Dynamic Spectrum Access in WLAN Channels: Empirical Model and Its Stochastic Analysis*

Stefan Geirhofer and Lang Tong
School of Electrical and Computer Engineering
Cornell University, Ithaca, NY 14853
{sg355, lt35}@cornell.edu

Brian M. Sadler
Army Research Laboratory
Adelphi, MD 20783-1197
bsadler@arl.army.mil

ABSTRACT

In this work we are concerned with dynamically sharing the spectrum in the time-domain by exploiting whitespace between the bursty transmissions of a primary user, represented by an 802.11b-based wireless LAN (WLAN). For deriving such schemes we need to establish a model of the WLAN's medium access as to predict its behavior accurately. Moreover, a balance between accuracy and complexity needs to be struck as to render the model useful in practice. We emphasize that our model is based on actual measurements at 2.4 GHz using a vector signal analyzer.

We have shown previously that a semi-Markov model is a viable approach for modeling the busy/idle durations. In the present paper we extend our results by (i) expanding the measurement setup and looking at more realistic traffic scenarios, (ii) providing a better approximation to the distribution of the idle durations, and (iii) fitting a phase-type approximation to arrive at a computationally simpler description. The goodness-of-fit of the proposed models is evaluated using the Kolmogorov-Smirnov test.

Categories and Subject Descriptors

C.3 [**Special Purpose and Application-Based Systems**]: Signal Processing Systems; I.6.4 [**Simulation and Modeling**]: Model Validation and Analysis

General Terms

Design, Measurement, Verification.

Keywords

WLAN Modeling, Dynamic Spectrum Access, Coexistence.

*Prepared though collaborative participation in the Communications and Networks Consortium sponsored by the U.S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-2-0011. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

1. INTRODUCTION

The spectrum relevant to wireless communications has gradually become a scarce resource, making it unlikely that regulators will be able to meet the bandwidth needs of emerging technologies. Actual measurements show, however, that at most times and locations, the spectrum is only lightly used. In fact, a typical utilization of only several percent or even less is reported [1]. This figure illustrates the weakness of static frequency allocations and has given rise to envisioning schemes that allow so-called secondary-users to access certain frequency bands dynamically, provided that they cause no (significant) interference to the actual licensees (referred to as primary users).

This paper investigates dynamic spectrum access in the time-domain by reusing whitespace between the bursty transmissions of a primary user [2], represented by a WLAN in the scope of this work. A reliable model of the primary user's channel access is indispensable if access strategies for the secondary user are to be found.

Dynamic spectrum access in the time-domain relies on the existence of sufficient whitespace between bursty transmissions. Indeed, we believe that this is the case in many practical scenarios. For illustration consider the baseband signal recorded by a signal analyzer for a Voice-over-IP conference session over WLAN shown in Fig. 1 (detailed setup parameters are discussed later). The packet transmissions can easily be discerned from the noise floor leaving large amounts of white space. In fact, the channel was found to be idle for 89% of the time.

This remaining whitespace can be used in numerous appli-

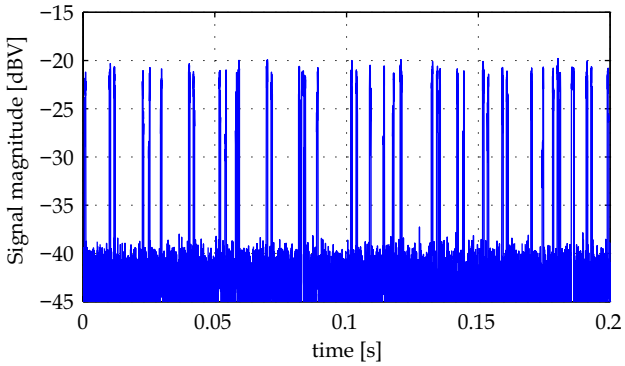

Figure 1: Baseband signal of a WLAN supporting a VoIP conference session.

cations. Consider wireless sensor networks whose heterogenous deployment and sporadic transmissions make an allocation of separate frequency bands unlikely. Nevertheless, given that most sensor networks communicate at low-rates, reusing the whitespace of other applications appears reasonable. Furthermore, our model can be applied to reduce interference between different standards in shared bands. Consider the coexistence of WLAN and Bluetooth in the 2.4 GHz ISM band as an important example.

1.1 Main Contribution

In this paper we propose a model for the busy and idle durations of a WLAN. Our contribution includes both a measurement component as well as the statistical analysis of the empirical data.

In terms of the measurement setup we present an antenna-based testbed as well as an isolated RF-setup, which guarantees our measurements to be free of interference. We can thus ensure that the collected data has not been corrupted by WLAN traffic in adjacent areas.

Furthermore, this paper addresses realistic traffic scenarios (FTP, Voice-over-IP) which are likely to arise in practice and compares these to synthetic traffic originating from a traffic generator.

In a previous paper [3] we have shown that a semi-Markov process is a viable model for our setup. In the present paper we significantly extend our results by considering a mixture distribution to differentiate between the idle periods due to the transmission standard (inter-frame spaces, contention window) and a truly unused ('free') channel. We show that augmenting the model in this way provides for a significantly better fit with the empirical data.

While the mixture distribution fits the data accurately, it might render subsequent analysis cumbersome given that heavy-tailed distributions are usually hard to analyze analytically [4]. As a consequence, we also fit a phase-type (hyper-Erlang) distribution, which has the nice property that it corresponds to the time-to-absorption in a continuous-time Markov process (CTMP) and usually allows for more tractable results.

The proposition of the above models is justified by employing the Kolmogorov-Smirnov test, showing an excellent goodness-of-fit.

1.2 Related Work

Dynamic spectrum access is an emerging area that has recently received considerable attention, both in terms of theoretical contributions as well as practical testbeds and implementations [5]. Sparked by projects such as the DARPA XG program [6] or the European DRiVE Project [7] various methodologies for efficiently sharing the spectrum have been explored. The majority of proposed techniques can be classified according to whether the spectrum is shared in the spatial or temporal domain, with each area having its own practical limitations and challenges.

In the spatial domain, the reliable sensing of primary users is of main concern, requiring the detection of very weak signals [8] possibly through cooperation among secondary users [9].

In the temporal domain, the main challenge is to identify spectral opportunities between bursty transmissions by predicting the primary user's medium access, and finding control schemes that efficiently exploit the remaining whitespace. In [2, 10] such optimal control policies are discussed within a partially-observable decision framework, relying on a Markovian assumption for the primary user's behavior.

2. MEASUREMENT SETUP

In this work we identify the primary user with an 802.11b based WLAN operating in the 2.4 GHz ISM band. Our results should, however, extend to other primary user's protocols given that they use multiaccess protocols similar to CSMA/CA (as used in the 802.11 standards). Different from related publications that capture packets by commerical WLAN adapter cards operating in a special mode we employ a vector signal analyzer to record raw complex baseband data which is subsequently processed to find the start and end times of packets. This approach guarantees an accurate and verifiable characterization of the channel's idle and busy periods.

