Interference-aware routing and spectrum allocation for millimeter wave backhaul in urban picocells

Maryam Eslami Rasekh*, Dongning Guo†, Upamanyu Madhow‡
Department of Electrical and Computer Engineering, University of California Santa Barbara
Email:{*merasekh, ‡madhow}@ece.ucsb.edu
Department of Electrical Engineering and Computer Science (EECS), Northwestern University
Email: †dGuo@northwestern.edu

***Abstract*—The exponential growth in demand for mobile data requires significant increases in spatial reuse, motivating an evolution towards picocellular architectures with densely deployed base stations. Providing backhaul for such a network is a key challenge, because of the high access link rates, and the cost and difficulty of running optical fiber to base stations that might be opportunistically placed on lampposts and rooftops. Wireless backhaul using millimeter (mm) wave spectrum is therefore an attractive and flexible approach, given the plentiful availability of spectrum and the possibility of synthesizing highly directive, steerable links. In this paper, we formulate and investigate the problem of joint resource allocation and routing on such a mm wave backhaul network, providing a linear programming formulation that accounts for mutual interference across simultaneously active links. While the number of variables grows exponentially in network size, it is possible to prune the problem size so that it is manageable for moderately sized networks. Numerical results and design implications are briefly discussed.**

I. Introduction

Picocellular architectures are a natural response to the capacity crisis faced by mobile data operators, especially in high-density urban environments. LTE promises to deliver peak rates of hundreds of Mbps in a picocell, reaching up to Gigabits per second with carrier aggregation. In next generation networks, such peak rates can be increased further, by an order of magnitude or more, by using millimeter (mm) wave spectrum (e.g., the 60 GHz unlicensed spectrum) directly from pico base station to the mobile [1], assuming that significant challenges due to blockage and mobility can be overcome. Indeed, a recent interference analysis for such networks [2] indicates that capacities of the order of terabits per second per kilometer (along a single urban canyon) can be obtained with only a few GHz of spectrum, because of the aggressive spatial reuse (e.g., 20 meter base station spacing) enabled by the highly directive links that can be synthesized in mm wave bands. Moreover, this capacity roughly adds up across parallel canyons, given the relative isolation provided by building blockage (mm waves are easily blocked by even small obstacles, given the small wavelength). However, utilizing such over-the-air capacity for end-to-end applications requires that the pico base stations have a sufficiently high-capacity connection to the Internet. For opportunistic picocellular deployments on lampposts and rooftops, it is unrealistic to expect optical fiber connectivity for each base station, so that wireless backhaul becomes a natural choice [3]–[5].

In this paper, we consider a mesh network with highly directive mm wave links as a means of providing backhaul for picocellular networks. The nodes in the network are the base stations (additional relay nodes could be added, but are not considered here), with a subset of nodes being gateways to the Internet. The objective is to route traffic to/from the base stations to an appropriately chosen (typically nearest) gateway, such that each picocell can support a given level of demand on the access links to mobile devices. We assume that directional antennas are used on each backhaul link, but there is enough residual interference that these links cannot be treated as wires. Thus, backhaul time-frequency resources must be allocated in an interference-aware fashion.

We formulate this here as a joint routing and resource allocation problem, whose goal is to maximize the access rate at the base stations, accounting for the mutual interference between simultaneously active links. Our framework applies to any mesh backhaul network, but in general, interference patterns depend on the propagation geometry in a complicated fashion. However, mm wave mesh networks in urban environments, as considered here, present an easier exercise in modeling: The highly directive nature of the links, and the ease of blockage of mm waves by buildings, imply that each link interferes with only a few others. Here we present numerical results for an example network derived from mapping an urban propagation geometry based on part of Manhattan, New York City to the corresponding mesh network.

Our optimization framework borrows ideas from recent work on optimization for HetNets and cognitive networks, which allocates resources to base stations by optimizing over all possible resource allocation patterns (i.e., all possible subsets of simultaneously active base stations) [6]–[8]. Blowing up the number of variables in the optimization problem in this fashion (for N base stations, there are $2^N - 1$ possible allocation patterns) leads to convex optimization problems with a unique global optimum, and even large problems can be solved with the advanced algorithms and computational tools available today. Another key observation from these papers is that Caratheodory's theorem guarantees that the number of resource allocation patterns actually employed in the optimal solution is at most N, which is not only practical, but also opens the way for the use of sparse optimization techniques.

