Future Millimetre-Wave Indoor Systems: A Blueprint for Joint Communication & Sensing

Mohammed Alloulah, Howard Huang

Abstract

Millimetre-wave (mm-wave) technology is emerging as a de facto enabler for next-generation high-rate communications, following recent increases in available spectrum up to a mighty 14GHz. The bountiful spectrum brings new opportunities indoors for passive high-resolution sensing of human motion and gestures using only mm-wave infrastructure and without additional sensors. We propose that a unified mm-wave system for combined communication and robust sensing will turbocharge the capabilities of application domains from in-home digital health to new possibilities for building analytics. For the design of such a dual-function network, we consider the system-level challenges of radio resource management, processing complexity, and inference. In particular, we propose a modified radio frame with a periodic transmission of a dedicated sensing signal. Simulation results show that centimetre-level sensing can be achieved using less than 5% of the radio resources, providing evidence that mm-wave communication networks with value-added sensing could be achieved with low cost.

I. Introduction

To meet the growing demands of wireless capacity, millimeter wave (mm-wave) spectrum in the 30GHz to 300GHz carrier frequency range will be used for future 5G cellular networks and local area networks based on the IEEE 802.11ad standard.

The higher frequency range, well-above the single-digit GHz frequencies used for previous generations, presents opportunities in terms of significantly higher data rates due to vast bandwidth resources. In 802.11ad networks, 2GHz bandwidth channels in the 60GHz frequency range will enable multi-gigabit per second data rates [1], and similar data rates and bandwidth allocations are expected in 5G systems [2].

One of the main challenges of mm-wave spectrum is higher atmospheric absorption that results in higher propagation losses. Using antenna arrays with tens or even hundreds of elements, beamforming can focus radio energy to improve the link budget. Because the element separation is on the order of fraction of a wavelength, dozens of antenna elements can fit on an array the size of a coin. Despite the beamforming gain, the propagation losses are so severe that communication range is typically limited to a few hundred meters. Hence for ubiquitous communication coverage, mm-wave systems are well suited for indoor networks where it is cost-effective to deploy infrastructure at a sufficiently high density.

In addition to communications, the radio spectrum can also be used for characterizing aspects of the physical environment via passive radar sensing. The principle is shown in Figure 1a, where a single infrastructure transceiver emits a signal that impinges on an object, and the receiver processes the reflected signal to characterize motion or estimate the range to the object using well-known signal processing techniques for radar [3]. These techniques can be implemented using sub-6GHz communication systems such as WiFi [4], e.g. for emotion recognition and vital signs [5]. In addition to radar processing, the channel state information (CSI) measurements from WiFi reflections can enable new applications for device-free human sensing [6]. In our proposal, we focus on mm-wave *radar* sensing, which we will often refer to as “sensing” (without the modifier).

In sensing applications, the vast mm-wave bandwidth resources enable finer spatial resolution so that even more subtle movements can be characterized. In addition, directional beamforming in mm-wave systems potentially allows multiple targets to be sensed simultaneously by multiple transceivers.

Mm-wave radar outdoors enables vehicles to more robustly sense objects under night or poor weather conditions when compared to lidar or cameras. Radar’s advantage is such that level 5 autonomous driving may almost wholly be achieved by a dense 4D radar scanning objects in range, speed, azimuthal and elevation angles. Indoors however, apart from isolated experiments, e.g. [7], mm-wave sensing remains largely uninvestigated, despite the potential advantages of higher resolution, multi-targeting.

Returning to Figure 1a, we recognise the opportunity for using a single mm-wave access point to achieve both high-rate communications and high-resolution sensing. Communication occurs between the access point and a mobile mm-wave transceiver carried or worn by a person, and passive radar sensing can provide insights on the motion and actions of the person. In general, a network of access points could be deployed to enable greater coverage and capabilities.

