Per-Frame Energy Consumption in 802.11 Devices and Its Implication on Modeling and Design

Pablo Serrano, Member, IEEE, Andres Garcia-Saavedra, Giuseppe Bianchi, Albert Banchs, Senior Member, IEEE, and Arturo Azcorra, Senior Member, IEEE

Abstract—This paper provides an in-depth understanding of the per-frame energy consumption behavior in 802.11 wireless LAN devices. Extensive measurements are performed for seven devices of different types (wireless routers, smartphones, tablets, and embedded devices) and for both UDP and TCP traffic. Experimental results unveil that a substantial fraction of energy consumption, hereafter descriptively named *cross-factor*, may be ascribed to each individual frame while it *crosses* the protocol stack (OS, driver, NIC) and is *independent* of the frame size. Our findings, summarized in a convenient energy consumption model, contrast traditional models that (implicitly) amortize such energy cost component in a fixed baseline cost or in a toll proportional to the frame size and raise the alert that, in some cases, conclusions drawn using traditional energy models may be fallacious.

Index Terms—Energy efficiency, energy measurements, energy modeling, power consumption, wireless networks.

I. INTRODUCTION

T HE INCREASE in energy density of current state of the art (lithium-ion) batteries is far from following Moore's Law, the current challenge being "just" a twofold density increase in the next 10 years [2]. This is not a good technological premise behind the energy greediness of wireless connectivity, second only to that required to backlight displays in most handheld devices. Moreover, battery-powered wireless devices are becoming ubiquitous and are frequently part of the network infrastructure itself, where relays or opportunistic intermediaries are widely considered in ad hoc, mesh, delay-tolerant network scenarios or emergency deployments. Additionally, for wireless systems powered via the grid, it is also important to reduce energy consumption to cut down the cost of operating the wireless infrastructure as well as its energy footprint.

It is hence not a surprise that a huge research effort has been dedicated to thoroughly understand the power consumption behavior of *real-world* wireless devices, as well as to find ways

Manuscript received August 24, 2013; revised March 03, 2014; accepted April 18, 2014; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor M. Meo. Date of publication May 22, 2014; date of current version August 14, 2015. This work was supported by the European Community under the CROWD Project (FP7-ICT-318115).

P. Serrano is with the University Carlos III of Madrid, 28911 Leganes, Spain. A. Garcia-Saavedra was with the University Carlos III of Madrid, 28911 Leganes, Spain. He is now with the Hamilton Institute, NUI Maynooth, Maynooth, Ireland.

G. Bianchi is with the University of Roma Tor Vergata, 00133 Rome, Italy.

A. Banchs and A. Azcorra are with the University Carlos III of Madrid, 28911 Leganes, Spain, and also with IMDEA Networks Institute, 28918 Barcelona, Spain.

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

Digital Object Identifier 10.1109/TNET.2014.2322262

for reducing energy consumption in the wireless access operation [3], [4]. For instance, with reference to the 802.11 WLAN (WiFi) technology [5], indeed the focus of this paper, energy efficiency improvements span very diverse aspects of the 802.11 operation, from management procedures [6], to usage of opportunistic relays [7] or infrastructure on demand [8], to PHY [9] and MAC [10] parameters' optimization, etc.

Obviously, a *quantitative* treatment of the attainable energy improvements is greatly simplified by the availability of realistic and accurate energy models, also considering that *fine-grained* per-frame experimental measurements (versus coarse aggregate power consumption statistics) are nontrivial to achieve. Most literature works, including but not limited to [10]–[18], base their analyses, optimizations, or algorithm/protocol designs on the widely accepted paradigm that the energy toll may be ascribed to two components: a baseline one, plus a second one linear with (transmission/reception) air time. The specific model's parameters can be gathered from data sheets [19] or experimental measurements [20], [21].

With such a widespread acceptance, questioning the abovementioned *classical* energy model seems tough. Actually, such a model makes perfect sense if we just focus on the network interface card consumption. However, in practice, processing in the host device drains energy as well. Thus, the question at stake is whether (and to what extent) there is some energy toll *in the device*, which is imputable to TX/RX processing, but is *improperly accounted in such a classical model*, e.g., because it can be neither considered: 1) independent of the radio operation and thus (implicitly) accounted in the fixed baseline energy consumption component; nor 2) strictly proportional to the traffic load in bytes, hence (implicitly) accounted in the linear air time energy cost component. Our paper not only raises this question (apparently unnoticed so far), but, more significantly, gives a (we believe) compelling answer, via extensive and tailored experiments that provide a detailed *anatomy* of the energy consumption in the protocol stack.

Two major findings appear to emerge. First, a *substantial* energy consumption occurs while a frame is delivered across the protocol stack, namely from the operating system to the driver to the NIC (and conversely for reception). Such "new" energy cost component, descriptively referred to as *cross-factor*, cannot be neglected; on the contrary, in some experiments it even accounts to *more than half* of the per-frame energy cost. Second, such cross-factor can be neither dealt with as an extra baseline component, nor (perhaps more surprisingly) as a cost proportional to the traffic load. Actually, this energy toll appears mostly *associated to the very fact that a frame is handled*, i.e., irrespective of the actual frame size in bytes. Traditional models that (implicitly) account for such cost as proportional to the load may work

1063-6692 © 2014 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications standards/publications/rights/index.html for more information.

as long as the frame size is fixed, but would yield incorrect results in the general case of variable frame sizes.

Our findings, wrapped into a new energy model, have a twofold implication. First, they suggest new energy reduction strategies, such as *batching* packets while they travel across the protocol stack, or avoiding stack crossing when possible. Second, the fact that a substantial amount of energy is drained by the processing of packet units, independent of their size or air time, may invalidate some specific optimizations proposed in the past. For instance, we show that energy-efficient optimizations leveraging relay nodes may yield *qualitatively different* conclusions when the cross-factor energy component is accounted for. In more detail, this paper contributes as follows.

Unveiling the Cross-Factor and New Energy Model: Targeted measurements reveal that a substantial fraction of energy is consumed by the processing of packets throughout the protocol stack. Most notably, such *cross-factor* appears to be a fixed per-frame toll, independent of the frame size.¹ We summarize these findings in a new and more realistic energy model that overcomes traditional models limited to NIC consumption [10]–[18].

Model Parametrization and Cross-Factor Quantification: A very extensive measurement campaign is performed for seven different devices, including qualitatively diverse ones, such as smartphones, tablets, embedded devices, wireless routers, and access points, using different OSs and HW. Experiments are performed for both UDP and TCP. On one side, a consistent *qualitative behavior* in the energy consumption is found for all devices, thus confirming the general applicability of our proposed model; on the other side, results show noticeable quantitative differences in the model's parameters, duly derived for *all* the seven devices.

Practical Implications: Focusing on selected use-case examples, we show that some energy optimizations proposed in the past may yield fundamentally different conclusions when revisited with the awareness of our more realistic energy consumption model. Our findings appear to raise the alert that there might be other cases where past conclusions should be reconsidered. Additionally, we discuss possible new means (entailing simplifications in the crossing of the protocol stack) to take advantage of our findings for improved energy efficiency. Every considered use case example is assessed for all the seven considered devices to gauge their quantitative impact.

Power Consumption Characterization: As a side contribution, we present a measurement methodology that, in contrast to previous works: 1) provides accurate fine-grained per-packet energy measurements; and 2) characterizes the *total* device power consumption instead of just that consumed by the wireless interface. For those setups in which the accuracy of the measurement device is not sufficient, we exploit techniques to reduce uncertainties due to measurements inaccuracy.

II. RELATED WORK

Energy Consumption of Devices: A number of previous works in the area analyze, like us, the consumption of

the complete device, either a laptop [21]-[23] or a mobile phone [24], [25]. Some of these works deal with specific issues, such as quantification of the consumption of components other than wireless interfaces (e.g., CPU, screen, memory) [24], power consumption measurements via available APIs for estimating the battery discharge state [23], assessment of tradeoff between CPU consumption due to data compression, and wireless consumption due to data transmission [22], but do not tackle the per-frame energy consumption domain. Only [21] briefly mentions that the energy consumption associated to packet processing might be nonnegligible, but does not provide any measurement or evidence. Reference [25] finds that message size can have a nonintuitive impact on the energy consumption, but their guess is either the existence of some power management threshold or a bug in the wireless firmware (indeed, energy bugs in mobile devices are a real concern [26]). We distinguish from these works in the fact that we perform a fine-grained per-packet energy consumption decomposition, versus their energy consumption analyses on a much coarser scale.²

Energy Consumption of Interfaces: Unlike the previous papers, most characterizations of the wireless interface consumption are done on a per-packet basis. The seminal work of [20] shows that transmission/reception of an 802.11 frame has a linear dependency on its length. This result is caused by the four different states a wireless NIC can be in, namely: sleep, idle, receiving, and transmitting. Reference [20] also identifies a fixed cost per frame, caused by control frames (e.g., RTS/CTS). The results are extended in [27] for different modulation and coding schemes and transmission power configurations, and a similar approach is followed in a recent work [28] for the case of 802.11n. While in these cases the 802.11 interface is treated as a whole, [29] distinguishes between the (approximately constant) application-specific integrated circuit (ASIC) consumption and the power amplifier (PA) consumption occurring only outside idle periods. None of these works analyzes the energy consumption of a frame as it is delivered to/from the NIC.

Energy Consumption Models: The (implicit or explicit) assumption of all previous energy consumption models [10]–[18] is that the PA operations dominate the consumption of the whole device, which allows to model consumption with a finite number of states, e.g., {active, idle} [11], [17], {transmission, reception, idle} [13]–[15], and so on. More specifically, the common approach followed by all these papers (as well as that recently included in the NS3 network simulator [30]) is to model the NIC consumption using data sheet parameters [19], and add to this a fixed amount to account for the nonwireless power consumption of the device. In [18], the authors propose an extended model that accounts for the power conversion efficiency of the PA, but eventually the model suffers from the same limitations. As we will see in this paper, these energy consumption models fail to capture crucial aspects of how energy is consumed in real world *devices*, and therefore their use might bias conclusions.

