Iris: A Directional MAC Protocol With Applications to Millimeter-Wave Mobile Ad-Hoc Networks

Bharadwaj Satchidanandan, Venkata Siva Santosh Ganji, and P. R. Kumar, *Fellow, IEEE*

***Abstract*—The scarcity of spectrum in sub-6 GHz frequency bands to meet projected wireless traffic demands has led to the wireless industry incorporating millimeter-wave technology in the design of next-generation wireless systems. The high path loss of millimeter-wave signals necessitates radios that operate in these frequencies to employ highly directional beams for transmission and reception. The transition from traditional omni-directional transmission and reception to highly directional links has drastic implications on the Medium Access Control (MAC) design, since it shifts the objective of the MAC layer from proactive interference avoidance to transmitter-receiver coordination to achieve beam alignment. In this paper, we present IRIS, a directional MAC protocol for mm-wave ad-hoc mobile networks that achieves this objective. Specifically, the Iris protocol is designed to distributedly coordinate the nodes so that transmitters and their intended receivers align their antenna boresights to establish a physical link between them when required. We establish certain performance guarantees that the Iris protocol provides, and illustrate the design process of a mm-wave ad-hoc MAC based on the Iris protocol.**

I. Introduction

Novel applications of wireless networks, such as in connected vehicles, Virtual Reality (VR), unmanned aerial systems, and Internet of Things (IoT), impose requirements such as low latency, high connection density, high data rate, and energy efficiency, on the communication network. The lack of available spectrum in sub-6 GHz frequency bands to meet these requirements has led to the third-generation partnership project (3GPP) incorporating, for the first time in any wireless standard, the use of Millimeter-Wave (mm-wave) communications in the fifth-generation (5G) cellular standard. Recognizing the scarcity of spectrum in sub-6GHz bands, the Federal Communications Commission (FCC) too, in the year 2016, released 10.85GHz of spectrum in the mm-wave frequency bands, 7GHz of which is unlicensed [1]. This move more than quadrupled the amount of licensed spectrum that the FCC had issued until that time [1], increased the amount of unlicensed spectrum by more than a factor of fifteen [1], and could potentially alleviate to a large extent the problem of "spectrum crunch" that the wireless carriers and the industry at large have been facing in the past decade.

Though mm-wave frequencies offer a large amount of bandwidth, they exhibit certain characteristics that render the design of communication systems operating at these frequencies radically different from those operating at sub-6GHz frequencies. Perhaps the most important of these is the high directionality of mm-wave links, which is used to combat the high path loss of these frequencies. To elaborate, recall from the Friis Transmission Equation that the path loss of a sinusoidal signal scales as f^2, where f is the frequency of the signal. This implies that for a given transmit power, antenna pattern, and transmitter-receiver separation, the received signal power at 60Ghz, which is the frequency band of interest in this paper, is 20dB lower than that at 6GHz. In addition to the free space path loss, occlusion and atmospheric absorption too cause signal attenuation. Specifically, most objects that occur in the environment, such as walls, buildings, human beings, trees, furniture, etc., block signals in the 60GHz frequency range. Also, the atmospheric absorption at 60GHz, which is about 20dB/Km, is about three orders of magnitude higher than that at sub-6GHz frequencies [2].

Now, the standard solution to overcome the high path loss is for the transmitters and the receivers to employ tens or even hundreds of antenna elements and beamform their transmissions and receptions, thereby amplifying the link gain. The required antenna dimensions at mm-wave frequencies make it feasible to embed such large antenna arrays in the form factor of a handheld device. Indeed, for a given antenna aperture, the transmit and receive antenna gains which feature in the Friis transmission equation scale with the operating frequency as f^2, so that all else remaining the same, the received signal power at 60GHz and in the presence of beamforming is 20dB higher than that at 6GHz with omni-directional transmission and reception [3]. The combination of higher received power and higher bandwidth at 60GHz makes mm-wave communications the technology of choice for a variety of emerging high-bandwidth applications such as virtual reality, wireless backhauling, high-definition video streaming, etc.

