Relay Selection and Scheduling for Millimeter Wave Backhaul in Urban Environments

Qiang Hu and *Douglas M. Blough*
School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, 30332

***Abstract*—Millimeter wave (mmWave) communication is a key enabling technology for 5G cellular systems. However, due to mmWave propagation characteristics, link length for very high rates is limited and will likely necessitate the use of relay nodes for longer-range ultra-high-speed backhaul communications. This paper investigates relay selection and scheduling to support high end-to-end throughput in mmWave relay-assisted backhaul networks in urban environments. A major challenge in urban environments is the presence of large obstacles (buildings) that block long line-of-sight paths, which are necessary for very high capacity mmWave links. Using a 3D model for buildings targeted at urban environments, we provide optimal and efficient algorithms both for scheduling communications along a single mmWave relay-assisted path and for choosing the relay-assisted path with maximum throughput among all candidate paths connecting a given base station pair. In addition to proving optimality of these algorithms, we evaluate their performance through simulations based on a real urban topology. Simulation results show that our algorithms can produce short relay paths with end-to-end throughputs of around 10 Gbps and higher that are capable of providing virtual mmWave links for a wireless backhaul use case. Our algorithms improve throughput from 23% to 49% over a range of settings, as compared to average relay paths, and throughput can be more than doubled compared to some relay path choices with similar numbers of relays.**

I. Introduction

Millimeter wave (mmWave) communication has the potential to provide ultra high speed wireless communication and it is being heavily researched for use in 5G cellular systems. However, there are a number of significant obstacles to adoption of mmWave communication. These include higher path loss compared to communication in the 700 MHz to 10 GHz range, which is the predominant portion of the wireless spectrum in use today. Other problems with mmWaves, e.g. scattering effects produced by rain drops [1], compound the path loss problem. Another significant problem for mmWave links is communication loss when an obstacle blocks the line-of-sight (LoS) path between transmitter and receiver, which is known as the blockage effect. These propagation issues limit the communication range for very high rate mmWave links to a few hundred meters or less.

There are positive aspects of mmWaves beyond the potential for high rates. One advantage is the ability to integrate a large number of antenna elements into wireless nodes, which can be used to produce narrow beamwidth directional antennas with high antenna gain. This helps alleviate signal propagation problems and reduces the interference footprint.

Due to the limited communication range of very high rate mmWave links, applications where high rate is desired over several hundreds of meters will likely require the use of relay nodes. We anticipate a sequence of relatively short but very high rate mmWave links combining to produce a long-distance high-rate mmWave path, which emulates a single long-range mmWave link. In this paper, we investigate the use of such relay paths to extend the range of wireless backhaul communications in an urban area. mmWave relay paths could also be useful in other applications that need long-range and high rate wireless communications, e.g., to deploy 5G service along highways or deliver high speed Internet to rural areas.

The specific problems addressed in this paper are relay selection and scheduling to maintain high end-to-end throughput in relay-assisted mmWave backhaul networks in urban areas. We propose a relay selection algorithm that can be applied in 3D urban scenarios to find available interference-free multi-hop relay paths that connect a pair of source and destination nodes and can sustain very high throughput. We show that the available paths are determined by several control parameters, such as the maximum number of hops, the link length range of individual relay links, the beamwidth of directional antennas, the locations of the source and destination pair, the density and height variation of buildings sitting between source and destination, and the density of candidate relay locations to be selected. Simulations incorporating the 3D urban area model built from a real urban scenario (i.e., downtown Atlanta, Georgia) are used to demonstrate the efficacy of our proposed algorithm for relay path selection.

The main contributions of our work include:

- a new mmWave backhaul network architecture for 5G, where base stations serve as mesh nodes in the backhaul, and base station pairs are connected by multi-hop paths consisting of dedicated inexpensive mmWave relays,
- a linear-time scheduling algorithm that achieves the maximum throughput possible along any multi-hop path that is subject to primary interference only, i.e. when there is no secondary interference, and
- a polynomial-time relay selection algorithm to find the interference-free path with maximum throughput given a maximum number of relays available to connect a given source, destination base station pair.

We also show that, over a range of settings, our path selection algorithm produces from 23% to 49% higher throughput than average interference-free paths satisfying the given criteria.

II. Related work

Optimizing performance of wireless mesh networks has been widely investigated. The prior work aimed to optimize performance through channel assignment, scheduling, routing, power control, and/or rate selection, e.g. [2]–[4]. Much of this work aims at controlling interference, whereas we show that

2155-6814/17 $31.00 © 2017 IEEE
DOI 10.1109/MASS.2017.48

interference is not an issue for mmWave networks in urban environments due to narrow-beam directional antennas and the presence of obstacles. We demonstrate the need for relays in high-throughput outdoor mmWave networks and consider the problems of relay location and scheduling in this novel environment.

In mmWave WLAN/WPAN indoor environments, use of relays has been considered in several papers. In [5]–[7], for each source and destination pair, a pre-existing node or candidate relay is chosen to serve as a relay, which provides the backup path in case of blockage. However, they only consider the short-range two-hop 2D indoor scenario which is different from our multi-hop 3D outdoor scenario here.

A few papers have addressed outdoor mmWave environments. Zhu, et al. [8], explores the feasibility of 60GHz picocells in typical urban environments. Their work focuses on the access of mobile users in picocell which is different from our multi-hop backhaul use case. Another work addresses throughput optimization in mmWave relay networks [9], but it does not consider the optimization of relay placement. In [10], the authors propose a general and tractable mmWave cellular model. But the theoretical results in this work are generated based on a stochastic 2D geometric model for the urban environment which is not practical. Moreover, their self-backhaul network architecture is different from our multi-hop relay assisted network achitecture. [11] address the performance of relay-assisted mmWave cellular networks considering only utilizes LoS in a stochastic 2D geometric model. Recently, Du. et al. [12], proposed a "straight-line" relay assisted mmWave backhaul to delivery Gbps data service; however, the frequencies used by different links are assumed as orthogonal, which is different from our work where the same band is used by all base stations and relays in the backhaul.

