Analysing Interface Bonding in 5G WLANs

Michael Dilmore, Angela Doufexi, and George Oikonomou,
Communication Systems & Networks,
Department of Electrical and Electronic Engineering,
University of Bristol, Merchant Venturers Building,
Woodland Road, Bristol, BS8 1UB
{M.Dilmore, A.Doufexi , G.Oikonomou}@bristol.ac.uk

***Abstract*—This work proposes a simple analytical model for interface bonding in 5G WLANs at the 2.4 GHz and 60 GHz ISM bands. Based on previous analysis of the IEEE 802.11 DCF by Bianchi and Chatzimisios, an expression for the predicted throughput of the bonded interface is given as a function of the number of competing wireless nodes in each network.**

The model is implemented and validated in MatLab using the Monte Carlo method. When applied to a practical interface bonding scenario, the model results suggest a practical limit of fifteen 2.4 GHz nodes when bonded with a 60 GHz interface, above which the resulting compound throughput is less than that of a single 60 GHz interface.

I. Introduction

State of the art mobile devices are increasingly capable of multi-homing, often featuring two, three, or more different wireless interfaces. The number of connectivity options is set to expand in the next generation of wireless networks with the introduction of several new WLAN (Wireless Local Area Network) standards, including IEEE 802.11ay [1], which provides speeds of up to 40 Gbps in the 60 GHz unlicensed Industrial-Scientific-Medical (ISM) band.

Wireless interface bonding, also known as wireless link aggregation or multi-streaming, exploits the multi-homing capability of modern mobile devices by aggregating interfaces in order to combine their data rates and/or provide redundancy, resulting in faster and more reliable links. The technique is of interest to cellular operators seeking to increase efficiency and exploit free-to-use resources, and also to network administrators at medium-to-large WLANs, wishing to add capacity at wireless bottlenecks.

The use of the technique over wireless interfaces presents several challenges. Factors such as device heterogeneity, the fast-fading wireless channel, Transport Layer re-ordering, and the stochastic delay of the IEEE 802.11 MAC, can have a damaging impact on throughput if improperly managed. In some cases bonding can under-perform the single-interface case. In order to exploit the increased opportunity for bonding presented by 5G, there is a need for better understanding of the circumstances in which its use is beneficial, and when it can be detrimental.

The main contributions of this work are the following:

(1) An accessible and flexible analytical model for interface bonding is proposed, including an expression for the predicted throughput of a bonded interface, which can be applied to any given bonding policy, and (2) the suggestion of a practical limit for the number of 2.4 GHz nodes when bonded with a 60 GHz wireless interface. The model is implemented in MatLab and validated using the statistical Monte Carlo method, and is shown to predict the compound throughput of a dual 2.4 GHz and 60 GHz bond with a reasonable degree of accuracy.

Few interface bonding analysis works exist, with Combes et al. [2] being the only notable example. However, the Combes model considers only the delay of the bond when applied to HD video streaming, and does not consider bond throughput.

The remainder of the paper is structured as follows. Section II reviews fundamental concepts and challenges in wireless bonding and provides an overview of current technology and research. Section III is a primer of closely related topics, including an overview of the IEEE 802.11 CSMA/CA algorithm and analytical works by Bianchi [3] and Chatzimisios [4]. Section IV presents the proposed bonding model in more detail. Section V is an evaluation and discussion of the results, and Section VI is a final conclusion.

II. Interface Bonding in 5G Wireless Networks

This section reviews fundamental concepts and challenges in wireless bonding and provides an overview of current related technology and research.

A. Definition

Interface bonding, link aggregation, multi-path streaming, trunking, bundling, striping, or teaming, is the technique of combining multiple network interfaces to present a single aggregated resource to the upper layers, and can be used for both wireless and wired interfaces. Using multiple interfaces to send a single stream of data, we can either: aggregate their data rates to increase performance, broadcast each packet over all links to improve delay, or use one link as a standby to improve reliability. This work distinguishes between wireless interface bonding, occurring at the Link, Transport, and Application Layers, and channel bonding, or carrier aggregation, which occurs at the Physical Layer.

B. Functional Components of Link Aggregation Systems

Based on [5] this work defines the three main functional components of a bonding solution that apply to both wired and wireless types: (1) The Link Monitoring Mechanism, responsible for monitoring the instantaneous capacity of links, (2) The Re-sequencing Unit, responsible for re-ordering frames at the receiver, and (3) The Scheduling Algorithm, which calculates an optimal bonding policy based on link state.