In this paper, we consider the problem of medium access control for mobile ad-hoc networks, where the nodes operate at a center frequency of about 60GHz. Such nodes, being highly directional, introduce the problem of "deafness" [3, 4] at the MAC layer. Deafness refers to the phenomenon in which a directional receiver cannot detect a transmission directed towards it by another node because of it not orienting its antenna boresight in the direction of the incoming transmission. As recognized in many prior works [3, 5, 6, 7, 4, 8], the phenomenon of deafness necessitates a complete rethinking of the medium access control layer, right from defining its functionality to protocol design. To elaborate, note that MAC protocols for ad-hoc networks with omni-directional transmitters and receivers,

This material is based upon work partially supported by ONR Contract N00014-18-1-2048, the U.S. Army Research Office under Contract No. W911NF-18-10331, NSF under Contract Nos. CNS-1646449, CCF-1619085, CNS-1719384, and Science & Technology Center Grant CCF-0939370.

such as CSMA/CA, Seedex [9], etc., are designed based on the paradigm of proactive interference avoidance, i.e., they are designed to induce distributed coordination among the nodes so as to prevent nodes that are in "close proximity" from transmitting simultaneously, thereby avoiding packet collisions at the receiving stations. On the other hand, if the transmitters and the receivers in the network are directional, then it is no longer necessarily the case that simultaneous transmissions by nodes in close proximity result in collisions. In fact, a prior study [5] has shown that for mm-wave networks with antenna beamwidths of about 10° and randomly located transmitter-receiver pairs, the probability of a packet collision due to interference from simultaneous transmissions is less than 4%, implying that interference avoidance is no longer the primary, or perhaps even a significant concern for MAC design; the small percentage of collisions that result from simultaneous transmissions can be handled in a reactive fashion rather than a proactive fashion [7], using retransmissions at the MAC layer or even at the transport layer. What, then, should the concern of the MAC layer be in a directional network, if not interference management?

At a high level, the functionality of the MAC layer can be interpreted as apportioning a shared resource among multiple nodes. Unlike in omni-directional networks, the shared resource in a highly directional network is not the physical medium, since the directionality of the nodes allows for its high spatial reuse, but rather the beams of the nodes, especially the receive beam. Specifically, unlike in omni-directional networks, the receive (and transmit) beam of a highly directional node can be "allocated" only to one other node in the network. Consequently, it is this resource that should be allocated by the MAC layer to each node across time. The functionality of the MAC layer in directional networks, therefore, is to induce distributed coordination among the nodes so that at all times, every transmitter and its intended receiver orient their antenna boresights in the appropriate directions so as to establish a physical link.

One of the factors which makes this problem challenging is the mobility of the nodes. Consider, for example, a network consisting of three nodes n_1, n_2, and n_3, and suppose that at time $t = 0$, each of them knows the direction in which it has to orient its antenna boresight in order to establish a physical link with any of the other nodes. Further suppose that the nodes n_1 and n_2 communicate at time $t = 0$. During the time that two nodes have between them an established physical link, they can perform adaptive beamforming to track the motion of each other using dedicated training sequences or by closing the loop around the received signal power. However, once this communication is complete, and either of the nodes, say node n_1, wishes to communicate with node n_3, it would not know the direction in which it has to orient its antenna boresight to again establish a physical link with that node. While node n_1 knows this direction at time $t = 0$, the nodes could have moved away from their initial locations as well as possibly rotated about their axes, thereby rendering the initial direction invalid. In order to avoid this problem, each node has to track the combined effect of translational and rotational motion of each of its neighbors persistently, even though it has the capability to establish a physical link with only one node at any given time.

In this paper, we present Iris, a directional MAC protocol for ad-hoc wireless networks that achieves efficient tracking. Specifically, the Iris protocol (i) is decentralized, (ii) allows each node to track the net effect of both translational and rotational motion of its neighbors persistently, (iii) allows for nodes to join or leave the network over time, (iv) induces coordination among the nodes in a distributed fashion, and (v) is Time Division Multiplexing (TDM)-based, and in an adaptive fashion modifies the TDM schedule as nodes enter or leave the system.

The rest of the paper is organized as follows. Section II gives a brief account of related work on this topic. Section III formulates the problem that is addressed in this paper. Section IV presents the Iris protocol for directional ad-hoc mobile networks, and establishes the performance guarantees provided by it. Section V presents a simple numerical design example to illustrate the design process of a mm-wave ad-hoc network based on the Iris protocol, and computes metrics such as throughput and overhead that are expected to result from such a design. Section VI concludes the paper.

II. Related Work

The use of directional antennas to reduce interference and increase the spatial reuse has been explored in many early works in the context of sub-6 GHz networks. One of the earliest such works is [10], which analyzes the average progress of a packet in a directional ad-hoc network for specific MAC protocols and routing policies. The transport capacity of an arbitrary omni-directional ad-hoc network is derived in [11] by bounding regions known as the "exclusion regions" of a transmission, and [12] derives the exclusion regions for arbitrary antenna patterns. This in turn can be used to derive the transport capacity of arbitrary directional ad-hoc networks. References [4, 8, 13, 14] are some other papers that address the problem of directional MAC.

