Scaling Millimeter-Wave Networks to Dense Deployments and Dynamic Environments

Claudio Fiandrino, *Member, IEEE*, Hany Assasa, Paolo Casari, *Senior Member, IEEE*, and Joerg Widmer, *Senior Member, IEEE*

***Abstract*—Millimeter wave (mmWave) communications have emerged as one of the most promising options to vastly increase wireless data rates due to the high bandwidth they offer. Given the high path loss at mmWave frequencies, such systems require directional antennas to achieve a good communication range. The communicating devices thus need to align the beam directions of their mmWave antennas. Due to the high penetration loss, the paths between the antennas also need to be free of blocking obstacles. This makes efficient and reliable operation of mmWave networks in dynamic environments very challenging. At the same time, the directionality reduces interference and allows to scale these networks to much higher access point and device densities.**

In this article, we discuss the above challenges and present techniques that allow mmWave networks to scale to high-density deployments, to adapt to dynamic and mobile environments, and to consistently achieve high data rates. This includes learning the environment to find different propagation paths, reacting timely to channel impairments such as blockage, and integrating mmWave networks with networks operating at lower-frequency for robustness. A key ingredient to enable these forms of adaptivity is the use of location information. Such mechanisms then turn a collection of very-high-speed but brittle mmWave links into an efficient, low-latency, and reliable network.

***Index Terms*—Millimeter wave communications, beam training, handover, location systems, 5G mobile networks, IEEE 802.11ad, IEEE 802.11ay.**

I. Introduction

MILLIMETER-wave (mmWave) communications systems have emerged as a key wireless technology for fifth-generation (5G) mobile networks and beyond [1], [2], as well as for high speed wireless local area networks (WLAN) [3], [4]. The limited spectrum available below 6 GHz makes it difficult for current systems to sustain the data rates required to meet traffic demands that are growing at a staggering pace [5]. mmWave bands use frequencies above 10 GHz and, due to the massive available bandwidth at these bands, they can provide orders-of-magnitude higher data rates than lower-frequency systems. However, higher propagation loss, atmospheric absorption, and penetration loss result in shorter coverage ranges and make mmWave communications highly susceptible to blockage. The unique dynamics of mmWave links, i.e., very high data rates combined with low ranges and high variability of the channel, present particular challenges to all the layers of the protocol stack, from the physical (PHY) and medium access control (MAC) layers to higher layers such as transport and even applications.

The use of highly directional mmWave antennas makes it possible to achieve multi-Gbit/s data rates over typical link distances, where beamforming gains compensate for the higher path loss at mmWave carrier frequencies [6]. At the same time, this complicates link establishment and maintenance especially under mobility and in dynamic environments, since the antenna beams of the sender and receiver must be well aligned in order to achieve a sufficient link margin. This alignment, the so-called beam training and tracking, has a significant impact on the performance of control procedures such as the initial network access [7], association schemes [8], medium access [9] and mobility management [10]. Hence, efficient mechanisms that limit the overhead of the beam training are essential [11], especially for dense deployments.

Given the shorter communication ranges, future mmWave networks will have a much higher access point (AP) density than lower-frequency networks, making efficient network management and control essential. To give an illustrating example, while the initial access is performed by signaling over omni-directional channels in current systems like Long Term Evolution (LTE) or 802.11ac, mmWave beamforming requires beam training with potentially many candidate access points, before the most suitable one can be selected.

The commercial potential of mmWave networks has given rise to standardization activities within the IEEE for wireless local/personal area networks (WLAN/WPAN) and within the 3rd Generation Partnership Project (3GPP) for 5G cellular networks. The Wireless Gigabit Alliance (WiGig) and the WiFi Alliance provided the first version of the IEEE 802.11ad amendment in 2012. 802.11ad builds upon prior 802.11 standards, but adapts operation to the specifics of mmWave, such as the integration of beam training for the directional antennas [3]. Studies show that current first generation devices provide stable links up to ranges of 20–30 m [12], which imposes limits on coverage and mobility. The IEEE 802.11ay standard extends 802.11ad to provide even higher throughput, better reliability, and improved range [4], [13]. The 5G New Radio (NR) standardization activities of 3GPP define how to use the new mmWave frequencies [14] for mobile networks in Release 15, TS 38.101-1. Fundamental components of cellular systems are undergoing modifications at all layers of the protocol stack. At the PHY and MAC layers, numerology, channel coding, multiple-input multiple-output (MIMO), initial access and handover are among the ones requiring adaptation [15]. Upper layers receive functionality extensions in the Radio Link Control (RLC), the Packet Data Convergence Protocol (PDCP), and the Radio Resource Control (RRC) [16].

In this paper, we survey mechanisms and techniques for efficient, low latency, and reliable mmWave networks. We outline viable solutions and provide a performance assessment based on results obtained from practical experiments. We

All the authors are with IMDEA Networks Institute, Madrid, Spain. Hany Assasa is also with University Carlos III, Madrid, Spain.
E-mail: {firstname.lastname}@imdea.org.

specifically focus on techniques enabling mmWave systems to adapt to complex and dynamic environments, and to scale to dense network deployments offering very high data rates.

We first discuss mmWave networking challenges with respect to beam training (Section II). Once a link is established, the beam alignment has to be preserved by means of tracking procedures, which is particularly challenging in dynamic and mobile environments. Furthermore, the aforementioned challenges are exacerbated for dense networks, where beam training with potentially many access points is necessary to determine the best one to associate with. The performance of beamforming, training and tracking can be significantly improved with the help of contextual information, and in particular through an in-band localization system able to capture the position and orientation of mobile devices with high accuracy and low overhead (Section III). The mmWave channel is quasi-optical, i.e., the channel consists of few paths that are either line-of-sight (LoS) or non-line-of-sight (NLoS) via low-order reflections, which makes it very amenable to the implementation of wireless location systems.

The high antenna directionality also has an impact on medium access. It reduces interference and improves spatial reuse. However, certain common MAC mechanisms such as carrier sense multiple access with collision avoidance (CSMA/CA) suffer from the incomplete carrier sensing caused by the antenna's directionality, and more efficient medium access is required to scale to high density networks (Section IV). Furthermore, simultaneous use of multiple interfaces at the MAC level, for example a mmWave interface complemented by a sub-6 GHz interface, can improve network reliability and resilience in a way that is transparent to higher layers.

In summary, mmWave networks constitute an extremely interesting solution to the spectrum shortage encountered at lower frequencies, but significant challenges remain in order to realize their full potential.

II. Scalability Challenges of mmWave Networks

A. Beam Training

High-gain directional antennas are essential to overcome the severe propagation issues at mmWave frequencies. In turn, this requires mmWave devices to perform beam training in order to determine suitable directions for transmission and reception. Beam training is the procedure that aligns antenna beams at the transmitter and receiver. It is one of the fundamental mechanisms that mmWave systems employ in order to adapt to changes in the environment. Its efficiency is thus crucial for the overall efficiency of the network itself. For analog beamforming, beam training involves selecting the most suitable beamforming vectors, usually from a predefined codebook [17]. The simplest method for beam training is an exhaustive search over the predefined antenna beams, after which the one that maximizes the signal-to-noise ratio (SNR) over the link is finally chosen to communicate. However, this approach is highly inefficient, and the overhead scales with the number of antenna configurations (and thus the number of antenna elements). To establish a single link with no omni-directional reception and N beams per-antenna, the beam search space is N^2. With omni-directional reception, the complexity is still linear. While such complexity is feasible for current designs, it will become inefficient for future devices with increasingly larger antenna arrays.

The search process can be divided into multiple phases with different antenna beamwidths, reducing the number of beam training steps by narrowing down the search space going from wide beams to successively narrower beams [18], [19]. A simple two-stage version of such hierarchical beam training is used in the 802.11ad standard. Given the sparseness of the mmWave multi-path channel, compressive beam training concepts can also provide a fast solution. The intuition is to reduce the search space by probing only a subset of the N possible patterns [20], [21] (see Fig. 1(a)) and estimating the most likely angle of arrival of the signal by analyzing the sequence of SNR values obtained for the different measurements. The estimated angle then determines the best pattern to be used for transmission and/or reception. Similarly, Agile-Link [22] determines the best alignment in a logarithmic number of measurements by creating so-called multi-armed beams to quickly assess the signal power over multiple spatial directions and estimate the most likely angle of arrival. The flexibility of hybrid and fully digital beamforming enables much better adaptivity, using multi-directional transmission and reception to probe several directions at the same time, so as to reduce the number of search steps even further [11].

B. The Impact of Dynamics

Once a link is established, the communicating devices need to continuously track each other to maintain antenna beams aligned in case of mobility. To this end, beam training can be performed periodically (as is done in 802.11ad) or whenever the link quality degrades below a certain threshold. While the approaches discussed in the previous section work well in conventional indoor settings and with moderate mobility, even the reduced overhead of the more advanced beam training schemes becomes a limiting factor for highly dynamic environments.

Rather than carrying out a full beam training, more sophisticated solutions can reduce such overhead to nearly zero by performing tracking on a per-packet basis [23]. This approach uses beam patterns with two strong adjacent lobes during part of the packet preamble, where the lobes have a 180° phase shift with respect to each other. By detecting the presence or absence of a phase shift during the preamble, it is possible to infer device movement and rotation as part of the packet detection process. The approach provides a much higher degree of adaptivity to dynamic environments.