Energy-Efficient Mechanisms: Energy-efficient schemes have been proposed at all layers of the 802.11 stack. Starting

 $^{^{1}}A$ posteriori, this fact might be eventually considered "obvious": Indeed, the most expensive OS/driver operations associated to the processing of a frame do not depend on the frame size. However, it seems fair to say that this fact was overlooked by the wireless community so far, and our merit is to raise attention on this, as well as provide its actual modeling and quantification.

²Indeed, setting up a measurement system able to capture the power consumed by the entire device with the required level of accuracy is very challenging and requires know-how on power measurements that networking researchers do not usually have.

Device	Name	WiFi NIC	Mem./CPU	OS
Wireless	Soekris	Atheros AR5414	128MB	Gentoo 10.0
router	net4826-48	802.11a/5GHz	233MHz	
Wireless Soekris		Atheros AR5414	128MB	OpenBSD 5.1
router	net4826-48	(802.11a/5GHz	233MHz	
Wireless	Soekris	Atheros AR9220	128MB	Gentoo 10.0
router	net4826-48	802.11n/5GHz	233MHz	(kernel 2.6.32)
Access	Linksys	Broadcom BCM4320	16MB	OpenWRT
Point	WRT54GL	802.11g/2.4GHz	200MHz	Backfire
Smart	HTC	Texas Inst. WL1273	384MB	Android 2.2
Phone	Legend	802.11g/2.4GHz	+600MHz	
Tablet	Galaxy	Broadcom BCM4334	2GB	Android 4.1.1
	Note 10.1	802.11a/2.4GHz	+1.4GHz	
Embedded	Raspberry	Ralink RT5370	512MB	Raspbian
Device	Pi	802.11g/2.4GHz	700MHz	Wheezy

TABLE I Platforms Considered

from the lowest layer, [9] precomputes the optimal rate-power configuration for each data frame. Several works aim at reducing the energy wastage in the WLAN by adapting the contention parameters [10], [14] or extending the backoff operation [12], [16]. Cooperative relaying for energy efficiency is analyzed in [7]; reference [31] exploits idle period predictions to switch from active to sleep states. Increasing the sleep state time is the main energy saving target in the standard power saving mechanism (PSM) [4] and in traffic management (and shaping) schemes such as NAPman [6] or in "infrastructure on demand" schemes [8] devised to (de)activate access points based on client load. All these proposals are based either on: 1) the energy consumption of the PA, which might be detailed but underestimates the consumption of the complete device; or on 2) the coarse-grained estimated consumption of the complete device, which precludes a thorough understanding of the per-packet delivery implications.

III. ENERGY MEASUREMENT TESTBED

802.11 Devices Under Study: In our experimental analysis, baseline results are obtained using a Soekris net4826-48 device, equipped with an Atheros AR5414-based 802.11a/b/g Mini-PCI card, and configured to use the 802.11a PHY. The hardware comprises a 233-MHz AMD Geode SC1100 CPU, two Mini-PCI sockets, 128 MB SDRAM and 256-MB compact flash circuits for data storage, extended with a 2-GB USB drive. The OS is a Gentoo 10.0 Linux (kernel 2.6.24), and the driver is MadWifi v0.9.4.

To rule out the possibility that our findings could be biased by the selected HW, OS/driver SW, or WLAN band/card/PHY, most of the paper's experiments have been performed for other platforms (see Table I). This set of platforms includes many different types of devices, such as wireless routers, access points, smartphones, tablets, and embedded devices.³

In all experiments, traffic was generated using mgen.⁴ Additional devices in monitor mode were employed to track wireless channel activity and confirm it was caused only by our experiments, and that no packets were dropped at any layer of the protocol stack, events which would have biased our findings.

Power Consumption Measurements: Depending on the platform, power consumption was measured via two instruments. The Monsoon FTA22D power meter⁵ supplies a stable voltage to the device and samples the power consumption at 5 kHz with high accuracy $(\pm 1\%)$, thus providing very reliable fine-grained power measurements. However, this meter only supplies voltages of up to 5.4 V and could thus be employed only for the tablet Samsung Galaxy Note 10.1, the smartphone HTC Legend, and the Raspberry Pi. Measurements for the Soekris, Alix, and Linksys platforms, which require higher voltage, were carried out using a PCE PA-6000 power meter,⁶ which has a much coarser granularity than the Monsoon meter (namely about three samples/second). The PCE meter was powered with 6 AA batteries (ac supply via the wall socket would reduce accuracy), and we employed a Protek 3033B device to power the wireless device. Both for the Monsoon and the PCE meters, each sample provides the result of integrating the power consumed over the corresponding sampling period, and hence the average power consumption can be obtained by averaging all the samples. This setting permits to perform measurements without dismantling the device, as required by some specialized equipments, which would restrict experimentation to, e.g., devices using card extenders.

A major practical challenge was to reduce the errors associated to the measurements of the PCE PA-6000 power meter, which natively provides an accuracy of $\Delta v = 0.1$ V for the voltage and $\Delta i = 0.01$ A for the current. Taking into account the well-known fact that the relative error for the product p = v $\cdot i$ is approximated by the sum of the relative errors for v and i, these inaccuracies yield a relative error in the measured power of above 5% for our typical baseline power measurements, undermining our ability to quantify small, but for our purposes meaningful, trends (e.g., power consumption variations for an increased frame size). The methodology that we followed to improve the accuracy of our measurements is to use, instead of a single device, K devices running the same experiment in parallel over different noninterfering wireless channels. With this, the relative error of the measurements is reduced by 1/K. For 802.11g, we used the three noninterfering channels, whereas we used K = 4 for 802.11a, thus improving accuracy to about 2% (see [1] for a more detailed explanation). This provides an accuracy comparable to that of [20], which only measures the power consumption of the wireless card.

IV. ENERGY CONSUMPTION ANATOMY

In order to characterize the power consumption of 802.11 devices, we have conducted an in-depth experimental investigation of the considered 802.11 devices. Even though we conducted our measurements for all the platforms under study, for space reasons here we only present the detailed measurements for one of the platforms (the Soekris device with the Linux OS, in Sections IV-A–IV-C) and then provide a summary of the results obtained with other platforms/OS (Section IV-E). All the measurements presented here are for the devices operating under the infrastructure mode. However, we verified that the devices show a very similar behavior when operating under the *ad hoc* and monitor modes.

⁴http://cs.itd.nrl.navy.mil/work/mgen

³In addition to the results reported here for different platforms, the generality of our model is further confirmed by the some additional results reported in [1] as well as the measurements carried out independently in [32]

⁵http://www.msoon.com/LabEquipment/PowerMonitor

⁶http://www.industrial-needs.com/technical-data/power-analyser-PCE-PA-6000.htm

Config.	Description	Cons. (W)
w/o card	no NIC connected	$2.29\pm2.2\%$
WiFi off	NIC connected	$ 2.58 \pm 2.0\% $
	driver not loaded	(+0.29)
Idle (ρ_{id})	NIC activated+associated to AP	$3.56 \pm 1.7\%$
	no RX/TX besides beacons	(+0.98)

TABLE II SOEKRIS BASELINE CONSUMPTION PROFILE

A prerequisite for the characterization of 802.11 devices is to quantify their "baseline" power consumption, i.e., when the devices neither send nor receive traffic. Table II reports measurements for the Soekris platform in three "baseline" configurations. Note that plugging the wireless card ("WiFi off") increases consumption by 0.29 W (+12.6%), whereas loading the driver and associating to an AP ("Idle") further increases the consumption by 0.98 W (an extra 25% increment). The power consumed in the "Idle" state, named ρ_{id} , will be used as baseline reference in what follows.

A. Understanding Transmission Costs

Results in this section aim at characterizing the energy cost of transmissions. To avoid biasing results with the ACK reception cost (separately quantified later on), the following results are obtained for unicast *unacknowledged* frames. ACKs have been disabled by setting the noackpolicy bit of the WMM (Wi-Fi Multimedia) parameters for the access point parameter set: This introduces an Information Element in beacon frames that prevents associated stations from replying with ACKs (confirmed by sniffed traces). Each result is obtained by measuring the power consumption over a 20-s experiment. Unless otherwise specified, all experiments are based on UDP traffic.

Transmission Power Consumption Patterns: Measurements of total device power consumption have been carried out spanning all the combinations of the following four parameters: 1) frame size L from 80 to 1500 in steps of 20 B; 2) modulation and coding schemes MCS \in {6, 12, 24, 48} Mb/s; 3) configured transmission power txpower \in {6, 9, 12, 15} dBm; and 4) frame generation rates λ_g from 100 to 2000 in steps of 100 frames per second (fps).

As we observe from the above measurements that, for a given txpower, λ_{g} , and MCS configuration, power consumption increases proportionally to the air time, it turns out that the most insightful way to represent such results is via a power/airtime *plot* shown in Fig. 1. Such a plot reports the average power consumed by the whole device, versus the channel airtime percentage $\tau_{tx} = \lambda_g T_L$, where λ_g is the frame generation rate (kept below saturation), and $T_L = T_{PLCP} + (H + L)/MCS$ is the time required to transmit a frame of size L using the modulation and coding scheme MCS, duly accounting for the Physical Layer Convergence Protocol preamble T_{PLCP} , and the MAC overhead H (MAC header plus FCS). The plot shows the values resulting from applying a simple linear regression to all the samples obtained for a given configuration, as well as the corresponding standard asymptotic error, which is computed with the assumption that the error of each individual sample is normally distributed [33]. For reference purposes, the plot also depicts the baseline power consumption ρ_{id} when the target device is in "Idle" state. From the plots, it clearly appears that the total



Fig. 1. Total power consumed by (unacknowledged) transmissions versus airtime percentage $\tau_{\rm tx}$.

device power consumption, denoted \mathbf{P} , is articulated into three main components

$$\mathbf{P} = \rho_{\rm id} + P_{\rm tx} + P_{\rm xg}(\lambda_{\rm g}) \tag{1}$$

where:

- The first component, ρ_{id} , is the (platform-specific) baseline power consumption.
- The second component, P_{tx} , is the classical one in traditional energy consumption models, which linearly grows with the airtime percentage τ_{tx} , i.e., $P_{tx} = \rho_{tx}\tau_{tx}$. The slope ρ_{tx} depends on the radio transmission parameters MCS and txpower: the greater MCS and/or txpower, the greater the slope.
- The third component, $P_{xg}(\lambda_g)$, accounts for the fact that the above linear trend *starts from a relatively large positive offset* over the baseline level ρ_{id} (+12% and +35% increment for 400 and 1200 fps, respectively); offset that is *not accounted by classical energy models* [10]–[18]. Moreover, Fig. 1 suggests that such component depends only on the frame generation rate λ_g .

