Ultrawideband Channel Modeling on the Basis of Information-Theoretic Criteria

Ulrich G. Schuster, *Student Member, IEEE,* and Helmut Bölcskei, *Senior Member, IEEE*

***Abstract*— We present results of two indoor ultrawideband channel measurement campaigns in the 2–5 GHz frequency band. In measurement campaign I (MC I), the channel is static and we sample it spatially, while in MC II the transmitting and receiving antennas are fixed and channel variation is induced by people moving in the environment. Transmitter and receiver are separated by up to 27 m in MC I and up to 20 m in MC II. To determine suitable small-scale fading distributions for the tap amplitudes of the discrete-time baseband-equivalent channel impulse response, we use Akaike's Information Criterion (AIC). Despite the large bandwidth, AIC supports the Rayleigh (MC I) or the Rice distribution (MC II). For data from MC II, we estimate the covariance matrix of the random channel impulse response and demonstrate that the number of corresponding significant eigenvalues, and hence the diversity order of the channel, scales approximately linearly with bandwidth. Contrary to the uncorrelated scattering assumption, we find that the channel taps are weakly correlated. The ergodic capacity predicted by the Ricean channel model with parameters estimated from MC II shows good agreement with the ergodic capacity obtained by direct evaluation of the measurement results, while the corresponding outage capacities show a worse fit for low outage probabilities because of shadowing.**

***Index Terms*— Ultrawideband (UWB) communication, channel modeling, Akaike's Information Criterion (AIC), Rayleigh channels, Ricean channels, uncorrelated scattering (US).**

I. Introduction and Overview of Results

WIRELESS communication systems that operate over *ultrawideband* (UWB) channels with several gigahertz of bandwidth promise ease of multiple access, efficient use of the radio spectrum through overlay techniques, and improved link reliability because of frequency diversity. Two key assumptions on the small-scale fading behavior of the channel are often made in theoretical analyses of UWB systems:

A1: The taps of the baseband-equivalent channel impulse response are circularly symmetric complex Gaussian distributed (Rayleigh fading assumption).

A2: The number of independent diversity branches, called *stochastic degrees of freedom* (DOF) in the following, scales linearly with bandwidth.

Manuscript received October 6, 2005; revised February 15, 2006; accepted March 5, 2006. The editor coordinating the review of this paper and approving it for publication was X. Shen. This work was supported in part by the Swiss Kommission für Technologie und Innovation (KTI) under grant 6715.2 ENS-ES and by the European Commission as part of the Integrated Project PULSERS under contract FP6-506897. Parts of this work were presented at the International Workshop on Convergent Technologies (IWCT) 2005, Oulu, Finland, and at the International Symposium on Information Theory (ISIT) 2005, Adelaide, Australia.

The authors are with the Communication Technology Lab., ETH Zurich, CH-8092 Zurich, Switzerland (e-mail: {schuster, boelcskei}@nari.ee.ethz.ch).

Digital Object Identifier 10.1109/TWC.2007.05857.

A1 can be justified by the central limit theorem: if each tap consists of a large number of contributions from different scatterers, the resulting tap distribution can be modeled as complex Gaussian [1]. The reasoning behind A2 is that the continuous-time channel satisfies the uncorrelated scattering (US) assumption.

Pierce [2] and Viterbi [3] relied on A1 to show that the capacity of an infinite-bandwidth fading channel under a constraint on the average input power is equal to the capacity of the infinite-bandwidth AWGN channel,[1] and Médard and Gallager [6], as well as Subramanian and Hajek [7], demonstrated that for an average-power constrained fading channel the mutual information obtainable with white-like signals equals zero in the wideband limit if A1 and A2 hold.

A1 and A2 are well justified for small bandwidths [1]. However, the high temporal resolution of UWB systems might result in a small number of partial waves [8] contributing to each channel tap [9], [10], which renders A1 questionable. The Nakagami [10], lognormal [11], and Weibull [12] distributions are proposed in the literature to model the channel tap amplitudes of UWB channels instead, but there are also measurement results that support the Rayleigh [13] or the Rice [14] amplitude distributions that arise from a complex Gaussian distribution of the individual channel taps. Satisfying A2 for increasing bandwidth requires arbitrarily rich scattering, which real-world environments might not provide. However, to the best of our knowledge, A2 has been addressed only in a single measurement campaign so far [15].

To obtain a channel model suitable for analysis and design of communication systems and to assess the validity of A1 and A2 at large bandwidths, we conducted two UWB channel measurement campaigns in the frequency band from 2 GHz to 5 GHz in an indoor public space. In measurement campaign I (MC I), the channel was static, and we sampled it spatially, while in MC II the transmitting and receiving antennas were fixed, and channel variation was induced by people moving in the environment. Our main contributions and conclusions are:

- We propose to use Akaike's Information Criterion (AIC) as a tool to compare the fit of the tap amplitude distributions put forward in the literature [10]–[14] to the measured data, and argue that AIC is better suited for this task than the goodness-of-fit (GOF) tests often used in this context.
- We find that, depending on the source of channel variation, the Rayleigh (MC I) or the Rice (MC II) amplitude distribution is still adequate to model small-scale fading of the

[1] Telatar and Tse [4] and Verdú [5] demonstrated that A1 is not necessary to obtain this result.

measured UWB channels. This result supports A1, despite the large bandwidth of 3 GHz.

- Under the assumption that the Rice amplitude distribution in MC II results from a jointly circularly symmetric complex Gaussian distribution of the channel taps around their respective means, we investigate the correlation between the taps. Although our results indicate that the taps are weakly correlated, A2 holds. Hence, the number of stochastic DOF increases approximately linearly with bandwidth.

Notation: Expectation with respect to (w.r.t.) the random variable (RV) X is denoted by $\mathbb{E}_X$. Cumulative distribution functions (CDFs) are denoted by the letters F and G, corresponding probability density functions (PDFs) by f and g. If a CDF, say G, depends on some parameter vector $\boldsymbol{\theta}$, we write $G_{\boldsymbol{\theta}}$. Estimated quantities are indicated by a hat $\hat{}$. The superscript ^{T} stands for transposition and ^{H} for conjugate transposition. All logarithms are w.r.t. the base e. A circularly symmetric complex Gaussian RV $Z \sim \mathcal{CN}(0,\sigma^2)$ is a RV $Z = X + jY$ where X and Y are independent and identically distributed (i.i.d.) $\mathcal{N}(0,\sigma^2/2)$.

II. Model Structure

A model is always an *approximation* of reality. Our goal is to find a channel model that approximates the measured channels as well as possible, while being analytically tractable. Throughout this paper, we restrict our attention to linear time-invariant (LTI) channels, and do not take into account Doppler dispersion. The effect of an LTI channel with impulse response $h(t)$ on a transmitted signal $s(t)$ is described by the baseband-equivalent input-output relation

$$y(t) = \int_{-\infty}^{\infty} h(\tau)s(t-\tau)d\tau \tag{1}$$

where $y(t)$ denotes the channel output signal. The complicated effects of the physical propagation environment that determine $h(t)$ can often be approximated sufficiently well if $h(t)$ is modeled as a random process. Our measurement setup allows only for the characterization of a *block-fading* model, where $h(t)$ remains constant for the duration of one block and changes to an independent realization in the next block. The prevalent physically motivated wideband block-fading model in the literature is given in baseband-equivalent form as [1]

$$h(t) = \sum_{i=1}^{M} a_i \delta(t-\tau_i) \tag{2}$$

where M is the number of multipath components, a_i is the complex-valued random gain of the ith multipath component, and τ_i is the corresponding random delay. The Dirac function $\delta(t)$ in (2) has a flat spectrum; therefore, it models ideal reflections from smooth surfaces. But other propagation mechanisms, such as reflections from rough surfaces, diffraction, and scattering, are frequency-dependent in general, and thus result in time dispersion even along a single propagation path [16]. Although this effect is negligible for narrowband systems, it might become important in UWB channels. A possible refinement of the model in (2) is to replace the Dirac impulses by general functions $d_i(t)$, potentially different for each propagation path [17]. We take yet another approach and *discretize* the input-output relation (2). Wireless communication systems always use effectively band-limited signals, so that the effective channel, consisting of the propagation channel [1], the antennas, and the transmit and receive filters, will be band limited. Sampling $s(t)$ and $y(t)$, we obtain

$$y[m] = \sum_{l=0}^{L-1} h[l]s[m-l]. \tag{3}$$

The discrete-time impulse response $h[l]$ of the effective channel consists of L taps, $s[m]$ denotes the discrete-time input signal, and $y[m]$ the discrete-time output signal. Note that the discrete-time input-output relation (3) does not entail a loss of information compared with (1), since we assume the input signal to be band limited. Directly characterizing the statistics of the channel taps $h[l]$ is more robust than specifying the statistics of the gain parameters a_i and the delay parameters τ_i in (2), because the tap statistics only need to reflect the aggregated effect of the propagation phenomena contributing to each tap [18]. Moreover, the direct statistical characterization of the taps in (3) results in a more accurate description of the effective channel, because the model structure imposes fewer assumptions about the effective channel than the model structure in (2), which is potentially ill-suited for UWB channels, as discussed above.

