Aggregate Interference Distribution From Large Wireless Networks With Correlated Shadowing: An Analytical–Numerical–Simulation Approach

Sebastian S. Szyszkowicz, *Student Member, IEEE*, Furkan Alaca, *Student Member, IEEE*, Halim Yanikomeroglu, *Member, IEEE*, and John S. Thompson, *Member, IEEE*

***Abstract*—As the number and variety of wireless devices sharing spectrum increases, it becomes increasingly important to characterize the sum interference that is produced by a large number of interferers. We show that, in the case of several interferers, the assumption of independent shadowing paths is very inaccurate and must be replaced by an appropriate correlation model. We choose one such model, which has desirable mathematical and physical properties, is tunable, and is particularly well suited for simulation, although our approach can also be used with other correlation models. In addition, we allow a very versatile channel and system model. The simulation cost of such systems quickly grows for large numbers of interferers due to the time and memory constraints of the Monte Carlo simulation algorithm using the classic matrix factorization (e.g., Cholesky) approach. We show how an alternative simulation approach using shadowing fields can significantly reduce the order of the computational cost. In addition, we show how judicious random sample reuse and extrapolation based on a numerical analysis of moments can be used to further simplify the simulation. Through the combination of these three approaches, using a mixture of simulation, numerical, and analytical techniques, we can obtain accurate approximations of the distribution of the total interference power while reducing computational time by factors of more than 1000. We can also make some mathematical statements about the problem, which may be useful for further developments. We argue that our model is complex enough to accommodate a good degree of realism and that our approach is a viable alternative to the pure analysis of such a complex and versatile problem.**

***Index Terms*—Algorithm optimization, cochannel interference, correlated shadowing.**

Manuscript received September 27, 2010; revised March 22, 2011; accepted April 26, 2011. Date of publication May 27, 2011; date of current version July 18, 2011. This work was supported in part by the National Sciences and Engineering Research Council of Canada through the Postgraduate Scholarship D. This work was presented in part at the IEEE Vehicular Technology Conference, Spring 2010, Taipei, under the title *Efficient Simulation Using Shadowing Fields of Many Wireless Interferers With Correlated Shadowing*. The review of this paper was coordinated by Prof. Y. Su.

S. S. Szyszkowicz, F. Alaca, and H. Yanikomeroglu are with the Department of Systems and Computer Engineering, Carleton University, Ottawa, ON K1S 5B6, Canada (e-mail: sz@sce.carleton.ca; falaca@sce.carleton.ca; halim@sce.carleton.ca).

J. S. Thompson is with the Institute for Digital Communications, School of Engineering, University of Edinburgh, EH8 9YL Edinburgh, U.K. (e-mail: John.Thompson@ed.ac.uk).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TVT.2011.2158012

I. Introduction

In RECENT YEARS, the simulation and analysis of interference that is caused by N distinct cochannel interferers, with N even in the hundreds [1]–[3], has received significant interest. The interference from a Poisson field of interferers with independent shadowing, which is a very similar problem, has been studied, at least, since 1992 (see [4]) and still receives attention [5]. The interfering nodes may be femtocells [1], sensor nodes, or any other devices that aggressively share spectrum, often in a noncoordinated and opportunistic manner. This type of scenario has become increasingly relevant as wireless communications move away from the traditional coordinated cellular model to more heterogeneous and distributed paradigms, e.g., ad-hoc networking and cognitive radio [2], [5]. As such, the study of interference from several cochannel interferers is essential for the design of future wireless systems. On the other hand, the total interference from a small number of interferers (usually in a cellular context with frequency reuse) has been well studied, often within the context of the sum-of-lognormals mathematical problem [6], [7]. This paper lies at the intersection of these fields, effectively bringing together correlated shadowing, large numbers of interferers, and arbitrary interferer layouts.

We will demonstrate that the independence assumption, although already known to give different results for small N [8], gives very inaccurate results for large N, indicating that shadowing correlation should be incorporated in simulation and analysis. Furthermore, we would argue that, given the general mathematical difficulty of such problems, developing analytical tools may not be worth the effort when the total interference may quickly be simulated through approaches such as Monte Carlo. This condition, however, does not hold true when the number of interferers is very large: Although the simulation time may then increase as quickly as $\mathcal{O}(N^3)$ (and the memory required as $\mathcal{O}(N^2)$), we will show that approximate numerical and analytical approaches may, in fact, significantly become simpler as $N \to \infty$.

In this paper, we will first demonstrate how the use of shadowing fields may reduce the time (and memory) cost to $\mathcal{O}(N)$. We will then examine how we can judiciously reuse and recombine random samples to obtain further gains in time. Finally, we will show how, through a combination of analytical and numerical techniques, we can extrapolate distributions for

very large N from distributions that are simulated at lower N, thus reducing the computational time to $\mathcal{O}(1)$. These three techniques combined reduce the computational time on a standard personal computer from 10.5 h (for $N = 1000$) to 16 s (for any $N \geq 500$).

In Section II, we describe a general physical model for wireless interference with several flexible parameters. In Section III, we argue why correlation in shadowing is indispensable for the analysis of interference from large networks and give several basic analytical results, many of which are later used in the extrapolation techniques. In Section IV, we first describe the standard simulation algorithm and then the shadowing fields, random sample reuse, and extrapolation—three techniques that are combined to make the simulations much more feasible for large N. We show the resulting simulated distributions and that they well correspond with the classical (slow) simulation approach. We conclude in Section V.

II. Problem Statement

The problem can fully and statistically be described by the statistical system and channel parameters and how the individual interference components are combined to form the aggregate interference power. Our model allows great flexibility in the choice of system and channel description.

A. System Model

Consider N interferers that are located at positions $\vec{r}_i$. We denote $r_i = \|\vec{r}_i\|$. Although other works [1], [2], [4], [5] often assume a particular interferer distribution, which is usually uniform over some area, we want to allow for as general a formulation as possible. We thus allow the locations $\vec{r}_i$ to be independent and identically distributed (i.i.d.) according to a valid probability density function (pdf) $g(\vec{r})$ in two dimensions such that, without loss of generality, the receiver is located at the origin. We define $r_{\min} > 0$ and $r_{\max}$ such that all positions always fall in between, i.e., $r_{\min} < r_i < r_{\max}$. Nevertheless, this modeling does not include the possibility of the clustering of the interferers (as done in [3]), but the model can be extended to include such an effect without too much difficulty.

We will assume that the interferers each transmit with random power T_i, *i.i.d.* according to some *pdf* $f_T(x)$, with only the conditions that $\mathbb{P}(T_i < 0) = 0$ and $\mathbb{VAR}\{T_i\} < \infty$. These variables may account for ON–OFF behavior, power control, and other factors.

B. Channel Model

The most significant channel model element in our problem is wireless shadowing. Given fixed propagation paths, the logarithms of the shadowing values may be modeled as jointly Gaussian random variables (RVs) [8]–[12]. Consider S_i the shadowing (in dB) that is experienced on path i. Although shadowing is often modeled as having approximately constant dB spread with distance, it can be argued that this is not the case for very short distances (see the model in [13]). We therefore consider the dB spread to be, in general, a function of distance $\sigma_{\rm s}(r)$, which is a positive continuous function.

We assume that the vector $\vec{S} = [S_i]_{i=1}^N$ is Gaussian when conditioned on $\vec{r}_1, \ldots, \vec{r}_N$, with

$$\begin{aligned} \mathbb{E}\{S_i\} &= 0 \\ \mathbb{E}\left\{S_i^2 | r_i\right\} &= \sigma_{\rm s}^2(r_i) \\ \mathbb{E}\{S_i S_j | \vec{r}_i, \vec{r}_j\} &= \sigma_{\rm s}(r_i)\sigma_{\rm s}(r_j) h(\vec{r}_i, \vec{r}_j) \end{aligned} \tag{1}$$

where h is the shadowing correlation model. We impose on h the condition that it is feasible [12] and that $h(\vec{r}_i, \vec{r}_j) \approx 1$ when $\vec{r}_i \approx \vec{r}_j$, specifically that

$$\forall\, \varepsilon > 0 \exists \delta > 0 : \|\vec{r}_i - \vec{r}_j\| < \delta \Rightarrow h(\vec{r}_i, \vec{r}_j) > 1 - \varepsilon. \tag{2}$$

Then, the correlation matrix of $\vec{S}$ conditioned on $\vec{r}_1, \ldots, \vec{r}_N$ is given by

$$\mathbf{K}_{N\times N} = \left[\sigma_{\rm s}(r_i)\sigma_{\rm s}(r_j) h(\vec{r}_i, \vec{r}_j)\right]. \tag{3}$$

Assuming that h is such that $\mathbf{K}$ is always a positive semidefinite (psd) matrix, it follows that the Gaussian vector $\vec{S}$ is always feasible [12], i.e., it can always be constructed, and is fully determined by (1).

Consider the angle of arrival separation

$$\theta = |\angle \vec{r}_i - \angle \vec{r}_j| \in [0^\circ, 180^\circ] \tag{4}$$

and the arrival distance ratio (in dB)

$$R = |10 \log_{10} r_i/r_j| = \frac{10}{\ln 10} |\ln r_i - \ln r_j|. \tag{5}$$

We choose a correlation model h that may be expressed in terms of θ and R only and is separable with respect to these dimensions:

$$h(r_i, r_j) = \max\{1 - \theta/\theta_0, 0\} \cdot \max\{1 - R/R_0, 0\} \tag{6}$$

with tunable parameters $0^\circ < \theta_0 \leq 180^\circ$ and $R_0 > 0$. This model is given in [12] and is based on [9].

We chose this model among several approaches for the following reasons.

