

Mobile power management for wireless communication networks

John M. Rulnick^a and Nicholas Bambos^b

^a *Department of Electrical and Computer Engineering, WPI, Worcester, MA 01609, USA*

^b *Department of Engineering Economic Systems and Operations Research, Stanford University, Stanford, CA 94305, USA*

For fixed quality-of-service constraints and varying channel interference, how should a mobile node in a wireless network adjust its transmitter power so that energy consumption is minimized? Several transmission schemes are considered, and optimal solutions are obtained for channels with stationary, extraneous interference. A simple dynamic power management algorithm based on these solutions is developed. The algorithm is tested by a series of simulations, including the extraneous-interference case and the more general case where multiple, mutually interfering transmitters operate in a therefore highly responsive interference environment. Power management is compared with conventional power control for models based on FDMA/TDMA and CDMA cellular networks. Results show improved network capacity and stability in addition to substantially improved battery life at the mobile terminals.

1. Introduction

Whereas overcoming interference has always been a central concern in the design of wireless networks, doing so while conserving energy is important for a growing class of users. Today's first and second generation networks are seeing rapid growth in the use of devices such as personal digital assistants (PDAs), palm-top and notebook computers, and lightweight mobile phones, while future micro-cell and personal communication services (PCS) networks are being designed for very low power and very high mobility. For users of battery-powered communication equipment it is advantageous – or necessary – to minimize energy consumption, provided, of course, that quality-of-service requirements are met. For example, it may be desirable to suspend transmission at moments when interference is unusually high. How to autonomously determine when, and at what power, a mobile terminal should attempt transmission is the subject of the present article.

By contrast, existing power control schemes generally focus on capacity and quality issues. Power control is employed as a means of balancing received power levels [7,10,16,22,31,32]; or balancing or guaranteeing signal-to-interference ratios (SIRs) [1–5,8,12–14,18–20,23,33,35,37–39], typically at the maximum possible common SIR; and it is sometimes integrated with other network management tasks such as base or channel assignment [9,15,26,36]. For cellular systems in particular, the point has been to minimize co-channel interference (as, for example, under TDMA, FDMA, PRMA, or related protocols¹) and/or near-far effects (as under spread-spectrum schemes such as code-division multiple access (CDMA)).

Here we consider a transmitter sending data to a remote terminal or base station via a communication channel subject to time-varying interference. The goal is to guarantee quality of service (expressed as an information transmis-

sion rate or average delay) while conserving energy, in order to extend the life of the battery [28]. Transmitting at a higher power yields a higher SIR and thus a higher success rate, but at the cost of more rapidly exhausting the energy supply – not to mention causing increased interference to other users. This idea translates directly to the optimization problem discussed in section 2. We first solve this problem under the assumption that interference is extraneous (i.e., unresponsive to the user's own transmissions) and time-stationary. We compare the optimal power management results to constant-SIR and constant-power solutions.

Then in section 3 we discuss the relationship between rates and delays in the context of random data arrival streams. Simple analysis of a special case provides us with a baseline for determining delays in complicated operating regimes. In section 4 we use the aforementioned extraneous-interference results as the basis for a distributed power management algorithm which dynamically selects power levels, and which works in responsive interference environments. In section 5 we test this algorithm, and verify our analysis, by simulation. Concluding remarks and future research directions are discussed in section 6.

2. Power management amidst stationary, unresponsive interference

Consider a single network node attempting to send information to another node at a particular rate while minimizing energy consumption. Its communication channel is subject to time-varying interference. If the received signal power is p during a time slot when received interference power is i then successful reception of the data occurs with probability $s(i, p)$, and failure occurs with probability $e(i, p) = 1 - s(i, p)$. Success or failure is denoted by $\chi_t = 1$ or 0, respectively, at discrete time t .

Unless otherwise noted we shall equate the transmitted and received power p , i.e., normalize so the desired link has gain 1, as shown in figure 1. The SIR at the receiver is de-

¹ TDMA = time-division multiple access, FDMA = frequency-division multiple access, PRMA = packet-reservation multiple access.




Figure 1. Illustration of transmitted and received powers before and after normalization.

Table 1
Some frequently used notation.

e	bit error probability function
γ	signal to interference ratio (SIR)
I_t	interference power at time t
i	value of interference power
p	transmitted or received signal power
r	target transmission rate
\bar{r}	actual transmission rate
s	instantaneous transmission rate ($1 - e$)
θ	threshold; $p = 0$ when $I_t > \theta$

noted by $\gamma = p/i$, assuming interference power $i > 0$. For concreteness we shall call these “per-bit” quantities, though in practice it is likely that the quantity of interest will be a larger unit of data (e.g., byte, packet). We assume erant bits are simply retransmitted. The interference power, which is for now assumed to be unresponsive to our transmitter power (i.e., generated extraneously), is modeled as a stationary, ergodic stochastic process $\{I_t: t = 1, 2, \dots\}$. Note that, in accord with the normalization illustrated in figure 1, when the desired link gain varies with time (as for fast-moving mobiles) the time variation must be accounted for in the process $\{I_t\}$. Please refer to table 1 for a summary of our notation.

2.1. Power management problem formulations

In general we seek a rule $p = p(i)$ which minimizes energy consumption subject to maintaining a transmission rate $\bar{r} \geq r$ for some fixed r , $0 < r < 1$. The quantity \bar{r} is defined by

$$\bar{r} = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t=1}^n \chi_t,$$

when the limit exists. Since interference is stationary and ergodic, and p is a function of i , we know the limit exists with probability one and $\bar{r} = E[\chi_1]$ (where E denotes expectation with respect to an underlying probability measure P). But $E[\chi_1] = E[s(I_1, p(I_1))]$, which in turn equals

$$\int_0^{\infty} s(i, p(i)) dP(I_1 \leq i)$$

(the latter is a Lebesgue–Stieltjes integral; see, e.g., [17, 27]). So our optimization problem takes the form

$$\min_{p \geq 0} \int_0^{\infty} p(i) dP(I_1 \leq i), \quad (1)$$

$$\int_0^{\infty} s(i, p(i)) dP(I_1 \leq i) \geq r. \quad (2)$$

The function p which solves this problem will be referred to as the (relatively) unrestricted power management function.

Before proceeding let us note an important assumption underlying this formulation: that signal processing at the receiver is sufficient to provide an accurate estimate of the interference power. In practice the interference power may need to be estimated from other parameters such as total received power, SIR, or recent error frequencies. Also, we assume network signaling makes this estimate available to the transmitter as needed, and the update cycles are shorter than the time over which interference power changes substantially. In other words, interference should not vary too much over a round-trip time for data, which, for example, in a cellular network is typically on the order of 10^{-5} to 10^{-6} seconds [11].

Although solving (1) and (2) yields our optimal power management solution, we would like to gauge system performance and energy savings relative to conventional power control. Moreover, not all practical transmitters will have complete freedom to select power levels. For these two reasons we consider “constant-SIR” and “constant-power” power management schemes below. Constant-SIR solutions are obtained by solving (1) and (2) subject to the additional constraint that $p(i)$ must be a constant multiple of i , i.e.,

$$p(i) = \hat{\gamma}i,$$

where $\hat{\gamma} > 0$ is a parameter of the minimization. Likewise we obtain “constant-SIR with threshold” solutions by solving (1), (2), and

$$p(i) = \begin{cases} \hat{\gamma}i, & \text{if } i \leq \theta, \\ 0, & \text{if } i > \theta, \end{cases} \quad (3)$$

where θ is the interference power threshold beyond which transmission is suspended.

