Aalborg Universitet

IEEE 802.11 Networks: A Simple Model Geared Towards Offloading Studies and Considerations on Future Small Cells

Garcia, Luis Guilherme Uzeda; Rodriguez, Ignacio; Catania, Davide; Mogensen, Preben

Link to publication from Aalborg University

IEEE 802.11 Networks: A Simple Model Geared Towards Offloading Studies and Considerations on Future Small Cells

Luis G. Uzeda Garcia[1], Ignacio Rodriguez[1], Davide Catania[1] and Preben Mogensen[1,2]

[1]Department of Electronic Systems, Aalborg University, Denmark, {lug, irl, dac, pm}@es.aau.dk
[2]Nokia Siemens Networks, Research Center Aalborg, Denmark, preben.mogensen@nsn.com

Abstract—**WiFi is the prevalent wireless access technology in local area deployments and is expected to play a major role in a mobile operator's data offloading strategy. As a result, having simple tools that are able to assess the offloading potential of IEEE 802.11 networks is vital. In this paper, we propose a simple closed-form solution to calculate down- and uplink throughput values per user under full-buffer traffic when small WiFi cells are used to offload macrocells. Extensive measurement campaigns and simulation results demonstrate that there is an excellent quantitative match between analytical model and data despite the simplicity of the former. Finally, in light of our observations we discuss some of the fundamental technological limitations that may have a significant impact on the future of small cells.**

Index Terms—**Small cells, WiFi, IEEE 802.11, Offloading.**

I. Introduction

Wireless Local Area Network (WLAN) products based on the IEEE 802.11 family of standards [1–3] have proven to be a tremendous commercial success. WiFi, as the technology became popularly known, is now ubiquitous and its fifth generation, 802.11ac, is expected to be ratified by late 2013.

Not surprisingly, the available literature on IEEE 802.11-based networks is monumental. Virtually, every single aspect related to WiFi has been extensively addressed. Arguably, the work by Bianchi [4] is the foremost contribution in this area. His paper analyzes the primary media access control (MAC) technique of 802.11, namely the distributed coordination function (DCF). His efforts yielded a remarkably accurate model based on a two-dimensional Markov chain to compute the 802.11 DCF throughput under saturation. The model became known as Bianchi's model and has been extended countless times to factor in certain missing features of the DCF and its successor, the Enhanced Distributed Channel Access (EDCA). The main drawback of this and similar models is that they do not possess a closed-form solution and hence need to be solved numerically.

In this short paper, we lay out a very simple model that tries to capture the essence of the IEEE 802.11 MAC. Contrary to the trend, we do not strive for a model that generalizes Bianchi's in some sense or one that outperforms it in terms of accuracy. In fact, our model is aimed at cellular system engineers to whom WiFi was until very recently the elephant in the room. In modern cellular systems, duplexing, multiplexing and interference coordination are typically thought of as three clearly distinct tasks or functionalities. In contrast, in a basic WiFi system the DCF (or the EDCA) plays all three roles simultaneously. This radical paradigm shift may result in some serious misconceptions, which, in turn, could lead to erroneous conclusions regarding the offloading potential of different technologies.

Our main goal is to provide a lightweight closed-form model that would allow a straightforward computation of the down- and uplink throughput values per user under full-buffer when WiFi networks are used to offload macrocells. Such model can be applied in system level simulation studies consisting of complex deployments with multiple layers of both WiFi and cellular, such as the ones found in [5; 6]. This line of research is becoming increasingly more important to network operators due to the surge of mobile data.

Strictly speaking, the model is valid for User Datagram Protocol (UDP) traffic, however we present one empirical variant that covers Transmission Control Protocol (TCP) traffic as well. It is also shown that the unmodified UDP version serves as an upper bound for TCP traffic. Throughout the paper, the interested reader will find references for further reading.