More recently, there has also been work on the topic of mm-wave MAC. Interference analysis of mm-wave networks using the protocol model and the physical model are presented in [5, 7], and it is shown that probability of collisions due to interference from simultaneous transmissions is lesser than 4%, implying that interference management is not a major concern for MAC design in such networks. This leads to the pseudo-wired abstraction that is developed in these papers, which is also used in this paper. A MAC protocol for infrastructure-mode mm-wave wireless network is presented in [6], with the assumption that the nodes in the network are stationary, and that each node knows the direction of other nodes in the network. A more comprehensive survey of prior work in mm-wave as well as directional MAC protocols can be found in [15].

All of the aforementioned papers assume either that the nodes in the network are stationary, with each node knowing the direction in which it has to orient its antenna boresight in order to communicate with a particular node, or that an omni-directional side channel is available to exchange some control information. Fewer papers address the problem of MAC design when the nodes are mobile. References [16, 17] address the MAC problem for mobile directional networks, with the assumption that only the transmissions are directional, and the reception is omni-directional. This assumption removes the problem of deafness, which is one of the critical features of mm-wave wireless networks. A Polling-based MAC (PMAC) protocol is presented in [18] in which the nodes periodically poll each of its neighbors at agreed upon slots so as to recalibrate their beams and track the nodes as they move. The tracking mechanism that we propose in this paper is based on the idea of PMAC's periodic polling, and also our own prior work [19] for infrastructure-mode mm-wave directional networks. While [18] presents a tracking mechanism, a general methodology for determining the design parameters of the MAC protocol, such as the frame duration to *guarantee* tracking, slot lengths for neighbor discovery phase, polling phase, and data transfer phase, physical layer pilot overheads, etc., are not addressed in [18]. In this paper, we build upon the approach presented in [18], and develop a detailed design framework by which various MAC parameters can be chosen as a function of the mobility parameters of the nodes, randomness of the environment, mm-wave channel characteristics, etc. We also analyze metrics such as the link throughput and training overheads that are expected of our design in the context of mm-wave networks.

III. Problem Formulation

We suppose that at time $t = 0$, certain nodes are "born" at arbitrary locations on a plane. At birth, the nodes do not know the locations of the other nodes, or even if there are are any other nodes in the plane or not. In order to discover other nodes, the nodes perform neighbor discovery as per the Iris protocol that is described in Section IV, establish links with them using a standard neighbor association procedure, and gradually build the network over time. We allow for the network to be open, in that nodes can join and leave the network over time. We do not assume the clocks of all nodes to be synchronized, so that they have a common view of time, but do assume that they have identical drifts. However, this assumption can be relaxed by incorporating guard intervals in each time slot (defined in Section IV), and making appropriate modifications to the protocol to account for these guard intervals.

The nodes operate at a center frequency of 60 GHz. We do not assume the presence of any omni-directional side channel for the nodes to exchange any control information. While the typical azimuth pattern of a 64-element or a 128-element antenna array has a narrow, high-gain main lobe and a number of low-gain side lobes, for the purposes of MAC design, we abstract the azimuth pattern as having a uniform gain over the main lobe, and having zero gain outside of it. This is known as the pseudo-wired abstraction, and is propounded in earlier works [3] as being an efficacious abstraction of the antenna pattern insofar as MAC design is concerned. We suppose in this paper that all nodes have the same, fixed beamwidth, and denote its value by ϕ. Beamwidths in the range of 10° to 20° are feasible for mm-wave nodes with small form factors. We suppose that the nodes are equipped with digital or hybrid beamforming capability which allows them to electronically steer their beam in any direction.

A. Topological Coherence Time

One of the central notions in our MAC design is that of the topological coherence time introduced in [19]. As mentioned before, the nodes in the network could be mobile, and we denote by v_{max} and ω_{max} the maximum translational and rotational speeds of the nodes. The topological coherence time of the network, roughly speaking, is the duration of time that the topology of the network remains "approximately constant." In what follows, we give a more precise definition of this quantity, employing the same notation that we had used in [15].

