In-Band Multi-Connectivity with Local Beamtraining for Improving mmWave Network Resilience

Nina Grosheva
nina.grosheva@imdea.org
IMDEA Networks Institute and
Universidad Carlos III de Madrid
Madrid, Spain

Rizqi Hersyandika
rizqi.hersyandika@kuleuven.be
KU Leuven
Leuven, Belgium

Joerg Widmer
joerg.widmer@imdea.org
IMDEA Networks Institute
Madrid, Spain

Sofie Pollin
sofie.pollin@kuleuven.be
KU Leuven
Leuven, Belgium

ABSTRACT

Multi-connectivity is considered a key enabler for 5G networks and beyond, aiming to enhance capacity by combining multiple communication links in the same or different bands. Similarly, in cell-free networks all Access Points (APs) jointly serve users in the same band, boosting capacity through enhanced spectral efficiency. Both approaches can be very effective in Millimeter-Wave (mmWave) networks by addressing key issues of reliability and robustness due to the multiple simultaneous links. Furthermore, the use of narrow directional beams in mmWave spatially separates the signals, allowing for in-band multi-connectivity through local beamtraining. Such in-band multi-connectivity would be an alternative design to traditional cell-free networks that does not rely on phase-coherent processing or centralized methods for interference suppression. The physical layer processing and resource allocation problem then simplifies to a local beamtraining challenge, making these networks easier and simpler to implement and deploy, as any connection just has to train and maintain the local beam. We validate this approach by designing a multi-connectivity mmWave network with minimal network synchronization, relying solely on analog beamforming for spatial separation. Our evaluation results demonstrate that in-band multi-connectivity with 4 asynchronous and independent links can provide uninterrupted service even in dense, high-traffic scenarios, compared to up to 20% of service loss in a standard single-connectivity deployment. Distributing the traffic across multiple APs also had throughput gains of up to 30%, showing that multi-connectivity mmWave networks can provide a high-throughput, reliable and stable service for next-generation applications.

CCS CONCEPTS

• **Networks** → **Network architectures**; **Network simulations**; **Network performance analysis**; **Network reliability**.

KEYWORDS

Millimeter Wave, Multi-Connectivity, IEEE 802.11ad/ay, ns-3

ACM Reference Format:
Nina Grosheva, Rizqi Hersyandika, Joerg Widmer, and Sofie Pollin. 2023. In-Band Multi-Connectivity with Local Beamtraining for Improving mmWave Network Resilience. https://doi.org/10.1145/3616388.3617529

1 INTRODUCTION

The exponential rise of wireless devices and new, high-performance applications continuously increase demands from wireless networks. The need for higher capacity, lower latency and higher reliability demand innovative designs and an evolution of current networks. Network densification and small-cell deployments have been used to increase capacity and data rates. However, it has led to problems with intra-cell interference and poor performance for cell-edge users [5]. The cell-free concept has been proposed as an alternative design to improve coverage and capacity and provide uniform service for all users [19, 20]. In cell-free networks, a large number of Access Points (APs) are densely distributed and jointly and simultaneously serve clients. The idea is that by maintaining simultaneous connections with multiple APs we can exploit the spatial diversity against shadow fading and improve network coverage. Additionally, by removing the division of the network into cells and having APs jointly serve clients, cell-free networks can provide more uniform service with a user-centric approach [15]. As the client simultaneously communicates with multiple APs, it is necessary to have tight coordination and synchronization across the whole network to enable phase-coherent signal processing across the geographically distributed APs [4]. This remains a key challenge for the further development of cell-free networks.

More practical approaches to exploiting link diversity have been standardized in 3GPP as multi-connectivity. In-band multi-connectivity similar to cell-free has been introduced as Coordinated Multi-Point (CoMP) where the network is divided into disjointed clusters that jointly serve users in their joint coverage area. CoMP has lower implementation complexity compared to a full cell-free network, as synchronization is only required between a cluster of APs in the same geographical area. However, this network-centric approach can result in sub-optimal coverage as compared to a cell-free design. In addition, joint transmission from multiple APs still requires phase-coherent processing, which can be challenging to implement.

A parallel direction for the evolution of wireless networks is the shift towards higher carrier frequencies, such as the Millimeter-Wave (mmWave) band, and even further towards Terahertz (THz) communication. This is motivated by the congestion of the sub-6 GHz bands, contrasted with the large amounts of unused spectrum available at higher frequencies. The large bandwidth of these systems can provide extremely high data rates [9] and low latency, making them key enablers for 5G and future 6G applications. Their unique propagation characteristics introduce both challenges to enable reliable communication, and opportunities to introduce new network paradigms. On the one hand, the increased attenuation, sensitivity to blockage and need for directionality lead to problems with low coverage ranges, unstable links and reduced reliability [22, 23]. On the other hand, the use of narrow directional beams allows for spatial multiplexing by localizing the interference in specific directions. Analog beamforming to generate directional beampatterns can therefore remove the need for fully coherent processing to cope with interference.

Applying in-band multi-connectivity to mmWave networks can be particularly beneficial. In addition to the spectral efficiency benefits, enabling users to simultaneously connect to multiple APs can address the key challenges of resilience and coverage, reducing link breakages. Prior studies demonstrate that CoMP enhances coverage and link reliability in mmWave networks [16, 17]. The impact of interference and blockage in a mmWave CoMP system was also experimentally studied in [12, 18], showing its ability to suppress interference and its robustness against link blockage. A full cell-free design or CoMP with phase-coherent processing, however, can be extremely challenging to practically implement at mmWave frequencies, due to the tight synchronization requirements and need for global channel state information if full interference suppression is desired. The use of narrow directional beams in mmWave introduces the opportunity for a different multi-connectivity design that relies on local spatial separation of signals for interference management. In contrast to solutions that rely on synchronization, this approach can be implemented in a decentralized manner with very low AP coordination. This design can also incorporate the benefits of cell-free networks - as there is no need for tight AP synchronization it is possible to have a user-centric design with no cell division, where users can be served by any AP in the network. In such network implementations, the local beamtraining which selects the analog beams becomes the key performance factor, as it needs to ensure sufficient interference suppression.