Analogous to the constant-SIR solutions are “constant-power” solutions to (1), (2), and

$$p(i) = \hat{p}, \quad (4)$$

for $\hat{p} > 0$. Finally, we also consider “constant-power with threshold” solutions, which satisfy (1), (2), and

$$p(i) = \begin{cases} \hat{p}, & \text{if } i \leq \theta, \\ 0, & \text{if } i > \theta. \end{cases} \quad (5)$$

2.2. Optimal power management solutions

It remains to consider the forms of s and $P(I_1 \leq \cdot)$. The appropriate choice for s will generally depend on physical conditions: the topology of the network, the type of protocol employed, the modulation scheme, the environment, etc. Consider the curves in figure 2, which show typical bit error rates versus SIR ($e = 1 - s$ versus γ) for different digital modulation schemes [25,34]. For example, the formula for the “NC-FSK, fade” curve is $1/(\gamma + 2)$; for “DPSK, no




Figure 2. Bit error rates (BERs) versus bitwise signal to interference ratio.

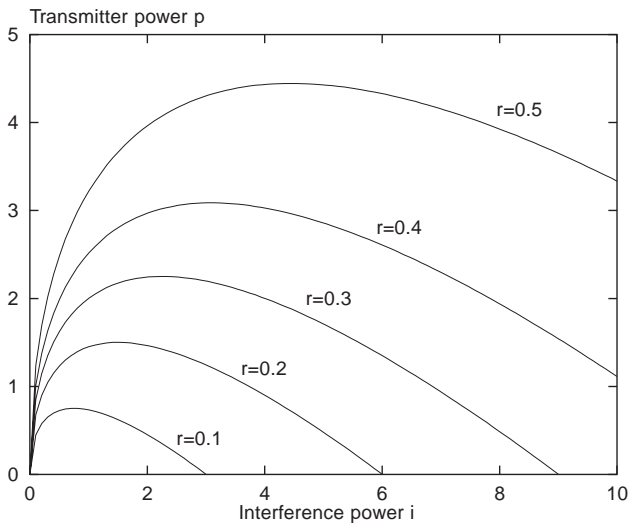


Figure 3. Optimal unrestricted transmitter power functions.

fade” we have $e = a \exp\{-b\gamma\}$; and for “BPSK, no fade”, $e = a \operatorname{erfc}\{\sqrt{b\gamma}\}$ ($a = 1/2$ and $b = 1$, as shown).²

We shall assume throughout section 2.2 that $e = 1/(\gamma + 1)$, so

$$s(i, p) = \frac{p}{p+i} = \frac{\gamma}{\gamma+1}, \quad (6)$$

and that interference is uniformly distributed, i.e.,

$$P(I_1 \leq i) = \frac{i}{I}, \quad 0 \leq i \leq I, \quad (7)$$

for some $I > 0$. We will eventually see that our most important conclusions are independent of the distribution of the interference process (see section 2.3).

Let us start by finding the solution to the relatively unconstrained problem defined by (1) and (2). Using calculus

²NC-FSK, fade = non-coherent frequency-shift keying over a fading channel; DPSK, no fade = differential phase-shift keying over a non-fading channel; BPSK, no fade = binary phase-shift keying over a non-fading channel.

of variations we find that the solution is a stationary point of the associated Lagrangian functional [6,30]

$$F = p - \lambda \frac{p}{i+p}, \quad (8)$$

for some constant (Lagrange multiplier) $\lambda > 0$. The stationary point satisfies

$$0 = \frac{\partial F}{\partial p} = 1 - \lambda \frac{i}{(i+p)^2}.$$

Taking the positive of two solutions we get

$$p(i) = \sqrt{\lambda i} - i. \quad (9)$$

First let $\lambda \leq I$; substituting (9) into, and taking equality in, eq. (2), yields $\lambda = 3Ir$, and hence $\lambda \leq I \Rightarrow r \leq 1/3$. On the other hand, taking $\lambda > I$ yields, by the same procedure, $\lambda = 4I/[9(1-r)^2]$, and $r > 1/3$. Summarizing,

$$p(i) = \begin{cases} \sqrt{3Ir} - i, & \text{if } r \leq 1/3 \text{ and } i \leq 3Ir, \\ 0, & \text{if } r \leq 1/3 \text{ and } i > 3Ir, \\ \frac{2\sqrt{Ii}}{3(1-r)} - i, & \text{if } r > 1/3. \end{cases} \quad (10)$$

This solution is shown in figure 3 for various values of r , using $I = 10$. Notice that when $r \leq 1/3$ we effectively use a threshold of $\theta = \lambda = 3Ir$, whereas for $r > 1/3$ we have $\theta = I$.

To better understand the form of this solution, let us contrast it with one obtained for an instantaneous transmission rate function such as

$$s(i, p) = \frac{p}{p+1} \frac{1}{i+1}.$$

Note that this s cannot be expressed strictly a function of constants and the SIR p/i . Omitting the details, we obtain $p(i) = \sqrt{\lambda/(i+1)} - 1$ which, unlike (10), is strictly decreasing in i . Whereas (10) indicates that transmitter power should go to 0 ($p \sim \sqrt{\lambda i}$) as interference $i \rightarrow 0$, here we have $p(i) \uparrow \sqrt{\lambda} - 1 > 0$. This is because here, for small i , s is approximately $p/(p+1)$, so p must be large in absolute terms to obtain a large s ; whereas in our analysis above $s = p/(p+i) \approx 1$ as long as p is a sufficiently large relative to i . Since error probabilities are usually (decreasing) functions of the SIR, we expect the form shown in figure 3 to serve as a useful indicator of general minimum-energy operation.

Using (10), the average energy consumption per unit time, or average transmitter power, is

$$E[p] = \begin{cases} \frac{3}{2}Ir^2, & \text{if } r \leq 1/3, \\ I\left(\frac{4/9}{1-r} - 1/2\right), & \text{if } r > 1/3. \end{cases}$$

We wish to compare this to the case of constant-SIR (with or without-threshold) transmission. In the threshold case we solve (1) subject to (2) and (3). We immediately find




Figure 4. Percent increase in useful battery life using unrestricted power management versus best-case constant-SIR operation.

$\hat{\gamma} = Ir/(\theta - Ir)$, and consequently that the solution calls for $\hat{\gamma} = 1$ if $r \leq 1/2$ and $\hat{\gamma} = r/(1 - r)$ otherwise, i.e.,

$$p(i) = \begin{cases} i, & \text{if } r \leq 1/2 \text{ and } i \leq 2Ir, \\ 0, & \text{if } r \leq 1/2 \text{ and } i > 2Ir, \\ \frac{ri}{1-r}, & \text{if } r > 1/2. \end{cases}$$

Notice that we have $\theta = 2Ir$ when $r \leq 1/2$, and otherwise $\theta = I$.