Since movement is typically continuous, past angle information is also useful to narrow down the beam tracking search space [24]. Other approaches use contextual information. In [25], [26], the 2.4/5.8 GHz bands available in multi-band devices are used to measure the relative angle between two devices, which is then used to bootstrap the beam training process. In addition, the low-frequency interface of a multi-band device provides a highly important communications fallback in case the mmWave link breaks, and can even be used to predict when the link will become available again [27]. Also

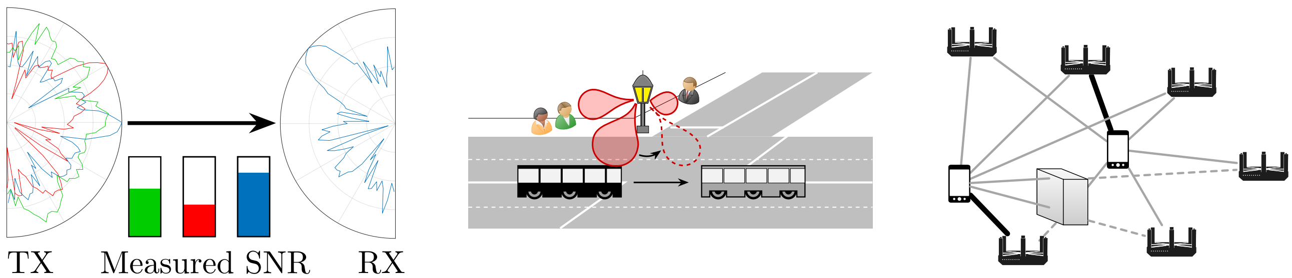

(a) Beam training (b) Beam tracking (c) Access point association

Fig. 1. Scalability challenges of (a) beam training, (b) tracking, and (c) access point selection and association.

the sensors (accelerometers, gyroscopes, and magnetometers) available, for example, in smartphones or vehicles can help determine rotation and movement [28], [29]. Similarly, for vehicular networks, the information provided by cameras, light detection and ranging systems (LIDARs), global positioning system (GPS) receivers, and the like help obtain the relative position of neighboring vehicles or access points, and thus speed up the beam training process [30]. It is even possible to take advantage of the blinking notification LEDs of an access point to track relative changes of the position and orientation of a mobile device, and thus steer the antenna towards the access point [31], [32].

Most importantly, movement is often highly predictable [33]. For example, a vehicle at the intersection shown in Fig. 1(b) can only move straight or turn left, and the speed with which these movements happen are typically bounded within a relatively narrow range. This makes it possible to reliably predict future beam steering angles from past information. Using machine learning, even complex scenarios can be predicted [34], and blockage when the users moves behind obstacles, such as a building, vegetation, vehicles, or other users, can be anticipated to perform a handover prior to link disruption. In summary, with these techniques, mmWave links can be maintained even in highly dynamic scenarios.

C. Access Point Selection and Handover in Dense Networks

Prediction- and learning-based techniques to deal with challenging scenarios can be applied whenever history information is available, for example when only the beam steering angle for an existing link requires updating. Otherwise, some form of beam training must performed.

Beam training is a necessary procedure for any mmWave device performing an initial access to the network. In 5G cellular networks, the initial access process entails three phases to register a user as active: downlink timing and frequency synchronization (cell search procedure), system information acquisition and uplink timing synchronization (random access procedure). IEEE 802.11ad organizes the initial access to the medium by having the access point transmit beacon messages with network information at certain intervals. Access schemes in cellular and WLAN mmWave networks face similar challenges. Due to channel characteristics and the use of directional antennas, there can be a mismatch between the discoverable and supportable coverage area. An omni-directional transmission from a device A may not reach device B, even though B is within the directional transmission range of A, which is particularly problematic for the omni-directional transmission of control messages such as primary and secondary synchronization signals for cell search [35] or request-to-send/clear-to-send messages [36]. Therefore, even control messages are typically sent directionally. In 5G networks, cell search procedures [7] can be based on (i) exhaustive search, (ii) iterative search (a process similar to the one mandated by 802.11ad), and (iii) context-information [37]. After cell search, users undergo a random access procedure consisting of 4 phases: (i) preamble transmission, (ii) response, (iii) connection request, and (iv) resolution. Factors limiting the scalability of random access are the duration and the number of usable receive beams [38]. When using directional communications for random access, the duration of the procedure can be long due to multiple preamble transmissions for each pair of transmit and receive beams.

Similar to initial access, whenever a handover is required there may be many potential access points to which a device may connect. In either case, beam training with just a single candidate access point is suboptimal: the channel quality to the access point may be low, the access point may already be serving many users, or the connection to that specific access point will quickly be lost again due to device mobility. To make an optimal decision, the device should beam-train with all access points in its vicinity but this results in combinatorial complexity for the optimal access point association problem (see Fig. 1(c)), and even for very fast beam training procedures the overhead becomes prohibitive for moderate to high access point and device densities. This is particularly problematic since access point densities for mmWave networks will be significantly higher compared to lower-frequency networks. There are two main reasons for this fact: to provide good coverage, and to improve capacity by reducing the number of devices simultaneously connected to an access point. Device-to-device communications further exacerbate the problem, as potentially all devices within range would need to carry out beam training with one another.

The key to making such scenarios scalable is location information. Even when a link is to be established without any prior history, simply knowing the location of the device and of the nearby access points makes it possible to directly select both the most suitable access point *and* the beam pattern that ensures the highest beamforming gain. In contrast to the combinatorial complexity of exhaustive access point probing, localizing a device typically requires only a few measurements with a small set of access points, and location information is inherently amenable to prediction. While location information is useful to reduce the beam training overhead in general and deal with high environment dynamics, its use is crucial to scale mmWave networks to very high densities.

III. The Importance of Location Information

As has become evident above, accurate location information is essential to improve initial access and mobility management and thus allow mmWave systems to scale and rapidly adapt to changes in the environment. For a versatile, stand-alone mmWave system, such location information needs to be provided fully in-band, without the need for additional devices or sensors such as GPS, accelerometers, gyroscopes, etc. Furthermore, to ensure scalability to high device densities, additional control messages for the purpose of localization should be minimized or avoided altogether, and ideally the system should just use information that is already available to the physical and MAC layers of the mmWave communication stack. Because localization algorithms need to run on mmWave end devices in real time, they need to be low-complexity both for the initialization as well as the run-time phases.

Client localization algorithms specifically designed for mmWave systems are relatively new in the literature. A recent work has shown that the potential performance of mmWave client localization is very high [39]: the packing of massive antenna arrays in a comparatively small space, the potentially high SNR regimes, and the very large bandwidth result in millimeter-level error bounds. For 5G cellular scenarios, this means that both downlink and uplink localization become possible, and that in the presence of a sufficiently multipath- and scattering-rich channel, the error performance of localization schemes is potentially sub-centimeter [40]. For sufficiently large arrays, the work in [41] shows that even random beamforming makes it possible to achieve sub-centimeter accuracy, which tends to improve with accurate beam pattern design. Even with smaller arrays it is still possible to exploit the angular sparsity of mmWave propagation to identify changes in a MIMO beamspace channel matrix, which can be leveraged to localize a mobile client. The main advantage of this approach is that line of sight LoS propagation is not strictly required [42].

While mmWave signals have been used for environment scanning radars [43], [44], such approaches hinge on specialized equipment which is rarely available on mmWave network devices. The results of the client scans also have to be communicated to the access points: this requires protocols deployed for this purpose, and thus is not desirable. Early contributions investigating the potential of simple client-based mmWave localization with a single anchor have achieved decimeter-level accuracy both in simulations and using real 60 GHz equipment in conjunction with the deployment of lab-grade hardware [45]. As a design constraint, the considered algorithms are sufficiently simple to be run on computationally limited devices.

Range-based mmWave localization can usually achieve better performance than purely angle-based algorithms [46]. This is the case of [47], which ranges each mmWave propagation path joining an access point and a client, and derives the location of the client with multilateration. This approach is called pseudo-lateration and requires only a single access point, although the ranging procedure requires the characterization of the environment's path loss model, which in turn imposes a high measurement burden, and needs to be carried out periodically in case the environment itself changes over time. The work in [48] exploits time-of-flight measurements and accurate mmWave pointing to estimate both the range and the angular coordinates of a client with respect to a known access point position, which makes it possible to localize the client with sub-meter errors in about 70% of the cases. The algorithm requires the extraction of channel state information from beam training measurements carried out by the IEEE 802.11ad standard.

Multipath propagation is seen as a resource for both client localization and environment estimation in [49]: the angle measurements obtained by the beam training process make it possible to distill angle-of-departure and/or -arrival information, which is then used to estimate the location of the client even in the presence of the very rough measurements provided by today's commercial mmWave devices. The work in [50] shows that the same input information can be leveraged to jointly estimate the location of the anchor access points and the client *even in the absence of any initial information about the environment and the access point locations*. As the process exploits multipath mmWave propagation, the method also estimates the location of walls and obstacles in a simultaneous localization and mapping (SLAM) fashion. The work in [51] merges angle-of-arrival, received signal strength and time of arrival information to infer both the location of the client and the shape of the environment in a SLAM fashion. The devised system yields good performance in simulations.