C. General Challenges

The theoretical maximum throughput of a bond is the sum of the throughputs of the individual interfaces, but in practice this is unlikely to be achieved for a number of reasons. The fast-fading nature of the wireless channel causes the instantaneous data rate and delay of a link to fluctuate randomly over small time intervals. Mobility, link adaptation, and the broadcast nature of the radio frequency (RF) channel further contribute to difficulties. Greater wireless hardware/software and performance heterogeneity make fully technology-agnostic approaches harder to implement. In addition, the inherently stochastic delay of the IEEE 802.11 Distributed Coordination Function (DCF) makes optimal bonding almost impossible in practice. These factors combine to make efficient wireless bonding generally more difficult compared to wired.

Aggregating heterogeneous links presents a particular problem for connection-oriented protocols such as the Transport Control Protocol (TCP). When links with different data rates are used to send a single stream of data, for example Wi-Fi and LTE, the slower paths have the effect of slowing down the faster paths. This leads to more out-of-order packets in the receive buffer, which adds an additional re-ordering overhead. TCP misinterprets these out-of-order packets as congestion and reduces its window size, further reducing throughput [6].

D. Overview of Current Wireless Bonding Technology and Research

In recent years, most link aggregation research has been concerned with three technologies: LTE Unlicensed, including LTE-U [7] and LTE-WLAN-Aggregation (LWA) [8], Channel Bonding in IEEE 802.11 [9], and Multipath TCP [10]. Other standards such as the Link Aggregation Control Protocol (LACP) [11] and the Multi-link Point-to-Point Protocol (MPPP) [12] also exist, which like MPTCP are designed specifically for Ethernet interfaces.

1) LTE Unlicensed: LTE Unlicensed is a group of 3GPP technologies that extend LTE into unlicensed ISM bands at 2.4 GHz and 5 GHz. LTE-U and LWA use Carrier Aggregation to bond together physical channels in both traditional LTE bands and the ISM bands, although each band is still subject to its own MAC entity. LWA is distinct in that it directly integrates Wi-Fi technology within the LTE infrastructure by bonding at the Packet Data Convergence Protocol (PDCP) Layer. All LTE Unlicensed technologies share the common goal of incorporating free-to-use spectral resources into the LTE infrastructure, but differ in the mechanisms used to access the medium and level of politeness. One of the biggest issues in WWAN-WLAN bonding is inter-RAT coexistence. When using the ISM bands in certain regulatory domains such as the UK, there a listen-before-talk (LBT) requirement, i.e. the nodes must sense the medium before transmitting. IEEE 802.11 satisfies this requirement by virtue of its DCF. LTE-U incorporates unlicensed ISM bands by integrating them directly at the Physical Layer without LBT. A number of works have demonstrated the detrimental effect of LTE-U on adjacent legacy 802.11 networks, including [13]. This issue has been the cause of a well-documented debate between the Wi-Fi Alliance and 3GPP over fair use of ISM resources. Several coexistence solutions have been proposed, such as as Wang et al. [14], which is based on cognitive channel switching and adaptive muting, and a Q-learning approach by Su et al. [15].

2) Multi-path TCP: The Multi-path Transmission Control Protocol (MPTCP) is an IETF standard for interface bonding at the Transport Layer that provides per-flow congestion control and two-tier data sequencing. MPTCP achieves aggregation by the use of multiple sub-flows, which are individually similar to regular TCP connections. Efficient MPTCP operation over heterogeneous wireless interfaces has been the subject of several works including [16], in which the author proposes a machine-learning based solution for aggregating multiple, heterogeneous wireless interfaces that maximises throughput using an autonomous parameter optimisation scheme.

3) Channel Bonding in IEEE 802.11: The IEEE 802.11n [9] standard defines a method for link aggregation called Channel Bonding that combines multiple, contiguous 20 MHz channels into a single 40 MHz channel in order to increase user throughput. As per the 802.11 DCF, Clear Channel Assessment (CCA) is performed on both channels before transmission, first on the primary channel, and then the secondary. IEEE 802.11ac is designed for 5 GHz and extends bonding capabilities to allow channel bandwidths of 80 MHz and 160 MHz. The channel access mechanism is also extended, with virtual sensing and CSMA-CA being used on the primary channel, and the secondary channels using only CCA. Upcoming 802.11-based technologies such as IEEE 802.11ay are likely to feature similar bonding mechanisms.