In this article, we explore the fundamental challenges of implementing a mm-wave system for joint communications and sensing. We outline challenges stemming from the spectrum characteristics, consider how the two functions interact, and propose a blueprint for an underlying radio system design. Given the value of passive, scalable, high-resolution sensing, we hope to show these benefits can be achieved with low marginal cost by leveraging the vast spectrum and common hardware resources of a mm-wave access point. We motivate the benefits of a joint communications-sensing mm-wave system with two use cases:

(a) **Digital Health.** An 802.11ad access point mounted on the ceiling of home environments or eldercare facilities could

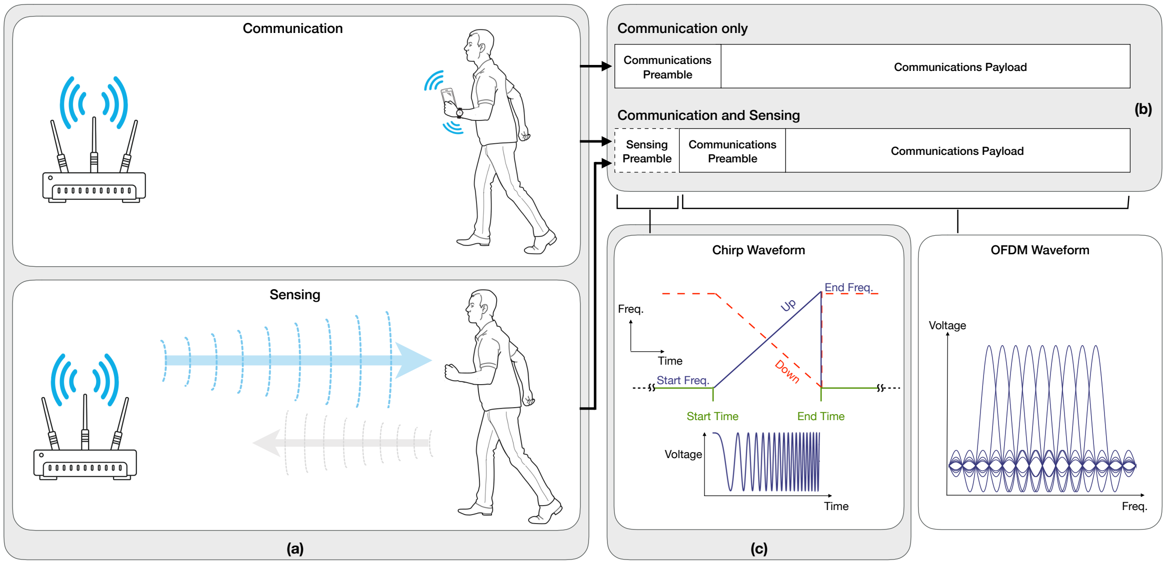

Fig. 1: (a) Communication requires two transceivers whereas radar sensing uses a single infrastructure transceiver. Under sensing, the router emits a signal that reflects off an object back to the router enabling it to characterize motion or estimate the range to the object. (b) Adding a sensing preamble to a classic communication frame. (c) Multiple access schemes for radar signalling.

be used to track motion and actions for multiple patients in order to support behavioural analytics at unprecedented granularities. Such analytics can then be fused with many other medication adherence technologies in order to tackle the US $290 billion annual costs of non-adherence to medications [8]. In similar vein and more broadly, non-invasive eldercare can also be enabled, complementing and ultimately precluding the need for wearable sensors. Alternatively, sensing could be performed *scalably* by personal robotic assistants equipped with 5G mm-wave modems. A key feature of 5G modems is ultra-reliable, low-latency communication which would be useful for controlling robots during critical emergencies.

(b) **Building Analytics.** An important class of current applications, such as residential security and building management, are especially suited for a network-centric, value-added sensing. This is because network-centric sensing substantially simplifies their implementation and long-term maintenance and evolution, with obvious implications on roll-out cost and coverage. However, state-of-the-art communication-only network sensing has proved hard to generalise with the guaranteed performance demanded by such real-world applications [6], [9]. This suggests that next generation networks should be built from the ground up with sensing in mind in order to ship ready to unlock future economical benefits that may not be fully conceived yet.

In Section II, we discuss the key characteristics of mm-wave systems and how they affect communication and sensing performance. Section III proposes a system architecture and considers the challenges in practical implementation. We conclude in Section IV with proposals for next steps in addressing the key challenges. While our discussion is for the most part relevant for general mm-wave systems, we will often focus on the IEEE 802.11ad system as an illustrative example.

II. A Joint Communication and Sensing Network

Our main thesis is: *transitioning to mm-wave frequencies will open up the possibility for architecting a wireless network that has robust sensing functions alongside its traditional data-carrier mandate*. That is, *joint* provisions at the transmitter, receiver, and protocol resources will allow the network to perceive its environment as well as transport data at much enhanced scales compared to conventional network architectures. Such network evolution will be made possible, in part, by a number of unique characteristics that mm-wave frequencies possess. We next review these characteristics and comment on their appealing implications on communication and sensing.