For recording the baseband data we used an Agilent 89640A vector signal analyzer (VSA) [11] which internally downconverted the RF signals to an intermediate frequency and then was configured to sample at a rate of 44 MHz.

In this work we consider both a WLAN communicating via antennas, as well as an RF-isolated setup that guarantees our measurements to be free of interference from other devices operating in adjacent frequency bands. The setup is illustrated in Fig. 2 and Fig. 3, respectively.

2.1 Antenna-based setup

The antenna-based propagation setup consists of a Netgear WGT624 wireless router and three computers with wireless adapter cards (two Netgear WG311T and one WG-511T; cf. Fig. 2). The setup operated in Channel 11, which represents a 22 MHz frequency band centered at 2.462 GHz. All the equipment was located in the same room, resulting in a high-SNR setup with no hidden terminals. Using the VSA, we verified that interference from adjacent channels was minimum although a completely interference-free setup could not be guaranteed.

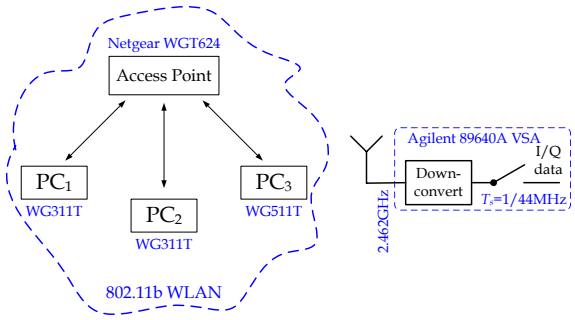

Figure 2: Antenna-based measurement setup.

2.2 Isolated RF-setup

Besides the antenna-based setup, we also considered the isolated RF-setup shown in Fig. 3. It consists of a Linksys WRT54GC wireless router and three workstations with Netgear WG311T wireless adapter cards. All the devices are connected to a Broadwave Technologies resistive power divider via RG174U coaxial cables and SMA connectors. The VSA is also connected to the divider resulting in a fully isolated setup. Strictly speaking there is still some residual

interference that couples directly via the workstations into the wireless adapter cards. However, given that all devices are connected with coaxial cables this interference is very small compared to the desired signal and can be neglected.

The Netgear router used for the antenna-based setup could not be used for the isolated measurements as well since its built-in antenna was non-detachable. The use of two different routers caused our setup to differ in terms of the type of synchronization preamble used. While the Netgear router could be configured to use only long-synchronization preambles, the Linksys router did not allow for specifying this option. As a consequence most of the time a short preamble was transmitted (given the high SNR setup). While this leads to slightly different packet durations, the qualitative behavior of our results remained unaltered.

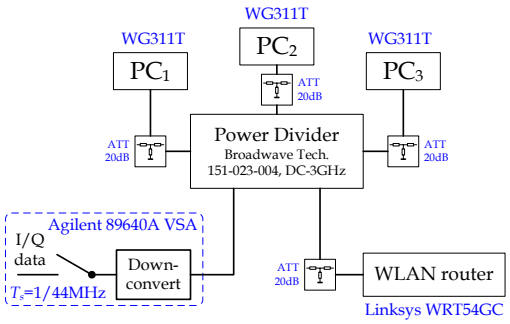

Figure 3: Isolated measurement setup.

2.3 Traffic generation

Each of the workstations was used to generate traffic using the Distributed Internet Traffic generator (D-ITG) [12]. The software allows for a flexible statistical characterization of the traffic, including varying packet lengths and inter-departure times. A detailed specification of the settings is provided with the measurement results in Sec. 4.

Additionally, we also investigate typical usage scenarios of WLAN by using the popular "Skype" voice-over-IP (VoIP) client to set up a conference call within the WLAN, using the traffic generator to simulate G.711 codec based voice communication, and using an SFTP client to download files from a central server. A detailed treatment of the results is again deferred to Sec. 4.

3. SENSING STRATEGIES

The two measurement setups described in the last section both yield time captures of the complex baseband signal. Given these data, we process the signals to determine the exact start and end of each packet. Clearly, this fully determines the channel's idle/busy durations.

We consider two different sensing strategies depending on whether the transmission standard of the primary user is assumed to be known. In the former case, the detection of the packets shall be based on energy. In the latter case we can exploit the standard-specifics to achieve better performance [3].

3.1 Energy-based detection

If the primary user's transmission standard is unknown, a natural approach for detecting the start and end of packets is based on the transmitted energy. In order to achieve satisfactory performance we consider blocks of N samples whose length is shorter than the smallest packet length [3]. The detection problem can then be formulated as

$$\mathcal{H}_0 \quad : \quad Y_i = V_i, \ i = 1, \ldots, N \tag{1}$$

$$\mathcal{H}_1 \quad : \quad Y_i = S_i + V_i, \ i = 1, \ldots, N, \tag{2}$$

where Y_i denotes the complex baseband samples, V_i are noise samples, $V_i \sim \mathcal{CN}(0, \sigma_0^2)$, and S_i denotes the signal samples drawn from a complex Gaussian, $S_i \sim \mathcal{CN}(0, \sigma_1^2)$. Lacking any information on the transmission standard of the primary user, the Gaussian assumption for S_i appears reasonable.

The hypothesis testing problem defined above is standard [13] and the optimal Neyman-Pearson detector is given by

$$T(\mathbf{y}) = \sum_{i=1}^{N} |y_i|^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma, \tag{3}$$

where the threshold γ is determined according to the probability of false alarm, which amounts to

$$\alpha = \Pr(T(\mathbf{y}) > \gamma | \mathcal{H}_0) = 1 - \tilde{\Gamma}_r(N, \frac{\gamma}{\sigma_0^2}), \tag{4}$$

where

$$\tilde{\Gamma}_r(N, \xi) = \frac{1}{\Gamma(N)} \int_0^{\xi} t^{N-1} e^{-t} dt \tag{5}$$

is the regularized gamma function and $\Gamma(N)$ is the complete gamma function. Similarly, the power of the detector is given by

$$\beta = \Pr(T(\mathbf{y}) > \gamma | \mathcal{H}_1) = 1 - \tilde{\Gamma}_r(N, \frac{\gamma}{\sigma_0^2 + \sigma_1^2}). \tag{6}$$

The above expressions show that the detection performance depends on the SNR $= \sigma_1^2/\sigma_0^2$ as well as the block length N. For our setup we chose $N = 44$ samples, which corresponds to 1 µs long blocks. If we demand $\alpha = 1 - \beta < 10^{-5}$ then we can see that we have to guarantee that the SNR is above 4.29 dB which is easily met in our setup.

Finally, it has to be noted that the Gaussian assumption for the noise V_i might not be appropriate if significant interference occurs. Indeed, this might be a limiting factor if we consider that the WLAN channels are partially overlapping. Suppressing this interference by a filter may thus be necessary in practice.