III. The Statistical Modeling Problem

Throughout this paper, we restrict our attention to small-scale fading only, i.e., the variation of the $h[l]$ due to constructive and destructive interference of the arriving waves at the receiving antenna. Our goal is to statistically characterize the L-dimensional channel vector $\mathbf{h} = \left[h[0]\ h[1] \ldots h[L-1]\right]^T$ by selecting suitable marginal tap distributions $F_{h[l]}$ and a joint distribution $F_{\mathbf{h}}$. We think of the taps $h[l]$ as being distributed according to the unknown *operating model*, defined as the nearest representation of the true situation that can be constructed by means of a probability model[2] [19]. The main goal in statistical channel modeling is to find a probability model for the channel taps that approximates the operating model as closely as possible; a measure of approximation quality is called a *discrepancy* [19]. The selection of the probability model should be based on *physical insight*, *mathematical tractability*, and *consistency* w.r.t. measured data.

The most widely used approach to characterize the marginal tap amplitude distributions on the basis of measurement data is through GOF tests [20]: The null hypothesis $\mathcal{H}_0$ is that a given parameterized function $G_{\boldsymbol{\theta}}$ equals the CDF F of the operating model, i.e., $\mathcal{H}_0 : G_{\boldsymbol{\theta}} = F$. The alternative hypothesis is the complementary event $\mathcal{H}_1 : G_{\boldsymbol{\theta}} \neq F$. If a suitable test statistic (i.e., a function of the measured data) exceeds a given threshold, $\mathcal{H}_0$ will be accepted. The probability that $\mathcal{H}_0$ is rejected although it is true is called the *significance level* of the test. In the context of channel modeling, GOF tests are commonly applied by first estimating the parameter vector $\boldsymbol{\theta}$ from measured data,

[2]We have an unfortunate collision of terms: a *model* in the channel modeling literature refers to an input-output relation as in (2) or (3), together with the corresponding parameterization. A *probability model* refers to the characterization of a quantity by means of a random variable. We will explicitly state what type of model we refer to if the meaning is not clear from the context.

and subsequently computing the corresponding test statistics. The test is performed with a common significance level (chosen ad hoc) for every candidate distribution, and the distribution with the highest passing percentage across several measurement locations, or taps, is selected [11], [12], [21]. However, because passing percentages of GOF tests are not well-defined discrepancies [22], comparing them for several candidate distributions does not constitute a valid model selection criterion. Direct application of GOF tests also fails, because the space of possible channel tap distributions is virtually infinite for complex operating models—the probability of selecting $G_{\boldsymbol{\theta}}$ so that $G_{\boldsymbol{\theta}} = F$ will be zero.

Pioneering work in the field of model selection through discrepancy minimization was done by Akaike, whose information criterion (AIC) [23] has found widespread use. One of the two main contributions of the present paper is to use AIC for UWB channel modeling to verify A1. To the best of our knowledge, model selection methods to determine the amplitude distribution of narrowband channels are proposed only in [24], where the authors use the principle of minimum description length (MDL) [25]. For linear models [26], e.g., regressive models, MDL is asymptotically efficient in the number of samples if the operating model is among the candidate models; if the operating model is not a candidate model, AIC asymptotically selects the model that minimizes the expected squared error between model and data [26], [27]. Because we cannot expect the operating model to be part of our candidate set, we heuristically chose to use AIC, although we are not dealing with linear models.

IV. A Brief Review of Model Selection Using AIC

Our review of the basics of model selection follows Linhart and Zucchini [19]. We restrict our discussion to univariate CDFs, corresponding to the characterization of the individual channel taps' marginal distributions. The model selection methodology described below can also be used to characterize multivariate distributions; for the problem at hand, however, this is hardly feasible because of the large bandwidth and the resulting large number of taps. We will provide further comments about the joint distribution of $\mathbf{h}$ in Section VII.

Denote the unknown CDF of the operating model by F and the set of all univariate CDFs by $\mathcal{M}$. A parametric *candidate family* $\mathcal{G}^j = \{G^j_{\boldsymbol{\theta}^j} \mid \boldsymbol{\theta}^j \in \mathcal{T}^j\}$ is a subset of $\mathcal{M}$, where individual CDFs $G^j_{\boldsymbol{\theta}^j}$ are identified by the U-dimensional parameter vector $\boldsymbol{\theta}^j \in \mathcal{T}^j$, with $\mathcal{T}^j \subset \mathbb{R}^U$. For notational convenience, we take $G^j_{\boldsymbol{\theta}}$ to mean $G^j_{\boldsymbol{\theta}^j}$ in the following. Candidate families need to be chosen *in advance* to reflect prior knowledge about the modeling problem [28]. The set of J candidate families $\mathcal{C} = \bigcup_{j=1}^{J} \mathcal{G}^j$ constitutes the *candidate set*. A *discrepancy* is a functional $\Delta : \mathcal{M} \times \mathcal{M} \to \mathbb{R}$ that satisfies $\Delta(G, F) \geq \Delta(F, F)$ for all $G \in \mathcal{M}$. A consistent estimator for the discrepancy $\Delta(G^j_{\boldsymbol{\theta}}, F)$ on the basis of N independent samples, distributed according to F, is called an *empirical discrepancy* and will be denoted by $\Delta_N(G^j_{\boldsymbol{\theta}}, F)$.

Our goal is to choose the distribution that minimizes the discrepancy among all members of the candidate set. The procedure consists of two steps: With the operating model unknown, we first estimate the parameter vector $\boldsymbol{\theta}^j$ for each candidate family $\mathcal{G}^j$ from N i.i.d. samples $\mathbf{x} = [x_1\, x_2 \ldots x_N]^T$ using the *minimum discrepancy estimator* $\hat{\boldsymbol{\theta}}^j(\mathbf{x}) = \arg\min_{\boldsymbol{\theta} \in \mathcal{T}^j} \Delta_N(G^j_{\boldsymbol{\theta}}, F)$; for simplicity, we write $\hat{\boldsymbol{\theta}}^j$ instead of $\hat{\boldsymbol{\theta}}^j(\mathbf{x})$ in the following. Because $\hat{\boldsymbol{\theta}}^j$ depends on samples $\mathbf{x}$ that are realizations of the RV $X \sim F$, the resulting discrepancy $\Delta(G^j_{\hat{\boldsymbol{\theta}}}, F)$ is a RV. A good probability model should lead to consistent predictions; hence, it must provide a good approximation to the operating model on average, not just for the actual samples $\mathbf{x}$. The second step is thus to find j so that the *expected discrepancy* $\mathbb{E}_X[\Delta(G^j_{\hat{\boldsymbol{\theta}}}, F)]$ is minimized over the candidate set. This two-step approach shows that the overall discrepancy consists of two distinct contributions: (i) the *approximation discrepancy* is the error induced by selecting a probability model different from the operating model, even if $\hat{\boldsymbol{\theta}}^j = \boldsymbol{\theta}^j$; (ii) the *estimation discrepancy* is the error caused by estimating the parameters of the distribution from a finite number of samples. A more complex probability model with more free parameters U will, in general, have a lower approximation discrepancy at the cost of a larger estimation discrepancy. A sensible model choice aims at balancing both discrepancies w.r.t. the number of samples available.

The discrepancy used in AIC is based on the Kullback–Leibler (KL) distance [28]. For two PDFs f and g, the KL-distance is defined as

$$D(f \,||\, g) = \mathbb{E}_Y[\log f(Y)] - \mathbb{E}_Y[\log g(Y)] \tag{4}$$

where Y is distributed according to the operating model with PDF f. The KL distance $D(f \,||\, g)$ is nonnegative and equals zero only if $f = g$. The first term on the right-hand side of (4) depends on the operating model only; to obtain a suitable discrepancy, it thus suffices to consider the second term, which is called *KL discrepancy*. Consequently, the expected KL discrepancy is [28]

$$-\mathbb{E}_{\hat{\boldsymbol{\theta}}}\Big[\mathbb{E}_Y\big[\log g_{\hat{\boldsymbol{\theta}}}(Y)\big]\Big] \tag{5}$$

where the inner expectation is w.r.t. the operating PDF f, and the outer expectation is w.r.t. the distribution of the parameter estimate $\hat{\boldsymbol{\theta}}$, which is a function of the data $\mathbf{x}$. AIC is an approximately unbiased estimator of the expected discrepancy (5) between the candidate families $\mathcal{G}^j$, $j = 1, 2, \ldots, J$, and the operating model, and is given by [23]

$$\mathrm{AIC}_j = -2\sum_{n=1}^{N} \log g_{\hat{\boldsymbol{\theta}}^j}(x_n) + 2U. \tag{6}$$

As a rule of thumb, we need $N/U \gtrsim 40$ to obtain useful AIC values [28, Ch. 2.4]. The corresponding minimum discrepancy estimator is the maximum likelihood (ML) estimator:

$$\hat{\boldsymbol{\theta}}^j = \arg\max_{\boldsymbol{\theta} \in \mathcal{T}^j} \frac{1}{N} \sum_{n=1}^{N} \log g_{\boldsymbol{\theta}}(x_n). \tag{7}$$

AIC estimates the approximation quality of different probability models. It can be used to rank the candidate distributions; the minimum AIC value indicates the best fit. To conveniently compare the relative fit of each distribution within the candidate set, we define the AIC differences [28] $\Phi_j = \mathrm{AIC}_j - \min_i \mathrm{AIC}_i$,

where $\min_i \mathrm{AIC}_i$ denotes the minimum AIC value over all J candidate families, and compute the *Akaike weights* [29]

$$w_j = \frac{e^{-\frac{1}{2}\Phi_j}}{\sum_{i=1}^{J} e^{-\frac{1}{2}\Phi_i}} \tag{8}$$

which satisfy $\sum_{j=1}^{J} w_j = 1$. AIC is an approximately unbiased estimator of the expected KL discrepancy (5), so that an estimate of the likelihood $\mathcal{L}(m_j \,|\, \mathbf{x})$ of the probability model m_j with CDF $G^j_{\hat{\boldsymbol{\theta}}} \in \mathcal{G}^j$ is obtained through the transformation $\mathcal{L}(m_j \,|\, \mathbf{x}) = \alpha e^{-\Phi_j/2}$, with normalization constant α [28]. Therefore, the weight w_j can be interpreted as an estimate of the probability that the CDF $G^j_{\hat{\boldsymbol{\theta}}}$ shows the best fit within the candidate set [29]. Consequently, the Akaike weights allow us not only to select the best distribution in the candidate set, but also provide information about the relative approximation quality of each distribution. If several distributions have similar Akaike weights, the evidence provided by the data is not sufficient to allow for a better differentiation [28, Ch. 4.5].

