

Dynamic Spectrum Access in the Time Domain: Modeling and Exploiting Whitespace

Stefan Geirhofer, *Student Member, IEEE*, Lang Tong, *Fellow, IEEE*,
and Brian M. Sadler, *Fellow, IEEE*

Abstract

Dynamic spectrum access is a promising approach to alleviate the spectrum scarcity wireless communications faces today. In short, it aims at reusing sparsely occupied frequency bands while causing no (or insignificant) interference to the actual licensees. This article focuses on applying this concept in the time domain by exploiting idle durations between bursty transmissions of multiaccess communication channels, addressing WLAN as an example of practical importance. A statistical model based on empirical data is presented and it is shown how to use this model for deriving access strategies. The coexistence of Bluetooth and WLAN is considered as a concrete example.

I. INTRODUCTION

Most of the radio spectrum relevant to wireless communications has been densely allocated by regulators, making it unlikely that the bandwidth needs of emerging technologies can be met. Actual measurements illustrate, however, that the scarcity is not a result of heavy usage of the spectrum; in contrast it is merely due to the inefficiency of the static frequency allocation pursued by regulators. A typical spectrum utilization of around 5% or even less is reported [1].

Dynamic spectrum access resolves this paradox by opening assigned but sparsely used frequency bands to secondary users, provided that interference to the actual licensee is kept insignificant. Different approaches to this problem have been suggested ranging from entirely opening certain bands (in the fashion of the unlicensed ISM bands) to hierarchical schemes requiring secondary users to meet interference

Stefan Geirhofer and Lang Tong are with the Department of Electrical and Computer Engineering, Cornell University, Ithaca NY 14853 (e-mail: {sg355,lt35}@cornell.edu). Brian Sadler is with the U.S. Army Research Laboratory, Adelphi, MD 20783 (e-mail: bsadler@arl.army.mil). This work is supported in part 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.

constraints [2]. In this paper we restrict ourselves to such hierarchical schemes, that is, we design the cognitive radio so that both systems are orthogonal.

Mutual interference can be avoided by exploiting different degrees-of-freedom in the design of the secondary system. In particular, orthogonality can be achieved by sufficient spatial separation [3], by OFDM techniques [4], or by interleaving transmissions in the time domain [1], [5], [6]. The majority of research, so far, has focused on DSA in the spatial domain, meaning that if the primary user's absence in a certain area can be detected reliably, and if we carefully confine our transmissions to this area, interference is limited by this sufficient spatial separation. Reusing TV-bands that are allocated but not broadcast across the entire country are a prominent example of this paradigm [1].

Dynamic spectrum access is, however, not limited to the spatial domain. As in [5], [6], this paper aims at achieving orthogonality among users in the time-domain by transmitting during the idle periods remaining between the primary user's packet transmissions. Clearly, the problem under consideration is of practical appeal only if the idle periods between packet transmissions are long enough. Although the amount of whitespace naturally depends on and varies according to the traffic characteristics, our results illustrate that in most cases sufficient whitespace will be available.

As a motivating example, consider a conference call using the popular VoIP client "Skype" (with three participating parties). The complex baseband signal for this example is shown in Fig. 1. Although the channel is continuously used for packet transmissions (cf. Fig. 1(a)), there are large gaps between consecutive packets as illustrated in Fig. 1(b). In fact, the channel was idle for 89% of the time. The situation is similar for other traffic scenarios [7], [8].

There are numerous applications that can take advantage of such transmission opportunities. As an example, consider building automation where sensor networks are used to provide a central controller with local measurements of environment parameters. By nature of this application such networks will have to coexist with high-bandwidth data networks (such as WLAN). Given the fact that such sensors communicate only sporadically and at low rate, it appears reasonable to assume that such systems could efficiently reuse remaining whitespace.

The existence of sufficient whitespace raises the question of how to exploit this resource in practice: How complicated is a model that provides adequate prediction performance? Moreover, given a model, how do we apply it to derive practical access schemes for the secondary system? Finally, are such schemes amenable to a real-time implementation possibly on a battery-powered device with processing limitations?

