Aalborg Universitet

![Aalborg University logo](AALBORG UNIVERSITY)

Radio-Aware Multi-Connectivity Solutions based on Layer-4 Scheduling for Wi-Fi in IIoT Scenarios

Fink, Andreas; Mogensen, Rasmus Suhr; Rodriguez, Ignacio; Kolding, Troels; Karstensen, Anders; Pocovi, Guillermo

Radio-Aware Multi-Connectivity Solutions based on Layer-4 Scheduling for Wi-Fi in IIoT Scenarios

Andreas Fink*, Rasmus Suhr Mogensen*, Ignacio Rodriguez*, Troels Kolding†,
Anders Karstensen*, Guillermo Pocovi†

*Wireless Communication Networks Section, Department of Electronic Systems, Aalborg University, Denmark
†Nokia Bell Labs, Aalborg, Denmark

Abstract—Due to mobility and interference, using enterprise Wi-Fi for communication in industrial networks can result in control loop latencies exceeding 100 ms for at least 0.1% of the time, even in those cases where Wi-Fi handover-specific parameters have been optimized, making the technology unfit for Industrial IoT (IIoT) with strict communication reliability requirements. To improve its performance, this paper presents a novel approach towards the design and implementation of a radio-aware multi-connectivity concept using a layer-4 scheduling mechanism. Two packet scheduling mechanisms are presented: packet duplication and best path scheduling. A mobility co-ordinator scheme is used to improve the performance of the packet schedulers by preventing simultaneous handovers and ensures the STAs connect to different APs. By using this multi-connectivity solution, a significant performance improvement was observed, cutting down the latencies of the system to 30-80 ms at the 99.9%-ile of reliability (depending on the operational conditions). Furthermore, by applying the proposed schemes, Wi-Fi handovers delays can be fully mitigated allowing for true seamless roaming in mobile conditions.

I. INTRODUCTION

One of the promises of Industry 4.0 is the increased coordination between various types of Industrial IoT (IIoT) equipment to both improve decision making and increase efficiency of the production in general. This functionality comes at the cost of having new and more sophisticated systems with more demanding communication requirements. With Autonomous Mobile Robots (AMR) and other dynamic equipment utilizing high bandwidth to share data and receive latency-sensitive critical control traffic over a wireless interface, existing wireless technologies will be challenged to support this. As IEEE 802.11 Wi-Fi is commonplace among industrial plants, a large portion of IIoT is designed to utilize this technology. In our previous work, we showed that Wi-Fi technology has challenges in meeting the requirements of latency-sensitive IIoT [1], especially due to interruptions in communication when having mobility across Access Points (APs). While the latest iteration of the technology, Wi-Fi 6, contains features that provide better support for critical traffic flows and Quality of Service differentiation, it does not solve the issues due to mobility, which may have a significant impact on IIoT latency-sensitive applications.

Wi-Fi is commonly deployed in infrastructure mode, with a single Wi-Fi Station (STA) interface on a device connecting to a nearby AP. If the signal between the STA and AP deteriorates, the STA will scan for nearby APs and perform a handover to establish a new connection. This results in a drastic increase of the latency until the new connection is established [1]. Several amendments have been introduced to the IEEE 802.11 standard with the goal of reducing handover gaps and allow for seamless roaming, such as IEEE 802.11r to reduce the handover duration and IEEE 802.11k which monitors nearby APs improving overall scanning time. However, the mobility latency levels achieved by applying these techniques might not be sufficient to support certain very demanding IIoT applications.

While one solution for this might be to implement a vendor-specific solution modifying the lower layers of the protocol, considering a higher layer multi-connectivity approach [2] that utilizes multiple simultaneous Wi-Fi STAs on the same device (multi-STA configuration), connected to different APs, might be a better choice, as it would ensure device and network interoperability. Simply increasing the number of active Wi-Fi STAs per device is not ideal, as if they are not properly managed, they might increase the collisions and network load, impacting the overall performance of the system [3]. However, when proper STA coordination is introduced, higher reliability and significantly decreased Packet Error Rate (PER) is expected [4]. Furthermore, by having multiple STAs available, new possibilities emerge in terms of AP connection steering, allowing for an optimized transition between APs, minimizing the impact of handovers in the performance. A similar concept was found to be successful in public LTE networks [5].