The rest of the paper is organized as follows: Section II is devoted to an explanation of the proposed model. The basic assumptions and its limitations are discussed as well. Section III is dedicated to the empirical validation of the model. The investigation encompasses a wide range of simulation and measurement scenarios. Section IV analyzes the numerical results and attempts to put them into perspective. The considerations run from important improvements covered by subsequent amendments to the standard (802.11e, 802.11n and 802.11ac) to fundamental limitations that have impact on regulatory issues and the future of small cells. Finally, Section V concludes the paper.

II. A Simple Throughput Model

A. Insight

The main idea behind the model is the macroscopic behavior of the Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) protocol. CSMA/CA is the linchpin of both the

DCF and EDCA[1] found in WiFi and was designed to achieve equal channel access probabilities for all nodes. This holds in the long term and as long as there are no hidden nodes, i.e. all devices within radio range of each other.

We can then assume that K devices, either user Stations (STAs) or Access Points (APs), will alternate transmissions and the total system bandwidth W will be employed by one device at a time. This can be characterized as a dynamic time division of the resources, whereby the k-th device ideally acquires a fraction τ_k of the total capacity:

$$C_k = \tau_k W \log_2 (1 + SINR) \qquad (1)$$

In (1), the $SINR$ term denotes the signal-to-interference-plus-noise ratio, while τ_k corresponds to a fraction of the total transmission time. In an idealized scenario with no overheads, in the absence of collisions, equal data rates and frame sizes, time shares will be identical; consequently the throughput per device is:

$$C_k = \left(\frac{1}{K}\right) W \log_2 (1 + SINR) \qquad (2)$$

Clearly, the situation is not that simple. First, one cannot expect that data rates will be the same for all devices. Second, the achievable throughput above the MAC layer is significantly lower than the nominal physical layer data rates due to collisions, random back-off timers and fixed overheads as exposed by the seminal paper [11]. Next we shall refine the model to account for the aforementioned aspects.

B. The Model

Ignoring fixed overheads for the moment, the fraction of time each station occupies the medium equals $\tau_k = F_k/PHY_k$, where F_k is the transmitted frame size and $PHY_k \equiv W \log_2 (1 + SINR_k)$. However if all frames have the same size (e.g. the maximum size specified by the 802.11 standard due to the full-buffer assumption), it is possible to model the system without incorporating τ_k explicitly. Instead, relying on the equal channel access probabilities for all nodes, one can posit that an effective long term throughput over the medium is shared equitably, thus $\tau_k = \tau \, \forall \, k \in K$. It is easy to demonstrate that such effective throughput (PHY_{EFF}) can be calculated as the harmonic mean of the data rates of all l competing links (data flows); hence:

$$PHY_{EFF} = \left(\frac{1}{L}\sum_{l=1}^{L} PHY_l^{-1}\right)^{-1} \qquad (3)$$

Fig. 1 depicts a case where $L = 6$ and (3) follows from the fact that payloads are equally sized but might take different amounts of air-time, being at the root of the well-known performance *anomaly*[2] discussed in [11].

---

[1]There is a vast amount of literature describing the inner workings of CSMA/CA, DCF and EDCA [4; 7–10]. For this paper the emergent behavior is much more relevant.

[2]The anomaly implies that a single slow (low data rate) device may limit the throughput of all fast devices.

Despite the great strides made by the recent incarnations of WiFi, the throughput at the MAC layer is always smaller than the data rate [11; 12]. In order to account for that, we propose the usage of a multiplicative constant ($0 < \eta < 1$), a fudge factor, to adjust the effective throughput that is shared among all devices within the same contention domain, i.e. all co-channel devices within sensing range of each other.

This so-called fudge factor lumps together all factors that contribute to the MAC inefficiency, such as: physical layer preamble, headers, inter-frame spacings, random back-off timers, acknowledgments, etc. Unfortunately, it is difficult to derive the value of η analytically. In this paper, we employ a simple linear least squares technique to fit it to the data available. In Section IV, we shall discuss further the nature of this constant.