The average power in this case is

$$E[p] = \begin{cases} 2Ir^2, & \text{if } r \leq 1/2, \\ \frac{Ir}{2(1-r)}, & \text{if } r > 1/2. \end{cases}$$

Since useful battery life is inversely proportional to $E[p]$, we find for example that at $r = 0.25$ or $r = 0.99$ the optimal power management scheme provides a 33% or 13% increase, respectively, in the useful life of the battery versus best-case constant-SIR-with-threshold operation.

The constant-SIR (without threshold) solution may be obtained from the above by taking $\theta = I$. This yields

$$p(i) = \frac{ri}{1-r}.$$

This solution will be used later as a basis for comparing unrestricted power management in a variety of interference environments. The average power is

$$E[p] = \frac{Ir}{2(1-r)}.$$

Figure 4 shows the expected extension of useful battery life to be obtained by switching from constant-SIR to the unrestricted scheme in the stationary, unresponsive interference case.

For the constant-power scheme (the solution to eqs. (1) and (2) with (4) if a threshold is allowed, or (5) if not), we use equality in (2) to find

$$\theta = \hat{p} \left(\exp \left\{ \frac{Ir}{\hat{p}} \right\} - 1 \right). \quad (11)$$

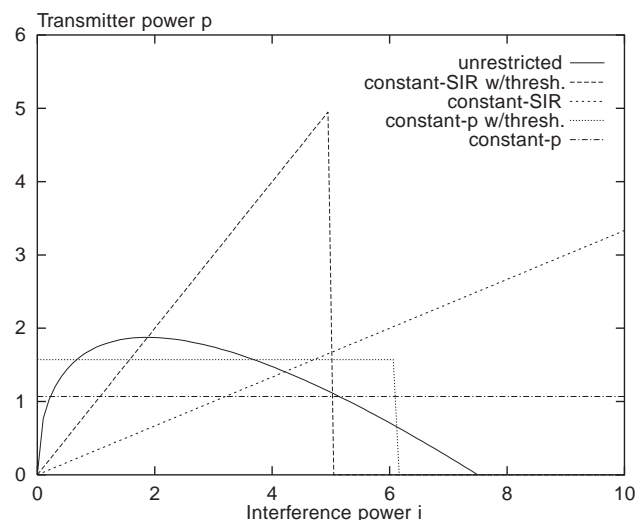


Figure 5. Transmitter power using various forms of power management, $r = 0.25$. The rate of energy consumption is proportional to the area under the curve.

From here it is necessary to use numerical root-finding methods. In the threshold case (when θ is free), an analytical solution proceeds only as far as the following relation:

$$\hat{p} \left(1 - \exp \left\{ -\frac{Ir}{\hat{p}} \right\} \right) = \frac{Ir}{2}.$$

In the other case ($\theta = I$ fixed) we may solve for \hat{p} in (11) using $\theta = I$. The average energy consumption per unit time in either case is

$$E[p] = \frac{\hat{p}\theta}{I}.$$

Now let us briefly compare methods. Figure 5 shows the optimal curves, taking $I = 10$ and $r = 0.25$, for each of the five cases we have considered. Recall that we have so far assumed $s = \gamma/(\gamma + 1)$ and uniformly distributed interference (what happens when these restrictions are relaxed is discussed in section 2.3). Because each I_t is uniform on $(0, 10)$ we can equate the (normalized) area under each curve with the average energy consumption per unit time for the corresponding method: 0.94 for the unrestricted case, 0.97 for constant-power with threshold, 1.07 for constant-power, 1.25 for constant-SIR with threshold, and 1.67 for constant-SIR. The constant-SIR schemes are harshly penalized at low values of r for using high power when interference is moderate to high. In this example the useful battery life is improved by 78% by using optimal unrestricted power management rather than constant-SIR transmission to obtain the specified transmission rate.

Increasing r to 0.99 (see figure 6) causes attempted transmission at all interference levels (i.e., $\theta = I$) no matter what the power management method, so the with- and without-threshold curves coincide. This yields an average power of 439.4 in the unrestricted case, and 495.0 in the constant-SIR and constant-power cases. At this higher transmission rate the constant-power scheme unnecessarily uses a lot of power at low interference levels, whereas the

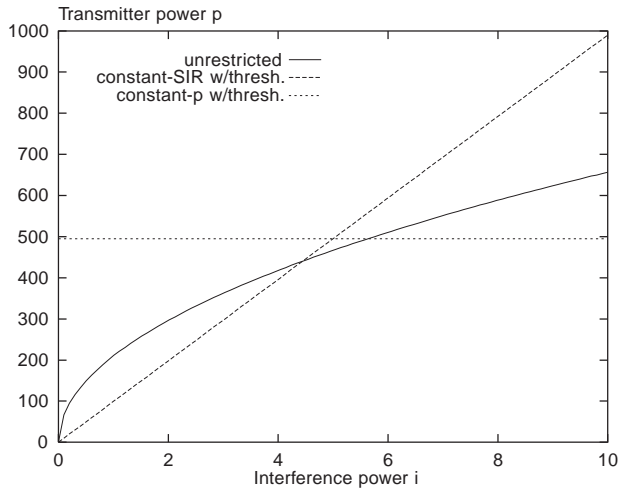


Figure 6. Transmitter power functions as in figure 5, but with $r = 0.99$.

disparity between the constant-SIR and unrestricted solutions is narrowed significantly at high interference levels, relative to the $r = 0.25$ case.

2.3. Arbitrary interference distributions

In section 2.2 we focused on specific forms of the success probability function and the interference power distribution, namely (6) and (7), respectively. What can we say about the optimal unrestricted power management function in the general case? Though the optimal p will always depend explicitly on s , it turns out that *the form of the optimal power management function p is independent of the distribution of interference power*. In fact, by calculus of variations and the method of Lagrange applied to (1) and (2) (see, e.g., [30, pp. 368–370]), one can show that the general minimum-energy solution must satisfy Euler’s equation

$$\frac{d}{di} \left(\frac{\partial F}{\partial p'} \right) = \frac{\partial F}{\partial p},$$

where $F = p - \lambda s$; cf. (8). It follows that the form of p is obtained simply by solving

$$\frac{\partial e}{\partial p} = -\frac{1}{\lambda}, \quad (12)$$

where $e = 1 - s$ is the bit error probability function.