A common approach for multi-connectivity is to utilize Multipath TCP (MPTCP) [6]–[9], an extension to TCP allowing a single-connection to establish multiple subflows over different paths. However, this comes at a cost of degraded throughput [10] and lack of data multiplexing and prioritization [11]. Another approach is through packet duplication, which can significantly reduce high delays and jitter [12]. Duplicating traffic using different technologies have likewise been previously considered, where [13] and [14] both demonstrate how Wi-Fi and 4G LTE can improve the communication latency by making use of the separate medium access control schemes.

Utilizing multi-connectivity, either over multiple technologies, or simply using redundancy in a single technology, can deliver improved latency performance which is more suitable for critical IIoT devices. In this paper, we aim at leveraging two low-latency mobile multi-connectivity implementations

over Wi-Fi that are further enhanced by considering information from the radio layer, such as signal strength or connection states. This paper presents a novel approach, validated with enterprise off-the-shelf Wi-Fi STAs and APs (so that it can be directly reused in other deployment scenarios), that utilizes contextual information from the Wi-Fi STAs to manage how the traffic is routed and to control to which APs the STAs are connected to. We describe the design, implementation, and experimental validation of two multi-connectivity scheduler schemes and complementary mobility coordinators. To determine the benefits of the proposed solutions, their performance is evaluated in a realistic industrial environment with emphasis on reliable latency (i.e., the latency achieved with high probabilities such as the 99.9%-ile).

The rest of the paper is structured as follows: Section II details our Wi-Fi multi-connectivity solutions. Section III introduces the experimental environment, the setups, and the different radio configurations tested for the multiple multi-connectivity schemes. Section IV presents the latency performance measurements results. Section V contains the discussion of the results and highlights areas of potential improvements. Finally, section VI concludes the paper.

II. Radio-aware Scheduling Schemes for Wi-Fi

The proposed approach to multi-connectivity is based on a customized radio-aware layer-4 (transport-layer) packet scheduler to control the traffic flow through two Wi-Fi STAs. This scheduler bases its decisions on radio properties such as connection state and Received Signal Strength Indicator (RSSI) to improve the Wi-Fi system performance. To further enhance the performance, a Mobility Coordinator (MC) is introduced to ensure AP diversity (by preventing the STAs from connecting to the same AP) and to avoid simultaneous handovers. As a reference, the multi-STA components are illustrated in Fig. 1. Two schemes were designed and implemented based on these elements: 1) Packet Duplication (PD), and 2) Best Path Scheduling (BPS). These schemes implement different scheduler configurations which are assisted by the mobility coordinator. The performance of the proposed schemes can be further increased by proper network planning.

A. Wi-Fi Packet Scheduler

1) PD: in this scheme, the same packet is sent over the two STA interfaces, yielding significant performance improvements in terms of both communication and reliability. However, introducing a significant amount of redundant information to be transmitted over the medium increases in turn the average time to access the medium for all devices on the network due to the Listen-before-Talk (LBT) channel access mechanism. With a goodput of less than 50% when also considering packet headers, this can severely harm the overall throughput, especially if multiple devices utilize this method. This issue will be mitigated with the assistance of the mobility coordinator.


Fig. 1: Overview of multi-STA components: STA 1 and 2, packet scheduler and mobility coordinator.

2) BPS: the device uses single-connectivity while taking advantage of the presence of a secondary STA. In contrast to MPTCP, where a shortest-RTT scheduler is used to determine the best path, our approach uses the RSSI of the STAs to steer the traffic. The device seamlessly switches between the STAs depending on the RSSI as both are connected to an AP, chosen based on a 5 dB margin. If the STA currently being used for communication abruptly loses its connection, the secondary STA can immediately take over. The margin of 5 dB is chosen to avoid excessive switches between the STAs for scenarios with similar RSSI values.