We adapt these ideas to our setting by considering an L-link mesh network, and optimizing over all $2^L - 1$ possible link activation patterns, subject also to flow constraints at all nodes but the gateway node. Millimeter wave networks in urban environments are sparse (due to directivity and link blockage), hence the number of links is only a small multiple of the number of nodes. Analogous to prior work, Caratheodory's theorem can be shown to guarantee the existence of an optimal solution with at most N activation patterns having nonzero allocations. Moreover, simple heuristics can prune out a very large number of inadmissible patterns to reduce the computational complexity. Our formulation, which seeks to maximize the access rate while minimizing the latency, is actually simpler than prior work on HetNets (e.g., [9]), and results in a linear programming problem rather than a convex optimization problem. In contrast to a convex optimization problem, a linear programming problem does not necessarily have a unique global optimum so we are not guaranteed to find the minimal allocation. This is often the case for the sparse interference present in our model, however we find that by perturbing the interference matrix by a small amount a sparse allocation is obtained. We loosely map these results to rough guidelines on backhaul and access link speeds which might be used in the short and long term.

II. System Model

The system considered here is a network of nodes in an urban environment, connected to each other through highly directive mm wave links. Each node is co-located with a picocellular base station, hence links are short (of the order of 100 meters or less). Each node is connected to one or more neighboring nodes through directional mm wave wireless links. We assume that each such link is associated with a separate transceiver. Thus, a node can transmit to multiple neighbors at a time, and can receive from multiple neighbors at a time. However, we do not allow a node to transmit and receive at the same time: bandwidth in the mm wave band is plentiful enough that there is no need to take on the technical challenges of full duplex operation, at least in the first generation of mm wave backhaul.

A fraction of the nodes in the network are gateway nodes with high-speed wired connectivity to the core network. We assume gateway nodes are distributed evenly throughout the network, and that each node uses the closest gateway node (in terms of number of hops) to route its data. Thus, each gateway supports a subset of the network consisting of the nodes in its vicinity, and the network is divided into "clusters" centered at gateway nodes. Since data transfer does not happen between clusters, and due to the relatively local confinement of interference, clusters can operate more or less independently of each other. Therefore, we assume here that the problem can be reduced to allocation and scheduling for a single cluster, rather than for the entire network, which can be large. However, quantifying the effect

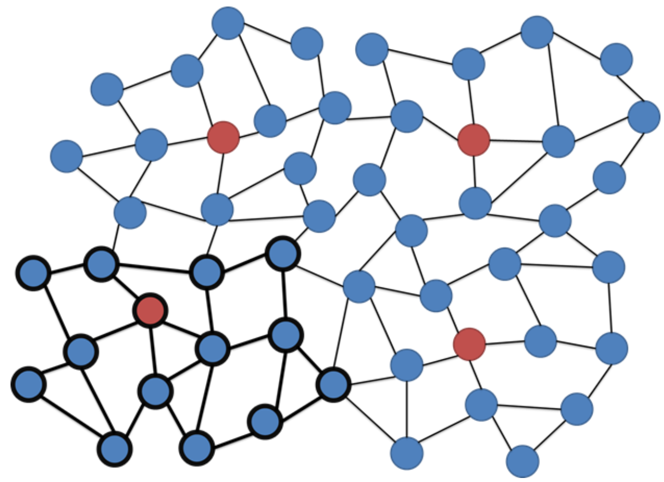

Fig. 1: Clusters and gateways in the backhaul network

of inter-cluster interference, and refining the optimization problem to account for it, is an important topic for future work. Fig. 1 demonstrates the decomposition of the larger mesh into smaller clusters around gateways. Note that in this figure each line between two nodes represents two links, one in each direction.

The purpose of the backhaul network is to deliver data from/to gateway nodes to/from regular nodes through multi-hop paths. The resource allocation problem can be solved separately for downlink and uplink data, hence we focus on supporting downlink data (which is typically of much higher volume), without loss of generality. The goal, then, is to calculate allocations that route downlink data to nodes providing them with the largest data rate possible.