A slightly different approach is taken in [52], where the authors adapt the structure of a deep neural network so that it can take as input the phase differences of a sub-6 GHz signal measured at the elements of an antenna array and relate them to a user location, even with a few tens of training samples. The accuracy is sufficient to point a directive beam accurately. An earlier work [25] demonstrated that using angle-of-arrival estimates from low-frequency (2.4/5.8-GHz) antenna arrays to reduce the overhead of mmWave beam training is in fact a promising approach, both for the inference of the optimal steering angles and for the detection of LoS paths, which are subject to lower attenuation than NLoS paths and thus yield higher throughput. More recently, independent measurements carried out in an anechoic chamber confirm that sub-6 GHz angle measurements can inform mmWave beam pointing to a degree that achieves near-ideal range measurements [53].

A location system fully integrated with the mmWave network has a number of advantages: besides supporting location-aided services, such a system can reduce the overhead of all network maintenance and adaptation steps, including beam training, user tracking, handover, load balancing and obstacle avoidance. However, in order to be scalable to high user densities, the system must be independent of other subsystems and eliminate any further interactions between the user and the access points. Ideally, the system should be able to operate independently on each client, without any sort of interaction among the clients (e.g., to reach mutual consensus on the location of one another), and with no information passing between the clients and the access points (except for the necessary standard-compliant beam training algorithms). The largest majority of the literature on mmWave localization does not allow this kind of flexibility and isolation of the location system. In our own work, we propose to implement this architecture both *i*) in an access point-driven "intelligent system" fashion and *ii*) in a fully distributed client-based fashion. In case *i*), the access points measure angles of departure to all clients and infer their location, thereby having sufficient information to drive optimal access point-client associations, beam training procedures, and handovers. The client performs periodic beam training with the single access point to which it is associated, while the remaining access points overhear the beam training messages to estimate angle information. These are then collected to compute location estimates. Thus, case *i*) scales well with the number of access points, but each client needs to perform beam training such that it can be overheard by all access points within range, i.e., only one client should beam-train at a given time. Scalability is therefore limited for very high client densities.

In case *ii*), the localization procedure is fully client-driven, and therefore scales very well with the number of clients. The complexity only increases linearly with the number of (visible) access points. We remark that having more access points in an indoor area improves coverage and localization accuracy, and in any case the client only needs to measure angles of arrival with the visible access points, not with all of them. For this reason, we argue that the additional complexity is acceptable.

In the following, we first present an actual implementation of an access point-driven location system (type *i*) implemented using consumer-grade standard-compliant Talon AD7200 routers, and then discuss the design of a fully-scalable system designed along the lines of type *ii*).

A. Implementing an Access Point-Driven Location System

For narrow antenna beams, beam training and tracking implicitly provide angle of departure (AoD) from the transmitter and angle of arrival (AoA) at the receiver, as each beam covers a well defined spatial angle. Given the location of the access points, locating mmWave clients under such a sector beam model is straightforward. However, commercial off-the-shelf (COTS) devices do not provide such a level of accuracy. The beam shapes of consumer-grade antennas are highly irregular and do not point towards a specific direction, but rather exhibit several lobes that are often equally-strong. This prevents a direct translation of sector identifiers into angle information. Moreover, the quasi-omnidirectional beam patterns employed by current 60 GHz hardware for reception are explicitly designed to avoid the complexity of receive beam training, which further limits the angular information that can be collected at receivers to achieve accurate localization. Finally, the phased antenna arrays of consumer-grade mmWave hardware often provide very coarse signal-to-noise ratio measurements and usually do not provide phase information.

The above constraints call for a different approach to the client localization process: instead of assuming that each sector identifier relates to a given angle, we leverage the knowledge of the device's transmission beam patterns to compute a sparse channel decomposition. This yields the power and angle of departure of each mmWave propagation path. Merging this information from all access points makes it possible to estimate the location of the user. Location errors due to blockage are mitigated by integrating a mechanism similar to dead reckoning in a modified particle filter.

Our method works in real time on commercial off-the-shelf 60-GHz access points with electronically steerable phased antenna arrays, and the only modification applied to the device's firmware is to make the SNR information of each beam training message available to the location system. Notably, the access point's operation itself is unchanged. Below, we provide the characteristics and the main steps of our method, and introduce some results. The full details are available in [49].

Our localization method hinges on the estimation of the angles-of-departure of the signal transmitted by the client. This is achieved indirectly, by processing the power measurements collected independently by each visible access point, for each of the client's transmit beam patterns. The process involves the following two functions: a channel decomposition that reliably estimates angles of departure from the client despite the imperfect beam patterns, and a modified particle filter that merges measurement data with system evolution models in order to estimate the location of the client.

Since we cannot rely on phase information from the access points or access antenna array weights, we cannot infer the mutual interference of different multipath components pertaining to the same signal. Moreover, the coarse dB-scale quantization of the SNR values sensed by the access points impedes the use of non-coherent path estimation approaches such as [54]. Thus, we choose to only constrain the measured amplitude to be less than the sum of the amplitudes of all paths, a required condition for the problem to have a physically representative solution. This leads to the formulation of the following linear estimation problem with variables $\alpha_i(\theta)$:

$$\min \sum_{\theta} \alpha_i(\theta) \left(\sum_{b \in \mathcal{B}_i} p_b(\theta)^2 \right)^{1/2}, \tag{1a}$$

$$\text{s.t.} \sum_{\theta} \alpha_i(\theta) p_b(\theta) \geq P_{R_i}^{(b)}, \quad \forall b \in \mathcal{B}_i, \tag{1b}$$

$$\alpha_i(\theta) \geq 0, \quad \forall \theta \in \Theta, \tag{1c}$$

where $p_b(\theta)$ is one of B beam patterns indexed by $b = 1, \ldots, B$, $\theta \in \Theta$ is the emission angle, $\mathcal{B}_i$ is the set of beam patterns used by access point i for which a measurement is available (which can be a subset of the total beam pattern set due to firmware inefficiencies). The cardinality of set Θ depends on the resolution of the angular domain quantization. Finally, $P_{R_i}^{(b)} = 10^{\gamma_i^{(b)}/20}$ is the signal amplitude corresponding to the SNR $\gamma_i^{(b)}$ (in dB) measured by access point i when the client transmits with beam pattern b.

The minimization formulation of the objective function in (1a) imposes that, for each access point, there exist only a limited number of angles that actually contribute to the received

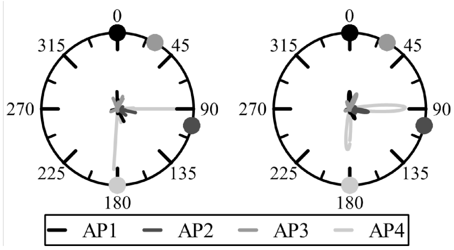

Fig. 2. Angle of departure identification via channel decomposition. Left: linear programming result from Eq. (1). Right: smoothed version for the computation of the angle goodness function. (Adapted from [49].)

power. By constraining the total amplitude of the signal received via all beam patterns along the angles that provide a non-zero contribution to be at least the total received signal amplitude output by the router's hardware (constraint (1b)), we impose that the solution becomes naturally sparse, as there are normally just a few paths that propagate from a client to each access point. Finally, constraint (1c) guarantees that each angle contribution is physically meaningful.

The result of the AoD estimation problem (1a) is shown in the left panel of Fig. 2, where the outer dots represent the actual AoD from each of four access points. The linear program solution includes several likely directions (whose total number is bounded by $|\mathcal{B}_i|$, i.e., by the number of beam patterns used for beam training by the hardware), among which at least one is in fact a good estimate of the LoS AoD. To actually localize the device we exploit the angles estimated via problem (1) to construct an angular goodness function that helps determine the coherence of the measurements with the location and orientation of the device. Such function cannot be non-zero for a discrete set of angle values and zero everywhere else, otherwise we would not be able to compute a soft confidence value in the presence of angle estimation errors. This problem is solved by convolving the variables $\alpha_i(\theta)$ with a smooth Gaussian kernel of predefined standard deviation $\sigma_e = 10°$, namely $g(\theta) = \exp\big(-\theta^2/(2\sigma_e^2)\big)/\sqrt{2\pi\sigma_e^2}$, to yield $v_i(\theta) = \alpha_i(\theta) \otimes g(\theta)$. An angular goodness function $\bar{\mathcal{L}}(\mathbf{x})$ can then be defined by checking how a position and orientation $\mathbf{x}$ fits the smooth angles of departure estimated $v_i(\theta)$ (details in [49]). Function $\bar{\mathcal{L}}(\mathbf{x})$ can be joined with a distance goodness function $D(\mathbf{x})$ which evaluates whether the distance between the access point and the user matches a reasonable propagation model. Note that this can be roughly estimated with a few measurements, and is *not* used for range-based client localization, only to estimate the fitness of a distance value. We finally define an angular and distance fitness function as

$$F(\mathbf{x}) = \bar{\mathcal{L}}(\mathbf{x})\, D(\mathbf{x}), \tag{2}$$

and employ it to check the correctness of a position estimate in light of the collected measurements.

The localization of the user can then be achieved via a real-time-capable filter such as a particle filter, which computes a number of likely user locations using a predefined mobility model, and tests the fitness of each particle using the previously introduced angular and distance goodness function. However, such a filter may diverge quickly in the presence of local fitness maxima, especially in the presence of the inaccurate angle and power measurements provided by real hardware. Therefore, we employ a modified particle filter, where the particle set is enriched by "informed" particles. These particles are generated at the locations obtained by solving angle-difference-of-arrival localization problems [45], when the measurements collected are sufficient to enable this (namely, when at least three angle estimates from three different access points are available for a given measurement epoch). This makes it possible to steer the particle filter to more accurate estimates.