A. Mm-wave characteristics for communications

The characteristics of mm-wave are often discussed in the context of communication systems. These characteristics are:
Path-loss. Mm-wave frequencies experience high free space path-loss and are less able to diffract around objects compared to traditional lower RF frequencies. Propagation range indoors is also in part limited by the high oxygen absorption in the unlicensed 60GHz band indoors. Such propagation phenomena make mm-wave quasi-optical in that focused beams between nodes are required in order to sustain high throughput communications.
Directionality. Owing to the high pathloss of mm-wave links, beamforming at the transmitter and/or receiver are required in order to boost signal-to-noise ratio (SNR). Such directional communication gives rise to an inherent and effective spatial multiplexing of the wireless medium. Spatial multiplexing denotes the ability to have co-located transceiver nodes potentially sharing the exact spectrum but being separated in space using quasi-optical beams that are pointed in different directions.
Wavelength. In the unlicensed 60GHz band, the wavelength is about 5 millimetre. Such fine wavelength allows for packing

dense antenna arrays in compact form factors. Dense arrays will in turn vastly enhance the ability to electronically steer communication and sensing beams in order to maximise the scanning coverage indoors without requiring impractically large antenna sizes.

Bandwidth. Transitioning to mm-wave bands opens up vast swathes of spectrum—up to 14GHz according to FCC. Protocols such as 802.11ad allocate larger 2.16GHz channels per co-located device. When coupled with spatial multiplexing, these large channels allow for unprecedented *scalable* data rates irrespective of how many devices co-exist in an environment.

B. Mm-wave for radar sensing

The implications of mm-wave characteristics on indoor radar sensing are comparatively less highlighted in literature. We thus provide commentary on (i) the appeal of mm-wave for human sensing in relation to three degrees of freedom: range, speed, and orientation, and (ii) a high-resolution imaging form by which human sensing can be realised atop mm-wave's capabilities.

Fundamental limits. The sensing range resolution is given by $c/2B$, where c is the speed of light and B is the utilised bandwidth. Range resolution is inversely proportional to bandwidth. For example, we expect range resolution enhancements in 802.11ad in excess of 12x that of sub-6GHz by virtue of available bandwidth. In common system incarnations, the Doppler resolution is given by $\lambda/2T$, where λ is the wavelength and T is how long we choose to analyse the Doppler effect of the target of interest. Doppler resolution is proportional to the wavelength. Thus we would also expect velocity estimation enhancements in 11ad in excess of 10x compared to that of sub-6GHz. The angular resolution is also proportional to wavelength i.e. $\propto \lambda/D$, where (again λ is the wavelength) and D is the array size. That is, for two antenna arrays of the same size, the angular resolution at 60GHz is greater than 12x that at 5GHz. Consequently, we can realise denser antenna arrays in compact form factors in 11ad which would allow us to more accurately estimate the orientation of objects indoors.

Environmental imaging. State-of-the-art radar techniques allow for the simultaneous estimation of (i) range to potential targets, (ii) the speed of these targets, and (ii) their orientation, all with respect to the probing radar. The result of this comprehensive sensing process is an *image* of the environment. Further processing on top of radar imaging enables the inference and tracking of environmental dynamics to varying complexities.

Having established the radar sensing and data-carrier superiority of the indoor 60GHz band, we next *sketch a blueprint for harnessing mm-wave's advantages in the context of a future mm-wave system incarnation.*

III. Radio Proposal & Challenges

Millimeter-wave communication systems typically use OFDM waveforms because of their high spectral efficiency and robustness in wireless channels. On the other hand, mm-wave radar systems typically use chirp waveforms because the receiver can be implemented with relatively low cost, which could be difficult to achieve otherwise given the high frequencies and large bandwidths in mm-wave systems. *For our joint sensing and communications system, we propose simply to time multiplex a chirp-based sensing transmission with an OFDM-based communication transmission, as shown in Figure 1b.* A naive implementation would have separate realizations of the two radio systems, but cost savings could be achieved by taking advantage of common components.