While some of these questions have been addressed by researchers [1], many issues remain to be inves-

tigated. Optimal access strategies have been derived in [5] based on partially observable Markov decision processes assuming that primary and secondary users share the same slotted transmission structure. This assumption, however, is not justified in some practical scenarios. In fact, WLAN transmissions do not have a slot structure but mandate the channel to be modeled continuously.

Finding a realistic yet tractable model to suit such standards is challenging. A bottom-up approach by directly incorporating the specifications of all devices involved is intractable due to the plethora of factors that affect the medium access (number of users, traffic characteristics, higher layer effects, and channel conditions to name only a few). Ultimately, we need to strike a balance between statistical accuracy on the one hand, and complexity on the other. We will see that a semi-Markov model is a promising candidate to trade off both extremes [7]. Lastly, we mention the Laycock-Gott model [9] which has been used to predict spectral occupancy in HF bands. Due to different modeling objectives, however, our work requires different statistical techniques.

This paper is organized as follows. The experimental testbed used to gather empirical data of the channel's statistics is presented in Sec. II together with possible sensing methods. The main contribution is the establishment of a semi-Markov model in Sec. III; we show how to arrive at estimates for the model parameters and discuss both stationary and nonstationary traffic scenarios. Finally, Sec. IV gives an example of how to apply this model to derive access strategies, focusing on Bluetooth/WLAN coexistence as an example of practical importance.

Throughout the paper we shall identify the primary user with an 802.11b-based wireless LAN (WLAN). We believe this not only is a choice of practical relevance but also can be extended to related multiaccess communications systems.

II. SENSING METHODS: AN EXPERIMENTAL TESTBED

The reliable detection of primary users is the main challenge in spatial DSA since interference is limited only by sufficient spatial separation. Possible shadowing effects and the hidden-node problem can make it difficult or impossible for a single node to reach the required performance. Cooperative sensing among secondary users becomes indispensable in those cases [10].

For the temporal DSA schemes considered in this article the sensing burden is alleviated for two reasons. First, we expect to be dealing with a different propagation environment. Most likely we will face an indoor propagation scenario and medium to high SNR conditions (as it is typical for WLAN applications). Second, in temporal DSA, the attention shifts from purely detecting the presence of a primary system to predicting its behavior. We emphasize that it is this prediction—and not the sensing

alone—that guarantees orthogonality between primary and secondary users.

This section briefly outlines the testbed setup used to gather the empirical data our model is based on. Subsequently, two sensing strategies are outlined and their practical relevance is addressed.

A. Testbed description

The data we used to propose and analyze our model have been gathered by a testbed consisting of a wireless router and several workstations with WLAN adapter cards (see Fig. 2). In contrast to measurement setups considered in other publications, we used a vector signal analyzer to capture the raw complex baseband data. In this way, different sensing strategies can be evaluated on the *same* data, giving valuable insights when it comes to developing a real-time implementation. Furthermore, by comparing the results of different sensing strategies, the validity of the data could be confirmed.

In addition to the antenna-based setup depicted in Fig. 2 we also tested an isolated setup, where all devices were connected via cables to a resistive power divider. In this way interference from adjacent WLAN hotspots (unrelated to our measurement) could be avoided [7].

We employed different traffic generators to guarantee repeatable measurement results, namely “D-ITG” to create synthetic, constant-payload UDP traffic [7] and “GenSyn” to generate realistic HTTP traffic [11].

B. Sensing strategies

The strategies for detecting busy and idle durations can be classified according to whether the primary user’s transmission standard is known. If not, a logical approach is energy-based detection. On the other hand, if we have information about the primary system, this knowledge can be exploited to achieve better performance. We refer to such methods as feature-based detection.

Energy-based detection is formulated mathematically as discerning the hypotheses of observing only noise and observing a signal in noise, respectively. We assume that a decision is based on a block of complex baseband samples. Given a target value for the error probability we have to choose the number of samples (the block length) such that the specification is met even for the worst-case SNR. In practice, if we require an error probability no greater than 10^{-5} , and if we consider blocks of $1\mu\text{s}$, a minimum SNR of 4.8 dB is necessary. This is a realistic value given usual WLAN setups.