The steps followed by the localization algorithm are summarized as follows:

- Access points collect data from 802.11ad's beam training procedure; this data includes SNR measurements quantized to 0.25 dB for each sector;
- the location system jointly processes the measurements to derive angle-of-departure estimates;
- a particle filter is evolved, and the likelihood of each particle is tested using a likelihood function based on the measurements collected by the access points;
- when angle-of-departure measurements are collected from at least 3 access points, an "informed" particle is fed into the particle filter;
- the locations of the surviving particles are weighed via their respective likelihood and averaged to yield the client's location estimate.

We demonstrate the capabilities of our access point-driven location system using the TP-Link Talon AD7200 router, that has 34 hard-coded antenna configurations which can be selected to steer the antenna array. The corresponding beam patterns were measured in an anechoic chamber to obtain the $p_b(\theta)$ data required for the channel decomposition in Eq. (1a). We consider two different indoor environments: *i*) an unfurnished 11×21 m^2 auditorium (see Fig. 3(a)), which allows to study our system in a blockage- and movement-free setting; and *ii*) a real-world 7.4×13.5 m^2 open office environment that includes typical office furniture, such as screens and metal cabinets that reflect mmWave signals, and obstructing structural elements such as columns. The office is actively being used by employees during the measurements (see Fig. 3(b)). In both scenarios, we deploy the access points to maximize the coverage within the measurement area, and take sector SNR and RSSI measurements at 32 locations arranged in an 8×4 grid.

Figures 3(c) and 3(d) depict our measurement results in the auditorium and office area, respectively, by mapping the color to the median error for all positions and orientations, including some examples of the estimated locations. The system achieves high accuracy in both scenarios, except for some degradation near the walls and corners of the indoor area due to imperfect access point illumination. Localization results in the auditorium are very accurate, as we obtain sub-meter accuracy in 70% of the cases. In the office space sub-meter accuracy is achieved in 60% of the cases.

B. Designing a Client-Based Location System

Considering the limitations of consumer-grade hardware, the above system provides good results. However, future mmWave

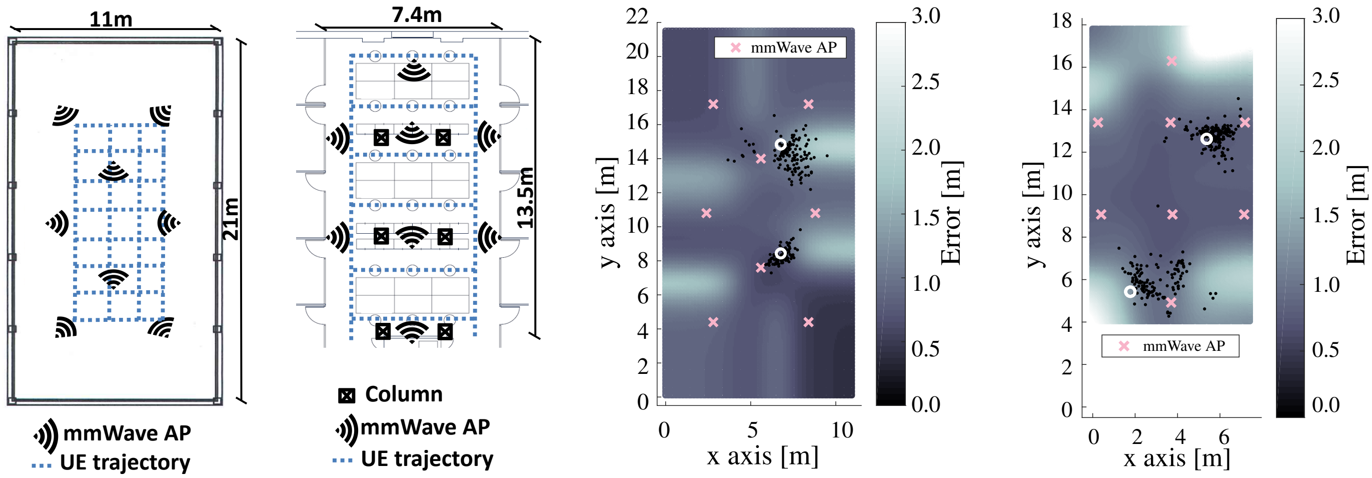

(a) Auditorium scenario (b) Office scenario (c) Auditorium error map (d) Office error map

Fig. 3. Performance of the access point-driven location system. (a)–(b) Measurement setups. (c)–(d) Error maps for each setup. The white circles in panels (c) and (d) represent two measured positions, whereas the black dots show their corresponding location estimates. (Adapted from [49].)

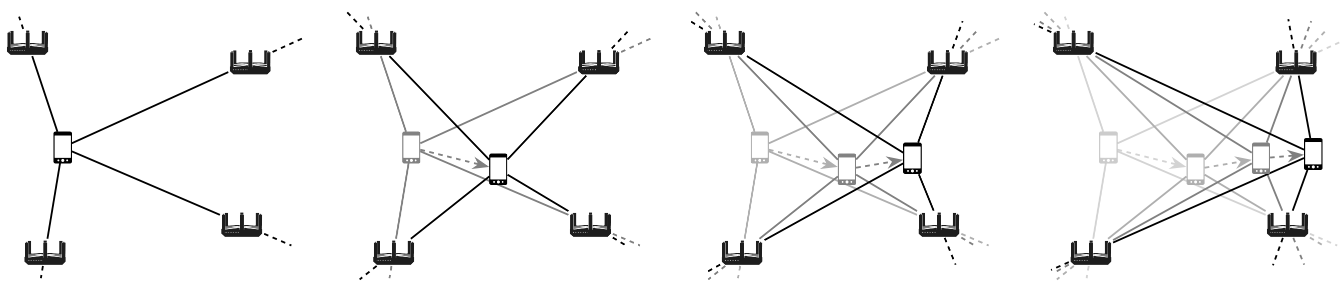

(a) Measurement 1 (b) Measurement 2 (c) Measurement 3 (d) Measurement 4

Fig. 4. Sequence of measurements required by the client-driven location system described in Section III-B. After the fourth measurement, it is possible to estimate both the location of the access points and the coordinates of the measurement points of the user, according to a relative coordinate system.

systems are expected to integrate much more directional antenna arrays [55] and thus can more easily and accurately estimate angle information. This makes it possible to design a location system that relies *only* on angle-difference-of-arrival (ADoA) measurements *carried out locally by each client* in order to estimate the location of a device. Most importantly, the localization process can be carried out with *zero initial information* about the map of the indoor environment, about the initial location of the device, and even about the location of the mmWave access points. All this information can be estimated as the device moves and receives additional ADoA measurements, in a range-free fashion, i.e., without requiring any distance information. The intuition behind this concept is provided in Fig. 4. If a client is fully unaware of the environment and measures, e.g., four angles of arrival from four different access points as in Fig. 4(a), the problem of jointly computing the location of the access points and of the client is under-determined. If we knew the distance traveled by the client between two subsequent measurements we would be able to triangulate the location of the access points and solve the rotation ambiguity with a third measurement (Figs. 4(b) and 4(c)). However, we advocate that the mmWave location system should be independent of other subsystems such as inertial movement sensors or GPS, therefore we have no way to measure the distance covered by the client. To solve the problem we observe that with just three access points, the problem of jointly estimating the location of the access points and of the user is under-determined, as for *any* three access point locations we could define a circumference, position the client at its center and measure the same angles of arrival to those access points. A fourth access point is then needed to "validate" the estimation, by checking that the angle of arrival from this fourth access point is in accordance with the estimated coordinates of all other access points and the client. Having up to 4 measurements with the same scenario as in Fig. 4(d) provides a sufficient number of angle of arrival measurements to eliminate all unknowns of the problem in a relative coordinate system.

Formally, call N_{AP} the number of access points and N_M the number of locations where the client gathers one angle of arrival measurement for each of the access points. Then, at each measurement location, the client takes $N_{AP} - 1$ angle difference of arrival measurements, for a total of $(N_{AP} - 1)N_M$ measurements. We assume relative 2-dimensional localization where we do not need to solve any rotation, translation and scale ambiguities.[1] We have $2N_{AP}+2N_M-4$ unknowns to determine, where the latter 4 unknowns refer to rotation (1 unknown), translation (2) and scale (also 1). We therefore need to have

$$(N_{AP}-1)N_M \geq 2N_{AP}+2N_M-4\,, \tag{3}$$

which can be solved only if $N_{AP} \geq 4$, for which we require at least $N_M = 4$ measurements. We remark that the measurements are collected independently by each device and no prior knowledge or interaction with the access points is required to localize a node, hence the system is inherently scalable and can adapt itself to any environment with sufficient access point illumination.

In the following, we briefly summarize the main steps required to carry out localization and environment estimation using ADoA location estimates. The details of the procedure can be found in [50]. First, we note that ADoAs have the advantage to being invariant to rotation, so that we do not need to trace the orientation of a device with respect to a given coordinate system. Our ADoA-based algorithm processes angle information related to both LoS and NLoS propagation paths emanating from the same mmWave transmission. NLoS components can be modeled as emanating from a "virtual source" corresponding to the mirroring of the physical mmWave source via a reflective surface [56], [57]. When a client localizes itself through the signal of one or more access points, virtual access points help enrich the set of available anchors, so that the client can be located even in areas with limited mmWave illumination.