To further reduce radio cost, one could imagine a system that uses a single waveform used for both sensing and communications. For example, the OFDM signal could be used for sensing as well. However, OFDM signals suffer from poor radar sensing performance because of their sensitivity to carrier-frequency offset using typical oscillators [10]. Conversely, if chirp-like signals were to be used for communications, they would have very low spectral efficiency, which is undesirable. While waveform design for integrated sensing-communication systems is a topic of current research interest, the performance of either sensing subsystem or communication is compromised. Thus it is our opinion that our proposed time-multiplexed system reaches a better compromise with respect to the overall complexity-performance tradeoff. Furthermore, we note that in addition to radar sensing, the OFDM-based communications system could leverage existing CSI-based techniques for device-free human sensing or other techniques for indirect, *low-fidelity* sensing [11].

To elaborate, CSI-based low-fidelity sensing struggles to generalise across end-users and environments without extensive training that attempts to extract domain-independent features [9]. For example, using activity data collected from as many as 22 pairs of *different* user-environment experiments, Jiang et al. report an achievable level of accuracy just over 70%. This level of accuracy falls short of supporting the current and emerging application requirements we outlined in Section I. In contrast, radar-based sensing measures directly the physical primitives necessary to construct well-behaved models which allow researchers to claim capabilities such as emotion recognition [5]. However, unlike sub-6GHz, mm-wave can enable multiple such capabilities to operate simultaneously owing to the much increased spectrum. As such, we would expect a step change in application sophistication once the mm-wave network is equipped with a radar sensing mode.

In what follows, we provide additional justification for our time-multiplexed proposal and address the key challenges regarding the radio stack. Perhaps the most pressing challenges are those pertaining to (i) radio resource management for sensing and communication functions, (ii) efficient processing, and (iii) inference and decision-making.

Quantitative results in this section have been generated using system-level simulations of indoor human movement traces. Twenty scenarios split equally between walking and running were imported from [12]. We divide the channel into timeslots such that we time-multiplex sensing and communications. Continuous, back-to-back transmission is assumed. To investigate sensing granularity effects within the context of human indoor movement, we sweep how many timeslots are allocated for sensing and how many are left of communications. The two extreme configurations correspond to sensing

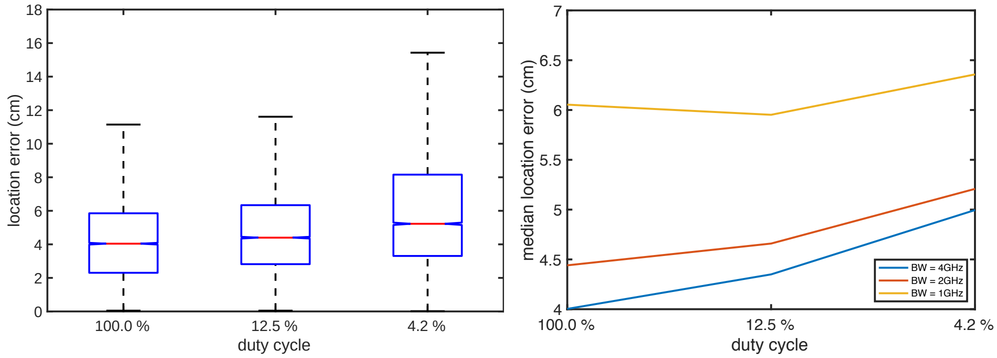

(a) Location error as a function of sensing duty cycle for indoor human movement using 4GHz bandwidth.

(b) Effect of bandwidth utilisation on median location error across sensing duty cycles for indoor human movement.

Fig. 2: Localisation performance of 60GHz radar sensing with 8-element array.

all the time or as little as 4% of the time. When tracking a moving person under a reduced sensing duty cycle, we interpolate between two sensing intervals in order to update our dynamic scene estimates. We assume a transmitter with 4mW power, 10dB gain, and 80μs chirp duration. In the interest of increased coverage, we also assume a *ubiquitous radar* architecture in which a separate transmit antenna is utilised and whose beam lobe is wide [13], or quasi-omnidirectional [14], [15].

A. Radio resource management

We now discuss ways by which a sensing function could be made to work efficiently within a communication protocol.