Results obtained for the remaining target platforms confirm the same behavior as the one described by (1) for all platforms, each of them with its specific parameter values (reported in Section IV-E).

Per-Frame Processing Toll: To more closely investigate the nature of such emerging power consumption offset $P_{xg}(\lambda_g)$, Fig. 2 plots its value obtained from several measurements taken for 12 different configurations of the NIC parameters (MCS, txpower). The plot clearly shows that $P_{xg}(\lambda_g)$ is *proportional* to the frame generation rate λ_g , whereas it is practically *independent of* the frame size or the radio settings.

Thus, if we denote with $\gamma_{xg} = P_{xg}(\lambda_g)/\lambda_g$ the proportionality constant, it appears that γ_{xg} is the *energy toll associated to the processing of each individual frame*, irrespective of its size or radio transmission parameters. Note that this energy toll is *not* associated to protocol operations such as RTS/CTS or ACKs, indeed disabled in such experiments. For reasons that will become clear throughout this section, we call this per-frame energy toll the *cross-factor*.



Fig. 2. Relation between $P_{xg}(\lambda_g)$ and λ_g .



Fig. 3. Interfaces/modules crossed during transmission.

Cross-Factor: To grasp a deeper insight on the reasons behind the cross-factor γ_{xg} , we have engineered tailored measurements devised to *quantify* how the energy toll splits across the frame processing chain along the protocol stack depicted in Fig. 3. Specifically, we have run three sets of experiments, where we discard packets at a given level of the stack and measure the power consumed up to that level.

- App.: Packets are generated by mgen, but are discarded before being delivered to the OS, i.e., at the mark (a) in Fig. 3, by sending them to the "sink device" (/dev/null).
- **TCP/IP**: Packets are discarded at the bottom of the TCP/IP stack [mark (b) in Fig. 3], by deactivating the ARP lookup function, so that the device cannot retrieve the MAC destination from the ARP cache and therefore must drop the frame.
- Driver: Packets are discarded after the MadWifi driver's processing [mark (c) in Fig. 3], by commenting the hard-start command that performs the actual delivery of the frame to the NIC.

Representative measurements (energy per frame) are shown in Fig. 4, along with the total energy consumption per properly transmitted frame ("Total") and the values predicted by applying (1) ("Model").

The figure clearly shows that the energy toll due to frame processing is practically *independent of* the frame generation rate *and* the frame size.⁷ Moreover, it shows that the energy consumed while crossing the host device stack (i.e., up to the



Fig. 4. Per-frame energy cost in transmission.

driver included) is *substantial*, around 0.75 mJ per frame, and may become the *major* energy cost in several scenarios (e.g., short packets and/or large MCS—in essence short airtime).

Finally, even if direct measurements cannot be attained below the driver level, Fig. 4 shows that a further constant per-frame energy drain occurs at the driver-to-NIC interface level and/or below. Its quantification may be estimated by measuring the energy consumed with very short packets and large MCS, as wireless transmission cost is marginal in this case (very small airtime). Summarizing, for the Soekris device, the cross-factor coefficient amounts to about 0.93 mJ/frame. Such per-frame processing cost appears to roughly split as follows: 24% application; 33% TCP/IP stack, 21% driver, and 22% NIC.

Retransmissions: The above results seem to suggest that retransmissions at the MAC layer, e.g., caused by an unacknowledged transmission, should not be affected by the cross-factor toll. This can be verified by *provisionally assuming* that this is the case, i.e., modeling retransmission cost as purely due to the over the air transmission cost component, and then checking whether the resulting model matches experimental measurements. Hence, let $P_{\rm retx}$ be the power drained by retransmissions and assume that

$$P_{\rm retx} = R \cdot \rho_{\rm tx} \tau_{\rm tx} = R \cdot \rho_{\rm tx} \lambda_{\rm g} T_L \tag{2}$$

where R is the number of retransmissions. Then, the *total* power consumed by packets retransmitted R times, P_R , is readily obtained as

$$P_R = P_{\rm xg}(\lambda_{\rm g}) + P_{\rm tx} + P_{\rm retx} \tag{3}$$

where the first addendum is the per-frame processing toll (paid once), the second addendum is the power consumed by the very first transmission, and the last addendum is the extra retransmission cost as per (2). Fig. 5 compares the modeling prediction of (3) with the power (additional to the baseline component ρ_{id}) consumed by a device configured to send 1400-B frames generated at a rate of 80 fps to fake addresses (to prevent the reception of ACKs). The number of allotted retransmissions R (configured via the ah_setupTxDesc driver's descriptor) was varied from 0 to 5, and, for simplicity (i.e., to avoid the need to nontrivially configure the driver so as to prevent MCS downgrade in front of persistent losses), frames were transmitted using the 6-Mb/s basic MCS. As shown in

⁷Although it is barely appreciable in the figure, there is actually a small component in the packet processing that is proportional to the packet size. In our model, this is captured in $\rho_{tx} \tau_{tx}$: Even though this term mainly relates to the power consumed by transmissions, it increases linearly with the packet size and hence indirectly captures any other power consumption component that follows such behavior.



Fig. 5. Impact of retransmissions on power consumption.



Fig. 6. Power consumed by (unacknowledged) reception versus airtime.

Fig. 5, theoretical results tightly match the experimental measurements, thus confirming that the cross-factor has (if any) a negligible impact on retransmissions.

B. Reception Power Consumption Analysis

The analysis of the power consumption of the device while *receiving* frames is somewhat dual to that carried out in depth for the previous transmission case, and hence will be presented with less detail for the sake of conciseness. We use the same configurations of MCS and txpower as in Section IV-A (ACK disabled as well), with different combinations of the frame length L and frame *reception* rate λ_r . The resulting *power/airtime* plot is shown in Fig. 6, airtime now given by $\tau_{rx} = \lambda_r T_L$. Results for the txpower parameter are not shown, as this parameter does not affect the receiver's power consumption (as indeed well known from [27]).

Fig. 6 exhibits the same qualitative pattern found in the transmission scenario. The increment of the power consumption over ρ_{id} is composed of two components: a first one linear with the airtime and accounting for the power required to receive frames, P_{rx} (indeed in line with traditional energy models), and a second one proportional to the number of frames received and accounting for the cross-factor energy toll, $P_{xr}(\lambda_r)$. The total power consumption at the receiving side is thus

$$\mathbf{P} = \rho_{\rm id} + P_{\rm rx} + P_{\rm xr}(\lambda_{\rm r}) = \rho_{\rm id} + \rho_{\rm rx}\tau_{\rm rx} + \lambda_{\rm r}\gamma_{\rm xr} \quad (4)$$



Fig. 7. Impact of sending ACKs on the receiver.

where γ_{xr} is the cross-factor in reception, i.e., the energy processing toll to deliver the received frame across the protocol stack, and $\rho_{rx}\tau_{rx}$ is the traditional reception cost proportional to the airtime. Again, Fig. 6 confirms that results from the above equation (lines) match well the experimental measurements (symbols).

C. Characterization of ACKs, Control Frames, Collisions

To complete our analysis, it remains to characterize the additional power consumed for sending/receiving acknowledgments (Sections IV-A and IV-B have considered unacknowledged operation) and the power consumption caused by overhearing a collision.

ACKs and Other Control Frames: Since ACK frames, like retransmissions, do not have to cross the stack but are internally generated by the NIC, we make the hypothesis that their power consumption can be characterized by just the cost of the relevant ACK transmission or reception. Under such hypothesis, the power consumed for replying with ACKs to received frames (arriving at rate λ_r) is trivially given by

$$P_{\rm tx,Ack} = \rho_{\rm tx} \lambda_{\rm r} T_{\rm Ack} \tag{5}$$

where $T_{Ack} = T_{PLCP} + ACK/MCS_C$ is the time required to transmit an ACK frame, i.e., a PLCP preamble plus the 14-B ACK frame transmitted at the modulation and coding scheme MCS_C configured for control traffic. Similarly, the power consumed to receive an ACK is readily computed as

$$P_{\rm rx,Ack} = \rho_{\rm rx} \lambda_{\rm g} T_{\rm Ack}.$$
 (6)

Fig. 7 shows the experimental validation for the ACK transmission case. Such experimental results, obtained with $\lambda_r = 1000$ fps, confirm that the measurements match the results predicted by the model, which includes the energy consumed by the reception of frames ($P_{xr} + P_{rx}$) and the transmission of the ACKs ($P_{tx,Ack}$). Similar findings hold for RTS and CTS frames (not shown here for space reasons).

Collisions and Other Transmissions: Finally, we analyze the impact on the energy consumption in reception when the medium is occupied by collisions or by transmissions addressed to another device (i.e., to another MAC address). For this purpose, we configured a communication between two nodes and set up another node to act as interferer. The sender and the interferer were configured with the carrier sense threshold at the

TABLE III Impact of Collisions On Reception

Frames/s	Interference	Frames/s	Measured	Model
sent	rate	received	power	
100 fps	0%	100 fps	3.67 W	3.68 W
200 fps	50%	100 fps	3.67 W	3.68 W
200 fps	100%	0 fps	3.56 W	3.56 W
200 fps	0%	200 fps	3.80 W	3.81 W
400 fps	50%	200 fps	3.80 W	3.81 W
400 fps	100%	0 fps	3.56 W	3.56 W

highest value, which practically results in no carrier sensing, and the interferer used the lowest values for the CW, SIFS, and MCS parameters while deactivating the use of ACKs, resulting in practically continuous transmission. To control the amount of time the interferer was sending data (i.e., the "interference rate"), we used the quiet element option to silence the interface for a given amount of time every beacon period. With this setting, the interferer transmits continuously during a long period, thus emulating the typical behavior of collisions with carrier-sensing that always affect frame transmissions from the beginning.⁸ and varying the relative position of the devices, to find a setting in which simultaneous transmissions resulted in all frames being lost (i.e., no capture effect).