As for interference analysis in mmWave networks, Madhow, et al. [13], conducted a probabilistic analysis of interference in a mmWave mesh network with uncoordinated transmissions. The paper concludes that secondary interference can be ignored with narrow beamwidth. However, the wireless backhaul in 5G considered herein requires both higher data rate and better robustness than the use case in [13]. We show that secondary interference can still have a significant impact.

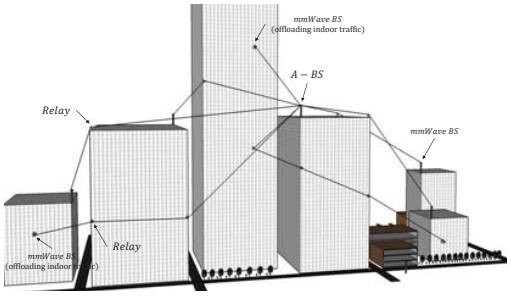

Fig. 1. Relay assisted mmWave backhaul network architecture.

III. Relay-assisted mmWave Backhaul Networks

Future 5G cellular system deployments will be based on small cells to provide better user access experience, such as higher data rates, in-building coverage, etc. A large number of small cell base stations (BSs) will be deployed in urban areas, which makes the traditional wired backhaul almost impossible considering the prohibitive cost and installation constraints. On the other hand, wireless, especially mmWave, backhaul networks are an ideal candidate for small cells. We propose a relay assisted mmWave mesh network architecture for backhaul in urban areas, in which mmWave relay nodes are used to assist in connecting BSs of the wireless mesh network.

A. Relay-assisted mmWave backhaul architecture

As shown in Fig. 1, in a typical urban area, there exist small cell BSs, macro cell BSs, and mmWave relays. It is assumed that all three types of entities can communicate in the mmWave band and they all have directional antennas with steerable beams to help compensate for high mmWave path loss. We assume that macro BSs and a small portion of small cell BSs will have wired connections, and we refer to these as anchored BSs (A-BSs), while the rest of the small cell BSs are non-anchored ones (i.e., no wired connections). The backhaul transfers data between non-anchored BSs and anchored BSs.

To achieve the required very high data rate of mmWave links (around 10 Gbps) in the backhaul network, LoS paths have to be utilized. Though some work tries to make use of NLoS paths to achieve a better coverage in mmWave cellular network, the propagation measurement at mmWave bands conducted in the urban environment has shown that typically every reflection would introduce more than 15dB signal attenuation [14], which is unacceptable in the backhaul use case, where very high signal to interference and noise ratio (SINR) must be achieved. However, due to the existence of buildings, walls, trees, and other obstacles, the LoS path between two BSs is usually blocked, even when they have a distance shorter than 100m. For example, small cell BSs deployed within buildings cannot connect with the outdoor BSs directly, due to the poor penetration property of mmWave signals. In fact, connecting indoor and outdoor BSs is a key feature in mmWave backhaul, because wireless operators prefer not to use wired Internet to offload their indoor wireless data traffic due to the high service cost. Thus, indoor mmWave BSs need an outdoor mmWave BS mounted on the surface of the building to act as the gateway connecting the indoor and outdoor scenarios in urban backhaul.

Above all, it is not always possible to build a self-backhaul network which only relies on the direct communications between BSs [10], or a centralized mmWave backhaul structure where single-hop connections are built between small cell BSs and macro cell BSs [15]. Thus, a relay assisted network architecture becomes a more favorable choice for mmWave backhaul in 5G era. The relay-assisted mmWave backhaul architecture we propose has the following characteristics:

- Any obstacles between two BSs are bypassed by a relay path. Every individual link along a relay path is LoS.
- The lengths of individual links are shortened, which increases the capacity of individual links along the path.

- The relays are dedicated devices in mmWave backhaul, which operate on a frequency different from the access frequency. The out-of-band backhaul solution further increases both the backhaul and access network capacities as compared with an in-band solution.
- With steerable beams and redundancy in the relay deployment, temporary blockage and link failure can be recovered locally, which reduces the burden of network layer rerouting and load balancing. Due to limited space, this part is not covered in here, but in our future work.
- The complexity of relays is much smaller than that of BSs, as discussed in the next subsection.

B. Relay use and deployment issues

As mentioned in the previous subsection, relays might be required in mmWave backhaul networks to connect two BSs that do not have a LoS path between them in urban areas. However, even if there are some situations where LoS paths exist, e.g. traversing a distance over a straight-line street, relays are still useful to improve performance. This is due to the poor mmWave signal propagation characteristics, which mean that achieving very high capacity with link lengths of more than a few hundred meters is not possible.

Fig. 2 shows the capacity of an individual mmWave link as the link length changes in an interference-free case. The channel model and the parameter choices used to generate this figure can be found in section IV-C and Tab. I respectively.

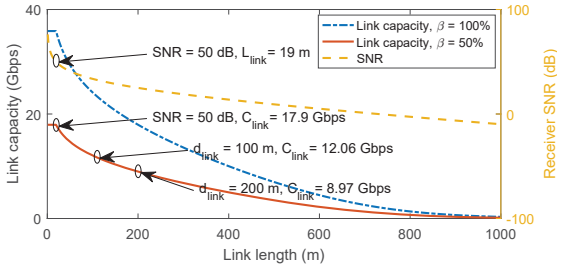

Fig. 2. LoS Link capacity vs. link length at 60 GHz

Note that as the link length decreases, the link capacity increases, which indicates that deploying more relays and keeping link lengths short might yield higher end-to-end throughput. In an extreme case, all hops along the path reach a maximum link capacity if the length of each hop is very short (in this setting, $C_{\max}$ is around 36 Gbps, when $L_{link} \leq 19$m). In this case, the maximum end-to-end throughput across a relay-assisted path is βC_{max}. The value β is called the link utility ratio, which represents the portion of a node's transmit time in the overall schedule. If paths are subject to primary interference, $\beta \leq 0.5$ and the maximum throughput in this setting is about 18 Gbps. Note that with a very long mmWave link, say around 600 meters, link capacity is significantly less (5 Gbps or less) even if $\beta = 1.0$. Fig. 2 shows the trend of βC_{link} when $\beta = 0.5$. The value of β greatly impacts the end-to-end throughput and, as mentioned, it is determined by the schedule of links along the path. Construction of an optimal schedule within a virtual long link will be addressed later.