The proposed in-band multi-connectivity with local beamtraining system requires additional resources compared to standard single-connectivity designs as all devices need multiple Radio Frequency (RF) chains. These types of architectures are already considered for traditional Multiple-Input and Multiple-Output (MIMO) deployments, with multiple independent streams to a single AP. Due to the sparsity of the mmWave channel, however, successfully establishing multiple independent streams for a MIMO setup can be challenging. Additionally, this approach only aims to increase the link data rate and can not prevent link blockage and loss of service, which are arguably much more relevant issues for mmWave networks where the large bandwidth already allows for high data rates even with a Single-Input and Single-Output (SISO) connection. Therefore, we consider multi-connectivity to a be highly interesting research direction for mmWave with multiple RF chains.

In this paper, we implement and evaluate a multi-connectivity mmWave network to show it can enhance the overall performance, but especially the reliability of mmWave networks. Our goal is to evaluate the benefit of investing more resources in a multi-RF architecture to improve the reliability and stability of mmWave networks. Unlike most previous work that studies multi-connectivity from an information theoretic or signal processing perspective, we look at how such a system would work when deployed with a full protocol stack and what network designs are needed for its functioning. For this purpose we use the network simulator ns-3 and implement a multi-RF, multi-connectivity mmWave network, comparing its performance with a standard small cell mmWave deployment. Our system design focuses on simple and practical implementation, taking advantage of the spatial sharing properties of mmWave to design networks with a very low level of network coordination. To avoid the need for phase-coherent processing and PHY layer cooperation, in our system, all links are considered interference and we rely only on local analog beamforming to perform interference suppression. Thus, each AP decodes the signals from the user without cooperation with the other APs. Furthermore, we also do not coordinate channel access on the Medium Access Control (MAC) layer, allowing for asynchronous, independent transmissions. While performance could be enhanced by using scheduled access to lower interference and collisions, this necessitates building a global schedule that implements spatial sharing, which is complex and costly, especially when considering multi-connectivity. Instead, we chose a simpler approach that still validated the benefits of multi-connectivity for mmWave. Therefore, we base our implementation on mmWave Wi-Fi [14] as Wi-Fi networks represent the perfect examples of dense distributed networks without PHY and channel access coordination.

Our work proves that this network design is viable and can outperform a traditional single-connectivity design. We observe large and consistent gains in reliability, reducing the time spent with no service from up to 20% to 0% in all simulation scenarios we ran, and improvements in the overall network throughput of up to 30%. Interestingly, we observe that mmWave multi-connectivity has a bigger benefit in denser scenarios, where the spatial diversity allows to avoid blockage and constant packet collisions from nearby nodes. Lastly, we study how network design choices such as the beamforming training method affect the performance, exploring different ways to enhance the multi-connectivity implementation.

2 SYSTEM MODEL AND DESIGN

Cell-free and multi-connectivity CoMP architectures exploit spatial diversity by enabling simultaneous service by multiple APs. The key challenge in implementing these designs at mmWave frequencies lies in the tight synchronization requirements for phase-coherent processing. However, by leveraging the spatial multiplexing property of directional communication it is possible to design simpler multi-connectivity implementations which instead rely on analog beamforming for interference suppression. The goal is to use beams that allow for high spatial reuse so that multiple links that are sufficiently spatially separated can be simultaneously active, making the local beamtraining the crucial performance factor. Such an approach also removes the need to design transmission schedules that incorporate spatial sharing, which is complex to do in a dynamic environment where the interference varies not only with the device locations but also with the transmit and receive beampatterns used. In this way, it is possible to design future small cell mmWave deployments to incorporate distributed features from current Wi-Fi.

Therefore, we base our design on mmWave Wi-Fi. We consider indoor scenarios with very dense AP deployment to ensure coverage from multiple APs. All multi-connectivity devices have multiple Phased Antenna Arrays (PAAs) and each PAA is connected to an RF chain. Stations (STAs) use the RF chains to associate and communicate with different APs simultaneously, splitting the data uniformly between the RF chains. To focus on the benefits of multi-connectivity spatial multiplexing only, we consider that APs serve only 1 client per RF chain. APs are only loosely coordinated by a central controller that manages association and packet allocation. This allows us to integrate the user-centric approach of cell-free networks, where clients can connect to any AP in the network, regardless of cell limits. Transmissions occur without any synchronization or cooperation on the MAC or PHY layers, relying only on analog beamforming for interference management. Thus, concurrent packets received at different RF chains on the same device are considered (self-)interference and we rely on local training to achieve good spatial separation to successfully receive them.

Our objective is to explore the viability of in-band mmWave multi-connectivity with local beamforming with a design that minimizes network synchronization for a more practical deployment. In this way, our results represent a lower bound for multi-connectivity, as hybrid beamforming with digital precoding, as well as additional network synchronization or channel access coordination, would improve interference management and result in higher data rates. We compare the multi-RF multi-connectivity system with a standard single-RF implementation where each STA is only associated with one AP at a time. As we are primarily interested in increasing mmWave network resilience, and not in increasing the per-link throughput, we believe this is a better comparison than a MIMO system. Additionally, we study data rates below the maximum capacity of a single mmWave Wi-Fi channel, meaning that there is no need to use MIMO to serve the offered traffic.

