this simple linear system failed to give a positive solution as indicated by the "—" entries in Table III. This can happen if entries of A are large enough to allow $\|Ax\| > \|x\|$ for some vectors x.

As can be seen from Table IV, the simple $\eta + Ax$ function gives different (and presumably less accurate) results even when it does yield a reasonable solution. Contrast this with the resource-saving transformations of Section IV which never had a significant effect on the solution.

TABLE IV
RELATIVE DIFFERENCES ($\|x - y\| / (\|x\| + \|y\|)$) BETWEEN SOLUTION VECTORS x AND y FOR EACH PAIR OF DIFFERENT F DEFINITIONS.

	$Ax + \eta$ vs. $F\S III$	$Ax + \eta$ vs. $F\S IV$	$F\S III$ vs. $F\S IV$
H57	0.012	0.012	0.0019
M39	0.29	0.29	0.013
M39h	—	—	0.029
M60	0.090	0.090	0.0036
M115	—	—	0.0065

Only the M39h scenario (a portion of which was shown in Figure 1) had any x_k values large enough to trigger reasonable noise rise limits. Qualitatively similar results were obtained from the simple composition of F with a function that smoothly limits each component to $\hat{\rho}\eta_k$ and from the $h(x)$ scheme. When $\bar{\rho}$ is set so both schemes reduce x_7 by roughly a factor of 2, other sectors are only modestly affected; e.g., the noise rise schemes reduce x_8 from 31.60pW to 28.11pW and 26.85pW, respectively. The main difference is that the $h(x)$ scheme makes it hard to find a fixed point unless there is a lot of smoothing in the $h(x)$ function. In fact, Newton steps and Picard iterations can both fail.

VII. Conclusions

We have given realistic models for uplink power control for voice and circuit data services under a variety of technologies. Like our previous work [6], [7], [8], it is motivated by practical optimization of wireless systems. The use of averages and fading probabilities smooths out discontinuities that could otherwise prevent the model from having a solution. Although incorporating outage probability leads to a nonlinear system for which we cannot guarantee convergence, we have ensured that F and its Jacobian are easy to evaluate, and the hybrid algorithm for finding the fixed point has proved to be very reliable in practice.

References

[1] W.-J. Oh and K. M. Wasserman, "Optimality of greedy power control and variable spreading gain in multi-class CDMA mobile networks," in *Proc. of the 5th Annual ACM/IEEE Int. Conf. on Mobile Computing and Networking*, 1999, pp. 102–112.

[2] J. S. Lee and L. E. Miller, *CDMA Systems Engineering Handbook*. Artech House, 1998.

[3] S. Kandukuri and S. Boyd, "Optimal power control in interference-limited fading wireless channels with outage-probability specifications," *IEEE Trans. Wireless Commun.*, vol. 1, no. 1, pp. 46–55, Jan. 2002.

[4] R. D. Yates, "A framework for uplink power control in cellular radio systems," *IEEE J. Sel. Areas in Commun.*, vol. 13, no. 7, pp. 1341–1347, Sept. 1995.

[5] C. Nuzman, "Note on convergence of a non-monotonic power control model," 2006, private communication.

[6] K. L. Clarkson and J. D. Hobby, "A model of soft handoff under dynamic shadow fading," in *VTC-2004-Spring: IEEE Vehicular Technology Conference*, vol. 3, 2004, pp. 1534–1538.

[7] ——, "A model of coverage probability under shadow fading," 2003, to appear. [Online]. Available: http://cm.bell-labs.com/cm/cs/who/clarkson/covprob/p.pdf

[8] ——, "Ocelot's knapsack calculations for modeling power amplifier and walsh code limits," 2006, to appear. [Online]. Available: http://cm.bell-labs.com/cm/cs/who/hobby/lossmodel.pdf