Table III presents the measured power consumption for different sending and interference rates. We observe that: 1) the power consumed in reception depends exclusively on the traffic actually received (see, e.g., when the interference rate goes from 0% to 50%); and 2) collisions have the same impact as an idle medium (e.g., the cases with 100% interference rate coincides with ρ_{id}). Based on this, we conclude that collisions have no practical impact on the energy consumption at the receiver (this is confirmed by the results of the model included in the table). Note that, for the transmitter, collisions have already been modeled in our analysis of retransmissions.

To analyze the impact of the transmissions addressed to another station, we configured a communication between two nodes and measured the energy consumption at a third node that was in the transmission range of this communication. We observed that the energy consumed by this node was the same as if the medium was idle, which shows that transmissions addressed to other stations practically do not consume energy. This is in agreement with our previous results: According to (4), the energy cost of listening to the PLCP plus headers is only 38 μ J/frame (for 6-Mb/s MCS), which has practically no impact on the overall consumption.

D. Impact of CPU Usage

All the experiments shown so far have been performed without any application running on the device. To understand whether the CPU usage by applications has any impact on the conclusions drawn, we performed an additional experiment with a tailored application running in parallel to the transmission/reception of packets. Fig. 8 shows the result of repeating the baseline experiment of Fig. 1 with this application running on the device, calibrated to induce a low (25%), medium (50%),



Fig. 8. Power consumption with an application running on the device.

and large (75%) load on the CPU, respectively. We observe that the energy consumption behavior is the same as in Fig. 1, with the only difference that the baseline consumption ρ_{id} is shifted due to the extra energy consumed by the application: On top of ρ_{id} , there is an offset with respect to the baseline consumption caused by the cross-factor (which takes the same value as in Fig. 1 independent of the CPU load) and then power consumption increases linearly with the airtime. This confirms that the CPU usage has no impact on the power consumption behavior of transmitting/receiving packets.

E. Other Devices

The results provided in Sections IV-A–IV-D have been obtained for one particular hardware platform (the Soekris device) running one specific OS (Linux Gentoo 10.0). To verify that the behavior observed in those sections is not specific of the chosen reference device, OS, or WLAN band/card/PHY, we repeated all the experiments for each of the platforms listed in Table I. Results, not reported for space reasons (a summary of findings per each device is provided in Table IV) fully confirm the qualitative behavior discussed before, thus suggesting that the conclusions drawn above respond to *general energy consumption patterns* of all 802.11 wireless devices (or at least a very wide range).

In addition to their qualitative behavior, it is also interesting to analyze the quantitative differences between the energy components of the different devices. In particular, one of the key results drawn from the previous sections is that the cross-factor has a substantial impact over the total energy consumed by a frame. In order to gain insight on the weight of the cross-factor in different platforms and OS, Fig. 9 depicts, for each of the devices of Table I, how the per-frame energy consumption is decomposed into the following two components: the *cross-factor* component (P_{xg} and P_{xr}) and the transmission/reception component (P_{tx} or P_{rx}). Results are provided for MCS = 48 Mb/s (except for the "Soekris 802.11n," for which we use MCS index 6), *txpower* = 15 dBm, and two different packet lengths (100 and 1500 B).⁹ For those devices that do not implement CPU scaling, we used the default CPU frequency configuration, while for those

⁸Note that, with hidden nodes, interfering frames could also arrive at a later point of the frame transmission. However, the measurements of [34] show that such interfering frames, which do not affect the preamble, need to have a very high power to cause a transmission error. Prior to our measurements, we performed extensive tests using different txpower configurations

⁹Since the transmission/reception component is proportional to the packet length, by evaluating very large (1500 B) and very small (100 B) packets, we cover the two extreme cases for the weight of the *cross-factor*.

Tx/Rx Cross-Factor 100 per-frame energy decomposition (%) 80 60 40 20 0 200 200 Soekris Soekris HTC Galaxy Note 10.1 Raspberry Pi Soekris Linux 802 11r Linksvs OpenBSD Legend

Fig. 9. Fraction of the cross-factor over the total per-frame energy consumption cost.

that implement it (HTC Legend and Galaxy Node), we used the largest CPU frequency.¹⁰

From the results obtained, we observe that the *cross-factor* has a very significant weight in the vast majority of the devices and configurations studied, as it accounts for *more than 50%* of the per-frame energy consumption in most cases. Only for the cases of transmitting large packets with the Galaxy Note, the HTC Legend, or the Raspberry Pi, the *cross-factor* has a smaller (but still very substantial) weight; in particular, in these cases the *cross-factor* accounts for more than 30% of the energy consumed by a frame.

From the results reported above, we draw the following conclusions: 1) the behavior observed along this section is not specific for the Soekris device, but is generalized to all the devices under study, which includes a very representative set of the types of 802.11 devices of practical interest; 2) the qualitative behavior with different operating systems (Linux and OpenBSD) is the same, and the *cross-factor* is of the same order for both (although, as expected, the specific values are slightly different as a result of the different protocol stack implementation); and 3) the weight of the *cross-factor* on the per-frame energy consumption is very substantial in all cases, even though as expected it is a bit smaller for those devices that rely on a more energy-efficient design.

V. ENERGY CONSUMPTION MODEL

The findings (1)–(6) gathered in Section IV can be conveniently summarized into a *complete model* for the power consumption of 802.11 devices. Indeed, the power consumed by an 802.11 device consists of the following components: 1) the idle consumption, ρ_{id} ; 2) the cross-factor for the packets generated by the application, P_{xg} ; 3) the power required to transmit them, P_{tx} ; 4) the power consumed in retransmissions, P_{retx} ; 5) the power spent in receiving frames, P_{rx} ; 6) the cross-factor for the received frames, P_{xx} ; and 7) the power spent on sending and receiving ACK frames, $P_{rx,Ack}$ and $P_{tx,Ack}$

$$\mathbf{P} = \rho_{\mathrm{id}} + P_{\mathrm{tx}} + P_{\mathrm{xg}} + P_{\mathrm{retx}} + P_{\mathrm{rx}} + P_{\mathrm{xr}} + P_{\mathrm{rx,Ack}} + P_{\mathrm{tx,Ack}}.$$
 (7)

¹⁰CPU scaling adaptively sets the CPU frequency based on the CPU load. As we observed that this leads to quite unpredictable consumption, we disabled this function in our experiments and set the CPU frequency to a fixed value. Table IV shows the impact of different CPU frequency values on energy consumption. By substituting the expressions obtained in Section IV for all the above components and regrouping the terms, we obtain

$$\mathbf{P} = \rho_{\rm id} + \rho_{\rm tx} (\lambda_{\rm g} T_L + \lambda_{\rm g} T_L R + \lambda_{\rm r} T_{\rm Ack}) + \rho_{\rm rx} (\lambda_{\rm r} T_L + \lambda_{\rm g} T_{\rm Ack}) + \gamma_{\rm xg} \lambda_{\rm g} + \gamma_{\rm xr} \lambda_{\rm r}.$$
(8)

By taking into account that $\lambda_g T_L + \lambda_g T_L R + \lambda_r T_{Ack}$ corresponds to the transmission airtime percentage τ_{tx} , and $\lambda_r T_L + \lambda_g T_{Ack}$ to reception airtime percentage τ_{rx} , the above equation can be rewritten as

$$\mathbf{P} = \rho_{\rm id} + \rho_{\rm tx} \tau_{\rm tx} + \rho_{\rm rx} \tau_{\rm rx} + \gamma_{\rm xg} \lambda_{\rm g} + \gamma_{\rm xr} \lambda_{\rm r}.$$
(9)

The above expression gives the model for the power consumption of an 802.11 device that we propose in this paper. Out of the nine variables in this expression, five of them (ρ_{id} , ρ_{tx} , ρ_{rx} , γ_{xr} , and γ_{xg}) are constant parameters that depend on the device and the configuration of its communication parameters, while the other four parameters (τ_{tx} , τ_{rx} , λ_r , and λ_g) are variables that depend on the number of stations in the WLAN and their traffic generation behavior.

As already mentioned in Section IV, the key difference between the above model and the 'traditional' one used in many previous works [10]–[18] is that the traditional model only includes the first three components (namely ρ_{id} , $\rho_{tx}\tau_{tx}$, and $\rho_{rx}\tau_{rx}$), while our model adds to these three components two additional ones ($\gamma_{xr}\lambda_r$ and $\gamma_{xg}\lambda_g$). As shown by our measurements, these two additional components account for a very significant portion of the power consumption, which renderss the traditional model highly inaccurate.

Parametrization of the Model: In order to characterize the power consumption of an 802.11 device, we need to parametrize the five constant parameters in (9) for the device. One of the contributions of this paper is to obtain the values of these five parameters for all the devices under study: Soekris (with Linux and OpenBSD), Alix, Linksys, HTC Legend, Samsung Galaxy Note and Raspberry Pi. The numerical values are given in Table IV for the different MCS and txpower configurations, as well as different CPU frequencies for those devices that support CPU scaling.