Finally, the average uplink (UL) throughput per station made available by the MAC to the upper layers under full-buffer conditions is modeled by the simple expression:

$$C_k^{UL} = \frac{\eta \cdot PHY_{EFF}}{K} \qquad (4)$$

Meanwhile the corresponding downlink (DL) prediction is:

$$C_k^{DL} = \frac{\eta \cdot PHY_{EFF}}{K S_k^n} \qquad (5)$$

where S_k^n denotes the number of stations served by the $n-$th AP (cell) serving the $k-$th station. The denominators of both equations deal with the distribution of the total throughput among all devices within the same contention domain, i.e. co-channel devices within sensing range of each other. Both equations are illustrated by Fig. 1, which displays the corresponding effective shares of the system resources per data flow. It also includes two extra devices sharing the same channel: Access Point 2 (*AP2*) and its served station, *STA3*.

Without loss of generality, if N represents the total number of contending APs and $S_k^n = S \, \forall \, n \in N$, then $K = NS + N$ which allows (5) to be rewritten as:

$$C_k^{DL} = \frac{\eta \cdot PHY_{EFF}}{N(S^2 + S)} \qquad (6)$$

Equations (6) highlights the fact that the DL throughput per station has an *inverse-square dependency* on the number of stations associated with an AP, whereas the UL throughput is simply inversely proportional to the number of devices. This arises because the AP is merely another device competing for the medium and its scant opportunities are further subdivided to serve one receiving station at a time.

The important question now is whether such a simplistic model relying on some gross simplifications is indeed able to capture the essence of CSMA/CA. As we shall see, the numerical and measurement results of the validation experiments described in Section III clearly state that the answer is yes when the assumptions presented next are held.
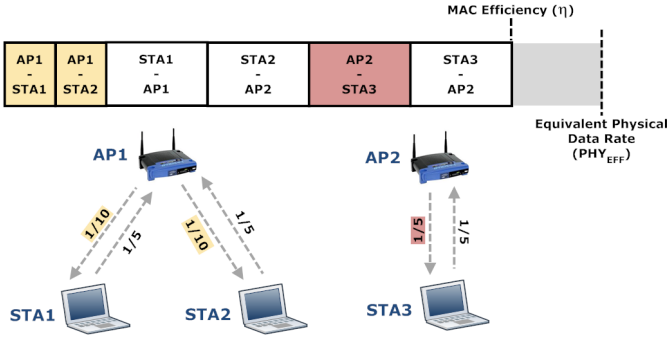
Fig. 1: An illustration of the resource partitioning predicted by the model. Two neighboring co-channel cells displayed.

Parameter	Value
Path Loss Model	Winner II Indoor
Transmit Power	20 dBm
Spectrum Allocation	Channel 13 on 2.4GHz ISM Band
AP-STA Configuration 1	1 AP, with 1-8 STA's
AP-STA Configuration 2	2 AP, with 1-3 STA's each
AP-STA Configuration 3	3 AP, with 1-2 STA's each
AP/STA Deployment	Random Position per simulation Drop
Carrier Sensing Threshold	-76dBm
PHY Data Rate	54Mbit/s
Packet Size	3x2304 bytes (Frame Burst)
Traffic Model	Full Buffer
Simulation Time	10s
Average Number of Simulation Samples	$\sim$500 per simulation

TABLE I: System Level Simulator Parameters

C. Key Assumptions and Limitations

As described previously, the model is expected to reflect the behavior of the WiFi MAC when senders *in both directions* always have data to transmit, i.e. DL and UL full-buffer traffic. This performance metric is widely used in the evaluation of cellular systems and especially relevant to offloading studies because it corresponds to the maximum amount of traffic that can be offered to a network. Moreover, since we are particularly interested in bulk data transfers, all flows are assumed to belong to the same access category (AC), thus essentially nullifying the prioritized QoS mechanism offered by the EDCA.

Our working assumption is that cells are fairly small and planned for very high (maximum) throughput indoor coverage akin to femtocells. In sum, the scenario could be seen as one where either (i) co-channel cells are isolated, (ii) or if two access points in neighboring residences choose the same channel then there is full connectivity; i.e. there are either no hidden nodes. Therefore the model is clearly not applicable to ad-hoc or multi-hop topologies.