The preceding discussion seems to indicate that deploying a very large number of relays with very short inter-relay distance is an optimal choice. However, since deploying more relays increases the cost, consumes more power, and increases end-to-end delay, in practice, it is necessary to limit the number of hops in a relay path.

Even if the hop count is limited, the total number of relays deployed across an urban area can become quite large. Thus, it is important to keep the cost of deploying and maintaining each relay low. Our architecture addresses this in several ways. First, contrary to prior work that proposed using relays in non-mmWave wireless backhaul, our relays are dedicated to the relay function and do not serve as access points for the user tier. We assume with small cells and the suitability of NLOS paths for the access tier, that good coverage can be achieved by the BSs alone. Second, each relay is part of only one path, which simplifies scheduling and other MAC aspects, compared again to more general ad hoc network relay structures used in non-mmWave wireless backhaul. Finally, we assume that relays are constrained by primary interference. Primary interference could possibly be avoided with multiple external antennas for each relay node that are mounted in a well separated fashion. However, this would signicantly increase the cost and complexity of deploying and maintaining relays and also limit the possible relay locations.

IV. 3D Modeling of Urban Area

In this section, we discuss the 3D modeling of urban area for analyzing the relay-assisted mmWave backhaul network. The LoS communication model used in our proposed relay selection algorithm is also covered.

A. 3D is essential in analyzing mmWave backhaul networks

In most previous work, the modeling of the deployment area is two-dimensional (2D) for simplicity. It is reasonable for 2D models to be used in analyzing the performance of previous generations of cellular system, such as 3G or 4G, since the signal at lower frequency band has better penetration and diffraction properties, and the cell size is much larger than the size of buildings. However, in mmWave band, the blockage effect is very significant, thus the height of a building and the height at which mmWave BSs and relays are deployed will largely determine whether possible paths can bypass the building. Skyscrapers may reach several hundred meters in a typical urban area. For example, there are 38 buildings taller than 122m (i.e., 400 feet) in downtown Atlanta, GA (see Fig. 3). As shown in Fig. 2, the maximum length of mmWave links in urban areas is considered around 100m–150m to achieve very high data rate (around 10 Gbps). Due to the similar scale of both building height and mmWave link length, the third dimension can have a big impact on the coverage of mmWave signals.

As mentioned above, the backhaul scenario discussed in our work is a mmWave backhaul network constructed by LoS paths. Thus, 3D modeling of the environment is more suitable than 2D, as it gives us a more practical view of the transmission environment, and the process of finding LoS

paths can be more accurate and reliable. On the other hand, 2D modeling may still be useful in the research on the cellular access tier, where NLoS path and multi-path effects will be considered. In addition, our 3D modeling uses the real building information. In a typical downtown area, buildings are arranged in blocks, and they have similar orientations, which is totally different from the assumption in previous stochastic 2D modeling, where Poisson Point Process is used to model the location of buildings and the orientations of buildings are assumed to follow a uniform distribution on $[0, 2\pi]$. For example, as we can see in Fig. 3, among the recorded 227 buildings in downtown Atlanta, most of them have an azimuth angle close to 0^o or 49^o, while only 7.9% have other orientations. In fact, the distribution of building orientations in Atlanta already appears more random than in other metropolitan areas such as Chicago or New York.

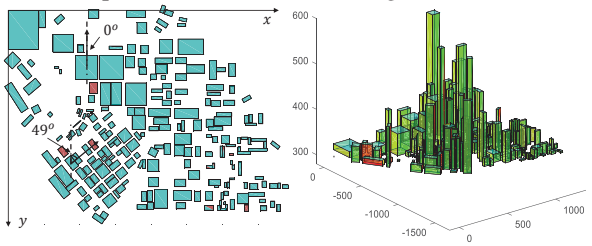

(a) (b)

Fig. 3. Building topology in downtown Atlanta. (a) Top view. (b) 3D view.

B. Building model

Considering the abundant existence of trees, moving vehicles, and other obstacles located at relatively low heights, it is a wise choice to deploy outdoor mmWave BSs and relays at a height higher than 5m. Otherwise, the blockage attenuation is high and temporary blockages can happen frequently. Thus, all possible LoS paths between mmWave nodes (i.e., BSs and relays) are above the level of 5m as well. Therefore, only buildings higher than 5m are modeled in our simulations. The blank areas in Fig. 3 are parks and ground parking lots.

As for a rectangular area to be modeled such as in Fig. 3, at the top-left corner, there is a reference point with a coordinate $(0, 0, 0)$. The horizontal and vertical lines are x-axis and y-axis, respectively, while the z-axis follows a left-hand coordinate system with x-axis and y-axis. There are 227 buildings higher than 5m in the area. The location of building $i, (1 \leq i \leq 227)$ is recorded as the coordinate (x_i, y_i, z_{gi}) of its center point from the top view, where z_{gi} is the height of its ground level. All buildings are modeled as cuboids, which is a common method used to reduce the complexity of modeling buildings. The orientation of building i is defined as the angle $\Theta_i \in [0^o, 90^o]$ between the normal direction of a building surface perpendicular to the XY plane and the Y^- direction.

C. Channel and antenna models

We make the standard assumption of additive white Gaussian noise channels. Link capacities are assumed to follow Shannon's Theorem with an upper limit, i.e.

$$C = B \log_2(1 + \min\{\text{SINR}, T_{max}\}) \ , \tag{1}$$

where B is channel bandwidth, SINR is signal to interference plus noise ratio, and T_{max} is the SINR that produces the network's maximum rate. In real networks, capacity cannot be increased without limit and this is captured by T_{max}. Since the LoS link is relatively short and the SINR at the receiver end is required to be high to support backhaul usage, it is assumed that the small scale fading is mild, thus the channel estimation penalty is negligible, and the maximum achievable rate of the channel can be closely approximated by eq. 1.

SINR is defined by the following relationship:

$$\text{SINR} = \frac{P_r}{N_T + I} \ , \tag{2}$$

where P_r is the power of the intended transmitter's signal when it reaches the receiver, N_T is the power of thermal noise, and I is the combined power of signals from any interfering transmitters.

