Optimizing mmWave Wireless Backhaul Scheduling

Edgar Arribas, Antonio Fernández Anta, Dariusz R. Kowalski, Vincenzo Mancuso,
Miguel A. Mosteiro, Joerg Widmer, and Prudence W. H. Wong

Abstract—Millimeter wave (mmWave) communication not only provides ultra-high speed radio access but is also ideally suited for efficient and flexible wireless backhauling. Specifically for dense deployments, a mmWave macro base station (MBS) that serves a large number of mmWave micro base stations (μBSs) is much more cost effective than legacy cellular architectures which connect μBSs to the core network through fibers. In addition, μBSs can cooperate with each other by acting as relay nodes. The directional nature of mmWave communication allows for spatial reuse, even in the presence of interference, which can be exploited to optimize mmWave wireless backhaul performance. The optimization opportunistically prioritizes the use of good connections at the MBS and further leverages compact and concurrent transmissions between μBS. Relays and directional antennas speed up communication, but increase the complexity of the scheduling problem. In this work, we study the mmWave backhaul scheduling problem and derive an MILP formulation for it as well as upper and lower bounds. We prove that the problem is NP-hard and can be approximated, but only if interference is negligible. By means of numerical simulations, we compare theoretical results with heuristics in small system sizes. Results validate the analysis and demonstrate the high performance of our heuristics in realistic cellular settings.

---

1 INTRODUCTION

NEXT generation cellular networks aim to increase wireless data rates to gigabits per second and reduce latencies to milliseconds or less [1]. Such networks will use a flexible architecture to reliably cope with an ultra high density of devices, up to tens or hundreds of devices per square meter. Millimeter wave (mmWave) communication on frequencies from 6 to 300 GHz is an extremely interesting technology to address these challenges. The unprecedented vast amount of available spectrum and the possibility of integrating antenna arrays with a high number of antenna elements for very directional communication allow for multi-gigabit link speeds and excellent spatial reuse [2]. At the same time, communication range is a key issue. mmWave signals are very vulnerable to shadowing and exhibit high frequency-related attenuation, which has to be accounted in the network architecture design. mmWave communication has been proposed for wireless backhauling of small cells as well as the actual radio access. It is particularly well suited for backhauling in extremely dense cell deployments, where other backhaul technologies are cost intensive [3].

In this context, mmWave can be used to build a wireless backhaul among a macro base station (MBS) and a large number of micro base stations (μBSs) spaced a few tens of meters apart, thus extending the coverage of the MBS. At the same time, the highly directional nature of mmWave
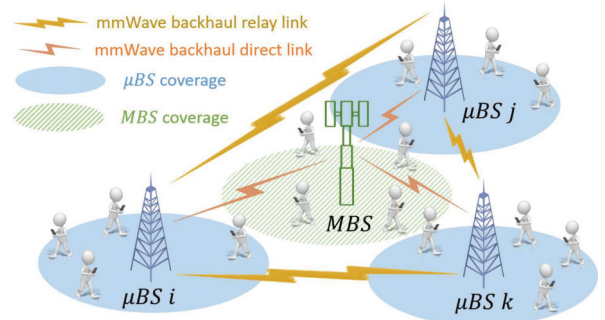

Fig. 1. Reference scenario: mmWave backhaul network.

makes cooperative relaying among μBSs much more useful than at lower frequencies, where interference offsets much of the potential gains. The rationale behind this approach is that the MBS is typically the bottleneck in the network, and it is therefore beneficial to offload traffic from the MBS to a nearby μBS as fast as possible. This μBS can then relay the traffic to the destination μBSs, while at the same time the MBS can already forward more traffic to the next suitable μBS. Such a relay schedule improves spatial reuse and reduces the overall time taken to distribute the traffic to the destination μBSs. The MBS can even have multiple RF chains—the electronic device used to transmit/receive radio signals—and so communicate with more than one μBS in parallel. Fig. 1 illustrates this approach, which is the reference scenario used in this paper. Since it can be adapted at millisecond time scales and includes costs and advantages of beamsteering in the loop, the mmWave relay case is very different from other relay optimization problems studied in the literature and involving, e.g., WLAN, cellular and satellite networks [4] or free-space optical links [5].

The possibility of relaying and the availability of multiple antenna elements make possible communication speed-ups as described above, but also increase the complexity of

- E. Arribas is with IMDEA Networks Institute and Universidad Carlos III de Madrid (UC3M), Madrid, Spain. E-mail: edgar.arribas@imdea.org.
- A. Fernández Anta, V. Mancuso, and J. Widmer are with IMDEA Networks Institute, Leganés (Madrid), Spain. E-mail: {antonio.fernandez, vincenzo.mancuso, joerg.widmer}@imdea.org.
- D. R. Kowalski is with Augusta University, USA and SWPS University of Social Sciences and Humanities, Warsaw, Poland. E-mail: dariusz.kowalski@swps.edu.pl.
- M. A. Mosteiro is with Pace University. E-mail: mmosteiro@pace.edu.
- P. W. H. Wong is with University of Liverpool, UK. E-mail: pwong@liverpool.ac.uk.

scheduling data for delivery, as one needs to choose whether to relay or not, to which μBSs, and which MBS links must be used. Therefore, understanding whether scheduling data delivery is NP-hard in this context, and if so, which approximations can be guaranteed even in the limit, is fundamental to gain insight on practical challenges such as scalability. Also, in such scenario, finding out which heuristics perform well, and for which system sizes, is crucial for practical purposes. In this work, we carry out such studies as follows.

Given a collection of data to deliver to a set of μBSs, we study the mmWave relay optimization problem of minimizing the time to complete the delivery, i.e. the makespan[1], in a network managed by an MBS. We consider both the case of interference-free links and the case of more realistic transmissions in the presence of directional cross-link interference. We call such optimization problem mmWave Backhaul Scheduling (MMWBS). Solving the problem results in a compact concurrent relaying schedule of links, which flexibly and opportunistically reuses mmWave resources over the backhaul links. However, (re-)configuring mmWave links brings with it a beam training and steering overhead that needs to be taken into account to implement a scheduling strategy that works efficiently at packet level.

Plenty of work has been done in scheduling communications in related models, including different wireless and optical networks. For instance, the work in [6] applies to wireless networks of arbitrary topology, but link activation cost, interference, and concurrent communication through multiple outgoing links are not taken into account. Even if communication models differ only in minor aspects, problems may be entirely different [7], [8], [9], [10]. To the best of our knowledge none of these solutions apply to our setting.

Roadmap. We first present an overview of the most relevant related work in Section 2. Then, we summarize the contribution and main findings of our work in Section 3 and present the system model in detail in Section 4. We formulate the problem in Section 5 and analyze it in Section 6. We discuss the design of heuristics in Section 7 and report on performance evaluation in Section 8 through numerical simulation. Finally, we discuss the lessons learnt in Section 9, and summarize and conclude the paper in Section 10.