Feature-based techniques improve the detection performance by exploiting standard specifics. In particular, for the 802.11b based WLAN considered in this work, we can find the start of packets by locking on to the synchronization preamble of the transmitted packets. Moreover, by decoding the LENGTH field within the preamble, we can find the exact packet duration [7].

Although feature-based detection has the ability to increase detection performance, energy-detection typically lends itself to a simpler implementation. In practice, additional considerations need to be taken into account. For instance, feature-based detection has the drawback of requiring a continuous monitoring of the channel. For this reason, energy-based detection should be preferred if the channel is only sensed at certain time instances (for instance at the beginning of every slot).

III. MODELING WHITESPACE: A STATISTICAL APPROACH

The last section discussed two sensing strategies for detecting the busy and idle durations of the channel. Based on these data we now develop a statistical model that allows us to predict the channel's behavior. We emphasize that since WLAN does not have a slot structure, a description in continuous time is necessary. Specifically, our model needs to incorporate two separate components: (i) the states of the channel and their transition behavior, as well as (ii) how long the system resides in each of the states.

The above observation motivates the use of a semi-Markov model, which can be viewed as an extension of a continuous-time Markov chain (CTMC). While both models describe the transition behavior in the same way, a semi-Markov model allows for specifying the occupancy durations for each state arbitrarily (they need not be exponential as in a CTMC).

A. Channel state classification

In defining the state space, consider once more the complex baseband signal shown in Fig. 1(b). The figure motivates to extend the model's state space beyond the trivial 'busy' and 'idle' states. In fact, the WLAN standard specifies that every transmission of a data packet must be followed by an acknowledgement (after a short gap called short interframe space (SIFS)). Data packets and acknowledgements serve different purposes and as a consequence it makes sense to separate their statistical characterization; ultimately it is very likely for an acknowledgement to follow a data packet. We thus introduce the states DATA, SIFS, and ACK to represent the possible channel conditions described above (see Fig. 3).

First insights into the model can be gained by estimating its transition probabilities. It turns out that the sequence of states DATA→SIFS→ACK is effectively deterministic, in the sense that the probability of observing this sequence is close to one. This does not come as a surprise since this sequence of states precisely corresponds to a successful packet transmission together with the obligatory acknowledgement. The model can thus be simplified by lumping these three states and defining a new TRANSMIT state that corresponds to such a successful transmission (see Fig. 3).

This simplification also makes sense from a practical viewpoint. Ultimately, our model serves the purpose of predicting the idle durations (the ‘whitespace’) of the channel. While the short gap between data packets and their acknowledgements corresponds to an idle channel, it is too short to be used for secondary transmissions.

The above paragraphs have introduced the TRANSMIT state consisting of the transmission of a data packet together with its acknowledgement. Successive transmissions are separated by an idle channel. However, it is important to further differentiate between two cases. First, after a successful transmission, a station cannot continue transmitting but has to follow a standardized backoff procedure in order to give all stations equal access to the medium. This backoff procedure is implemented by a so-called contention period: all stations draw random numbers within a certain range and defer access for this number of slot durations. A station can go ahead and transmit only after having waited for this time period, provided that no other station has accessed the medium beforehand. Ideally this means that the station which draws the smallest number gets to access the medium first.

Not all idle periods are due to the contention window, though. The channel can also be idle because none of the stations has any packets to transmit. In fact, this is precisely the situation we aim to exploit in dynamic spectrum access. We shall refer to such a channel as being “free”.

B. Occupancy durations

The simplified model shown in Fig. 3 guides us to the second component of our model: the occupancy duration or sojourn time in each state. To specify this component we need to find (a parametric) distribution that captures this statistical behavior.

The sojourn time in the TRANSMIT state is primarily affected by the traffic characteristics and the scheduler employed in the adapter cards and the wireless router. Transfer of large files generally leads to rather long packets, while streaming traffic such as VoIP leads to shorter packets. From our measurement results for different traffic scenarios we infer that, in general, several discrete components are observed.