Denote by G_* the minimum link gain required for two nodes to establish a physical link, and suppose that two nodes n_1 and n_2 have, at time t, their antenna boresights oriented such that the gain $G_{1,2}(t)$ of the link between them exceeds G_* by at least d dB. Let $\theta_1(t)$ be the angle along which n_1's antenna boresight is aligned at time t with respect to a reference ray in its local frame, and let $\theta_2(t)$ be defined likewise. Suppose further that "just after" time t, the nodes disconnect this physical link and orient their antenna boresights in a different direction, perhaps to communicate with some other node. Then, the $(d, q)-$topological coherence time T_{TC} is defined as the maximum duration since time t by which the nodes n_1 and n_2 should reorient their antenna boresights in directions $\theta_1(t)$ and $\theta_2(t)$, respectively, with respect to their local frames so that with "high" probability, a physical link can be established between them, i.e., the resulting link gain is at least G_*. More precisely, the topological coherence time $T_{TC}(t)$ is defined as the solution to

$$\begin{aligned} &\sup\ \tau \\ &s.t\ \ \mathbb{P}\{G_{1,2}(t+\tau) < G_{1,2}(t) - d\} \leq q. \end{aligned} \tag{1}$$

In the above definition, the probability is taken with respect to the random motion of the nodes as well as of the reflectors in the environment. We assume that the network is such that the $(d, q)-$topological coherence time is independent of t, and that the probability in (1), viewed as a function of τ, is a non-decreasing function. We are interested in the regime where d and q are "small" quantities. Once d and q are fixed, we drop the qualifier (d, q), and refer to the $(d, q)-$topological coherence time as simply the topological coherence time. The case when the topological coherence time is infinite is of limited interest since it corresponds to a situation where there is no mobility. We assume henceforth that the topological coherence time is finite.

At an operational level, the topological coherence time can be thought of as the maximum time that can elapse between two successive service periods of any node to any of its neighbors so that the node, by orienting its antenna boresight in the last known direction of that neighbor, can re-establish a physical link with it with high probability, thereby not "losing track" of that neighbor. Once a physical link of a possibly low gain is established, the nodes perform beam refinement to recalibrate their antenna boresights and improve the link gain.

IV. IRIS: A DIRECTIONAL MAC PROTOCOL FOR AD-HOC MOBILE NETWORKS

The Iris protocol divides time into contiguous slots of μ time units each. Given a particular slot duration, the topological coherence time can be expressed in discrete time as being K slots, where $K := \lfloor T_{TC}/\mu \rfloor$. Each node, whenever it communicates with a neighboring node, does so for the duration of a time slot. Consequently, the slot duration μ must be at least as large as the time required to transmit various pilot symbols to establish a physical link and exchange control information. Other than this, the value of μ can be chosen in an arbitrary fashion. For a reason that will be apparent shortly, the value μ is chosen such that K is an odd integer. Given a topological coherence time of K slots, the slots are aggregated into "frames" consisting of K slots each.

As is common, we model the network as a graph, with vertices denoting the nodes in the network and edges denoting neighbor relations, i.e., an edge (i, j) is present if node i includes node j as its neighbor and vice versa. Since the network is open, and nodes can join or leave the network over time, the graph is time-varying.

At the heart of the Iris protocol is a decentralized graph coloring algorithm wherein each pair of neighboring nodes colors the link that they share in such a manner that no two adjacent edges share the same color. In what follows, we describe the protocol in detail in several steps.

Let $\Delta := \frac{K+1}{2}$, and $\mathcal{C} := \{1, \ldots, 2\Delta - 1\}$ be a "color set" consisting of $2\Delta - 1$ "colors." The Iris protocol colors each slot in a frame using one of the $2\Delta - 1$ colors such that no two slots in a frame share the same color. In this paper, we assume that the same permutation of $\mathcal{C}$ is used to color the slots in every frame, resulting in a periodic slot color sequence with period $K = 2\Delta - 1$. However, it is straightforward to modify the protocol to allow for different frames to be colored using different permutations of $\mathcal{C}$. In any case, the slot color sequence is fixed, and is known to every node at its birth.

The Iris protocol consists of three main phases.

A. The Neighbor Discovery Procedure

If a node wishes to perform neighbor discovery in a particular time slot, it randomly orients its antenna sector in one of the M predefined sectors on the plane, where $M := \frac{2\pi}{\phi}$. After a random back-off, the node transmits beacon symbols to inform the nodes that may be present in that sector of its presence. The random back-off ensures that it's own beacon transmission doesn't interfere with the beacon transmitted by a node that may be present in that sector and attempting to perform neighbor discovery. Suppose that two nodes discover each other in the process. Each node then expands its set of neighbors to include the newly discovered node. After that, the nodes spend some time to establish a physical link, as a part of which they transmit the necessary training symbols such as channel estimation pilots, symbol timing recovery pilots, frame synchronization pilots, beam refinement pilots, etc. Once a link is established, the nodes color the link that they share, or equivalently, color each other using a common color, by initiating the "color assignment phase."