2 RELATED WORK

The use of relays in cellular networks has been proposed to extend cellular and ad-hoc/WLAN coverage and improve user throughput. Many proposals focus on the use of orthogonal relay resources, to not interfere with the direct communication link. Authors of [11] and [12] discuss how to implement and optimize cellular relay with opportunistic features, using legacy 802.11 and LTE bands, while authors of [13] apply the D2D paradigm to mmWave relays in the 60 GHz band. They only use simple heuristics and model interference without considering beamforming gains due to steerable antennas used for mmWave.

While relays on sub-6 GHz bands cause and suffer from significant interference due to their omnidirectional transmissions, the directionality of mmWave antennas mitigates

interference, especially in backhaul systems [14], [15]. Multiple links can be active simultaneously as long as their beams do not overlap. This motivates our backhaul interference model which is based on the *protocol model* [16]. We assume that two links can be scheduled in the same time slot as long as a receiver does not experience interference from a significant side or main lobe of the other transmitter, i.e., only if the strongest interference is weak enough. As widely discussed in [17], this is a very accurate interference model for mmWave backhaul systems with strong directionality. Furthermore, such a simple model is almost as accurate as a more complex SINR-based one, while being mathematically tractable. Recent works have also applied similar models. The authors of [15] optimize routing in mmWave cellular networks by means of maximizing spatial reuse thanks to the limited interference from remaining transmissions. In [18], the authors apply D2D features to access and backhaul mmWave networks so to exploit the directionality of mmWave antennas and mitigate interference from multiple sources, and in [19], energy efficiency of mmWave backhaul networks is studied in order to establish concurrent flows such that interference between beams is lower than a threshold. Therefore, the interference model considered in this paper is based on realistic assumptions and is commonly used in the literature.

There is also a large body of work in optimization, scheduling, and relay selection in OFDMA/TDD cellular networks and WLANs using the same resources for direct transmissions and for relay [20], [21], [22], [23], [24], [25] including several that propose dynamic TDD algorithms designed to exploit the new LTE-B enhanced Interference Mitigation and Traffic Adaptation (eIMTA) capabilities [26], [27], [28]. Authors of [29] compute the capacity of a relay-assisted OFDMA-based cellular network, in which interference is the main limiting factor. However, they do not account for the intrinsic characteristics of mmWave transmission to model and optimize transmission qualities and interferences. Indeed, the existing works focus mainly on scheduling and ICIC under the assumption of an interference-limited regime, whereas we assume constant interference due to directional isolation. We also take centralized scheduling, unlike works on distributed MAC schemes for mesh networks, as in [30]. Also, while we base our analysis on average rate values, stochastic geometry analysis of self-backhaul mmWave networks has recently been shown in [31] as well as a scaling law analysis in [32]. However, our formulations and analysis are novel and allow to shed light on the tradeoff between compact parallelism of mmWave link scheduling and establishment costs.

The use of mmWave for backhauling small cells in a dense cellular environment enables cost-effective and flexible replacement of the expensive and time-consuming deployment of fiber for gateway access. As discussed in [33], [34], the IEEE 802.11ay amendment, which is the successor to IEEE 802.11ad, includes several modifications that make mmWave suitable for wireless backhauling, among other use-cases. IEEE 802.11ay includes new techniques such as channel bonding and aggregation, non-uniform constellations and enhanced beamforming training that enable peak rates of tens of gigabits per second and allow to build high-speed wireless backhaul networks. Authors in [35]

---

1. The makespan is the time needed to complete the delivery of all files. Hence, the makespan corresponds to the elapsed time until all files arrive at their destinations.

introduce mmWave backhaul for heterogeneous networks, in which they use joint scheduling and resource allocation schemes based on spatial-division multiple access. In contrast to our approach, they maximize the flow throughput while selecting which paths the content should follow from the BS to the user, without relaying among APs. In [18], the authors also propose a joint transmission scheduling scheme for radio access and wireless backhaul using mmWave D2D communication, where the decision is whether to use a backhaul path or transmit locally among D2D users in case of sufficient proximity. However, although APs relay content through mmWave links, the routes are predetermined by some criterion, instead of minimizing delivery time. Finally, authors in [3] design a mmWave framework for wireless backhaul where flows can follow multiple paths or be served concurrently between two devices. They aim to maximize the aggregated transmission rate, different from us.

Besides cellular networks and WLAN-like approaches, relay has been studied for satellite and mixed satellite-terrestrial communication networks using GHz bands. In that framework, the relay is typically considered as a tool for spreading the information broadcast by a satellite [36], [37], [38]. In case of small satellites on low orbits, inter-satellite links have been proposed for relay with directional links. However, the main issue studied in that case is the short life of inter-satellite links, which drastically reduces relay opportunities and requires the adoption of delay-tolerant solutions [4]. Thereby, solutions proposed for scheduling of satellite communications and relay cannot be used for fast reconfigurable mmWave relay, where there are neither broadcast elements nor connectivity time lapses.

3 SUMMARY OF MAIN RESULTS

We start modeling the problem as a Mixed-Integer-Linear Program (MILP) to obtain some preliminary insight on MMWBS (cf. Section 5). Then, we pursue a more advanced study in Section 6 obtaining the following main results.

- We show that the combination of interference with the possibility of relaying makes the problem very hard, proving in Theorem 1 that not even an approximation to the optimal makespan of MMWBS with interference can be guaranteed in the worst case.
- We also show that, even without interference, MMWBS is NP-hard in Theorem 2.
- Knowing from Theorems 1 and 2 that from a theoretical standpoint we can only aim for an approximation to the optimal MMWBS schedule in interference-free channels, we find it in Section 6.3. We present Algorithm 2 to compute such schedule and we provide theoretical guarantees of the approximation in Theorem 3. The above results combined expose the challenges of MMWBS.
- Theorem 3 also upper bounds the makespan of MMWBS. We establish another upper bound in Observation 1 for the natural schedule that routes all data without relaying, using only one RF chain (cf. Section 6.4).
- By formulating a simplified version of MMWBS in a Linear Program and using other mathematical argumentations, we prove lower bounds on the makespan of MMWBS in Section 6.5. We summarize our theoretical upper and lower bounds in Table 1.

TABLE 1
Summary of makespan upper and lower bounds.