1) In [12], we have shown that this model always yields *psd* correlation matrices $\mathbf{K}$. This is not the case for several of the existing models.
2) Furthermore, in [12], we have argued that, among all models that always give *psd* correlation matrices $\mathbf{K}$, this model seemed most physically realistic. In particular, we contrast this model with techniques that are expressible in the form $h(\vec{r}_i, \vec{r}_j) = f(\|\vec{r}_i - \vec{r}_j\|)$, notably with $f(x) = e^{-x/d_0}$ [14]. We have argued in [12] that these models are difficult to reconcile with the propagation arguments given in [9].
3) The selected model has two tunable parameters and can therefore approximate a wide range of correlation models with reasonable accuracy, as in [9].
4) The mathematical form of this model lends itself particularly well to quick simulation using shadowing fields, as we will demonstrate in Section IV-C.

The average path-loss is a deterministic function of distance, $p(r)$, which is a positive continuous nonincreasing function, with bounded variation over useful values of r. Path-loss is

often modeled as a power law $p(r) = r^{-\beta}$ for $2 \leq \beta \leq 6$. However, to allow for the incorporation of both very near and very far-away interferers, more flexibility in the model will be useful.

C. Aggregate Interference

It is usually assumed [8], [10], [11], [15] that the total interference power is the sum of the individual interference powers, as explained by incoherent signal addition. We are then interested in finding the statistics of the total interference power given by

$$I = \sum_{i=1}^{N} I_i, \qquad I_i = kp(r_i)e^{\lambda S_i} T_i \tag{7}$$

where $\lambda = 0.1 \ln 10$, and k is the common constant gain that accounts for multiplicative constants, e.g., antenna gains, reference distance, and transmit power. Without loss of generality, we set $k = 1$.

It is primarily the generation of the correlated S_i that will represent the cost of the simulation for high N and will therefore be our primary focus of study.

III. Toward a Numerical–Analytical Solution

Finding the distribution of I in closed form is a daunting task. In [10], we have shown that it may approximately be lognormal for large N and under some more restricted conditions. However, in general, it may be impossible to determine what shape the distribution of I takes or how we can find it in some manageable analytical or numerical fashion. Nevertheless, we may use analysis to make some useful observations about the distribution of I.

A. Exchangeable RVs

The theory of exchangeable RVs [16], [17] is a useful framework for thinking about the quantities I_i. RVs are said to be exchangeable if their joint distribution remains unchanged when their indices are permuted. The individual interferences I_i are, indeed, exchangeable, because the assignment of the indices is arbitrary. Although the set $\{S_i|\vec{r}_i\}$ is not exchangeable, because the RVs may have different variances and pairs may have different correlation coefficients, the set $\{S_i\}$ (as well as $\{I_i\}$) is exchangeable.

B. Behavior of Moments

In interference analysis, we are interested in studying the statistical behavior of the total interference power I. The natural approach is to first establish its mean and variance. We first evaluate the following quantities:

$$\begin{aligned} A &= \mathbb{E}\left\{p(r_1)e^{\lambda S_1}\right\} \\ B &= \mathbb{E}\left\{p^2(r_1)e^{2\lambda S_1}\right\} \\ C &= \mathbb{E}\left\{p(r_1)p(r_2)e^{\lambda(S_1+S_2)}\right\} \\ D &= \mathbb{E}\{T_1\} \\ E &= \mathbb{E}\left\{T_1^2\right\}. \end{aligned} \tag{8}$$

Then, because I_i's are exchangeable, we may easily formulate the moments as

$$\begin{aligned} \mathbb{E}\{I\} &= NAD \\ \mathbb{E}\{I^2\} &= NBE + (N^2 - N)CD^2 \\ \mathbb{V}\mathrm{AR}\{I\} &= N(BE - CD^2) + N^2(C - A^2)D^2. \end{aligned} \tag{9}$$

We observe that, while for independent shadowing we have $C = A^2 \Rightarrow \mathbb{V}\mathrm{AR}\{I\} = \mathcal{O}(N)$, in general for correlated shadowing $\mathbb{V}\mathrm{AR}\{I\} = \mathcal{O}(N^2)$. The mean power of I remains the same, regardless of the correlation. Thus, what was already observed for small N [8] will even more be significant for large N: adding correlation significantly changes (particularly, broadens) the distribution of I. It follows that, given a sufficiently realistic shadowing correlation model, the distribution of I that is obtained using correlated shadowing will be much more realistic than using independent shadowing.

In addition, because of the asymptotic behavior of the mean and variance, analyzing I as $N \to \infty$ requires the study of the convergence of I/N (rather than $(I - \mathbb{E}\{I\})/\sqrt{N}$ in the independent case). Because of the existence of correlation, the classical central limit theorem cannot be applied for large N. Indeed, I/N does not necessarily converge to a Gaussian distribution and may, in fact, converge to a distribution that is close to a lognormal with a significant spread [10].

We can therefore conclude that correlation in shadowing becomes a dominating factor in the distribution of I as N becomes large.

C. Numerical Evaluation of Moments

Although we may estimate the expectations of (8) through Monte Carlo simulation, we find it faster and more exact to evaluate them through numerical integration. By first conditioning on r_1, we have

$$A = \int_{r_{\min}}^{r_{\max}} \bar{g}(r)p(r)e^{\frac{1}{2}\lambda^2\sigma_{\mathrm{s}}^2(r)}dr \tag{10}$$

where $\bar{g}(r)$ is the *pdf* of r_i, which is evaluated from $g(\vec{r})$ as follows:

$$\bar{g}(r) = r\int_0^{2\pi} g(\vec{r})d\angle\vec{r}. \tag{11}$$

Similarly, we have

$$B = \int_{r_{\min}}^{r_{\max}} \bar{g}(r)p^2(r)e^{2\lambda^2\sigma_{\mathrm{s}}^2(r)}dr. \tag{12}$$

Finally, we have

$$\begin{aligned} C = &\iint_{\mathbb{R}^2}\iint_{\mathbb{R}^2} g(\vec{r_1})g(\vec{r_2})p(r_1)p(r_2)r_1r_2 \\ &\times e^{\lambda^2\left(\frac{1}{2}\sigma_{\mathrm{s}}^2(r_1)+\frac{1}{2}\sigma_{\mathrm{s}}^2(r_2)+\sigma_{\mathrm{s}}(r_1)\sigma_{\mathrm{s}}(r_2)h(\vec{r}_1,\vec{r}_2)\right)}d\vec{r}_1d\vec{r}_2. \end{aligned} \tag{13}$$

This integral needs to be evaluated in four dimensions for a general correlation model h and is not easy to separate in the case of (6).

With regard to D and E, they are simply obtained from $f_T(x)$.

D. Limiting Behavior and Mean Matching

Based on the exchangeability of I_i [17] and the variance behavior of I given by (9), we may conclude that I/N converges to a particular distribution, although we do not know what that distribution is. This condition leads us to formulate the following scaling approximation for interference from large networks under correlated shadowing:

$$\sum_{i=1}^{NM} I_i \overset{\mathcal{D}}{\approx} M \sum_{i=1}^{N} I_i, \qquad M \geq 1, \quad N \text{ large} \tag{14}$$

where $\overset{\mathcal{D}}{\approx}$ indicates that the two quantities have approximately the same cumulative distribution function (cdf). This case allows for the following two approaches for efficiently finding a good approximation of the distribution of I for high N.

1) We may simulate I for a smaller N and extrapolate its behavior for NM interferers by simply scaling the result by M, as shown in (14).
2) Based on the distribution of I/N as $N \to \infty$ (assuming that it can be found), the distribution of I for N large is well approximated as follows:

$$\sum_{i=1}^{N} I_i \overset{\mathcal{D}}{\approx} N \left(\lim_{N' \to \infty} \frac{1}{N'} \sum_{j=1}^{N'} I_j \right), \quad N \text{ large.} \tag{15}$$

Because the mean of I is simply proportional to N, this approach is equivalent to matching the mean between the available and the desired distributions.

E. Variance Matching

A different linear transformation may be obtained by matching not the mean but only the variance between the available and the desired distributions. In this case, we also have a finite and an infinite version as follows.

1) If we have the distribution of I for N and wish to have it for NM interferers, we may use (9) to obtain

$$\sum_{i=1}^{NM} I_i \overset{\mathcal{D}}{\approx} c \sum_{i=1}^{N} I_i, \qquad M \geq 1, \quad N \text{ large} \tag{16}$$

where

$$c = \sqrt{\frac{NM(BE - CD^2) + N^2M^2(C - A^2)D^2}{N(BE - CD^2) + N^2(C - A^2)D^2}}. \tag{17}$$

2) If we know the distribution of I/N as $N \to \infty$, we may also use variance matching as follows:

$$\mathbb{VAR}\{I/N\} \longrightarrow (C - A^2)D^2 \quad \text{as } N \to \infty. \tag{18}$$

We may use this to obtain

$$\sum_{i=1}^{N} I_i \overset{\mathcal{D}}{\approx} c' \left(\lim_{N' \to \infty} \frac{1}{N'} \sum_{j=1}^{N'} I_j \right), \quad N \text{ large} \tag{19}$$

where

$$c' = \sqrt{N^2 + N \frac{BE - CD^2}{(C - A^2)D^2}}. \tag{20}$$

Variance matching is not very different from mean matching, because they both multiply the data by a constant factor, and $c \approx M$ and $c' \approx N$ for high N. Finally, we may obtain a more refined approximation by matching the first two moments.

F. Moment Matching

Because the first two moments of I are always easily obtainable from (9), the scaling approach can be refined using both moments in the following two respective ways.