V. The Channel Measurement Campaigns

We conducted two different UWB channel measurement campaigns to characterize the small-scale fading behavior of channels for several envisioned UWB application scenarios. Both campaigns were carried out in the lobby of the ETZ building at ETH Zurich, a typical public open-space environment. A floor plan is shown in Fig. 1. The lobby has large windows over the entire length on one side, and concrete walls, covered by metal plates, on the opposite side. The floor is tiled with natural stones, the ceiling is made of concrete; there is a row of concrete pillars in parallel to the windows. The office in the top right-hand corner of Fig. 1 is separated from the lobby by windows, the conference room at the bottom-right is separated by a brick wall with rough surface. The raw measurement data of both measurement campaigns are available for download [30].

A. Measurement Campaign I (MC I)

1) Measurement Setup: In wireless access point scenarios or peer-to-peer systems, at least one of the terminals moves w.r.t. the static environment. To assess the impact of the corresponding spatial variations of the channel, we positioned the transmit antenna on a rectangular 9×5 grid, while the receive antenna was fixed. The grid spacing was 7 cm in both dimensions—slightly smaller than half the wavelength at the lowest measured frequency (2 GHz)—to obtain samples that can be considered approximately independent.[3] We took measurements for two grid positions displaced by approximately 50 cm, resulting in a total of $N = 90$ samples. To separate small-scale effects from large-scale variation of the received power, the overall area over which the antenna was displaced was small compared with the distance d between the transmitter and the receiver grid. We recorded 90 transfer functions for various distances d in three different settings: (i) five line-of-sight (LOS) measurements with distances indicated in Fig. 1, (ii) three measurements where the row of concrete pillars partly obstructed the line-of-sight (OLOS), with $d = 27.5\,\mathrm{m}$, $d = 24.6\,\mathrm{m}$, and $d = 21.7\,\mathrm{m}$, and (iii) two non-line-of-sight (NLOS) measurements, with distances $d = 22.1\,\mathrm{m}$ and $d = 22.7\,\mathrm{m}$.

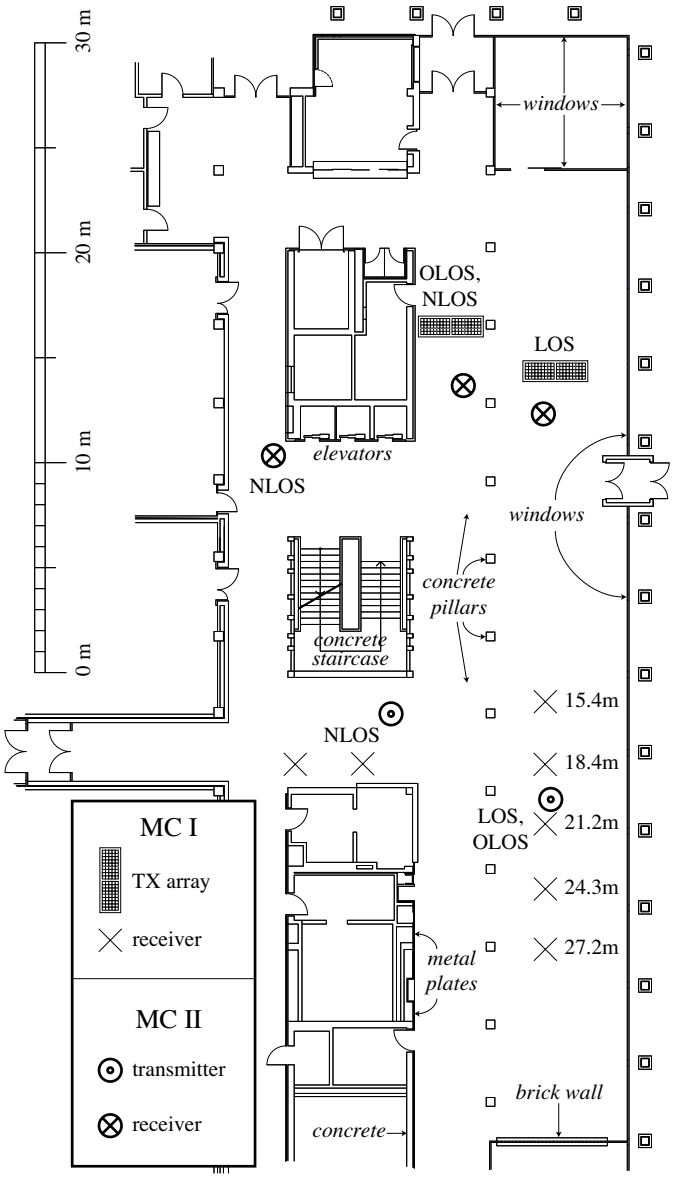

Fig. 1. Floor plan of the lobby in the ETZ building at ETH Zurich.

A schematic of the measurement setup is shown in Fig. 2. We recorded one channel transfer function for each of the 90 positions of the transmit antenna, using an HP 8722D vector network analyzer (VNA) operating in stepped frequency mode and with an IF bandwidth of 300 Hz; the sweep time was 9.8 s [31]. Every transfer function was recorded at 1601 equally-spaced frequency points in the band from 2 GHz to 5 GHz. The resulting frequency resolution of 1.875 MHz implies a maximum resolvable delay of 533 ns. The sounding signal was transmitted at a power level of 25 dBm; to overcome the loss of the long cable between the receive antenna and the VNA, we placed an additional amplifier close to the receive antenna. Because we are interested in the characterization of effective channels as typical UWB communication systems would see them, we used a prototype version of an off-the-shelf UWB antenna (Skycross SMT-3TO10M) for both transmitter and receiver, and did not

[3]Independence over space is not a property of the channel, but a modeling assumption. To ensure that channel samples vary enough between grid points to be *modeled* as independent, these samples should result from the superposition of waves whose phases differ by more than π between neighboring grid points.

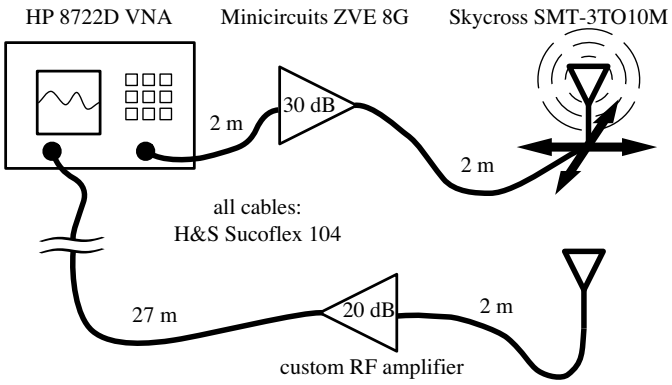

Fig. 2. Measurement setup for MC I.

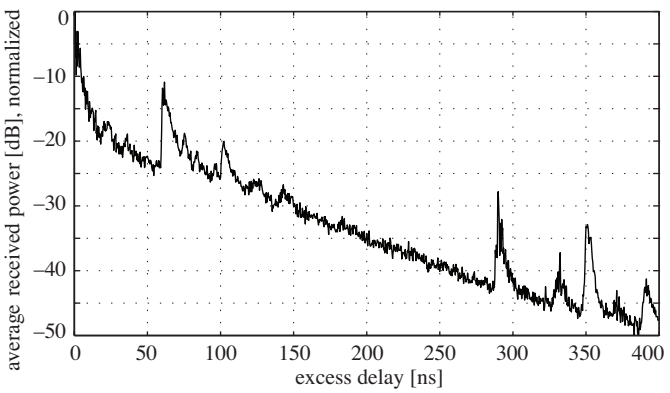

Fig. 3. MC I: power-delay profile, LOS setting with $d = 27.2\,\mathrm{m}$.

remove the effect of the antenna transfer characteristic from the measured data. At the transmitter side, we mounted the antenna on a positioning table, 1.9 m above the floor; the receive antenna was fixed on a tripod at a height of 1.75 m. To prevent temporal variation in the channel, we ensured that nobody was moving in the lobby, switched off all fluorescent lights, and shut down all IEEE 802.11 base stations.

2) Calibration and Postprocessing: On the basis of an initial back-to-back calibration measurement without antennas, the VNA performed automatic calibration during subsequent measurements. We used a 1601-point inverse Discrete Fourier Transform (IDFT) to convert the calibrated measurements of the transfer functions to the time domain in equivalent baseband representation. To obtain a set of impulse responses as a receiver synchronized to the corresponding transmitter might see them, we aligned all 90 recorded impulse responses for each distance d by DFT-based interpolation with eightfold oversampling, visual alignment, and subsequent downsampling of the aligned impulse responses.