Inter-ference	Makespan $\|\mathcal{S}\|$	cf.
Yes	$\|\mathcal{S}\| \geq \left\lceil \alpha + \frac{\sum_{j \in R} d_j}{C} \right\rceil$ $C = \max_{\substack{R' \subseteq R: \|R'\| \leq K: \\ \forall a,b \in R': I_{(s,a),(s,b)}=1}} \sum_{i \in R'} c_{(s,i)}$	Fact 1
IND[2]	$\|\mathcal{S}\| \geq d + \max\left\{0, \left\lceil \frac{n' - \|D\| - (d(d-1)/2)K}{\|D\|} \right\rceil\right\}$ $d = \lceil \|D\|/K \rceil$, $D \subseteq R$ μBSs receiving directly from MBS.	Lemma 1
IND	$\|\mathcal{S}\| \geq \left\lceil \sqrt{\frac{1}{4} + 2\left(\frac{n'}{K} + 1\right)} - \frac{1}{2} \right\rceil$ $n' \leq n$ destination μBSs.	Thm. 4
Yes	$\|\mathcal{S}\| \leq \sum_{\substack{i \in R, \\ d_i \neq 0}} \left\lceil \alpha + \frac{d_i}{c_{(s,i)}} \right\rceil$	Obs. 1
No	$\|\mathcal{S}\| \leq \sum_{\substack{i \in R: \\ t_{si} > 0}} \lceil \alpha + t_{si} \rceil + \frac{3}{2}\left(\left\lceil \frac{T}{1-\alpha} \right\rceil + \sqrt{3 \left\lceil \frac{T}{1-\alpha} \right\rceil}\right)$ T and $\{t_{si}\}_{i \in R}$: as given by the LP of Fig. 4.	Thm. 3
No	$\frac{\|\mathcal{S}\|}{T_{OPT}} \leq \left(K + \frac{3}{2}\right)\left(\frac{1}{1-\alpha} + \frac{1}{T_{OPT}}\right) + \frac{3}{2}\sqrt{3\left(\frac{1}{1-\alpha} + \frac{1}{T_{OPT}^2}\right)}$ T_{OPT}: optimal makespan.	Thm. 3

- Leveraging the insight gained from the analysis for worst case scenarios, we design simple yet effective heuristics for MMWBS (cf. Section 7).
- Finally, we carry out realistic numerical simulations to compare the optimization, the theoretical bounds, and the heuristic approximations (cf. Section 8). The experimental evaluation shows that, on average and for small testable systems, these heuristics find near-optimal solutions, both with and without interference.

Bounds in Table 1 correspond to the following. The lower bound in Fact 1 shows the minimum time needed to deliver all data through the fastest interference-free links that can be active simultaneously. The lower bounds in Lemma 1 and Theorem 4 correspond to the minimum time taken by link activations even maximizing parallelism. The first is existential (corresponds to any given fixed schedule), whereas the second is universal. The upper bound in Observation 1 corresponds to delivering all data without relaying, using one MBS link at a time, and the upper bounds in Theorem 3 are makespan and approximation guarantees for Algorithm 2.

In our experiments, heuristics perform better than the constant-approximation schedule of Algorithm 2, which can be used only in absence of interference. Nevertheless, these heuristics perform better for typical settings (as we try to model for simulations) whereas theoretical results give provable guarantees that good solutions are possible for all scenarios (even if they seem to be worse by a constant factor than the proposed heuristics when run in typical settings). Moreover, we give the function of how performance scales up asymptotically (cf. Theorem 3). In other words, heuristics and Algorithm 2 are of independent interest.

All in all, our theoretical and experimental results show that compact concurrent relaying is a powerful tool for wireless backhauling in mmWave networks.

2. Acronym IND stands for "Independent".

4 MODEL

We consider a backhaul system formed by an MBS s and a set $R = [1, n]$ of n static μBSs that may act as relays for s. We denote the set of nodes in the mmWave backhaul network as $V = R \cup \{s\}$, and the set of links as $E = V \times R$.

Although the main potential feature of mmWave links is the directional communication and interference mitigation for spatial reuse, it has been experimentally observed [39], [40] that commercial beam-patterns offer geometries where transmissions may potentially interfere in some regions of the space, as depicted in Fig. 2. Such beam-patterns may have non-negligible sidelobes with high power, which indeed spoil the received signals from the μBSs positioned in the direction of such sidelobes. We model such effect in our theoretical framework by introducing an interference between transmissions via pairs of links—which could be set up arbitrarily—as follows. The binary parameter $I_{\ell_1, \ell_2} \in \{0, 1\}$, known by the MBS, tells a priori if links ℓ_1 and ℓ_2 can be active simultaneously. Our interference model is a particular case of conflict graphs used in previous works (e.g. [41], [42]), and it is justified by the fact that an active mmWave link is very sensitive to even small interference from other nearby transmissions [2]. Therefore, the binary parameter I_{ℓ_1, ℓ_2} states that links ℓ_1 and ℓ_2 cannot be active simultaneously in case their transmission beams interfere.

Time is slotted, and the capacity of each link ℓ is given as the number of bits that can be sent in one time slot, denoted as c_ℓ. Also this quantity is known by the MBS. Each link ℓ has a cost of activation $0 < \alpha < 1$ modelling the portion of a slot used to activate a link (antenna steering delay, potential preamble, and header overhead). Once active, a link can be used during any number of consecutive slots without incurring further activation cost. In fixed backhaul systems—as the one discussed in this paper—the activation cost related to beam training can be saved since link end points are static. However, changing the configuration of the phased antenna array to the known setting for a link still requires a short but non-zero time. Hence, α remains positive, which is relevant for our theoretical analysis. In addition, novel designs are being considered that have non-negligible activation cost. In [43], authors test a proof-of-concept of mmWave phase shifters with miniaturized liquid crystal. Further, in [44], liquid crystal polymers are proposed as an efficient solution for future flexible 5G mmWave devices. Whereas such new antenna designs have higher activation time, they have desirable features such as higher gains and better beam shapes, that improve performance. Hence, it is important to analyze the whole framework and the system performance with positive values for the activation cost α.

We assume that each μBS, i.e., each node in R, has one RF chain, so that they can only communicate with one other node in V, and only in one direction, in any time slot. On the MBS side, we assume that up to K links from the MBS can be active in the same time slot, where $K > 0$ is a parameter. That is, the MBS can communicate with several μBSs simultaneously, leveraging complicated architectures with multiple antennas and RF chains. This is to avoid a bottleneck at the MBS for file deliveries, given that the MBS is the orchestrator of all traffic arriving to the wireless backhaul network. We assume that channel state information
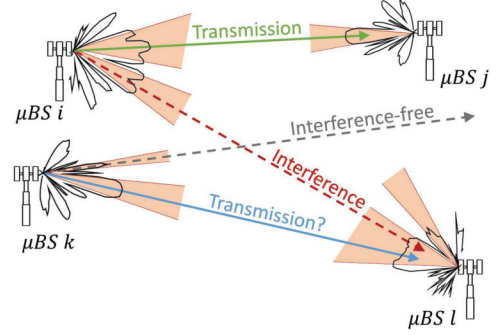

Fig. 2. Reference scenario for the interference model.

from each pair of links is available at the MBS, following the standards defined by the 3GPP in Release 13 [45]. Hence, the network uses the sidelink transmission mode in which the network orchestrator, i.e., the MBS, manages the resources for sidelinks, as well as schedules transmissions according to the channel feedback received from μBSs control channels. Retransmissions are handled by the MAC procedure, such as the HARQ protocol specified in the 3GPP Release 8 [46], or any other available procedure for physical resource access as CSMA/CA protocols like 802.11ad.