The statistical characterization of the IDLE channel is more challenging since, as mentioned before, the channel can be idle either due to the contention window or a truly “free” channel. This suggests a mixture distribution, incorporating the effects from both states. Indeed, it will turn out that the contention period has an almost uniform sojourn time (ranging from 0 to 0.7 ms), while a free state shows heavy-tailed behavior and is well approximated by a generalized Pareto distribution.

The empirical distribution function for the idle duration of the channel is shown in Fig. 4 for the case of constant payload UDP traffic. The effects of the mixture distribution—uniformly distributed body

but heavy-tail—are clearly visible. Note the bend in the distribution at approximately 0.7 ms. We can also see that the contention window’s influence becomes more pronounced as the channel gets busier (upper curve). This is intuitive, since the contention procedure will be applied more often as traffic load increases.

In addition to this mixture fit, we have used phase-type distributions to capture the data statistically. Such distributions have the advantage that they can be viewed as the time-to-absorption of a continuous-time Markov chain with multiple states, generally allowing for some simplification in applying the model. A hyper-Erlang distribution turned out to lead to a good fit. Furthermore, an efficient EM type method exists to estimate the parameters of this model [12]. The fit is again shown in Fig. 4.

We have also fitted an exponential distribution to the data. While it does not approximate the empirical data as well as the distributions discussed above, it makes the model easier to apply in practice. Statistically, such an approach is not entirely justified, but worthwhile to identify heuristic algorithms.

C. Goodness-of-fit analysis

The accuracy of the above fitting methods can intuitively be evaluated by looking at the plots in Fig. 4. Beyond this intuitive assessment an analytical analysis can be based on the Kolmogorov-Smirnov (K-S) test, which is a well-established technique in statistics [13].

The K-S test examines whether the empirical data is drawn independently from a distribution $F(t)$, namely the one we want to postulate in our model. Intuitively, if this were the case we would expect the empirical CDF $F_e(t)$, based on the data alone, to resemble $F(t)$ closely. Indeed, the K-S test is based on the maximum aberration between $F_e(t)$ and $F(t)$; a small aberration suggests that the samples come from the proposed distribution.

Although this maximum aberration, also referred to as K-S statistic, reflects the goodness-of-fit quantitatively, it is not normalized by the sample size N . Ultimately, if $F_e(t)$ is constructed from a large set of observations, we expect a better fit than from a small sample. This is accounted for by the significance level α , which specifies the maximum tolerable aberration in relation to N [13]. Usually, a significance level of approximately $\alpha = 0.1$ is deemed appropriate.

1) *Stationary UDP traffic scenario:* We first consider the stationary UDP traffic scenario mentioned previously. The sojourn time in the TRANSMIT state is deterministic due to the constant payload. The statistical distribution for the IDLE state is shown in Fig. 4 for the rates $\lambda = 25$ pkts/s and $\lambda = 100$ pkts/s, respectively. The fact that the traffic is stationary enables us to accurately evaluate the goodness-of-fit based on long trace data.

We can see that both the mixture and the hyper-Erlang distribution show an excellent fit with the empirical cumulative distribution function (E-CDF). The K-S test confirms our model’s accuracy for both distributions. Although Fig. 4 suggests a decent fit for the exponential distribution also, this hypothesis is rejected by the test. Nevertheless, an exponential distribution might still be interesting in practice, because it simplifies the derivation of access schemes.

2) *Nonstationary HTTP traffic*: In the stationary traffic scenario above we could evaluate the goodness-of-fit without having to account for time variations. In practice, however, most traffic scenarios are nonstationary. Therefore, we consider more realistic HTTP traffic generated by “GenSyn”. This traffic generator accesses a list of specified webpages, but in a way that mimics the behavior of a real user and thus accounts for HTTP’s nonstationary behavior [11].

We again consider a hyper-Erlang distribution and approximate the nonstationary traffic with blocks of length L . Within a block the traffic parameters are assumed constant. In choosing L we have to strike a balance between approximating the time variations and being able to estimate the model parameters accurately. A good tradeoff was obtained for $L = 100$.