Resource partitioning. A sensing preamble can be prepended to a communication frame as depicted in Figure 1b. Such sensing preamble supplies “physical environment” information that would benefit both high-level stack applications in need of sensing primitives and the communication protocol itself. Mm-wave communication systems require extensive physical searches in order to deliver high data rates owing to the unique characteristics discussed earlier. However, it is important to make sure that the benefits this sensing preamble may have are not offset by a gross reduction in the spectral efficiency of the overall system. The spectral efficiency is defined as the percentage of time the protocol spends in actual data transfer versus other setup and configuration overhead. We ask: *Are we able to add a radar scan interval to a communication protocol with an acceptable overall system trade-off.*

Results are shown in Figure 2a. It is evident that a drastic reduction in the sensing duty cycle within an mm-wave frame is not accompanied with a drastic penalty on estimation fidelity. Specifically, the difference between the median errors of the two extreme cases of 100% and 4.2% sensing duty cycle is a mere 1cm—although other error quantiles show a slightly steeper divergence. Nonetheless, this experiment suggests that powerful native radar sensing of the physical medium is able to deliver a nuanced tracking of environmental dynamics at a minimal reduction in communication spectral efficiency. Additionally, such environmental tracking may in fact complement current mm-wave communication systems and help them sustain their theoretical performance under mobility scenarios. It has been shown in prior art that 11ad is especially susceptible to human mobility-induced channel stresses [1]. Bielsa et al. show that under mobility scenarios, TCP throughput can be nullified because the protocol struggles to maintain a beam pointed in the direction of the mobile node [1]. Native radar sensing can *acquire* and *track* the trajectory of mobile devices in order to ensure that directional communications is maintained at all times—hence potentially doing without lengthy beam setup preambles.

Channelisation. The process of channelisation refers to the management of spectrum access amongst a group of contending users. In terms of managing channelisation for sensing, it is perhaps most straightforward to restrict sensing to the channel bandwidth already allotted by the communication protocol. As such, co-located devices become frequency multiplexed which allows sensing to seamlessly inherit communication configurations such as transmission power limits. However, there are different multiple access provisions that sensing can utilise should the need arise for extended bandwidth. For instance, an access point (AP) may provide a “gesture” interface for an end user say for configuration purposes or within the context of virtual reality applications. For enhanced gesture recognition, an AP may utilise an extended bandwidth e.g. 7GHz [13]. To minimise signalling collisions with co-located nodes that periodically conduct sensing preambles of their own, chirps can be further coded in: (i) frequency sweep direction (up and down), (ii) start and end timings that form a sweep interval, (iii) and start and end frequencies that form a frequency “ramp.” Depicted in Figure 1c, these mechanisms provide many degrees of freedom to randomise the sensing medium access thereby minimising interference.

In Figure 2b, we repeat our duty cycle simulations for three bandwidth options: 2GHz in keeping within the 11ad channel size, 4GHz doubling range resolution, and 1GHz halving range resolution. It can be seen that for macro human motion indoors, doubling the bandwidth has limited effect on location error. Specifically, 5mm median accuracy enhancement is gained by doubling the radar bandwidth for 100% sensing duty cycle. This tiny margin reduces further at low duty cycles owning to interpolation being the dominant source of error. For

1GHz bandwidth, sensing duty cycle has less effect because of the reduced range resolution.

We hope to have informed the discussion on the potential upside of integrating native sensing functions in future 60GHz indoor networks from, at least, a communication spectral efficiency and bandwidth perspectives.

B. *Efficient radar processing*

Integrating sensing functions in a future 11ad incarnation has also implications on memory and compute resources. Specifically, workloads associated with radar processing require careful considerations within the context of an 11ad commodity AP. To examine such considerations, let us consider a uniform linear array (ULA), whose radar data is arranged in a cube. The cube dimensions corresponds to target distances, target velocities, and target orientations i.e. respectively range, Doppler, and antenna as shown in Figure 3a. The target detection problem reduces to 3D FFT followed by peak analysis in 3D for the simultaneous estimation of range, Doppler, and direction-of-arrival (DOA) [3]. While building intuition gradually, we touch on three avenues which we believe would facilitate processing complexity reduction, thereby making radar workloads more amenable to integration in a commodity AP.

Sequential detection. Sequential detection refers to performing a target parameter estimation using a sequence of *dependent* processing steps, as opposed to a computationally intensive "blanket" processing that would also waste compute resources on irrelevant noise i.e. no targets. We can make assumptions on the sparsity of the sensing scene as opposed to considering all possible combinations of range, Doppler and DOA spectral bins. This sparsity allows for significant reduction in computational complexity. That is, we only proceed to analyse the Doppler properties of range peaks. Similarly, we only run DOA estimation algorithms for Doppler and range peaks. As illustrated in Figure 3a, this flavour of detection is thus *sequential* starting with range peaks, then Doppler peaks, then angle estimates.