For example, we can generalize our above discussion to consider a bit error probability function

$$e = \frac{b - a}{\gamma + b}, \quad a < b,$$

– we have used $a = 0, b = 1$ so far – and get, via (12),

$$p(i) = \sqrt{\lambda(b - a)i} - bi$$

(cf. eq. (9)). The “NC-FSK, fade” curve in figure 2 is given by $1/(\gamma + 2)$, so the optimal p using NC-FSK modulation




Figure 7. Markov chain for bits of data present, with interference-dependent transition rates.

in a fading environment would be $p(i) = \sqrt{\lambda i} - 2i$. For DPSK over a non-fading channel we have

$$e = a \exp\{-b\gamma\}, \quad a > 0, b > 0,$$

and for BPSK over a non-fading channel,

$$e = a \operatorname{erfc}\{\sqrt{b\gamma}\}, \quad a > 0, b > 0.$$

The corresponding solutions to (12) are found to be

$$p(i) = -\frac{i}{b} \ln\left(\frac{i}{ab\lambda}\right) \quad (\text{DPSK, no fading}),$$

and, via Leibniz’s rule, the $p = p(i)$ which satisfies

$$p^{-1/2} \exp\left\{-\frac{bp}{i}\right\} = \frac{1}{a\lambda} \sqrt{\frac{\pi i}{b}} \quad (\text{BPSK, no fading}).$$

In any case, after solving (12), λ is obtained (when it exists, i.e., when the transmission rate r is not unattainable) by substituting the value of p obtained from (12) into the constraint (2) with ‘ $\geq r$ ’ replaced by ‘ $= r$ ’. So naturally the constant λ (constant relative to i) will in general depend on r and the distribution of the interference power. Fortunately, this dependence does not hamper implementation of a dynamic power management algorithm, as we shall see in section 4.

3. Quality of service: Rates versus delays

Thus far our analysis has used a guaranteed mean transmission rate, r , as the main performance constraint, with the implicit assumption that there are always data to transmit. Let us now turn to delays, and assume that a single unit of data (again, say, a bit) arrives with probability $a > 0$ during any particular time slot, independent of others, and is queued there in an infinite buffer until successfully transmitted. Under certain circumstances the state of the transmitter may then be modeled as the birth–death Markov chain shown in figure 7, where state k represents the presence of k bits. In general, under power management, the rates r_k will depend on the statistics of the interference, and we obtain a stationary probability distribution $\{\pi_k: k = 0, 1, \dots\}$ of the form

$$\pi_k = \frac{a^k(1 - r_1)(1 - r_2) \cdots (1 - r_{k-1})}{(1 - a)^k(r_1 r_2 \cdots r_k)} \pi_0, \quad k = 1, 2, \dots, \quad (13)$$

as long as the chain is stable, i.e.,

$$\sum_{k=1}^{\infty} \frac{a^k(1-r_1)(1-r_2)\cdots(1-r_{k-1})}{(1-a)^k(r_1r_2\cdots r_k)} < \infty.$$

To get some analytical insight into the relationship between delays and transmission rates (and to justify the Markov chain model), we need to impose some kind of structure on the interference. In [29] we use matrix-geometric methods [24] to obtain an expression for delays when interference is a stationary continuous-time Markov process unaffected by transmitter power. With a stronger assumption, namely that $\{I_t\}$ is an independent, identically distributed (i.i.d.) sequence (again unaffected by our transmitter power), we can obtain a simple expression for the mean delay. In this case $r_k = r$ for all k , and (13) becomes

$$\pi_k = \frac{a^k(1-r)^{k-1}}{[(1-a)r]^k} \pi_0, \quad k = 1, 2, \dots$$

Solving $1 = \sum_{k \geq 0} \pi_k$ for π_0 yields $\pi_0 = (r-a)/r$. From this it follows that the mean number of bits present will be $a(1-a)/(r-a)$ and hence, by Little's law, that the expected bit delay is given by

$$E[D] = \frac{1-a}{r-a}, \quad r > a. \quad (14)$$

Eq. (14), which assumes one-at-a-time arrivals and i.i.d. interference, provides a baseline for calculating delays. In the admittedly limited situations where these assumptions are strictly valid, we may express the optimization problems based on (1) and (2) as delay-constrained problems rather than rate-constrained. We generally expect delays to be greater than (14) when arrivals are bursty or batched, as will be the case in many real wireless communication systems. The effects of a non-i.i.d. or – more importantly – transmission-dependent environment are not so clear. In order to examine this situation we will, via simulation, move our energy-saving transmitter to a realm where other, interfering transmitters are operating and, moreover, are themselves assumed to be using power management.

4. A distributed dynamic power management algorithm

Two key assumptions about interference made in the preceding analysis were that it is (i) time-stationary and (ii) unresponsive to the optimizer's transmissions. It turns out that, fortunately for the sake of practical implementation, power management can still be employed when these assumptions are substantially relaxed. And while at least asymptotic stationarity will always be required for meaningful power management constrained by time-average parameters, our unresponsive-interference restriction can be completely relaxed. The simulator described in section 5 allows us to model environments which include any number of transmitters, any subset of which may be using some form of power management. Here we present a dynamic

power management algorithm (DPMA) which is intended to work in such environments, and which is based on the analytical results obtained in section 2.

The algorithm requires no prior knowledge of the operating environment; the only foreknowledge necessary is a good estimate of the error probability function e , since the algorithm will seek a solution satisfying (12). However, like the analysis on which it is based, the DPMA will assume that interference power can be measured reliably at the receiver and communicated rapidly to the transmitter; in some real situations this may not be possible, or at least not without additional overhead. Also, for concreteness, we will assume (6) holds, and hence so does (9), giving us explicit expressions in step 4 of the main algorithm and step 3 of the sub-algorithm, below. For other forms of e or s simply substitute the appropriate expressions into those steps.

Briefly, the algorithm works as follows: Starting at $t = 0$, interference levels are measured and the values used to update a frequency vector F ; the i th component, F_i , is incremented by one if the interference is in the interval $[i\rho, (i+1)\rho)$, where ρ is the (perhaps adaptively determined) "resolution". Based on F , r (the required transmission rate), and p (as determined by (12)), the Lagrange multiplier λ is numerically estimated to within some specified tolerance; put another way, (2) is solved, with $F_i/(t+1)$ in place of $dP(I_1 \leq i)$. Then this value of λ is used to obtain a numerical value for p , and transmission is attempted at power level p , assuming a bit (or appropriate unit of data) is ready for transmission.

Dynamic power management algorithm

Notes: The parameters ρ ("resolution") and r are assumed fixed in advance. The symbol ':= ' is used to denote assignment.

0. START: $F_i = 0$ for all i , $\lambda = 1$, $t = 0$.
1. Measure (or estimate) interference I .
2. Set $F_i := F_i + 1$ for $i = \lfloor I/\rho \rfloor$.
3. Update λ (see sub-algorithm below).
4. Set $p = \max\{0, \sqrt{\lambda I} - I\}$ and go to 1.
5. END.

Sub-algorithm for updating λ

Notes: This recursive algorithm is called from the main algorithm with initial values $\text{sign} = 0$ and $\text{step} > 0$. Also, the variable tol is assumed to be fixed and positive.

0. START: $i = 0$, $j = 0.5\rho$, $s = 0$.
1. If $j \geq \lambda$ go to 10.
2. If $F_i = 0$ go to 9.
3. Set $s := s + [F_i/(t+1)][1 - \sqrt{j/\lambda}]$.
4. If $s < r$ go to 9.
5. If $\text{sign} \neq 1$ go to 8.
6. If $\text{step} \leq \text{tol}$ go to 13.
7. Set $\text{step} := \text{step}/2$.
8. Set $\text{sign} = -1$ and go to 12.
9. Set $i := i + 1$, $j := j + \rho$, and go to 1.