This approach is illustrated in Fig. 5 for a hyper-Erlang and an exponential fit, respectively. The figure shows how the goodness-of-fit evolves with the time variations of the traffic. Similar to the stationary scenario, the hyper-Erlang distribution shows a satisfactory goodness-of-fit over the entire sample, while the exponential distribution is clearly rejected by the K-S test.

IV. DERIVING ACCESS SCHEMES: A PRACTICAL EXAMPLE

In the last section we have proposed a statistical model for the primary user’s channel access and verified its goodness-of-fit. We are now giving a concrete example on how to apply this model in practice. Specifically, we show how to enhance the coexistence of Bluetooth and WLAN in the unlicensed ISM band, identifying the primary user with WLAN and the secondary system with Bluetooth.

A. *Bluetooth/WLAN coexistence*

The coexistence between Bluetooth and WLAN in the unlicensed ISM band is of significant practical concern because mutual interference can severely limit the performance of both systems [14]. The interference arises from the fact that both standards operate in the ISM band around 2.4 GHz, spanning about 80 MHz of bandwidth. The 802.11 based WLAN systems employed today utilize 11 to 13 channels (depending on national standardization), each having a bandwidth of 22 MHz. We stress that most of these

channels are overlapping and that there is, in fact, only one configuration of three orthogonal channels (cf. Fig. 6). We shall focus on this configuration in the following.

Bluetooth employs frequency hopping transmissions in 1 MHz wide channels across the same frequency band. In basic operation, the Bluetooth's hopping pattern is pseudo random throughout the entire band and completely oblivious of the WLAN's transmissions. Consequently, collisions between WLAN packets and Bluetooth transmissions occur frequently, degrading the performance of both systems.

Given the practical importance of this problem, several propositions on how to mitigate the mutual interference have been made. One way of reducing interference is to adapt the Bluetooth's hopping sequence such that it avoids channels with large interference. Such methods are referred to as adaptive frequency hopping [14] and consist of two main components, namely (i) the sensing and classification of channels according to their interference and (ii) the adaptation of the hopping sequence such that 'bad' channels are avoided whenever possible.

Adaptive frequency hopping has been considered previously based on classifying the Bluetooth channels according to their long-term behavior (for instance, by counting the number of bit errors for each channel). While such an approach is able to mitigate interference from static sources (*e.g.* microwave ovens), it is not well suited to capture the time-varying behavior of WLAN interferers. In contrast, by incorporating the statistical model proposed in this paper, it is possible to predict the WLAN's behavior and thus to improve the performance significantly.

Our analysis focuses on the representative configuration shown in Fig. 6. There are three orthogonal WLAN channels, each overlapping in frequency with 22 Bluetooth channels. We shall refer to the WLAN channels as "bands" to avoid confusion with the Bluetooth channels. Fig. 6 illustrates that the Bluetooth's transmissions are slotted while WLAN stations can start transmitting at any time. Furthermore, we emphasize that all Bluetooth channels within the same WLAN band show the same access behavior due to the larger bandwidth of the WLAN. It thus suffices to model statistically the behavior of the three WLAN bands rather than each of the 80 Bluetooth channels. The computational burden is thus significantly reduced.

B. Enhanced hopping scheme

We propose the following adaptive frequency hopping operation for the Bluetooth transmitter. At the beginning of each slot, the current channel is sensed. If the presence of a WLAN packet is detected, no transmission takes place since this would inevitably lead to a collision. On the other hand, even if no WLAN packet is detected, a collision can still occur if the WLAN becomes active during the subsequent

slot, see slot 6 in Fig. 6(a). The probability of this event can be kept small, however, by initiating a transmission only with a certain probability $\gamma \leq 1$, even if the channel is sensed idle.

The probability of collision (and thus the design of γ) depends on the WLAN's traffic characteristics. In particular, we employ our model to find the probability that conditioned on the channel having been idle for k slots, no WLAN packet will be transmitted during the subsequent slot. By obtaining this probability we can design γ such that collisions with the WLAN occur with probability smaller than some constraint.