To obtain these values, we have applied the expressions for the simple linear regression and the standard asymptotic

 TABLE IV

 PARAMETRIZATION OF THE ENERGY MODEL FOR ALL THE DEVICES UNDER STUDY

	MCS			6 Mbps	12 Mbps	24 Mbps	48 Mbps
	ρ_{rx} (W)			$0.16\pm8\%$	$0.27 \pm 5.6\%$	$0.6 \pm 11\%$	$1.14 \pm 3.5\%$
Soekris	$ ho_{tx}$ (W)		6 dBm 9 dBm 12 dBm 15 dBm	$\begin{array}{c} 0.52 \pm 3.1\% \\ 0.57 \pm 2.1\% \\ 0.70 \pm 1.7\% \\ 0.86 \pm 2.2\% \end{array}$	$\begin{array}{c} 0.55 \pm 4.6\% \\ 0.59 \pm 1.8\% \\ 0.73 \pm 2.2\% \\ 0.89 \pm 2.3\% \end{array}$	$\begin{array}{c} 0.81 \pm 5.3\% \\ 0.88 \pm 2.3\% \\ 1.02 \pm 2.8\% \\ 1.17 \pm 2.5\% \end{array}$	$\begin{array}{c} 1.2 \pm 1.6\% \\ 1.24 \pm 2.7\% \\ 1.37 \pm 3.1\% \\ 1.58 \pm 3.3\% \end{array}$
	ρ_{id} (W)	Linux OpenBSD	$\begin{array}{ } 3.56 \pm 0.6\% \\ 3.48 \pm 0.8\% \end{array}$	γ_{xg} (mJ)	$\begin{array}{c} 0.93 \pm 1.2\% \\ 1.27 \pm 1.9\% \end{array}$	γ_{xr} (mJ)	$\begin{array}{c} 0.93 \pm 2.2\% \\ 1.26 \pm 2.0\% \end{array}$
	MCS			6 Mbps	12 Mbps	24 Mbps	48 Mbps
	ρ_{rx} (W)			$0.19 \pm 5.3\%$	$0.29 \pm 3.4\%$	$0.53 \pm 2.3\%$	$0.74 \pm 4.4\%$
Linksys	$ ho_{tx}$ (W)		6 dBm 9 dBm 12 dBm 15 dBm	$\begin{array}{c} 0.70 \pm 1.1\% \\ 0.77 \pm 1.4\% \\ 0.84 \pm 1.2\% \\ 0.97 \pm 0.9\% \end{array}$	$\begin{array}{c} 0.72 \pm 2.2\% \\ 0.81 \pm 2.6\% \\ 0.85 \pm 1.5\% \\ 1.0 \pm 1.5\% \end{array}$	$\begin{array}{c} 0.75 \pm 2.0\% \\ 0.84 \pm 2.3\% \\ 0.92 \pm 2.4\% \\ 1.04 \pm 2.1\% \end{array}$	$\begin{array}{c} 0.81 \pm 3.7\% \\ 0.88 \pm 3.4\% \\ 0.99 \pm 4.0\% \\ 1.08 \pm 3.7\% \end{array}$
	ρ_{id} (W)		$2.73 \pm 0.4\%$	$ \gamma_{xg} \text{ (mJ)}$	$0.46 \pm 3.3\%$	$ \gamma_{xr} \text{ (mJ)}$	$0.43 \pm 4.2\%$
	MCS*			MCS 0 (6.5 Mbps)	MCS 6 (58.5 Mbps)	MCS 8 (13 Mbps)	MCS 10 (39 Mbps)
1n)	ρ_{rx} (W)			$0.23 \pm 11.7\%$	$1.08 \pm 8.3\%$	$0.61 \pm 7.4\%$	$1.10 \pm 6.4\%$
cekris (802.1	$ ho_{tx}$ (W)		6 dBm 9 dBm 12 dBm 15 dBm	$\begin{array}{c} 1.46 \pm 2.2\% \\ 1.47 \pm 2.0\% \\ 1.52 \pm 2.4\% \\ 1.59 \pm 2.0\% \end{array}$	$2.67 \pm 3.4\% \\ 2.68 \pm 6.7\% \\ 2.73 \pm 6.6\% \\ 2.80 \pm 3.9\%$	$\begin{array}{c} 1.46 \pm 2.5\% \\ 1.47 \pm 3.3\% \\ 1.52 \pm 2.5\% \\ 1.60 \pm 2.5\% \end{array}$	$\begin{array}{c} 2.13 \pm 4.2\% \\ 2.16 \pm 4.2\% \\ 2.24 \pm 3.6\% \\ 2.37 \pm 3.4\% \end{array}$
Š	ρ_{id} (W)		$3.23 \pm 0.18\%$	$\mid \gamma_{xg} \ ({ m mJ})$	$1.41 \pm 0.71\%$	$\mid \gamma_{xr}$ (mJ)	$0.94\pm3.3\%$
	MCS			6 Mbps	12 Mbps	24 Mbps	48 Mbps
-	ρ_{rx} (mW)			$52.02 \pm 19.6\%$	$96.54 \pm 10.8\%$	$155.94 \pm 13.4\%$	$325.08 \pm 11.3\%$
HTC Legend	$ ho_{tx}$ (mW)		6 dBm 9 dBm 12 dBm 15 dBm	$\begin{array}{c} 412.65 \pm 3.6\% \\ 422.26 \pm 2.8\% \\ 468.26 \pm 2.1\% \\ 500.98 \pm 2.2\% \end{array}$	$\begin{array}{c} 456.60 \pm 4.2\% \\ 471.50 \pm 5.9\% \\ 519.60 \pm 6.0\% \\ 570.78 \pm 2.1\% \end{array}$	$\begin{array}{c} 534.40 \pm 2.4\% \\ 541.62 \pm 3.0\% \\ 559.96 \pm 2.2\% \\ 603.03 \pm 2.1\% \end{array}$	$\begin{array}{c} 665.35 \pm 4.0\% \\ 673.64 \pm 4.2\% \\ 682.50 \pm 4.4\% \\ 695.37 \pm 4.6\% \end{array}$
	$ ho_{id}$ (mW)	245MHz 480MHz 600MHz	$ \begin{vmatrix} 495.32 \pm 1.8\% \\ 558.40 \pm 0.7\% \\ 635.27 \pm 1.6\% \end{vmatrix} $	γ_{xg} (mJ)	$\begin{array}{c} 0.0825 \pm 13.6\% \\ 0.1111 \pm 6.2\% \\ 0.1295 \pm 8.7\% \end{array}$	γ_{xr} (mJ)	$\begin{array}{c} 0.0761 \pm 8.4\% \\ 0.0802 \pm 9.9\% \\ 0.1212 \pm 2.6\% \end{array}$
	MCS			6 Mbps	12 Mbps	24 Mbps	48 Mbps
te 10.1	ρ_{rx} (mW)		600MHz 1GHz 1.4GHz	$\begin{array}{c} 16.16 \pm 13.8\% \\ 24.51 \pm 13.1\% \\ 54.08 \pm 12.0\% \end{array}$	$\begin{array}{c} 21.22 \pm 12.0\% \\ 40.70 \pm 9.9\% \\ 58.20 \pm 11.5\% \end{array}$	$\begin{array}{c} 38.93 \pm 16.5\% \\ 48.47 \pm 6.1\% \\ 82.35 \pm 6.6\% \end{array}$	$\begin{array}{c} 51.82 \pm 14.12\% \\ 71.72 \pm 9.5\% \\ 124.19 \pm 13.3\% \end{array}$
galaxy Not		6 dBm	600MHz 1GHz 1.4GHz	$\begin{array}{c} 604.14 \pm 1.2\% \\ 604.20 \pm 0.9\% \\ 605.15 \pm 1.0\% \end{array}$	$\begin{array}{c} 608.35 \pm 1.3\% \\ 608.82 \pm 1.5\% \\ 609.73 \pm 1.5\% \\ 616.52 \pm 1.4\% \end{array}$	$\begin{array}{c} 613.06 \pm 2.0\% \\ 621.31 \pm 1.9\% \\ 637.82 \pm 2.2\% \end{array}$	$\begin{array}{c} 614.67 \pm 3.6\% \\ 642.05 \pm 3.6\% \\ 677.06 \pm 3.4\% \end{array}$
Samsung ($ ho_{tx}$ (mW)	9 dBm 12 dBm	1GHz 1.4GHz 600MHz 1GHz 1.4GHz	$\begin{array}{c} 610.06 \pm 1.2\%\\ 611.20 \pm 1.1\%\\ 616.22 \pm 1.1\%\\ 620.18 \pm 0.7\%\\ 621.83 \pm 0.8\%\\ 626.49 \pm 1.5\%\end{array}$	$\begin{array}{c} 616.32 \pm 1.4\% \\ 617.30 \pm 1.3\% \\ 617.39 \pm 1.5\% \\ 628.51 \pm 1.2\% \\ 629.79 \pm 1.4\% \\ 630.65 \pm 1.6\% \end{array}$	$\begin{array}{c} 611.56 \pm 1.7\% \\ 621.46 \pm 1.9\% \\ 639.26 \pm 2.2\% \\ 631.53 \pm 1.8\% \\ 633.12 \pm 1.6\% \\ 641.54 \pm 1.6\% \end{array}$	$\begin{array}{c} 620.24 \pm 2.3\%\\ 649.95 \pm 3.2\%\\ 679.23 \pm 3.4\%\\ 640.60 \pm 3.3\%\\ 653.55 \pm 3.1\%\\ 698.42 \pm 3.0\%\end{array}$
		15 dBm	600MHz 1GHz 1.4GHz	$\begin{array}{c} 629.65 \pm 0.8\% \\ 632.96 \pm 1.2\% \\ 643.94 \pm 1.2\% \end{array}$	$\begin{array}{c} 636.94 \pm 1.4\% \\ 649.65 \pm 1.4\% \\ 653.71 \pm 1.7\% \end{array}$	$\begin{array}{c} 641.66 \pm 2.0\% \\ 654.90 \pm 1.8\% \\ 674.27 \pm 1.8\% \end{array}$	$\begin{array}{c} 645.45 \pm 2.5\% \\ 662.94 \pm 2.9\% \\ 708.58 \pm 4.3\% \end{array}$
	$ ho_{id}$ (mW)	600MHz 1.0GHz 1.4GHz	$\begin{array}{c} 436.30 \pm 0.52\% \\ 506.74 \pm 0.28\% \\ 591.59 \pm 0.3\% \end{array}$	γ_{xg} (mJ)	$\begin{array}{c} 0.069 \pm 6.5\% \\ 0.074 \pm 3.8\% \\ 0.088 \pm 4.4\% \end{array}$	γ_{xr} (mJ)	$\begin{array}{c} 0.048 \pm 16.8\% \\ 0.059 \pm 5.2\% \\ 0.098 \pm 6.7\% \end{array}$
	MCS			6 Mbps	12 Mbps	24 Mbps	48 Mbps
	ρ_{rx} (mW)			$5.1 \pm 35.3\%$	$6.5 \pm 23.6\%$	$31.6\pm23.7\%$	$63.4 \pm 8.06\%$
Raspberry P	ρ_{tx} (mW)		6 dBm 9 dBm 12 dBm 14 dBm	$593.6 \pm 0.45\% \\ 627.7 \pm 0.37\% \\ 687.8 \pm 0.70\% \\ 692.2 \pm 1.85\%$	$583.3 \pm 0.82\% \\ 611.8 \pm 1.37\% \\ 674.6 \pm 0.84\% \\ 716.3 \pm 0.85\%$	$565.2 \pm 0.96\% \\ 587.5 \pm 1.28\% \\ 653.82 \pm 1.17\% \\ 748.8 \pm 3.7\%$	$\begin{array}{c} 599.7 \pm 0.84\% \\ 621.9 \pm 2.14\% \\ 693.7 \pm 1.56\% \\ 806.6 \pm 2.56\% \end{array}$
1	$\alpha \in (\mathbf{mW})$		$\pm 2220.3 \pm 0.13\%$	γ_{ra} (mJ)	$0.126 \pm 3.34\%$	$ \gamma_{mn}$ (mJ)	$0.049 \pm 5.78\%$

* The Guard Interval is set to 800 ns, and the channel bandwidth to 20 MHz; MCS 0 and 6 use a single spatial stream, while MCS 8 and 10 use two.

error [33]. Some observations revealed by our parametrization are the following.