Let t be the current time, $\mathbf{x}_i$ the location of anchor i, and $\mathbf{y}^{(t)}$ the location of the client at time t. Call $\mathbf{v}_i^{(t)} = 2\left\|\mathbf{y}^{(t)}-\mathbf{x}_i\right\|^{-2}(\mathbf{y}^{(t)}-\mathbf{x}_i)$. Finally, let $\theta_{ij}^{(t)} = \phi_j^{(t)} - \phi_i^{(t)}$ be the ADoA corresponding to signals from anchors i and j (located at $\mathbf{x}_i$ and $\mathbf{x}_j$, respectively), and let $\zeta_{ij}^{(t)} = \pi/2 - \theta_{ij}^{(t)}$. In our algorithm, each client solves the following problem in a distributed fashion:

$$\underset{\{\mathbf{v}_i^{(t)}\},\{\mathbf{x}_i\}}{\arg\min} \sum_{(i,t),(j,t)\in\mathcal{V}} \left(\mathbf{v}_i^{(t)\top}\mathbf{R}_{\zeta_{ij}^{(t)}}(\mathbf{x}_j-\mathbf{x}_i) - 2\sin\theta_{ij}^{(t)}\right)^2, \tag{4}$$

where $\mathbf{R}_\alpha$ is the matrix that rotates a vector in the counter-clockwise direction by an angle α in $\mathbb{R}^2$. In Eq. (4), the set $\mathcal{V}$ contains all pairs (i,t) such that at time t the client can measure visible paths corresponding to the anchor i, and the sum is computed over all i, j, and t such that $(i,t),(j,t)\in\mathcal{V}$. We remark that the solution to this problem is complex and depends on the location of the anchor nodes $\mathbf{x}_i$, $i\in\mathcal{A}$, which is also unknown. We iteratively solve problem (4) by estimating the location of the user and the location of the access points in turns. In particular, it can be shown that the location of the client is the solution to the following minimum mean-square error (MMSE) problem:

$$\mathbf{v}_i^{(t)} = (\mathbf{M}_i^{(t)}\mathbf{M}_i^{(t)\top})^{-1}\mathbf{M}_i^{(t)}\mathbf{b}_i^{(t)\top}\,, \tag{5}$$

[1]We remark that the resolution of these ambiguities is not needed for network optimization operations such as beam training and handovers: in order to know in which direction to point a beam, a device only only needs relative coordinates. In any event, all ambiguities can be fully removed by knowing the absolute location of any two reference points in the space, e.g., the coordinates of two access points.

where $\mathbf{M}_i^{(t)}$ is the $2\times|\mathcal{J}^{(t)}|$ matrix whose columns are

$$[\mathbf{M}_i^{(t)}]_{:,k} = \mathbf{R}_{\zeta_{ik}^{(t)}}(\mathbf{x}_j-\mathbf{x}_i)\,, \tag{6}$$

$\mathbf{b}_i^{(t)} \in \mathbb{R}^{1\times|\mathcal{J}|}$ is the column vector whose kth entry is $[\mathbf{b}_i^{(t)}]_k = 2\sin\theta_{ik}^{(t)}$, for $k\in\mathcal{J}^{(t)}$, and $\mathcal{J}^{(t)}$ is the set of visible anchors for which the client can collect angle measurements at time t.

Subsequently, we solve a second MMSE problem to find the location of the anchors given the location of the user. After some algebra, it can be shown that the location of the access points is the solution to the problem

$$\widehat{\mathbf{x}} = \underset{\mathbf{x}}{\arg\min} \sum_{(i,t),(j,t)\in\mathcal{V}} \left(\mathbf{q}_{ij}^{(t)\top}\mathbf{x} - 2\sin\theta_{ij}^{(t)}\right)^2, \tag{7}$$

where $\mathbf{x} = \left[\mathbf{x}_1^\top\ \mathbf{x}_2^\top\ \cdots\ \mathbf{x}_{|\mathcal{A}|}^\top\right]^\top$ is the anchor location vector. The solutions in (5) and (7) can be computed iteratively until the solution converges, or until a maximum number of iterations has been reached. To initialize the computation, we provide an initial estimate of the anchor locations by rearranging (4) into a form amenable to grid search.

As the measurements collected include physical and virtual anchors alike, multipath propagation can be exploited for environment estimation by pairing each physical access point with its virtual counterparts that model the source of NLoS propagation paths (i.e., reflections). This makes it possible to estimate the location of the environment boundaries, but is also sensitive to angle and access point estimation errors. Since our client-based location system estimates the position both of the access points and of the client, we filter the boundary estimates by requiring that the same boundary point has been seen under different angles and using different physical and virtual access point pairs, before it can be accepted as a feasible estimate.

The main steps of the algorithm are summarized as follows:

- The client measures the angles-of-arrival to the visible access points, and uses them to compute angle-difference-of-arrival values;
- The client initializes and iteratively solves problems (5) and (7) to localize itself and the access points in a relative coordinate system;
- The corresponding information is employed to estimate the location of obstacles and room boundaries via geometrical considerations.

We simulate the performance of our algorithm by using a ray tracer in a complex indoor area with several internal walls that separate a few offices from an open space, including reflective office materials, and featuring a dense access point deployment. Fig. 5(a) shows the cumulative distribution function of the localization error for our algorithm, compared to two range-free, angle-based benchmarks from the literature, one based on triangulation (named TV) and a second one based on the geometrical ADoA algorithm [45]. We remark that both benchmarks require some additional knowledge: for example, TV needs information about the device orientation, and both the TV and ADoA benchmarks require to know the floor map and the location of the access points, in contrast to our algorithm that works with zero initial knowledge, and is thus much more scalable and adaptable. Our algorithm achieves sub-meter accuracy in more than 90% of the cases, 80% of the

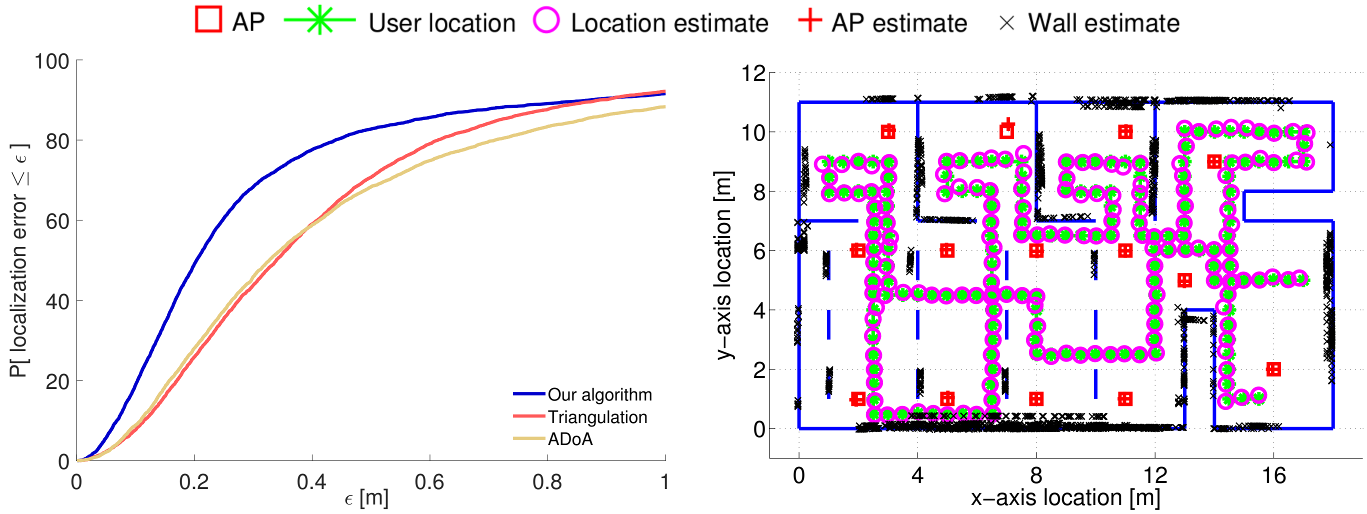

(a) Cumulative distribution function of the localization error of our algorithm, compared to two existing range-free, angle-based approaches.

(b) Simultaneous localization and mapping performance in a complex indoor office area (see key on top).

Fig. 5. Performance of the client-based location system in a complex indoor scenario, compared to two benchmarks making less realistic assumptions on the knowledge of the environment. (Adapted from [50].)

estimates are within less than 40 cm from the true location of the client, and the median error is less than 20 cm. This is in contrast with the two benchmark range-free algorithms, which achieve much worse accuracy, confirming that our algorithm is a feasible approach.

Fig. 5(b) shows the true trajectory (green) of the client that visits several spaces within the office area, superimposed to the estimates found by our algorithm (magenta). We also show the location of the physical access points (red square for the true location and red cross for the estimate). Virtual anchor nodes are not included for clarity but are used by the algorithm to improve the accuracy of the estimates. We observe that in most cases our algorithm achieves an excellent reconstruction of the client's trajectory as well as of the locations of the access points, and that the largest errors are typically within tens of centimeters from the actual locations.