Instead, we evaluate the trade-off between the cost of a multi-RF device and the improvements in reliability that can be gained. Fig. 1 shows a comparison of the two implementations. In the single-RF case, STAs have multiple PAAs, but only one is active and connected to the RF chain through which all data gets transmitted. We explore two options on how to select the active PAA: 1) random selection or 2) choosing the PAA that gets coverage from the highest number of APs to maximize the probability of good coverage. In the multi-connectivity, multi-RF case, APs jointly serve STAs within each AP's coverage, enabling STAs to establish multiple directional links with multiple serving APs. The number of directional links is upper-bounded by the number of RF chains of the STA. All APs are connected to a central controller via reliable fronthaul links, enabling AP selection for each STA. In this way, STA benefit from the overlapped coverage of multiple APs.

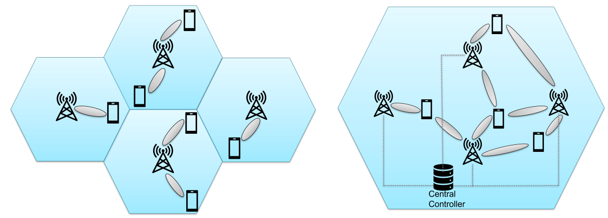

(a) (b)

Figure 1: Comparison of (a) Single-Connectivity and (b) Multi-Connectivity multi-RF implementation

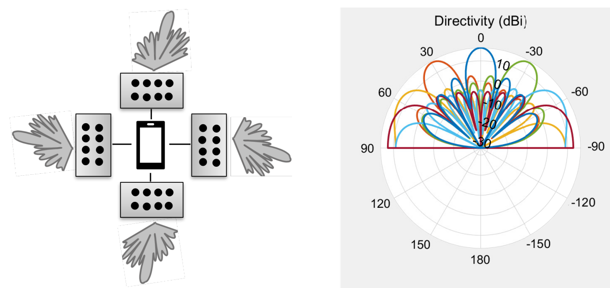

(a) (b)

Figure 2: (a) PAA orientation (b) Directivity of the antenna beampatterns

2.1 Antenna Model

We consider devices with 4 PAAs, each oriented in a different direction, as shown in Fig. 2a. In the multi-connectivity implementation, each PAA is connected to a separate RF chain, while in the single-RF case only one PAA is active at a time. We use a 2x8 element Uniform Rectangular Array (URA) architecture to generate narrow beams to spatially separate signals. Fig. 2b shows examples of the generated directional beampatterns. We notice that the beampatterns are back baffled, meaning that the PAAs do not generate a response on the back side of the antenna. This means that STAs can lose connection if there is no available AP in their field of view. During communication, we use both directional transmission and reception, as they are crucial both for communication and interference management.

2.2 Beamforming Training for Multi-Connectivity Operation

Local analog beamforming training is performed independently by each STA. In the training, devices test different beampatterns from a pre-determined codebook to find the optimal one for communication. We are considering dense networks where each device has multiple PAAs that need to be trained, which can result in very high training overhead. To avoid excessive overhead, we investigated different training methods, focusing on efficiency, scalability, and sensitivity to interference. In the end, we implement a mechanism based on the Group Beamforming principle introduced in IEEE 802.11ay [10], adapting it to our needs. The key change was to make the training user-initiated as users require training with multiple APs. Thus, a user-centric approach can reduce overhead and improve efficiency and accuracy. We train each device PAA individually and do not take into account the inter-PAAs interference within a device when selecting the beams. We note that the performance can be improved by jointly training the PAAs to minimize the interference between them; however, this adds complexity and overhead. Not only is the number of possible combinations of receive and transmit beampatterns too large to do an exhaustive search but AP coordination is needed to measure the joint interference. As the goal of this work is to evaluate a system with minimal coordination we leave joint-PAAs training for future work.

In our modified Group Beamforming, a STA initiates the training by transmitting multiple packets, each with a different directional beampattern. APs that receive the packets measure the Signal-to-Interference-plus-Noise Ratio (SINR) of each packet, determining the quality of each trained beampattern. Thus, each AP can decide which of the trained transmit beampatterns the STA should use to communicate with them. The novel aspect of Group Beamforming is that the STA additionally appends an element called a Training (TRN) field to each training packet. The TRN field has multiple subfields and each subfield is composed of Golay sequences which have good correlation properties. This allows multiple additional beampatterns to be trained in the same packet, by fast beampattern switching between subfields. APs that receive the packets use a different receive beampattern for each subfield. By measuring the SINR of each subfield they can determine the optimal beampattern for reception from the STA. If we consider antenna reciprocity, meaning that transmit and receive beampatterns are equivalent, we can get the full transmit and receive configuration for the STA and all APs. As multiple beampatterns are trained in the same packet, this approach minimizes the training overhead. Moreover, by relying on receive rather than transmit training we can both train many APs at the same time and reduce the training interference since APs do not transmit during the training.

This method results in low overhead and high accuracy. However, in some dense scenarios, it suffers from errors caused by interference, particularly on the STA side. The same behavior was also observed in [10], where a proposed modification that used TRN training for both APs and STAs resolved most errors. Therefore, we also implement and evaluate this modification, to observe how the beamforming training accuracy affects network performance in interference-challenged scenarios.

2.3 Association

We consider that the association is determined by the central controller based on the maximum received power between two devices. Each AP RF chain only serves a single client for even load distribution and clients associate to a different APs with each RF chain. STAs try all APs within their field of view, in descending order of the received power, until they find an available AP. Some STAs (or STA RF chains in the case of multi-RF multi-connectivity devices) will not have a connection because they are in a bad location with no available APs in their field of view. To try to ensure minimal service for each client, priority during the selection is given to the clients with the lowest maximum received power. The same association process is used for the single-RF implementation, with the distinction that STAs only associate to one AP with the single RF chain connected to the active PAA.