Chen et al. [17] show how the traditional RTS/CTS handshake performs poorly when channel bonding is utilised. In the 802.11ac standard, transmit signal power levels are restricted for different channel bandwidths, which has the effect of raising the sensing threshold and SNR, leading to an increase in the likelihood of collisions due to hidden nodes and receive errors respectively.

4) Other Technologies: LACP and MPPP are both examples of wired interface bonding. The Linux Bonding Driver [18] bonds together multiple network interfaces so they appear as a single aggregate resource to the kernel and upper protocol layers. The driver offers 7 different modes of transmission, including round robin, active-backup, 802.3ad, and broadcast. The Red Hat Teaming driver is a proprietary version of the Linux Bonding Driver that provides the same functionality but with a more modular, extensible architecture. Open vSwitch [19] offers similar bonding functionality. These solutions were designed primarily for wired devices but can be used for wireless interfaces with severe performance limitations. So et al. in [18] propose a module for the Linux Bonding Driver called New Load Balancing (NLB) designed specifically for wireless interfaces. The module measures the inter-arrival time of packets at the receiver using the packet-pair method and uses this to calculate the appropriate bonding schedule. The work shows that in clear channel conditions the algorithm achieves a throughput that is near to the sum of those of the individual interfaces. However, the NLB scheduler is slow to react to changes in channel load, taking 10 seconds to reach its peak throughput of 35 Mbps after the addition of co-channel interference on one of the interfaces, and even longer (almost 30 seconds) to reach its previous maximum once the interference has been removed.

III. Prerequisites

This section contains a primer of concepts and analytical models on which the proposed interface bonding model relies.

A. Overview of IEEE 802.11 Distributed Coordination Function

The IEEE 802.11 standard provides a framework for data communication over the free-to-use ISM bands at 2.4 GHZ, 5 GHz, sub 1 GHz, and mmWave ranges. The standard provides two channel access mechanisms. The Centralised Coordination Function (CCF), and the Distributed Coordination Function (DCF). In the CCF all data transmissions are scheduled by the controller using a super frame structure and a polling mechanism. The DCF uses a variation of the CSMA algorithm known as CSMA with Collision Avoidance (CSMA/CA). A node wishing to transmit data must wait a random amount of time between 0 and $(2^{BE}-1)$ before sensing the medium. If the medium is busy, the node picks a new random number between 0 and $(2^{BE}-1)$ before sensing the medium again. If the medium is idle the node may transmit its data. In the event of a lost ACK, a collision is inferred and the value of BE is incremented up to a maximum value. The process repeats until the transmission is successful or the maximum number of retransmissions is reached.

B. Related Analytical Works

Bianchi was the first to analyse the performance of the IEEE 802.11 CSMA/CA algorithm using a Discrete Time Markov Chain (DTMC). The model is comprised of a DTMC model representing node state and associated transition probabilities, and a system throughput formula derived from the node model. Bianchi noted that the backoff timer process of the IEEE 802.11 CSMA/CA, $b(t)$, is non-Markovian, as its value depends on the transmission history of the node, i.e. the number of transmission attempts. However, the embedded chain representing the backoff stage process, $s(t)$, is Markovian. If it is assumed that each packet collides with constant and independent probability p, then a bi-dimensional DTMC, $\{s(t), b(t)\}$, can be developed. Bianchi provides an expression for the probability that a station transmits in a random slot time:

$$\tau = \frac{2(1-p)}{(1-2p)(W+1)+pW(1-(2p)^m)} \qquad (1)$$

τ depends on the conditional collision probability p, which is the probability that, in a given time slot, at least one of the remaining n-1 nodes transmits. In steady state this gives:

$$p = 1-(1-\tau)^{n-1} \qquad (2)$$

Bianchi also defines the throughput as the average number of payload bits transmitted per slot time and provides the following formula:

$$S = \frac{P_s P_{tr} P}{(1-P_{tr})\delta + P_{tr}P_sT_s + P_{tr}(1-P_s)T_c} \qquad (3)$$

Chatzimisios et al define the access delay of the IEEE 802.11 DCF as the length of time between a frame arriving at the head of the source node transmission queue and it being successfully received at the sink, which includes the time taken for contention, acknowledgement, and inter-frame spacing:

$$D = E[X]E[Slot] \qquad (4)$$

Where $E[Slot]$ is the expected length of a slot time and $E[X]$ is given by:

$$E[X] = \frac{(1-2p)(W+1)+pW(1-(2p)^m)}{2(1-2p)(1-p)} \qquad (5)$$

Eq. 2 and Eq. 4 are central to the bonding model proposed in the following section.