3) Measurement Results: For the LOS setting with distance $d = 27.2\,\mathrm{m}$, Fig. 3 shows the power-delay profile (PDP) obtained by averaging over the squared magnitude of all 90 aligned impulse responses. Because the effective channel is band limited, the PDP and all corresponding impulse responses are "dense", i.e., there are no empty taps as would be predicted by a model of the form in (2). We can observe several clusters protruding from the otherwise linearly (in dB) decaying PDP. The delay relative to the first arrival, called *excess delay*, of prominent cluster peaks matches the path distance of reflections from the window between the lobby and the office at the top right in Fig. 1, and multiple reflections from the office and the brick wall at the bottom. The PDPs for the different OLOS settings not shown here have only one cluster, which can be attributed to a reflection from the windows on the right in Fig. 1. The direct path is obstructed by the row of concrete pillars, and, therefore, does not convey the largest part of the energy. This phenomenon is even more pronounced in the NLOS setting, where the direct path is blocked by the elevator pit and the concrete staircase.

B. Measurement Campaign II (MC II)

1) Measurement Setup: In wireless personal area networks and wireless sensor networks, e.g., for building automation or industrial applications, the terminals are typically static, and variations in the channel are mainly caused by moving scatterers, e.g., moving persons. Time-variability of the channel precludes frequency-domain sounding with a VNA because of its long sweep time. Instead, we used the digital sampling oscilloscope (DSO) Agilent DSO81204A to record the channel output in real time. Results of other UWB time-domain measurement campaigns are reported in [10], [32]. Our sounding signal was a periodically repeated pseudonoise (PN) sequence of length $2^{15} - 1$, clocked at 10 GHz, so that the received signal consisted of a convolution of the PN sequence and the impulse response of the effective channel [33], [34]. The measurement system is shown schematically in Fig. 4. The repetition rate of the PN sequence of over 300 kHz was several orders of magnitude higher than the maximum Doppler shift in the channel; hence, we can assume that the channel remained static for the duration of several periods of the PN sequence. The DSO sampled the signal received at port 1 at a rate of 40 GHz, with an analog bandwidth of 12 GHz. Limited DSO memory allowed us to record the channel output only for approximately 3.3 μs at a time; we call the corresponding data a *channel snapshot*. Our measurement system was able to record approximately one snapshot per second. Because the DSO sampled port 1 and port 2 synchronously, we used the copies of the PN sequence transmitted over a long cable and recorded at port 2 as *reference snapshots* to infer the timing of the channel snapshots, as will be detailed later. For accurate timing, we used phase-stable semirigid cables to connect the antennas to transmitter and receiver. Other components used in MC II, such as amplifiers and antennas, were identical to those in MC I. Both antennas were mounted on tripods, approximately 1.6 m above the floor. We chose similar measurement locations as in MC II, although we had to interchange the position of transmitter and receiver, as indicated in Fig. 1. In the LOS setting, with a distance between transmitter and receiver of approximately $d = 20\,\mathrm{m}$, we acquired $N = 1011$ snapshots, in the OLOS setting $N = 2722$ snapshots, also for $d = 20\,\mathrm{m}$, and $N = 1256$ snapshots in the NLOS setting with $d = 13\,\mathrm{m}$.

2) Calibration and Postprocessing: Each channel snapshot contained two contiguous periods of the PN sequence and additional incomplete periods of the sequence at the beginning and at the end of the snapshot. To extract the impulse response of the effective channel from 2 GHz to 5 GHz, we applied the following postprocessing steps to the measured data: (i) We

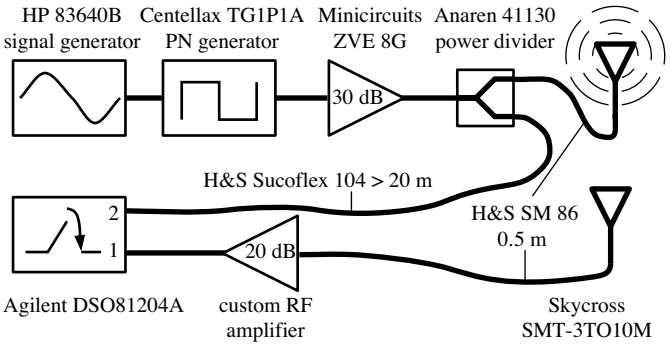

Fig. 4. Measurement setup for MC II.

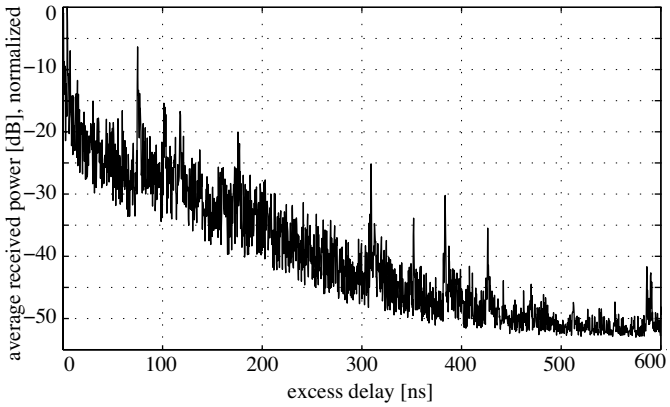

Fig. 5. MC II: power-delay profile, LOS setting with $d = 20\,\mathrm{m}$.

correlated each channel snapshot and each reference snapshot with the PN sequence, so that the resulting snapshots contained two contiguous recordings of the corresponding impulse responses, called *raw samples* in the following. (ii) As in MC I, we applied DFT-based interpolation with eightfold oversampling. (iii) Because transmitter and receiver were not synchronized, the start of the first raw sample in each reference snapshot was offset by an unknown number b of taps from the start of the snapshot; the first raw sample in the channel snapshot was offset by $b + c$, were c denotes the offset due to the delay difference in the channel being measured and the reference channel. This offset was constant for a given measurement setting (LOS, OLOS, NLOS), and was determined manually. Because the reference channel was constant over time, we could automatically determine b for each reference snapshot, and thus extract the two raw channel samples. (iv) Due to a technical problem in the DSO, c was sometimes off by an additional two taps. We manually identified the affected raw channel samples and corrected for this offset. (v) As the channel was virtually constant over two subsequent periods of the PN sequence, we averaged the corresponding raw channel samples to improve the measurement SNR. (vi) We transformed the averaged raw channel samples to the frequency domain using the DFT, extracted the band from 2 GHz to 5 GHz, and applied the IDFT to obtain baseband-equivalent representations of the impulse responses of the effective channel. (vii) We decimated and truncated the baseband-equivalent impulse responses to 2001 taps at a sampling rate of 3 GHz. (viii) To calibrate the measured data, we equalized each impulse response, using the response of the measurement system without antennas. The measurement system response was estimated from over 100 back-to-back measurements. (ix) Because MC II was conducted at daytime, we could not deactivate the IEEE 802.11 base stations installed in the lobby. To remove the interference, we filtered the impulse responses with a notch filter to suppress all frequencies between 2.4 GHz and 2.485 GHz.

3) Measurement Results: The PDPs obtained for each setting in MC II are less smooth than the corresponding PDPs in MC I. This can be seen, e.g., for the LOS setting, by comparing Fig. 3 with Fig. 5. An explanation for this difference is the presence of a strong mean component in *each* tap of the MC II impulse responses. If we subtract the mean and compute the PDP from the zero-mean impulse responses, the resulting PDP more closely resembles the PDP in Fig. 3.

VI. Discussion of A1: Marginal Tap Statistics

We applied AIC to our measurement data to evaluate the fit of the tap amplitude distributions proposed in [10]–[14]. Our candidate set $\mathcal{C}$ consists of the single-parameter ($U = 1$) Rayleigh family and the two-parameter ($U = 2$) Rice, Nakagami, lognormal, and Weibull families. The Rice, Nakagami, and Weibull families contain the Rayleigh family as a special case. Rayleigh and Rice amplitude distributions are widely used for narrowband fading models, and have a clear physical interpretation [1], [8]: if the received field is modeled as the superposition of many independent partial waves of comparable average power that originate from diffuse reflections by random rough surfaces, the central limit theorem will state that the individual distributions of the channel impulse response taps approach a complex Gaussian distribution, circularly symmetric around a potentially present mean. The corresponding tap amplitudes will be Rayleigh distributed if the mean is zero, or Rice distributed otherwise. If we assume that the partial waves result from partially diffuse and partially specular reflections from random rough surfaces, the partial waves will be correlated, with unequal variance of the real and imaginary components [8]. In this case, the Nakagami distribution is a good approximation for the resulting amplitude distribution [35], [8]. It has been found to fit UWB measurements [10]. The lognormal distribution is often used to model shadowing, but there seems to be no physical interpretation for its use as a small scale fading model. Nevertheless, some researchers report good agreement with UWB measurement data [11], [32], [36], and the widely used IEEE 802.15.3a channel model also uses the lognormal distribution [37], [9]. Although physical motivation for the Weibull distribution is scarce, a good fit was reported in several measurement campaigns as well [12], [38].

A. Measurement Campaign I

Fig. 6 shows the normalized PDP and Akaike weights for each candidate family, computed for the LOS setting with $d = 27.2\,\mathrm{m}$. Our findings can be summarized as follows:

1) The Rayleigh distribution shows the best fit, followed by the Rice, Nakagami, and Weibull distributions in no particular order.
2) The lognormal distribution exhibits a bad fit, with the exception of taps in the vicinity of cluster peaks.
3) The variability of the Akaike weights across taps is high.