The approach described above constrains the collision probability for each band in Fig. 6. At the same time, however, we desire to maximize the Bluetooth's throughput. This maximization is tightly linked with the optimal hopping across the WLAN bands since each band provides a parallel transmission opportunity. Hardware limitations make it unrealistic, however, that the Bluetooth transmitter can observe the state of all bands at the same time. The problem is thus *partially-observed* making the optimal design of the system a challenging problem that needs to be addressed within a decision-theoretic framework [5], [15]. While finding the optimal solution to this problem goes beyond the scope of this paper, we illustrate the possible performance gain along the lines of a heuristic scheme.

Consider a scheme that deterministically switches to the next band, whenever the current channel becomes busy. The performance of such a scheme for the UDP traffic scenario considered previously is shown in Fig. 6(b). Its performance is compared to a scheme that can observe all WLAN bands at the same time. While this assumption is not necessarily practical it presents an upper bound on the performance of any partially observed scheme. The figure shows that finding the optimal switching between bands is a worthwhile problem. Ultimately, a possible performance improvement of up to a factor of two can be expected. We also obtained the collision probability and the throughput of a regular (completely oblivious) Bluetooth transmitter. It achieved a (relative) throughput of 70% while showing a collision probability of 35%. By looking at the performance evaluation in Fig. 6(b) we can see that the regular Bluetooth operation is clearly outperformed.


V. CONCLUSIONS

In conclusion we have addressed the problem of dynamically accessing spectrum in the time-domain by taking advantage of whitespace that remains between bursty packet transmissions of a WLAN. The proposed model captures the WLAN's medium access statistically, but remains tractable enough to be used for deriving practical access schemes for the secondary user. As an example, we have illustrated how our model could be used to enhance the coexistence between WLAN and Bluetooth in the unlicensed

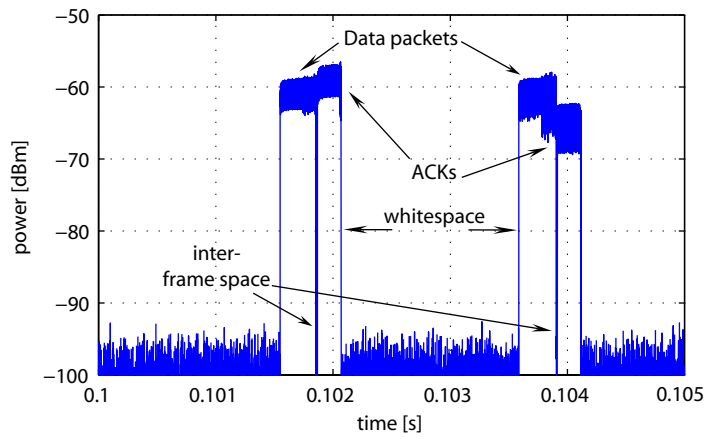
ISM band. We emphasize that our model is based on empirical data gathered from snapshots of the raw complex baseband signal of the WLAN. This enables us to test sensing algorithms on the same data and hence to gain insight into challenges that arise when it comes to a real-time implementation.