IV. Proposed Analytical Model

This section describes the proposed WLAN interface bonding model in detail, including the main assumptions and limitations, and how the model is constructed.

A. Model Assumptions

The model inherits the assumptions and limitations of the original Bianchi model: (1) each frame collides with constant and independent probability, (2) the only source of frame loss is due to collisions (i.e. losses from congestion or path errors are ignored), and (3) an infinite number of retransmission attempts. The model also inherits the flexibility of the Bianchi model and can easily be modified to work with other IEEE 802.11 technologies by changing relevant PHY and MAC parameters accordingly. The model may also be updated to accommodate more advanced DCF models, such as those with finite retransmission attempts and Poisson traffic.

B. Model Construction

The maximum theoretical throughput of the bond is given as the sum of the throughputs of the individual interfaces. This is the highest possible throughput assuming no slow-down effects due to the difference in delay between interfaces. The equation for the maximum throughput is as follows:

$$S_{Bond} = S_1 + S_2 \qquad (6)$$

Where S_i represents the theoretical throughput of the i^{th} interface, as provided by Eq. 3.

In practical situations involving heterogeneous wireless interfaces and connection-oriented Transport protocols, the slower link has the effect of slowing down the faster link. This is because the Transport Layer must wait for all segments to

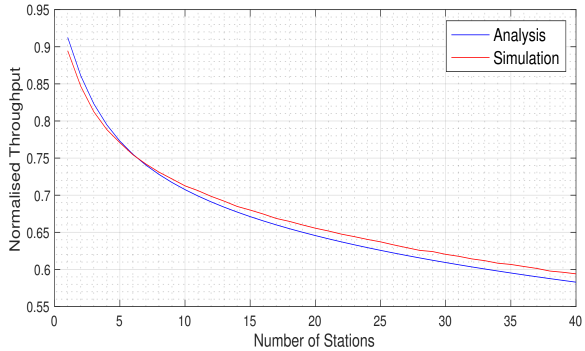

Fig. 1. Throughput vs. Number of Nodes for IEEE 802.11n.

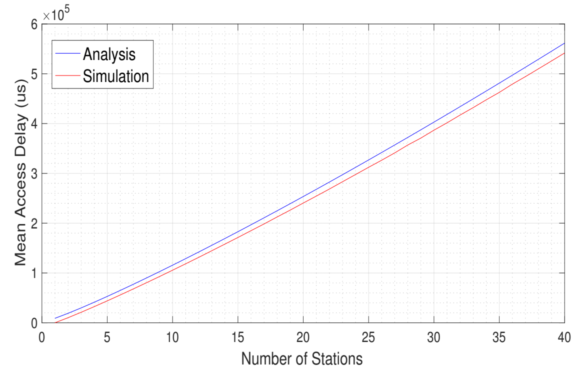

Fig. 2. Expected Delay vs. Number of Nodes for IEEE 802.11n.

be correctly received before it can pass the ordered data to the upper layers. The bigger the difference in the delay between interfaces, the longer the faster link has to wait and the bigger the overall impact on the compound throughput. Therefore the compound throughput is proportional to the difference in delay between links, which is represented in this work by the following equation:

$$S_{Pred} = D(S_1 + S_2) \tag{7}$$

Where D denotes the difference in the expected delay between the individual interfaces, as provided by equation Eq. 4. The above equation attempts to accurately predict the throughput of a given bond given any number of wireless nodes.

V. Evaluation

The first step in validating the proposed model is to ensure the validity of the existing analysis by Bianchi and Chatzimisios. Using the parameters in Table. I a Monte Carlo simulation was conducted in which a total of 300,000 frames were transmitted and the number of nodes varied from 1 to 40. The MAC parameters are taken directly from the relevant IEEE 802.11 standards, and the frame payload size chosen to reflect that in the original Bianchi experiment.