B. Wi-Fi Mobility Coordinator

When the connection between a STA and an AP is degraded significantly (measured either through the RSSI or by detecting a connection loss), the STA will scan for other eligible APs nearby and then roam to the one with highest RSSI. As the antennas of both STA are approximately collocated, if no effort is put into coordinating the two STAs, they may experience very similar channel conditions, choosing to scan and roam between APs simultaneously. This will hurt the overall performance of either of the presented packet scheduler configurations and, thus, improvements are desirable.

Both STAs contain a list of eligible Basic Service Set IDs (BSSIDs) which are used during network scans to choose the serving AP. To introduce coordination between the two STAs, a blacklist is maintained by the mobility coordinator to prevent both STAs from connecting to the same AP. The coordinator will furthermore periodically check if the STAs need to roam to a new AP with a fixed periodicity (2.5 s in our case). In our implementation, roaming events will be triggered based on a RSSI threshold of -85 dBm to minimize the number of handover events. When the threshold is reached for a STA, the BSSID of the previous AP is added to the blacklist temporarily to force the disconnection. This threshold was chosen to allow for a full scan and handover before the STA would lose connection to its current AP.

The two algorithms for the combined packet scheduler and mobility coordinators are illustrated in Fig. 2 and 3 for the PD and BPS packet schedulers, respectively. The objective of the mobility coordinator differs slightly depending on the scheduler. For the PD scheme, the coordinator prioritizes uptime on both interfaces while keeping track of associations and dissociations for each STA to maintain the blacklist. For

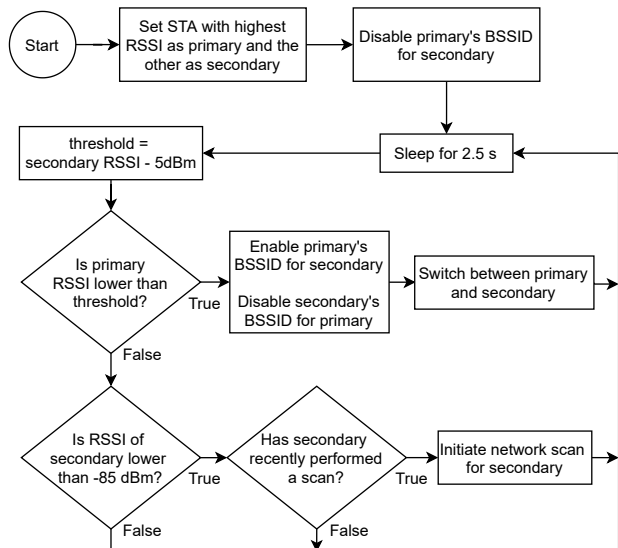Fig. 2: Flowchart for the packet duplication (PD) scheme.

Fig. 3: Flowchart for the best path scheduling (BPS) scheme.

the BPS algorithm, the coordinator will only initiate roaming events for the secondary (idle) interface.

C. Wi-Fi Network Planning

Deploying APs such that there are overlaps in coverage areas is necessary to fully utilize the mobility coordinator. Because of the LBT mechanisms of Wi-Fi, the average time until the medium can be accessed will be highly dependent on the number of active devices. The optimal performance of the mobility coordinator under the PD scheme will be achieved when two overlapping APs utilize different frequency

channels. Under other spectrum configuration circumstances, both STAs might experience the same average time until medium access.

III. EXPERIMENTAL TEST SETUP

The performance evaluation of the designed and implemented Wi-Fi multi-connectivity schemes was performed at the AAU 5G Smart Production Lab at Aalborg University, Denmark [15]. This industrial environment (shown in Fig. 4) is equipped with three ceiling-mounted CISCO MR36 Enterprise Wi-Fi 6 APs [16] deployed throughout the lab as illustrated in Fig. 5. In order to trigger the mobility aspects of our Wi-Fi solution performance evaluation, a MiR200 AMR (also shown in Fig. 4) was used. The AMR was configured to follow a specific route through the lab as illustrated in Fig. 5, carrying the implemented multi-STA device around at a maximum speed of 1.5 m/s.