- 1) Most cards show moderate differences of ρ_{tx} for different MCS.
- 2) In most cases, receiving, ρ_{rx} , consumes substantially less power than transmitting, ρ_{tx} . However, in some cases, receiving is almost as costly as transmitting or even more.¹¹
- While the differences on the energy consumed by different devices are relatively small in transmission mode, they are much higher in reception mode.
- The CPU frequency has little impact on the power consumed by the NIC-specific parts.
- 5) The cross-factor for sending and receiving is of the same order of magnitude for most devices.
- 6) For 802.11n, the values of ρ_{tx} are similar for MCS 0 and 8 even though they use a different number of spatial streams, which may indicate that the driver/HW does not turn off the RF chains that are unused.

Model Validation With Multiple Stations: To validate our model in a general scenario with multiple sending and receiving stations, we consider a WLAN with one AP and three stations. Each station generates unicast traffic to the AP at a rate G, while the AP sends unicast traffic at the same rate G to each station. To apply the model of (9), we need to obtain the parameters τ_{tx} , τ_{rx} , λ_r , and λ_g . These can be obtained from typical statistics recorded by the wireless driver, namely, number of generated frames (N_g) , successful frames (N_{tx}) , transmissions attempts $(N_{\rm at})$, and received frames $(N_{\rm rx})$. With these, if the experiment is run for a duration of $T, \lambda_{\rm g}$ and $\lambda_{\rm r}$ are computed as $\lambda_{\rm g} = N_{\rm g}/T$ and $\lambda_{\rm r} = N_{\rm rx}/T$. To compute $\tau_{\rm tx}$, we account for all transmission attempts of the device plus the time spent sending the ACKs, i.e., $\tau_{tx} = (N_{at}T_L + N_{rx}T_{Ack})/T$. Similarly, to compute τ_{rx} , we need to take into account the frames and the ACKs received, $\tau_{\rm rx} = (N_{\rm rx}T_L + N_{\rm tx}T_{\rm ack})/T$.

We compare the energy consumption given by our model, fed with the above driver statistics, against the energy consumption measurements. Fig. 10 depicts these results for various combinations of L and MCS, sweeping along different traffic generation rates G in the x-axis. We conclude from the figure that the proposed model is able to accurately predict the power consumption in a general scenario.

Model Extension for TCP: The experiments performed so far, as well as the proposed model, have been restricted to UDP. Arguably, we expect quantitative differences in the cross-factor emerging with TCP traffic due to the increased stack processing complexity. To assess these differences, we separately study TCP segments and TCP ACKs. For measuring the power consumed by a TCP data packet, we have modified the TCP stack as follows: 1) at the receiver side, we do not send TCP ACKs; and 2) at the sender side, we have deactivated the TCP timers as well as the checks of congestion and receive window. With these modifications, the energy consumed by the sender is caused by the processing of TCP segments only, which are sent at the rate given by the application layer (since congestion control is not activated). By performing a similar experiment to the one of Fig. 1 with this modified TCP stack, we have measured the cross-factor associated to the processing of TCP segments,



Fig. 10. Model validation with multiple stations.

which is of $\gamma_{xg_tcp} = 1.38 \text{ mJ}$ (all other energy components remain the same), indeed almost 50% larger than that previously measured for UDP.

To further evaluate the energy consumed by processing the TCP ACKs, we have repeated the same experiment as above with the unmodified TCP stack, with results in a cross-factor of 2.1 mJ. Since the main difference between the modified TCP stack and the unmodified one is the processing of TCP acknowl-edgments, by subtracting from this cross-factor the one measured above, we obtain the energy toll for receiving TCP ACKs, i.e., $\gamma_{xr_tcpack} = 2.1 - 1.38 = 0.72$ mJ. Note that this result is in line with the one above since, as compared to TCP segments, TCP ACKs require less processing and do not involve the application layer (which, as shown in Fig. 4, accounts for a significant portion of the cross-factor).

Following the above findings, we can extend model (9) as follows in order to account for all traffic types:

$$\mathbf{P} = \rho_{\rm id} + \rho_{\rm tx} \tau_{\rm tx} + \rho_{\rm rx} \tau_{\rm rx} + \sum_{i \in \mathcal{K}} \gamma_i \lambda_i \tag{10}$$

where γ_i and λ_i are the cross-factor and sending rate, respectively, and \mathcal{K} is the set of different frame types, which includes transmitted and received UDP packets as well as TCP data packets and TCP ACKs, i.e., $\mathcal{K} = \{xg, xr, xg_tcp, xr_tcp, xg_tcpack, xr_tcpack\}$.

In order to validate the extended model, we have measured the energy consumed by a TCP session between two stations for different configurations of the MCS, packet size, and sending rate. The results on the energy consumed by the sender, depicted in Fig. 11, show that the energy consumption predicted by the model closely matches our measurements, which confirms the accuracy of the proposed extension.

VI. IMPLICATIONS ON DESIGN

The new energy consumption insights gathered in this paper may have significant implications on the design of energy-efficient mechanisms. On the one hand, existing schemes may need to be revisited so as to properly account for the impact of the cross-factor component. Indeed, according to traditional power consumption models (i.e., only baseline component plus a toll proportional to the airtime), mechanisms yielding shorter airtimes would *surely* bring about energy gains. With the cross-factor, this *might not be anymore the case*, when

¹¹A possible explanation for the behavior observed is the consumption from the adaptive amplifier at the receiver, which places the power of the receiving signal within two thresholds where the AD converter works well.



Fig. 11. Model validation for TCP.

the power savings attained at the radio interface are paid with an increased frame handling and its associated (*nonmarginal*) power consumption. On the other hand, the gained knowledge that a frame crossing the stack brings about a *fixed* penalty unrelated to the frame size may be exploited to devise techniques to *avoid* or *reduce* such energy toll.

In the following, with no pretense of completeness, we present *quantitative* examples that show how our new insights may affect existing energy efficient mechanisms as well as inspire novel approaches. As in Section IV, we first conduct experiments with the Soekris/Linux baseline platform, and then extend them to all other devices under study. Some additional examples are presented in [1].

A. Reconsidering Existing Schemes

Packet Relaying: Packet relaying in WLANs is commonly used to improve performance [35] and energy efficiency [7]. The rationale is that the use of a relay permits shorter transmission times, which compensate the impact of the extra number of hops, thus introducing a net gain. However, classical energy-efficiency analyses do not balance the airtime energy saving with the energy drain introduced by the additional frame processing, a penalty that may *fundamentally affect the relevant conclusions*.

To quantitatively support this claim, we deployed a two-hop scenario comprising three nodes (sender, relay, and receiver) and compared the power consumption in two different configurations (taken from [35]): 1) traffic directly sent to the receiver (1-hop, at 6 Mb/s), and 2) relay node used (2-hops, both at 48 Mb/s). Traffic is generated at a rate of $\lambda_g = 400$ fps with txpower = 15 dBm and different frame sizes L. Packet forwarding in the relay is performed at the routing layer. In both configurations, the relay node is always active (note that in most of the analyses on energy efficiency of relaying, the relay does not use the "sleep mode" [7], [36]).

Three types of results are shown in Fig. 12(a): 1) experimental measurements; 2) theoretical predictions using a traditional model that neglects the impact of crossing the protocol stack ("old"); and 3) predictions using the model presented in this paper ("new").¹² Not (anymore) surprisingly, results for the two models are *qualitatively* different. According to the traditional model, packet relaying always provides a gain since the energy consumption of the 2-hops case is always smaller than that of the 1-hop case. In contrast, according to the actual measurements and our model, we only gain from using the relay when packets are sufficiently long (i.e., when the airtime cost becomes dominant over the cross-factor penalty).

Data Compression in Multihop Networks: In wireless multihop networks, data compression has been proposed to reduce the information relayed; with such techniques, an intermediary node receives several frames, compresses them into a single frame, and sends it to the next hop [37], [38].

According to traditional energy models, these approaches *surely* save energy, whereas our new energy consumption insights suggest that this may not be always true. To analyze this, we used a three-node testbed consisting of a source, a sink and a relay, all using MCS = 48 Mb/s and txpower = 15 dBm. The source node generates 500-B packets at 1200 fps and sends them to the relay. The relay runs an application that receives these packets and *emulates* compression by forwarding one frame for every m frames received. Thus, our experiments do not capture the processing toll of the compression, and hence results reflect the best possible case for the performance of this scheme.

Fig. 12(b) shows total power consumption results (experimental ones, as well as predictions from old and new energy model), for different values of the compression ratio m, when data are compressed (and forwarded) at the application layer. These results are compared against the case where data are not compressed at the relay node but simply forwarded toward the sink at the routing layer.

As anticipated, the old model (top curve) predicts that compression is always advantageous. However, experimental results, matched by the new model predictions (bottom curve), show that data compression does not provide any gain in terms of energy consumption, not even for compression rates as high as 10. The reason is that the energy savings resulting from the data compression are outweighed by the extra cost of handling the packets at the application layer (cross-factor of 0.93 mJ for sending and 0.93 mJ for receiving) instead of the routing layer (cross-factor of 0.8 mJ for forwarding). This example thus shows that mechanisms devised on the basis of traditional energy models may not only fail to provide the expected gains, but may even *worsen* the actual energy consumption.