For the calculation of both P_r and I, the Friis transmission equation is used:

$$P_r(d) = P_t \times G_t \times G_r \times (\frac{\lambda}{4\pi d})^{\eta} \times e^{-\alpha d} \tag{3}$$

where P_t is the transmit power, G_t and G_r are antenna gains at transmit and receive ends, respectively, λ is the wavelength of the signal, d is the propagation distance, α is the attenuation factor due to the atmospheric absorption, and η is the path loss exponent. Because of the short LoS link used in the backhaul and the high SINR at the receiver, ignoring the relatively small random attenuation due to the shadowing effect does not hurt the effectiveness of our analysis. However, due to the implementation loss (5 dB), noise figure (5 dB), and the heavy rain attenuation (10 dB/km), an additional link margin $L_m = 10\text{dB} + 10\text{dB/km} \times d$ should be deducted when calculating the receive power.

Different from the 2D antenna model used in the previous related work, such as [13], the antenna model used in 3D scenarios is also 3D. With directional antennas, the antenna gain is not the same for all azimuth and elevation angle values. Here we employ an ideal 3D-flat-top antenna model, which has a constant high gain G_{high} within the beam and a constant low gain G_{low} outside the beam.

V. Optimal Scheduling Algorithm for Interference Free Multi-Hop Paths

We first consider the problem of maximizing throughput on a single mmWave relay path. As mentioned in Section III-B, the end-to-end throughput is determined not only by the capacity of each individual link, but also by the link schedule. Due to the use of narrow beam directional antennas in the mmWave backhaul network, secondary interference-free multi-hop paths can be formed in urban areas through the relay selection method covered in Section VI. Thus, in this section, we present an algorithm to find the optimal schedule with maximum throughput for a given multi-hop path subject to primary interference only and with a single flow. Despite the strong constraints of our setting, this specific problem has not been studied to date in the literature. In the next section, we use the algorithm to find an overall best path given a set of candidate relay locations and a source, destination BS pair.

A. Formal problem statement

The scheduling problem in the proposed multi-hop interference-free path (virtual long link) for mmWave backhaul networks can be formally described as follows. A finite set of mmWave nodes is represented by $V = \{N_1, N_2, ..., N_{|V|}\}$. Due to the linearity of the multi-hop path, the nodes in V are ordered in space (see the abstract model in Fig. 4(a)). A link e_i is an ordered pair (N_i, N_{i+1}), where $1 \leq i \leq |V|-1$ and N_i, N_{i+1} are the transmitter and receiver of the link, respectively. The $|V|-1$ links form a set E, and the pair (V, E) forms a graph consisting of a single path from N_1 to $N_{|V|}$.

We say that a link is active if its transmitter is transmitting data to its receiver. A set of links $M \in E$ satisfies the primary interference constraint if no two links in M have a node in common. Such a set M is called a matching and allows all links in M to be simultaneously active. A schedule $\mathcal{S}$ is an indexed family $\mathcal{S} = (M_\alpha, \tau_\alpha : \alpha \in \mathcal{A})$, where the index set $\mathcal{A}$ is an arbitrary finite set, $\tau_\alpha \geq 0$ for each α and M_α is a matching for each α. The indicator vector $I(F)$, for any $F \subseteq E$, is the vector in R^E defined by

$$I_e(F) = \begin{cases} 1 & \text{if } e \in F \\ 0 & \text{if } e \in E - F \end{cases} \tag{4}$$

Using the above notations, the length t and the demand vector $f \in R^E$ of the schedule $\mathcal{S}$ can be defined by

$$t = \sum_\alpha \tau_\alpha \tag{5}$$

$$f = \sum_\alpha \tau_\alpha I(\mathcal{M}_\alpha) \tag{6}$$

where each element f_e in vector f is the total amount of time link e is active over the entire schedule. Thus, the optimal scheduling problem is equivalent to finding a schedule S for graph (V, E) with the minimum schedule length $t_{min}(f)$, where each element f_e satisfies the demand of the link e.

In [16], an algorithm is provided that computes the optimal schedule in polynomial time for any network that is subject to primary interference only through linear programming. However, [16] considers more general network structures, and the algorithm is relatively complicated. The time complexity for computing $t_{\min}(f)$ is $O(|V|^5)$ and the computation of the corresponding schedule requires $O(|E| \cdot |V|^5)$. Here, we provide a much simpler $O(|V|)$ algorithm to compute the optimal schedule for the single path case.

B. Optimal scheduling algorithm

In a relay path, all traffic is originated at the first node in the path and all traffic is consumed by the last node in the path.[1] Thus, the demand in bits, D, is the same for every link. Since different link lengths result in different link capacities, the time demand f_{e_i} of link e_i is determined by D and the capacity of e_i, which we denote by C_i. Considering a linear network, we use f_i in place of f_{e_i} to refer to the demand of link e_i for simplicity. Then,

$$f_i = D/C_i; \tag{7}$$

[1]In line with our virtual long link concept, we assume that traffic flows in only one direction along the path at a time.

The following theorem shows how the minimum schedule length can be calculated in single path networks with only primary interference. While this theorem can be proved by adapting some of the analysis in [16] to the single-path case, we instead prove it in a simpler way using induction.

Theorem 1. *The minimum schedule length in a multi-hop interference-free relay path is equal to the maximum demand sum of two consecutive links, i.e.,*

$$t_{\min} = \max_{1 \leq i \leq |V|-1} f_i + f_{i+1} \tag{8}$$

Proof. Due to the path's secondary-interference-free nature, the scheduling on the two sides of a pair of consecutive links are independent. Let the two consecutive links with maximum demand sum be e_j and e_{j+1}. We must prove that $t_{\min} = f_j + f_{j+1}$. It is clear that e_j and e_{j+1} cannot transmit at the same time due to the primary interference constraint and, therefore $t_{\min} \geq f_j + f_{j+1}$. It remains to show that $t_{\min} \leq f_j + f_{j+1}$, i.e. that all links' demands can be scheduled within $f_j + f_{j+1}$. We use induction to prove this.

1) *Basis*: It is clear that e_j and e_{j+1} can be scheduled within $f_j + f_{j+1}$. Without loss of generality, let e_j's transmission time be $[0, f_j)$, and e_{j+1}'s transmission time be $[f_j, f_j + f_{j+1})$. Also, denote $f_j + f_{j+1}$ by t_{end}.