The multi-STA was configured using wpa_supplicant v. 2.9 [17] which is used to communicate with the driver for the Wi-Fi STAs and is furthermore used to handle roaming and key negotiation. The number of frequency channels which the STAs scan was optimized to match the number of APs in the testing environment (three), ensuring no overlap of networks with identical frequencies. To enable the multi-connectivity aspect of the setup, an improved version of the multi-access gateway presented in [13] was used. In uplink, the gateway encapsulates traffic from an end-device (i.e. mobile robot) and transmits it through specified interfaces, in our case two Intel Wi-Fi 6 AX200 network cards, to another gateway on the network side from which the traffic is decapsulated and transmitted to another end-device (i.e., network server). The procedure is similar in downlink but with encapsulation happening at the network-side and decapsulation happening at the device-side. This is illustrated in Fig. 6. During transmissions between gateways, an additional 44 bytes of headers and metadata from the encapsulation are added per frame, excluding technology-specific headers. The use of the gateways in this setup will furthermore introduce a calibrated delay of ∼0.12 ms per frame from the two Ethernet transmissions and processing in the gateway.


Fig. 4: Overview of the industrial environment and the multi-STA measurement setup, including one of the ceiling-mounted APs and the AMR used for mobility.

Fig. 5: Floor plan of the test environment, including AP locations (green, blue, and red dots) and coverage areas, the AMR measurement route and the location of the static traffic sources for the generation of background traffic (white dots).

The multi-connectivity performance was evaluated between two end-devices (a PC mounted on the AMR at the device-side and a server at the network-side) using the Linux ping functionality, providing an insight on round-trip time (RTT) and packet error rate (PER) statistics. A packet size of 64 B and an inter-packet interval of 50 ms was used. The packet interval was chosen to obtain sufficient samples while capturing the impact of handovers on the performance, for which typical AMR traffic models are unsuitable. Simultaneously, RSSI was monitored. The following four configurations were examined in the experimental testing:

1) BPS with mobility coordination over dedicated channels.
2) PD without mobility coordination over dedicated channels.
3) PD with mobility coordination over dedicated channels.
4) PD with mobility coordination and frequency re-use.

Furthermore, all configurations were tested in idle networks with no other traffic than the one generated by the multi-STA device, and also with background traffic where two static STAs were used to load each AP with a constant controlled traffic load of 10 Mbit/s uplink and 10 Mbit/s downlink. The traffic load was chosen to reflect a low-medium usage of the network and to observe an impact on the latency without reaching
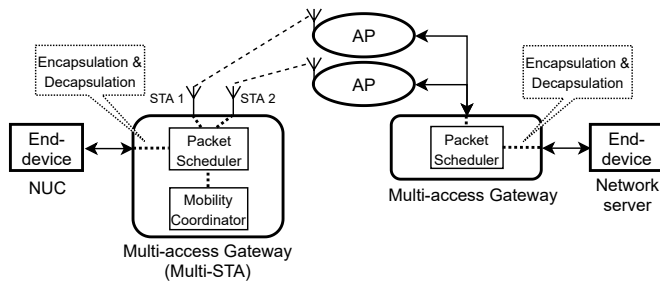

Fig. 6: Test setup using multi-access gateways. Solid connections represent Ethernet connections, while dashed connections represent Wi-Fi.

congestive conditions. For configurations with BPS scheme, the gateway on the network-side was configured to steer traffic to the STA that it had last received data from to ensure single-connectivity behavior in both uplink and downlink. In those cases, where the mobility coordinator was not used (the PD case without mobility coordination over dedicated channels), the STAs were configured to initiate network scans when they reach -85 dBm RSSI, from which they will search for another suitable AP nearby and initiate the handover. This is done to reduce the stickiness of the connection to an AP, as it would otherwise remain connected until the connection is lost (at ∼ -93 dBm).