B. Novel Ways to Tackle Energy Efficiency

Packet Batching: As emerged in our work, energy consumption across the protocol stack relates to the handling of frame *units* and is practically independent of the frame size. This suggests a straightforward energy-saving strategy: *Batch packets* into bundles at the highest suitable layer for a considered scenario, deliver the bundle across the stack, thus paying the energy price associated to a *single* unit, and then restore the original frames as late as possible down the stack. Unlike previous aggregation schemes for wireless networks, this mechanism: 1) does not change the packets that are actually sent, but only modifies the way they are handled *within* the device [39]; and 2) does not save energy by reducing the cumulative *tail energy*

¹²For the model, we account for a cross-factor of 0.8 mJ to forward a packet at the relay, which has been obtained by measuring the energy consumed by forwarding a packet at the routing layer.



Fig. 12. Performance of mechanisms to reduce energy consumption. (a) Power consumption with and without relay as a function of the frame size. (b) Power consumed as a function of the compression factor. (c) Energy consumption for different "aggregation factor" values.

 TABLE V

 GAINS OF EXISTING APPROACHES WITH THE "NEW" AND "OLD" MODELS FOR ALL DEVICES

scheme	Soekris Linux		Soekris OpenBSD		Linksys Soekris 802.11n		s 802.11n	HTC Legend		Samsung Galaxy Note 10.1		Raspberry Pi		
seneme	old	new	old	new	old	new	old	new	old	new	old	new	old	new
relay	28.4%	-34.3%	28.4%	-36.2%	62.6%	-17.5%	28.7%	-35.78%	31.1%	-22.2%	71.7%	13.8%	71.8%	13.2%
compress	47.2%	-20.2%	47.2%	-21.7%	54.0%	-14.8%	66.8%	-1.1%	54.7%	-5.4%	69.9%	-8.4 %	74.1%	2.1%

consumed as a result of lingering in high power states after completing a transmission [40], [41].¹³

We quantified the attainable energy savings by implementing a scheme that consists of: 1) an "aggregator" at the application layer, which waits for n packets to generate a bundle and pass it to the TCP/IP stack; and 2) a "de-aggregator" at the wireless driver, which splits the bundle back into the original frames. Experimental measurements are reported for 100-B packets, bundled up to an "aggregation factor" n = 10, and for various (application layer) frame generation rates λ_g . Frames are transmitted over the wireless channel at MCS = 48 Mb/s and txpower = 15 dBm.

Results, shown in Fig. 12(c), have a twofold implication. First, they provide further evidence that the cross-factor toll is practically independent of the frame size: The model matches well the measurements, and the use of an *n*-bundle reduces the energy toll above the driver by *n*. Second, energy savings are notable: With 1000 fps, an aggregation factor of 10 yields a saving of almost 0.8 W, and even the aggregation of just two packets may yield considerable savings (e.g., from 4.5 to 4.15 W). Obviously, casting the above described scheme into target applications (or even more general frameworks) is not straightforward¹⁴ and is out of the scope of this paper. Nevertheless, the above results suggest that such effort may be rewarded with notable energy saving.

Driver-Level Packet Generation: As a significant part of the energy consumption comes from the crossing of the protocol stack, another approach to save energy could be to generate packets as low in the protocol stack as possible. For instance, some applications periodically generate dummy packets that

carry no information from the application layer. Hence, such packets could be easily generated at some lower layer in the protocol stack, thereby saving the power corresponding to crossing the stack.

In order to evaluate the above idea, we emulated the behavior of Skype, which during silent periods periodically generates packets that carry no information. Thus, one option would be that when Skype turns into the silent mode, instead of generating such packets at the application layer, it just instructs the driver to generate them at the lowest level of the protocol stack. By implementing this mechanism and measuring the resulting energy consumption for the typical case in which silence periods account for 60% of the time [43], we observed that this technique saves as much as 44% of the per-packet energy consumed by Skype, which confirms the effectiveness of the proposed technique.

C. Other Devices

In Sections VI-A and VI-B, we have shown that, for the Soekris device, many existing approaches may not deliver the expected gains, while substantial savings may be achieved with novel approaches that reduce the energy consumed by packet processing in the protocol stack. To gain insight into the generality of such results, we repeated the experiments of the above sections for all the devices under study.

To evaluate the performance of the existing and novel approaches of Sections VI-A and VI-B, we consider the following scenarios: 1) for all approaches, packets have a 100-B payload, which is the case, e.g., of the Skype application using the preferred codec [44]; 2) for the "data compression" experiment, we use a compression factor of m = 10; and 3) for the "packet batching," we set the aggregation factor n equal to six packets.

Table V shows (for all the devices) the energy gains delivered by the three approaches of Section VI-A, where the gain is the percentage of per-frame energy saved by the corresponding approach, i.e., $(P_{std} - P_{app})/P_{std}$, where P_{app} is the per-frame

¹³Note that the TCP Segment Offload (TSO) technique [42], which has been proposed in the context of Ethernet to achieve throughput gains, implements a similar idea to the packet batching proposed here for TCP.

¹⁴Further technical problems must be dealt with, including the interaction with the TCP/IP protocol stack (e.g., if the target application requires data to be delivered as independent TCP/IP packets) and the application's requirements (e.g., the target application scenario must tolerate the extra batching delay introduced).

TABLE VI GAINS OF NOVEL APPROACHES FOR ALL DEVICES

scheme	Soekris Linux	Soekris OpenBSD	Linksys	Soekris 802.11n	HTC Legend	Samsung Galaxy Note 10.1	Raspberry Pi
batching	79.36%	80.38%	77.85%	79.64%	71.85%	67.9%	72.24%
driver gen	44.69%	44.95%	44.29%	45.01%	42.8%	42.4%	41.71%

power consumption with the approach and $P_{\rm std}$ is the consumption with the standard stack. In particular, the table provides: 1) the gains predicted by the classical energy model in terms of per-frame energy cost ("old"); and 2) the gains measured from our experiments, which coincide with the new energy model proposed in this paper ("new").

We observe from the results of Table V that, as with the Soekris device, the actual measurements show a very different behavior from that predicted with the classical model. With the packet relaying approach ("relay"), for some of the devices, the difference with the classical model is not so drastic as to yield a performance loss, but still the gain is negligible and/or very far from the classical model. Similar conclusions hold for the data compression approach ("compress").

Table VI further shows measurements of the performance gains obtained with the novel approaches of Section IV-B. We observe that these gains are substantial for all devices: For the packet batching approach ("batching"), gains range from 68% (for the Samsung Galaxy Note 10.1) to 80% (for the Soekris), while for the driver-level packet generation approach ("driver gen"), we also obtain very significant gains (around 45%).

From the above results, we conclude that while the new energy model derived in this paper has a more drastic impact on those devices with a larger *cross-factor*, it has very strong implications for all the devices under study since in all cases: 1) the gains of existing approaches are very different from those predicted by the classical model; and 2) new approaches lead to very substantial savings.

VII. CONCLUSION

In this paper, we have conducted a thorough measurement analysis of the power consumption of 802.11 devices that, in contrast to previous works, provides a detailed anatomy of the per-packet consumption and characterizes the total consumption of the device, and not only of its wireless interface. Our analysis is, to the best of our knowledge, the first one to reveal that: 1) a substantial fraction of energy is consumed when packets cross the protocols stack (the *cross-factor*); and 2) the energy consumed by such protocol stack operations does not depend on the frame size, as opposed to the network-related operations accounted for in most of the literature. Based on our findings, we have proposed a new energy consumption model that accurately predicts the power consumption of WLAN devices. Building on our model, we have shown that some schemes targeting energy efficiency may not provide the expected gains, and even worsen performance, when the cross-factor is taken into account. We have further shown some illustrative examples where the understanding gained with our analysis can be used to devise novel algorithms that save energy by reducing the cross-factor, either by bundling packets, skipping parts of the protocol stack, or operating at the MAC layer. The lessons learned from these experiments provide some guidelines for applications developers pursuing energy-efficient operation in WLANs.

ACKNOWLEDGMENT

The authors thank the reviewers for their valuable feedback, which has been very helpful in improving the paper.

REFERENCES

- A. Garcia-Saavedra, P. Serrano, A. Banchs, and G. Bianchi, "Energy consumption anatomy of 802.11 devices and its implication on modeling and design," in *Proc. ACM CoNEXT*, Dec. 2012, pp. 169–180.
- [2] J.-M. Tarascon, "Key challenges in future Li-battery research," *Phil. Trans. Royal Soc. A*, vol. 368, no. 1923, pp. 3227–3241, Jul. 2010.
- [3] G. Y. Li *et al.*, "Energy-efficient wireless communications: Tutorial, survey, open issues," *IEEE Wireless Comm.*, vol. 18, no. 6, pp. 28–35, Dec. 2011.
- [4] S.-L. Tsao and C.-H. Huang, "A survey of energy efficient MAC protocols for IEEE 802.11 WLAN," *Comput. Commun.*, vol. 34, no. 1, pp. 54–67, Jan. 2011.
- [5] IEEE 802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, IEEE Std. 802.11, 2007.
- [6] E. Rozner, V. Navda, R. Ramjee, and S. Rayanchu, "NAPman: Network-assisted power management for WiFi devices," in *Proc. ACM MobiSys*, Jun. 2010, pp. 91–106.
- [7] X. He and F. Li, "Throughput and energy efficiency comparison of onehop, two-hop, virtual relay and cooperative retransmission schemes," in *Proc. Eur. Wireless*, Apr. 2010, pp. 580–587.
- [8] A. P. Jardosh et al., "Green WLANs: On-demand WLAN infrastructures," Mobile Netw. Appl., vol. 14, no. 6, pp. 798–814, Dec. 2009.
- [9] D. Qiao, S. Choi, and K. Shin, "Interference analysis and transmit power control in IEEE 802.11a/h wireless LANs," *IEEE/ACM Trans. Netw.*, vol. 15, no. 5, pp. 1007–1020, Oct. 2007.
- [10] J.-C. Chen and K.-W. Cheng, "EDCA/CA: Enhancement of IEEE 802.11e EDCA by contention adaption for energy efficiency," *IEEE Trans. Wireless Commun.*, vol. 7, no. 8, pp. 2866–2870, Aug. 2008.
- [11] R. Bruno, M. Conti, and E. Gregori, "Optimization of efficiency and energy consumption in p-persistent CSMA-based wireless LANs," *IEEE Trans. Mobile Comput.*, vol. 1, no. 1, pp. 10–31, Mar. 2002.
- [12] V. Baiamonte and C.-F. Chiasserini, "Saving energy during channel contention in 802.11 WLANs," *Mobile Netw. Appl.*, vol. 11, no. 2, pp. 287–296, Apr. 2006.
- [13] M. Ergen and P. Varaiya, "Decomposition of energy consumption in IEEE 802.11," in *Proc. IEEE ICC*, Jun. 2007, pp. 403–408.
- [14] A. Garcia-Saavedra, P. Serrano, A. Banchs, and M. Hollick, "Energyefficient fair channel access for IEEE 802.11 WLANs," in *Proc. IEEE WoWMoM*, Jun. 2011, pp. 1–9.
- [15] X. Wang, J. Yin, and D. P. Agrawal, "Analysis and optimization of the energy efficiency in the 802.11 DCF," *Mobile Netw. Appl.*, vol. 11, no. 2, pp. 279–286, Apr. 2006.
- [16] E.-S. Jung and N. H. Vaidya, "An energy efficient MAC protocol for wireless LANs," in *Proc. IEEE INFOCOM*, Jun. 2002, pp. 1756–1764.
- [17] M. Carvalho, C. Margi, K. Obraczka, and J. Garcia-Luna-Aceves, "Modeling energy consumption in single-hop IEEE 802.11 ad hoc networks," in *Proc. ICCCN*, Oct. 2004, pp. 367–372.
- [18] D. Qiao, S. Choi, A. Jain, and K. G. Shin, "MiSer: An optimal low-energy transmission strategy for IEEE 802.11a/h," in *Proc. ACM MobiCom*, Sep. 2003, pp. 161–175.
- [19] S. Chiaravalloti, F. Idzikowski, and L. Budzisz, "Power consumption of WLAN network elements," Tech. Univ. Berlin, Berlin, Germany, TKN-11-002, Aug. 2011.
- [20] L. Feeney and M. Nilsson, "Investigating the energy consumption of a wireless network interface in an ad hoc networking environment," in *Proc. IEEE INFOCOM*, Apr. 2001, pp. 1548–1557.
- [21] Atheros Communications, San Jose, CA, USA, "Power consumption and energy efficiency comparisons of WLAN products," White paper, Apr. 2004.
- [22] Intel Corp., Santa Clara, CA, USA, "Data transfer over wireless LAN power consumption analysis," White paper, Feb. 2009.
- [23] E. Lochin et al., "Energy consumption models for ad-hoc mobile terminals," in Proc. Med-Hoc Net, Jun. 2003.