1) If we have the distribution of I for N and wish to have it for NM interferers, we may use (9) to obtain

$$\sum_{i=1}^{NM} I_i \overset{\mathcal{D}}{\approx} b + a \sum_{i=1}^{N} I_i, \qquad M \geq 1, \quad N \text{ large} \tag{21}$$

where

$$a = \sqrt{\frac{M\left(BE - CD^2 + NM(C - A^2)D^2\right)}{BE - CD^2 + N(C - A^2)D^2}}$$
$$b = NAD(M - a). \tag{22}$$

2) If we know the distribution of I/N as $N \to \infty$, we may also use moment matching. Indeed, the limiting moments are

$$\begin{aligned} \mathbb{E}\{I/N\} &\longrightarrow AD \\ \mathbb{E}\left\{(I/N)^2\right\} &\longrightarrow CD^2 \\ \mathbb{VAR}\{I/N\} &\longrightarrow (C - A^2)D^2 \quad \text{as } N \to \infty. \end{aligned} \tag{23}$$

We may use this technique to obtain

$$\sum_{i=1}^{N} I_i \overset{\mathcal{D}}{\approx} b' + a' \left(\lim_{N' \to \infty} \frac{1}{N'} \sum_{j=1}^{N'} I_j \right), \quad N \text{ large} \tag{24}$$

where

$$\begin{aligned} a' &= c' \text{ in (20)} \\ b' &= AD(N - a'). \end{aligned} \tag{25}$$

Again, this matching is not very different from the previous approaches, because for large N, we have a' approaching N and b' approaching zero. One important difference is that adding the b' term does not allow the approximating distribution to take values in the neighborhood of zero, which may cause high relative error in the lower tail.

These three types of moment-matching each offer the following two possibilities of approximating the *cdf* of I for large N: 1) to approximate larger simulations from smaller simulations and 2) to approximate large simulations from the limiting distribution of I/N (however, this second approach is not readily available to us, because we do not know how we can efficiently find this limiting distribution). We will explore this approach in detail in Section IV-G.

IV. Optimizing the Simulation of Correlated Shadowing

To simulate I (and, in fact, $\{I_i\}$), the following two very different approaches can, nevertheless, give very close results: 1) matrix factorization and 2) shadowing fields. A simulation can fully be specified by the parameters listed in Table I.

A. Standard Approach: Matrix Factorization

Generating $\vec{S}$ in Monte Carlo simulations is traditionally [9] done by solving for $\mathbf{C}_{N\times N}$ in

$$\mathbf{K} = \mathbf{C}^T\mathbf{C} \tag{26}$$

for each particular realization of $\mathbf{K}$. We will write $\mathbf{C} = \sqrt[*]{\mathbf{K}}$. The next step is to generate a vector $\vec{Z} = [Z_i]_{i=1}^N$ of independent standard Gaussian $\mathcal{N}(0,1)$ RVs. Then, $\vec{S}$ is obtained from

$$\vec{S} = \vec{Z}\mathbf{C}. \tag{27}$$

This approach is implemented as follows:

Algorithm That Uses Matrix Factorization

Ensure: The histogram of $I[k]$ approximates the *pdf* of I.

for $k = 1$ to K **do**
  for $i = 1$ to N **do**
    $\vec{r}_i \Leftarrow i.i.d.$ random from $g(\vec{r})$
    $Z_i \Leftarrow i.i.d.$ random $\mathcal{N}(0,1)$
    $T_i \Leftarrow i.i.d.$ random from $f_T(x)$
  end for
  for $i = 1$ to N **do**
    $\bar{K}[i,i] \Leftarrow 1$
    for $j = 1$ to $i-1$ **do**
      $\bar{K}[i,j] = \bar{K}[j,i] \Leftarrow h(\vec{r}_i, \vec{r}_j)$
    end for
  end for
  $\bar{\mathbf{C}} \Leftarrow \sqrt[*]{\mathbf{K}}$
  $\vec{S} \Leftarrow \mathrm{diag}(\sigma_{\mathrm{s}}(\vec{r}_1), \ldots, \sigma_{\mathrm{s}}(\vec{r}_N)) \cdot \vec{Z} \cdot \bar{\mathbf{C}}$
  $I[k] \Leftarrow \sum_{i=1}^N p(r_i) e^{\lambda S_i} T_i$
end for

Solving $\sqrt[*]{\mathbf{K}}$ can efficiently be performed by the Cholesky factorization[1] with complexity $\mathcal{O}(N^3)$ [10], [18].

[1] Although the Cholesky factorization fails for singular matrices $\mathbf{K}$ [12], [18], we have observed in the simulations here and in [10] that this event is extremely rare when using the model (6) with double-precision arithmetic.

TABLE I
Simulation Parameters

Parameter	Description
K	Total number of Monte Carlo trials of I
N	Number of interferers
$K_{\vec{r}}$	Number of position realisations of the N interferers
K_{Ch}	Number of channel realisations
Physical	
$g(\vec{r})$	Statistical density of interferer positions
$r_{\min}, r_{\max}$	Limits on interferer distances from receiver: $g(\vec{r}) = 0\ \forall \vec{r} : \|\vec{r}\| \notin [r_{\min}, r_{\max}]$
$p(r)$	Average pathloss at distance r
$\sigma_{\mathrm{s}}(r)$	Shadowing spread in dB at distance r
$f_T(x)$	*pdf* of interferer power T_i
θ_0, R_0	Parameters of shadowing correlation model (6)
Shadow Fields	
D_Θ	Number of discrete field points in θ dimension
D_{R}	Number of discrete field points in R dimension
F_Θ	Length of digital FIR filter in θ dimension
F_{R}	Length of digital FIR filter in R dimension
$\mathbf{F}$	$F_\Theta \times F_{\mathrm{R}}$ filtering kernel

B. Shadowing Fields

A shadowing field is a Gaussian random process in two dimensions (properly, a random field) with a specific autocorrelation function. This autocorrelation is such that, when interferers with positions $\vec{r}_i$ are placed on the field and the value of the field at the point $\vec{r}_i$ is taken as the value of S_i, $\vec{S}|\vec{r}_1, \ldots, \vec{r}_N$ has the desired correlation matrix $\mathbf{K}$. This case can be compared, by analogy, to gravitational (or electric) fields, where the field gives the acceleration of a mass that is placed at any point, whether or not there is, in fact, a mass at that point.

The idea of generating shadowing fields has been explored [11], [19]–[21] with correlation functions of the form $h(\vec{r}_i, \vec{r}_j) = f(\|\vec{r}_i - \vec{r}_j\|)$. We have argued in [12] that such models may not reflect true shadowing spatial correlation characteristics. Furthermore, correlation as a separable function of θ and R can easily be simulated using a geometric transformation. The accuracy of our method is limited only by the quantization level.

Consider a random field (a 2-D random process) $\bar{\mathbf{M}}$ of continuous parameters (x, y). Let the random field be stationary [22], with an autocorrelation function such that the correlation between the field at two points $\bar{\mathbf{M}}(x_i, y_i)$ and $\bar{\mathbf{M}}(x_j, y_j)$ corresponds to the desired shadowing correlation under some transformation.

Consider what we will call the log-polar[2] transformation

$$\begin{aligned}\mathcal{T}_{\mathrm{LP}} &: (\theta, R) \longmapsto 10^{0.1R}(\cos\theta, \sin\theta)\\ \mathcal{T}_{\mathrm{LP}} &: [0, 2\pi] \times [10\log_{10} r_{\min}, 10\log_{10} r_{\max}]\\ &\longmapsto \{\vec{r} : r_{\min} \le r \le r_{\max}\}.\end{aligned} \tag{28}$$

[2] A polar representation of certain shadowing fields is suggested in [19].

Let us choose the autocorrelation of $\bar{\mathbf{M}}$ as

$$\eta_x(\xi) = \begin{cases} 1 - |\xi|/\theta_0, & |\xi| \le \theta_0 \\ 0, & \theta_0 \le |\xi| \le 2\pi - \theta_0 \\ 1 + (|\xi| - 2\pi)/\theta_0, & 2\pi - \theta_0 \le |\xi| \le 2\pi \end{cases}$$

$$\eta_y(\upsilon) = \begin{cases} 1 - |\upsilon|/R_0, & |\upsilon| \le R_0 \\ 0, & |\upsilon| > R_0 \end{cases}$$

$$\mathbb{E}\left\{\bar{\mathbf{M}}(x+\xi, y+\upsilon)\bar{\mathbf{M}}(x,y)\right\} = \eta_x(\xi)\eta_y(\upsilon) \tag{29}$$

for $\bar{\mathbf{M}}$ defined on $[0, 2\pi] \times [10\log_{10} r_{\min}, 10\log_{10} r_{\max}]$. We find that the field $\bar{\mathbf{M}}$, under transformation $\mathcal{T}_{\mathrm{LP}}$, has the correlation properties of (6), i.e.,

$$\mathcal{T}_{\mathrm{LP}} : (x_i, y_i) \longmapsto \vec{r}_i, i = 1, 2, \ldots$$
$$\mathbb{E}\left\{\bar{\mathbf{M}}(x_1, y_1)\bar{\mathbf{M}}(x_2, y_2)\right\} = h(\vec{r}_1, \vec{r}_2). \tag{30}$$

Therefore, we may write

$$S_i = \sigma_{\mathrm{s}}(r_i)\bar{\mathbf{M}}\left(\mathcal{T}_{\mathrm{LP}}^{-1}(\vec{r}_i)\right) \tag{31}$$

and S_i's will have the same correlation matrix as in (27).

For numerical purposes, $\bar{\mathbf{M}}$ can be approximated by a discrete-parameter matrix $\mathbf{M}_{D_\Theta \times D_{\mathrm{R}}}$, with a regularly spaced quantization grid along θ and R. The correlation of the form (6), which is triangular in both dimensions, can be obtained by using a 2-D finite impulse response (FIR) uniform square filter $\mathbf{F}_{F_\Theta \times F_{\mathrm{R}}}$, ideally choosing F_Θ and F_{R} so that we exactly have

$$F_\Theta/D_\Theta = \theta_0/2\pi$$
$$F_{\mathrm{R}}/D_{\mathrm{R}} = R_0/10\log_{10}(r_{\max}/r_{\min}) \tag{32}$$

with $\mathbf{F}$ equal everywhere to $1/\sqrt{F_\Theta F_{\mathrm{R}}}$.