IV. WI-FI PERFORMANCE RESULTS

We first take a look to the performance of the PD scheme with and without mobility coordination. The RSSI measured at each STA is illustrated in Fig. 7 for both schemes. The data confirms that if the STAs are used without mobility coordination, the RSSI will, as expected, be very similar due to the low spatial diversity of the setup. Spatial diversity in the setup could be improved by separating as much as possible the antennas of the different STAs, but that would not be realistic as, in practice, industrial hardware imposes restrictive constraints on the communication modules and antenna location. With the mobility coordinator enabled, the RSSI-traces for the different STAs became uncorrelated due to each STA connecting to a different AP. It is, however, also observed that when the signal strength degrades to -85 dBm (corresponding to the interval between 55-85 s) and if no
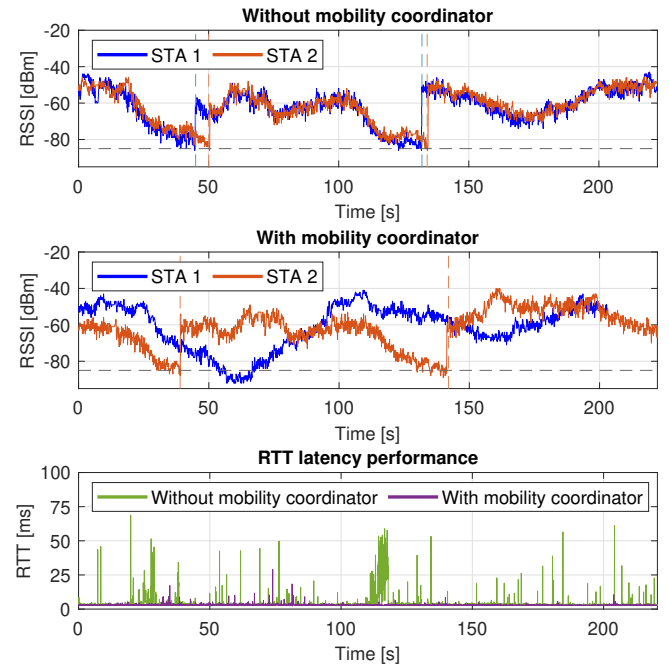

Fig. 7: RSSI traces and overall Wi-Fi RTT latencies for two different PD schemes: with and without mobile coordinator. The -85 dBm threshold is highlighted with the horizontal line, and handovers are highlighted with vertical lines.

Fig. 8: Empirical CCDFs of the RTT for single- and multi-connectivity configurations operating over dedicated channels.

TABLE I: Summary of RTT latency statistics and PER measurement results for the single- and multi-connectivity configurations operating over dedicated channels.

Test Configuration	Coordination scheme	Channel condition	Min	Avg	99.9%-ile	Jitter	PER
Optimized single-connectivity [1]	Without mobility coordinator	Idle network	1.48 ms	3.23 ms	57.60 ms	1.56 ms	0.044%
		Background traffic	1.29 ms	6.26 ms	143.00 ms	4.78 ms	0.065%
Best path scheduling	With mobility coordinator	Idle network	2.23 ms	3.36 ms	55.60 ms	1.00 ms	0.005%
		Background traffic	2.24 ms	4.80 ms	80.70 ms	2.94 ms	0.028%
Packet duplication	Without mobility coordinator	Idle network	2.32 ms	3.88 ms	71.10 ms	1.42 ms	0.018%
		Background traffic	1.99 ms	4.87 ms	102.00 ms	2.66 ms	0.041%
	With mobility coordinator	Idle network	2.14 ms	3.15 ms	30.80 ms	0.58 ms	0%
		Background traffic	2.06 ms	4.12 ms	77.10 ms	1.75 ms	0.001%

eligible AP can be reached by the secondary STA (this happens when the AMR is only in coverage with AP 3 in Fig. 5), the STA will fully disconnect and either remain in a searching state until a suitable AP is found, or it will ping-pong between reconnecting to the previous AP and disconnecting due to low RSSI. When considering the RTT latency performance, the results illustrate that using the mobility coordinator translates into significantly higher stability (reduced amount of latency spikes, and spikes of shorter duration) than for the uncoordinated configuration.