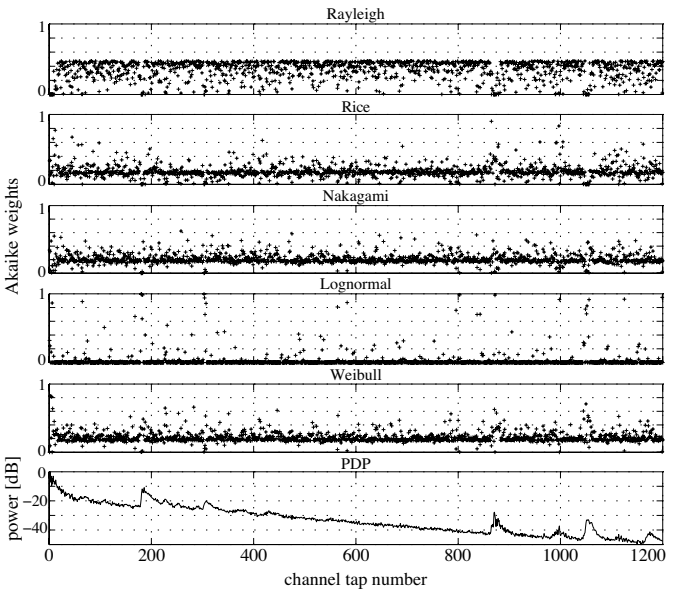

Fig. 6. MC I: PDP and Akaike weights, LOS setting with $d = 27.2\,\mathrm{m}$.

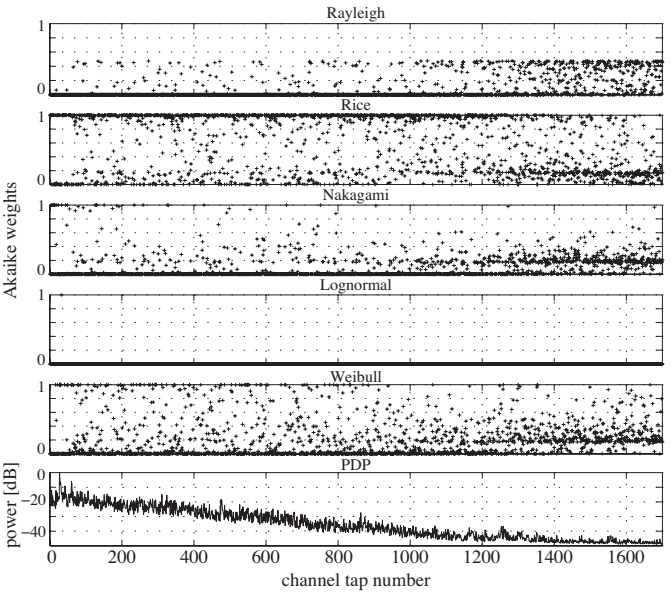

Fig. 7. MC II: PDP and Akaike weights, OLOS setting with $d = 20\,\mathrm{m}$.

To double-check our results, we applied AIC to NLOS data from UWB measurements conducted in an office [14], and obtained similar results.

Interpretation of the Results: The parameter estimates $\hat{\boldsymbol{\theta}}^j$ for the Rice, Nakagami, and Weibull distributions are close to the values that reduce the respective distributions to the Rayleigh distribution. The Ricean K-factor, for example, is well below 5 dB for taps not in the vicinity of a cluster peak; for such small values, however, reliable estimation of the K-factor is impossible [39]. Hence, the variability of the Akaike weights across taps results from the high sensitivity of the w_j to the parameter estimates $\hat{\boldsymbol{\theta}}^j$ of distributions that are close to each other, indicating that there is not enough evidence in the data to support a better differentiation. The lognormal distribution provides a bad fit for most taps, and is thus not suited to model the tap fading of the measured channel. The difference of the Akaike weights between the Rayleigh distribution ($U = 1$) and the Rice, Nakagami, and Weibull distributions ($U = 2$) is mainly due to the term $2U$ in (6). The reason for this behavior lies in the nature of AIC, which measures approximation and estimation discrepancy: the estimation discrepancy increases with the number of free parameters U. Hence, the fit of the Nakagami, Rice, and Weibull distributions to our measured data is evidently not good enough to warrant this additional complexity.

The better fit of the lognormal, Weibull, and Rice distributions for taps at the beginning of clusters in the LOS setting shown in Fig. 6 is probably an artifact of our modeling assumptions: The peak of each cluster probably results from a specular reflection of a large, even surface, like a window in the lobby. The excess delay of such a reflection will vary across the 90 impulse responses, because they are measured at different positions on the grid. Consequently, a specular reflection will not always contribute to the same tap, but it will appear as an outlier in several adjacent taps instead, resulting in a better fit of more heavy-tailed PDFs. Contrary to our modeling assumption, these PDFs do not accurately describe the small-scale fading that results from constructive and destructive interference of partial waves at different points on the measurement grid; rather, they reflect the impact of a single specular reflection.

B. Measurement Campaign II

Fig. 7 depicts the normalized PDP and the Akaike weights obtained for the OLOS setting of MC II. Results obtained for the LOS and NLOS settings are similar, and will hence not be shown here. Our findings can be summarized as follows:

1) The Rice distribution shows the best fit for most of the first 1200 taps.
2) The Rayleigh distribution provides a good fit for most of the taps with indices larger than 1200.
3) The lognormal distribution does not fit the measurement data at all,[4] while the Nakagami distribution is suitable only for the first few taps.
4) The Weibull distribution shows a good fit for some taps.

Interpretation of the Results: Ricean fading is often attributed to a strong mean component in the impulse response [1], which renders the individual channel taps less variable compared with Rayleigh fading. While in MC I, large scattering objects, like windows and walls, move relative to the position of the antennas, the moving scatterers in MC II, i.e., people in the lobby, are much smaller. Hence, many significant contributions in almost all MC II channel taps are due to static scatterers, consistent with the physical interpretation of the Rice distribution, our observations of the large fluctuations of the PDP in Section V-B, and with results of fixed-antenna measurements of narrowband indoor channels in [40] and [38]. Because transmit and receive antennas were mounted approximately 1.6 m above the floor, people moving in the environment sometimes blocked the LOS or propagation paths of some of the dominant reflections. Consequently, it is not possible to perfectly separate small-scale fading and shadowing in MC II, as

[4]The Akaike weights for the lognormal distribution are not strictly equal to zero, but they are too small to be distinguishable from zero on the scale of Fig. 7.

will be discussed in more detail in Section VIII-B. The Weibull distribution often fits best for shadowed taps. A recent study of temporal variation in a UWB channel with fixed antennas also finds a good fit of the Weibull distribution [12]. However, the measurement methodology in [12] is different from MC II.

Approximately from tap 1200 on, the Rayleigh distribution exhibits a better fit than the Rice distribution. This effect results from the low measurement SNR of these taps; effectively, we are fitting Gaussian distributed noise.

C. General Comments

The above analysis shows that even for bandwidths of up to 3 GHz our measurement results support the Rayleigh and Rice distributions. However, the differences to the Nakagami and Weibull distributions in terms of the Akaike weights are often small, especially in MC I. Consequently, the data do not provide enough evidence to unequivocally select a single distribution. We can thus either resort to multimodel inference [28], or use extrinsic criteria to narrow down our selection. One important extrinsic criterion is the mathematical tractability of a given channel model, which leads us to advocate the use of the Rayleigh and Rice distributions. As the Nakagami or Weibull distributions are often close to the Rayleigh or Rice distributions in the sense of AIC, performance predictions of UWB communication systems on the basis of any of these distributions should yield similar results.

We did not apply model selection tools to determine the distribution of the phases of the complex-valued channel taps, as there does not seem to be a physically motivated alternative to the almost exclusively used uniform-phase assumption. To see if this assumption holds for our measured data, we examined scatter plots of several measured taps. For MC I taps, we find good agreement with the uniform-phase assumption, except for taps corresponding to cluster peaks; likewise, scatter plots of MC II taps show that the tap values are well localized in the complex plane, circularly symmetric around a nonzero mean.

The combination of the Rayleigh or Rice amplitude distribution with a uniformly distributed phase of the zero-mean component results in the complex Gaussian distribution for each individual tap. Marginally Gaussian distributions do not imply joint Gaussianity of the channel impulse response vector $\mathbf{h}$. Unfortunately, selecting the joint PDF of $\mathbf{h}$ using AIC is a hopeless endeavor: to characterize the large number of taps, we would need many more samples than can be practically obtained by any measurement procedure. As a simple heuristic test, we verified through AIC that the tap amplitudes of the DFT of $\mathbf{h}$ are Rayleigh, respectively Rice, distributed. If some vector $\mathbf{a}$ is jointly Gaussian distributed, circularly symmetric around its mean, the DFT of $\mathbf{a}$ will result in another jointly complex Gaussian vector with circularly symmetric zero-mean component. Hence, with no evidence against the joint Gaussianity assumption, and referring again to the analytical tractability of a good model, we advocate the use of the jointly complex Gaussian distribution for $\mathbf{h}$.

VII. Discussion of A2: The Uncorrelated Scattering Assumption

Correlation between channel taps in the discrete-time model (3) can result from correlated scattering in the underlying continuous-time propagation channel [1] or from the effect of the antennas and the transmit and receive filters. Separating these two sources of correlation on the basis of measurements is difficult. Because the discrete-time model is commonly used for analysis and design of communication systems, and most often the channel taps are assumed uncorrelated [18], we consider the discrete-time effective channel only and do not attempt to assess the continuous-time US assumption.