Due to the high directionality of antennas, the interference graph in such networks is sparse. In urban environments, street canyons are shielded from each other due to building blockage, but links aligned along the street can mutually interfere when they are transmitting in the same direction. In this case, the receiving beam of one link is aimed at the transmitting beam of another, and both antennas amplify the undesired signal. Even links that are not precisely aligned in this fashion can interfere with each other, and the interference can be predicted by tracing the most significant rays from transmitters to receivers, including the line-of-sight and reflected paths. Such interference calculations then feed into the resource allocation problem formulated in the next section.

Given that backhaul links are quasi time-invariant, and the sparse multipath fading on such links can be effectively combated with relatively small link margins using spatial and frequency diversity techniques [10], we assume that the data rate for a link is governed only by the signal-to-interference-plus-noise ratio (SINR) corresponding to a given set of simultaneously active links, which can be computed based on the interference graph. These SINR computations, together with the constraint on not simultaneously transmitting and receiving, are the main modeling ingredients in the optimization formulation in the next section.

Our aim is to assign the available resources to links, taking into account their interference behavior, such that the backhaul data rate provided for each node is maximized. In the next section, this problem is formulated as a linear program (LP).

III. Problem Formulation

We consider a cluster with L links, denoted as a set $\{1, \ldots, L\}$. It is assumed that the time-frequency resources are homogeneous on the timescale that resource management is concerned with. For simplicity, it is also assumed that a transmitter applies fixed flat power spectral density regardless of the bandwidth allocation.[1] Thus the interference condition and the SINR of all links over any slice of time-frequency resource depends only on the subset of links that are active over the given resource. Let γ_l denote the received signal-to-noise ratio (SNR) of link l in the absence of interference. Let $I_{k\to l}$ denote the interference-to-noise ratio caused by link k on link l if both are active.

We express all of our results in terms of spectral efficiency, which can be translated to data rates simply by multiplying by the bandwidth. The spectral efficiency of link l under pattern P ($l \in P$) is

$$r_{l,P} = \log\left(1 + \frac{\gamma_l}{1 + \sum_{k\in P\setminus\{l\}} I_{k\to l}}\right). \quad (1)$$

There are $2^L - 1$ non-empty subsets of links, which are referred to as *link activation patterns* or simply *patterns*. The problem boils down to allocating all available time-frequency resources to those patterns in order to maximize the data rate delivered to each node. Let x_P denote the fraction of time-frequency resources allocated to subset $P \subset \{1, \ldots, L\}$. Evidently,

$$\sum_{P\subset\{1,\ldots,L\}} x_P = 1. \quad (2)$$

Let N denote the number of nodes in the cluster, and set node 1 to be the gateway node. Let d denote an average "network-level" spectral efficiency parameter: node i sees a date rate of $\alpha_i B d$, where α_i is a relative allocation parameter (specified by the system designer) and B is the backhaul bandwidth. Note that $\alpha_i \equiv 1$ corresponds to delivering the same data rate to each picocell, and this is what we assume in our numerical examples later. However, allowing α_i to vary across i gives us the ability, for example, to repurpose backhaul resources to address hotspots. Since we work with spectral efficiencies, the backhaul bandwidth B does not enter into the formulation of the optimization problem below. We use the terms data rate and spectral efficiency interchangeably in the discussions in this section.

For simplicity, we decouple resource allocation for downlink and uplink data, and focus on downlink data (which constitutes the bulk of data traffic) for concreteness. The downlink data is delivered through nodes over a multi-hop path starting from the gateway (node 1). Thus, in order to ensure that the rate delivered to each node can sustain the outgoing data, as well as the data rate consumed by the node itself, it is required that, for each node $i \geq 2$ (i.e., for each node other than the gateway, which is the source) we have:

$$\sum_{P\subset\{1,\ldots,L\}} \left(x_P \sum_{l\in P} r_{l,P} f_{li} \right) \geq \alpha_i d, \quad i = 2, 3, \ldots, N \quad (3)$$

where

$$f_{li} = \begin{cases} +1 & \text{link } l \text{ runs into node } i \\ -1 & \text{link } l \text{ runs out of node } i \\ 0 & \text{link } l \text{ is not connected to node } i. \end{cases} \quad (4)$$