For each μBS $r \in R$, there is a certain amount $d_r \geq 0$ of data (in bits) whose destination is r stored at the MBS s. This data corresponds to the downlink traffic for the mobile terminals associated with that μBS, and the objective is to route it to r as quickly as possible. To this end, the MBS can send the data d_r over the direct link $\ell = (s, r)$ or via an indirect path. Path lengths are limited to two hops, i.e., there can be at most one intermediate relay r', resulting in a path $\{(s, r'), (r', r)\}, r' \in R \setminus \{r\}$. We leave the study of multi-hop relay to future research, as here we focus on unveiling the potential of relaying in its simplest form. We consider that R is split into two disjoint sets, $R = R^R \dot{\cup} R^D$, where R^R is the set of μBSs that can relay (and may have their own data to receive) and R^D is the set of μBSs that are only destinations.

We define MMWBS first informally: Given a source MBS, a set of μBSs, the links between them, the interference setting, and a collection of data to deliver from source to destinations within the model described above, find a schedule of communication so that all the data is routed from the source MBS to the destination μBSs in the smallest number of time slots. The length of such a schedule is called the makespan. In this paper, we use $\mathcal{S}$ to indicate a schedule and $|\mathcal{S}|$ for its makespan.

A formal definition of the problem, its input (communication network, the interference parameter, and collection of data) and output (a schedule of links usage) is given in the following section as an MILP.

5 MILP FOR MMWBS

We formulate the decision problem of MMWBS as an MILP. Given an integer T, we want to decide whether or not there is a communication schedule of length T slots such that all the data is routed from the source to the destinations. Accordingly, our MILP has only a set of constraints, i.e., we do not require a utility function. Hence, we search for the minimum T such that the MILP has a feasible solution. We provide now a description of the MILP. Recall that in our

model, c_ℓ is the capacity of link ℓ in bits per time slot, d_r is the amount of bits with destination μBS r, and α is the link activation cost. We use the following decision variables:

- $d_{ir}(t)$ is the number of bits for destination $r \in R$ stored in node $i \in V$ at the beginning of time slot t, that is, $d_{ir}(t)$ is an integer number such that $d_{ir}(t) \geq 0$.
- $f_{\ell r}(t)$ is the fraction of the capacity of link ℓ used to send bits for destination $r \in R$ during time slot t, that is, $f_{\ell r}(t)$ is a real number such that $0 \leq f_{\ell r}(t) \leq 1$.

We now define the constraints on such variables imposed by the parameters. We start with data-flow constraints. All data is initially at the source, (Eqs. (1)–(2)), the data stored in each μBS after each time slot is updated considering the fractions of capacities used, (Eqs. (3)–(5)), and all the data must be delivered to the corresponding μBSs after T slots (Eqs. (6)–(7)). Namely, $\forall r, i \in R, \forall j \in R^R \backslash \{r\}, \forall t \in [1, T]$, we have:

$$d_{sr}(1) = d_r; \tag{1}$$

$$d_{ir}(1) = 0; \tag{2}$$

$$d_{sr}(t+1) = d_{sr}(t) - f_{(s,r)r}(t)\,c_{(s,r)} - \sum_{\substack{u \in R^R \\ u \neq r}} f_{(s,u)r}(t)\,c_{(s,u)}; \tag{3}$$

$$d_{jr}(t+1) = d_{jr}(t) - f_{(j,r)r}(t)\,c_{(j,r)} + f_{(s,j)r}(t)\,c_{(s,j)}; \tag{4}$$

$$d_{rr}(t+1) = d_{rr}(t) + f_{(s,r)r}(t)\,c_{(s,r)} + \sum_{u \in R^R} f_{(u,r)r}(t)\,c_{(u,r)}; \tag{5}$$

$$d_{sr}(T+1) = 0; \tag{6}$$

$$d_{rr}(T+1) = d_r. \tag{7}$$

For convenience, we define a binary variable that indicates when a link is active in a given time slot:

- $a_\ell(t) \in \{0, 1\}$ is an indicator variable such that $a_\ell(t) = 1$ when link ℓ is active during time slot t.

For any link $\ell \in E$ and time slot $t \in [1, T]$,

$$a_\ell(t) \geq \sum_{r \in R} f_{\ell r}(t); \tag{8}$$

where Eq. (8) ensures that $a_\ell(t) = 1$ if some data flows through link ℓ during time slot t (recall that $f_{\ell r}(t)$ are fractions of the capacity of ℓ, hence the aggregated flow is at most 1, as constrained in Eq. (10)). Now, we constrain the link activations of each time slot t in order not to interfere among themselves, according to the binary interference parameter I_{ℓ_1, ℓ_2}. For all links $\ell_1, \ell_2 \in E$ and $\forall t \in [1, T]$:

$$(1 - I_{\ell_1, \ell_2}) \cdot (a_{\ell_1}(t) + a_{\ell_2}(t)) \leq 1. \tag{9}$$

Eq. (9) ensures that, in case two links ℓ_1, ℓ_2 cannot be used in the same time slot t due to interference, i.e., $I_{\ell_1, \ell_2} = 0$, there exists a constraint that permits to activate only one link in each time slot. In case links ℓ_1, ℓ_2 do not interfere, such constraint does not exist. Since the binary parameter I_{ℓ_1, ℓ_2} is an input of the problem, Eq. (9) is linear.

Finally, we constrain the decision variables so that the schedule obtained does not violate link capacities, including the overhead for link activation (Eq. (10)) and the maximum number of simultaneously active links (Eqs. (11)–(12)), due to the nodes' RF chains. For any link $\ell \in E$, μBS $r \in R$, and time slot $t \in [1, T]$ we have:

$$\sum_{r \in R} f_{\ell r}(t) \leq 1 - \alpha \cdot (1 - a_\ell(t-1)); \tag{10}$$

$$a_{sr}(t) + \sum_{j \in R} (a_{rj}(t) + a_{jr}(t)) \leq 1; \tag{11}$$

$$\sum_{j \in R} a_{sj}(t) \leq K. \tag{12}$$

- Input parameters:
  T: Number of time slots.
  d_r: Amount of bits with destination $r \in R$.
  c_ℓ: Capacity of link ℓ, given as the number of bits that can be sent in one time slot.
  α: Cost of activation relative to one time slot, $\alpha \in]0, 1[$.
  I_{ℓ_1, ℓ_2}: Binary interference parameter that states if links ℓ_1, ℓ_2 can be active concurrently.
  K: Number of RF chains at the MBS.