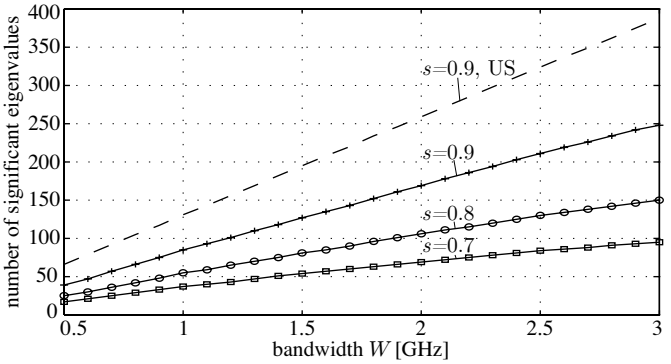

Fig. 8. MC II: Number of significant eigenvalues of $\widehat{\mathbf{K}}$ as a function of bandwidth, OLOS setting.

We decompose the random channel vector $\mathbf{h}$ according to $\mathbf{h} = \mathbf{m} + \tilde{\mathbf{h}}$, with $\mathbf{m} = \mathbb{E}[\mathbf{h}]$. Under the assumption that $\mathbf{h}$ is jointly complex Gaussian distributed with circularly symmetric $\tilde{\mathbf{h}}$, the joint distribution of $\mathbf{h}$ is completely specified through $\mathbf{m}$ and the covariance matrix $\mathbf{K} = \mathbb{E}[\tilde{\mathbf{h}}\tilde{\mathbf{h}}^H]$. To analyze A2, we need to estimate $\mathbf{K}$ for different bandwidths W. For each setting (LOS, OLOS, NLOS), we therefore truncated all measured impulse responses $\mathbf{h}_n$, $n = 1, 2, \ldots, N$ after L taps, where L depends linearly on the bandwidth ($L = 701$ at $W = 3$ GHz). The empirical $L \times L$ covariance matrix is given by $\widehat{\mathbf{K}} = 1/(N-1)\sum_{n=1}^{N}(\mathbf{h}_n - \widehat{\mathbf{m}})(\mathbf{h}_n - \widehat{\mathbf{m}})^H$, where $\widehat{\mathbf{m}} = 1/N\sum_{n=1}^{N}\mathbf{h}_n$. Accurate estimation of $\mathbf{K}$ requires a large number of samples, because L is large. MC I provides only $N = 90$ samples, so that we have to resort to MC II, with $N = 1011$ (LOS), $N = 2722$ (OLOS), and $N = 1256$ (NLOS) samples available. Because the channel can be modeled as Ricean in all three settings in MC II, we are effectively investigating the properties of the random part $\tilde{\mathbf{h}}$ of the channel vector $\mathbf{h}$. The concept of diversity order in such channels is well defined as the rank of $\mathbf{K}$ [41].[5]

A. Scaling of the Number of Stochastic Degrees of Freedom

The US assumption leads to a linear scaling of the number of "significant" eigenvalues of $\mathbf{K}$ as a function of W. We denote the kth normalized eigenvalue of $\widehat{\mathbf{K}}$ as $\hat{\lambda}_k$, with normalization $\sum_{k=1}^{L}\hat{\lambda}_k = 1$, and arrange the eigenvalues in decreasing order, i.e., $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \cdots \geq \hat{\lambda}_L$. We declare all eigenvalues $\hat{\lambda}_k$ with index $k \leq L_s$ to be significant, where L_s is the largest integer satisfying $\sum_{k=1}^{L_s}\hat{\lambda}_k \leq s$, with $0 \leq s \leq 1$. This criterion essentially measures the number of diversity branches with an effective branch receive SNR above a certain threshold,

[5]To be precise, the diversity order of a frequency-selective Ricean channel is given by $\mathrm{rank}(\mathbf{K})$ only if $\mathbf{m}$ lies in the range space of $\mathbf{K}$; otherwise, the diversity order is infinite [41].

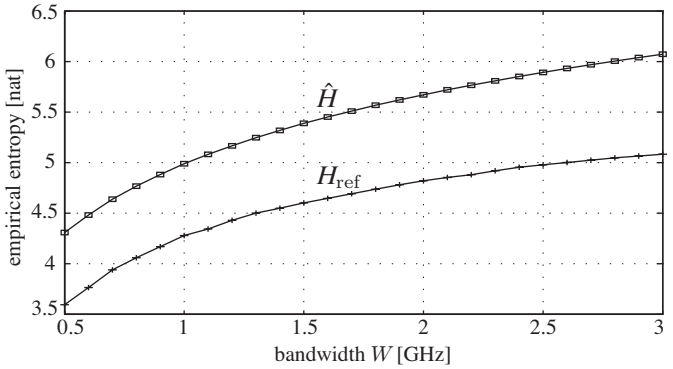

Fig. 9. MC II: $\hat{H}$ and H_{ref} as a function of bandwidth, OLOS setting.

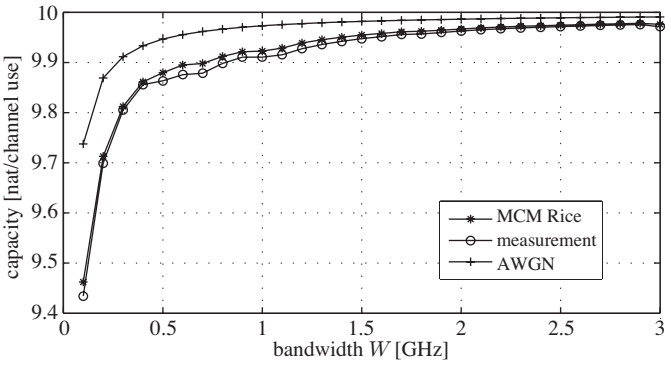

Fig. 10. Estimated capacity $\hat{C}$ according to (10) for MC II NLOS data, and Monte Carlo estimate of the capacity of the corresponding synthetic channel, $P/N_0 = 10\,\mathrm{dB}$.

and has already been used to determine the effective diversity order of channels at 900 MHz and 40 GHz [42]. For the OLOS setting, Fig. 8 depicts the scaling behavior of the eigenvalues as a function of W for $s = 0.9$, 0.8, and 0.7, and for a hypothetical discrete-time US channel with the same PDP as estimated from our data. The number of stochastic DOF scales approximately linearly for all three values of s. The corresponding observations for the LOS and NLOS settings are similar. For nonvanishing L/N, small $\hat{\lambda}_k$ are biased down and large $\hat{\lambda}_k$ are biased up [43], which could lead to a downward bias of L_s for large $L \propto W$. However, the linear scaling in Fig. 8 and the negligible bias-correction term of the eigenvalue estimator in [43] for our data suggest that N might be large enough to obtain sensible estimates for all bandwidths. To double-check our results, we used the model order selection criteria AIC and MDL [44], adapted to take into account the noise color introduced through the calibration process [45], and observed the same scaling behavior. Different from our results, a sublinear scaling was found in [15], with a measurement setup similar to our MC I. To check if the different scaling behavior results from sampling jitter introduced by the DSO, we used data from [15], and artificially added sampling jitter. We found that jitter with standard deviation of 10 ps, more than given in the DSO specifications, is too small to have any noticeable effect upon the number of significant eigenvalues. A fundamental cause for the difference in scaling behavior between MC II and [15] might be the different sources of channel variability discussed in Section VI-B; we do not have a physical explanation to support this conjecture, though.

B. Uncorrelated Scattering

Linear scaling of the number of significant eigenvalues with bandwidth W alone is not sufficient to conclude that the discrete-time US assumption holds. To answer this question, we estimated the normalized correlation coefficients between taps i and j as $\hat{\rho}_{ij} = [\widehat{\mathbf{K}}]_{ij}/([\widehat{\mathbf{K}}]_{ii}[\widehat{\mathbf{K}}]_{jj})^{1/2}$, with $i, j = 1, 2, \ldots, L$. On average, the correlation is small, around 0.1, but some taps show strong correlation of up to 0.7. In the NLOS setting, the correlation is, in general, somewhat higher than in the LOS and OLOS settings.

To quantify the intertap correlation, we evaluated the empirical entropy $\hat{H} = -\sum_{k=1}^{L} \hat{\lambda}_k \log \hat{\lambda}_k$ and compared it with $H_{\mathrm{ref}} = -\sum_{k=1}^{L} p_k \log p_k$, where p_k is obtained by uniformly sampling the empirical PDP at rate W after removing the mean component. For the discrete-time US assumption to be strictly satisfied, we need $\hat{H} = H_{\mathrm{ref}}$. Fig. 9 shows, however, that $\hat{H} < H_{\mathrm{ref}}$ with a gap of up to 1 nat. Furthermore, the curves for the measured channel and the hypothetical US channel in Fig. 8 also differ significantly. Hence, even though the intertap correlation is often low, its effect on the empirical entropy and the number of significant eigenvalues is not negligible, so that we cannot conclude that the taps are essentially uncorrelated.