3 PERFORMANCE EVALUATION

3.1 Simulation Scenarios

Unlike previous work on cell-free and multi-connectivity which mostly looks at analytical performance bounds simplifying practical protocol or training aspects, we consider a more complete implementation that takes into account aspects like channel access, packet collisions and interactions between APs. For this purpose, we expand the IEEE 802.11ad/ay ns-3 implementation [2] to support multi-connectivity, specifically, our system model as described in Section 2. This involves the creation of STAs that can associate to different APs with each RF chain and independently communicate with them. In addition, we enable STAs to use the Group Beamforming mechanism for our user-initiated beamforming training. Finally, we generate a variety of simulation scenarios to evaluate multi-connectivity under different conditions.

Table 1: Simulations Parameters

Parameter Name	Parameter Value
Transport Protocol	UDP
Aggregation Type	A-MSDU and A-MPDU
A-MSDU Max. Size	7935 Bytes
A-MPDU Max. Size	4 194 303 Bytes
MAC Protocol	CSMA/CA
Number of Codebook Sectors	30 Sectors
Transmit Power	10 dBm
Operating Frequency	60.48 GHz
Simulation Time	50 s [1]

We study indoor environments using the Quasi-Deterministic (Q-D) channel realization software [3]. The default scenario is a rectangular room (7.4 $m \times 13.5\ m \times 3\ m$) and we also study a larger room (29.6 $m \times 54\ m \times 3\ m$). We simulate dense deployments with 10, 20 and 40 APs for sufficient coverage. To isolate the effects of multi-connectivity the number of STAs is reduced to 2, 4 and 8 respectively. Devices are randomly placed in a fixed position and we average the results over 50 simulations with different device locations. We study both downlink and uplink traffic, with varying per-STA data rates of 800 Mbps, 1.6 Gbps, 3.2 Gbps and 6.4 Gbps. Other simulation parameters are listed in Table 1.

3.2 Simulation Results

We analyse multi-connectivity under different levels of network load and interference by varying both the network density and data rate. As discussed in Section 2, we compare the performance to a single-RF design, using two strategies to select the active PAA - randomly or to maximize coverage. Analysing both the individual and aggregate throughput we identified three network states of low, medium and high channel load with similar behaviour and performance trends, discussed below.

3.2.1 Low Network Load. In our environment, an aggregate network load of up to 6.4 Gbps under-saturated the wireless medium, as the maximum capacity of a single 2.16 GHz IEEE 802.11ay wireless channel using the highest single-carrier Modulation and Coding Scheme (MCS)-21 is 8 Gbps. Channel access procedures and packet collisions lower the effective achievable rate, however, traffic up to 6.4 Gbps does not overload the network and causes no congestion.

Fig. 3 shows the Cumulative Distribution Function (CDF) of the per-STA throughput for two such scenarios: a network of 10 APs,

[1]We also tested longer simulation times (2 min, 5 min) and found no changes in the statistics of the performance as the effects we observe happen on the sub-second level.

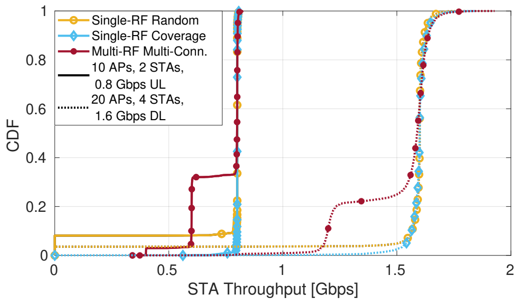

Figure 3: STA Throughput CDF, low network load.

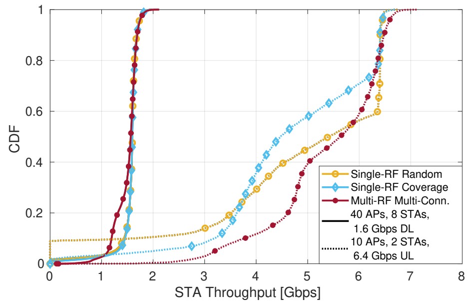

(a) STA throughput CDF

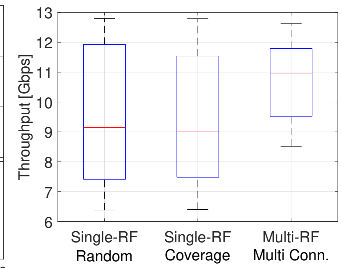

(b) Aggregate network throughput, 10 APs, 2 STAs, 6.4 Gbps UL

Figure 4: Performance in medium network load scenarios.

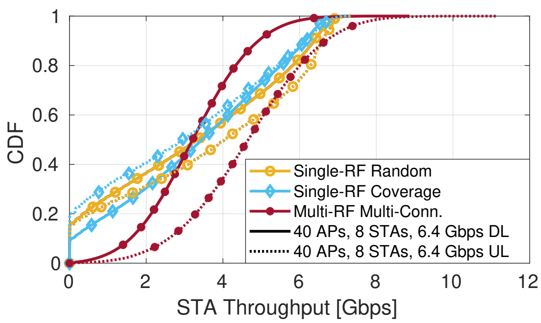

(a) STA throughput CDF

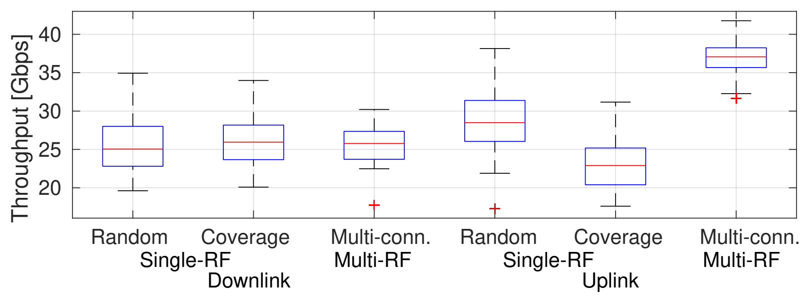

(b) Aggregate network throughput

Figure 5: Performance in high network load scenarios.