Additionally, the model does not intend to cover cases where either the Point Coordination Function (PCF) or the Hybrid Coordination Function (HCF) Controlled Channel Access (HCCA) are employed. Neither of these optional features has been widely implemented, and would lead to a pseudo-framed scheduled access if employed [13].

III. EMPIRICAL VALIDATION

A comprehensive set of experiments has been carried out in order to validate the model. It consisted of simulation campaigns and over 100 hours of real-world measurements. Three basic scenarios were considered:

- Single cell: a single AP with 1-8 STA's connected.
- Two co-channel cells: 2 AP's with 1-3 STA's connected to each.
- Three co-channel cells: 3 AP's with 1-2 STA's connected to each.

Neglecting obtuse channel selections, the last two cases can easily occur in very dense uncoordinated deployments where the number of neighboring WiFi cells is much larger than the number of orthogonal channels available, e.g. on the 2.4 GHz

band. Even though multiple cells sharing the same contention domain scenario is a much less probable event on the 5 GHz band due to the large number of non-overlapping 20 MHz channels, the same cannot be stated when talking about future 160 MHz and even 80 MHz wide channels.

A. System Level Simulations

In order to cross-validate the measurements and the analytical formula provided, a proprietary system level simulator was used. The system level simulator makes use of a detailed, nearly standard-compliant WiFi MAC 802.11g implementation, including MAC layer headers and acknowledgments as well as physical layer preambles defined in the standard. Link adaptation is not implemented and all frames are assumed to be transmitted and correctly received at 54 Mbit/s if a target SNR is reached, otherwise the frame is dropped and no acknowledgment is sent to the transmitting side. The scenario being considered is small enough to ensure that all nodes can listen to each other and packets will only be dropped if they collide, given that their back-off timers expired at the same time. The chosen simulation scenario therefore ensures that hidden nodes are not present. A summary of the key parameters used is found in Table I.

B. Measurement Campaign

The measurement campaign was performed on channel 13 of the 2.4 GHz ISM band in a highly isolated location (basement) in order to avoid/minimize the influence of the existing university's WiFi network deployed on channels 1, 6 and 11. Nine ordinary and identical 802.11g Linksys WRT54GL v1.1 routers with firmware Tomato 1.28 [14] were used as wireless interfaces. In all the cases considered, frame bursting was enabled[3] and the distance between serving AP and STAs was 1m in order to guarantee that data rates of 54 Mbit/s were attainable. In the multi-cell cases, all the nodes were placed in the same 10mx6m room and were able to listen to each others transmissions, thus ensuring full connectivity.

---

[3]Frame bursting allows one station to transmit 3 frames before relinquishing the channel. Disabling this feature had no impact on the overall conclusions. The only observable and expected effect was a reduction in terms of MAC efficiency, consequently leading to a lower fudge factor η.
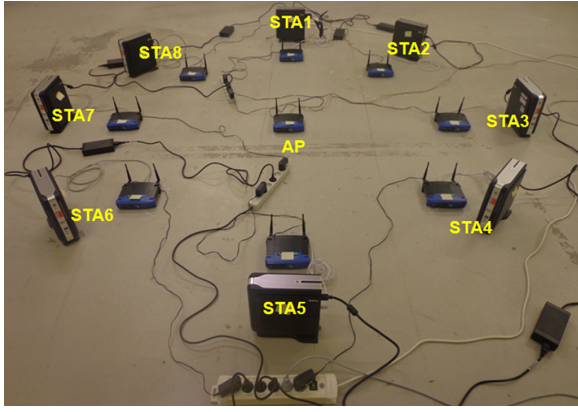
Fig. 2: A picture of the single-cell deployment scenario.