VIII. Capacity Estimates

Statistical channel models are typically used to predict the performance of communication systems. Because the most fundamental performance measure is the channel capacity C, the goal of this section is to assess the approximation quality of the model (3) with complex Gaussian taps of nonzero mean w.r.t. the capacity directly estimated from MC II data. To this end, we used the Monte-Carlo Method (MCM) to synthesize complex-Gaussian-distributed impulse responses with either uncorrelated or correlated taps; mean and covariance in both cases were ML-estimated from MC II data. Throughout, we assume no channel state information (CSI) at the transmitter and perfect CSI at the receiver. Then, the mutual information of a channel with random transfer function $H(f)$ and additive white Gaussian noise is a random variable I, given as [46]

$$I = \int_{-W/2}^{W/2} \log(1 + \mathrm{SNR}\,|H(f)|^2) df \qquad \text{[nat/s]}. \qquad (9)$$

A. Ergodic Capacity

An estimate of the ergodic capacity $C = \mathbb{E}[I]$ on the basis of N samples of the discrete-time channel impulse response $\mathbf{h}$ can be obtained as [47], [48]

$$\hat{C} = \frac{1}{N} \sum_{n=1}^{N} \sum_{k=0}^{K-1} \log\Big(1 + \frac{P}{KN_0}\,|H_n[k]|^2\Big) \qquad \text{[nat/s]} \quad (10)$$

where $\mathbf{H}_n = \big[H_n[0]\, H_n[1] \ldots H_n[K-1]\big]^T$ is the length-K DFT of the nth channel vector sample $\mathbf{h}_n$. The average transmit power is constrained to P, allocated uniformly over all

K parallel channels, and N_0 is the noise variance per channel. This estimate does not assume a specific distribution of the taps, nor does it impose any correlation structure between taps. We computed the capacity estimate (10) from the measured and the synthesized impulse responses as a function of bandwidth W. All impulse responses were normalized to unit average power and truncated after $L = 701$ taps. The result for the NLOS setting, along with the capacity of the AWGN channel with the same receive SNR, is shown in Fig. 10, where we set $P/N_0 = 10\,\mathrm{dB}$, and used $K = 5608$ parallel channels at $W = 3\,\mathrm{GHz}$. The channel synthesized according to the uncorrelated Ricean model predicts the ergodic capacity of the measured channel very accurately, with a maximum mean squared error at $P/N_0 = 10\,\mathrm{dB}$ of less than 0.3% in the NLOS setting, less than 0.2% in the OLOS setting, and less than 0.07% in the LOS setting.

B. Outage Capacity

For slow-fading channels [18], the ϵ-outage capacity C_{out}, defined as $\mathbb{P}(I < C_{\mathrm{out}}) \leq \epsilon$, is a more sensible performance measure than the ergodic capacity [46]. Fig. 11 shows empirical CDFs of I for the measured LOS channel and two synthetic channels. One of the synthetic channels takes into account intertap correlation according to $\widehat{\mathbf{K}}$. Measured and synthetic channels behave quite differently at low outage probabilities. For the measured channel, the large change in mutual information from 8.8 nat to 9.3 nat at nearly constant outage of 0.12 indicates a reduction in total received power for some channel realizations, which is most likely a result of shadowing. Although it is common practice in channel modeling to separate shadowing and small-scale fading, this separation is not an intrinsic property of the channel, but a modeling assumption. Fig. 11 shows that this assumption might not always be valid for the measured channels. A likely reason for the drop in received power in some channel realizations is that persons moving in the lobby sometimes block the LOS or some other propagation paths that convey a significant fraction of the total energy. We used the synthesized channel impulse responses to simulate the shadowing effect by attenuating several taps in a fraction of all generated channel realizations. It turned out that attenuating the strongest four taps is already sufficient to obtain a behavior similar to the measured one. The distribution of the multiplicative shadowing random variable was bimodal, to model presence or absence of a person blocking a dominant propagation path; this is in stark contrast to the commonly used lognormal distribution for shadow fading. Scatter plots for individual taps support the finding that only a small number of taps are affected by shadowing. The AIC-best distribution for these taps is most often the Weibull distribution, consistent with our findings in Section VI-B. Shadowing has an even larger impact on the CDF of the mutual information in the NLOS setting, because the overall received SNR is lower, while the impact in the OLOS setting is smaller than in the LOS setting depicted in Fig. 11.

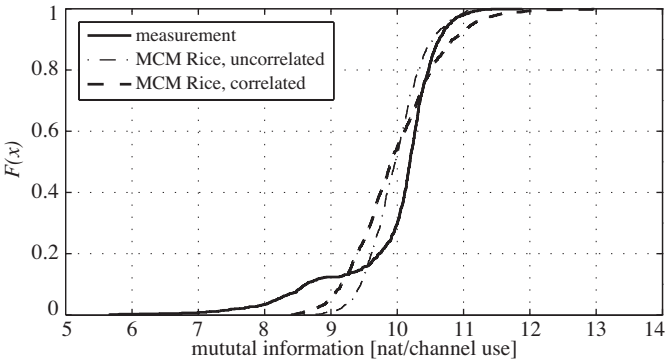

Fig. 11. Empirical CDF of the mutual information for MC II LOS data, and Monte Carlo estimate of the CDF for the corresponding synthetic channels, $P/N_0 = 10\,\mathrm{dB}$.

IX. Conclusion

On the basis of indoor UWB channel measurements in the frequency band from 2 GHz to 5 GHz, we found that Akaike's Information Criterion supports the Rayleigh, respectively Rice, tap amplitude distribution, except for taps in the vicinity of cluster peaks in MC I and taps with shadowed specular reflections in MC II. This result is somewhat surprising, as it is often argued that for large bandwidths the number of partial waves contributing to each tap is not high enough to invoke the central limit theorem, which would justify the complex Gaussian assumption underlying the Rayleigh and Rice distributions. We have also demonstrated that the differences (in the sense of AIC) between the Rayleigh distribution and the Rice, Nakagami, and Weibull distributions in MC I, and between the Rice and the Weibull distribution in MC II are often minor. Although we could not verify if the assumption of a *jointly* Gaussian distribution of the channel vector $\mathbf{h}$ holds, we advocate the use of the complex Gaussian distribution, because it is mathematically tractable, and we did not find evidence against it. The differences in the mean component between MC I and MC II show that it is important to determine if a UWB communication system will operate with mobile or fixed terminals.

Our second main result is the observation that the number of stochastic DOF scales approximately linearly with bandwidth. Consequently, the diversity order of the channel shows the same scaling behavior, a common assumption in information-theoretic studies of UWB systems. Nevertheless, there seems to be residual correlation between the individual channel taps, which indicates that the discrete-time US assumption does not hold strictly. While our primary focus was on the small-scale fading behavior, analysis of the outage capacity showed that shadowing cannot always be completely separated from small-scale fading in MC II. Shadowing affects a small but important set of taps, which correspond to the LOS or specular reflections. As these taps convey a large part of the total energy, they might possibly determine the overall performance of a UWB communication system with static terminals.

Acknowledgment

The authors would like to thank G. Durisi, who was of great help in MC II, H.-R. Benedickter and C. Schmid for their support regarding RF measurement and instrumentation issues, H.-R. Künsch and his staff for helpful discussions, and A. Menouni Hayar, J. Kunisch, and J. Pamp for access to their measurement data.

REFERENCES

[1] R. Vaughan and J. B. Andersen, *Channels, Propagation and Antennas for Mobile Communications*. London, U.K.: The Institution of Electrical Engineers, 2003.

[2] J. R. Pierce, "Ultimate performance of M-ary transmission on fading channels," *IEEE Trans. Inf. Theory*, vol. 12, no. 1, pp. 2–5, Jan. 1966.

[3] A. J. Viterbi, "Performance of an M-ary orthogonal communication system using stationary stochastic signals," *IEEE Trans. Inf. Theory*, vol. 13, no. 3, pp. 414–422, July 1967.

[4] I. E. Telatar and D. N. C. Tse, "Capacity and mutual information of wideband multipath fading channels," *IEEE Trans. Inf. Theory*, vol. 46, no. 4, pp. 1384–1400, July 2000.

[5] S. Verdú, "Spectral efficiency in the wideband regime," *IEEE Trans. Inf. Theory*, vol. 48, no. 6, pp. 1319–1343, June 2002.

[6] M. Médard and R. G. Gallager, "Bandwidth scaling for fading multipath channels," *IEEE Trans. Inf. Theory*, vol. 48, no. 4, pp. 840–852, Apr. 2002.

[7] V. G. Subramanian and B. Hajek, "Broad-band fading channels: signal burstiness and capacity," *IEEE Trans. Inf. Theory*, vol. 48, no. 4, pp. 809–827, Apr. 2002.

[8] W. R. Braun and U. Dersch, "A physical mobile radio channel model," *IEEE Trans. Veh. Technol.*, vol. 40, no. 2, pp. 472–482, May 1991.

[9] A. F. Molisch, J. R. Foerster, and M. Pendergrass, "Channel models for ultrawideband personal area networks," *IEEE Wireless Commun. Mag.*, vol. 10, no. 6, pp. 14–21, Dec. 2003.

[10] D. Cassioli, M. Z. Win, and A. F. Molisch, "The ultra-wide bandwidth indoor channel: From statistical models to simulations," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 6, pp. 1247–1257, Aug. 2002.

[11] J. Keignart and N. Daniele, "Channel sounding and modelling for indoor UWB communications," in *Proc. Int. Workshop Ultra-Wideband Syst. (IWUWBS)*, Oulu, Finland, June 2003.

[12] P. Pagani and P. Pajusco, "Experimental assessment of the UWB channel variability in a dynamic indoor environment," in *Proc. IEEE Int. Symp. Personal, Indoor, Mobile Radio Commun. (PIMRC)*, Barcelona, Spain, Sept. 2004, pp. 2973–2977.

[13] S. S. Ghassemzadeh, R. Jana, C. W. Rice, W. Turin, and V. Tarokh, "Measurement and modeling of an ultra-wide bandwidth indoor channel," *IEEE Trans. Commun.*, vol. 52, no. 10, pp. 1786–1796, Oct. 2004.

[14] J. Kunisch and J. Pamp, "Measurement results and modeling aspects for the UWB radio channel," in *IEEE Conf. Ultra Wideband Syst. Technol. (UWBST) Dig. Tech. Papers*, Baltimore, MD, U.S.A., May 2002, pp. 19–23.