2 STAs and a per-STA uplink data rate of 0.8 Gbps and a network of 20 APs, 4 STAs and a per-STA downlink data rate of 1.6 Gbps. We can see identical performance trends for both configurations, also observed in all other scenarios below 6.4 Gbps. First, the targeted data rate is achieved by a large majority of STAs in both the single-RF and multi-RF deployments. However, we also observe an interesting pattern for multi-connectivity, where approximately 20% of STAs only achieve 75% of the targeted rate. This reflects the increased risk that one of the 4 active RF chains is blocked or does not have coverage. Multi-connectivity performance in this case can be improved by redistributing the data from the blocked chain to the other chains, with the added complexity of detecting the loss of service and determining how the data should be distributed. In the single-RF case, this probability is much lower. However, the harm if it happens is much larger - when the RF chain is blocked, the STA completely loses service. We can see that with random PAA selection, this happens approximately 8% of the time in the 2 STA network and 3% in the 4 STA network. In low-load scenarios, selecting the PAA to maximize coverage helped us cope with these issues. Finally, we observe that the traffic direction does not affect the performance. As each AP only serves a single STA, uplink transmissions do not have added channel contention and the performance is equivalent to downlink scenarios.

Lessons learned: When the network is lightly loaded, the main challenge is coverage. In this case, needing coverage for multiple RF chains from different directions can lead to lower achievable data rates as there is simply no AP available in the field of view. However, loss of coverage in the single-RF case is more harmful as it leads to complete loss of service. There are scenarios where multi-connectivity pays off, but in others, it can have disadvantages. This demonstrates the complexity and sensitivity of mmWave networks and why they require tailored networking approaches.

3.2.2 Medium Network Load. From the tested rates, we found that scenarios with an aggregate load of 12.8 Gbps represented medium network load, where the channel is saturated and spatial sharing is necessary to achieve the desired data rate. Therefore, the performance varies based on the network topology, traffic parameters and especially the location of interfering devices. We can see this in Fig. 4a, which shows the throughput CDF in two medium-load scenarios - a network with 40 APs, 8 STAs and a per-STA downlink data rate of 1.6 Gbps and a network with 10 APs, 2 STAs and a per-STA uplink data rate of 6.4 Gbps. We can observe clear differences in multi-connectivity performance in the two scenarios. We found that this data rate represented a transition region where performance depends on the network configuration, with the first scenario (solid lines) representing the worst multi-connectivity performance and the second (dotted lines) the best one. All other tested configurations with an aggregate rate of 12.8 Gbps were between these two scenarios. In the first scenario, the behaviour is similar to the low-load case, as multi-connectivity performance is slightly disturbed by the higher probability of blocked RF chains. However, as the single-RF performance is also slightly degraded by interference, all three deployments have similar performance.

The second scenario is significantly different. Here we begin to see how a multi-connectivity design can improve mmWave networks. By distributing the load among 4 RF chains, the multi-connectivity implementation manages to cope with interference better and especially to ensure uninterrupted service. STAs receive half of the desired data rate only 5% of the time, compared with 10% and 15% in the single-RF deployments. We also highlight how in the single-RF case where the active PAA is randomly chosen, STAs are left without a network connection 9% of the time due to no coverage or high interference. Multi-connectivity, instead, manages to offer a much more uniform service to users, which leads to a gain of over 0.8 Gbps in the average per-STA throughput. Fig. 4b presents the aggregate network throughput for this scenario, showing that not only does the multi-connectivity improve the median aggregate by approximately 2 Gbps, but the gains are achieved by boosting the lowest-performing scenarios. This was a consistent trend we observed, where even when multi-connectivity had lower average performance, it had higher lower performance bounds. This validates both the benefit for increased reliability and the ability to provide more uniform service to all the clients in the network.

Another interesting trend is when comparing single-RF performance with active PAA selection to maximize coverage and randomly. In Fig. 4a we can observe that although selecting the PAA to maximize coverage still reduces the probability that the STA loses coverage, the performance degrades above 3.7 Gbps. Here, the random PAA selection performs much better and STAs can achieve the wanted data rate an additional 16% of the time. We can see in Fig. 4b that the resulting aggregate network throughput is equal in both cases. The reason for this poor performance when choosing the PAA to maximize coverage is that it can inadvertently select a PAA that is transmitting towards a zone with multiple APs and thus high interference in the uplink direction. Therefore, we found that this PAA selection had poor performance in uplink scenarios, an unexpected performance trend that showed the complexity of interaction in dense mmWave networks.

Finally, we found that multi-connectivity was also affected by the traffic direction, with an opposite trend. In this case, downlink traffic caused interference problems since the STA is receiving with all 4 RF chains simultaneously. While analog beamforming proved to be surprisingly robust, high-interference environments were still challenging and performance suffered, as compared to uplink scenarios. To overcome this, it is necessary to perform more complex beamforming training, which attempts to minimize the interference between the RF chains of the STA or to use hybrid beamforming by adding digital precoding.

Lessons learned: At higher network loads, interference begins to affect the performance. In this case, spatially distributing the load across multiple RF chains leads to gains in performance. In addition, it's important to consider issues like inter-RF interference in multi-RF chain systems, or selecting APs from low interference zones for single-RF systems.

3.2.3 High Network Load. The highest network loads of 25.6 Gbps and 51.2 Gbps oversaturated the wireless medium, pushing the network to operate in overloaded mode. In these cases, the desired data rate was extremely difficult to achieve due to high interference. The performance in these scenarios matched the observations of the second scenario with medium network load. There were two main trends, with a difference between the uplink and downlink direction, as seen in Fig. 5which shows the performance in a network with 40 APs, 8 STAs and a per-STA data rate of 6.4 Gbps.