Downlink (UDP)						
# Co-channel cells →	1		2		3	
# Stations/cell	Mean	Std	Mean	Std	Mean	Std
1	18.24	0.54	9.50	0.58	6.23	0.91
2	6.84	0.58	3.70	0.65	2.50	0.50
3	3.64	0.42	1.83	0.40	-	-
4	2.30	0.48	-	-	-	-
5	1.65	0.33	-	-	-	-
6	1.20	0.24	-	-	-	-
7	0.93	0.18	-	-	-	-
8	0.76	0.14	-	-	-	-
Uplink (UDP)						
# Co-channel cells →	1		2		3	
# Stations/cell	Mean	Std	Mean	Std	Mean	Std
1	19.76	0.50	9.86	0.69	6.83	0.68
2	12.20	0.62	5.98	0.84	4.05	0.53
3	9.21	0.67	4.41	0.60	-	-
4	7.21	0.56	-	-	-	-
5	5.89	0.37	-	-	-	-
6	5.08	0.40	-	-	-	-
7	4.43	0.50	-	-	-	-
8	3.92	0.37	-	-	-	-

TABLE II: Summary of UDP Throughput Measurements (Mbit/s).

All wireless devices were connected via 100 Mbit/s Ethernet cables to their individual control terminals where UDP/TCP traffic was generated and measured by Iperf [15]. The terminals used were Dell Optiplex SX270 with Intel Pentium 4, 2.4 GHz x2 processors, 2 GB of RAM running Linux Ubuntu 12.04 LTS. Moreover, Iperf was configured such that all wireless nodes always had data ready to be transmitted in their buffers. The system was controlled remotely, and the room was empty while all the measurements were performed.

Fig. 2 depicts one of the measurement configurations. For all the different cases, the measurement procedure consisted of 2 different stages:

- Calibration: all the STA's were tested individually during 5 min to verify compliance.
- Measurement: for each subcase, 10 realizations 5 min long each were run. Average values extracted from this ensemble.

IV. NUMERICAL RESULTS AND DISCUSSIONS

A. UDP Performance

Figure 3 condenses the results from our extensive experiments, while Table II presents a more detailed summary of the data acquired. The former depicts the evolution of average UL as well as DL throughput values as a function of the number of STA served by each AP. The quantitative match between model and both simulation and measurements is surprisingly good when one considers the simplicity of the analytical model.

It is noteworthy to observe that it is always better to have a larger number of APs to serve the same total number of STAs when downlink throughput is the performance metric of interest. This comes from the fact that the subdivision of AP resources is reduced.

Moreover, it is remarkable that a fixed multiplicative factor $\eta = 0.68$ (best-fit) suffices. Conventional wisdom dictates that the contention windows dilate as the number of active devices increases in order to decrease the number of collisions. Longer waiting periods lead to a larger implied overhead and therefore a lower and variable MAC efficiency.

On the other hand, when the number of active stations is low, there is a higher likelihood that the channel will remain idle after the end of a transmission. Recall that devices cannot initiate a transmission until their back-off counters expire. With a higher number of devices, these random gaps will tend to become shorter, thereby increasing the overall efficiency. These two effects seem to counterbalance each other leading to a constant η value for the observed range.

It is also worth stressing that using an incorrect value for η does not change the predictions of the model qualitatively, but it obviously has a quantitative impact. Tabulating η for other versions of the 802.11 standard (e.g. b or n) is suggested for future work.

Fig. 4 focuses on the DL performance and tests the model in the presence of a single slow station. The data rate of one station was set at 6 Mbit/s to emulate the presence of a node near the cell edge. Once again the predictions match the empirical data. The damaging effects of the 802.11 anomaly become evident since a substantially smaller "pie" of resources gets subdivided. From those results, the importance of planning cells for high data rate coverage should become clear.