Our first attempt at characterizing the optimal allocation of $\{x_P\}$ is via the following problem:

$$\underset{\{x_P\}}{\text{maximize}} \quad d \quad (5a)$$
$$\text{subject to} \sum_{P\subset\{1,\ldots,L\}} x_P = 1 \quad (5b)$$
$$\sum_{P\subset\{1,\ldots,L\}} x_P \sum_{l\in P} r_{l,P} f_{li} - \alpha_i d \geq 0, \quad (5c)$$
$$i = 2, \ldots, N.$$

The problem can be expressed in standard matrix form by defining the variable 2^L-vector:

$$X = [x_1, x_2, \ldots, x_{2^L-1}, d]^T \quad (6)$$

and constant vectors of compatible lengths:

$$b = [1,\, 0,\, \cdots\, ,\, 0]^T \quad (7)$$
$$c = [0, 0, \ldots, 0, 1] \quad (8)$$

as well as

$$A_{(N\times 2^L)} = \begin{bmatrix} 1 & \cdots & 1 & 0 \\ \ddots & & & \alpha_2 \\ & -A_{ij} & & \vdots \\ & & \ddots & \alpha_N \end{bmatrix} \quad (9)$$

where, for $i = 2, \ldots, N; j = 1, \ldots, 2^L - 1$,

$$A_{ij} = \sum_{l\in P_j} r_{l,P_j} f_{li} \quad (10)$$

where P_j denotes the j-th subset in the power set of $\{1, \ldots, L\}$. Problem (5) is then equivalent to the following:

$$\text{maximize} \ \ cX \quad (11a)$$
$$\text{subject to} \ \ AX \leq b. \quad (11b)$$

The solution to this linear optimization problem presents an optimal allocation that maximizes delivered data rates, d, subject to flow conservation. The solution is not unique, which is not by itself a drawback. More importantly, however, the formulation in (11) imposes no restraint on latency and power preservation, and can result in allocations with unnecessarily long paths from nodes to the gateway, especially in

[1] Choosing which portion of the spectrum to use is equivalent to on-off power control in the frequency domain. More refined power control is left for future work.

networks with sparse interference graphs, as is often the case for highly directive mm wave links.

In order to limit path lengths in the allocation while maintaining the linearity of the problem, we now add a second term to the objective that implicitly accounts for delay. Consider the total data rate of each link (sum of x_P on all subsets containing that link) as a unit of delay for that portion of data. We can therefore decrease the net delay in the network by minimizing the sum of link data rates. To this end, we modify the objective as follows:

$$\text{maximize} \ \ (c - \lambda s)X \tag{12}$$

where cX represents the throughput and the vector s is chosen so that sX is the sum of the data rate transmitted on all links. In other words, the k-th element of s is the sum of the data rate of all links active in the k-th subset of $\{1, \ldots, L\}$ subject to interference from other active links in that subset:

$$s_k = \sum_{l \in P_k} r_{l,P_k}, \quad k = 1, \ldots, 2^L - 1. \tag{13}$$

While we wish to decrease the number of link activations and the delay, we do not wish to sacrifice throughput. For this reason, the weighting factor, λ, must be sufficiently small to ensure that priority is given to throughput over delay. Even if $\lambda > 0$ is small, the corresponding term in the objective function does ensure that we obtain the smallest net delay among all possible allocations providing the optimal throughput. In order to obtain insight into how small λ must be, consider two allocations X_1 and X_2. We require a guarantee that:

$$d_1 > d_2 \Rightarrow (c - \lambda s)X_1 > (c - \lambda s)X_2 \tag{14}$$

or, equivalently,

$$d_1 - d_2 = cX_1 - cX_2 > \lambda(sX_1 - sX_2). \tag{15}$$

Now note that by increasing d from d_2 to d_1, a total of $(N-1)(d_1 - d_2)$ units of added data rate must be routed from the gateway to the nodes. Due to multihop routing, this added load can show up on several links in the network, increasing the net data rate on each link. An upper bound on this additional load is obtained by assuming that the extra load is added to the load carried by all L links in the network. This implies that the net link data rate, sX, is increased by $(N-1)L(d_1 - d_2)$. In order to guarantee (15), λ must satisfy (16):

$$\lambda < \frac{1}{(N-1)L} \tag{16}$$

in order to ensure priority of throughput maximization over net link activation.