We can observe that in the uplink case, the multi-RF multi-connectivity deployment is able to significantly outperform the single-RF deployments, with the highest gains in the low throughput areas below 2 Gbps, leading to an average aggregate network throughput gain of over 10 Gbps (30%). Crucially, even in the highest-interference scenarios, our multi-connectivity implementation is able to ensure that STAs never lose connectivity and the per-STA throughput is always above 0.5 Gbps. This shows a remarkable improvement in the stability and reliability of mmWave networks, where link breaks and service interruptions are identified as key challenges. This is noticeable in the single-RF implementations, where STAs experience outage up to 15% and 20% of the time. We can also see that in this case, selecting the active PAA to maximize coverage consistently results in suboptimal performance. Due to the oversaturation of the wireless channel, interference is the dominant problem and selecting the active PAA to maximize coverage no longer lowers the probability that the STA will not experience breaks in service.

Fig. 5 also shows downlink performance in high-load scenarios. Multi-connectivity maintains its advantage in the low throughput areas below 3 Gbps, ensuring that STAs can communicate with at least one RF chain at all times. However, the simultaneous reception with all 4 RF chains exposes STAs to too much interference and prevents them from achieving as high throughput as is possible with a single-RF. The resulting aggregate network throughput achieved with the single and multi-RF deployments is approximately the same. As discused above, digital precoding or a combining process from the received signals can help address this issue, although it comes at the cost of increased complexity.

Lessons learned: In high interference scenarios the spatial diversity of multi-connectivity can lead to high gains in throughput. In addition, service outage becomes a crucial issue for performance and multi-connectivity can provide continuous coverage and greatly increase mmWave network resilience.

3.2.4 Large room scenarios. In addition to our default environment, we also study a 4 times larger room. This is a rather interesting scenario because there are two conflicting effects. On one hand, the larger room allows for more spatial sharing since devices are further apart which can improve network performance. On the other hand, since devices are further apart this can lead to lower Signal-to-Noise Ratio (SNR) between a client and its AP, lowering the achievable MCS and consequently user data rate. We found that in different scenarios the trade-off between these two effects varied, leading to better or worse performance as compared to our baseline, smaller room, scenario. In Fig. 6 we show the throughput CDF for exemplary scenarios with low, medium and high network load. A very interesting trend was a clear difference between the multi-connectivity implementation and the single-RF deployments.

In the multi-connectivity case, performance in the small and large rooms is rather similar, as seen in Fig. 6a. When the network load is low, the performance is equal. In some medium and high-load scenarios, the performance is better in the large room due to the increased spatial sharing, as in the medium-load scenario in Fig 6a. However, in most cases, the performance slightly degrades

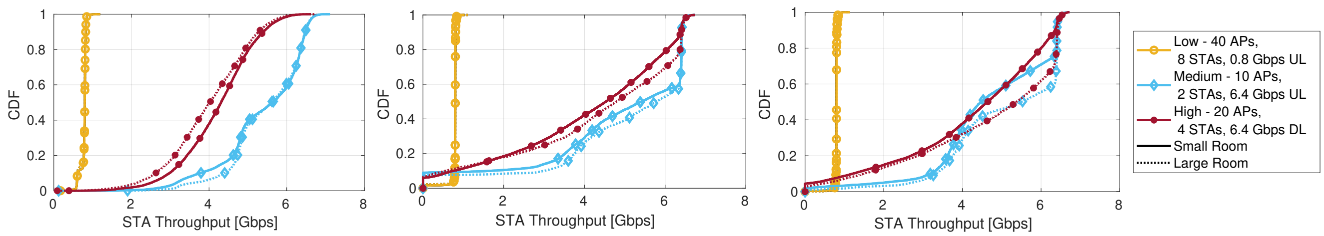

(a) Multi-RF multi-connectivity

(b) Single-RF Random

(c) Single-RF Coverage

Figure 6: Throughput CDF comparison of small and large room scenarios.

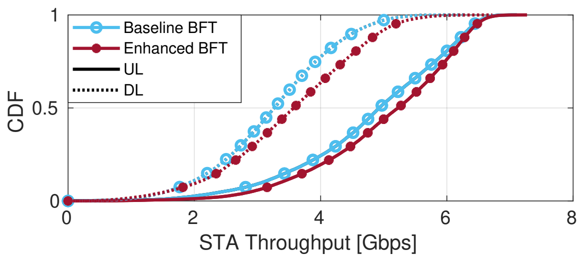

Figure 7: Throughput CDF with enhanced BFT, large room, 40 APs, 8 STAs, 6.4 Gbps.

in the large room, due to the larger distance between STAs and APs, as can be observed in the high-load scenario. This is consistent with previous results that show that the main challenge in our multi-connectivity system is the simultaneous reception at the different RF chains without any interference suppression at the PHY or channel access optimization at higher layers. Since that interference is not significantly affected by the increased spatial sharing, the larger link distances predominately affect the performance negatively. Lastly, these results demonstrate how multi-connectivity architectures offer consistent service in different environments, especially in terms of reliability and connectivity.

In single-RF deployments (Fig. 6b and 6c), however, we observed both that the effect of the room size was much larger and that performance was consistently improved in the large room. In this case, the reduced interference from neighbouring nodes and the increased spatial sharing allowed for larger performance gains, up to 25 %. This was particularly pronounced when the PAA was chosen to maximize coverage. The issues with interference in the downlink were significantly improved, leading to an equal performance in the two single-RF implementations. In addition, since single-RF implementations require coverage for fewer RF chains, the likelihood that one will have to associate to an AP at a large distance is lower than in the multi-RF scenarios. Therefore the larger distances between nodes affected the performance less.