B. The Color Assignment Procedure

Two nodes p and j that wish to assign a common color to each other first elect at random, perhaps by seeing who has generated a higher random number, a leader among themselves. Denote the elected leader by l, and the node that is not the leader by s. After leader election, node s transmits a set F_s, known as its feasible color set, to node l. The feasible color set of a node at a particular time slot is essentially the set of (available) colors that none of its neighbors are colored with at that time. For a node that has no neighbors, the feasible color set is the entire color set $\mathcal{C}$. Once node l receives the feasible color set F_s, it computes the intersection $F_s \cap F_l$ to find the set of colors feasible to both node p and node j. As detailed in Section IV-C, the Iris protocol ensures that no node has more than Δ neighbors at any time, and consequently, this intersection is guaranteed to be nonempty. Node l chooses one of the colors $\widehat{c}$ in the intersection uniformly at random, and informs node s of that color. Each node then colors the link that they share, or equivalently, colors each other, using $\widehat{c}$, and removes that color from its feasible color set.

C. The Link Establishment Procedure

We now describe the procedure that a generic node p follows at a generic time slot n under the Iris protocol. A node p (i) checks the color $c[n]$ of that time slot, (ii) checks if it has any neighbor that has been colored using the color $c[n]$, and if so, orients its antenna boresight in the last known direction of that neighbor. If, on the other hand, it has no neighbor colored using $c[n]$, *and it has lesser than* Δ *neighbors*, then it performs neighbor discovery in that time slot.

Since each of the $2\Delta - 1$ colors is used once in every frame, it follows that whenever a node wishes to communicate with a particular neighbor, the number of time slots that would have elapsed since its last communication with that neighbor would be $K - 1$, which is lesser than the topological coherence time of the network. Therefore, with "high" probability, it can re-establish a physical link with that neighbor. In the event that a node cannot re-establish a link with a neighbor in a particular frame, so that a time duration larger than the topological coherence time elapses without any communication with that neighbor, it assumes that that neighbor has exited the network, removes it from its list of neighbors, and adds the color of that neighbor back to its feasible color set F_p. The probability that a node cannot re-establish a physical link with a neighbor in a frame, given that that neighbor has not exited the network, can be controlled by controlling the value of T_{TC}, as described in Section III-A. If a physical link between two nodes is "lost" in a particular frame, it can be re-established only when they perform neighbor discovery. In networks where the nodes have very narrow beamwidths, the average initial access delay is high, and in such cases, it could be of interest to choose T_{TC} in a manner that renders the aforementioned probability to be of a "small" value.

A pseudocode of the Iris protocol is presented in Algorithm 1. The following proposition, stated without proof, follows immediately, and shows that directional nodes that implement the Iris protocol eventually discover other nodes within range and form a network (proposition 1(i)), and that each node schedules each of its neighbors once every K time slots (proposition 1(ii)), thereby ensuring that with high probability, the nodes don't lose track of their neighbors.

Proposition 1. *Consider an ad-hoc network in which nodes can enter and leave the network at arbitrary times, and suppose that every node employs Algorithm 1. Then,*

(i) If p and j are any two nodes in the network that are within each other's range, and if the number of neighbors of each of the nodes is lesser than Δ in all of the slots in any T frames, then,

$$\mathbb{P}(\textit{Nodes } p \textit{ and } j \textit{ do not discover each other in any of the } T \textit{ frames}) \leq \left(\frac{M^2-1}{M^2}\right)^T. \quad (2)$$

(ii) Every link in the network is colored by the nodes composing that link in such a manner that no two adjacent links share the same color. Consequently, every node schedules every one of its neighbors once every K slots.

V. A Numerical Design Example

In this section, we present a simple numerical example to illustrate the design process of a mm-wave ad-hoc network based on the Iris protocol, and compute the throughput and overhead that result from such a design. The following design is based on certain rules of thumb and conservative estimates of the relevant system parameters. An optimized design based on actual measurements of the system parameters is likely to yield higher throughput and lower overheads than those reported here.

Consider an ad-hoc network of mm-wave nodes where each node operates at a center frequency of $f_c = 60$ GHz, bandwidth of $W = 2$ GHz, and has a beamwidth of $\phi = 10°$. Each link is assumed to operate in Time Division Duplex (TDD) mode, and the fraction of time that each node occupies the link is determined by the nodes themselves in accordance with some policy as a part of control information exchange.