In the presented formulation, the size of the problem grows exponentially with L, so that even in a single cluster, the problem size can quickly become unmanageable. We therefore add a pruning step that dismisses all subsets that correspond to simultaneous transmit and receive on any node. This greatly decreases the problem size, and allows us to scale to larger networks. For instance, the number of subsets for a network with 22 links (11 two-way links) considered in our numerical results (see Fig. 5) is reduced from $2^{22} - 1 \approx 4$ million subsets to around 10,000.

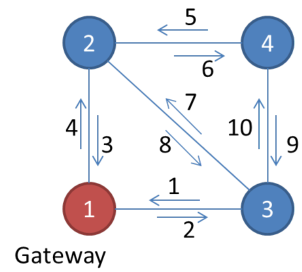

Fig. 2: Topology of small network

Link index	1	2	3	4	5	6	7	8	9	10	%
Active subsets	0	**1**	0	**1**	0	0	0	0	0	0	53.21
	0	**1**	0	0	0	**1**	**1**	0	0	0	10.89
	0	**1**	0	0	0	**1**	0	0	0	0	10.89
	0	0	0	**1**	0	0	0	**1**	0	**1**	10.89
	0	0	0	**1**	0	0	0	0	0	**1**	10.89
	0	0	0	0	0	**1**	0	0	0	**1**	3.21
Total link activation (%)	0	75	0	75	0	25	10.89	10.89	0	25	

Fig. 3: Resource allocation solution for the network of Fig. 2 with objective defined in (11)

IV. Numerical Results

For simplicity, we set the SNR for each link in the network in the absence of interference to a fixed value of 10 dB (our formulation allows for variation of SNR across links, as would be needed to accommodate variations in range). Since we have restricted attention to a single cluster, we pick the value of SNR to be relatively small in order to account for unmodeled interference from other clusters. We assume that the data rate demand is uniform across nodes ($\alpha_i \equiv 1$). We present our results in terms of spectral efficiency, but translate this to some example data rates towards the end of this section. For reference, the link-level spectral efficiency corresponding to 10 dB SNR is 3.46 bits per second per Hz (bps/Hz).

We begin with presenting results for the simple topology in Fig. 2. We choose this example to highlight the contrast between our final objective function (12), which implicitly

Link index	1	2	3	4	5	6	7	8	9	10	%
Active subsets	0	**1**	0	**1**	0	0	0	0	0	0	57.02
	0	**1**	0	0	0	**1**	0	0	0	0	17.98
	0	0	0	**1**	0	0	0	0	0	**1**	17.98
	0	0	0	0	0	**1**	0	0	0	**1**	7.02
Total link activation (%)	0	75	0	75	0	25	0	0	0	25	

Fig. 4: Resource allocation solution for the network of Fig. 2 with objective defined in (12)

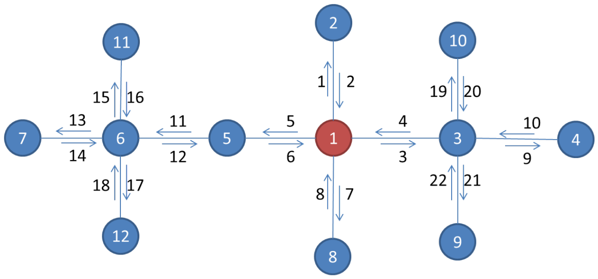

Fig. 5: Topology of larger cluster with several aligned links

Link #	3	4	5	6	9	10	11	12	13	14
3	0	∞	0	∞	∞	0	0	**0.2**	0	**0.125**
4	∞	0	∞	0	0	∞	0	0	0	0
5	0	∞	0	∞	0	**0.2**	∞	0	0	0
6	∞	0	∞	0	0	0	0	∞	0	**0.2**
9	∞	0	0	**0.2**	0	∞	0	**0.125**	0	**0.08**
10	0	∞	0	0	∞	0	0	0	0	0
11	0	**0.2**	∞	0	0	**0.125**	0	∞	∞	0
12	0	0	0	∞	0	0	∞	0	0	∞
13	0	**0.125**	**0.2**	0	0	**0.08**	∞	0	0	∞
14	0	0	0	0	0	0	0	∞	∞	0