2) *Forward inductive step*: Assume e_i is scheduled during $[0, f_i)$ and e_{i+1} is scheduled during $[t_{end} - f_{i+1}, t_{end})$. (Note that this holds for e_j and e_{j+1} from the base case.) Now, consider links e_{i+2} and e_{i+3}. e_{i+2} can be scheduled during $[0, f_{i+2})$. Note that e_{i+2} will be done transmitting before e_{i+1} starts transmitting, because $f_{i+1} + f_{i+2} \leq f_j + f_{j+1} = t_{end}$. Now, e_{i+3} can be scheduled during $[t_{end} - f_{i+3}, t_{end})$. Since, it is again true that $t_{end} = f_j + f_{j+1} \geq f_{i+2} + f_{i+3}$, then $t_{end} - f_{i+3} \geq f_{i+2}$ and the transmissions of e_{i+2} and e_{i+3} do not overlap and are completed within the total scheduling length of $f_j + f_{j+1}$.

The inductive step continues until the last link in the path, $e_{|M|}$. If $e_{|M|}$ is the 1st link in an inductive pair, i.e. $|M| = j + 2k$ for some integer k, then we must add an imaginary link $e_{|M|+1}$ with zero demand and schedule $e_{|M|}$ and $e_{|M|+1}$ according to the inductive rule in order to complete the induction.

3) *Backward inductive step*: Here, we must start from e_j and e_{j+1} and go down to e_1. The scheduling is the same as in the forward step and the logic is the same to show the scheduling does not violate the primary interference constraint. Here, if $e_1 = e_{j+1} - 2k$ for some integer k, then we must add an imaginary edge e_0 with demand zero and schedule it with e_1 to complete the induction.

Thus, all links along the linear network can be scheduled within $[0, f_j + f_{j+1} = t_{\min})$, and the proof is completed. □

Following the idea of the scheduling scheme outlined in the proof, once $t_{\min}$ has been calculated, the schedule for each link e_i in the network can be specified as:

$$\begin{cases} t = [0, f_i) & \text{if } i \text{ is } odd \\ t = [t_{min} - f_i, t_{min}) & \text{if } i \text{ is } even \end{cases} \tag{9}$$

The scheduling algorithm goes through the link demands one by one to find the largest consecutive pair and then assigns time slots to each link according to (9). The time complexity is thus $O(R)$, where R is the number of relays in the path.

An example network is shown in Fig. 4 with $C_2 > C_4 > C_3 > C_1$, and $\frac{1}{C_3} + \frac{1}{C_4} > \frac{1}{C_1} + \frac{1}{C_2} > \frac{1}{C_2} + \frac{1}{C_3}$. From Theorem 1, the minimum schedule length is $t_{min} = D(\frac{1}{C_3} + \frac{1}{C_4})$, and one possible link schedule is shown in Figure 4(b) following the scheme used in the proof.

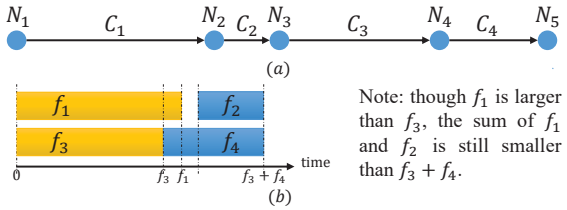

Fig. 4. (a) Abstract model of a multi-hop interference-free virtual long link. $|\mathcal{N}| = 5$. (b) Minimum schedule example.

VI. Relay Selection for Connecting Two Base Stations in Urban Area

As in the previous section, we focus on optimal path selection from one BS to the other. As long as channels are symmetric, which should be the case particularly since only LoS paths are used in our approach, the optimal path in the reverse direction will be the same. There are several ways that traffic in both directions could be handled. The simplest would simply be to alternate traffic flows in the two directions according to some schedule. This would obviously cut the overall throughput in each direction roughly in half. If that result is not satisfactory and relay cost is not a limiting factor, a network provider could choose to create separate paths in the two directions and achieve full duplex operation across the virtual long link. This would be possible as long as the endpoint BSs can create a reasonable separation between antennas for the two paths. Since our focus is solely on constructing optimal relay paths, we assume all traffic flows in the same direction in the remainder of the discussion.

A. Interference avoidance along the multi-hop relaying path

Considering a multi-hop relaying path connecting two BSs in the urban area, it is intuitive to think that the directions of all transmitters along the path are correlated with the direction from the source to the destination. Additionally, side/back lobes always exist in practice. Thus, mutual interference may exist due to the concurrent transmissions along the path.

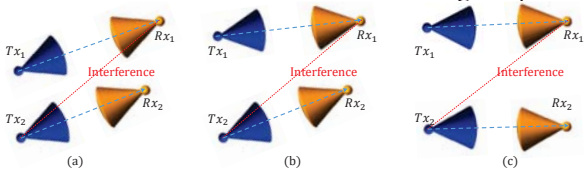

Fig. 5. Three different interference scenarios. (a) Most interference. (b) Medium interference. (c) Least interference.

Fig. 5 shows three different interference cases: (a) depicts the most interference case where the interference signal experiences G_{high} at both Tx_2 and Rx_1. In (b), the antenna gains on interference signal are G_{high} and G_{low}, while in (c), both gains are G_{low}.

If we assume the intended link length is 100 m, an interferer is 300 m away from the intended receiver, and $\beta = 50\%$, the achieved link rate is 1.3 Gbps, 7.2 Gbps, and 12.6 Gbps in case (a-c) respectively. Since in case (c), the amount of interference is smaller than the noise level, it is regarded as *interference free*. Note that the reflected signal is ignored in the interference analysis due to its low power and its very low possibility of reaching the intended receiver in a 3D scenario; however, the interference detection of one-time reflection signal can be easily incorporated into our algorithm if needed.

Mutual interference will significantly reduce the link capacity and the network throughput. Thus interference has to be avoided when selecting relays to form a multi-hop path in mmWave backhaul, due to its high throughput demand.

B. Finding the path with maximum throughput among all available interference-free paths

Algorithm 1 shows the pseudocode for our optimal path selection algorithm, which does a search of possible paths subject to several constraints, which are discussed below.