To obtain the value of the discretized field $\mathbf{M}$ at some point, we must round the coordinates $\mathcal{T}_{\mathrm{LP}}^{-1}(\vec{r}_i)$ to the nearest quantization point. The algorithm is thus limited in precision by the finite spatial quantization. On the other hand, the computational cost of generating one field grows $\mathcal{O}(F_\Theta F_{\mathrm{R}} D_\Theta D_{\mathrm{R}}) = \mathcal{O}(D_\Theta^2 D_{\mathrm{R}}^2)$, and therefore, it is critical to choose the number of quantization points $D_\Theta D_{\mathrm{R}}$ to balance precision and computational time. Fig. 1 shows realizations of the same shadowing field at different resolutions.

Shadowing fields have the additional advantages of requiring only $\mathcal{O}(N)$ memory (rather than $\mathcal{O}(N^2)$ for matrix factorization) and easily accommodating interferer mobility; indeed, although matrix factorization only gives shadowing values at the N specified locations, shadowing fields give the value of the (potential) shadowing *everywhere*. This condition is also useful when N is uncertain, random, or variable.

C. Efficient Filtering for Separable Triangular Correlations

The choice of the correlation model (6) is particularly fortunate from the point of view of computational efficiency, because it benefits from the following two properties: 1) separability and 2) a triangular shape.

1) Separability: The nature of the correlation model (6) is such that it can be expressed as the product of a function of θ and a function of R. It follows [22] that the resulting 2-D process in the $\theta - R$ plane is also separable. It can therefore be simulated by separately filtering over each dimension, which reduces the general filtering cost from $\mathcal{O}(F_\Theta F_{\mathrm{R}} D_\Theta D_{\mathrm{R}})$ (e.g., as in [20]) to $\mathcal{O}((F_\Theta + F_{\mathrm{R}}) D_\Theta D_{\mathrm{R}})$.

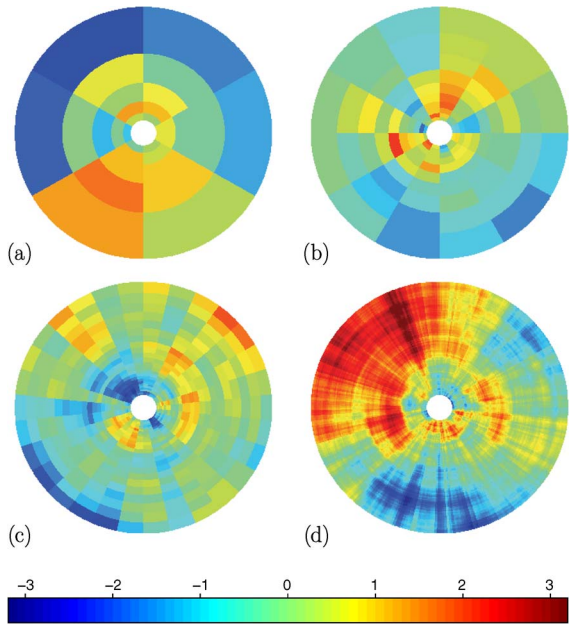

Fig. 1. Shadowing field realizations for $r_{\min} = 50$, $r_{\max} = 500$, $\theta_0 = 60°$, and $R_0 = 6$ dB with increasing resolution: $D_\Theta = 6n$, $D_{\mathrm{R}} = 5n$, $F_\Theta = n$, and $F_{\mathrm{R}} = 3n$, where (a) $n = 1$, (b) $n = 2$, (c) $n = 5$, and (d) $n = 50$. The color of the areas corresponds to the value of $S_i/\sigma_{\mathrm{s}}(\vec{r}_i)$ therein.

2) Optimized Box Filters: The triangular form in θ and R of the correlation expression in (6) requires the use of rectangular (box) filters applied to a white Gaussian process. Computationally, this approach is very efficient, because the filtering requires no multiplications. In addition, it can even more efficiently be implemented [23], with the number of additions now approximately $2D_\Theta$ and $2D_{\mathrm{R}}$ in each respective dimension, rather than $(F_\Theta - 1)D_\Theta$ and $(F_{\mathrm{R}} - 1)D_{\mathrm{R}}$, respectively. This condition is because adjacent outputs of a box filter differ only by two input values. All this makes the total computation cost for one field realization $\mathcal{O}(D_\Theta D_{\mathrm{R}})$, which is independent of the filter size and, hence, of the correlation distances θ_0 and R_0.

3) Optimized Shadowing Fields Algorithm: A general shadowing fields algorithm will have the following form.

BASIC SHADOWING FIELDS ALGORITHM

Ensure: The histogram of $I[k]$ approximates the *pdf* of I.
  for $k = 1$ to K **do**
    for $i = 1$ to N **do**
      $\vec{r}_i \Leftarrow i.i.d.$ random from $g(\vec{r})$
      $T_i \Leftarrow i.i.d.$ random from $f_T(x)$
    end for
    $\mathbf{M} \Leftarrow$ shadowing field realization†
    for $i = 1$ to N **do**
      $S_i \Leftarrow \sigma_{\mathrm{s}}(r_i)\mathbf{M}[\mathcal{T}_{\mathrm{LP}}^{-1}(\vec{r}_i)$ (nearest index)
    end for
    $I[k] \Leftarrow \sum_{i=1}^{N} p(r_i) e^{\lambda S_i} T_i$
  end for

The correlation model (6) will benefit from separability in the R and θ domains, as well as from the triangular form of both autocorrelations. It is also important to remember to wrap the field in the θ direction to ensure circular continuity. These elements produce the following detailed implementation.

---

†FAST SHADOWING FIELDS GENERATION

Ensure: $\mathbf{M}$ is Gaussian and approximately correlated according to (6).

$\mathbf{Z}_{D_\Theta \times (D_\mathrm{R}+F_\mathrm{R}-1)} \Leftarrow i.i.d.$ random $\mathcal{N}(0,1)$
Initialize a temporary matrix $\mathbf{W}_{D_\Theta \times (D_\mathrm{R}+F_\mathrm{R}-1)}$.
for $m = 1$ to $D_\mathrm{R} + F_\mathrm{R} - 1$ **do**
  $\mathbf{W}[1,m] \Leftarrow \sum_{n=1}^{F_\Theta} \mathbf{Z}[n,m]$
end for
for $n = 1$ to $D_\Theta - 1$ **do**
  $n^* \Leftarrow (n + F_\Theta - 1) \bmod D_\Theta + 1$
  for $m = 1$ to $D_\mathrm{R} + F_\mathrm{R} - 1$ **do**
  $\mathbf{W}[n+1,m] \Leftarrow \mathbf{W}[n,m] - \mathbf{Z}[n,m] + \mathbf{Z}[n^*,m]$
end for
end for
for $n = 1$ to D_Θ **do**
  $\mathbf{M}[n,1] \Leftarrow \sum_{m=1}^{F_\mathrm{R}} \mathbf{W}[n,m]$
end for
for $m = 1$ to $D_\mathrm{R} - 1$ **do**
  for $n = 1$ to D_Θ **do**
  $\mathbf{M}[n,m+1] \Leftarrow \mathbf{M}[n,m] - \mathbf{W}[n,m] + \mathbf{W}[n,m+F_\mathrm{R}]$
end for
end for

D. Reusing Random Samples

Converting the simulation algorithm from matrix factorization to shadowing fields can bring a great gain in simulation time for large N. We may, however, augment these gains by using an additional technique, based on reusing random samples, that can be applied to both algorithms to greatly decrease their computational time at a negligible cost in accuracy.

We begin with the observation that both algorithms, although quite different, are also fundamentally similar; both algorithms require generating NK positions from $g(\vec{r})$, as well as generating a large set of independent Gaussian RVs. These RVs are then linearly combined to obtain the required correlation structure, after which, we may compute I_i and I according to (7). We further observe that the generation of the independent Gaussian RVs and the interferer positions are two separate tasks.

This case leads to the idea of generating fewer random quantities of both kinds and pairing them in different combinations to achieve a similar amount of randomness as the case when all realizations are different. We will call $K_{\vec{r}}$ the number of times that the positions of the interferers are generated for a total sample size of $NK_{\vec{r}}$ from $g(\vec{r})$. We will also call K_Ch the total number of channel realizations generated, although this approach will have a somewhat different meaning in both algorithms. We impose the following conditions:

$$K/K_{\vec{r}} \in \mathbb{N}, \quad K/K_\mathrm{Ch} \in \mathbb{N}, \quad K_{\vec{r}}K_\mathrm{Ch}/K \in \mathbb{N}.$$

One important implementation consideration is whether to first generate the positions or the Gaussian RVs, as one of these two will need to be stored in memory during the simulation to allow for reuse. Let us examine separately the details of implementing reuse for each algorithm.

1) Reuse in Matrix Factorization: Matrix factorization requires the generation and factorization of $K_{\vec{r}}N \times N$ matrices but the generation of only $N \times K_\mathrm{Ch}$ independent Gaussian RVs. It is therefore more judicious, for large N, to first generate the Gaussian RVs and store them for reuse and then to generate and discard the interferer positions and the corresponding correlation matrices one by one. This approach can be implemented as follows.