The environment mapping performance of our algorithm is also very good, considering the lack of any initial knowledge about the access point deployment or the floor plan. As the client moves through the area and reflected paths are found, more and more wall sections and mmWave-reflective obstacles such as computer screens are discovered. These elements are marked as black crosses in Fig. 5(b). In some cases, there exists a slight error in the estimation of the anchor locations, especially within the offices in the top section of the area, where the algorithm must often rely on the illumination from a single access point. These access point localization errors translate into slight room boundary estimation errors (e.g., in the left-middle section of the bottom wall). Still, the reconstruction of the client's path and of the environment is very good, especially considering that our algorithm only uses angle information extracted from beam training procedures.

C. Discussion

In conclusion, the results shown above indicate that newer generations of mmWave communication devices with better antenna characteristics in terms of beam pattern directionality and more fine-grained measurements of the complex baseband signals received by the arrays will enable highly accurate and scalable location estimation and environment estimation algorithms. These algorithms can be used for a variety of network adaptation and optimization tasks, and will enable complex location-based applications. Since the systems are based purely on angle information and no coordination among nodes is needed, each node uses its own local coordinate system that is invariant to rotation, translation and scaling. For the network adaptation and optimization tasks discussed below, this is irrelevant, but for location-based services where absolute location is needed, either two reference points must be specified that determine location and scale, or a few SNR or time-of-flight based distances measurements can be used to provide an absolute reference.

As an example, an accurate and distributed location and environment estimation system enables clients to know when obstructions to LoS mmWave propagation paths will occur, so that fast handover schemes can be initiated *before* the client loses connectivity, and without requiring explicit beam training to the new access point (since the relative angle to that access point is known). Moreover, rather than having to probe several candidate access points to find the most suitable one, the client can use its location and that of the access points to directly select the best candidate that provides a good LoS link and at the same time is expected to remain in LoS, given the current trajectory of the client. Also, if an obstruction occurs, e.g., by a human traversing the LoS link, the client can immediately re-associate itself to another access point without any prior beam training. Conversely, if location estimation is not carried out locally by the clients, but rather achieved collaboratively by the access points of the network, the network itself can be managed as a dynamic system. Such system can organize itself to distribute client-access point associations the the objective

(a) Deployment with multiple access points and clients

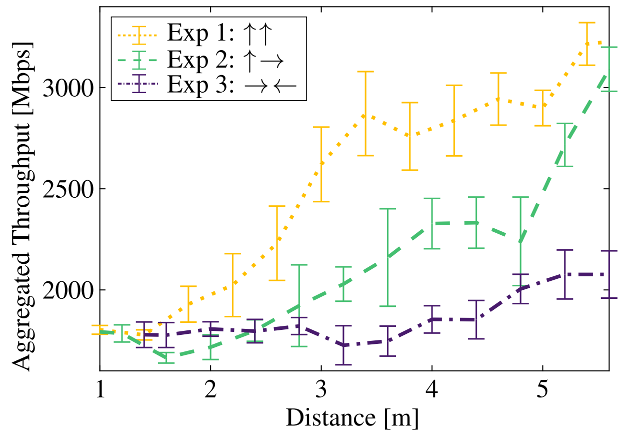

(b) Spatial sharing for different placements of two links.

Fig. 6. CSMA/CA performance for a dense network deployment. (Adapted from [9].)

of maximizing the network throughput and balance the load across the available access points. Moreover, it can preemptively avoid disconnections caused by the mobility of a client or of the environment around it. This leaves such time-consuming operations as beam training and tracking as a fallback solution for the limited number of cases where location information cannot be used (e.g., due to an excessive localization error).

In addition to the above advantages related to mmWave network adaptation, management and optimization, an accurate and scalable location system is a tool of great importance for a number of applications, including personalized healthcare and assisted living [46], customized experiences in specific environments (shopping malls, hospitality, etc.), follow-me services, indoor robotics, as well as more lay applications such as social networking and automated appliance control [58].

IV. MAC Layer Operation

This section discusses the implications of mmWave characteristics on MAC-layer operations, including resource allocation and medium access, and how these operations need to be adapted. Again, location information can significantly help enable efficient designs for large network deployments.

A. Spatial Reuse

In theory, the high directionality of the communication links limits interference, allowing multiple transmissions to take place simultaneously. To this end, both 802.11ad and 802.11ay include mechanisms for spatial sharing and interference mitigation. Access points can measure the interference among individual links in order to simultaneously schedule links of stations that do not interfere.

However, the beam patterns of low-cost commercial devices are far from perfect, and have significant side lobes that may create interference [20]. Thus, analyzing the actual degree of spatial reuse through experiments is of interest. To assess spatial reuse, we consider the following scenario. Two pairs of wireless links (referred to as the “left” and “right” link) operate in the same channel. The placement and the direction of the communication from the two clients to the two access points occurs according to the following topologies: i) parallel vertical links (↑↑), ii) vertical and horizontal links (↑→), and iii) horizontal links (→←). As illustration, in topology (iii), the two clients face each other and the two access points are placed between them, each one facing its own client, and away from the other access point. An example configuration is given in Fig. 6(a). Fig. 6(b) shows the aggregated throughput of both links together with the 95% confidence interval for different values of the distance between each access point and the client connected to it. A single link achieves a throughput of around 1.6 Gbit/s. Hence perfect spatial reuse, where both links operate simultaneously without impairing each other, should yield an aggregate throughput of 3.2 Gbit/s. This is only achieved for parallel links (topology i) and a separation distance of more than 5 m. For all configurations, the aggregated throughput increases with the separation distance. However, the beam steering (and thus the node placement) clearly has a significant impact, and spatial sharing is not feasible for the links with clients facing one another (topology iii) even for very large distances. This suggests that opportunities for spatial reuse critically depend on the beam configurations of the involved devices, and thus on the current scheduling decisions.

With precise location information for access points and devices, as well as knowledge of the available beam patterns of the devices, it is possible to determine potential interference a priori. Therefore, location information can help identify promising candidate links for spatial reuse, rather than having to try all link combinations to identify which ones can operate simultaneously. Furthermore, knowledge of the device locations allows to design custom beam patterns with a strong main lobe towards the communication partner and very low gain towards interferers. The higher the device and access point densities become in future networks, the larger the importance of efficiently detecting spatial reuse opportunities and of using beam patterns specifically designed to enhance spatial reuse.

B. Taking Advantage of Multiple Bands

The problems of blockage and limited range of mmWave networks require efficient recovery mechanisms. For this reason, in both WLAN and 5G cellular networks there exist methods for fast seamless handover and joint use of mmWave and low-frequency interfaces. To take advantage of multiple bands, IEEE 802.11ad includes a fast session transfer (FST) technique. An IEEE 802.11ad device can seamlessly change its operational band from 60 GHz to 2.4/5.8 GHz in a manner that is transparent to higher layers, i.e., while maintaining the same interface and MAC address. As a result, a device can extend its coverage area without interrupting the currently active flows. For the cellular context, 3GPP Release 12 introduced the concept of dual-connectivity, where a device can be simultaneously connected to two base stations, one master base station and a secondary base station. The master handles both the control plane and the user plane, whereas the secondary typically handles only the user plane. For 5G systems, the concept of dual-connectivity evolved into multi-connectivity where a device can be simultaneously attached to an LTE base station and a next generation 5G base station operating at mmWave frequencies. The 3GPP Release 15 provides the basis for inter-networking among LTE and 5G new radio (NR) [59], allowing user-plane traffic to be split or duplicated at the aggregation point [60].

The objective of the above mechanisms is to enhance adaptability and utility of the system by combining the high data rates offered by the mmWave bands with the reliability and resilience offered by sub-6 GHz bands. At the same time, the joint use of multiple bands brings significant challenges to the network. Signals may penetrate obstacles that are present in a given environment, be (partially) reflected, attenuated, or completely blocked, depending on the frequency of the signals and the material and thickness of the obstacle. If the problem of optimal access point selection is already combinatorial by considering only multiple access points at mmWave frequencies, this problem becomes even harder for multi-band systems. Theoretically it would be necessary to assess the channel quality to each access point in each band, resulting in a prohibitively high overhead.

Also for such multi-band systems, accurate location information is of immense help, as it allows the network to determine where obstacles are located. A simple approximation (that holds most of the times) is to assume that obstacles only attenuate low frequency signals, whereas they block mmWave signals. Given the predicted trajectory of a client, a location-aided multi-band system can determine when a client will not see any mmWave access points due to obstacles, and therefore would have to switch to a low frequency alternative. Furthermore, whenever this does not hold (i.e., an access point is assumed to be visible but in reality it is not, or vice versa), this information can be recorded for the given location and be taken into account in future decisions.

In summary, location information is not only essential for beam training and access point association, but also for scheduling, spatial reuse, and deciding when to fall back to lower frequency systems to avoid outage. Similar optimizations may be feasible using machine learning based on large volumes of past channel measurements instead of explicit location information. However, due to the quasi-optical mmWave channel characteristics, location information is an excellent performance predictor for mmWave, requires significantly fewer measurements than machine learning to work reliably, and gives more useful information for network performance troubleshooting.

V. Concluding Remarks

Millimeter-wave communications systems offer multi-Gbit/s data rates and limited interference thanks to highly directional antennas. However, mmWave channel properties such as high propagation loss and blockage require highly efficient network management and control algorithms to achieve a good overall *network-level* performance.