3.2 Feature-based detection

The energy-based detection scheme described in the last section is based on the assumption that the primary user's transmission standard is unknown. In some applications, however, it is reasonable to assume that the transmission specifics are known to the primary user. This knowledge can in turn be exploited to improve the detection of packets.

The layout of an 802.11b physical layer (PHY) frame is shown in Fig. 4. It consists of a PLCP preamble, split into a block of scrambled '1's ('0's for the short-preamble) and the start-of-frame delimiter (SFD) indicating the beginning of the PLCP header. The SFD can be used to precisely detect the start of the packet. The information provided in the header consists of a SIGNAL, SERVICE, and LENGTH field as well as a CRC protecting these three blocks.

From our viewpoint the SFD and the LENGTH field are most interesting; the former determines the start of the

packet while the latter provides the duration (and thus the end) of the packet.

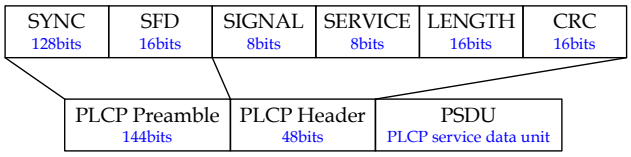

Figure 4: Physical layer preamble in 802.11b (long preamble).

The receive processing for the feature-based detection scheme is depicted in Fig. 5. The complex baseband data collected at a rate of 44 MHz is first passed through a Gaussian pulse shaping filter with a bandwidth-symbol time product of $BT_s = 1/2$. In order to obtain chip-synchronization the filtered signal is correlated with the 11-sample Barker sequence specified by the standard [14]. The resulting signal shows periodic peaks whenever the spreading sequence lines up with the input signal. We detect these peaks and downsample the signal to the symbol rate of 11 Mbps. Subsequently, we despread and demodulate the DBPSK/DQPSK encoded preamble. The frequency offset at the receiver is noticeable but can be neglected since the signals are differentially encoded. After successful decoding, the resulting bit stream is descrambled and the start-of-frame delimiter (SFD) is detected. In the same way, the SIGNAL, SERVICE, and LENGTH field are extracted and the CRC check is performed to ensure that the extracted information is correct.

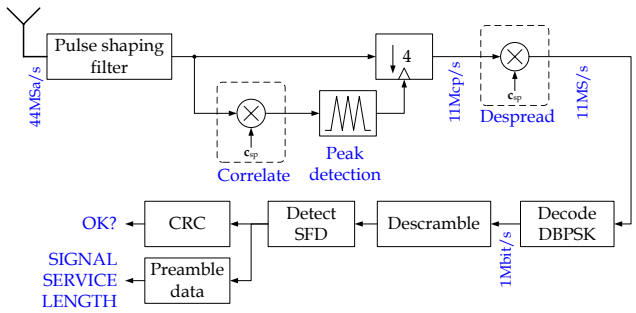

Figure 5: Receive processing for feature-based detection.

4. MEASUREMENT RESULTS

In this section we present the measurement results for the statistics of the busy/idle durations of the channel. We investigate different traffic scenarios as pointed out in Sec. 2.3. In particular, we first consider constant length UDP traffic with Poisson distributed inter-arrival times. This allows us to parameterize the 'business' of the channel by increasing the rate parameter λ of this distribution. Second, we consider FTP and Voice-over-IP traffic to investigate whether our idealized setup extends to practical traffic scenarios. For a better understanding of the results we start this section with a brief illustration of WLAN's medium access in order to keep this paper self-contained.

4.1 WLAN Medium Access Protocol

The 802.11 standard for WLAN [14,15] uses Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) to control the station's access to the medium (cf.Fig. 6). This implies that before transmitting a packet, the station has to first sense the medium.

If the channel is free, the station continues sensing for the distributed coordination function inter-frame space (DIFS). If the channel remains idle during the entire period, the station can go ahead and start transmitting.

After a packet transmission, the receiver has to confirm reception immediately by transmitting an acknowledgement. Only a short inter-frame space (SIFS) is necessary as to give priority to the (required) transmission of acknowledgements (cf. Fig. 6).

If the channel is busy in the first place the station has to defer access until the medium becomes idle again. Then, after a DIFS, a contention window is used to avoid collision between the multiple stations trying to access the medium. Specifically, each station generates a uniform random number $i \in \{0, \ldots, 31\}$ and defers transmission for $iT_{\text{slot}} = i \cdot 20\,\mu$s before accessing the channel (given that no other station has already started to access the channel before).

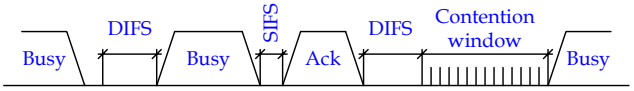

Figure 6: Medium access in an 802.11b-based WLAN.

The standard provides some more technical details that are not addressed above. In particular, if collisions occur the length of the contention window is increased. These specifics, however, do not manifest themselves in our measurement results and shall thus not be addressed in this paper.

4.2 Measurement validation

We first look at a simple measurement scenario to further illustrate the specifics of the medium access and to validate our measurement setup. In particular, we consider the isolated measurement setup depicted in Fig. 3 with only one PC and the wireless router turned on (the other ports of the resistive power divider we terminated to eliminate reflections). The traffic generator was then used to generate UDP packets of constant length 512 B with constant inter-arrival times at a rate of 10^5pkts/s. This rate is too high to be transmitted across the channel but ensures that the workstation's transmit buffer is never empty.

Using the setup described above we used the VSA to capture 100 blocks of complex baseband data, of duration 0.25 s each. The blocks were then processed using both sensing strategies discussed in Sec. 3. The results of energy- and feature-based detection match nicely leading to the histograms for the busy/idle durations shown in Fig. 7.

The histograms indeed reflect the characteristics of the standard. First, the histogram of the busy durations depicted in Fig. 7(a) shows only three components, corresponding to the transmission of acknowledgement packets ($t \approx 0.11$ ms), data packets ($t \approx 0.51$ ms) and beacon frames ($t \approx 0.76$ ms), respectively. Given that we forced the data

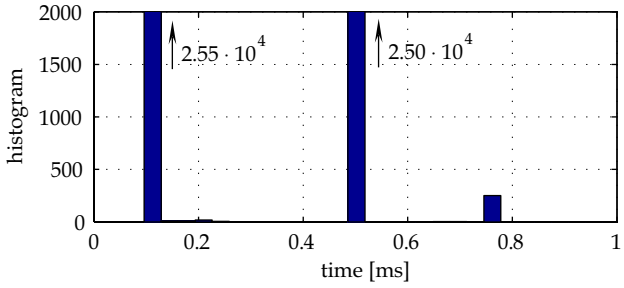

(a) Histogram for the busy durations.