We first estimate the topological coherence time, which is a function of the mobility parameters of the nodes composing the network. We suppose in this example that the nodes to move at pedestrian speeds. As a conservative estimate, we

Algorithm 1 The Iris Protocol for Node p

1: Initialize arrays $N_p \leftarrow [\ ]$, $C_p \leftarrow [\ ]$, and initialize set $F_p \leftarrow \{1, \ldots, 2\Delta - 1\}$
2: **while** *CurrentTime* $\neq (k-1)h$ for some $k \in \mathbb{N}$ **do**
3: Wait
4: *SlotColor* $=$GETSLOTCOLOR(k)
5: **if** $\exists\ i \in \{1, \ldots, 2\Delta - 1\}$ such that $C_p(i) = SlotColor$ **then**
6: Beamform towards the i^{th} neighbor using the last known beamforming weights for that neighbor and attempt to re-establish physical link with that node.
7: **if** A link with neighbor i cannot be established **then**
8: $F_p \leftarrow F_p \cup \{SlotColor\}$.
9: Remove neighbor i from the list of neighbors.
10: Set $C_p(i)$ to be empty.
11: **GoTo** line 2
12: **else**
13: **if** $Deg(p) \leq \Delta - 1$ **then**
14: NEIGHBORDISCOVERY
15: **procedure** NEIGHBORDISCOVERY
16: Point the antenna beam in one of the M pre-defined sectors at random and search for new neighbors in that sector by listening for beacon symbols or transmitting beacon symbols after a random back-off.
17: **if** A new neighbor is discovered with address *ID* **then**
18: Include *ID* in entry j of N_p where j is some index such that $N_p(j)$ is not occupied by any other node
19: *SetColor*(j)
20: Communicate with neighbor j till the end of the current time slot.
21: **else**
22: **GoTo** Line 2 of Algorithm 1 at the end of the current slot.
23: **procedure** SETCOLOR(j)
24: $Leader =$ LEADERELECTION$(p, N_p(j))$
25: **if** $Leader = p$ **then**
26: Receive the set $F_{N_p(j)}$ from neighbor j
27: $\widehat{c} =$ *Unif* $(F_p \cap F_{N_p(j)})$
28: $C_p(j) = \widehat{c}$
29: $F_p \leftarrow F_p \setminus \{\widehat{c}\}$
30: Send message "Change color to $\widehat{c}$" to neighbor j
31: **else if** $Leader = N_p(j)$ **then**
32: Transmit the set F_p to neighbor j
33: **while** A color change message is not received from neighbor j **do**
34: Wait
35: **if** "Change color to $\widehat{c}$" message is received from neighbor j **then**
36: $C_p(j) = \widehat{c}$
37: $F_p \leftarrow F_p \setminus \{\widehat{c}\}$

suppose that they move at maximum $v_{max} = 5m/s$. The rotation of the nodes also causes misalignment of the antenna boresights, and consequently affects the topological coherence time. As a conservative estimate, we suppose that the nodes rotate at a maximum speed of $\omega_{max} = 2\pi$ rad/s. Owing to the fact that mm-wave signals, as mentioned before, are (i) transmitted in a directional manner, and (ii) occluded by most objects in the environment, such as buildings, walls, trees, indoor objects, etc., we suppose that the channel is predominantly line-of-sight, and that the multipath components are sparse and of low power. As outlined in [19], the topological coherence time of such a network is expected to be in the order of tens of milliseconds. We suppose in this example that it is 10ms.

The next step is to determine the slot duration μ. Recall that the slot duration must satisfy two constraints, viz., it must be at least as long as the time duration required to transmit the necessary pilots such as channel estimation pilots, symbol timing recovery pilots, and beam refinement pilots, as well as control information, and that $\lfloor T_{TC}/\mu \rfloor$ is an odd integer. We estimate below the number of symbols required for each of these pilots.

The number of channel estimation (CE) pilots required depends on the delay spread of the channel. While this could vary significantly across different environments, for the purposes of this example, we assume that the delay spread is no more than $L = 500$ns. Prior mm-wave channel measurements [20] have shown that typical indoor channels exhibit peak delay spreads that are in this range. While it is sufficient to allocate for CE pilots a time duration that matches the channel delay spread, allocating a larger time duration allows the nodes to perform smoothing of the channel impulse response estimate in the presence of noise. As a conservative design, we allocate a CE time duration that is ten times the delay spread of the channel. This yields a CE pilot duration of 5μs.