- Decision variables:
  $d_{ir}(t)$: number of bits for destination $r \in R$ stored in node $i \in V$ at the beginning of time slot t.
  $f_{\ell r}(t)$: fraction of the capacity of link ℓ used to send bits for destination $r \in R$ during time slot t.
  $a_\ell(t)$: indicator variable such that $a_\ell(t) = 1$ when link ℓ is active during time slot t.

Set of constraints, $\forall r, i \in R, \forall j \in R^R \backslash \{r\}, \forall t \in [1, T]$:

- Data-flow constraints:
  $d_{sr}(1) = d_r;$
  $d_{ir}(1) = 0;$
  $d_{sr}(t+1) = d_{sr}(t) - f_{(s,r)r}(t)\,c_{(s,r)} - \sum_{\substack{u \in R^R \\ u \neq r}} f_{(s,u)r}(t)\,c_{(s,u)};$
  $d_{jr}(t+1) = d_{jr}(t) - f_{(j,r)r}(t)\,c_{(j,r)} + f_{(s,j)r}(t)\,c_{(s,j)};$
  $d_{rr}(t+1) = d_{rr}(t) + f_{(s,r)r}(t)\,c_{(s,r)} + \sum_{u \in R^R} f_{(u,r)r}(t)\,c_{(u,r)};$
  $d_{sr}(T+1) = 0;$
  $d_{rr}(T+1) = d_r.$

- Activation constraints:
  $a_\ell(t) \geq \sum_{r \in R} f_{\ell r}(t);$
  $(1 - I_{\ell_1, \ell_2}) \cdot (a_{\ell_1}(t) + a_{\ell_2}(t)) \leq 1.$

- Link capacity constraints:
  $\sum_{r \in R} f_{\ell r}(t) \leq 1 - \alpha \cdot (1 - a_\ell(t-1));$
  $a_{sr}(t) + \sum_{j \in R} (a_{rj}(t) + a_{jr}(t)) \leq 1;$
  $\sum_{j \in R} a_{sj}(t) \leq K.$

- Range constraints:
  $d_{ir}(t) \in \mathbb{Z}_{\geq 0}, \quad \forall r \in R, \forall i \in V, \forall t \in [0, T]$
  $f_{\ell r} \in [0, 1], \quad \forall l \in V \times R, \forall r \in R, \forall t \in [1, T];$
  $a_\ell \in \{0, 1\}, \quad \forall \ell \in V \times R.$

Fig. 3. MILP to solve the decision version of the mmWave Backhaul Scheduling (MMWBS) problem.

In Fig. 3 we present the formulation of the decision problem described in this section. Please note that we have formulated the decision version of the problem. Hence, the goal is to decide whether it is possible to route the data in T time slots, so that there is no utility function.

6 THEORETICAL ANALYSIS

In this section, we present our theoretical study of MMWBS. We show first that, in face of interference, MMWBS cannot be approximated (Theorem 1). That is, no algorithm to compute a MMWBS schedule can have a guaranteed approximation to the optimal makespan in the worst case. The natural question that follows is what is the complexity of MMWBS without interference. We answer such question showing that even in such simpler scenario MMWBS is NP-hard (Theorem 2). In other words, we can only aim for approximations to the optimal schedule without interference. We present such algorithm (cf. Algorithm 2) and prove its approximation in Theorem 3. All these results combined expose the intrinsic challenges of solving optimally MMWBS. Nevertheless, lower and upper bounds on the makespan are of interest. Formulating a simplified version of MMWBS as a Linear Program and using other mathematical argumentations, we prove various lower bounds. For comparison,

we also establish upper bounds for a schedule without relaying or spatial reuse, and we bound the makespan of our approximation algorithm (Theorem 3). The details follow.

6.1 MMWBS cannot be approximated

We show here that the MMWBS Problem (with interference) cannot be approximated. We do so by reducing the maximum clique problem [47] to MMWBS, as follows.

Theorem 1. *For all $\varepsilon > 0$, approximating the MMWBS problem to within $\sqrt{n^{1-\varepsilon}/2}$ is NP-hard.*

Proof. Consider an instance $G = (W, E_W)$ of the maximum clique problem with $|W| = n$ nodes. We build an instance of MMWBS from G as follows. The activation cost is $\alpha \leq 0.5$. The set of μBSs will be $R = W \cup \{d\}$, where $R^R = W$ and $R^D = \{d\}$. The MBS can send to up to $K = n$ μBSs simultaneously. The links from the MBS s to the μBSs in R^R interfere with each other as follows: For every $i, j \in W$, links $\ell_i = (s, i)$ and $\ell_j = (s, j)$ have $I_{\ell_i, \ell_j} = 1$ if and only if $(i, j) \in E_W$. The rest of links can only be used alone, i.e., for each link $\ell \in \{(s, d)\} \cup \{(i, j) : i, j \in R\}$, $I_{\ell, \ell'} = 0$, for all ℓ'. This interference setting guarantees that the only parallelism in the communication can come from the MBS sending to the nodes in R^R, and there can be as many simultaneous transmissions as the size C of the maximum clique in G.

There is a single file F to be sent from the MBS to d of size $f > n^2$. The links from the MBS s to the μBSs in R^R have capacity 1. The link (s, d) has extremely low capacity. The links from the μBSs in R^R to d have capacity $\frac{f}{1-\alpha}$.

It can be seen that the optimal makespan of this instance of MMWBS is $T^* = \alpha + f/C + C$, where C is the size of the maximum clique in G. The schedule uses $\alpha + f/C$ time slots[3] to get the file F to C nodes of R^R. Then, these μBSs send their corresponding portion of file to d in the next C time slots. Observe that d cannot receive anything while the file is being sent by s given the interference between links.