B. TCP Performance

In Fig. 5, TCP results come into the picture. A single AP scenario is considered. TCP is particularly hard to model analytically due to its self-clocked closed-loop nature. An in-depth analysis of the behavior of TCP over WiFi is beyond the scope of this contribution, but readers can find insightful discussions in [16–18]. Here, we limit ourselves to a few cautious observations.
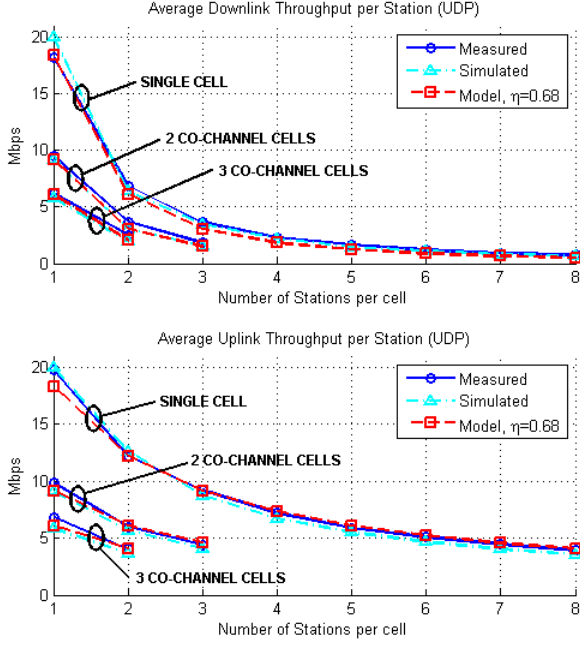
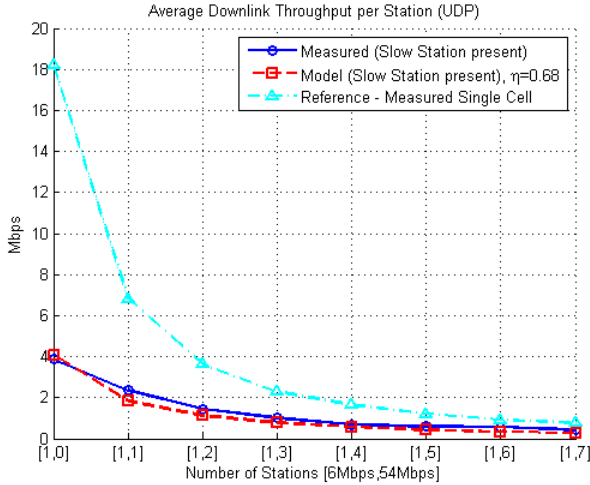Fig. 3: Comparison between measurements, simulations and throughput model.

Fig. 4: Model predictions in the presence of the well-known 802.11 anomaly.

When TCP is used, a new type of packet carrying the transport layer acknowledgements (ACKs) is introduced [19]. This overhead has a negative impact on the overall efficiency. Nonetheless, simply reducing η does not lead to match between the measured data and predictions from (5) that is as good as the one seen with UDP traffic.

However, when $N = 1$ AP and the dependency on the number of DL flows per cell is made *cubic rather than quadratic*, in other words, (6) is modified to:
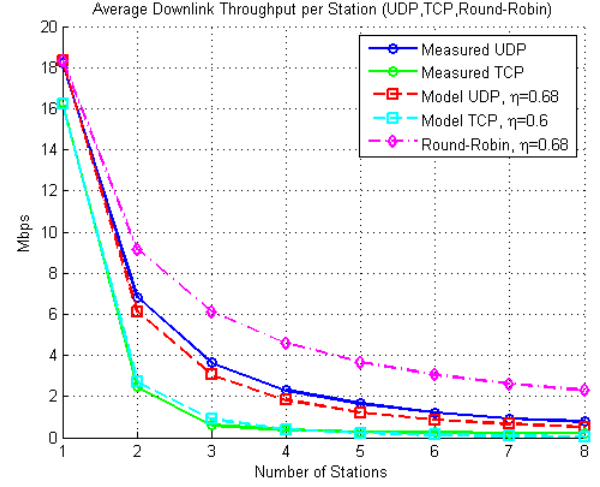

Fig. 5: Comparison between UDP, TCP and scheduled models.

$$C_k^{DL} = \frac{\eta \cdot PHY_{EFF}}{(S^3 + S^2)} \tag{7}$$

the match becomes the one observed in Fig. 5. This is an interesting empirical finding. It can be observed that the DL starvation (a *duplexing anomaly*) is exacerbated as anticipated by the seminal work in [16]. Unfortunately, it was not possible carry out additional experiments in order to increase the confidence in our empirical TCP model, because our simulation tool does not currently include an implementation of the TCP/IP protocol stack and our measurement campaign was confined to a single kind of WiFi interface, namely 802.11g Linksys WRT54GL v1.1 routers.