Fig. 8 displays the empirical complementary cumulative distribution functions (CCDF). Each of the presented results sets were computed from more than 100,000 RTT latency samples of each tested configuration. Key latency statistics and PER are summarized in Table I. As a reference in both Fig. 8 and Table I, results from the optimized single-connectivity mobility performance test presented in [1] are also included. The results indicate that by having two STAs available in a single device, using either BPS or PD with mobility coordinator, significant latency improvements can be achieved for the 95%-ile and above as compared to the single-connectivity reference. While taking the RSSI of the STAs into account by using the mobility coordinator will result in a more robust connection, the main improvement stems from the multi-STA being able to fully mitigate the latency impact

of handovers. Apart from latency, this has also a positive effect on the PER which is reduced significantly in BPS, and almost completely in PD. In the cases where background traffic was present, the increase in latency is significantly lower in BPS and PD as compared to the single-connectivity configuration. In the case of PD multi-connectivity without mobility coordination, the performance is very similar to the one for single-connectivity for the idle network case, while a gain is observed for the case with background traffic. At 99.9%-iles, the performance of the BPS and PD schemes with mobility coordination and idle network can be as low as 56 and 31 ms, respectively, as compared to 58 ms in the single-connectivity case (4-46% gains). When background traffic is present, the gains for the multi-connectivity schemes are even larger (43-46%).

All previous results assumed some level of network planning and were obtained using individual dedicated frequency channels for each of the APs. However, the impact of uncoordinated deployments where the APs operate under frequency re-use was also evaluated. This evaluation was done for the PD scheme and the results are shown in Fig. 9. The results indicate that if dedicated frequency channels cannot be guaranteed for each AP to avoid overlapping regions, the performance of the PD multi-connectivity scheme will be slightly degraded in the presence of background traffic.

Fig. 9: Empirical CCDFs of the RTT for the packet duplication multi-connectivity schemes.

V. DISCUSSION

Using mobility coordination will result in increased performance when multiple APs on dedicated channels are available, with PD performing better than BPS both in terms of latency and PER. However, as PD causes an increase of the overall load of the network, thus causing interference for other devices, BPS could be a more appealing scheme to be applied for multi-connectivity. This will depend on both the network setup and amount of active Wi-Fi devices, but the performance may differ if a different type of traffic load is applied, e.g. exponential traffic patterns as compared to a constant load.

While the presented algorithms do improve the latency and PER performance, some areas for potential improvement have been identified. In regions where only one the primary STA remains operational due to impossibility of the secondary STA to find a suitable AP (such as in the PD scheme with mobility coordination presented in Fig. 7), it would be beneficial to force the secondary STA to connect to the same AP as the primary STA has, despite of having correlated RSSI, a performance gain could be achieved, especially in the presence of background traffic. As an alternative, fine tuning of the RSSI thresholds for the mobility coordinator could also result in an improved performance. If end-device location information is available from an external positioning system, the scheduling algorithms and mobility coordinator could be enriched by including this information in their decisions, eliminating the need for scanning for channels and neighbor APs.

VI. CONCLUSION

In this paper, a novel approach to introducing multi-connectivity using off-the-shelf Wi-Fi hardware configurations is presented. We present a custom transport-layer packet scheduler located at the edges, therefore not requiring any changes to the network itself, as opposed to e.g. proprietary solutions. This approach uses knowledge of the Wi-Fi connection (e.g. connection state and RSSI) to improve both how the traffic is steered and to introduce a mobility coordinator between the multiple STAs in the device, such that they connect to different APs and avoid simultaneous handovers. Two schemes are evaluated: 1) best path scheduling utilizing a primary STA with highest RSSI and seamlessly switching

to the secondary STA when the signal degrades, and 2) packet duplication over the two STAs simultaneously.

The experimental performance evaluation showed that both multi-connectivity schemes improve the Wi-Fi performance as compared to the single-connectivity case. Latency improvements of up to 46% were observed at the 99.9-%iles, lowering the latency from 143 ms to 77-81 ms in the presence of background traffic. Using the mobility coordinator ensures correlated links for the different STAs, which translates into non-simultaneous handovers thus, fully mitigating the impact of mobility between different APs in the communication, resulting in a seamless roaming operation. Packet duplication was found to be the best performing scheme, but comes at the operational cost of having increased load in the system, thus the operational conditions should be carefully analyzed before prioritizing it over the best path scheduling scheme.