Doppler-centric targets. Certain indoor applications aim to only supply human-centric sensing such as crowd counting. A unique property of human targets indoors is their Doppler content in case of movement and/or due to breathing vital signs. This unique Doppler property of human targets can then be used to gate unnecessary processing of inanimate objects, resulting in further scope for processing optimisations. Similar human signatures arising from the mixing of multi-body parts Doppler content—e.g. from swinging arms and a more rigid torso—have been shown to allow for pedestrian discrimination in short range automotive radar [16].

Scene stationarity. Assumptions on the stationarity of the sensing scene can also be made in order to yield potentially large scope for optimisation. Specifically, human movement is much slower than even a duty-cycled mm-wave sensing interval. Therefore, the sensing scene undergoes slow variations in time. After the initial acquisition of range, Doppler, and DOA bins, the subsequent frequency grid of an adjacent radar sensing interval will necessarily evolve from its predecessor. As such, we may be able to efficiently search for updated target bins deterministically by projecting human targets forward in time using their current range, speed, and orientation estimates. Sparse frequency estimation techniques have been shown to be applicable to 3D radar applications in [14]. It would be interesting if such formulation can be made to work recursively as outlined.

We remark that using best-in-class silicon design and integration practices—from analogue, digital, to algorithms—aggressive miniaturisation suitable for a wearable form factor and power consumption has been demonstrated in the 60GHz band in prior art [13].

C. *ML-aided inference*

The last component for adding sensing functions to a future 11ad system incarnation is inference. Inference is needed in order to convert primitive measurements—i.e. a person's range, speed, and orientation—into activity classification. That is, inference is a local intelligence and decision-making component necessary to endow robotic eldercare assistants with autonomy. In adjacent domains, machine learning (ML) techniques have recently demonstrated commendable ability to capture *latent* featurisations beyond what expert designers can model and devise. One particularly relevant adjacent domain that has benefited from ML and bears close resemblance to radar is computer vision; both radar and vision assign intensity scores to 3D space voxels. Specifically, deep learning ML techniques represent the state-of-the-art in visual classification tasks. By extension, we would expect ML to afford the problem domain of radar sensing equal revolutionising advantages. Indeed, using an ML model that captures the nuances of multi-body part reflections, recent pioneering results from indoor RF radar have set a new bar for device-free indoor localisation [17]. Zhao et al. developed an ML model that tracks a torso centre of mass surrounded by lesser body parts which jointly exhibit a certain kinematic relationship in space and across time.

However, there are complexities that surround hand-crafting generalisable models for mm-wave radar. As an example of such complexities, we next consider atmospheric attenuation across the 14GHz spectrum allotted for 60GHz indoor use. We aim to expose yet another vector of variation intrinsic to the radar process itself in order to posit the need for ML.

Non-uniform attenuation. Figure 3b depicts the atmospheric attenuation across the indoor 60GHz band evaluated at three distances: 1, 10, and 100 metres. We note that the attenuation across the 14GHz band is non-uniform. Specifically, beyond 63GHz, we enter a signal regime in which higher frequencies have effectively increased SNR compared to lower frequencies. Such frequency-dependent attenuation also varies in range. While these variations are by no means severe, we would expect such behaviour to interact with the radar sensing model in subtle ways resulting in measurement artefacts. That is, the channel-dependent attenuation necessitates regularisation in order to harmonise the radar sensing model across the 60GHz band irrespective of the utilised channel.

To further study the subtle effects of a channel-dependent radar operation, we conduct radar simulations in two channels: a flat 60GHz and a sloping 67GHz—see Figure 3b.