The results in Fig. 1 and Fig. 2 show that both models closely follow the simulation output. The minor divergence between the analytical and simulation results is caused by the decoupling assumption made by Bianchi that each frame collides with constant and independent probability, which does not capture the decreasing likelihood of collision as the contention window range becomes bigger. The Chatzimisios model is also based on Bianchi and makes the same decoupling assumption.

To test the predicted throughput hypothesis in Eq. 7, another simulation was conducted to measure the impact of the slower link on the resulting compound throughput using the same simulation parameters for a dual interface 2.4 GHz and 60 GHz bond. Again, a total of 300,000 frames were transmitted, with the same number being sent over both interfaces to model a simple round-robin bonding policy. The results of

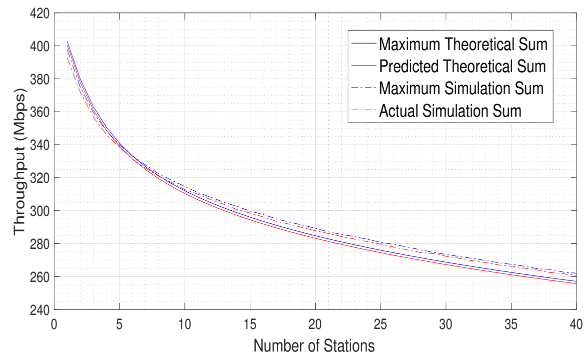

Fig. 3. Bond Throughput vs. Number of Nodes.

this experiment are shown in Fig. 3 and are plotted alongside the maximum theoretical throughput and predicted bond throughput as per Eq. 6 and Eq. 7 respectively.

In order to demonstrate the usefulness of the proposed model, the performance of a 2.4 GHz and 60 GHz bond is evaluated. The number of 60 GHz nodes is kept constant at 10, while the number of 2.4 GHz nodes is varied from 1 to 40. Again, a total of 300,000 frames were sent over both interfaces to simulate a round-robin policy. The analytical results are presented in Fig. 4 and are plotted alongside the theoretical throughput of a single 60 GHz interface for comparison.

TABLE I
Simulation Parameters

	IEEE 802.11n	IEEE 802.11ad
Slot Time (microseconds)	20	5
CW Minimum (slots)	15	15
CW Maximum (slots)	1023	1023
Frame Payload Size (bits)	8184	8184
MAC Header Size (bits)	288	240
PHY Header (bits)	192	192
ACK Header Size (bits)	112	112
SIFS Duration (microseconds)	10	3
DIFS Duration (microseconds)	30	13
Propagation Time (microseconds)	1	0.1
PHY Rate (Mbps)	54	385

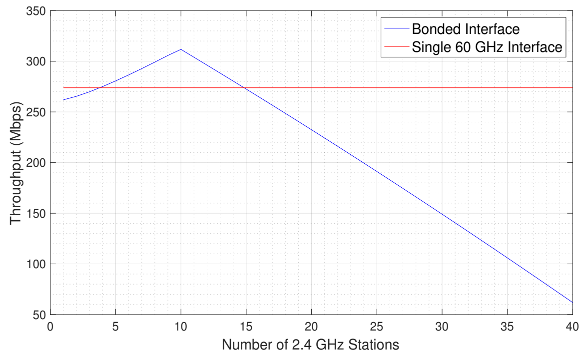

Fig. 4. Bond Throughput vs. Number of Nodes (60 GHz nodes = 10)

A. Discussion

The results of the bonding simulation in Fig. 3 show that the proposed model predicts the throughput of the bond with a reasonable degree of accuracy. The minor divergence between analytical and simulation results can be explained by the Bianchi decoupling assumption described above.

Applying the model to a dual interface 2.4 GHz and 60 GHz bond in Fig. 4, shows that the compound throughput falls off as more 2.4 GHz nodes are added. At approximately 14 nodes, the resulting compound throughput is roughly equal to that of a single 60 GHz interface, and throughput continues to degrade as more 2.4 GHz nodes are added, which suggests a practical upper limit of 14 nodes before the use of bonding becomes detrimental

VI. Conclusion

This work presents a simple and accessible model for analysing the performance of interface bonding over multiple IEEE 802.11-based technologies. The model is validated using Monte Carlo simulations and is shown to accurately predict the compound throughput for any number of competing nodes. When applied in a practical 5G bonding scenario, the model can help identify when interface bonding is beneficial and when its use is detrimental, by providing practical upper limits on the number of wireless nodes present in each network.