REFERENCES

[1] A. Fink *et al.*, "Empirical Performance Evaluation of Enterprise Wi-Fi for IIoT Applications Requiring Mobility," unpublished, submitted to European Wireless 2021.

[2] E. J. Khatib *et al.*, "Multi-Connectivity for Ultra-Reliable Communication in Industrial Scenarios," in *IEEE Vehicular Technology Conference (VTC2019-Spring)*, 2019, pp. 1–6.

[3] M.-T. Suer *et al.*, "Multi-connectivity as an enabler for reliable low latency communications—an overview," *IEEE Communications Surveys Tutorials*, vol. 22, no. 1, pp. 156–169, 2020.

[4] J. J. Nielsen *et al.*, "Reliability and Error Burst Length Analysis of Wireless Multi-Connectivity," in *International Symposium on Wireless Communication Systems (ISWCS)*, 2019, pp. 107–111.

[5] M. Lauridsen *et al.*, "Reducing Handover Outage for Autonomous Vehicles with LTE Hybrid Access," in *IEEE International Conference on Communications (ICC)*, 2018, pp. 1–6.

[6] C. Raiciu *et al.*, "How Hard Can It Be? Designing and Implementing a Deployable Multipath TCP," in *USENIX Conference on Networked Systems Design and Implementation (NSDI)*, 2012.

[7] A. Croitoru *et al.*, "Towards Wifi Mobility without Fast Handover," in *USENIX Conference on Networked Systems Design and Implementation (NSDI)*. USENIX Association, May 2015, pp. 219–234.

[8] K. Nguyen *et al.*, "Improving WiFi networking with concurrent connections and multipath TCP," in *IEEE International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM)*, 2013, pp. 1–3.

[9] D. Rico *et al.*, "A Survey of End-to-End Solutions for Reliable Low-Latency Communications in 5G Networks," *IEEE Access*, vol. 8, pp. 192 808–192 834, 2020.

[10] M. R. Palash *et al.*, "MPWiFi: Synergizing MPTCP Based Simultaneous Multipath Access and WiFi Network Performance," *IEEE Transactions on Mobile Computing*, vol. 19, no. 1, pp. 142–158, 2020.

[11] R. S. Mogensen *et al.*, "Selective Redundant MP-QUIC for 5G Mission Critical Wireless Applications," in *IEEE Vehicular Technology Conference (VTC2019-Spring)*, 2019, pp. 1–5.

[12] H. Zhang *et al.*, "WiFi and Multiple Interfaces: Adequate for Virtual Reality?" in *IEEE International Conference on Parallel and Distributed Systems (ICPADS)*, 2018, pp. 220–227.

[13] R. S. Mogensen *et al.*, "Implementation and Trial Evaluation of a Wireless Manufacturing Execution System for Industry 4.0," in *IEEE Vehicular Technology Conference (VTC2019-Fall)*, 2019, pp. 1–7.

[14] M. Suer *et al.*, "Evaluation of Multi-Connectivity Schemes for URLLC Traffic over WiFi and LTE," in *IEEE Wireless Communications and Networking Conference (WCNC)*, 2020, pp. 1–7.

[15] I. Rodriguez *et al.*, "5G Swarm Production: Advanced Industrial Manufacturing Concepts Enabled by Wireless Automation," *IEEE Communications Magazine*, vol. 59, no. 1, pp. 48–54, 2021.

[16] "CISCO MR36 Cloud-managed Wi-Fi 6 (802.11ax)." [Online]. Available: https://meraki.cisco.com/product/wi-fi/indoor-access-points/mr36/

[17] "Linux WPA Supplicant (IEEE 802.1X, WPA, WPA2, RSN, IEEE 802.11i)." [Online]. Available: https://w1.fi/wpa_supplicant/