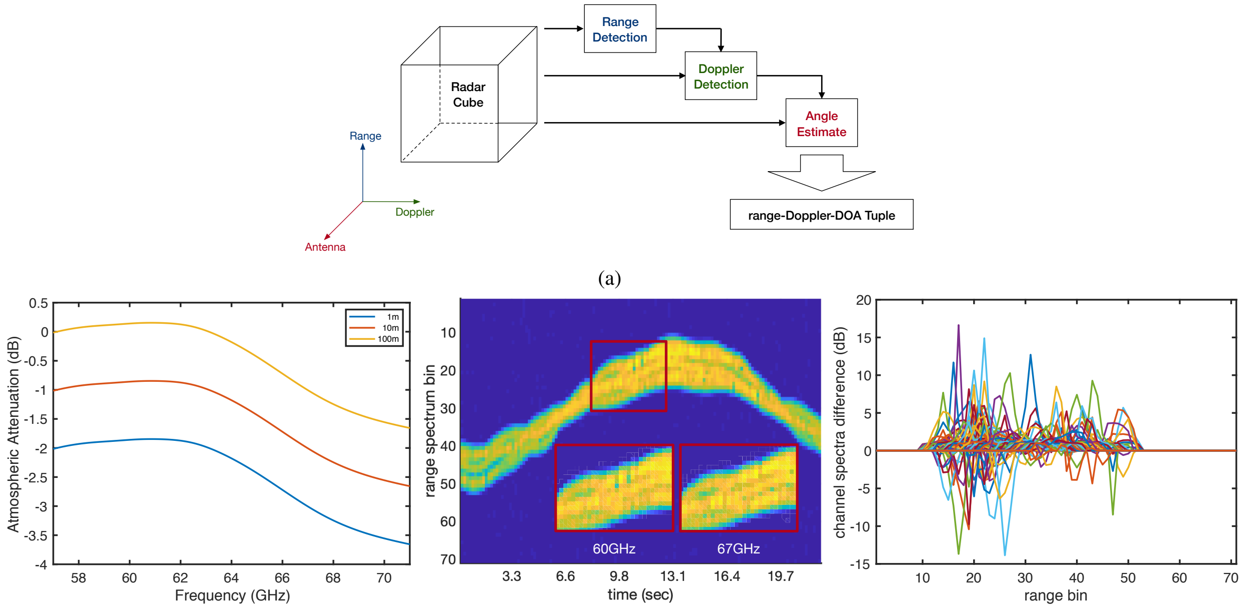

(a) (b) (c) (d)

Fig. 3: (a) Compute-efficient sequential radar detection. (c) 60GHz & 67GHz range spectra and their (d) discrepancies arising from the subtle indoor 60GHz band's channel-specific atmospheric attenuation behaviour in (b).

Identical 2GHz bandwidth around these centre frequencies is utilised. The simulation scenarios is a specialised repeat of one of the human traces making up the results presented in Figure 2. Specifically, the simulation scenario incorporates multi-body part traces of a walking person [12]. The body targets correspond to: a dominant torso, 2 arms, and 2 legs of lesser RCS's. Due to the mixing of the multi-body part tones, we aim to better expose the aforementioned subtle atmospheric attenuation-induced artefacts as observed within the radar sensing model.

Figures 3c & 3d illustrate the subtle channel-dependent artefacts. The intensity of the multi-body part range cluster for the 60GHz and 67GHz channels is plotted throughout simulation time in Figure 3c. Zooming into the red rectangles corresponding to a particular section of the cluster at both the flat 60GHz and sloping 67GHz channels, it is hard to discern the subtle differences in two channels by eye. To better expose such range spectrum channel artefacts, we take the difference and plot it overlaid in Figure 3d. With the exception of the operation channel, simulation settings were identical, including the receiver noise random number generator seed i.e. deterministic noise. Comparing the magnitudes of the two spectra, we can clearly see that these transient errors are confined to the cluster zone and are in the 10's of decibels.

The above example serves to motivate the need to regularise for these effects using a robust approach—be them intrinsic to the radar measurement process (e.g. atmospheric spectra artefacts), or extrinsic arising from users performing identical tasks in individualised ways. Thus irrespective of the vector of variation, we would remark that ML techniques using empirical data represent the best known approach to building practical systems that have built-in intelligence and resilience in highly complex sensing tasks such as radar and computer vision. The downside, however, are labour-intensive data collection campaigns to cover the sampling space in order to ensure that the ML model generalises in the real-world.

IV. Conclusion

In this paper, we present our blueprint of a future mm-wave network that supplies sensing functions alongside traditional data communications. Compared to current sub-6GHz systems, the vast spectral resources of mm-wave systems will enable higher rate communications and higher resolution radar sensing capabilities. Because of inherent beamformed nature of mm-wave systems, spatial multiplexing will allow the simultaneous sensing of multiple objects. Such a perceptive future network will make possible new transformative applications for digital health and/or building analytics. We demonstrate the feasibility of integrating robust radar sensing techniques within a communication protocol and show that the associated challenges—chiefly radio resource management, processing complexity, and inference—are within the reach of current technologies. We hope our blueprint proves informative and serves to motivate the community to press ahead with tackling the design and implementation of next generation mm-wave networks.

References

[1] G. Bielsa, A. Loch, I. Tejado, T. Nitsche, and J. Widmer, "60 GHz Networking: Mobility, Beamforming, and Frame Level Operation From Theory to Practice," *IEEE Transactions on Mobile Computing*, 2018.