Improvements in the antenna arrays will further increase data rates and improve spatial reuse, but may increase the beam training and tracking overhead. As a consequence, the use of context information (and in particular location information) becomes essential, as it makes it possible to significantly reduce the overhead of such procedures. For this purpose, we present efficient in-band mmWave location systems that will allow future mmWave networks to adapt more rapidly to changing environment conditions and scale to high device densities. The high antenna directionality also brings about challenges for the medium access control. We discuss these aspects through practical measurements and simulations, and highlight promising solutions to improve performance and adaptivity.

Acknowledgment

This work was supported by the ERC project SEARCHLIGHT (grant no. 617721), the Spanish Ramon y Cajal grant RYC-2012-10788, and the Madrid Regional Government through the TIGRE5-CM program (S2013/ICE-2919).

The authors would like to thank Guillermo Bielsa, Danilo de Donno, Pablo Jiménez, Adrian Loch, and Joan Palacios for their contributions to the research that forms the basis of this article.

References

[1] S. Rangan, T. S. Rappaport, and E. Erkip, “Millimeter-wave cellular wireless networks: Potentials and challenges,” *Proceedings of the IEEE*, vol. 102, no. 3, pp. 366–385, Mar 2014.

[2] T. S. Rappaport, S. Sun, R. Mayzus *et al.*, “Millimeter wave mobile communications for 5G cellular: It will work!” *IEEE Access*, vol. 1, pp. 335–349, May 2013.

[3] T. Nitsche, C. Cordeiro, A. B. Flores *et al.*, “IEEE 802.11ad: directional 60 GHz communication for multi-gigabit-per-second Wi-Fi,” *IEEE Communications Magazine*, vol. 52, no. 12, pp. 132–141, Dec 2014.

[4] Y. Ghasempour, C. R. C. M. da Silva, C. Cordeiro *et al.*, “IEEE 802.11ay: Next-generation 60 GHz communication for 100 Gb/s Wi-Fi,” *IEEE Communications Magazine*, vol. 55, no. 12, pp. 186–192, Dec 2017.

[5] M. Shafi, A. F. Molisch, P. J. Smith *et al.*, “5G: A tutorial overview of standards, trials, challenges, deployment, and practice,” *IEEE JSAC*, vol. 35, no. 6, pp. 1201–1221, Jun 2017.

[6] W. Roh, J. Y. Seol, J. Park *et al.*, “Millimeter-wave beamforming as an enabling technology for 5G cellular communications: theoretical feasibility and prototype results,” *IEEE Communications Magazine*, vol. 52, no. 2, pp. 106–113, Feb 2014.

[7] M. Giordani, M. Mezzavilla, and M. Zorzi, “Initial access in 5G mmwave cellular networks,” *IEEE Communications Magazine*, vol. 54, no. 11, pp. 40–47, November 2016.

[8] S. Goyal, M. Mezzavilla, S. Rangan *et al.*, “User association in 5G mmwave networks,” in *Proc. IEEE WCNC*, Mar 2017, pp. 1–6.

[9] H. Assasa, S. K. Saha, A. Loch *et al.*, "Medium access and transport protocol aspects in practical 802.11ad networks," in *Proc. IEEE WoWMoM*, June 2018, pp. 1–9.

[10] M. Polese, M. Giordani, M. Mezzavilla *et al.*, "Improved handover through dual connectivity in 5G mmwave mobile networks," *IEEE JSAC*, vol. 35, no. 9, pp. 2069–2084, Sep 2017.

[11] D. de Donno, J. Palacios, and J. Widmer, "Millimeter-wave beam training acceleration through low-complexity hybrid transceivers," *IEEE Trans. on Wireless Communications*, vol. 16, no. 6, pp. 3646–3660, Jun 2017.

[12] S. K. Saha, V. V. Vira, A. Garg *et al.*, "A feasibility study of 60 GHz indoor WLANs," in *Proc. ICCCN*, Aug 2016, pp. 1–9.

[13] C. R. C. M. da Silva, J. Kosloff, C. Chen *et al.*, "Beamforming training for IEEE 802.11ay millimeter wave systems," in *Proc. Inf. Theory and Appl. Workshop*, Jan 2018, pp. 1–6.

[14] 3GPP, "NR - User Equipment (UE) radio transmission and reception," *Group Radio Access Network*, vol. Technical Specification 38.101-1, Release 15, pp. 1 – 53, Dec 2017.

[15] H. Shokri-Ghadikolaei, C. Fischione, G. Fodor *et al.*, "Millimeter wave cellular networks: A MAC layer perspective," *IEEE Trans. on Communications*, vol. 63, no. 10, pp. 3437–3458, Oct 2015.

[16] P. Rugeland, M. Tercero, Y. Li *et al.*, "Architectural enablers and concepts for mm-wave RAN integration," *mmMAGIC White Paper*, pp. 1 – 26, Mar 2017.

[17] A. Alkhateeb, Y. H. Nam, M. S. Rahman *et al.*, "Initial beam association in millimeter wave cellular systems: Analysis and design insights," *IEEE Trans. on Wireless Communications*, vol. 16, no. 5, pp. 2807–2821, May 2017.

[18] J. Wang, Z. Lan, C. W. Pyo *et al.*, "Beam codebook based beamforming protocol for multi-Gbps millimeter-wave WPAN systems," *IEEE JSAC*, vol. 27, no. 8, pp. 1390–1399, Oct 2009.

[19] B. Li, Z. Zhou, H. Zhang *et al.*, "Efficient beamforming training for 60-GHz millimeter-wave communications: A novel numerical optimization framework," *IEEE Trans. on Vehicular Technology*, vol. 63, no. 2, pp. 703–717, Feb 2014.

[20] D. Steinmetzer, D. Wegemer, M. Schulz *et al.*, "Compressive millimeter-wave sector selection in off-the-shelf IEEE 802.11ad devices," in *Proc. ACM CoNEXT*, 2017, pp. 414–425.

[21] S. Sur, I. Pefkianakis, X. Zhang *et al.*, "Towards scalable and ubiquitous millimeter-wave wireless networks," in *Proc. ACM MobiCom*, New Delhi, India, Oct. 2018.

[22] H. Hassanieh, O. Abari, M. Rodriguez *et al.*, "Fast millimeter wave beam alignment," in *Proc. ACM SIGCOMM*, 2018, pp. 432–445.

[23] A. Loch, H. Assasa, J. Palacios *et al.*, "Zero overhead device tracking in 60 GHz wireless networks using multi-lobe beam patterns," in *Proc. ACM CoNEXT*, 2017, pp. 224–237.

[24] A. Patra, L. Simić, and M. Petrova, "Experimental evaluation of a novel fast beamsteering algorithm for link re-establishment in mm-wave indoor WLANs," in *Proc. IEEE PIMRC*, Sep 2016, pp. 1–7.

[25] T. Nitsche, A. B. Flores, E. W. Knightly *et al.*, "Steering with eyes closed: Mm-Wave beam steering without in-band measurement," in *Proc. IEEE INFOCOM*, April 2015, pp. 2416–2424.

[26] A. Ali, N. González-Prelcic, and R. W. Heath, "Millimeter wave beam-selection using out-of-band spatial information," *IEEE Transactions on Wireless Communications*, vol. 17, no. 2, pp. 1038–1052, Feb 2018.

[27] S. Sur, I. Pefkianakis, X. Zhang *et al.*, "WiFi-assisted 60 GHz wireless networks," in *Proc. ACM MobiCom*, 2017, pp. 28–41.

[28] T. Wei and X. Zhang, "Pose information assisted 60 GHz networks: Towards seamless coverage and mobility support," in *Proc. ACM MobiCom*, 2017, pp. 42–55.

[29] Z. Qi and W. Liu, "Three-dimensional millimetre wave beam tracking based on handset MEMS sensors with extended kalman filtering," in *Radio Propagation and Technologies for 5G*, Oct 2016, pp. 1–6.

[30] J. Choi, V. Va, N. Gonzalez-Prelcic *et al.*, "Millimeter-wave vehicular communication to support massive automotive sensing," *IEEE Communications Magazine*, vol. 54, no. 12, pp. 160–167, Dec 2016.

[31] M. K. Haider, Y. Ghasempour, and E. W. Knightly, "SearchLight: Tracking device mobility using indoor luminaries to adapt 60 GHz beams," in *Proc ACM Mobihoc*, 2018, pp. 181–190.

[32] M. K. Haider, Y. Ghasempour, D. Koutsonikolas *et al.*, "LiSteer: Mmwave beam acquisition and steering by tracking indicator LEDs on wireless APs," in *Proc. ACM MobiCom*, New Delhi, India, Oct. 2018.

[33] S. Sur, X. Zhang, P. Ramanathan *et al.*, "BeamSpy: Enabling robust 60 GHz links under blockage," in *Proc. USENIX NSDI*, 2016, pp. 193–206.

[34] G. H. Sim, S. Klos, A. Asadi *et al.*, "An online context-aware machine learning algorithm for 5G mmWave vehicular communications," *IEEE/ACM Transactions on Networking*, pp. 1–14, Sep 2018.

[35] H. Shokri-Ghadikolaei, C. Fischione, G. Fodor *et al.*, "Millimeter wave cellular networks: A MAC layer perspective," *IEEE Trans. on Communications*, vol. 63, no. 10, pp. 3437–3458, Oct 2015.

[36] A. Akhtar and S. C. Ergen, "Directional MAC protocol for IEEE 802.11ad based wireless local area networks," *Ad Hoc Networks*, vol. 69, pp. 49 – 64, 2018.