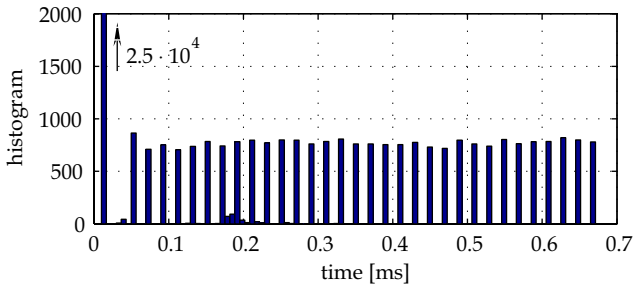

(b) Histogram for the idle durations.

Figure 7: Measurement validation using 1PC (cf. Sec. 4.2)

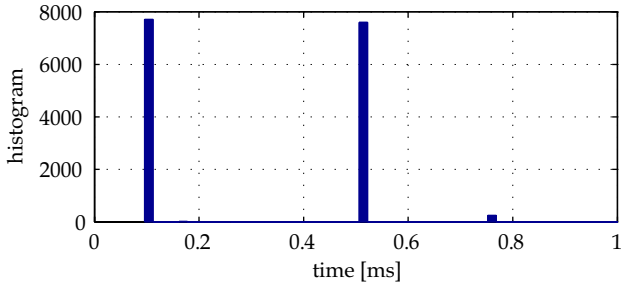

(a) Histogram for the busy durations.

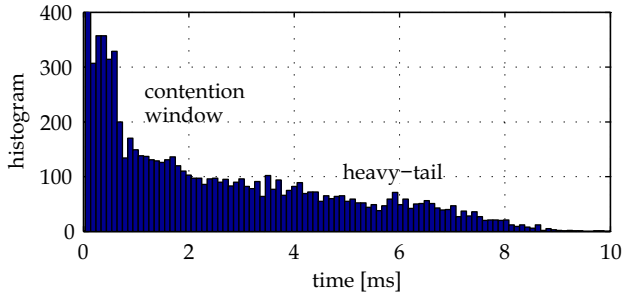

(b) Histogram for the idle durations.

Figure 8: Histograms for the UDP traffic scenario (cf. Sec. 4.3)

packets to be of constant length, this result is in accordance with our expectations.

The histogram of the idle durations reflects the standard as well. We see a discrete component at $t \approx 10\,\mu s$, which nicely corresponds to the SIFS. Furthermore, we see 32 discrete components, each spaced $20\,\mu s$ apart. These correspond to the contention window as described in Sec. 4.1.

4.3 UDP Traffic with Poisson Inter-arrivals

In the last section we have validated the measurement setup using a simplified traffic scenario. In this section we are now using all three workstations together with the wireless router (cf. Fig. 2 and Fig. 3). The traffic generator was again used to generate UDP packets of constant length of 512 B but the inter-arrival rates for each workstation were now drawn from independent Poisson distributions with common but varying rate parameter λ. As λ increases, the number of transmitted packets per unit time increases and consequently the amount of whitespace decreases.

The histograms for the busy and idle periods are shown for $\lambda = 100\,\text{pkts/s}$ in Fig. 8. In particular, the busy durations are again discrete as in Fig. 7(a) with the components corresponding to the acknowledgement packets, the data packets, and the router's beacons, respectively. The idle durations on the other hand allow for two preliminary conjectures. First, there is a significant component around 0.7 ms (corresponding to the effect of the contention window and the DIFS). Second, the tail of the histogram appears to decay slower than exponentially, suggesting that a heavy-tailed distribution might be a good fit.

Given the above observations as well as the standard specifics it makes sense to define the following set of states depending on the state that the channel is in.

Data The channel is busy due to the transmission of a data packet. The sojourn time in this state is deterministic and amounts to the time required to transmit the 512 B size packet, i.e. 0.51 ms.

Sifs The channel is idle due to the short inter-frame space required between a data packet and its subsequent acknowledgement. The sojourn time in this state is $10\,\mu s$.

Ack The channel is busy due to the transmission of an acknowledgement packet. The sojourn time is deterministic and amounts to 0.11 ms.

Cw The channel is idle but there are primary users contending for the medium. The sojourn time in this state can be (approximately) derived from the standard. We assume a finite support from $[0, 0.7\,\text{ms}]$ (the size of the contention window). The type of the distribution depends on how many terminals are contending for the medium at the same time. Given that we are mainly concerned with a lightly used channel a uniform distribution will turn out to be a good fit.

Free The channel is idle since none of the primary users has packets to transmit. From the viewpoint of dynamic spectrum access the time spent in this state is essentially defining to what extent the channel can be reused. A generalized Pareto distribution will turn out to be a good fit for the sojourn time in this state.

The Sifs, Cw, and Free state each correspond to an idle medium. In our statistical analysis we will focus on the latter two since the Sifs duration is purely deterministic and too short to be used for dynamic spectrum access (only $10\,\mu s$).

While the histograms depicted in Fig. 8 give a first impression on the distribution of the idle durations, more insight can be gained by looking at the empirical distribution function, which is defined as the fraction of observations smaller than t [16]

$$F_e(t) = \frac{\#i : y_i \leq t}{n}, \tag{7}$$

where $y_i, i = 1, \ldots, n$ correspond to n independent samples. The empirical distribution function is shown in Fig. 9 for several values of the rate parameter λ. We can make two important observations. First, the idle duration (whitespace) decreases with λ, *i.e.* for $\lambda_1 < \lambda_2$, $F_e(t;\lambda_1) < F_e(t;\lambda_2)$, $\forall t$. Second, for $\lambda \leq 200$pkts/s we can clearly see that the distribution of the idle times is a mixture of the contention window and the distribution of the truly 'free' channel (note the bend in the curves at 0.7 ms). Furthermore, the vertical line in Fig. 9 illustrates the finite support of the contention window's distribution. We can see that the slope of F_e within that region is approximately constant, suggesting that a uniform distribution as an appropriate fit (this is also suggested by the standard specifics). The tail distribution corresponding to the free channel shows heavy-tailed behavior and will be analyzed in detail in Sec. 5.

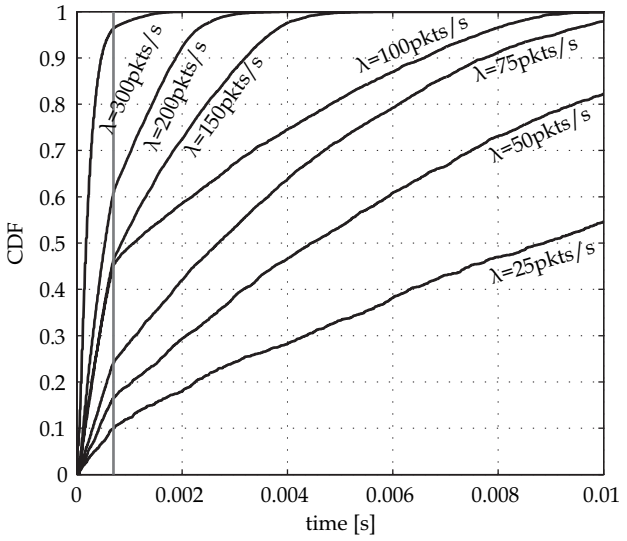

Figure 9: Empirical cdf for the idle durations.