Algorithm 1 Finding the path with maximum throughput

Input: src, dst, $relays$, $numHop$, bW
Output: $maxPath$

$nodes$.add($relays, src, dst$); // $nodes$ stores all nodes
for $node$ in $nodes$ **do**
  $iNbs$.add(findAllLoSNbs(node, nodes));
  $iNbsList$.add($iNbs$); // store indices of neighbors
Initialize $allPaths$ as an empty list of paths;
$minSL$ = Inf; // minimum schedule length;
for $maxHop = 1 : numHop$ **do**
  Initialize $path$ as an empty list of nodes' indices;
  $path$.add($iSrc$); // $iSrc$ is the index of src in $nodes$
  $iHop = 1$; // the index of current hop;
  searchNextNode($path, iHop, maxHop$);
$maxPath = allPaths$.get(end);
Function: void searchNextNode($path, iHop, maxHop$)
if $iHop \leq maxHop$ **then**
  $iPreNode = path.get(iHop - 1)$;
  $iCurNbs = iNbsList.get(iPreNode)$;
  for iNb in $iCurNbs$ **do**
    Calculate the sum of demand of the current link and the previous link as $sumDemand$.
    if ($sumDemand \geq minSL$) **then**
      return
    if (!intfTest($iNb, path, iHop, bW$)) **then**
      $path$.add(iNb);
      if ($nodes.get(iNb) == dst$) **then**
        $allPaths$.add($path$);
        $minSL$ = calMinSL($path, iHop$);
      if ($nodes.get(iNb)! = dst$) **then**
        searchNextRelay($path, iHop + 1, maxHop$);
return

First, only the LoS neighbors of each $node$ (Lines 2–4) are considered as candidate nodes to be selected in one step. The start of a path is the given source node (Line 9). The main loop of the algorithm does a depth-first search for a maximum throughput path using a progressively larger maximum number of hops, $maxHop$, at each iteration (Lines 7–11). This is done because the paths found at one iteration can be very helpful in bounding the search done at the next iteration.

In the depth-first search (Line 13–27), when the index of current hop does not exceed $maxHop$ (Line 13), the selection of a candidate node makes the current path invalid in the following cases: (1) the sum of link demand of two consecutive links, the current link (i.e., from the previous node to the selected node) and the previous link in the path, exceeds the current minimum schedule length (Lines 18–19); (2) the candidate node will be interfered by previously selected nodes in the path, which means the relationship between two concurrent transmitting links is the most or medium interference case shown in Fig. 5 (Line 20). If neither of the above cases occurs, the node is added to the current path (Line 21). If the selected node is the destination, the path is added to the path list and the minimum schedule length is updated (Lines 22–24). Otherwise, the search continues to the next hop (Line 25–26). At the end, the last path added to the list is the maximum throughput path, as discussed next.

Note that condition (1) above comes from Theorem 1 in section V-B. When the first path to the destination is found, the minimum schedule length of that path is recorded. From that point forward, each partial path is tested against the minimum schedule length at each step. If at any point, its schedule length would become greater than the minimum already found, the path is abandoned. Thus, when a new path actually reaches the destination, it means that its schedule length is better than the best previously found. Therefore, at the end of the algorithm, the last path added to the path list is the optimal path. The minimum schedule length value is propagated from one iteration of the main loop to the next and, as just discussed, this helps eliminate many possible paths in performing the depth-first search to a given maximum number of hops.

The time complexity of Algorithm 1 is $O(D^M)$, where D is the maximum degree of the connectivity graph of candidate relay locations and M is the maximum hop count for a path. Two relay locations are neighbors in the connectivity graph if they have a LoS path between them. In the worst case, the time complexity is $O(N^M)$, where N is the total number of candidate relay locations but, due to the large number of obstacles in the urban environment, the maximum degree will typically be much smaller than N. As long as M is constant, the algorithm runs in polynomial time in either case, although the polynomial degree can be fairly large. In practice, it runs much faster than this bound in most cases due to the constraints that limit the number of paths searched.

VII. Numerical results and simulations

In this section, simulation results are provided to evaluate our relay path selection algorithms. The fixed parameters used throughout this section are shown in Tab. I. These parameter choices represent typical values taken from a survey of the mmWave literature. The topology of buildings in downtown Atlanta (see Fig. 3) is used. In our model, a large number of candidate relay locations are uniformly distributed on the surfaces of each building in the topology. The density δ_r is defined as the number of candidate relay locations per m^2. This topology contains 227 buildings higher than 5 meters, and for each building with a height between 20 and 200 meters, one of its rooftop corners is randomly picked as a candidate BS position (130 positions in total). BSs are expected to be deployed at positions with a good coverage of other relays mounted on the surfaces of surrounding buildings.

TABLE I
Parameters of simulation environment

BW	2.16 GHz	P_t	1 W	$G_{tx,rx}$	21.87 dBi
f_c	60 GHz	ϕ	16^o	η	2.0
L_m	10 dB	α	16 dB/km	T_{max}	50 dB

A. End-to-end throughput of mmWave virtual long links

We randomly chose BS pairs separated by a distance in the range of $[20, 200)$, $[200, 400)$, $[400, 600)$, $[600, 800)$, and $[800, 1000)$. 100 BS pairs were picked for each distance range. The maximum end-to-end throughput and minimum scheduling length of available interference-free paths with a limited number of hops was computed using our relay selection algorithm. One candidate relay location was placed randomly on every building surface plus an additional $0.002/m^2$ candidate relay locations were uniformly distributed over all surfaces.

TABLE II
Minimum number of hops for each BS pair ($\delta_r = 0.0020/m^2$)

BS distance (m)	Min number of hops (percentile)
$[20, 200)$	1 (48%), 2 (51%), 3 (1%)
$[200, 400)$	1 (13%), 2 (83%), 3 (14%)
$[400, 600)$	2 (30%), 3 (61%), 4 (9%)
$[600, 800)$	3 (35%), 4 (62%), 5 (3%)
$[800, 1000)$	4 (43%), 5 (49%), 6 (6%), 7 (2%)

Since deploying more relays along a path increases both cost and end-to-end delay, we are interested in performance at or near the minimum number of hops. Tab. II shows the distribution of the minimum number of hops for BS pairs over different distances, which is quite consistent for different BS pairs in the same range. In all cases, a difference of one in the minimum number of hops covers more than 90% of all BS pairs. One interesting note from this table is that, even for short mesh links ($< 200m$), more than half of the BS pairs do not have a LoS path connecting them due to blockages (minimum number of hops ≥ 2), which means that relays are absolutely necessary even across these short distances.