Lessons learned: When considering the deployment size there exists a trade-off between the spatial sharing potential and the average SNR of the links. Multi-connectivity deployments were shown to be less sensitive to environment changes, as the spatial diversity of the links helps ensure consistent and uniform service.

3.2.5 Beamforming training improvements. Once we analyzed multi-connectivity in various scenarios and validated the benefits to reliability and resilience, we further tested various options that may improve the performance. This included trying to improve the carrier sensing thresholds, selecting the beampattern not only based on the measured SNR but also on the achieved throughput to maximize spatial sharing, as well as looking at different Rate Adaptation Algorithm (RAA) mechanisms. Ultimately, however, none of these factors had any major impact on performance. This was due to the fact that the main limitation of our multi-connectivity design came from the interference between the multiple RF chains that had to simultaneously receive data.

The one attempted improvement that had a positive impact was an enhancement to the beamforming training. Previous work has shown that in high-interference scenarios the Group Beamforming principle can lead to wrong beampattern choices if the training is corrupted by interference. We observed the same behaviour in our results, with the majority of beam selection errors on the STA side. Therefore, we used the proposed modification from [10] which uses the TRN approach for training both APs and STAs. We found that this was successful at fixing most errors, and furthermore resulted in gains in throughput. The largest gains were in the densest scenario with 8 STAs and 40 APs, where interference caused the most problems with the beamforming training. In addition, the gains were slightly larger in the larger room scenarios, where selecting the correct beampattern is crucial to overcoming the larger distance between devices. Fig. 7 shows a comparison of multi-connectivity performance in the large room with a per-STA data rate of 6.4 Gbps.

Lessons learned: Beamforming is a crucial factor in our multi-connectivity design, as the beampattern configuration affects both the signal strength and the received interference. Improving the analog beamforming, or using hybrid architectures, particularly to cope with intra-RF interference can greatly improve performance.

4 RELATED WORK

Previous work on multi-connectivity most commonly uses an analytical approach by deriving theoretical performance bounds. In this context, different works demonstrate the benefits of multi-connectivity to mmWave networks reliability [8, 16, 17, 24]. [8] focuses on the degree of multi-connectivity to prevent outage, while in [16] a blockage-aware beamformer design is presented. The ability of multi-connectivity to enhance the robustness against link blockage is experimentally verified in [12, 22]. Not much work has been done, however, on protocol aspects of multi-connectivity networks, including evaluation scenarios that consider MAC and network layer procedures. In addition, these works all consider a phase-coherent processing approach to joint transmission, requiring tight AP synchronization. Similarly, work on cell-free mmWave networks [1, 7, 11, 21] focuses on PHY layer aspects like power control and pilot allocation, as well as the beamforming architecture. The design and performance of hybrid beamforming is studied, as they also use a phase-coherent processing design. Recently, practical aspects like handover [6] and association [13] have been investigated. Our work differs significantly from all of these works, both in the design which relies on analog beamforming for spatial separation and our evaluation methodology which uses a full-stack multi-connectivity implementation in ns-3.

5 CONCLUSIONS

In this paper, we present a network design for multi-connectivity mmWave networks. Our main goal was to show how adding extra resources in terms of RF chains can help solve the key challenges of reliability and resilience. Unlike a standard MIMO approach which only aims to increase the overall spectral efficiency, multi-connectivity designs allow to simultaneously maintain links with multiple APs and thus have increased coverage and resilience to blockage. Moreover, mmWave networks are particularly suited to multi-connectivity designs, as they can be implemented with a much lower level of network synchronization. The use of narrow directional beams allows for interference management through spatial separation of signals, rather than phase-coherent signal processing. We validate this in our design, which is based on the mmWave Wi-Fi protocol and relies solely on analog beamforming for interference management, with no PHY or channel access coordination. Unlike most previous work on multi-connectivity, we implement our design in ns-3 for a performance evaluation with a full protocol stack. Our results show that multi-connectivity maintains extremely high reliability even in high-interference scenarios. In fact, clients could always maintain a connection to at least one AP, contrasted with up to 20% of time spent without service in a standard single-RF deployment. In addition, distributing the load across multiple RF chains also allowed for a gain of up to 30% of aggregate network throughput in dense scenarios. Furthermore, multi-connectivity provides a more uniform service to users, with consistent performance in varying conditions. However, issues with interference caused by simultaneous reception at the multiple RF chains sometimes limited the achievable throughput in the downlink. In addition, finding multiple serving APs for each client sometimes led to issues with coverage and bad connection. Therefore, in future work, we will look at how interference-aware beamtraining, digital precoding and increased network synchronization can help overcome these issues to fully realize the potential of multi-connectivity for mmWave networks.

6 ACKNOWLEDGMENTS

This work has been funded by the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 861222 (MINTS).

REFERENCES

[1] Mario Alonzo, Stefano Buzzi, Alessio Zappone, and Ciro D'Elia. 2019. Energy-Efficient Power Control in Cell-Free and User-Centric Massive MIMO at Millimeter Wave. *IEEE Transactions on Green Communications and Networking* (2019).

[2] Hany Assasa, Nina Grosheva, Tanguy Ropitault, Steve Blandino, Nada Golmie, and Joerg Widmer. 2021. Implementation and evaluation of a WLAN IEEE 802.11ay model in network simulator ns-3. *Workshop on ns-3* (2021), 9–16.

[3] H. Assasa, J. Widmer, T. Ropitault, and N. Golmie. 2019. Enhancing the ns-3 IEEE 802.11ad Model Fidelity: Beam Codebooks, Multi-antenna Beamforming Training, and Quasi-deterministic mmWave Channel. In *Workshop on ns-3*. Florence, Italy.

[4] Emil Björnson and Luca Sanguinetti. 2019. Making Cell-Free Massive MIMO Competitive With MMSE Processing and Centralized Implementation. *IEEE Transactions on Wireless Communications* 19 (2019), 77–90.