4.4 Representative Traffic Scenarios

In addition to the UDP traffic, we have also looked at a variety of typical traffic scenarios, including file transfers and Voice-over-IP sessions over the WLAN. The resulting empirical cdfs for the idle durations are shown in Fig. 10 and are discussed separately in the following.

First, consider file transfer via secure-FTP from a remote server. In order to collect enough baseband data a text file of approximately 100 kB was transferred 1000 times using a secure-FTP client. The resulting curve shows that there is little remaining whitespace. The effect of the contention window is well-visible by the bend in the empirical cdf at 0.7 ms.

Second, we used D-ITG to generate traffic according to the G.711 codec (used in some VoIP clients). We consider the case of one and three codecs running simultaneously on each of the workstations. The resulting curves show an almost idle channel in the case of one active codec, while the channel appears quite busy in the case of three.

Finally, we used the popular "Skype" client to set up a conference call within the WLAN. A prerecorded audio sample was used to simulate the speech conversation on each of the workstations. The resulting empirical cdf shows that the channel is mostly idle.

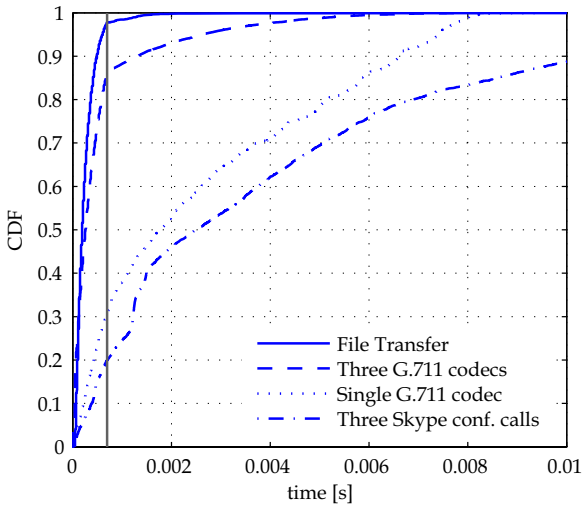

Figure 10: Empirical cdf for FTP and VoIP traffic.

In summary, the empirical cdfs for the traffic scenarios shown in Fig. 10 show a similar behavior compared to the UDP traffic considered before. Specifically, the tails of the distribution appear to be heavy-tailed again. In contrast to the UDP setup, however, the contention window's distribution does not seem to be uniform. While we have (mostly) considered stationary traffic scenarios in this work, we believe that our results extend to the nonstationary case by tracking the variations in traffic.

5. SEMI-MARKOV MODEL

The definition of states given in Sec. 4.3 allows us to convert the processed measurement data to a sequence of states. Note however, that the states CW and FREE are not observable since we can only detect an idle medium but not conclude whether the system is in either of the states. We shall refer to the lumped version of CW and FREE as the IDLE state for brevity.

In a previous paper [3] we have shown that a continuous-time semi-Markov model is an appropriate fit. We shall briefly review this result and then provide an extension to the model with a better goodness-of-fit across λ.

A semi-Markov model can be viewed as an extension to a continuous-time Markov process with separate statistical specification of the transition behavior and sojourn time within each state [17]. The transition behavior in a semi-Markov process retains the Markovian property with transitions from state i to j occurring with probability p_{ij}. In contrast to a CTMP though, given that a transition $i \to j$ occurs, the sojourn time t in state i (before transitioning to j) can be specified arbitrarily according to some cdf $Q_{ij}(t)$ [17].

Recall that in a CTMP the sojourn times in all states need to be exponentially distributed [18].

For specifying the parameters of the semi-Markov process we treat these two parts separately. First, the transition behavior is estimated using the observation sequence and then distributions for the sojourn times are fit to each state.

5.1 Estimating transition probabilities

First, we need to find the transition probabilities given the sequence of states obtained by measurement. To this end we can use the well-known likelihood estimator for the transition probability [17, 19]

$$p_{ij} = \frac{n_{ij}}{n_i}, \tag{8}$$

where n_{ij} is the number of transitions $i \to j$ in our observation sequence, and n_i is the total number of state i occurring in the sequence. Using the above estimator, we have shown in [3] that the sequence of states

$$\textsc{Data} \to \textsc{Sifs} \to \textsc{Ack} \tag{9}$$

is essentially deterministic since its transition probabilities are very close to one. In fact, this does not come as a surprise provided that our system is operating at high SNR and the above sequence simply corresponds to a successful transmission. It should be noted that while collisions still occur infrequently in our setup, their effect appears negligible. The transition diagram resulting from the above analysis is depicted in Fig. 11.

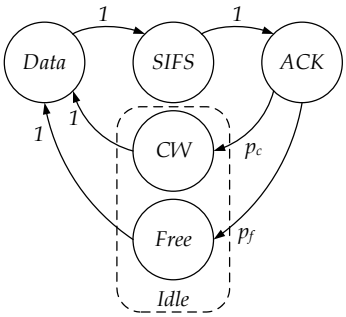

Figure 11: Transition diagram of the semi-Markov model.

5.2 Specifying the sojourn times

So far, we have arrived at the transition diagram shown in Fig. 11. Since the transitions Data → Sifs → Ack are deterministic and the sojourn time in each of these states is deterministic as well, we only need to fit the sojourn time the Idle state (and the substates Cw and Free).

In [3] we have shown that a generalized Pareto distribution provides a good fit to the tail of the idle distribution. However, especially for increasing λ the fit is rather poor for small values of t, since the effect of the contention window is not taken into account. In the present paper we show that the overall goodness-of-fit can be significantly improved by separating the effects of the contention window and the free channel.

As a matter of fact, the sojourn time $F(t)$ in the Idle state (the lumped version of Cw and Free) is a mixture distribution,

$$F(t;\boldsymbol{\theta}) = p_c F_c(t) + p_f F_f(t;\boldsymbol{\theta}), \tag{10}$$

where $F_c(t)$ is the cdf of the contention window (assumed uniform on $[0, T_c]$ and $F_f(t;\boldsymbol{\theta})$ denotes the generalized Pareto cdf of the unused channel depending on the unknown parameters $\boldsymbol{\theta}$. The transition probabilities p_c and p_f are also shown in Fig. 11.

There are several approaches for estimating the unknown parameters of a mixture distribution. The Expectation-Maximization (EM) algorithm [20] is a well-known technique, but computationally expensive to apply in our case. We shall hence pursue a slightly different approach but will come back to the EM-algorithm when fitting a phase-type distribution in Sec. 6.