[37] I. Filippini, V. Sciancalepore, F. Devoti *et al.*, "Fast cell discovery in mm-wave 5G networks with context information," *IEEE Trans. on Mobile Computing*, vol. 17, no. 7, pp. 1538–1552, Jul 2018.

[38] C. Jeong, J. Park, and H. Yu, "Random access in millimeter-wave beamforming cellular networks: issues and approaches," *IEEE Communications Magazine*, vol. 53, no. 1, pp. 180–185, Jan 2015.

[39] D. Wang, M. Fattouche, and X. Zhan, "Pursuance of mm-level accuracy: Ranging and positioning in mmwave systems," *IEEE Syst. J.*, 2018, online-first.

[40] Z. Abu-Shaban, X. Zhou, T. Abhayapala *et al.*, "Performance of location and orientation estimation in 5G mmWave systems: Uplink vs downlink," in *Proc. IEEE WCNC*, Barcelona, Spain, Apr. 2018.

[41] A. Guerra, F. Guidi, and D. Dardari, "On the impact of beamforming strategy on mm-wave localization performance limits," in *Proc. IEEE ICC Workshops*, Paris, France, May 2017.

[42] H. Deng and A. Sayeed, "Mm-wave MIMO channel modeling and user localization using sparse beamspace signatures," in *Proc. IEEE SPAWC*, Toronto, Canada, Jun. 2014.

[43] S. Hantscher, M. Hägelen, S. Lang *et al.*, "Localisation of concealed worn items using a millimeter wave FMCW radar," in *Proc. Asia-Pacif. µwave Conf.*, Melbourne, Australia, Dec. 2011.

[44] A. Guerra, F. Guidi, A. Clemente *et al.*, "Application of transmitarray antennas for indoor mapping at millimeter-waves," in *Proc. EuCNC*, Paris, France, June 2015.

[45] A. Olivier, G. Bielsa, I. Tejado *et al.*, "Lightweight indoor localization for 60-GHz millimeter wave systems," in *Proc. IEEE SECON*, Jun 2016.

[46] K. Witrisal, P. Meissner, E. Leitinger *et al.*, "High-accuracy localization for assisted living: 5G systems will turn multipath channels from foe to friend," *IEEE Signal Processing Magazine*, vol. 33, no. 2, pp. 59–70, Mar. 2016.

[47] J. Chen, D. Steinmetzer, J. Classen *et al.*, "Pseudo lateration: Millimeter-wave localization using a single RF chain," in *Proc. IEEE WCNC*, San Francisco, CA, Mar. 2017.

[48] I. Pefkianakis and K.-H. Kim, "Accurate 3d localization for 60 ghz networks," in *Proc. ACM SenSys*, Shenzhen, China, Nov. 2018.

[49] G. Bielsa, J. Palacios, A. Loch *et al.*, "Communication-driven localization and mapping for millimeter wave networks," in *Proc. IEEE INFOCOM*, Apr 2018.

[50] J. Palacios, P. Casari, and J. Widmer, "JADE: Zero-knowledge device localization and environment mapping for millimeter wave systems," in *Proc. IEEE INFOCOM*, May 2017.

[51] A. Yassin, Y. Nasser, A. Y. Al-Dubai *et al.*, "MOSAIC: Simultaneous localization and environment mapping using mmwave without a-priori knowledge," *IEEE Access*, vol. 6, pp. 68 932–68 947, 2018.

[52] M. Z. Comiter, M. B. Crouse, H. T. Kung *et al.*, "A structured deep neural network for data-driven localization in high frequency wireless networks," *AIRCC Intl. J. Comp. Networks & Comm.*, vol. 9, no. 3, pp. 21–39, May 2017.

[53] N. Maletic, V. Sark, J. Gutiérrez *et al.*, "Device localization using mmwave ranging with sub-6-assisted angle of arrival estimation," in *Proc. IEEE BMSB*, Jun. 2018.

[54] M. E. Rasekh, Z. Marzi, Y. Zhu *et al.*, "Noncoherent mmWave path tracking," in *Proc. Hotmobile*, Feb. 2017.

[55] M. Xiao, S. Mumtaz, Y. Huang *et al.*, "Millimeter wave communications for future mobile networks," *IEEE J. Sel. Areas Commun.*, vol. 35, no. 9, pp. 1909–1935, Sep. 2017.

[56] P. Meissner, C. Steiner, and K. Witrisal, "UWB positioning with virtual anchors and floor plan information," in *Proc. WPNC*, Mar. 2010.

[57] C. Zhang, F. Li, J. Luo *et al.*, "iLocScan: Harnessing multipath for simultaneous indoor source localization and space scanning," in *Proc. ACM SenSys*, Nov. 2014.

[58] L. Barkhuus and A. K. Dey, "Location-based services for mobile telephony: a study of users' privacy concerns," in *Proc. IFIP Interact*, Sep. 2003.

[59] J. Sachs, G. Wikstrom, T. Dudda *et al.*, "5G radio network design for ultra-reliable low-latency communication," *IEEE Network*, vol. 32, no. 2, pp. 24–31, Mar 2018.

[60] J. Rao and S. Vrzic, "Packet duplication for URLLC in 5G: Architectural enhancements and performance analysis," *IEEE Network*, vol. 32, no. 2, pp. 32–40, Mar 2018.

Claudio Fiandrino (S'14–M'17) is a postdoctoral researcher at IMDEA Networks Institute, Madrid, Spain. He joined the Wireless Networking Group (WNG) in December 2016 after having obtained his Ph.D. degree at the University of Luxembourg in November 2016. He received the Bachelor Degree in Ingegneria Telematica in 2010 and the Master Degree in Computer and Communication Networks Engineering in 2012 both from Politecnico di Torino. Claudio also holds the 2016 SmartICT Certificate on standardization for business innovation from the joint program of University of Luxembourg and ILNAS, the National Standardization Agency. He has been awarded the Spanish "Juan de la Cierva" grant and the Best Paper Awards in IEEE Cloudnet 2016 and in ACM WiNTECH 2018. He served as Publication and Web Chair in IEEE CloudNet 2014, Publicity chair in ACM/IEEE ANCS, and as TPC member in several IEEE and ACM conferences and workshops. His primary research interests include transport layer performance in mmWave networks, mobile edge/fog computing and mobile crowdsensing.

Hany Assasa is a Ph.D. student at IMDEA Networks Institute, Madrid, Spain. He joined the Wireless Networking Group (WNG) in February 2015 after having obtained a double degree master in Information and Communication Technologies (ICT) from Politecnico di Torino (Italy) and KTH Royal Institute of Technology (Sweden). In his Ph.D studies, Hany focuses on building efficient, robust, and reliable millimeter-wave wireless networks. Hany spent 5 months at NEC Laboratories Europe, Heidelberg, Germany where he worked in the 5G group on the topic of multi-RAT mmWave access. In addition, Hany spent one year at Ericsson Research AB (Sweden) in the Core and Architecture Unit where he worked on the topic of Service Mobility in Mobile Networks.

Paolo Casari (S'05–M'08–SM'13) received the PhD in Information Engineering in 2008 from the University of Padova, Italy. He was on leave at the Massachusetts Institute of Technology in 2007, working on underwater communications and networks. He collaborated to several funded projects including CLAM (FP7), RACUN (EDA), and several US ARO, ONR and NSF initiatives. He is the PI of the NATO SPS project ThreatDetect, and the scientific coordinator of the EU H2020 RECAP and SYMBIOSIS projects. Previously, he was Technical Manager of the NAUTILUS and WISEWAI projects. In 2015, he joined the IMDEA Networks Institute, Madrid, Spain, where he leads the Ubiquitous Wireless Networks group. He served in the organizing committee of several international conferences. He serves on the editorial boards of the IEEE TRANSACTIONS ON MOBILE COMPUTING and of the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS, and has been guest editor of a special issue of IEEE ACCESS on "Underwater Acoustic Communications and Networking," as well as of an earlier special issue of the *Hindawi Journal of Electrical and Computer Engineering* on the same topic. He received two best paper awards. His research interests include many aspects of networked communications, such as channel modeling, network protocol design, localization, simulation, and experimental evaluations.

Joerg Widmer (M'06–SM'10) is Research Professor as well as Research Director of IMDEA Networks in Madrid, Spain. His research focuses on wireless networks, ranging from extremely high frequency millimeter-wave communication and MAC layer design to mobile network architectures. From 2005 to 2010, he was manager of the Ubiquitous Networking Research Group at DOCOMO Euro-Labs in Munich, Germany, leading several projects in the area of mobile and cellular networks. Before, he worked as post-doctoral researcher at EPFL, Switzerland on ultra-wide band communication and network coding. He was a visiting researcher at the International Computer Science Institute in Berkeley, USA, University College London, UK, and TU Darmstadt, Germany. Joerg Widmer authored more than 150 conference and journal papers and three IETF RFCs, and holds 13 patents. He serves or served on the editorial board of the IEEE TRANSACTIONS ON MOBILE COMPUTING, the IEEE TRANSACTIONS ON COMMUNICATIONS, of *Elsevier Computer Networks* and collaborated with the program committees of several major conferences. He was awarded an ERC consolidator grant, the Friedrich Wilhelm Bessel Research Award of the Alexander von Humboldt Foundation, a Mercator Fellowship of the German Research Foundation, a Spanish Ramon y Cajal grant, as well as eight best paper awards.