[5] Ming Ding, David López-Pérez, Guoqiang Mao, Peng Wang, and Zihuai Lin. 2014. Will the Area Spectral Efficiency Monotonically Grow as Small Cells Go Dense? *2015 IEEE Global Communications Conference (GLOBECOM)* (2014), 1–7.

[6] Carmen D'Andrea, Giovanni Interdonato, and Stefano Buzzi. 2021. User-centric Handover in mmWave Cell-Free Massive MIMO with User Mobility. *2021 29th European Signal Processing Conference (EUSIPCO)* (2021), 1–5.

[7] Guillem Femenias and Felip Riera-Palou. 2019. Cell-Free Millimeter-Wave Massive MIMO Systems With Limited Fronthaul Capacity. *IEEE Access* 7 (2019).

[8] Margarita Gapeyenko, Vitaly Petrov, Dmitri Moltchanov, Mustafa Riza Akdeniz, Sergey D. Andreev, Nageen Himayat, and Yevgeni Koucheryavy. 2019. On the Degree of Multi-Connectivity in 5G Millimeter-Wave Cellular Urban Deployments. *IEEE Transactions on Vehicular Technology* 68 (2019), 1973–1978.

[9] Yasaman Ghasempour, Claudio R. C. M. da Silva, Carlos de M. Cordeiro, and Edward W. Knightly. 2017. IEEE 802.11ay: Next-Generation 60 GHz Communication for 100 Gb/s Wi-Fi. *IEEE Communications Magazine* 55 (2017), 186–192.

[10] Nina Grosheva, Hany Assasa, Tanguy Ropitault, Pablo Jiménez Mateo, Joerg Widmer, and Nada Golmie. 2022. A Comprehensive Analysis and Performance Enhancements for the IEEE 802.11ay Group Beamforming Protocol. *2022 IEEE 23rd International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)* (2022), 194–203.

[11] Ji-Chong Guo, Qiyue Yu, Wenbin Sun, and Wei-Xiao Meng. 2021. Robust Efficient Hybrid Pre-Coding Scheme for mmWave Cell-Free and User-Centric Massive MIMO Communications. *IEEE Transactions on Wireless Communications* (2021).

[12] Rizqi Hersyandika, Juan Sanchez, Yang Miao, Sofie Pollin, and Fredrik Tufvesson. 2022. Measurement-Based Blockage and Intra-Cluster Interference Analysis in MmWave Multi-Point Connectivity Networks *(EmergingWireless '22)*. Association for Computing Machinery, New York, NY, USA, 12–17.

[13] Rizqi Hersyandika, Qing Wang, and Sofie Pollin. 2021. Association in Dense Cell-Free mmWave Networks. *ICC 2021* (2021), 1–6.

[14] IEEE 802.11 working group. 2021. Wireless LAN MAC and PHY Specifications-Amendment 2: Enhanced Throughput for Operation in License-Exempt Bands Above 45 GHz. *IEEE Standard 802.11ay* (2021).

[15] Giovanni Interdonato, Emil Björnson, Hien Quoc Ngo, Pål K. Frenger, and Erik G. Larsson. 2018. Ubiquitous cell-free Massive MIMO communications. *EURASIP Journal on Wireless Communications and Networking* 2019 (2018), 1–13.

[16] Dileep Kumar, Jarkko Kaleva, and Antti Tölli. 2021. Blockage-Aware Reliable mmWave Access via Coordinated Multi-Point Connectivity. *IEEE Transactions on Wireless Communications* 20, 7 (2021), 4238–4252.

[17] Diana Maamari, Natasha Devroye, and Daniela Tuninetti. 2016. Coverage in mmWave Cellular Networks With Base Station Co-Operation. *IEEE Transactions on Wireless Communications* 15, 4 (2016), 2981–2994.

[18] George R. MacCartney and Theodore S. Rappaport. 2019. Millimeter-Wave Base Station Diversity for 5G Coordinated Multipoint (CoMP) Applications. *IEEE Transactions on Wireless Communications* 18, 7 (2019), 3395–3410.

[19] Hien Quoc Ngo, Alexei E. Ashikhmin, Hong Yang, Erik G. Larsson, and Thomas L. Marzetta. 2015. Cell-Free Massive MIMO: Uniformly great service for everyone. *2015 IEEE 16th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)* (2015), 201–205.

[20] Hien Quoc Ngo, Alexei E. Ashikhmin, Hong Yang, Erik G. Larsson, and Thomas L. Marzetta. 2016. Cell-Free Massive MIMO Versus Small Cells. *IEEE Transactions on Wireless Communications* 16 (2016), 1834–1850.

[21] Nhan Thanh Nguyen, Kyungchun Lee, and Huaiyu Dai. 2020. Hybrid Beamforming and Adaptive RF Chain Activation for Uplink Cell-Free Millimeter-Wave Massive MIMO Systems. *IEEE Transactions on Vehicular Technology* 71 (2020).

[22] Theodore Ted S. Rappaport, Shu Sun, Rimma Mayzus, Hang Zhao, Yaniv Azar, Kevin Wang, George N. Wong, Jocelyn K. Schulz, Mathew Samimi, and Felix Gutierrez. 2013. Millimeter Wave Mobile Communications for 5G Cellular: It Will Work! *IEEE Access* 1 (2013), 335–349.

[23] Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, and Parmesh Ramanathan. 2015. 60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling. In *Measurement and Modeling of Computer Systems*.

[24] Huiming Wang, Ke-Wen Huang, and Theodoros A. Tsiftsis. 2018. Base Station Cooperation in Millimeter Wave Cellular Networks: Performance Enhancement of Cell-Edge Users. *IEEE Transactions on Communications* 66 (2018), 5124–5139.