Standard sequences such as the Golay sequence, which exhibit certain autocorrelation properties, are used to perform symbol timing recovery and frame synchronization. As described in [19], a 2048-length sequence provides robust timing recovery. For a symbol duration of $T = \frac{1}{W} = 0.5$ns, this translates to a synchronization pilot duration of roughly 1μs. Finally, we allocate about 10μs for beam refinement pilots as in [19], and 10μs for exchange of control information every time a physical link is established. This allows for a node to transmit 20 kilobits of control information in a time slot using binary pulse amplitude modulation. Adding the above pilot and control information durations, we obtain a total pilot-cum-control overhead of about 25μs. Since the link is birectional TDD, the total pilot-cum-control overhead taking into account the pilots transmitted by both the nodes is 50μs. It follows that the value of μ must be chosen to be larger than 50μs.

We consider a specification that each node must be able to support, say, thirteen neighbors at most, so that the number of slots per frame amounts to $K = 2\Delta - 1 = 25$. This in turn yields a slot duration of $\mu = T_{TC}/K = 400\mu$s. For a node speed of $v_{max} = 5$m/s and operating frequency of $f_c = 60$ GHz, the doppler spread is $D = f_c v/c = 1$kHz, which in turn yields a channel coherence time T_c in the order of $T_c = \frac{1}{D} = 1$ms. Since this is significantly larger than the slot duration, it is sufficient to transmit CE pilots just once per slot. It follows that the slot duration that has been designed is an order of magnitude larger than the minimum required slot duration to transmit the pilot symbols and exchange control information, thereby satisfying the main design constraint for μ. The resulting pilot and control overhead is $50\mu s/400\mu s$ or 12.5%.

Observe that the Iris protocol forces the nodes to remain idle $\Delta - 1$ time slots per frame, which results in a forced idling time fraction of at least $\frac{\Delta-1}{2\Delta-1}$. For moderate or large values of Δ, this fraction is about 50%, and is largely insensitive to the value of Δ. Hence, no choice of the design parameters can cause significant reduction of the idling time. In light of this, decentralized coloring algorithms that employ fewer colors, or those that adapt the number of colors to the (time-varying) chromatic index of the graph, could be of interest.

Finally, we compute the per-link throughput attainable by this design. Note that each link in the network is utilized for 400μs every 10ms, of which 350μs is used to transmit payload, and 50μs is used to transmit pilots and control information. Assuming that the nodes employ rate-1/2 forward error correction and 16-QAM signal constellation, each link can provide a coded data rate of 140Mbps. This can be apportioned between the nodes composing the link in any proportion that they desire. Other important quality of service metrics such as end-to-end delay and throughput depend on certain higher layer functionalities such as the routing policy that is used.

VI. Conclusion

In this paper, we have addressed the problem of medium access control for mm-wave mobile ad-hoc networks. A distinguishing feature of such networks, one which renders the problem of MAC design for such networks radically different from that for traditional sub-6 GHz networks, is the high directionality of the nodes. The fact that the nodes can only detect signals whose angle-of-arrival lies in a narrow beam around its antenna boresight shifts the focus of the MAC layer from collision avoidance to transmitter-receiver coordination for antenna boresight alignment. This problem is compounded by the mobility of the nodes, since it necessitates that nodes persistently track other nodes in order to be able to establish a physical link with them when required, even though they can establish a physical link with only one node at a time. This brings to the fore a host of novel questions for system design, especially MAC design, such as efficient tracking mechanisms, tracking overheads, initial access mechanisms, initial access delay, efficient strategies for quick hand-offs, etc. In this paper, we have proposed a possible approach to address some of these questions. Specifically, we have presented the Iris protocol that allows various nodes in a directional mobile ad-hoc network to (i) discover peers, establish links with them, and gradually build a network, (ii) track the neighboring nodes that are already in the network, and (iii) modify in an adaptive fashion the nodes' schedules as existing nodes leave the network or new links are formed. We have also illustrated, using a numerical example, the design process of a mm-wave ad-hoc MAC based on the Iris protocol, and computed metrics such as link throughput and overheads that the Iris protocol yields. Future work includes implementing the tracking mechanism on a mm-wave testbed and evaluating its performance.

REFERENCES

[1] "Fact sheet: Spectrum frontiers rules identify, open up vast amounts of new high-band spectrum for next generation (5g) wireless broadband," 2016. [Online]. Available: https://docs.fcc.gov/public/attachments/DOC-340310A1.pdf

[2] T. S. Rappaport, J. N. Murdock, and F. Gutierrez, "State of the art in 60-ghz integrated circuits and systems for wireless communications," *Proceedings of the IEEE*, vol. 99, no. 8, pp. 1390–1436, Aug 2011.

[3] S. Singh, R. Mudumbai, and U. Madhow, "Distributed coordination with deaf neighbors: efficient medium access for 60 ghz mesh networks," in *INFOCOM, 2010 Proceedings IEEE*. IEEE, 2010, pp. 1–9.