Fig. 6: Part of the interference matrix of the topology shown in Fig. 5 containing aligned links

Without interference modeling		With interference modeling	
P	x_P (%)	P	x_P (%)
5	11.11	3,5	33.33
3,5	11.11	11	22.22
5,9	11.11	1,11	11.11
3,11	22.22	3,5,13,15,17	7.75
3,7,11	11.11	3,5,7,13,15,17	3.36
5,13	11.11	11,19	11.11
1,11,19	11.11	5,19,21	3.36
5,15,17,21	11.11	5,7,9,13,21	7.75
Expected: **d = 0.3844 bps,** **d/r_{link} = 11.11%** **Achieved:** **d = 0.2264 bps,** **d/r_{link} = 6.54%**		**d = 0.3844 bps,** **d/r_{link} = 11.11%**	

Fig. 7: Allocation and throughput for the network of Fig. 5, (a) with and (b) without aligned LoS interference

penalizes delay, and the initial objective function (11). While both solutions achieve the same network-level spectral efficiency, the latter gives the allocation depicted in Fig. 3 in which links 7 and 8 are activated in a portion of the resources while these links are not on the shortest path between the gateway and any node and effectively only move data from node 2 to 3 and back. These redundant link activations are a result of ignoring multihop delays in the objective. On the other hand, solving for the objective defined by (12), the resulting allocation is zero for all but four subsets of links. The non-zero subsets are shown in Fig. 4. Note that the links that are not on the shortest paths from the gateway to the nodes are not activated. The network-level spectral efficiency obtained by each node is $d = 1.73$ bps/Hz, which is half the value of the maximum link-level spectral efficiency of 3.46 bps/Hz.

We now consider a larger example with nontrivial interference, as shown in Fig. 5. In this network, many of the links are aligned and cannot rely on antenna directivity for isolation. The normalized interference between aligned links, shown in Fig. 6, shows the portion of the interference graph corresponding to the aligned links, with entries corresponding to the normalized interference between the links. The relative interference levels are assumed to be proportional to squared distances, plus 3 dB as a pessimistic margin to account for variations due to multipath fading. (It has been shown in prior work [11] that 3 dB is more than enough margin to account for fading in the sparse multipath characteristic of mm wave backhaul channels, with suitably designed frequency and/or spatial diversity.) For reference, we also consider a setting where the links do not interfere with each other (essentially treating each link as a wire), so that the only constraint is that a node cannot send and receive at the same time.

We only report now on the final objective function (12) with implicit delay penalty. Due to the sparsity of the interference matrix in both cases, the optimal solution to (12) is not unique and we may arrive at allocations that activate a large number of subsets. To prevent these large allocations, small random perturbations are added to the interference matrix to discard the extra optimums and the problem is solved using the perturbed interference matrix to obtain a minimal result. The exact solution to the unperturbed problem can then be found by optimizing the original problem only over the subsets activated in the perturbed case instead of the entire space. In the optimal allocation, the highest interference level seen by active links is 0.222. This corresponds to SINR of 4.9 dB and link-level spectral efficiency of 2.04 bps/Hz. Many active links see no interference and operate at SNR of 10 dB and spectral efficiency of 3.46 bps/Hz.

In Fig. 7 we see the optimal allocation and resulting throughput for the network with and without aligned link interference, observing no performance loss relative to the interference-free case when optimizing for interference. This is due to the structure of the network: links closer to the gateway that bear a higher load do not receive interference from other links in the downlink path, and hence do not suffer rate loss. On the other hand, interference from highly loaded links on lightly loaded links (e.g., from link 5 on link 13) can be compensated for by increasing the bandwidth used by link 13, since it transfers the load of only one node, and can therefore operate a low spectral efficiency.