Nonetheless, from the findings in [16] it can be stated with certainty that the different qualitative and quantitative trends stem from the congestion control mechanism having to cope with a single buffer at the AP being shared by the STA_{perAP} DL queues. This leads to frequent packet drops due to buffer overflows. Although not depicted here, the UL the predictions from (4) remain valid. Recall that in the UL, the buffers are unique to each STA.

C. Final Remarks

By design, WiFi relies on its MAC to ensure an ideally interference-free channel. The problem lies in the fact that in terms of capacity, the scale tips favorably towards more bandwidth rather than higher signal-to-noise-plus-interference ratio [20]. That is to say, judicious reuse of resources, in spite of increased interference, is actually beneficial. WiFi is simply oblivious to that. Therefore, even if one downplays the aforementioned duplexing anomaly by considering the extreme case where UL traffic is nearly non-existent, the transmission opportunities granted to each AP in the same contention domain will become progressively scarcer as deployments of small high-throughput cells become ever denser.

So far, two solutions have come to the rescue when the capacity requirements cannot be met: (i) wider swaths of

bandwidth coupled with higher bands and (ii) proprietary and cleverly engineered centralized solutions that virtually replace the standardized MAC. Neither approach is exactly future-proof because spectrum is first and foremost a valuable commodity and incompatible proprietary solutions might cover overlapping areas in the future. In view of a looming capacity crunch and in an era of cognitive radio systems, it is valid to raise the question of whether granting more spectrum to a proven yet fundamentally limited technology is the most sensible strategy. On the bright side, significant progress has been made by subsequent amendments to remedy the anomalies cited in this paper.

For example, HCCA allows the AP to take control over the channel virtually at any time. It also introduced the concept of transmission opportunities which limits the air time granted to slow devices. In the absence of legacy devices, these two elements can effectively mitigate both anomalies discussed previously. The single-cell performance could then roughly resemble that shown by the idealized dashed magenta curve in Fig. 5, where DL flows are scheduled in a round-robin fashion using 50% of the transmission opportunities. 802.11 ac promises to take this even one step further by bringing Multi-User MIMO (MU-MIMO) to the table, allowing multiple STAs to be served simultaneously. Nevertheless, it remains to be seen whether such solutions will be sufficient to turn WiFi into the ideal offloading solution.

V. CONCLUSIONS

The main contribution of this paper was to show that a very simple analytical model is able to successfully capture the essence of the prevalent media access control technique of IEEE 802.11 networks. Extensive real-world measurement campaigns and simulation results reassert that a set of basic equations can assist cellular system engineers in predicting the offloading potential of small WiFi cells.

The model also makes one of WiFi's well known shortcomings intuitively understandable. WiFi is a lopsided system because the uplink can acquire a disproportionately large share of the resources at the expense of a starving downlink. This is in stark contrast to the typical asymmetry of traffic. Finally, the last subsection of the paper was devoted to some thought-provoking discussions addressing technology limitations that might have a significant impact on the future of small cells.

REFERENCES

[1] IEEE Standard for Information Technology - Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks - Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. *IEEE Std 802.11-2007 (Revision of IEEE Std 802.11-1999)*, pages 1 –1076, 12 2007.

[2] IEEE Standard for Information Technology - Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks - Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements. *IEEE Std 802.11e-2005 (Amendment to IEEE Std 802.11, 1999 Edition (Reaff 2003)*, pages 1 –189, 2005.

[3] IEEE Standard for Information technology– Local and metropolitan area networks– Specific requirements– Part 11: Wireless LAN Medium Access Control (MAC)and Physical Layer (PHY) Specifications Amendment 5: Enhancements for Higher Throughput. *IEEE Std 802.11n-2009 (Amendment to IEEE Std 802.11-2007 as amended by IEEE Std 802.11k-2008, IEEE Std 802.11r-2008, IEEE Std 802.11y-2008, and IEEE Std 802.11w-2009)*, pages 1 –565, 2009.