Fig. 6 shows the averages of the maximum throughput and minimum schedule length (for transferring 100 Gb data) over the BS pairs within each distance range, for the best paths both of minimum length and of minimum length plus one. These simulations were carried out with high candidate relay location density, meaning that a large number of choices exist for each link along the relay path. Note that for short mesh links ($< 200m$), the maximum throughput is around 15–18

Gbps, which exceeds the capacity of a 100 meter link with 50% utility rate (see Fig. 2). Here, shortening the individual link distances increases throughput compared to a longer 50% utility link due to the propagation loss of mmWave signals. When the distance increases, the end-to-end throughput of the optimal path decreases. However, throughput is maintained close to 10 Gbps, which is the typical requirement for backhaul links, for BS separations up to 600 meters. Above 600 meters, optimal throughput drops to around 8 or 9 Gbps, which is still quite good. Note that, in these cases, single-hop paths do not exist between any BS pairs, but if they did exist, even their 100% utility rates would be well below what is achieved by the best relay paths due to the mmWave propagation effects.

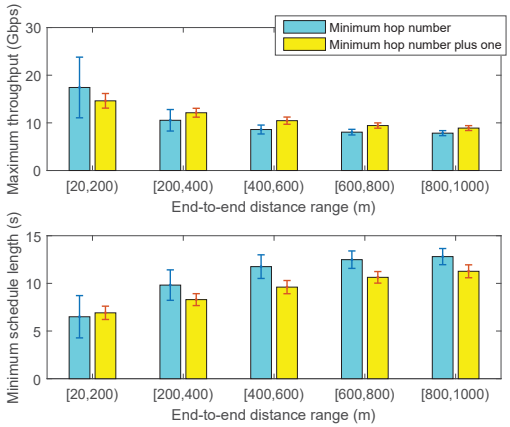

Fig. 6. Maximum end-to-end throughput and minimum scheduling length (SL) among all available interference-free paths. ($\delta_r = 0.0020/m^2$)

We are also interested in how throughput and delay change when more hops are allowed in a path. We simulated a scenario where the distance range between two BSs is $[400, 600)$ meters with the same relay density as above. 50 random pairs of BSs were chosen and up to 9 more hops were allowed beyond the minimum number required to connect each pair. We make a simple assumption that each relay adds 0.5 ms delay. We first note that, if a network provider can afford a few extra relays (+3) and the additional delay is tolerable to the applications using the network, throughput can be improved to almost 12 Gbps within this range, which is about a 20% improvement compared to Fig. 6. Beyond 3 additional relays, however, throughput does not improve due to bottleneck effects, while delay continues to increase. Thus, we expect that using a number of relays close to the minimum number needed for connectivity will be the best choice in practice.

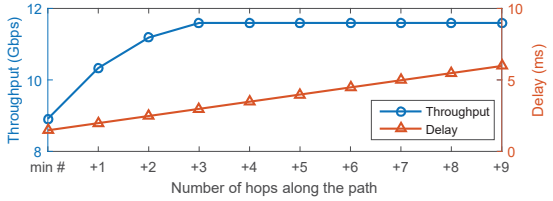

Fig. 7. Average throughput and estimated delay as hop number increases.

We also investigated the impact of the density of candidate relay locations on the results. Surprisingly, we found that the maximum throughput and minimum number of hops hardly varied over a fairly wide range of candidate relay location densities and, hence, we do not show the results herein. The distance of links connecting different relays on the same two buildings are similar. Thus, the results seem to show that as long as different paths use the same sequence of buildings, the throughputs will be similar. Due to the large number of buildings in the simulated urban area, one relay per building surface already produces good results. In the simulated area, one relay per surface corresponds to 1100 candidate relay locations, which is already a large number.

B. Impact of optimal scheduling

To investigate the benefit of finding the path with maximum throughput among all available paths using our optimal scheduling algorithm, we ran a brute-force version of the relay selection algorithm to obtain all possible interference-free paths. Only ten pairs of BSs were simulated in each distance category, due to the high complexity of finding all possible paths. Since the number of available paths in the minimum number of hops case is relatively small, to have a better sample size, we studied the minimum number plus one hop case, where thousands to millions of paths can be found in the scenarios that we simulated. The results are shown in Fig. 8. Note that the error bars in this figure represent the minimum gain and maximum gain, whereas the error bars in other figures of this paper represent the standard deviation.

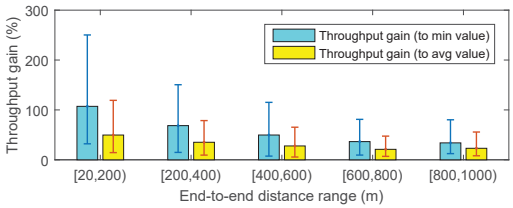

Fig. 8. Throughput gain due to optimal scheduling. ($\delta_r = 0.001/m^2$)

Fig. 8 shows that the throughput gain from our optimal scheduling algorithm can exceed 200% when the BS separation is less than 200 meters, which reflects the large dynamics of link length in this case. On average, the throughput gain compared to the maximum throughput of the worst path ranges from 33% to 107%, and the gain compared to the maximum throughput of an average path ranges from 23% to 49%. Clearly, the use of optimal scheduling to guide path selection can provide large improvements compared to random paths satisfying the maximum hop and link length constraints.

C. Comparison between 3D and 2D deployments

To investigate whether consideration of all 3 dimensions of the urban environment are really necessary, we also simulated a 2D scenario. To do this, we maintained the same height for all BS locations and all candidate relay positions. In order to provide a fair comparison, we kept the total number of candidate relay positions the same for each building surface in the 3D and 2D cases. Note that this effectively produces a higher density of possible relay positions for the 2D scenario since the relay positions are confined to a smaller area. For

both 3D and 2D scenarios, we simulated 100 sets of BS pairs with separations in the same 5 ranges as in the earlier simulations. The throughput and minimum number of hops required in both scenarios are shown in Fig. 9.