MATRIX FACTORIZATION ALGORITHM WITH REUSE

Ensure: The histogram of $I[k]$ approximates the *pdf* of I.

for $k = 1$ to K_Ch **do**
  for $n = 1$ to N **do**
  $Z_n[k] \Leftarrow i.i.d.$ random $\mathcal{N}(0,1)$
  end for
end for
for $k = 1$ to $K_{\vec{r}}$ **do**
  for $i = 1$ to N **do**
  $\vec{r}_i \Leftarrow i.i.d.$ random from $g(\vec{r})$
  end for
  for $i = 1$ to N **do**
  $\bar{K}[i,i] \Leftarrow 1$
  for $j = 1$ to $i - 1$ **do**
  $\bar{K}[i,j] = \bar{K}[j,i] \Leftarrow h(\vec{r}_i, \vec{r}_j)$
  end for
  end for
  $\bar{\mathbf{C}} \Leftarrow \sqrt[*]{\bar{\mathbf{K}}}$
  for $l = 1$ to $K/K_{\vec{r}}$ **do**
  $k^* \Leftarrow l + (k-1)K/K_{\vec{r}}$
  $l^* \Leftarrow (k^* - 1) \bmod K_\mathrm{Ch} + 1$
  $\vec{S} \Leftarrow \mathrm{diag}(\sigma_\mathrm{s}(\vec{r}_1), \ldots, \sigma_\mathrm{s}(\vec{r}_N)) \cdot \vec{Z}[l^*] \cdot \bar{\mathbf{C}}$
  for $i = 1$ to N **do**
  $T_i \Leftarrow i.i.d.$ random from $f_T(x)$
  end for
  $I[k^*] \Leftarrow \sum_{i=1}^{N} p(r_i) e^{\lambda S_i} T_i$
  end for
end for

Here, the calculation of l^* establishes the way in which reused random values are assigned to each other. Of course, there are several equivalent such mappings, because all the random values are *i.i.d.* and, thus, exchangeable. Note that we have not used reuse for the generation of T_i. In fact, we could do the reuse of T_i in the following three ways: 1) associating them with the values of $\vec{r}_i$; 2) associating them with the values of Z; or 3) dissociating them from both and having a more complicated mapping. For simplicity, we have not done this approach here, and all our simulations, except in Fig. 10, set $T_i = 1$.

2) Reuse in Shadowing Fields: Applying reuse in shadowing fields requires a reversed approach; because the shadowing fields algorithm is asymptotically faster than matrix

factorization, it is better to first generate the channel realizations (i.e., the fields), whose memory cost is $D_\Theta D_\mathrm{R} K_\mathrm{Ch}$ and does not depend on N. One possible implementation is given as follows.

BASIC SHADOWING FIELDS ALGORITHM WITH REUSE

Ensure: The histogram of $I[k]$ approximates the *pdf* of I.
for $k = 1$ to K_Ch **do**
  $\mathbf{M}[k] \Leftarrow$ shadowing field realization[†]
end for
for $k = 1$ to $K_{\vec{r}}$ **do**
  for $i = 1$ to N **do**
    $\vec{r}_i \Leftarrow i.i.d.$ random from $g(\vec{r})$
    $[x_i, y_i] \Leftarrow$ discrete coordinates in the shadowing field that corresponds to position $\vec{r}_i$
  end for
  for $l = 1$ to $K/K_{\vec{r}}$ **do**
    $k^* \Leftarrow l + (k-1)K/K_\mathrm{Ch}$
    $l^* \Leftarrow (k^* - 1) \bmod K_{\vec{r}} + 1$
  for $i = 1$ to N **do**
    $T_i \Leftarrow i.i.d.$ random from $f_T(x)$
    $S_i \Leftarrow \mathbf{M}_{[x_{l^*}, y_{l^*}]}[l^*]$
  end for
  $I[k^*] \Leftarrow \sum_{i=1}^{N} p(r_i) e^{\lambda S_i} T_i$
 end for
end for

Again, reuse has not been applied to T_i, which is independent of all other quantities. Reuse for T_i can be implemented in the same manner as discussed for matrix factorization.

E. Simulator Calibration

To apply shadowing fields and random sample reuse to a simulation problem, decisions need to be made with regard to some parameters that did not exist in the initial problem statement but are needed to make use of the approximating methods. We call this process calibration and require that the simulation parameters be such that the resulting approximating distributions do not differ in most places by more than 1 dB from the simulations that use matrix factorization and no reuse, which are considered exact. Although increasing these parameters will give more accurate results, simulation time will suffer; thus, a compromise must be struck.

The first step is to choose a suitable resolution for the shadowing fields. For simplicity, we will jointly and proportionally change D_Θ and D_R and allow only values that give integer values of $D_\Theta \theta_0 / 2\pi$ and $D_\mathrm{R} R_0 / 10 \log_{10}(r_\mathrm{max}/r_\mathrm{min})$. In the case of the parameters chosen in Table II, we may set $D_\Theta = 6n$, $D_\mathrm{R} = 5n$, which gives $F_\Theta = n$, $F_\mathrm{R} = 3n$, and $n \in \mathbb{N}^*$.

Fig. 1 shows the effect of increasing n on the appearance of realizations of shadowing fields, showing greater detail with increasing n, but at a cost of computational time $\mathcal{O}(n^2)$. In Fig. 2, we observe the effect of the resolution on the *cdf*[3] of I simulated using shadowing fields, with $n = 1, 2, 3$. First, we observe that, for $N = 1$, there is no significant difference between any of the simulations, because there is, in fact, no correlation, and we have merely shown that all algorithms produce the same marginal distribution for I_i. For $N = 10$, we observe some significant distortion in the lower tail, which seems little affected by changing the field resolution. For higher N, however, we clearly see an improvement with increasing resolution: although $n = 1$ gives rather poor results in both tails, $n = 2$ or 3 gives very accurate results across the whole range of values. Thus, we consider $n = 2$ sufficient at this point, which leads to the simulation parameters listed in Table II.

The next step is to choose the amount of reuse for both channel and interferer position realizations. We set $K_\mathrm{Ch} = K_{\vec{r}}$ for simplicity, and we can subsequently define a reuse factor $m = K/K_{\vec{r}} = K/K_\mathrm{Ch}$. We now examine the effect of the reuse factor on the distortion of the simulated distribution of I. Figs. 3 and 4 show the effect of reuse on matrix factorization and shadowing fields, respectively. We first observe that reusing random samples causes distortion in the far tails. This condition can be predicted from the fact that far tails are associated with

TABLE II
SIMULATION SETTINGS

Parameter	Value
K	1 000 000
$K_{\vec{r}}$	10 000
K_Ch	10 000
$g(\vec{r})$	uniform over an annular region with radii:
r_min	50
r_max	500
$p(r)$	$r^{-2}(1 + r/150)^{-2}$ (based on [24])
$\sigma_\mathrm{s}(r)$	$10(1 - \exp(-3r/200))$ dB (taken from [13])
θ_0	60°
R_0	6 dB
D_Θ	12
D_R	10
F_Θ	2
F_R	6

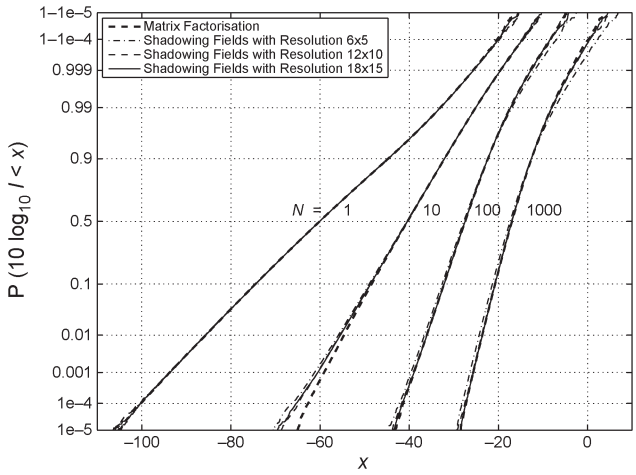

Fig. 2. Calibration of field resolution parameters $D_\Theta = 6n$ and $D_\mathrm{R} = 5n$.

[3]It is convenient to plot the *cdf* of I on lognormal paper [15], which we will do throughout this paper.

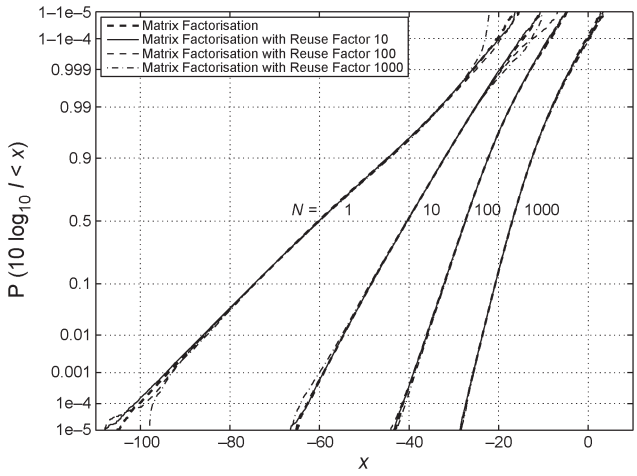

Fig. 3. Effect of random sample reuse on the matrix factorization algorithm, where $K = 1\,000\,000$, and $K_{\vec{r}} = K_{\mathrm{Ch}} = K/m$.

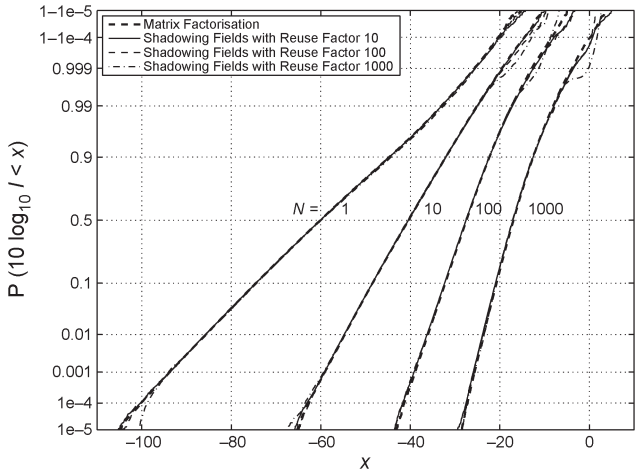

Fig. 4. Effect of combining random sample reuse with the calibrated $(D_\Theta = 12, D_\mathrm{R} = 10)$ shadowing fields algorithm. $K = 1\,000\,000$, and $K_{\vec{r}} = K_{\mathrm{Ch}} = K/m$.

rare events, the occurrence of which is significantly affected by repetition of random values. We also observe that the maximum possible reuse of $m = 1000$ causes too much distortion in certain cases, but $m = 100$ gives very reasonable accuracy. We thus choose $K_{\mathrm{Ch}} = K_{\vec{r}} = 10\,000$.