[2] M. R. Akdeniz, Y. Liu, M. K. Samimi, S. Sun, S. Rangan, T. S. Rappaport, and E. Erkip, "Millimeter wave channel modeling and cellular capacity evaluation," *IEEE journal on selected areas in communications*, vol. 32, no. 6, pp. 1164–1179, 2014.

[3] G. L. Charvat, *Small and Short-Range Radar Systems*, 1st ed. Boca Raton, FL, USA: CRC Press, Inc., 2014.

[4] I. Milani, F. Colone, C. Bongioanni, and P. Lombardo, "WiFi emission-based vs passive radar localization of human targets," in *Proceedings of the 2018 IEEE Radar Conference*, 2018, pp. 1311–1316.

[5] M. Zhao, F. Adib, and D. Katabi, "Emotion recognition using wireless signals," in *Proc. of the 22Nd Annual Int'l. Conf. on Mobile Computing and Networking*, ser. MobiCom '16. New York, NY, USA: ACM, 2016, pp. 95–108.

[6] D. Wu, D. Zhang, C. Xu, H. Wang, and X. Li, “Device-free WiFi human sensing: From pattern-based to model-based approaches,” *IEEE communications magazine*, vol. 55, no. 10, pp. 91–97, 2017.

[7] S. Bakhtiari, T. W. Elmer, N. M. Cox, N. Gopalsami, A. C. Raptis, S. Liao, I. Mikhelson, and A. V. Sahakian, “Compact millimeter-wave sensor for remote monitoring of vital signs,” *IEEE Transactions on Instrumentation and Measurement*, vol. 61, no. 3, pp. 830–841, 2012.

[8] L. Rosenbaum and W. H. Shrank, “Taking Our Medicine — Improving Adherence in the Accountability Era,” *New England Journal of Medicine*, vol. 369, no. 8, pp. 694–695, 2013, pMID: 23964931.

[9] W. Jiang *et al.*, “Towards environment independent device free human activity recognition,” in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom ’18. New York, NY, USA: ACM, 2018, pp. 289–304.

[10] Y. L. Sit, B. Nuss, and T. Zwick, “On mutual interference cancellation in a MIMO OFDM multiuser radar-communication network,” *IEEE Transactions on Vehicular Technology*, vol. 67, no. 4, pp. 3339–3348, 2018.

[11] J. Palacios, P. Casari, and J. Widmer, “JADE: Zero-knowledge device localization and environment mapping for millimeter wave systems,” in *IEEE INFOCOM 2017*, pp. 1–9.

[12] T. Theodoridis, “UCI machine learning repository,” 2011. [Online]. Available: https://archive.ics.uci.edu/ml/datasets/Vicon+Physical+Action+Data+Set

[13] J. Lien, N. Gillian, M. E. Karagozler, P. Amihood, C. Schwesig, E. Olson, H. Raja, and I. Poupyrev, “Soli: Ubiquitous Gesture Sensing with Millimeter Wave Radar,” *ACM Trans. Graph.*, vol. 35, no. 4, pp. 142:1–142:19, Jul. 2016.

[14] S. Wang, V. M. Patel, and A. Petropulu, “The Robust Sparse Fourier Transform (RSFT) and Its Application in Radar Signal Processing,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 53, no. 6, pp. 2735–2755, 2017.

[15] M. Skolnik, “Systems aspects of digital beam forming ubiquitous radar,” Naval Research Lab, Washington DC, Tech. Rep., 2002.

[16] V. C. Chen, F. Li, S.-S. Ho, and H. Wechsler, “Micro-Doppler effect in radar: phenomenon, model, and simulation study,” *IEEE Transactions on Aerospace and electronic systems*, vol. 42, no. 1, pp. 2–21, 2006.

[17] M. Zhao, Y. Tian, H. Zhao, M. A. Alsheikh, T. Li, R. Hristov, Z. Kabelac, D. Katabi, and A. Torralba, “RF-based 3D Skeletons,” in *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, ser. SIGCOMM ’18. New York, NY, USA: ACM, 2018, pp. 267–281.

Mohammed Alloulah is a member of technical staff at Nokia Bell Labs. Contact him at mohammed.alloulah@nokia-bell-labs.com.

Howard Huang is a distinguished member of technical staff at Nokia Bell Labs. Contact him at howard.huang@nokia-bell-labs.com.