Let us assume there is an algorithm A with polynomial time complexity that can approximate MMWBS to within a factor $\rho = \sqrt{n^{1-\varepsilon}/2}$, for some $\varepsilon > 0$. Then, the algorithm finds a value $T \in [T^*/\rho, \rho T^*]$. Applied to the above instance of MMWBS, the obtained approximation satisfies $T \in [(f/C+C)/\rho, \rho(f/C+C)]$. Since the size of any clique is at most n, we have that $C < f/C$, and hence $T \in \left(\frac{f}{\rho C}, \frac{2\rho f}{C}\right)$. This implies that $C \in \left(\frac{f}{\rho T}, \frac{2\rho f}{T}\right)$. Then, we can use A to obtain an approximation of C to within $\frac{2\rho f}{T} / \frac{f}{\rho T} = 2\rho^2 = n^{1-\varepsilon}$, which is not possible unless $P = NP$ [47]. $\square$

6.2 NP-hardness

In this section, we prove that the decision version[4] of the MMWBS problem is NP-hard, even if there is no interference and $K = 1$[5]. The proof is via a reduction from the partition problem with equal cardinality (PEC): Given $2n$ natural numbers $a_1, a_2, \ldots, a_{2n}$ such that $\sum_{1 \leq i \leq 2n} a_i = 2B$, the question is whether there exists a partition into two subsets

---

3. We disregard ceiling and floors for simplicity.
4. That is, whether there exists a schedule of a given makespan or not.
5. In terms of the MILP of Section 5, these assumptions imply that Restriction 9 disappears, and Restriction 12 becomes $\sum_{j \in R} a_{sj}(t) \leq 1$.

---

of n numbers each, such that the sum of each subset is exactly B. Let us consider that $a_{\min} = \min_{1 \leq i \leq 2n}\{a_i\} \geq \frac{B}{n+1}$.[6] This problem is NP-hard [48].

Theorem 2. *The decision version of the MMWBS problem is NP-hard for any value $0 < \alpha < 1$ of the link activation cost α, even if there is no interference and $K = 1$.*

Proof. Given an instance $\mathcal{I}$ of the PEC problem, we construct an instance $\mathcal{I}'$ for MMWBS as follows. There is the MBS s and a set $R = R^R = \{v_0, v_1, \ldots, v_{2n}\}$ of $2n + 1$ μBSs. The link capacities are $c_{(s,v_0)} = \frac{B}{1-\alpha}$; $\forall 1 \leq i \leq 2n : c_{(s,v_i)} = c_{(v_0,v_i)} = 1$; and $\forall 1 \leq i, j \leq 2n : c_{(v_i,v_j)} = 0$. For all $1 \leq i \leq 2n$, $d_i = a_i$ and $d_0 = 0$. The decision problem instance $\mathcal{I}'$ is whether it all data can be sent in $T = B+n+1$ time slots.

If there is a solution for $\mathcal{I}$ then we can create a solution for $\mathcal{I}'$ by routing B data to one subset of n μBSs via v_0 and B data to the other subset of n μBSs from s directly. The communication from s to v_0 takes $\alpha + \frac{B}{B/(1-\alpha)} = 1$ time slots. The remaining communication can take place in parallel and requires $B + n$ time slots. (The communication of either v_0 or s with v_i requires $\lceil a_i + \alpha \rceil = a_i + 1$ time slots.) Therefore, the makespan of the schedule is $B+n+1$, which is exactly T.

We now consider the other direction. Assume that there is a solution to $\mathcal{I}'$. Observe that this solution must use μBS v_0 as relay to allow parallel communications, otherwise, the makespan would be $2B+2n\lceil \alpha \rceil = 2B+2n > T$, because the a_i numbers are natural numbers and link capacities are 1, thus, at least one extra slot is needed for each link activation. Communication from s to v_0 takes at least $\lceil \alpha \rceil = 1$ time slots just to activate the link. We claim that s must send to n μBSs via v_0 and to the other n directly. Otherwise, either s or v_0 sends to at least $n+1$ μBSs. Then, communicating with these μBSs would require at least $(n+1)(a_{\min}+1) = a_{\min}(n+1)+n+1 \geq B+n+1$ slots, since $a_{\min} \geq \frac{B}{n+1}$, and the makespan would be at least $B+n+2 > T$. This means that s sends to n μBSs via v_0 and to n μBSs directly. Finally, both v_0 and s send to their respective n μBSs B bits of data each. Otherwise, since both together send $2B$ data, one of them, say s, sends more than B bits. Let $i_1, i_2, \ldots, i_n$ be the μBSs served by s directly. Then, $\sum_{j=1}^{n} a_{i_j} > B$. Since a_{i_j} is a natural number, sending to μBS i_j takes $\lceil a_{i_j} + \alpha \rceil = a_{i_j} + 1$ slots. Hence, the makespan would be $1 + \sum_{j=1}^{n}(a_{i_j}+1) > T$. $\square$

6.3 Constant-approximation schedule for MMWBS without interference

In this section, we present an algorithm to obtain a schedule for the MMWBS problem that achieves a constant approximation of the optimal makespan when links do not interfere with each other. Recall from Theorem 2 that this especial case of the MMWBS problem is still NP-hard, and that in the presence of interference even approximating the optimal makespan is NP-hard (cf. Theorem 1).

In Algorithm 1 we show the `Direct Download` schedule. `Direct Download` computes a schedule that uses up to K RF chains without relaying but taking into consideration interference, and will be used in Algorithm 2.

The first step of our algorithm is to solve the LP in Fig. 4, removing the restrictions that take interference into account (i.e., the line labelled (*)), which will be used later.

---

6. Otherwise, we can simply add $B - (n + 1)a_{\min}$ to each value a_i.

Algorithm 1 (`Direct Download`)

We define the schedule inductively in time, as follows.

For time slot $t = 1$, we select the fastest subset (up to size K) of links from the MBS (according to their downlink rate) that do not interfere. We do this one link at a time, from fast to slow, to avoid combinations. The selected set is activated in time slot $t = 1$ to download as much data as possible, while having into consideration the cost of activation of the links, α.

Then, for each time slot $t > 1$, we select first the links to μBSs that have received partially their files before t. Since they were not interfering among themselves in $t-1$, they neither interfere in t.

Additionally, we select the fastest subset of links to μBSs that did not start their download, up to a maximum K links counting the ongoing downloads, and restricted to non-interfering links (again one by one to avoid combinations). For these newly selected μBSs we consider the cost of activation of the link, α.

Finally, in case there is some μBS out of the MBS-coverage, this μBS downloads its file through one relay, without spatial reuse.

The last time slot when some link is active following this procedure, call it t_{ub}, is an upper bound to the optimal schedule $\mathcal{S}$, i.e.: $t_{ub} \geq |\mathcal{S}|$.

- Decision variables:

    t_{ij}: time (in slots) link (i,j) is transmitting.
    f_{sij}: flow from s to j through relay i.
    f_{si}: flow from s to i without relaying.
    T: bound on the makespan.