We also observe that the performance with reuse is usually worse for smaller $N(= 10)$ and improves with increasing N. This case can be interpreted from the fact that, given a particular $K_{\vec{r}}$ and K_{Ch}, the number of total independent RVs generated by the simulator increases with N, and thus, there is more "randomness" in the system, which leads to more accurate distributions. This trend is encouraging, because it confirms that our approach is well-suited for large N.

F. Time Performance Comparison

Once we have calibrated the simulator to produce accurate results, we perform a comparison of the computation time against N required for a simulation using the Cholesky factorization versus using shadowing fields. In addition, we show the performance of both algorithms with reuse of random samples

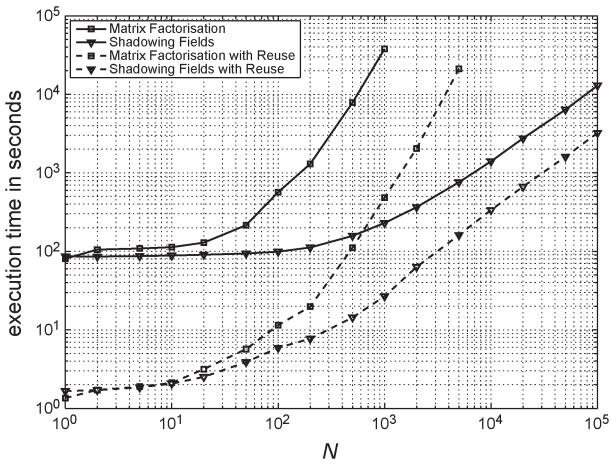

Fig. 5. Execution time performance of the matrix (Cholesky) factorization versus shadowing fields and of no reuse versus random sample reuse. Simulation parameters are given in Table II.

(each sample is reused $K/K_{\vec{r}} = K/K_{\mathrm{Ch}} = 100$ times). The performances of all four algorithms are shown in Fig. 5.

The first observation is that shadow fields always outperform matrix factorization, except for $N = 1$, where there is no correlation, and generating shadow fields is redundant. Furthermore, we observe that, for N beyond about 20, the time that is required for the Cholesky factorization drastically increases and takes on a sharper trend than that for shadowing fields, confirming that shadowing fields will always eventually outperform matrix factorization. Indeed, as can be observed from the algorithm, shadowing fields require an initial investment to generate the field realizations, after which the computational cost is linear in N, which is confirmed by the asymptotic behavior of both shadowing fields curves. On the other hand, both Cholesky factorization curves show an asymptotic trend between $\mathcal{O}(N^2)$ and $\mathcal{O}(N^3)$. The fact that the observed growth is less than $\mathcal{O}(N^3)$ can be explained in part by the heavy computational cost that is associated with constructing the correlation matrices (cost of $\mathcal{O}(N^2)$) and in part by the fact that the correlation matrix is relatively sparse.

Turning to the use of random sample reuse, we observe the most significant gain for matrix factorization. Indeed, while using $K/K_{\vec{r}} = K/K_{\mathrm{Ch}} = 100$ fewer random values, we observe a time gain of a factor of 78 at $N = 1000$. This case is because both generating the correlation matrix and factorizing it, the two most costly computations, are now done 100 fewer times, giving nearly the same gain in time. On the other hand, the shadowing fields approach presents more modest gains; although, for small N, the time gain is of a factor of 50, it asymptotically tapers off to only a factor of 4. The bottleneck in this case is extracting the values from the shadowing fields realizations based on the interferer positions. This operation, by its very nature, does not benefit from reuse.

Although the exact performance of each algorithm is dependent on software implementation and the hardware platform, the general trends and conclusions hold, because the results are consistent with the predicted complexity of each algorithm.

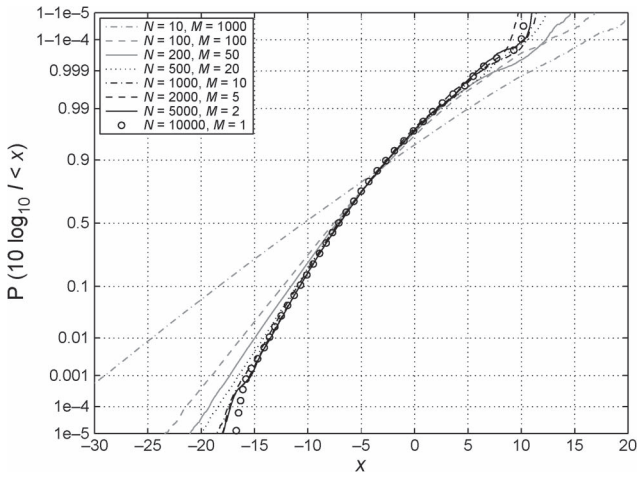

Fig. 6. Using mean matching (14) to extrapolate the distribution of I for very high N. All simulations are done using shadowing fields with reuse.

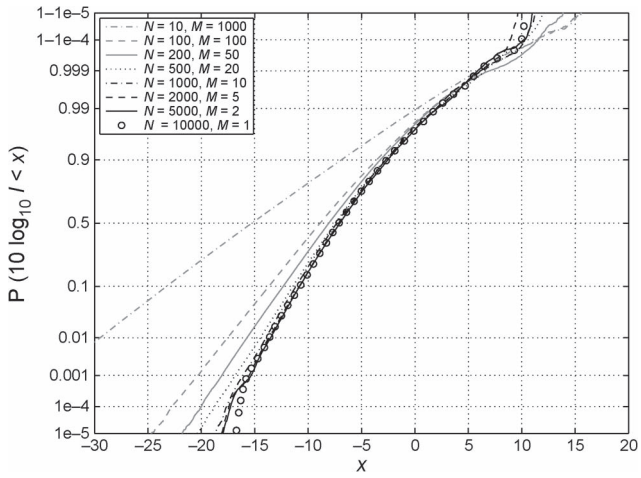

Fig. 7. Using variance matching (16) to extrapolate the distribution of I for very high N. All simulations are done using shadowing fields with reuse.

The hardware and software specifications of the simulation platform are given as follows:

- Intel® Core™ i7 860, 2.8 GHz, eight logical cores (four physical cores with hyper-threading);
- 8-GB random access memory (RAM);
- Microsoft® Windows® 7 Professional, 64 bit;
- The MathWorks™ MATLAB® version 7.9.0.529 (R2009b).

We extensively worked to write time-optimized MATLAB code for all algorithms, so as to have a good idea of the true computational costs of each algorithm. We observed that MATLAB automatically uses several of the logical cores in parallel to execute large repetitive tasks. The repetitive nature of the algorithms lends itself particularly well to hardware parallelization.

G. Moment-Corrected Extrapolations for High N

As we have seen in Section III, under proper normalization, the distribution of I converges as N grows. It does not matter to what exact distribution this distribution converges, only that convergence occurs. In fact, due to the complexity of the problem, it is not even certain that the limit distribution has any particular closed form (although we showed in [10] that approximate convergence to a lognormal distribution occurs in some cases).

This convergence can be exploited in accelerating simulations. Indeed, if we want the distribution of I for NM interferers such that N, $NM \in \mathbb{N}$, N large, and $M > 1$, we may merely simulate I with N interferers and then use the knowledge of the moments of I as a function of N to correct the simulated distribution to match the desired distribution. Because of convergence, the shape (in the linear domain) of the distribution does not change when N is very high; only the scale and offset parameters do.

We have given three formulas for extrapolating distributions. Each formula comes in the following two versions: 1) a finite version, where the distribution for NM interferers is obtained from the distribution for N interferers, and 2) an infinite version, where the distribution for N interferers is obtained from the distribution of I/N as $N \to \infty$. Because we do not have an efficient method for obtaining the limiting distribution as $N \to \infty$ (although we hope that such a method will be developed in the future), we will only use the first, i.e., finite, approach.

The first approach, given in Section III-D, is based on simply matching the mean of I, which is exactly proportional to N. This method is the simplest, does not require the numerical computation of A, B, and C, and requires only the multiplication of the simulated values of I by M, equivalently shifting the distribution to the right by $10 \log_{10} M$ dB on lognormal paper. Fig. 6 shows the distribution of I for $N = 10\,000$ (circles), as well as various simulations for lower N, with the appropriate correction factor $M = 10\,000/N$.

The second approach, given in Section III-E, is based, instead, on matching only the variance of I. This method requires the numerical computation of A, B, and C to obtain c in (17). It is similar to the first approach, because only a multiplying factor is applied to the simulated values of I, equivalently shifting the distribution to the right by $10 \log_{10} c$ dB on lognormal paper. Fig. 7 is made along the same lines as in Fig. 6 but with variance matching instead.

The third approach, given in Section III-F, is based on matching both the mean and the variance of I. This method again requires the numerical computation of A, B, and C to obtain a and b in (22). This last method is different from the first two methods, because it proposes not merely a linear but an affine transformation on the values of I, which leads to distortion, particularly in the lower tail, when plotted on lognormal paper. Fig. 8 is made along the same lines as in the previous two figures but matching both the mean and the variance (equivalently, the first two moments).