10. If $\text{sign} = -1$ set $\text{step} := \text{step}/2$.
11. Set $\text{sign} = 1$.
12. Set $\lambda := \lambda + \text{sign} \cdot \text{step}$ and go to 0.
13. END.

Note that the sub-algorithm returns with a value of λ which satisfies $\lambda^* \leq \lambda < \lambda^* + \text{tol}$, where λ^* is the “true” value.

Implementation notes

1. *Finite memory and adaptive resolution.* In any real implementation there will be some upper limit on the size of the vector F being stored. This finiteness can be accommodated by allowing the resolution, ρ , to adapt to peak interference levels. In fact, we use just such a device in our simulation of section 5.

2. *Interpolation of the interference frequency vector.* In the sub-algorithm, the parameter j determines the centering of F 's quantization of interference. At the start of the sub-algorithm we use $j = 0.5\rho$. A more conservative design might use $j = \rho$.

3. *Constant-SIR or constant-power DPMA.* It is of course possible to use the DPMA without using the optimal, unrestricted solution for p . The expression for p in the main algorithm would be adjusted accordingly. Below, when we refer to DPMA, we mean the unrestricted form used in the algorithm as stated above, unless we explicitly indicate otherwise.

5. Network simulation

In this section we test the results of our analysis from sections 2 and 3, and the algorithm presented in section 4, by computer simulation. Note that the simulator makes the same optimistic assumption as discussed in section 2.1: that signalling and signal processing are sufficient to provide all necessary information to each transmitter at the times it adjusts its power. This means that total receiver interference is known and each transmitter knows (and compensates for) its own path loss; or, equivalently, that the received SIR and received signal power are known. While this is certainly an idealization in absolute terms, it is applied to both DPMA and constant-SIR simulations, so the numerical results discussed below should still provide a reasonable guide to the *relative* performances of power management and conventional power control.

Unless otherwise noted, data arrivals are one-at-a-time (i.e., non-batch) and simulations are run for 25,000 time steps. We use $\text{tol} = 0.125$ and a simple adaptive resolution scheme (see implementation note 1, above). Finally, in the multiple transmitter case we assume transmitter powers are updated in round-robin fashion at each time instant.

5.1. Unresponsive interference

We start with a very simple case to illustrate the functioning of the DPMA: a single transmitter in a stationary,

unresponsive, sinusoidal-power interference environment. Ideally one would like the algorithm to “learn” the sinusoidal environment within a few cycles and, if r is not too large, transmit only when the interference power is near its periodic minima. Taking $a = 1$ and $r = 0.25$ yields figure 8. The average power used in this 500-step window is only 7.26, while peak interference power is 100. The transmission rate is 0.28, which meets our specification of 0.25 or greater. Longer simulation (duration 10,000) yields average power of only 5.01 and transmission rate 0.25. By comparison, the minimum-energy constant-SIR (without threshold) scheme uses average power 16.7 to attain the same rate, 0.25. Hence the energy savings using DPMA is no less than 70%, and battery life is improved by a factor of at least 3.3, relative to any constant-SIR scheme which achieves this rate.

Increasing r to 0.5 yields figure 9. Notice that the transmitter is essentially always “on” in this case. The average power used here is 35.7, yielding a transmission rate of




Figure 8. Transmitter using DPMA in an unresponsive, sinusoidal interference environment, $r = 0.25$.




Figure 9. Transmitter using DPMA in the same interference environment as figure 8 (note change of scale), but with $r = 0.5$.

0.52. These values become average power 31.4 and rate 0.50 at time 10,000. The minimum-energy constant-SIR values under the same circumstances are average power 50.0 and rate 0.50; i.e., the transmission power precisely tracks the sinusoid. The energy saved by DPMA in this case is 37%, and thus battery life is improved by 59%. At $r = 0.99$ the energy savings is still a respectable 19%, and battery life is extended 23%.

In fact, simulations of a wide range of situations for a single DPMA-based transmitter operating in an unresponsive interference environment – whether sinusoidal, bi-level, constant, i.i.d. uniform, or varying according to a finite birth–death Markov chain – demonstrate agreement with the analysis of section 2. Specified transmission rates are attained at the expected, minimal average transmission power levels. We also find delay values matching our analytical computations for i.i.d. interference (see eq. (14)) when we incorporate random data arrivals into the model. However, a penalty is incurred when using DPMA for low rs in the Markovian birth–death interference case: Delays can become substantially greater when interference lingers near high levels, during which the transmitter is turned down or off. For example, with $a = 0.05$ and $r = 0.1$, delays under DPMA are greater than those under constant-SIR by a factor greater than five. On the bright side, the battery life in this case is extended by a factor greater than ten. We shall omit further detail for the unresponsive-interference case and proceed now to the question of transmitter interaction.

5.2. Multiple transmitters: Responsive interference

We shall find it useful to use the cellular paradigm in describing our multiple transmitter simulations. Consider figure 10, which shows a typical uplink FDMA/TDMA cellular model. Two mobiles (M1 and M2) are using one particular frequency/time slot in different cells. The link gains $g_{j,k}$ from mobile k to base j are shown, with $g_{1,2}p_2$ representing co-channel interference at base station 1 (B1) if M2 uses power p_2 . The total interference at B1 is the sum of co-channel interference and noise, shown by dashed lines, and thus the SIR is $g_{1,1}p_1/(g_{1,2}p_2 + \xi_1)$, where we denote the noise at receiver j by ξ_j .




Figure 10. Typical FDMA/TDMA cellular system with uplink paths.

For a gain matrix $G = [g_{j,k}]$ given by

$$G_1 = \begin{bmatrix} 1 & 0.1 \\ 0.1 & 1 \end{bmatrix}$$

and i.i.d. noise uniform on $(0, 100]$ at each base, we obtain the data shown in table 2. The delay entries are blank when $a = 1$, since then $E[D] \rightarrow \infty$ with time; and the rate \bar{r} entries are blank when $a < r$, since then $\bar{r} \rightarrow a$ if the system is stable. The performance measures shown are averages for the two transmitters.

For five transmitters and a gain matrix

$$G_2 = \begin{bmatrix} 1 & 0.1 & 0.1 & 0.1 & 0.1 \\ 0.1 & 1 & 0.1 & 0.1 & 0.1 \\ 0.1 & 0.1 & 1 & 0.1 & 0.1 \\ 0.1 & 0.1 & 0.1 & 1 & 0.1 \\ 0.1 & 0.1 & 0.1 & 0.1 & 1 \end{bmatrix}$$

we obtain the data shown in table 3. In some cases the system saturated (i.e., the adaptive resolution, ρ , exceeded some upper bound) before the simulation was complete. In such cases the system is considered unstable. We will discuss stability in more detail below.

In general, these simulations show that the DPMA significantly extends battery life by using much less power on average, with no visible tradeoff in performance (rates

Table 2
Simulation results for two transmitters and gain matrix G_1 , with i.i.d. uniformly distributed receiver noise.