- [24] A. Carroll and G. Heiser, "An analysis of power consumption in a smartphone," in *Proc. USENIX ATC*, Jun. 2010, p. 21.
- [25] A. Rice and S. Hay, "Measuring mobile phone energy consumption for 802.11 wireless networking," *Pervasive Mobile Comput.*, vol. 6, no. 6, pp. 593–606, Dec. 2010.
- [26] A. Pathak, Y. C. Hu, and M. Zhang, "Bootstrapping energy debugging on smartphones: A first look at energy bugs in mobile devices," in *Proc. ACM HotNets*, Nov. 2011.
- [27] J.-P. Ebert, B. Burns, and A. Wolisz, "A trace-based approach for determining the energy consumption of a WLAN network interface," in *Proc. Eur. Wireless*, Feb. 2002, pp. 230–236.
- [28] D. Halperin, B. Greenstein, A. Sheth, and D. Wetherall, "Demystifying 802.11n power consumption," in *Proc. HotPower*, Oct. 2010.
- [29] E. Rantala, A. Karppanen, S. Granlund, and P. Sarolahti, "Modeling energy efficiency in wireless Internet communication," in *Proc. ACM MobiHeld*, Aug. 2009, pp. 67–68.
- [30] H. Wu, S. Nabar, and R. Poovendran, "An energy framework for the network simulator 3 (NS-3)," in *Proc. SIMUTools*, Mar. 2011, pp. 222–230.
- [31] A. W. Min, R. Wang, J. Tsai, M. A. Ergin, and T.-Y. C. Tai, "Improving energy efficiency for mobile platforms by exploiting low-power sleep states," in *Proc. Comput. Frontiers*, May 2012, pp. 133–142.
- [32] S. Keranidis *et al.*, "Energy consumption monitoring of wireless testbed infrastructure through the NITOS EMF framework," in *Proc. ACM WiNTECH*, Sep. 2013, pp. 73–80.
- [33] J. S. Milton and J. C. Arnold, *Introduction to probability and statistics*, 4th ed. New York, NY, USA: McGraw-Hill, 2003.
- [34] J. L. Lee et al., "An experimental study on the capture effect in 802.11a networks," in Proc. ACM WiNTECH, Sep. 2007, pp. 19–26.
- [35] P. Bahl et al., "Opportunistic use of client repeaters to improve performance of WLANs," *IEEE/ACM Trans. Netw.*, vol. 17, no. 4, pp. 1160–1171, Aug. 2009.
- [36] C. Sun and C. Yang, "Is two-way relay more energy efficient?," in *Proc. IEEE GLOBECOM*, Dec. 2011, pp. 1–6.
- [37] A. B. Sharma, L. Golubchik, R. Govindan, and M. J. Neely, "Dynamic data compression in multi-hop wireless networks," in *Proc. ACM SIG-METRICS*, Jun. 2009, pp. 145–156.
- [38] S. J. Baek, G. de Veciana, and X. Su, "Minimizing energy consumption in large-scale sensor networks through distributed data compression and hierarchical aggregation," *IEEE J. Sel. Areas Commun.*, vol. 22, no. 6, pp. 1130–1140, Aug. 2004.
- [39] A. Jain and M. Gruteser, "Benefits of packet aggregation in ad-hoc wireless network," Univ. Colorado, Boulder, CO, USA, Aug. 2003.
- [40] F. Qian et al., "Characterizing radio resource allocation for 3G networks," in Proc. ACM Internet Meas. Conf., Nov. 2010, pp. 137–150.
- [41] N. Balasubramanian, A. Balasubramanian, and A. Venkataramani, "Energy consumption in mobile phones: A measurement study and implications for network applications," in *Proc. ACM Internet Meas. Conf.*, Nov. 2009, pp. 280–293.
- [42] A. Menon and A. Zwaenepoel, "Optimizing TCP receive performance," in *Proc. USENIX ATC*, Jun. 2008, pp. 85–98.
- [43] Artificial Conversational Speech, ITU-T Recommendation P.59, 1993.
- [44] D. Bonfiglio, M. Mellia, M. Meo, N. Ritacca, and D. Rossi, "Tracking down Skype traffic," in *Proc. IEEE INFOCOM*, Apr. 2008, pp. 843–851.



Pablo Serrano (M'09) received the Telecommunication Engineering and Ph.D. degrees from the University Carlos III of Madrid (UC3M), Leganes, Spain, in 2002 and 2006, respectively.

He has been with the Telematics Department of UC3M since 2002, where he currently holds the position of Associate Professor. He was a Visiting Researcher with the Computer Network Research Group, University of Massachusetts, Amherst, MA, USA, in 2007, and with the Telefonica Research Center, Barcelona, Spain, in 2013. He has over 50

scientific papers in peer-reviewed international journal and conferences.

Dr. Serrano serves on the Editorial Board of the IEEE COMMUNICATIONS LETTERS, has been a Guest Editor for *Computer Networks*, and has served on the TPC of a number of conferences and workshops including IEEE INFOCOM, IEEE WoWMoM, and IEEE GLOBECOM.



Andres Garcia-Saavedra received the B.Sc. degree in telecommunication engineering from the University of Cantabria (UNICAN), Santander, Spain, in 2009, and the M.Sc. and Ph.D. degrees in telematics engineering from the University Carlos III of Madrid (UC3M), Leganes, Spain, in 2010 and 2013, respectively.

Currently, he works as a Research Fellow with the Hamilton Institute, National University of Ireland, Maynooth (NUIM), Maynooth, Ireland. His research interests lie in the application of fundamental math-

ematics to real-life computer communications systems—in particular, resource allocation problems, performance evaluation, and prototyping of wireless network protocols.



Giuseppe Bianchi received the Laurea degree in electrical engineering from Politecnico di Milano, Milan, Italy, in 1990, and the specialist degree in information technology from the CEFRIEL Research Center of Milan, Milan, Italy, in 1991.

He has been a Full Professor of networking with the University of Roma Tor Vergata, Rome, Italy, since 2007. His research activity includes wireless LANs, IP networking, privacy and security, and traffic monitoring and is documented in about 180 peer-reviewed international journal and conference



Prof. Bianchi has served as an Editor for the IEEE/ACM TRANSACTIONS ON NETWORKING, IEEE TRANSACTIONS ON WIRELESS COMMUNICATION (Area Editor), and *Computer Communication*. He has chaired several international networking conferences, including IEEE INFOCOM 2014, ACM SRIF 2013, ACM WINTECH 2011, IEEE WoWMoM 2010, etc.



Albert Banchs (M'04–SM'12) received the M.Sc. in telecommunications engineering and Ph.D. degree in telematics engineering from the UPC BarcelonaTech, Barcelona, Spain, in 1997 and 2002, respectively.

He was with the International Computer Science Institute (ICSI), Berkeley, CA, USA, in 1997; Telefonica I+D, Barcelona, Spain, in 1998; and NEC Europe, Ltd., Heidelberg, Germany, from 1998 to 2003. Since 2003, he has been with the University Carlos III of Madrid, Leganes, Spain. Since 2009, he has a double affiliation as Deputy Director of IMDEA Net-

works, Madrid, Spain.

Dr. Banchs has served as an Editor for *Computer Communications* and the IEEE COMMUNICATIONS LETTERS, and as TPC Chair of EW 2010, IEEE WoWMoM 2012, ACM WINTECH 2013, IEEE Online GreenComm 2013.



Arturo Azcorra (SM'02) received the M.Sc. and Ph.D. degrees in telecommunications engineering from the Universidad Politécnica de Madrid, Madrid, Spain, in 1986 and 1989, respectively. He also received the M.B.A. from the Instituto de Empresa, Madrid, Spain, in 1993.

He has a double appointment as Full Professor (with chair) with the Telematics Engineering Department, University Carlos III of Madrid, Leganes, Spain, and as Director of IMDEA Networks, Madrid, Spain. He has coordinated and served as a Program

Committee Member in numerous international conferences and has published over 100 scientific papers in books, international journals, and conferences.