Minimize T,
subject to:

- Flow constraints:

    $f_{si} + \sum_{j \in R^R \setminus \{i\}} f_{sji} = d_i, \qquad \forall i \in R;$
    $f_{si} + \sum_{j \in R \setminus \{i\}} f_{sij} \leq c_{(s,i)} \cdot t_{si}, \qquad \forall i \in R;$
    $f_{sij} \leq c_{(i,j)} \cdot t_{ij}, \qquad \forall i \in R^R, \forall j \in R \setminus \{i\}.$

- Disjoint-intervals constraints:

    $t'_{si} \triangleq t_{si}(1 + \alpha \cdot c_{(s,i)}/(\sum_{j \in R} d_j)), \qquad \forall i \in R;$
    $t'_{ij} \triangleq t_{ij}(1 + \alpha \cdot c_{(i,j)}/d_j), \qquad \forall i \in R, \forall j \in R \setminus \{i\};$
    $\sum_{i \in R} t'_{si} \leq K \cdot T;$
    $t'_{si} + \sum_{j \in R \setminus \{i\}} (t'_{ij} + t'_{ji}) \leq T, \qquad \forall i \in R;$
    $\sum_{\ell \in L} t'_\ell \leq T, \qquad \forall \text{maximal } L \subseteq E: \quad (*)$
    $\qquad\qquad\qquad\qquad\qquad \forall \ell_1, \ell_2 \in L, I_{\ell_1, \ell_2} = 0.$

- Range constraints:

    $t_{ij} \geq 0, \qquad \forall i, j \in V, i \neq j;$
    $t_{is} = 0, \qquad \forall i \in R;$
    $t_{ij} = 0, \qquad \forall i \in R^D, j \in R;$
    $f_{sij} \geq 0, \qquad \forall i \in R^R, \forall j \in R \setminus \{i\};$
    $f_{sij} = 0, \qquad \forall i \in R^D, \forall j \in R \setminus \{i\};$
    $f_{si} \geq 0, \qquad \forall i \in R.$

Fig. 4. LP to obtain how much data should be routed on each link and how much time each link must be active to minimize makespan. The LP does not give the schedule, i.e., a mapping from slots to link activations.

The objective function of this LP is simply to minimize the makespan, whereas the variables indicate the amount of data (flow) that has to be sent through each path (of at most two hops), and the amount of time each link has to be used. The flow and time-period constraints are the following.

1) The usage of each link does not exceed its capacity.
2) The amount of time a node is sending or receiving over any link is not more than the makespan.
3) The aggregated amount of time a set of mutually interfering links are sending or receiving is not more than the makespan. This constraint has been marked as (*) in Fig. 4. Please note that obtaining the sets L described in constraint (*) is analogue to enumerating all maximal cliques in the graph $\mathcal{G} = (V \times R, \mathcal{F})$, where $(\ell_1, \ell_2) \in \mathcal{F}$ if and only if $I_{\ell_1, \ell_2} = 0$. The enumeration of maximal cliques in a graph is an NP-hard problem [49]. In practice, since the inclusion of any enumeration of maximal sets L in the restriction (*) provides a lower bound, we include a greedy enumeration of sets in the following way: given a link ℓ, we build L_ℓ as a maximal set $L_\ell \subseteq E$ such that $\ell \in L_\ell$. For this purpose, we sequentially select all links in E and check if they interfere with all links in L_ℓ. If they do, they are included in L_ℓ. Our numerical results show that this greedy enumeration of maximal sets $\{L_\ell\}_{\ell \in E}$ actually has an impact on the lower bound when the activation cost is analyzed (see Section 8.2).

After removing the interference restrictions, this LP has a polynomial number of restrictions. This LP can be solved optimally with standard interior-point methods [50]. For a given MMWBS input, the LP outputs the amount of data (flow) that minimizes the makespan. However, although the makespan of the LP is a lower bound on the optimal solution for the MMWBS problem, it does not solve it. In

fact, the LP only outputs the period of time t_{ij} each link is active, but not how this time is distributed over slots and when links are activated. Moreover, these times do not take into account the cost of the link activation (the values t'_{ij} include only partially the activation cost). Finally, in our model, at most one link incident to each node may be active in any given time slot, but the solution obtained from the LP may violate this restriction. In our algorithm, we address these issues by modifying the schedule as follows.

We define a vertical phase in which first all the data held by the MBS is downloaded through each link $\ell = (s, r), r \in R$ according to the solution of the LP. Once a downlink ℓ is active, all the data that has to go across ℓ is scheduled, hence the cost of activation is incurred only once. In this phase the schedule of the activation of the links and the transmission of the data is done as referred to in Algorithm 1 (`Direct Download`). The length of this phase is upper bounded by $\sum_{i \in R: t_{si} > 0} \lceil \alpha + t_{si} \rceil$, although in practice it is expected to be smaller due to the possibility of K parallel transmissions.

Now we define a horizontal phase, when the data is sent among μBSs only. To guarantee that we activate at most one link incident to each node, we create virtual links. Then, for each link that has to be active during an interval t (in slots, but maybe not an integer number of slots), we create $\lceil t/(1-\alpha) \rceil$ virtual links between the same pair of nodes. This yields a multigraph on the set of nodes R. We then apply an edge-coloring algorithm in this multigraph so that each edge incident to the same node gets a different color. We modify the schedule accordingly assigning each color to a different slot. Regarding activation costs, given that the virtual links corresponding to the same physical link might not be scheduled consecutively, we upper bound the makespan assuming that a link activates each time it is used.

Algorithm 2 (`Constant Approximation Schedule`)

Input: an instance of the MMWCS problem

Ouput: a mapping of data to links for each time slot.

- Solve the LP of Fig. 4 on the given input to obtain the values $t_{si}, t_{ij}, f_{sij}, f_{si}$ for each $i, j \in R$.

- Use `Direct Download` schedule (cf. Algorithm 1) to transmit f_{si} and f_{sij} data from the MBS s to the μBS i over link (s, i) with one single link activation.

- Create a multigraph $\{V', E'\}$, where $V' = R$ and E' is a multiset of edges containing $\lceil t_{ij}/(1 - \alpha) \rceil$ copies of the edge (i, j), for each $i, j \in R$.

- Run an edge-coloring algorithm on $\{V', E'\}$ and map each color to one successive time slot.

- For each of the following time slots, for each $i, j \in R$, if there is an edge (i, j) in $\{V', E'\}$ corresponding to the current time slot (color), schedule the next block of f_{sij} data, including the link-activation header if needed.

We summarize the described procedure in Algorithm 2 and prove the constant approximation in Theorem 3.

Theorem 3. *Given a communication system with an MBS s, a set R of static μBSs and a collection of data to deliver from s to the μBSs within the model described in Section 4, the following holds:*

1) *Algorithm 2 (`Constant Approximation Schedule`) outputs a schedule $\mathcal{S}$ such that*

$$|\mathcal{S}| \leq \sum_{\substack{i \in R: \\ t_{si} > 0}} \lceil \alpha + t_{si} \rceil + \frac{3}{2} \left(\left\lceil \frac{T}{1-\alpha} \right\rceil + \sqrt{3 \left\lceil \frac{T}{1-\alpha} \right\rceil} \right),$$

where T and the $\{t_{si}\}_{i \in R}$ are as given by the LP of Fig. 4.