$N = 2, G = G_1$		Unrestricted DPMA			Constant-SIR		
arrival rate a	target rate r	average power	actual rate \bar{r}	mean delay	average power	actual rate \bar{r}	mean delay
1	0.1	1.53	0.098		5.62	0.100	
1	0.5	44.9	0.502		55.7	0.505	
1	0.9	3,620	0.899		4,500	0.901	
0.05	0.1	0.726		17.8	2.82		19.1
0.25	0.5	20.7		2.97	26.2		3.02
0.45	0.9	336		1.23	426		1.22
0.1	0.5	7.97		2.27	10.1		2.22
0.1	0.9	48.2		1.13	56.1		1.13

Table 3
Simulation results for five transmitters and gain matrix G_2 , with i.i.d. uniformly distributed receiver noise.

$N = 5, G = G_2$		Unrestricted DPMA			Constant-SIR		
arrival rate a	target rate r	average power	actual rate \bar{r}	mean delay	average power	actual rate \bar{r}	mean delay
1	0.1	1.71	0.102		5.82	0.101	
1	0.5	74.2	0.502		83.3	0.500	
1	0.9	saturated at $t = 7$			saturated at $t = 3$		
0.05	0.1	0.801		20.4	2.90		19.9
0.25	0.5	26.0		3.00	31.4		2.97
0.45	0.9	saturated at $t = 138$			saturated at $t = 6$		
0.1	0.5	8.61		2.24	10.9		2.23
0.1	0.9	65.9		1.12	85.9		1.12

Table 4

Simulation results for two transmitters and gain matrix G_1 , with birth-death Markovian receiver noise.

$N = 2, G = G_1$		Unrestricted DPMA			Constant-SIR		
arrival rate a	target rate r	average power	actual rate \bar{r}	mean delay	average power	actual rate \bar{r}	mean delay
1	0.1	0.0813	0.0996		5.51	0.101	
1	0.5	44.3	0.514		55.5	0.504	
1	0.9	3,510	0.898		4,370	0.900	
0.05	0.1	0.213		134	2.76		16.8
0.25	0.5	21.6		5.45	27.0		3.06
0.45	0.9	340		1.22	443		1.21
0.1	0.5	8.57		2.88	10.4		2.28
0.1	0.9	43.7		1.13	53.2		1.14

Table 5

Simulation results for five transmitters and gain matrix G_2 , with birth-death Markovian receiver noise.

$N = 5, G = G_2$		Unrestricted DPMA			Constant-SIR		
arrival rate a	target rate r	average power	actual rate \bar{r}	mean delay	average power	actual rate \bar{r}	mean delay
1	0.1	0.487	0.106		5.91	0.102	
1	0.5	74.1	0.506		82.4	0.499	
1	0.9	saturated at $t = 6$			saturated at $t = 3$		
0.05	0.1	0.246		118	2.99		20.5
0.25	0.5	26.5		4.06	31.0		2.96
0.45	0.9	saturated at $t = 1168$			saturated at $t = 27$		
0.1	0.5	9.17		2.74	11.0		2.26
0.1	0.9	65.8		1.13	87.8		1.13

or delays). Delays are very nearly the same as those predicted by eq. (14). We can compare these results to cases where the noise power is not i.i.d. uniform but instead varies according to a birth-death Markov chain (BDMC). If the BDMC is balanced and has roughly the same range as the i.i.d. uniform process then the steady state distributions will be approximately the same; however, the sample paths will differ in that the BDMC will tend to dwell for longer periods at high or low interference levels. The result, similar to the single transmitter case and shown clearly in tables 4 and 5, is that, for low r , delays under the “patient” DPMA will be much greater than under the relatively “impatient” constant-SIR scheme.

Now consider figure 11, which shows a typical uplink CDMA cellular model. Two mobiles (M1 and M2) are using one particular frequency/time slot in the same cell, and signal separation is maintained by coding. In this case, since all mobiles communicate with the same base station, $g_{j,k} = g_k$ for all j and the gain matrix G will have identical rows.³ Hence the system quality and capacity will depend on near-far effects rather than co-channel interference, and interfering-link gains can be greater than the signal-link

³ A more general representation would allow the multiplication of elements along the main diagonal by some processing gains greater than or equal to one.




Figure 11. Typical CDMA cellular system with uplink paths.

Table 6

Simulation results for two transmitters and gain matrix G_3 , with i.i.d. uniformly distributed receiver noise.

$N = 2, G = G_3$		Unrestricted DPMA			Constant-SIR		
arrival rate a	target rate r	average power	actual rate \bar{r}	mean delay	average power	actual rate \bar{r}	mean delay
1	0.1	9.73	0.099		34.4	0.100	
1	0.5	2,680	0.480		saturated at $t = 3881$		
1	0.9	saturated at $t = 3$			saturated at $t = 2$		
0.05	0.1	4.66		18.7	16.5		19.7
0.25	0.5	198		3.29	433		2.95
0.45	0.9	saturated at $t = 301$			saturated at $t = 10$		
0.1	0.5	53.0		2.27	86.1		2.36
0.1	0.9	384		1.13	saturated at $t = 3343$		

gain for any mobile.

For a gain matrix $G = [g_{j,k}]$ given by

$$G_3 = \begin{bmatrix} 1 & 0.1 \\ 1 & 0.1 \end{bmatrix},$$

and i.i.d. uniform noise at the base station (analogous to the TDMA results in table 2) we obtain the data shown in table 6. For five transmitters and a gain matrix

$$G_4 = \begin{bmatrix} 1 & 0.5 & 0.25 & 0.125 & 0.0625 \\ 1 & 0.5 & 0.25 & 0.125 & 0.0625 \\ 1 & 0.5 & 0.25 & 0.125 & 0.0625 \\ 1 & 0.5 & 0.25 & 0.125 & 0.0625 \\ 1 & 0.5 & 0.25 & 0.125 & 0.0625 \end{bmatrix}$$

we obtain the data shown in table 7. Again, values are averages over all N transmitters. Finally, for the CDMA model with noise which varies according to a birth-death Markov chain (BDMC), we have tables 8 and 9. Notice the substantially lower average powers, and the tradeoff with higher delays for low r . As one would expect, additional simulation results (not shown) indicate that batch arrivals increase delays further.

Concerning stability, it is interesting to note that in the cases where interfering link gains are high, transmitters which fail (saturate) while using conventional constant-SIR

Table 7

Simulation results for five transmitters and gain matrix G_4 , with i.i.d. uniformly distributed receiver noise.

$N = 5, G = G_4$		Unrestricted DPMA			Constant-SIR		
arrival rate a	target rate r	average power	actual rate \bar{r}	mean delay	average power	actual rate \bar{r}	mean delay
1	0.1	24.7	0.103		62.0	0.101	
1	0.5	saturated at $t = 88$			saturated at $t = 2$		
1	0.9	saturated at $t = 0$			saturated at $t = 1$		
0.05	0.1	7.01		18.3	23.1		20.0
0.25	0.5	saturated at $t = 1266$			saturated at $t = 16$		
0.45	0.9	saturated at $t = 1$			saturated at $t = 2$		
0.1	0.5	134		2.50	saturated at $t = 380$		
0.1	0.9	10,400		1.22	saturated at $t = 19$		

Table 8

Simulation results for two transmitters and gain matrix G_3 , with birth-death Markovian receiver noise.