[4] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *Selected Areas in Communications, IEEE Journal on*, 18(3):535 –547, March 2000.

[5] Liang Hu, C. Coletti, Nguyen Huan, P. Mogensen, and J. Elling. How Much Can Wi-Fi Offload? A Large-Scale Dense-Urban Indoor Deployment Study. In *Vehicular Technology Conference (VTC Spring), 2012 IEEE 75th*, pages 1 –6, may 2012.

[6] Liang Hu, C. Coletti, Nguyen Huan, István Kovács, and Benny Vejlgaard. Realistic Indoor Wi-Fi and Femto Deployment Study as the Offloading Solution to LTE Macro Networks. In *Vehicular Technology Conference (VTC Fall), 2012 IEEE 76th*, pages 1 –6, September 2012.

[7] F. Cali, M. Conti, and E. Gregori. Dynamic tuning of the IEEE 802.11 protocol to achieve a theoretical throughput limit. *Networking, IEEE/ACM Transactions on*, 8(6):785 –799, dec 2000.

[8] David Malone, Ken Duffy, and Doug Leith. Modeling the 802.11 Distributed Coordination Function in Nonsaturated Heterogeneous Conditions. *Networking, IEEE/ACM Transactions on*, 15(1):159 –172, feb. 2007.

[9] S. Mangold, Sunghyun Choi, G.R. Hiertz, O. Klein, and B. Walke. Analysis of IEEE 802.11e for QoS support in wireless LANs. *Wireless Communications, IEEE*, 10(6):40 – 50, dec. 2003.

[10] Paal E. Engelstad and Olav N. Østerbø. Non-saturation and saturation analysis of IEEE 802.11e EDCA with starvation prediction. In *Proceedings of the 8th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems*, MSWiM '05, pages 224–233, New York, NY, USA, 2005. ACM.

[11] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda. Performance anomaly of 802.11b. In *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, volume 2, pages 836 – 843 vol.2, march-3 april 2003.

[12] Yang Xiao and J. Rosdahl. Throughput and delay limits of IEEE 802.11. *Communications Letters, IEEE*, 6(8):355 –357, aug. 2002.

[13] Qiang Ni. Performance analysis and enhancements for ieee 802.11e wireless networks. *Network, IEEE*, 19(4):21 – 27, july-aug. 2005.

[14] Polarcloud. Tomato replacement firmware for linksys wrt54gl v1.28. http://www.polarcloud.com/tomato_128.

[15] National Laboratory for Applied Network Research (NLANR) and Distributed Applications Support Team (DAST). Iperf, the udp and tcp bandwidth performance measurement tool v2.0.5. http://iperf.sourceforge.net.

[16] S. Pilosof, Ramachandran Ramjee, D. Raz, Y. Shavitt, and Prasun Sinha. Understanding TCP fairness over wireless LAN. In *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, volume 2, pages 863 – 872 vol.2, march-3 april 2003.

[17] M. Seyedzadegan, M. Othman, S. Subramaniam, and Z. Zukarnain. The TCP fairness in WLAN: A review. In *Telecommunications and Malaysia International Conference on Communications, 2007. ICT-MICC 2007. IEEE International Conference on*, pages 644 –648, may 2007.

[18] N. Blefari-Melazzi, A. Detti, I. Habib, A. Ordine, and S. Salsano. TCP Fairness Issues in IEEE 802.11 Networks: Problem Analysis and Solutions Based on Rate Control. *Wireless Communications, IEEE Transactions on*, 6(4):1346 –1355, april 2007.

[19] A.S. Tanenbaum. *Computer networks*. Prentice Hall PTR, 2003.

[20] Nihar Jindal, Jeffrey G. Andrews, and Steven Weber. Bandwidth-SINR Tradeoffs in Spatial Networks. In *Information Theory, 2007. ISIT 2007. IEEE International Symposium on*, pages 836 –840, june 2007.