In order to simplify the analysis we can exploit some structure in (10). In fact, we know that the support of $F_c(t)$ is limited to $[0, T_c]$ (cf. Sec. 4 and Fig. 9, $T_c \approx 0.7\,\text{ms}$). Hence, if we discard all observations $y_i \in [0, T_c]$ (whether or not they are really coming from $F_c(t)$) we are no longer dealing with a mixture but can estimate the parameters of $F_f(t;\boldsymbol{\theta})$ directly.

According to the above we are concerned with estimating the parameters of the generalized Pareto distribution from left-truncated data. Let the truncated data gained by discarding all idle times smaller than the threshold T_c be denoted by $\tilde{y}_i,\ i = 1, \ldots, N_t$. Assuming a generalized Pareto distribution we have the following expression for the pdf

$$f_f(t;k,\sigma) = \frac{1}{\sigma}\left(1 + k\frac{t}{\sigma}\right)^{-1-1/k}, \tag{11}$$

where k denotes the shape, and σ denotes the scale parameter [21]. The cdf is given by

$$F_f(t;k,\sigma) = 1 - \left(1 + k\frac{t}{\sigma}\right)^{-1/k}. \tag{12}$$

Provided that we can only use the left-truncated samples $\tilde{y}_i$ for estimating the parameters the maximum likelihood estimate of the parameter vector $\boldsymbol{\theta} = [k, \sigma]^T$ is given by [22]

$$\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}} \prod_{i=1}^{N_t} \frac{f_f(y_i;\boldsymbol{\theta})}{1 - F_f(T_c;\boldsymbol{\theta})}, \tag{13}$$

where the term in the denominator is due to the left-truncation of the data at T_c. The maximization in the above formula was performed numerically, using an initial value obtained by a moment estimate for the non-truncated data [21].

Given that we have estimated one of the terms in the mixture distribution (10), and realizing that $F_c(t)$ is a uniform distribution on $[0, T_c]$ we can find p_c and p_f, thus fully specifying the desired approximation to the empirical cdf. The fitted distribution as well as the empirical cdf are shown in Fig. 9 for $\lambda = 25\,\text{pkts/s}$ and $\lambda = 100\,\text{pkts/s}$, respectively. The mixture distribution shows a very good fit, which will be assessed quantitatively in the next section. The fitted parameters are shown in Tab. 1.

5.3 Kolmogorov-Smirnov Test

In order to assess the goodness-of-fit of the fitted distribution we employ the Kolmogorov-Smirnov (K-S) test. This

technique allows to discern whether N independent observations are drawn from some given distribution $F(t)$, [23]

$$\mathcal{H}_0 \quad : \quad Y_i \sim F(t),\ i = 1, \ldots, N \tag{14}$$
$$\mathcal{H}_1 \quad : \quad Y_i \nsim F(t),\ i = 1, \ldots, N. \tag{15}$$

It turns out that for analyzing the above hypothesis testing problem, the K-S statistic can be employed

$$D = \max_t |F_e(t) - F(t)|, \tag{16}$$

where $F_e(t)$ is the empirical distribution constructed from the N observations as defined in (7). While the D-value already allows for a quantitative assessment of the goodness-of-fit, it is still influenced by the number of observations N we used in constructing $F_e(t)$. For that reason, a final assessment is usually based on the p-value defined by

$$p = \Pr(D \geq d | \mathcal{H}_0), \tag{17}$$

where d denotes the realization of D constructed from the data. It turns out [23] that the p-value is independent of the distribution $F(t)$ so that (17) can easily be evaluated by Monte-Carlo simulation or by using an appropriate table [22, 24]. Usually, a value of $p \approx 0.1$ is deemed high enough to consider the observations coming from $F(t)$.

We applied the K-S test to our problem. The resulting d and p-values are shown in Tab. 1 together with the parameter estimates $\hat{\boldsymbol{\theta}}$. The fitted distributions are also plotted in Fig. 12 showing an excellent fit with the empirical cdf. Moreover, we can see that ignoring the effect of the contention window and simply fitting a generalized Pareto distribution (direct fit) to the non-truncated data results in a rather poor fit as λ increases.

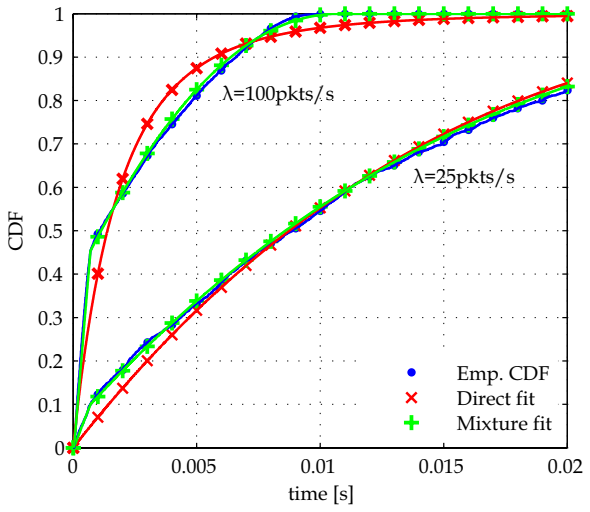

Figure 12: Empirical cdfs and fitted distributions.

6. PHASE-TYPE APPROXIMATION

While the mixture distribution fitted to the empirical cdf in the last section provides a very good fit especially at small to medium λ (that is at those values that are of most concern to us), there are other worthwhile approaches to this problem.

λ [pkts/s]	25	50	75	100	150	200
Direct fit						
k	-0.2705	-0.3225	-0.3171	0.2552	-0.2180	0.0146
σ	0.0138	0.0072	0.0045	0.0018	0.0016	$7.88 \cdot 10^{-4}$
d-value	0.0550	0.0727	0.0950	0.1486	0.1019	0.0547
p-value	0.0936	0.0103	0.0002	0	0.0001	0.0962
Mixture fit						
k	-0.3014	-0.3604	-0.4014	-0.4893	-0.3199	0.0164
σ	0.0149	0.0080	0.0053	0.0053	0.0021	$8.09 \cdot 10^{-4}$
d-value	0.0162	0.0178	0.0137	0.0311	0.0301	0.0499
p-value	0.9992	0.9996	1.0	0.7067	0.7405	0.1603
Hyper-Erlang fit						
d-value	0.0293	0.0167	0.0163	0.0210	0.0316	0.0240
p-value	0.7724	0.9987	0.9991	0.9770	0.6878	0.9295

Table 1: Parameter estimates and goodness-of-fit for a direct fit, a mixture fit, and a hyper-Erlang fit to the empirical data.

Although it is apparent from the empirical data that the channel's idle periods show heavy-tailed behavior, fitting such a distribution directly is not the only option. In fact, from the perspective of applying the model, heavy-tailed distributions have the disadvantage of being difficult to analyze analytically [4, 25]. For this reason, we consider fitting *phase-type* distributions in the following.