[4] R. R. Choudhury and N. H. Vaidya, "Deafness: a mac problem in ad hoc networks when using directional antennas," in *Proceedings of the 12th IEEE International Conference on Network Protocols, 2004. ICNP 2004.*, Oct 2004, pp. 283–292.

[5] R. Mudumbai, S. K. Singh, and U. Madhow, "Medium access control for 60 ghz outdoor mesh networks with highly directional links," in *IEEE INFOCOM 2009*, April 2009, pp. 2871–2875.

[6] S. Singh, F. Ziliotto, U. Madhow, E. Belding, and M. Rodwell, "Blockage and directivity in 60 ghz wireless personal area networks: from cross-layer model to multihop mac design," *IEEE Journal on Selected Areas in Communications*, vol. 27, no. 8, pp. 1400–1413, October 2009.

[7] S. Singh, R. Mudumbai, and U. Madhow, "Interference analysis for highly directional 60-ghz mesh networks: The case for rethinking medium access control," *IEEE/ACM Trans. Netw.*, vol. 19, no. 5, pp. 1513–1527, Oct. 2011. [Online]. Available: http://dx.doi.org/10.1109/TNET.2011.2122343

[8] R. R. Choudhury, X. Yang, R. Ramanathan, and N. H. Vaidya, "On designing mac protocols for wireless networks using directional antennas," *IEEE Transactions on Mobile Computing*, vol. 5, no. 5, pp. 477–491, May 2006.

[9] R. Rozovsky and P. R. Kumar, "Seedex: A mac protocol for ad hoc networks," in *Proceedings of the 2nd ACM international symposium on Mobile ad hoc networking & computing*. ACM, 2001, pp. 67–75.

[10] C.-J. Chang and J.-F. Chang, "Optimal design parameters in a multi-hop packet radio network using random access techniques," *Computer networks and ISDN systems*, vol. 11, no. 5, pp. 337–351, 1986.

[11] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Transactions on information theory*, vol. 46, no. 2, pp. 388–404, 2000.

[12] A. Agarwal and P. R. Kumar, "Improved capacity bounds for wireless networks," *Wireless communications and mobile computing*, vol. 4, no. 3, pp. 251–261, 2004.

[13] R. R. Choudhury, X. Yang, R. Ramanathan, and N. H. Vaidya, "Using directional antennas for medium access control in ad hoc networks," in *Proceedings of the 8th annual international conference on Mobile computing and networking*. ACM, 2002, pp. 59–70.

[14] Y.-B. Ko, V. Shankarkumar, and N. H. Vaidya, "Medium access control protocols using directional antennas in ad hoc networks," in *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 1. IEEE, 2000, pp. 13–21.

[15] B. Satchidanandan, S. Yau, S. S. Ganji, P. R. Kumar, A. Aziz, A. Ekbal, and N. Kundargi, "A Directional Medium Access Control Protocol for 5G Millimeter-Wave Local Area Networks," in *Lecture Notes on Computer Science*, to appear.

[16] A. Nasipuri, S. Ye, J. You, and R. E. Hiromoto, "A mac protocol for mobile ad hoc networks using directional antennas," in *Wireless Communications and Networking Confernce, 2000. WCNC. 2000 IEEE*, vol. 3. IEEE, 2000, pp. 1214–1219.

[17] T. Korakis, G. Jakllari, and L. Tassiulas, "Cdr-mac: A protocol for full exploitation of directional antennas in ad hoc wireless networks," *IEEE Transactions on Mobile Computing*, vol. 7, no. 2, pp. 145–155, 2008.

[18] G. Jakllari, W. Luo, and S. V. Krishnamurthy, "An integrated neighbor discovery and mac protocol for ad hoc networks using directional antennas," *IEEE Transactions on Wireless Communications*, vol. 6, no. 3, pp. 1114–1024, March 2007.

[19] B. Satchidanandan, S. Yau, P. R. Kumar, A. Aziz, A. Ekbal, and N. Kundargi, "TrackMAC: An IEEE 802.11ad-Compatible Beam Tracking-Based MAC Protocol for 5G Millimeter-Wave Local Area Networks," in *10th International Conference on Communication Systems Networks (COMSNETS)*, Jan 2018, pp. 185–182.

[20] P. F. M. Smulders and A. G. Wagemans, "Frequency-domain measurement of the millimeter wave indoor radio channel," *IEEE Transactions on Instrumentation and Measurement*, vol. 44, no. 6, pp. 1017–1022, Dec 1995.