REFERENCES

- [1] *Proceedings of the First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Nov. 2005.
- [2] Q. Zhao and B. M. Sadler, “Dynamic Spectrum Access: Signal Processing, Networking, and Regulatory Policy,” *accepted to the IEEE Signal Processing Mag.*, 2006.
- [3] 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.
- [4] T. A. Weiss and F. K. Jondral, “Spectrum pooling: an innovative strategy for the enhancement of spectrum efficiency,” *IEEE Commun. Mag.*, vol. 42, no. 3, pp. S8–14, Mar. 2004.
- [5] Q. Zhao, L. Tong, and A. Swami, “Decentralized Cognitive MAC for Dynamic Spectrum Access,” in *Proc. First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Nov. 2005, pp. 224–232.
- [6] P. Papadimitratos, S. Sankaranarayanan, and A. Mishra, “A Bandwidth Sharing Approach to Improve Licensed Spectrum Utilization,” *IEEE Commun. Mag.*, vol. 43, no. 12, pp. S10–S14, Dec. 2005.
- [7] S. Geirhofer, L. Tong, and B. M. Sadler, “Dynamic Spectrum Access in WLAN Channels: Empirical Model and Its Stochastic Analysis,” in *Proc. First International Workshop on Technology and Policy for Accessing Spectrum*, 2006.
- [8] —, “A Measurement-Based Model for Dynamic Spectrum Access,” in *Proc. IEEE Conference on Military Communications (MILCOM)*, 2006.
- [9] C. A. Pantjarios, P. J. Laycock, G. F. Gott, and S. K. Chan, “Development of the Laycock-Gott occupancy model,” *IEE Proc.-Commun.*, vol. 144, no. 1, pp. 33–39, Feb. 1997.
- [10] R. Tandra and A. Sahai, “Fundamental limits on detection in low SNR under noise uncertainty,” in *Proc. International Conference on Wireless Networks, Communications and Mobile Computing*, 2005.
- [11] P. E. Heegaard, “GenSyn - a Java based generator of synthetic Internet traffic linking user behaviour models to real network protocols,” in *Proc. ITC Specialist Seminar on IP Traffic Measurement, Modeling, and Management*, 2000.
- [12] A. Thümmler, P. Buchholz, and M. Telek, “A Novel Approach for Fitting Probability Distributions to Real Trace Data with the EM Algorithm,” in *Proc. International Conference on Dependable Systems and Networks (DSN)*, June 2005, pp. 712–721.
- [13] R. B. D’Agostino and M. A. Stephens, *Goodness-of-fit techniques*. Marcel Dekker, Inc., 1986.
- [14] N. Golmie, N. Chevrollier, and O. Rebala, “Bluetooth and WLAN Coexistence: Challenges and Solutions,” *IEEE Trans. Wireless Commun.*, vol. 10, no. 6, pp. 22–29, Dec. 2003.
- [15] Q. Zhao, S. Geirhofer, L. Tong, and B. M. Sadler, “Optimal Dynamic Spectrum Access via Periodic Channel Sensing,” to appear in *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, Mar. 2007.



(a) Complex baseband signal. Enlarged view in Fig. 1(b).



(b) Enlarged view of two subsequent packet transmissions.

Fig. 1. Complex baseband signal of an 802.11b-WLAN supporting a Skype conference call.

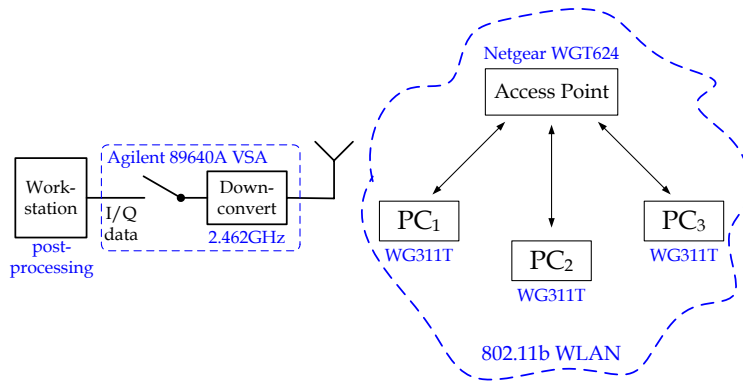


Fig. 2. Testbed configuration.

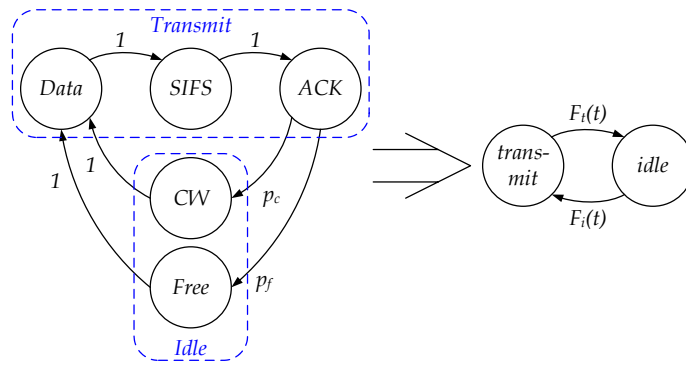


Fig. 3. The proposed semi-Markov model.