The second and third methods require the numerical evaluations of the integrals in (10), (12), and (13). The first two integrals are 1-D and are performed using basic Riemann integration with 100 000 uniform points, with negligible computational time. The integral in (13) is 4-D and is performed again using Riemann integration with 50^4 uniform points, with a computational time of approximately 1.25 s. The values that

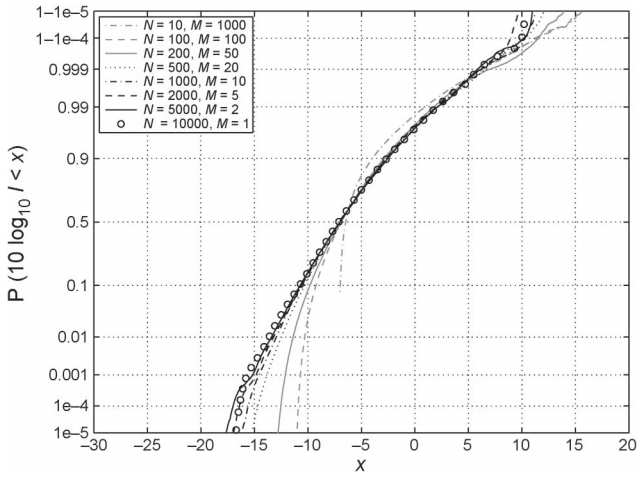

Fig. 8. Using two-moment matching (21) to extrapolate the distribution of I for very high N. All simulations are done using shadowing fields with reuse.

are obtained for the scenario described in Table II and used in Figs. 7 and 8 are

$$A = 2.972 \cdot 10^{-5}, \quad B = 1.256 \cdot 10^{-7}, \quad C = 2.342 \cdot 10^{-9},$$

and $D = E = 1$, because we do not consider randomness for T_i in these simulations.

For all three methods, we observe that good convergence (within 1 dB) can be obtained after about $N \geq 500$. Looking back at Fig. 5, we observe that a simulation with $N = 500$ requires about 14.4 s of computational time when using both shadowing fields and random sample reuse. With the additional computational cost of A, B, and C, the distribution of I for *any* N can be well approximated in under 16 s using extrapolation.

H. Varying Physical Parameters

To further validate our approach, we will vary some of the physical parameters given in Table II and observe the resulting distributions.

1) Changing the Layout of Interferers: We now consider a layout of interferers that, rather than surrounding the victim receiver, are located on one side and laid out in a 2-D Gaussian distribution, truncated at two standard deviations. Such a layout can model, for example, the layout of interferers in a city, with the highest activity in the downtown core. The distribution is given by

$$g(\vec{r}) = \begin{cases} \frac{\exp\left(-\frac{\|\vec{r}-\vec{r}_0\|^2}{225^2/2}\right)}{225\pi(1-e^{-2})}, & \|\vec{r}-\vec{r}_0\| < 225 \\ 0, & \|\vec{r}-\vec{r}_0\| > 225 \end{cases} \quad (33)$$

where $\vec{r}_0 = (275, 0)$. This layout is a rotationally symmetric Gaussian distribution centered at (275, 0) with a standard deviation of 112.5 and truncated at two standard deviations, i.e., 225. Because this layout is entirely contained inside $50 \leq \|\vec{r}\| \leq 500$, we may use the same shadowing fields as in the previous layout. The numerical parameters for this layout (with the given shadowing correlation function) are

$$A = 3.361 \cdot 10^{-4}, \quad B = 1.421 \cdot 10^{-6}, \quad C = 1.096 \cdot 10^{-8}.$$

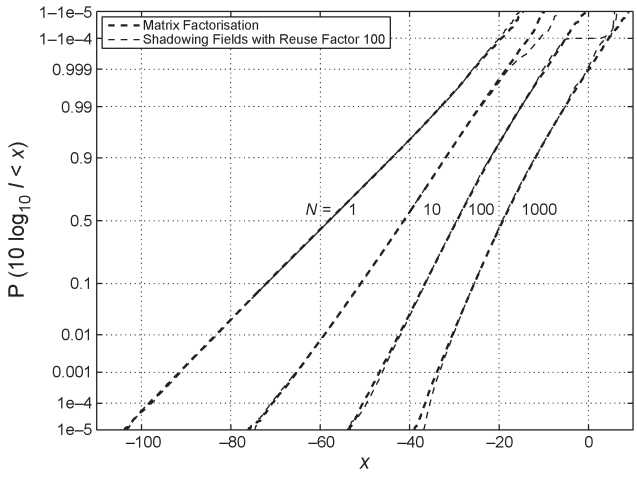

Fig. 9. Validation of the optimized simulation algorithm when the layout of interferers is a truncated Gaussian that is located away from the receiver.

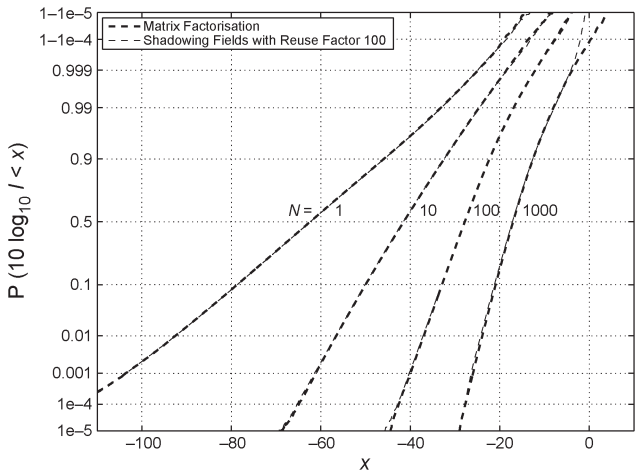

Fig. 10. Validation of the optimized simulation algorithm when the transmit powers of the interferers are *i.i.d.* exponential.

Fig. 9 illustrates the accuracy of shadowing fields with reuse for the Gaussian layout, where we see good agreement, except occasionally at extreme tail values. Of interest is that the shape of the obtained distributions is closer to the lognormal (i.e., a straight line on lognormal paper) than in the previous figures. This condition is consistent with our results in [10].

2) Incorporating i.i.d. *Random Transmit Power:* We now go back to the previous layout but add variability in the transmit powers T_i by modeling them as exponential RVs with mean 1 (see Fig. 10). Note that this approach gives the parameters $D = 1$ and $E = 2$. We again evaluate the accuracy of shadowing fields with reuse and find a good match between the distributions, except occasionally in the far tails.

The choice of the exponential distribution was arbitrary at this point, simply to validate the simulator. Note that the quantity in T_i need not only represent transmit power but may also incorporate any other phenomenon that is *i.i.d.* across all interferers, notably independent small-scale fading.

I. Optimizing Other Correlation Models

Although we have argued in [12] and herein for the several advantages of the model in (6), several other models can be

realized with shadowing fields with a little creativity, and of course, reuse and extrapolation are methods that do not depend on the correlation model. To generate shadowing fields according to other correlation functions, we must first find a transformation (e.g., $\mathcal{T}_{\mathrm{LP}}$) under which the underlying field is stationary. Then, we may exploit the particular form of the correlation (e.g., autoregressive) for an efficient implementation or, in general, execute convolutions with the appropriate kernel using the fast Fourier transform (FFT) algorithm, as shown in [25]. The complexity of FFT is $\mathcal{O}(D_\Theta D_{\mathrm{R}} \ln D_\Theta D_{\mathrm{R}})$ [23], which is not much higher than for separable triangular correlations. Further discussions on generating stationary Gaussian fields can be found in [22] and [26].

Furthermore, the FFT algorithm can be executed with significant time gains on specialized hardware, notably on several supported graphics processing units (GPUs) [27], which are commercially available and already included in several current desktop computers.

Finally, one important consideration is to ensure that the correlation model used is feasible, as explained in [12]. Otherwise, generating shadowing fields will be, by definition, impossible.

J. Model Extensions

Although the framework given in Section II still offers restrictions on the systems and channels that we can model, it should be emphasized that our approach extends well for more complex models. This is because the method focuses on the simulation of correlated shadowing, which is unaffected by other considerations when using shadowing fields. Furthermore, the techniques of random sample reuse and extrapolation are based on the very general assumption of exchangeability, which may hold (at least approximately) for a wide variety of scenarios.

Possible extensions that need not violate exchangeability are listed as follows:

- a random (e.g., Poisson distributed) number N of interferers;
- correlated (e.g., self-clustered) interferer positions;
- small-scale fading;
- directional victim receiver antenna;
- correlation between interferer positions and their transmit power.

We thus hope that our approach will be extended to help study large and complex wireless networks while allowing a greater degree of realism and computational efficiency than was formerly possible.

V. Conclusion

We have shown that the techniques of shadowing fields, random sample reuse, and extrapolation can be combined to reduce the computational time by several orders of magnitude. We can thus obtain the total interference distribution from large networks where position-dependent correlation in shadowing is employed, without significant loss in accuracy.

We have shown that, as the number of interferers increases, shadowing correlation becomes a dominating factor and needs to be included in the simulation and analysis of future interference-intensive scenarios. Furthermore, we have shown that the problem can be interpreted as a sum of exchangeable RVs and that this sum converges in distributions when properly normalized. This has led us to justify the extrapolation of simulation results from large to very large N and suggests the possibility of numerically finding the limiting distribution as an approximation for the distribution for large N, although how we can achieve this goal remains an open problem.

The shadowing correlation model used, which has been shown to have desirable mathematical and physical properties, was furthermore shown to be particularly well-suited to the technique of shadowing fields, which is, however, not restricted to this model.

By combining analytical, numerical, and simulation-based techniques, we can obtain accurate distributions in a very short time (no more than 16 s on a standard personal computer) for very detailed simulation models with the following flexible parameters: 1) any physical layout, 2) any path-loss and shadowing spread function of distance, 3) any transmit power distribution, and 4) a wide flexibility of shadowing correlation functions. Given the analytical complexity and versatility of the problem, a fully analytical approach is probably too complicated to attempt, and a mixed approach as presented here is probably the best solution to this very involved problem.

References

[1] S. R. Hall, A. W. Jeffries, S. E. Avis, and D. D. N. Bevan, "Performance of open access femtocells in 4G macrocellular networks," presented at the Wireless World Research Forum (WWRF) Meeting 20, Ottawa, ON, Canada, Apr. 2008, Tech. Rep.

[2] R. Hekmat and P. Van Mieghem, "Interference power sum with lognormal components in ad hoc and sensor networks," in *Proc. Int. Symp. Model. Optim. Mobile, Ad Hoc, Wireless Netw.*, Apr. 2005, pp. 174–182.