Phase-type distributions form a class of distributions that describe the time to absorption of a CTMP with a set of transient and a single absorbing state. These distributions are frequently used in statistical modeling [4, 25, 26] , since they allow to approximate heavy-tailed behavior within the framework of CTMP by expanding the state space.

In practice the usefulness of fitting a phase-type distribution ultimately depends on how many states (each having exponential sojourn time by definition) are needed to approximate the heavy-tailed behavior accurately. Ultimately, we face a tradeoff between obtaining a 'nice' distribution for each of the states, and having a larger number of them.

Phase-type distributions can be classified into numerous groups, including but not limited to hyper-exponential, Erlang, and hyper-Erlang distributions [25]. While, hyper-exponentials (a mixture of exponential distributions) are frequently using in statistical modeling [4] they are not an appropriate choice in our case. In fact, it can be shown that using hyper-exponential distributions we can only approximate distributions that have a Coefficient of Variation (CoV) greater than one [25]. From the empirical data, however, we can infer that we are dealing with a CoV smaller than one.

Instead, we shall consider a hyper-Erlang distribution, that is a distribution that is modeled by the CTMP depicted in Fig. 13. This class of distributions can approximate empirical distributions with arbitrary CoV and has furthermore received interest lately [26]. In particular, an efficient EM-based algorithm has been proposed in [25], which is used to fit the distribution and shall be briefly outlined in the following.

6.1 Expectation-Maximization Algorithm

In this section we show how to fit an M-component mixture distribution to the empirical distribution gained by measurement, that is we need to find a set of parameters $\boldsymbol{\mu} = [\mu_1, \ldots, \mu_M]$ and mixture coefficients $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_M]$

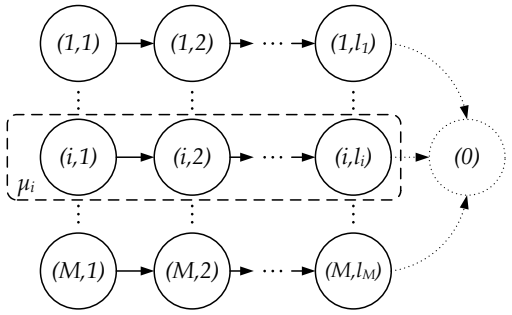

Figure 13: Transition diagram of CTMP resulting in an overall hyper-Erlang distribution. The absorbing state is denoted by '(0)'.

such that

$$f(t;\boldsymbol{\alpha},\boldsymbol{\mu}) = \sum_{i=1}^{M} \alpha_i f_{\mathrm{Er}}(t,\mu_i), \tag{18}$$

where

$$f_{\mathrm{Er}}(t;\mu_i) = \frac{(\mu_i t)^{l_i-1}}{(l_i-1)!}\mu_i e^{-\mu_i t} \tag{19}$$

is the pdf of an Erlang distribution with shape parameter l_i (assumed to be known) and rate parameter μ_i. Each Erlang distribution (19) in the mixture (18) can be viewed as generated by one of the chains in the CTMP depicted in Fig. 13.

For estimating the unknown parameters $\boldsymbol{\theta} = [\boldsymbol{\alpha},\boldsymbol{\mu}]^T$ given observations $y_i,\ i = 1,\ldots,N$, we can adopt a maximum likelihood approach arriving at

$$\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}} \sum_{k=1}^{N} \log\left[\sum_{i=1}^{M} \alpha_i f_{\mathrm{Er}}(y_k;\mu_i)\right]. \tag{20}$$

Unfortunately, the maximization in the above formula is not easy to carry out since the expression involves the logarithm of a sum [25]. However, we would be able to simplify (20) if we knew from which mixture each expectation were drawn from. Clearly, we are not given this additional information but we can interpret it as missing data and find its distribution.

More precisely, let us associate every observation y_i with an index z_i indicating from which mixture y_i is drawn from. It can then be shown [25] that the probability mass function $p_Z(z)$ of z_i is found by Bayes' rule given some initial estimate $\hat{\boldsymbol{\theta}} = [\hat{\boldsymbol{\alpha}},\hat{\boldsymbol{\mu}}]^T$,

$$p_Z(z|y_i;\hat{\boldsymbol{\theta}}) = \frac{\hat{\alpha}_z f_{\mathrm{Er}}(y_i;\hat{\mu}_z)}{\sum_{k=1}^{M}\hat{\alpha_k} f_{\mathrm{Er}}(y_i|\hat{\mu_k})} \tag{21}$$

Given the above pmf we can evaluate the expected value of (20) and maximize this function with respect to the unknown parameters,

$$\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}} \mathbb{E}_z\left\{\sum_{i=1}^{N} \log\left[\alpha_{Z_i} f_{\mathrm{Er}}(y_i;\mu_{Z_i})\right]\right\}. \tag{22}$$

It is shown in [25] that the maximization can be carried out in closed-form yielding

$$\hat{\alpha}_i = \frac{1}{N}\sum_{k=1}^{N} p_Z(i|y_k;\hat{\boldsymbol{\theta}}) \tag{23}$$

$$\hat{\mu}_i = \frac{l_i \sum_{k=1}^{N} p_Z(i|y_k;\hat{\boldsymbol{\theta}})}{\sum_{k=1}^{N} p_Z(i|y_k;\hat{\boldsymbol{\theta}}) y_k}. \tag{24}$$

The iterative procedure for finding the unknown parameters $\boldsymbol{\alpha}$ and $\boldsymbol{\mu}$ can now be summarized at follows [25]. First, start with an initial estimate $\boldsymbol{\theta} = [\boldsymbol{\alpha},\boldsymbol{\mu}]^T$ and compute the pmf p_Z of the unobserved data z_i using (21). Then, maximize the log-likelihood function averaged over p_Z by using the closed form expressions (23) and (24). This yields a new estimate for $\boldsymbol{\theta}$ and continue to apply the procedure iteratively.

It remains to specify how to find optimal values for the number of mixtures M, as well as the integer-valued shape parameter $l_i,\ i = 1,\ldots,M$ for each of the Erlang distributions. We have explored different values and obtained the best goodness of fit using $l_1 = 2$, $l_2 = 2$, and $l_3 = 3$ for a reasonable choice of $M = 3$ mixture components.

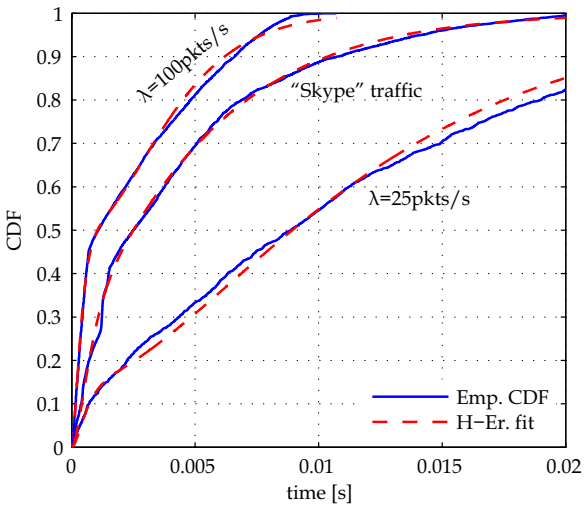

Figure 14: Hyper-Erlang fit for the empirical data.