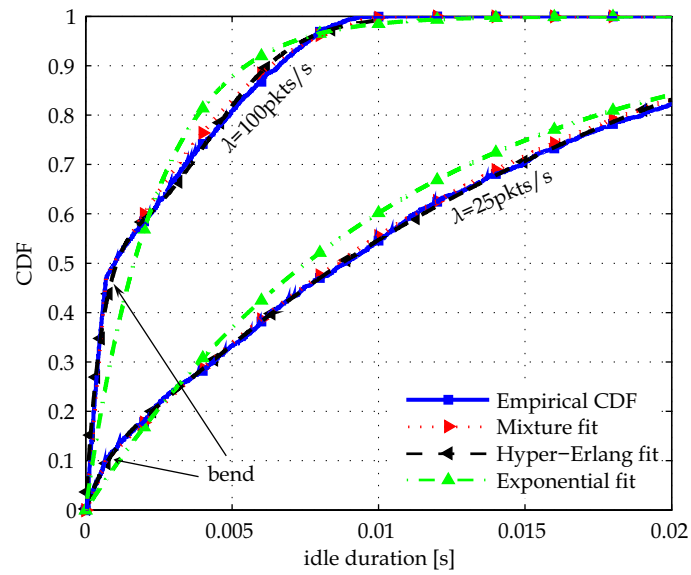


Fig. 4. Goodness-of-fit analysis for constant payload UDP traffic.


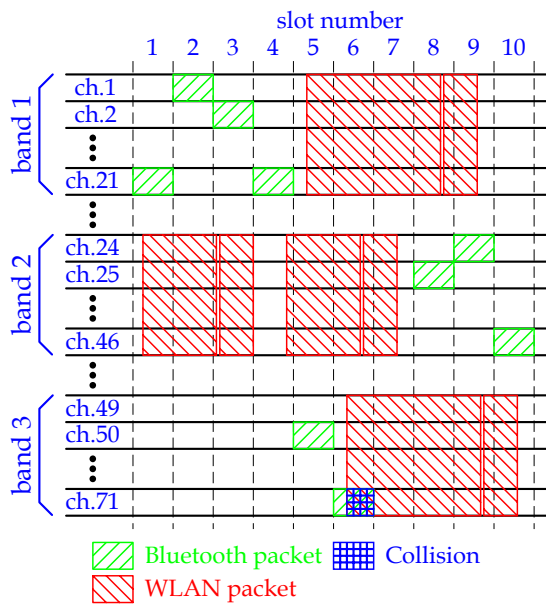
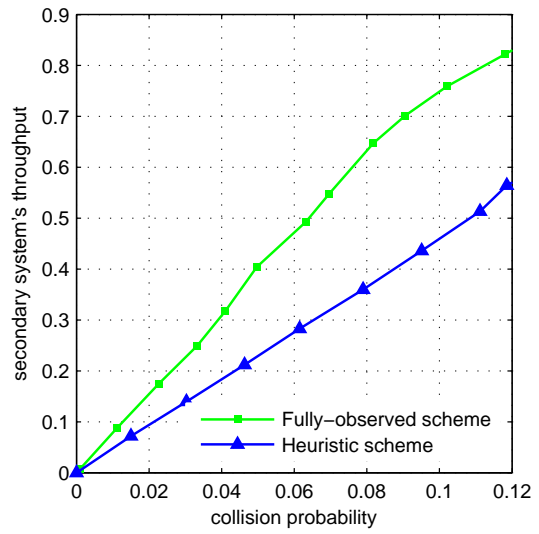


Fig. 5. Nonstationary traffic. While the hyper-Erlang fit shows an appropriate goodness-of-fit the exponential distribution again exceeds the threshold defined by the K-S test.



(a) Adaptive frequency hopping



(b) Performance evaluation

Fig. 6. Adaptive frequency hopping. The Bluetooth transmitter preferably hops to bands currently not occupied by the WLAN.