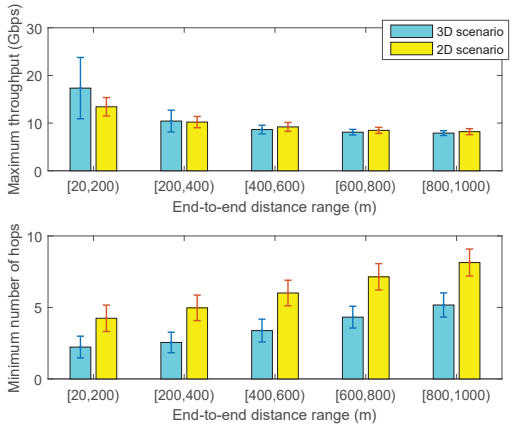

Fig. 9. Minimum number of hops required and the corresponding maximum end-to-end throughput in 3D and 2D scenarios. ($\delta_r = 0.0010/m^2$)

Fig. 9 shows that when BS separation is small ($< 400m$), the throughput in the 3D scenario is better than that in 2D. For other separations, the throughputs are comparable but the 2D throughput is actually slightly higher than in the 3D case. However, the critical difference between the two scenarios is in the minimum path length shown in the bottom part of Fig. 9. The figure shows that in 2D networks, the number of relays needed per path is approximately double the number needed for the 3D case. Deploying relays in 2D will therefore significantly increase hardware and deployment costs while also producing substantially higher delays, as compared to our proposed 3D deployment. As we will see in the next subsection, increasing path length also has a very negative impact on running time of the optimal scheduling algorithm.

D. Running time of relay path selection algorithm

Fig. 10 shows the CPU time taken by our path selection algorithm on an Intel 3.5 GHz i-7 processor with 16 GB of RAM on a single core for 3 representative source, destination BS pairs with separation in the range of 400 to 600 meters. Recall that the time complexity is $O(D^M)$, where D is the maximum degree of the candidate relay location graph and M is the maximum number of hops. Although, in theory, the polynomial degree increases with M, we see that in some cases, the running time does not increase significantly with M. Here, the minimum number of hops is either 2 or 3, so M can be as large as 12. In spite of this large M, the running time for the red triangle plot is less than 1 second for all cases and the running time for the yellow square plot peaks at around 10 seconds. As discussed in Section VI-B, this is because, in many cases, the minimum schedule length determined with a smaller maximum number of hops can eliminate a very large number of paths from consideration when doing the depth-first search with the maximum number of hops equal to $|M|$. There are some BS pairs, however, where the running time does increase substantially with M, e.g. the blue circle plot. Even for this case, however, the running time is only a few minutes with the largest M. Since the relay path selection algorithm will be run in the network planning stage, a few minutes running time is acceptable.

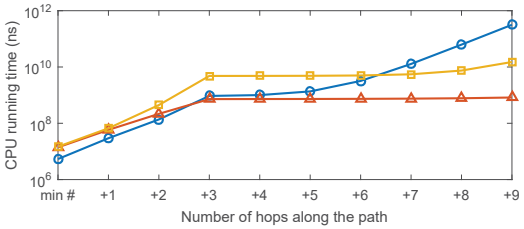

Fig. 10. Running time of path selection alg. vs. max. no. of hops

ACKNOWLEDGEMENT

This research was supported in part by the National Science Foundation through Award CNS–1513884.

REFERENCES

[1] T.L. Frey, "The Effects of the Atmosphere and Weather on the Performance of a mm-Wave Communication Link," *Applied Microwave and Wireless,* Vol. 11, pp. 76–81, 1999.

[2] M. Alicherry, R. Bhatia, and L. Li, "Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks," *Proc. Mobicom,* pp. 2196–2211, 2005.

[3] G. Brar, D.M. Blough, and P. Santi, "Computationally efficient scheduling with the physical interference model for throughput improvement in wireless mesh networks," *Proc. Mobicom,* pp. 2–13, 2006.

[4] J. Luo, C. Rosenberg, and A. Girard, "Engineering wireless mesh networks: joint scheduling, routing, power control, and rate adaptation," *IEEE/ACM Transactions on Networking,* 18.5 (2010): 1387–1400.

[5] G. Zheng, et al., "Toward Robust Relay Placement in 60 GHz mmWave Wireless Personal Area Networks with Directional Antenna," *IEEE Transactions on Mobile Computing* 15.3 (2016): 762-773.

[6] Y. Niu, et al., "Blockage robust and efficient scheduling for directional mmWave WPANs," *IEEE Trans. Vehicular Tech.,* 64.2 (2015): 728-742.

[7] J. Qiao, et al. "Efficient concurrent transmission scheduling for cooperative millimeter wave systems," *Proc. GLOBECOM,* pp. 4187–4192, 2012.

[8] Y. Zhu, et al., "Demystifying 60GHz outdoor picocells," *Proc. Mobicom,* pp. 5–16, 2014.

[9] C. Perfecto, et al., "Beamwidth optimization in millimeter wave small cell networks with relay nodes: A swarm intelligence approach," *Proc. European' Wireless,* pp. 1–6, 2016.

[10] S. Singh, et al., "Tractable model for rate in self-backhauled millimeter wave cellular networks," *IEEE Journal on Selected Areas in Communications,* 33.10 (2015): 2196–2211.

[11] B. Xie, Z. Zhang, and R. Hu, "Performance study on relay-assisted millimeter wave cellular networks," *Proc. Vehicular Technology Conference,* pp. 1–5, 2016.

[12] J. Du, et al., "Gbps user rates using mmWave relayed backhaul with high gain antennas," *IEEE Journal on Selected Areas in Communications,* Vol. 35, No. 6, pp. 1363–1372, 2017.

[13] 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. Networking*, 19.5 (2011): 1513–1527.

[14] E. Violette, et al., "Millimeter-wave propagation at street level in an urban environment," *IEEE Transactions on Geoscience and Remote Sensing* 26.3 (1988): 368–380.

[15] X. Ge, et al., "5G wireless backhaul networks: challenges and research advances," *IEEE Network,* 28.6 (2014): 6–11.

[16] B. Hajek and G. Sasaki, "Link scheduling in polynomial time," *IEEE transactions on Information Theory,* 34.5 (1988): 910–917.