References

[1] Y. Ghasempour, C. R. C. M. da Silva, C. Cordeiro, and E. W. Knightly. IEEE 802.11ay: Next-generation 60 GHz communication for 100 Gb/s Wi-Fi. *IEEE Communications Magazine*, 55(12):186–192, DECEMBER 2017.

[2] R. Combes, H. B. A. Sidi, and S. Elayoubi. Multipath streaming: Fundamental limits and efficient algorithms. *IEEE Journal on Selected Areas in Communications*, 35(1):188–199, Jan 2017.

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

[4] P. Chatzimisios, A. C. Boucouvalas, and V. Vitsas. Packet delay analysis of IEEE 802.11 MAC protocol. *Electronics Letters*, 39(18):1358–1359, Sept 2003.

[5] Karim Habak, Khaled A. Harras, and Moustafa Youssef. Bandwidth aggregation techniques in heterogeneous multi-homed devices: A survey. *Computer Networks*, 92:168 – 188, 2015.

[6] C. M. Arthur, A. Lehane, and D. Harle. Keeping Order: Determining the effect of TCP packet reordering. In *International Conference on Networking and Services (ICNS '07)*, pages 116–116, June 2007.

[7] S. A. A. Termizi, S. Rahmat, and M. A. Qadeer. LTE-U over Wi-Fi. In *2016 Thirteenth International Conference on Wireless and Optical Communications Networks (WOCN)*, pages 1–7, July 2016.

[8] P. Nuggehalli. LTE-WLAN aggregation [Industry Perspectives]. *IEEE Wireless Communications*, 23(4):4–6, August 2016.

[9] 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, Oct 2009.

[10] A. Ford, C. Raiciu, M. Handley, and O. Bonaventure. TCP extensions for multipath operation with multiple addresses. RFC 6824, RFC Editor, January 2013.

[11] IEEE standard for local and metropolitan area networks–Link Aggregation. *IEEE Std 802.1AX-2008*, pages 1–163, Nov 2008.

[12] K. Sklower, B. Lloyd, G. McGregor, D. Carr, and T. Coradetti. The PPP Multilink Protocol (MP). RFC 1990, RFC Editor, August 1996.

[13] H. Ben Hafaiedh, I. El Korbi, L. A. Saidane, and A. Kobbane. LTE-U and WiFi coexistence in the 5 GHz unlicensed spectrum: A survey. In *2017 International Conference on Performance Evaluation and Modeling in Wired and Wireless Networks (PEMWN)*, pages 1–7, Nov 2017.

[14] Z. Wang, H. M. Shawkat, S. Zhao, and B. Shen. An LTE-U coexistence scheme based on cognitive channel switching and adaptive muting strategy. In *2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, pages 1–6, Oct 2017.

[15] Y. Su, X. Du, L. Huang, Z. Gao, and M. Guizani. LTE-U and Wi-Fi Coexistence Algorithm Based on Q-Learning in Multi-Channel. *IEEE Access*, 6:13644–13652, 2018.

[16] Y. Kon, M. Ito, N. Hassel, M. Hasegawa, K. Ishizu, and H. Harada. Autonomous parameter optimization of a heterogeneous wireless network aggregation system using machine learning algorithms. In *2012 IEEE Consumer Communications and Networking Conference (CCNC)*, pages 894–898, Jan 2012.

[17] Y. Chen, C. Chuang, and K. Shih. A channel bonding MAC protocol for IEEE 802.11ac WLANs. In *2017 Ninth International Conference on Ubiquitous and Future Networks (ICUFN)*, pages 839–841, July 2017.

[18] J. So, J. Park, and A. Eliab. Wireless Interface Bonding Supporting In-Order Delivery and Automatic Load Balancing. *IEEE Transactions on Vehicular Technology*, 66(6):5190–5203, June 2017.

[19] D. T. Bui and K. Aberkane. A generic interface for Open vSwitch. In *2016 IEEE NetSoft Conference and Workshops (NetSoft)*, pages 53–57, June 2016.