$N = 2, G = G_3$		Unrestricted DPMA			Constant-SIR		
arrival rate a	target rate r	average power	actual rate \bar{r}	mean delay	average power	actual rate \bar{r}	mean delay
1	0.1	1.51	0.104		33.0	0.098	
1	0.5	1,970	0.468		saturated at $t = 5258$		
1	0.9	saturated at $t = 3$			saturated at $t = 2$		
0.05	0.1	1.23		110	16.1		18.1
0.25	0.5	201		3.91	477		3.01
0.45	0.9	saturated at $t = 5$			saturated at $t = 21$		
0.1	0.5	51.8		2.70	79.3		2.23
0.1	0.9	400		1.13	saturated at $t = 2584$		

Table 9

Simulation results for five transmitters and gain matrix G_4 , with birth-death Markovian receiver noise.

$N = 5, G = G_4$		Unrestricted DPMA			Constant-SIR		
arrival rate a	target rate r	average power	actual rate \bar{r}	mean delay	average power	actual rate \bar{r}	mean delay
1	0.1	18.8	0.115		61.7	0.100	
1	0.5	saturated at $t = 83$			saturated at $t = 2$		
1	0.9	saturated at $t = 1$			saturated at $t = 1$		
0.05	0.1	1.51		122	22.3		19.6
0.25	0.5	saturated at $t = 1991$			saturated at $t = 14$		
0.45	0.9	saturated at $t = 1$			saturated at $t = 1$		
0.1	0.5	128		2.68	saturated at $t = 38$		
0.1	0.9	11,300		1.21	saturated at $t = 146$		

power control succeed while using the DPMA (see tables 6–9). There are two reasons. First, for a large set of operating parameters, the optimal unrestricted power management function calls for reduced or zero transmitter power when interference is high. This tends to moderate the competition for high SIRs.

Second, note that, under (6), transmitter j 's instantaneous successful transmission probability is

$$s_j = \frac{p_j}{p_j + g_{j,j}^{-1}(\sum_{k \neq j} g_{j,k} p_k + \xi_j)}$$




Figure 12. Response of constant-SIR and DPMA schemes to an unattainable data transmission rate. The two transmitters rapidly saturate under constant-SIR, but do so only very gradually under DPMA.

Assume, for instance, that N transmitters with identical target rates are operating simultaneously. Using constant-SIR, the limiting value for s_j (when $p_k = p_j \rightarrow \infty$ for all k) is easily seen to be

$$r^* = \frac{1}{1 + g_{j,j}^{-1} \sum_{k \neq j} g_{j,k}}$$

So, for example, with $N = 2$, $g_{j,j} = 1$, and $g_{j,k} = 0.1$ (-10 dB) for $k \neq j$, the upper limit for s and hence for \bar{r} is $r^* = 0.909$. This in turn means that transmission rates $r > 0.909$ are unattainable using constant-SIR transmission. Yet, using unrestricted power management, the DPMA algorithm is able to yield stable operation and near-target rates or delays for very long times. The main reason is that the constant-SIR scheme attempts to reach its target (SIR or rate) immediately, while the more patient DPMA methodically accumulates data on its environment, and hence only gradually increases the transmitter power in persistently high-interference regimes. This is illustrated in figure 12 for $r = 0.91$. Meanwhile, as our simulation results show, the transmitter can operate at near-target performance. While the sensitivity of the DPMA can be easily adjusted if increased responsiveness is desired, the DPMA's measured approach clearly introduces a “slow-blow” kind of failure, as opposed to the rapid saturation and drop-out potentially seen by users of conventional power control. So enhanced stability and capacity may be important side benefits of using power management.

6. Conclusion and future directions

The motivation for the introduction of system-level power management is clear enough: a vision of highly mobile, always ready, data/voice/video communications for personal communication systems (PCS), micro-cellular networks, and other applications on the horizon. Indeed,

power management may be seen as “the heart of mobility”, since battery technology is presently so limiting [21].

We have attempted to illustrate the potential benefits of low-power, rate- or delay-constrained operation. By contrast, conventional power control, because of its use of received SIRs and power levels (and ultimately network capacity) as performance measures, did not consider this way of exploiting fluctuations in interference to save energy. (To be sure, if interference power approaches a constant – a common assumption in power control analyses – the optimal unrestricted solution converges to the constant-SIR solution.) The kind of power management introduced here explicitly capitalizes on variations in interference power, behaving (like carrier-sense multiple access (CSMA)) as an enhanced “listen-first” kind of protocol. Still, we are presently considering other ways for mobile devices or the entire network to save energy, such as network management which encourages cooperation between users.

While our results indicate the possibility of substantially reducing energy consumption without sacrificing quality of service, and possibly enhancing network stability and capacity along the way, it is important to consider the simplifying assumptions we have made. Perhaps most importantly, we have assumed accurate measurements or estimates of interference power are available as needed. More realistically, two problems will arise: Interference power (and/or path losses) will be coarsely estimated, and the estimates will arrive late. The latter problem – communication delay in a feedback control system – is of course not at all unique to power management or power control, nor are its ramifications. It is possible that these delays will substantially erode system performance or introduce instabilities; further study is needed.

The former problem – difficulty in precisely determining the interference power and/or desired-link gain – is well known in the realm of power control. What if one or both of these parameters are unknown? It may be necessary to guess them from recent error rates. Alternatively, it is often assumed in power control analyses that at least the received SIR at time t , γ_t , can be calculated accurately and communicated quickly to the transmitter. In this case it is reasonable to use the ratio p_t/γ_t – where p_t is the transmitter power at time t – as an approximation or prediction of I_{t+1}/g_{t+1} , the anticipated receiver interference normalized by the link gain. Finally, in circumstances where interference power varies too rapidly for the transmitter to react, or simply does not vary, the present brand of power management offers little or no advantage over conventional power control, and other energy-saving methods will have to be found.

Acknowledgements

The authors greatly appreciate the comments and suggestions of the anonymous referees who reviewed versions of this work for IEEE Infocom '96 and *Wireless Networks*; the content and presentation improved substantially, thanks

to them. We also gratefully acknowledge the support of the Advanced Research Projects Agency via contract ARPA-CSTO-93-112 and the National Science Foundation via National Young Investigator Award NSF-NCR-9258507.