[3] A. Behnad, H. Purmehdi, and F. Lahouti, "Probability of outage in a clustered Poisson field of interfering nodes," in *Proc. IEEE Int. Conf. Telecommun.*, Apr. 2010, pp. 784–789.

[4] E. Sousa, "Performance of a spread spectrum packet radio network link in a Poisson field of interferers," *IEEE Trans. Inf. Theory*, vol. 38, no. 6, pp. 1743–1754, Nov. 1992.

[5] M. Win, P. Pinto, and L. Shepp, "A mathematical theory of network interference and its applications," *Proc. IEEE*, vol. 97, no. 2, pp. 205–230, Feb. 2009.

[6] M. Di Renzo, F. Graziosi, and F. Santucci, "Further results on the approximation of lognormal power sum via Pearson type IV distribution: A general formula for log-moments computation," *IEEE Trans. Commun.*, vol. 57, no. 4, pp. 893–898, Apr. 2009.

[7] C. Tellambura and D. Senaratne, "Accurate computation of the MGF of the lognormal distribution and its application to sum of lognormals," *IEEE Trans. Commun.*, vol. 58, no. 5, pp. 1568–1577, May 2010.

[8] A. Safak and R. Prasad, "Effects of correlated shadowing signals on channel reuse in mobile radio systems," *IEEE Trans. Veh. Technol.*, vol. 40, no. 4, pp. 708–713, Nov. 1991.

[9] T. Klingenbrunn and P. Mogensen, "Modeling cross-correlated shadowing in network simulations," in *Proc. IEEE Veh. Technol. Conf.*, Sep. 1999, vol. 3, pp. 1407–1411.

[10] S. S. Szyszkowicz and H. Yanikomeroglu, "Analysis of interference from large clusters as modeled by the sum of many correlated lognormals," in *Proc. IEEE Wireless Commun. Netw. Conf.*, Mar./Apr. 2008, pp. 741–745.

[11] D. Catrein and R. Mathar, "Gaussian random fields as a model for spatially correlated lognormal fading," in *Proc. Australasian Telecommun. Netw. Appl. Conf.*, Dec. 2008, pp. 153–157.

[12] S. S. Szyszkowicz, H. Yanikomeroglu, and J. S. Thompson, "On the feasibility of wireless shadowing correlation models," *IEEE Trans. Veh. Technol.*, vol. 59, no. 9, pp. 4222–4236, Nov. 2010.

[13] D. Kitchener, M. Naden, W. Tong, P. Zhu, G. Senarnath, H. Zhang, D. Steer, and D. Yu, "Correlated lognormal shadowing model," Tech. Rep., IEEE 802.16 Session 44, Jul. 2006.

[14] M. Gudmundson, "Correlation model for shadow fading in mobile radio systems," *Electron. Lett.*, vol. 27, no. 23, pp. 2145–2146, Nov. 1991.
[15] N. Beaulieu and Q. Xie, "An optimal lognormal approximation to lognormal sum distributions," *IEEE Trans. Veh. Technol.*, vol. 53, no. 2, pp. 479–489, Mar. 2004.
[16] B. de Finetti, "La prévision: Ses lois logiques, ses sources subjectives," *Annales de l'institut Henri Poincaré*, vol. 7, no. 1, pp. 1–68, 1937.
[17] R. L. Taylor, P. Z. Daffer, and R. F. Patterson, *Limit Theorems for Sums of Exchangeable Random Variables*. Totowa, NJ: Rowman & Allanheld, 1985.
[18] B. Alkire, "Cholesky factorization of augmented positive definite matrices," Elect. Eng. Dept., Univ. Calif., Los Angeles, CA, Dec. 2002.
[19] K. Kumaran, S. E. Golowich, and S. Borst, "Correlated shadow fading in wireless networks and its effect on call dropping," *Wireless Netw.*, vol. 8, no. 1, pp. 61–71, Jan. 2002.
[20] I. Forkel, M. Schinnenburg, and M. Ang, "Generation of two-dimensional correlated shadowing for mobile radio network simulation," in *Proc. Int. Symp. Wireless Pers. Multimedia Commun.*, Sep. 2004, vol. 2, pp. 314–319.
[21] Z. Wang, E. Tameh, and A. Nix, "Joint shadowing process in urban peer-to-peer radio channels," *IEEE Trans. Veh. Technol.*, vol. 57, no. 1, pp. 52–64, Jan. 2008.
[22] G. Storvik, A. Frigessi, and D. Hirst, "Stationary space–time Gaussian fields and their time autoregressive representation," *Statist. Model.*, vol. 2, no. 2, pp. 139–161, Jul. 2002.
[23] A. Lukin, "Tips and tricks: Fast image filtering algorithms," in *Proc. Int. Conf. Comput. Graph. Vis.*, Jun. 2007, pp. 186–189.
[24] P. Harley, "Short distance attenuation measurements at 900 MHz and 1.8 GHz using low antenna heights for microcells," *IEEE J. Sel. Areas Commun.*, vol. 7, no. 1, pp. 5–11, Jan. 1989.
[25] J.-J. Wu, "Simulation of rough surfaces with FFT," *Tribol. Int.*, vol. 33, no. 1, pp. 47–58, Jan. 2000.
[26] B. Kozintsev, "Computations with Gaussian random fields," Ph.D. dissertation, Univ. Maryland, College Park, MD, 1999.
[27] T. Preis, P. Virnau, P. Wolfgang, and J. J. Schneider, "GPU accelerated Monte Carlo simulation of the 2D and 3D Ising model," *J. Comput. Phys.*, vol. 228, no. 12, pp. 4468–4477, Jul. 2009.

Sebastian S. Szyszkowicz (S'09) received the B.A.Sc. degree in electrical engineering (specializing in communications) from the University of Ottawa, Ottawa, ON, Canada, in 2003 and the M.A.Sc. degree in electrical engineering and the Ph.D. degree in electrical and computer engineering from Carleton University, Ottawa, in 2007 and 2011, respectively.

In winter 2008, he was a Visiting Researcher with the University of Edinburgh, Edinburgh, U.K. He has given several invited talks about his research at universities in Europe and Canada. He is currently with the Department of Systems and Computer Engineering, Carleton University. His research interests include wireless interference and propagation and the mathematical challenges therein.

Dr. Szyszkowicz is the recipient of the National Sciences and Engineering Research Council of Canada Postgraduate Scholarship awards for the Master's and Ph.D. levels.

Furkan Alaca (S'11) received the B.Eng. degree (with high distinction) in communications engineering in 2010 from Carleton University, Ottawa, ON, Canada, where he is currently working toward the M.A.Sc. degree with the Department of Systems and Computer Engineering. He currently holds the National Sciences and Engineering Research Council of Canada (NSERC) Canada Graduate Scholarship at the Master's level.

Mr. Alaca received the NSERC Undergraduate Student Research Award for three consecutive years from 2008 to 2010. He is also the recipient of the Senate Medal for Outstanding Academic Achievement.

Halim Yanikomeroglu (M'98) was born in Giresun, Turkey, in 1968. He received the B.Sc. degree in electrical and electronics engineering from the Middle East Technical University, Ankara, Turkey, in 1990, the M.A.Sc. degree in electrical engineering [now electrical and computer engineering (ECE)], and the Ph.D. degree in ECE from the University of Toronto, Toronto, ON, Canada, in 1992 and 1998, respectively.

From January 1993 to July 1994, he was with the R&D Group of Marconi Kominikasyon A.S., Ankara. Since 1998, he has been with the Department of Systems and Computer Engineering, Carleton University, Ottawa, ON, where he is currently a Full Professor. He is also an Adjunct Professor with the Advanced Technology Research Institute, King Saud University, Riyadh, Saudi Arabia. His research interests include the physical, medium access, and networking layers of wireless communications. In recent years, his research has been funded by Huawei, Research In Motion, Samsung, the Communications Research Centre of Canada, Nortel, and the National Sciences and Engineering Research Council of Canada.

Dr. Yanikomeroglu is a member of the Carleton University Senate and a registered Professional Engineer in the province of Ontario. He is a member of the Steering Committee of the IEEE Wireless Communications and Networking Conference (WCNC) and has been involved in the organization of this conference over the years, including serving as a Technical Program Cochair of WCNC in 2004 and the Technical Program Chair of WCNC in 2008. He has been involved in the organization of more than 100 conferences in various capacities and has also given about 20 tutorials at such conferences. He was a General Cochair of the 2010 IEEE Vehicular Technology Conference 2010–Fall, which was held in Ottawa. He has served on the Editorial Boards of the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS and the IEEE COMMUNICATIONS SURVEYS AND TUTORIALS. He was the Chair of the IEEE's Technical Committee on Personal Communications. He received the Carleton University Faculty Graduate Mentoring Award in 2010, the Carleton University Graduate Students Association Excellence Award in Graduate Teaching in 2010, and the Carleton University Research Achievement Award in 2009.

John S. Thompson (M'03) received the B.Eng. and Ph.D. degrees from the University of Edinburgh, Edinburgh, U.K., in 1992 and 1996, respectively.

From July 1995 to August 1999, he was a Postdoctoral Researcher with the University of Edinburgh, funded by the U.K. Engineering and Physical Sciences Research Council (EPSRC) and Nortel Networks. Since September 1999, he has been a Lecturer with the School of Engineering, University of Edinburgh, where he was promoted as a Reader in October 2005. He is the Founding Editor-in-Chief of the *IET Signal Processing Journal*. He has published approximately 200 papers, including a number of invited papers, book chapters, and tutorial talks. He is a coauthor of an undergraduate textbook on digital signal processing. His research interests include signal processing algorithms for wireless systems, antenna array techniques, and multihop wireless communications.

Dr. Thompson was a Technical Program Cochair of the IEEE International Conference on Communications (ICC) held in Glasgow in June 2007 and the 2010 IEEE Global Communications Conference (GLOBECOM) held in Miami.