Neglecting interference, however, can lead to loss of throughput for some nodes. As shown in Fig. 7, an interference-free allocation, when applied to a network with interference, fails to deliver the optimal throughput to all nodes. For some nodes (node 7 in this case), throughput decreases from 0.3844 bps/Hz to nearly half that amount, 0.2264 bps/Hz. Therefore, while aligned interference does not adversely affect the achievable performance of the network in this particular situation when modeled correctly and

Fig. 8: Clusters in a typical mesh backhaul in an urban environment mapped from part of New York City. Triangles and circles represent gateway non-gateway nodes respectively.

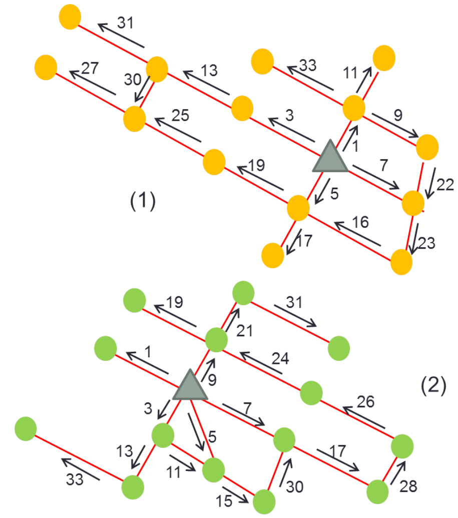

Fig. 9: Link numbering in the clusters shown on New York map

optimized for, ignoring it in resource allocation leads to poor performance.

In order to model a more realistic network, we now consider a portion of New York City shown in Fig. 8. We assume that the stations represented by triangles have access to fiber backhaul and act as gateways, while the circles are the added picocells that rely on the mesh for backhaul. By choosing the nearest gateway for routing, nodes are divided into clusters around gateways. Two of these clusters have been highlighted in Fig. 8 and are considered for analysis. The interference matrix for these clusters is obtained based on the network geometry, and the allocation problem is solved for each cluster independently assuming uniform traffic at nodes. The resulting allocations are reported in Fig. 10, with the link indexes displayed in Fig. 9. We see that in both clusters a considerable fraction (more than 17%) of the raw link data rate is delivered to each node, indicating the efficacy of wireless mesh backhaul for relatively small clusters.

Example data rates: In order to support a peak data rate of R_{cell} in a picocell, we need a backhaul bandwidth of R_{cell}/d. Thus, for our example cluster in Fig. 5, LTE picocell data rates of up to 500 Mbps can be accommodated with less than 2 GHz of backhaul link bandwidth, which easily fits into a single channel in the unlicensed 60 GHz band. The peak data rates on the backhaul links at 10 dB SNR would be of the order of 5 Gbps in this case. On the other hand, if we now consider 60 GHz links to the mobile with picocell data rates of 2 Gbps, then we need backhaul bandwidths of the order of 5-6 GHz, with peak backhaul link data rates of the order of 10-20 Gbps, requiring us to reach into higher mm wave bands.

P (cluster 1)	x_P (%)	P (cluster 2)	x_P (%)
1,5,7,13,16	7.17	1,5,9,13,17	4.35
1,3,7,19,22	3.36	1,5,9,11,13,17	8.70
1,3,19,23	0.14	9,13,15,17	13.04
3,9,19,23	0.44	3,15,17,19,21	8.70
1,3,17,19,23	17.25	1,5,9,11,13,17,24	4.35
1,3,5,7,25	6.87	3,5,9,17,26	4.35
1,5,7,13,25	13.89	3,7,15,21,26	4.35
1,5,7,13,27	3.22	3,15,17,21,26	8.70
1,5,7,13,16,27	17.10	13,15,17,21,26	4.35
1,3,5,7,25,30	3.36	3,7,15,19,21,28	8.70
1,3,5,7,22,25,30	3.51	3,5,7,9,28,30	30.43
3,9,19,23,31	6.43	1,5,9,13,17	4.35
3,9,11,19,23,31	17.25	1,5,9,11,13,17	8.70
d = 0.5968 bps **d/r_{link} = 17.25%**		**d = 0.6016 bps** **d/r_{link} = 17.39%**	

Fig. 10: Allocation and throughput for the two clusters shown in Fig. 9

V. Conclusions

We have shown that interference-aware resource allocation and routing on a mm wave backhaul can be reduced to a relatively tractable linear program. While the number of link activation patterns grows exponentially with network size, the problem can be reduced to manageable proportions for moderate-sized networks by pruning link activation patterns based the constraint on not simultaneously transmitting and receiving. The results that we obtain indicate that backhaul using a single channel in the 60 GHz unlicensed band, with peak link data rates of the order of 1-2 Gbps, would suffice to support dense deployment of LTE picocells. On the other hand, picocells using Gbps 60 GHz links to the mobile require backhaul link data rates of the order of 10-20 Gbps (e.g., using mm wave spectrum at 100+ GHz).