[15] A. Menouni Hayar, R. Knopp, and R. Saadane, "Subspace analysis of indoor UWB channels," *EURASIP J. Appl. Signal Process.*, vol. 2005, no. 3, pp. 287–295, Mar. 2005.

[16] D. A. McNamara, C. W. I. Pistorius, and J. A. G. Malherbe, *Introduction to the Uniform Geometrical Theory of Diffraction*. Boston, MA, U.S.A.: Artech House, 1990.

[17] R. C. Qiu, "A study of the ultra-wideband wireless propagation channel and optimum UWB receiver design," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 9, pp. 1628–1637, Dec. 2002.

[18] D. N. C. Tse and P. Viswanath, *Fundamentals of Wireless Communications*. Cambridge, U.K.: Cambridge Univ. Press, 2005.

[19] H. Linhart and W. Zucchini, *Model Selection*. New York, NY, U.S.A.: Wiley, 1986.

[20] R. B. D'Agostino and M. A. Stephens, Eds., *Goodness-of-Fit Techniques*. New York, NY, U.S.A.: Marcel Dekker, 1986.

[21] H. Hashemi, "Impulse response modeling of indoor radio propagation channels," *IEEE J. Sel. Areas Commun.*, vol. 11, no. 7, pp. 967–978, Sept. 1993.

[22] L. Guttman, "What is not what in statistics," *The Statistician*, vol. 26, no. 2, pp. 81–107, June 1977.

[23] H. Akaike, "Information theory and an extension of the maximum likelihood principle," in *Breakthroughs in Statistics*, S. Kotz and N. L. Johnson, Eds. New York, NY, U.S.A.: Springer, 1992, vol. 1, pp. 610–624, originally published in *Proc. Int. Symp. Inf. Theory*, Budapest, Hungary, 1973.

[24] M. A. Taneda, J. Takada, and K. Araki, "The problem of the fading model selection," *IEICE Trans. Commun.*, vol. E84-B, no. 3, pp. 660–666, Mar. 2001.

[25] M. H. Hansen and B. Yu, "Model selection and the principle of minimum description length," *J. Am. Statist. Assoc.*, vol. 96, no. 454, pp. 746–774, June 2001.

[26] J. Shao, "An asymptotic theory for linear model selection," *Statistica Sinica*, vol. 7, pp. 221–264, 1997.

[27] Y. Yang, "Can the strengths of AIC and BIC be shared? A conflict between model identification and regression estimation," *Biometrika*, vol. 92, no. 4, pp. 937–950, Dec. 2005.

[28] K. P. Burnham and D. R. Anderson, *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*, 2nd ed. New York, NY, U.S.A.: Springer, 2002.

[29] H. Akaike, "On the likelihood of a time series model," *The Statistician*, vol. 27, no. 3/4, pp. 217–235, Dec. 1978.

[30] U. G. Schuster and H. Bölcskei, "UWB measurement database," Mar. 2005. [Online]. Available: www.nari.ee.ethz.ch/commth/uwb-measurements

[31] A. M. Street, L. Lukama, and D. J. Edwards, "Use of VNAs for wideband propagation measurements," *IEE Proc. Commun.*, vol. 148, no. 6, pp. 411–415, Dec. 2001.

[32] Q. Li and S. Wong, "Measurement and analysis of the indoor UWB channel," in *Proc. IEEE Veh. Technol. Conf. (VTC)*, Orlando, FL, U.S.A., Oct. 2003, pp. 1–5.

[33] P. J. Cullen, P. C. Fannin, and A. Molina, "Wideband measurement and analysis techniques for the mobile radio channel," *IEEE Trans. Veh. Technol.*, vol. 42, no. 4, pp. 589–603, Nov. 1993.

[34] J. D. Parsons, *The Mobile Radio Propagation Channel*, 2nd ed. Chichester, U.K.: Wiley, 2000.

[35] M. Nakagami, "The m-distribution—a general formula of intensity distribution of rapid fading," in *Statistical Methods in Radio Wave Propagation*, W. C. Hoffman, Ed. New York, NY, U.S.A.: Pergamon, 1960, pp. 3–36.

[36] R. M. Buehrer, W. A. Davis, A. Safaai-Jazi, and D. Sweeney, "Characterization of the ultra-wideband channel," in *IEEE Conf. Ultra Wideband Syst. Technol. (UWBST) Dig. Tech. Papers*, Reston, VA, U.S.A., Nov. 2003, pp. 26–31.

[37] J. R. Foerster, "Channel modeling sub-committee report final," IEEE 802.15 SG3a, Tech. Rep. P802.15 02/490r1, Feb. 2003.

[38] H. Hashemi, M. McGuire, T. Vlasschaert, and D. Tholl, "Measurement and modeling of temporal variations of the indoor radio propagation channel," *IEEE Trans. Veh. Technol.*, vol. 43, no. 3, pp. 733–737, Aug. 1994.

[39] C. Tepedelenlioğlu, A. Abdi, and G. B. Giannakis, "The Ricean K factor: Estimation and performance analysis," *IEEE Trans. Wireless Commun.*, vol. 2, no. 4, pp. 799–810, July 2003.

[40] R. J. C. Bultitude, S. A. Mahmoud, and W. A. Sullivan, "A comparison of indoor radio propagation characteristics at 910 MHz and 1.75 GHz," *IEEE J. Sel. Areas Commun.*, vol. 7, no. 1, pp. 20–30, Jan. 1989.

[41] R. U. Nabar, H. Bölcskei, and A. J. Paulraj, "Diversity and outage performance in space-time block coded Ricean MIMO channels," *IEEE Trans. Wireless Commun.*, vol. 4, no. 5, pp. 2519–2532, Sept. 2005.

[42] F. Patenaude, J. Lodge, and J.-Y. Chouinard, "Eigen analysis of wide-band fading channel impulse responses," *IEEE Trans. Veh. Technol.*, vol. 48, no. 2, pp. 593–606, Mar. 1999.

[43] O. Ledoit and M. Wolf, "A well-conditioned estimator for large-dimensional covariance matrices," *SIAM J. Math. Anal.*, vol. 88, no. 2, pp. 365–411, Feb. 2004.

[44] M. Wax and T. Kailath, "Detection of signals by information theoretic criteria," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 33, no. 2, pp. 387–392, Apr. 1985.

[45] L. C. Zhao, P. R. Krishnaiah, and Z. D. Bai, "On detection of the number of signals when the noise covariance matrix is arbitrary," *J. Multivariate Anal.*, vol. 20, no. 1, pp. 26–49, Oct. 1986.

[46] E. Biglieri, J. Proakis, and S. Shamai (Shitz), "Fading channels: Information-theoretic and communications aspects," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2619–2692, Oct. 1998.

[47] W. Hirt and J. L. Massey, "Capacity of the discrete-time Gaussian channel with intersymbol interference," *IEEE Trans. Inf. Theory*, vol. 34, no. 3, pp. 380–388, May 1988.

[48] H. Bölcskei, D. Gesbert, and A. J. Paulraj, "On the capacity of OFDM-based spatial multiplexing systems," *IEEE Trans. Commun.*, vol. 50, no. 2, pp. 225–234, Feb. 2002.

Ulrich G. Schuster (S'03) studied electrical engineering at Aachen University of Technology, Aachen, Germany, and as a J. William Fulbright Scholar at the University of California at Berkeley, U.S.A., where he received the M.S. degree in 2003. Since then, he has been a research assistant at ETH Zurich, working towards the Dr. sc. degree. Mr. Schuster was a visiting researcher at the Centre for Wireless Communications at the University of Oulu, Finland, in 2004, and at Princeton University, Princeton, NJ, U.S.A., in 2006. His research interests are in the field of communication and information theory.

Helmut Bölcskei (M'98–SM'02) was born in Austria on May 29, 1970, and received the Dipl.-Ing. and Dr. techn. degrees in electrical engineering/communication theory from Vienna University of Technology, Vienna, Austria, in 1994 and 1997, respectively. From 1994 to 1998 he was with Vienna University of Technology. From 1999 to 2001 he was a postdoctoral researcher in the Information Systems Laboratory, Department of Electrical Engineering, Stanford University, Stanford, CA. He was in the founding team of Iospan Wireless Inc., a Silicon Valley-based startup company (acquired by Intel Corporation in 2002) specialized in multiple-input multiple-output (MIMO) wireless systems for high-speed Internet access. From 2001 to 2002 he was an Assistant Professor of Electrical Engineering at the University of Illinois at Urbana-Champaign. He has been with ETH Zurich since 2002, where he is Professor of Communication Theory. He was a visiting researcher at Philips Research Laboratories Eindhoven, The Netherlands, ENST Paris, France, and the Heinrich Hertz Institute Berlin, Germany. His research interests include communication and information theory with special emphasis on wireless communications, signal processing and quantum information processing.

He received the 2001 IEEE Signal Processing Society Young Author Best Paper Award, the 2006 IEEE Communications Society Leonard G. Abraham Best Paper Award, the ETH "Golden Owl" Teaching Award, and was an Erwin Schrödinger Fellow (1999–2001) of the Austrian National Science Foundation (FWF). He was a plenary speaker at several IEEE conferences and served as an associate editor of the *IEEE Transactions on Signal Processing*, the *IEEE Transactions on Wireless Communications*, and the *EURASIP Journal on Applied Signal Processing*. He is currently on the editorial board of "Foundations and Trends in Networking," serves as an associate editor for the *IEEE Transactions on Information Theory*, and is TPC co-chair of the 2008 IEEE International Symposium on Information Theory.