References

- [1] J. Aein, Power balancing in systems employing frequency reuse, COMSAT Technical Review 3(2) (1973).
- [2] M. Almgren, H. Andersson and K. Wallstedt, Power control in a cellular system, in: *Proceedings of the IEEE Vehicular Technologies Conference* (1994) pp. 833–837.
- [3] S. Ariyavisitakul, SIR-based power control in a CDMA system, in: *Proceedings of the IEEE Globecom* (1992) pp. 868–873.
- [4] S. Ariyavisitakul, Achievable performance of autonomous SIR-based power control, *Electronics Letters* (1993).
- [5] N. Bambos, S. Chen and G. Pottie, Radio link admission algorithms for wireless networks with power control and active link quality protection, in: *Proceedings of IEEE Infocom* (1995).
- [6] M. S. Bazaraa, H. D. Sherali and C. M. Shetty, *Nonlinear Programming: Theory and Algorithms*, Wiley-Interscience Series in Discrete Mathematics and Optimization (Wiley, New York, 2nd ed., 1993).
- [7] F. Bock and B. Ebstein, Assignment of transmitter powers by linear programming, *IEEE Transactions on Electromagnetic Compatibility* 6 (July 1964) 36–44.
- [8] G. J. Foschini and Z. Miljanic, A simple distributed autonomous power control algorithm and its convergence, *IEEE Transactions on Vehicular Technology* 42(4) (November 1993) 641–646.
- [9] G. J. Foschini and Z. Miljanic, Distributed autonomous wireless channel assignment algorithm with power control, *IEEE Transactions on Vehicular Technology* 44(3) (August 1995) 420–429.
- [10] T. Fujii and M. Sakamoto, Reduction of cochannel interference in cellular systems by intra-zone channel reassignment and adaptive transmitter power control, in: *Proceedings of the IEEE Vehicular Technology Conference* (1988) pp. 668–672.
- [11] D. Goodman, R. Valenzuela, K. Gayliard and B. Ramamurthi, Packet reservation multiple access for local wireless communications, *IEEE Transactions on Communications* 37(8) (August 1989) 885–890.
- [12] S. Grandhi, R. Vijayan and D. Goodman, A distributed algorithm for power control in cellular radio systems, in: *Proceedings of the Thirtieth Annual Allerton Conference on Communications, Control, and Computing* (1992).
- [13] S. A. Grandhi, R. Vijayan and D. J. Goodman, Distributed power control in cellular radio systems, *IEEE Transactions on Communications* 42(2–4) (February–April 1994) 226–228.
- [14] S. A. Grandhi and J. Zander, Constrained power control in cellular radio systems, in: *Proceedings of the IEEE Vehicular Technology Conference* (1994) pp. 824–828.
- [15] S. V. Hanly, An algorithm for combined cell-site selection and power control to maximize cellular spread spectrum capacity, *IEEE Journal on Selected Areas in Communications* 13(7) (September 1995) 1332–1340.
- [16] J. M. Holtzman, CDMA power control for wireless networks, in: *Third Generation Wireless Information Networks*, eds. S. Nanda and D. Goodman (Kluwer, 1992) pp. 299–311.
- [17] A. Kolmogorov and S. Fomin, *Introductory Real Analysis* (Dover, New York, 1970).
- [18] T.-H. Lee, J.-C. Lin and Y. T. Su, Downlink power control algorithms for cellular radio systems, *IEEE Transactions on Vehicular Technology* 44(1) (February 1995) 89–94.
- [19] H. J. Meyerhoff, Method for computing the optimum power balance in multi-beam satellites, COMSAT Technical Review 4(1) (Spring 1974) 139–146.
- [20] D. Mitra, An asynchronous distributed algorithm for power control in cellular radio systems, in: *Proceedings of the 4th WINLAB Workshop on 3rd Generation Wireless Information Networks* (1993).

- [21] H. Morbitzer, Guide to a longer life: Power management in mobile communications, *Philips Telecommunications Review* 52(1) (March 1994) 69–74.
- [22] T. Nagatsu, T. Tsuruhara and M. Sakamoto, Transmitter power control for cellular land mobile radio, in: *Proceedings of the IEEE Globecom* (1983) pp. 1430–1434.
- [23] R. W. Nettleton and H. Alavi, Power control for a spread spectrum cellular mobile radio system, in: *Proceedings of the IEEE Vehicular Technology Conference* (1983) pp. 242–246.
- [24] M. F. Neuts, *Matrix-Geometric Solutions in Stochastic Models: An Algorithmic Approach* (Dover, New York, 1994).
- [25] K. Pahlavan and A. H. Levesque, *Wireless Information Networks*, Wiley Series in Telecommunications and Signal Processing (Wiley, New York, 1995).
- [26] A. N. Rosenberg, Simulation of power control and voice-channel selection in cellular systems, in: *Proceedings of the IEEE Vehicular Technology Conference* (1985) pp. 12–15.
- [27] H. Royden, *Real Analysis* (Macmillan, New York, 3rd edn., 1988).
- [28] J. M. Rulnick and N. Bambos, Mobile power management for maximum battery life in wireless communication networks, in: *Proceedings of IEEE Infocom* (1996).
- [29] J. M. Rulnick and N. Bambos, Performance evaluation of power-managed mobile communication devices, in: *Proceedings of IEEE International Communications Conference (ICC)* (1996).
- [30] G. F. Simmons, *Differential Equations with Applications and Historical Notes* (McGraw-Hill, 1972).
- [31] W. Tschirk, Effect of transmission power control on the cochannel interference in cellular radio networks, *Elektrotechnik und Informationstechnik* 106(5) (1989) 194–196.
- [32] J. F. Whitehead, Signal-level-based dynamic power control for co-channel interference management, in: *Proceedings of the IEEE Vehicular Technology Conference* (1993) pp. 499–502.
- [33] V. Wong and C. Leung, A transmit power control scheme for improving performance in a mobile packet radio system, *IEEE Transactions on Vehicular Technology* 43(1) (February 1994) 174–180.
- [34] M. D. Yacoub, *Foundations of Mobile Radio Engineering* (CRC Press, Boca Raton, FL, 1993).
- [35] R. D. Yates, A framework for uplink power control in cellular radio systems, *IEEE Journal on Selected Areas in Communications* 13(7) (September 1995) 1341–1347.
- [36] R. D. Yates and C.-Y. Huang, Integrated power control and base station assignment, *IEEE Transactions on Vehicular Technology* 44(3) (August 1995) 638–644.
- [37] J. Zander, Distributed cochannel interference control in cellular radio systems, *IEEE Transactions on Vehicular Technology* 41(3) (August 1992) 305–311.
- [38] J. Zander, Performance of optimum transmitter power control in cellular radio systems, *IEEE Transactions on Vehicular Technology* 41(1) (February 1992) 57–62.
- [39] J. Zander, Transmitter power control for co-channel interference management in cellular radio systems, in: *Proceedings of the 4th WINLAB Workshop on 3rd Generation Wireless Information Networks* (1993).

John Rulnick earned the S.B. degree from the Massachusetts Institute of Technology in 1986 and M.S. and Ph.D. from the University of California, Los Angeles, in 1991 and 1996, respectively. He is currently an Assistant Professor in the Department of Electrical and Computer Engineering at Worcester Polytechnic Institute. His professional experience includes the design and implementation of various engineering systems and research consulting for industry, government agencies, and the RAND Corporation. His research interests include design, control, performance and economic analysis of networks and parallel processing systems, as well as applied probability and game theory. He is active in several engineering and mathematics professional societies, including ACM.

E-mail: rulnick@ece.wpi.edu

Nicholas Bambos earned the Ph.D. degree from the University of California, Berkeley, in 1989. He is currently an Associate Professor in the Department of Engineering Economic Systems and Operations Research at Stanford University. His research interests include communication networks, high-performance computing, operations research and optimization in communication networks, parallel and distributed processing systems and multiprocessor networks, queueing modeling and simulation, stochastic processes, network dynamics and optimization.

E-mail: bambos@leland.stanford.edu