In future work, we wish to extend the optimization framework. Our current formulation, which penalizes delay while maximizing throughput, does lead to the elimination of long paths, but the optimal solution is still not unique, and may have more active patterns than strictly necessary. We find that small random perturbations in the interference matrix help reduce the number of active patterns significantly, but it is of interest to explore the structure of the LP solution in more detail to formalize this, and to explore alternate mechanisms for promoting sparsity. Another important topic is to develop a systematic framework for handling inter-cluster interference, without an excessive increase in computational complexity.

VI. Acknowledgments

This research was supported in part by the National Science Foundation under grants CCF-1423040, CNS-1317153, and CNS-1518812, and gifts from Facebook and Futurewei Technologies.

References

[1] Y. Zhu, Z. Zhang, Z. Marzi, C. Nelson, U. Madhow, B. Y. Zhao, and H. Zheng, "Demystifying 60ghz outdoor picocells," in *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '14. New York, NY, USA: ACM, 2014, pp. 5–16. [Online]. Available: http://doi.acm.org/10.1145/2639108.2639121

[2] Z. Marzi, U. Madhow, and H. Zheng, "Interference analysis for mm-wave picocells," in *IEEE Globecom 2015*, December 2015, to appear.

[3] S. Singh, M. Kulkarni, A. Ghosh, and J. Andrews, "Tractable model for rate in self-backhauled millimeter wave cellular networks," *IEEE J. Sel. Areas Commun.*, 2015, early Access. [Online]. Available: http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7110547

[4] K. Zheng, L. Zhao, J. Mei, M. Dohler, W. Xiang, and Y. Peng, "10 Gb/s HetSNets with millimeter-wave communications: access and networking – challenges and protocols," *IEEE Commun. Mag.*, vol. 53, no. 1, pp. 222–231, 2015. [Online]. Available: http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7010538

[5] R. Taori and A. Sridharan, "Point-to-multipoint in-band mmwave backhaul for 5g networks," *IEEE Commun. Mag.*, vol. 53, no. 1, pp. 195–201, 2015. [Online]. Available: http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7010534

[6] B. Zhuang, D. Guo, and M. L. Honig, "Traffic-driven spectrum allocation in heterogeneous networks," *IEEE J. Sel. Areas Commun. Special Issue on Recent Advances in Heterogeneous Cellular Networks*, vol. 33, no. 10, pp. 2027–2038, 2015. [Online]. Available: http://arxiv.org/abs/1408.6011

[7] Q. Kuang, W. Utschick, and A. Dotzler, "Optimal joint user association and resource allocation in heterogeneous networks via sparsity pursuit," *CoRR*, vol. abs/1408.5091, 2014. [Online]. Available: http://arxiv.org/abs/1408.5091

[8] B. Zhuang, D. Guo, and M. L. Honig, "Energy-efficient cell activation, user association, and spectrum allocation in heterogeneous networks," *submitted to IEEE J. Sel. Areas Commun. Special Issue on Energy-Efficient Techniques for 5G Wireless Communication Systems, under revision*, 2015.

[9] O. Semiari, W. Saad, Z. Dawy, and M. Bennis, "Matching theory for backhaul management in small cell networks with mmwave capabilities," *CoRR*, vol. abs/1501.02410, 2015. [Online]. Available: http://arxiv.org/abs/1501.02410

[10] M. E. Rasekh, H. Zhang, and U. Madhow, "Diversity for sparse multipath channels: design of robust highly directional outdoor 60 ghz links," *unpublished*.

[11] E. Torkildson, H. Zhang, and U. Madhow, "Channel modeling for millimeter wave mimo," in *Information Theory and Applications Workshop (ITA), 2010*. IEEE, 2010, pp. 1–8.