6.2 Kolmogorov-Smirnov Test

We again investigated the goodness-of-fit for the hyper-Erlang distribution using the K-S test. The results are shown in Tab. 1 and plotted in Fig. 14. Again, the K-S test confirms the hypothesis that the fitted distribution is an accurate model across λ. Furthermore, we have included the empirical distribution for the VoIP setup using "Skype". The hyper-Erlang distribution also provides for an excellent fit in this case.

7. CONCLUSIONS

In conclusion we have proposed stochastic models that can be used to predict the idle durations between the bursty transmissions of a WLAN. Our contribution both involves a measurement-based component as well as the statistical analysis of the data.

Since the measurements are based on raw data gathered by a vector signal analyzer we were able to validate the setup and guarantee accurate results using both antenna-based and isolated-RF setups. Two sensing strategies were implemented to identify problems that would arise in an actual implementation and to verify the correct operation of the sensing algorithms.

The statistical analysis of the gathered data significantly extends our previous results and provides a very good fit with the empirical distribution. Furthermore, we propose to fit a hyper-Erlang distribution which might represent a good tradeoff between modeling accuracy and tractability of the model. Goodness-of-fit techniques are used to validate both fitting approaches.

Finally, we plan to further investigate our model's fit in practical traffic scenarios. In the paper, we have already addressed VoIP and FTP-based traffic, yet there are many more scenarios to consider (HTTP-traffic, video-streaming, etc.).

8. REFERENCES

[1] D. Čabrić, S. M. Mishra, D. Willkomm, R. Broderson, and A. Wolisz, "A Cognitive Radio Approach for Usage of Virtual Unlicensed Spectrum," in *Proc. IST Mobile Wireless Communications Summit*, 2005.

[2] Q. Zhao, L. Tong, and A. Swami, "Decentralized Cognitive MAC for Dynamic Spectrum Access," in *Proc. First IEEE Internatial Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Nov. 2005, pp. 224–232.

[3] S. Geirhofer, L. Tong, and B. M. Sadler, "A Measurement-Based Model for Dynamic Spectrum Access," submitted to IEEE Conference on Military Communications (MILCOM), 2006.

[4] A. Feldmann and W. Whitt, "Fitting Mixtures of Exponentials to Long-Tail Distributions to Analyze Network Performance Models," in *Proc. INFOCOM*, vol. 3, Apr. 1997, pp. 1096–1104.

[5] *Proceedings of the First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Nov. 2005.

[6] "DARPA: The Next Generation (XG) Program," http://www.darpa.mil/ato/programs/xg/index.htm.

[7] L. Xu, R. Tonjes, T. Paila, W. Hansmann, M. Frank, and M. Albrecht, "DRiVE-ing to the Internet: Dynamic Radio for IP Services in Vehicular Environments," in *Proc. of the 25th Annual IEEE Conference on Local Computer Networks*, Nov. 2000, pp. 281–289.

[8] R. Tandra and A. Sahai, "Fundamental limits on detection in low SNR under noise uncertainty," in *Proc. WirelessCom 05 Symposium on Signal Processing*, 2005.

[9] A. Sahai, N. Hoven, S. Mishra, and R. Tandra, "Fundamental tradeoffs in robust spectrum sensing for opportunistic frequency reuse," *submitted to IEEE J. Select. Areas Commun.*, 2006.

[10] Q. Zhao, L. Tong, A. Swami, and Y. Chen, "Decentralized Cognitive MAC for Opportunistic Spectrum Access in Ad Hoc Networks: A POMDP Framework," *submitted to IEEE J. Select. Areas Commun.*, Feb. 2006.

[11] Agilent Technologies, "Agilent 89611A 70MHz IF Vector Signal Analyzer," data sheet, Oct. 2001.

[12] S. Avallone, A. Botta, D. Emma, S. Guadagno, and A. Pescape, "D-ITG V.2.4 Manual," University of Napoli "Federio II", Tech. Rep., Dec. 2004.

[13] H. V. Poor, *An Introduction to Signal Detection and Estimation*, 2nd ed. Springer-Verlag, 1994.

[14] ANSI/IEEE Standard 802.11b-1999 (R2003), "Wireless lan medium access control (mac) and physical layer (phy) specifications: Higher-speed physical layer extension in the 2.4ghz band," IEEE SA Standards Board, Tech. Rep., 1999.

[15] ANSI/IEEE Standard 802.11, 1999 Edition (R2003), "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," IEEE/SA Standards Board, Tech. Rep., 1999.

[16] R. B. D'Agostino and M. A. Stephens, *Goodness-of-fit techniques*. Marcel Dekker, Inc., 1986.

[17] S. M. Ross, *Applied Probability Models with Optimization Applications*. Dover Publications, 1970.

[18] A. Papoulis and S. U. Pillai, *Probability, Random Variables, and Stochastic Processes*, 4th ed. McGraw Hill Publishing Company, 2002.

[19] P. Billingsley, "Statistical Methods in Markov Chains," *The Annals of Mathematical Statistics*, vol. 32, no. 1, pp. 12–40, Mar. 1961.

[20] A. Dempster, N. Laird, and D. Rubin, "Maximum Likelihood from Incomplete Data via the EM Algorithm," *Journal of the Royal Statistical Society, Series B (Methodological)*, vol. 39, no. 1, pp. 1–38, 1977.

[21] S. Kotz and S. Nadarajah, *Extreme Value Distributions. Theory and Applications*. Imperial College Press, 2000.

[22] A. C. Cohen, *Truncated and Censored Samples. Theory and Applications*. Marcel Dekker, Inc., 1991.

[23] S. M. Ross, *Simulation*, 3rd ed. Academic Press, 2002.

[24] F. J. Massey Jr., "The Kolmogorov-Smirnov Test for Goodness of Fit," *Journal of the American Statistical Association*, vol. 46, no. 253, pp. 68–78, Mar. 1951.

[25] A. Thümmler, P. Buchholz, and M. Telek, "A Novel Approach for Fitting Probability Distriubtions to Real Trace Data with the EM Algorithm," in *Proc. International Conference on Dependable Systems and Networks (DSN)*, June 2005, pp. 712–721.

[26] R. El Abdouni Khayari, R. Sadre, and B. Haverkort, "Fitting World-Wide Web Request Traces with the EM-algorithm," *Performance Evaluation*, vol. 52, no. 2, pp. 175–191, Apr. 2003.

[1]The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government.