2) *With respect to the optimal makespan T_{OPT}, the makespan $|\mathcal{S}|$ entails an approximation of at most*

$$\frac{|\mathcal{S}|}{T_{OPT}} \leq \left(K + \frac{3}{2} \right) \left(\frac{1}{1-\alpha} + \frac{1}{T_{OPT}} \right) + \frac{3}{2} \sqrt{3 \left(\frac{1}{1-\alpha} + \frac{1}{T_{OPT}^2} \right)}.$$

3) *The running time of the Algorithm 2 is*

$$poly \left(|R|, \ \log K + \sum_{\substack{i, j \in V: \\ c_{(i,j)} > 0}} \log c_{(i,j)} + \sum_{\substack{r \in R: \\ d_r > 0}} \log d_r, \ \sum_{r \in R} \left\lceil \frac{d_r}{1-\alpha} \right\rceil \right).$$

Proof. The first term of the upper bound on makespan $|\mathcal{S}|$, that is $\sum_{i \in R: t_{si} > 0} \lceil \alpha + t_{si} \rceil$, upper-bounds the time taken by the vertical phase, and it is simply the aggregation of data-delivery times and activations as if done sequentially, which is the worst case.

The second term corresponds to the horizontal phase and it is obtained as follows. As worst case scenario, we upper bound the cost of activations in this phase assuming that links are activated in each slot. That is, the makespan of the LP including the link-activation cost is at most $\lceil T/(1-\alpha) \rceil$. Additionally, we have to add the overhead cost of the coloring. It is known that the optimal number of colors (i.e., the chromatic index) is $\chi' \leq 3\Delta/2$ (cf. [51]), where Δ is the maximum degree of the graph. Moreover, it has been also shown in [52] how to find a coloring with $\chi' + \sqrt{9\chi'/2}$. We do not know the maximum degree of the multigraph, but we can bound it by the number of steps of the horizontal phase, which in turn is at most $\lceil T/(1-\alpha) \rceil$. Thus, using this coloring algorithm, Algorithm 2 finds a coloring of at most

$3(\lceil T/(1-\alpha) \rceil + \sqrt{3\lceil T/(1-\alpha) \rceil})/2$ colors, and the claimed schedule length follows.

To see why the claimed approximation factor holds, notice that T, i.e., the makespan of the LP, is a lower bound on the optimal makespan T_{OPT}, and that $\sum_{i \in R: t_{si} > 0} \lceil \alpha + t_{si} \rceil \leq K\lceil T/(1-\alpha) \rceil$. Then we have the following:

$$\frac{1}{T_{OPT}} \sum_{\substack{i \in R: \\ t_{si} > 0}} \lceil \alpha + t_{si} \rceil + \frac{3}{2T_{OPT}} \left(\left\lceil \frac{T}{1-\alpha} \right\rceil + \sqrt{3 \left\lceil \frac{T}{1-\alpha} \right\rceil} \right)$$

$$\leq \left(K + \frac{3}{2} \right) \left(\frac{1}{1-\alpha} + \frac{1}{T_{OPT}} \right) + \frac{3}{2} \sqrt{3 \left(\frac{1}{1-\alpha} + \frac{1}{T_{OPT}^2} \right)}.$$

Finally we show the running time of Algorithm 2. The first step can be carried out with an LP solver. There is a wealth of interior-point methods that can be used for this purpose, for instance, Karmarkar's $\mathcal{O}(m^{3.5} B^2)$ algorithm [50], where m is the number of variables and B is the number of the bits in the input. Our LP has $2|R|^2$ variables and $\log_2 K + \sum_{i,j \in V: c_{(i,j)} > 0} \log_2 c_{(i,j)} + \sum_{r \in R: d_r > 0} \log_2 d_r$ bits in the input. Hence, this step can be completed in $\mathcal{O}(|R|^7 + (\log K + \sum_{i,j \in V: c_{(i,j)} > 0} \log c_{(i,j)} + \sum_{r \in R: d_r > 0} \log d_r)^2)$ time.

For the vertical phase, adding the activation times takes $\mathcal{O}(|R|)$, as $|R|$ is the number of links outgoing from the MBS. The sorting of the μBSs takes $\mathcal{O}(|R| \cdot \log |R|)$, and the scheduling at most $\mathcal{O}(|R|^2)$. For the horizontal phase, we create the multigraph $\{R, E'\}$ where R is the set of μBSs and E' is the multiset of edges. For each μBS $r \in R$ the number of incoming virtual links is $\lceil d_r/(1-\alpha) \rceil$. Then, $|E'| \leq \sum_{r \in R} \lceil d_r/(1-\alpha) \rceil$ and the total time to create the multigraph is in $\mathcal{O}(|R| + \sum_{r \in R} \lceil d_r/(1-\alpha) \rceil)$. The next step is the edge-coloring algorithm of [52], which runs in $\text{poly}(\nu, \log \mu)$ time for a multigraph of ν nodes and maximum multiplicity μ. Then, this step takes $\text{poly}(|R|, \log \max_{r \in R} \lceil d_r/(1-\alpha) \rceil))$. Combining all the running times, the claim follows. □

6.4 Makespan upper bound

An upper bound on the MMWBS makespan is given by a scheme that routes all data without relaying or spatial reuse.

Observation 1. *Given a communication system with an MBS s and a set R of n static μBSs, and a collection of data to deliver from s to the μBSs within the model described in Section 4, there exists a schedule $\mathcal{S}$ such that*

$$|\mathcal{S}| \leq \sum_{\substack{i \in R, \\ d_i \neq 0 \\ c_{(s,i)} \neq 0}} \left\lceil \alpha + \frac{d_i}{c_{(s,i)}} \right\rceil + \sum_{\substack{i \in R, \\ d_i \neq 0, \\ c_{(s,i)} = 0}} \left\lceil \alpha + \frac{d_i}{c_{(s,j_i)}} \right\rceil + \left\lceil \alpha + \frac{d_i}{c_{(j_i, i)}} \right\rceil. \quad (13)$$

The makespan in Observation 1 comes from a schedule that delivers data sequentially to μBSs under MBS-coverage using only one RF chain in each time slot (first term). For those μBS i out of MBS-coverage, it sequentially considers indirect paths by relaying over a μBS j_i that relays d_i to i through the fastest path, but without spatial reuse (second term). Therefore, interference has no impact on this bound.

6.5 Lower bounds on the makespan

To solve the MMWBS problem, the MBS has to send to μBSs all the data using a scheduling $\mathcal{S}$. Let $|\mathcal{S}|$ be the makespan of $\mathcal{S}$ in time slots. Sending all the data through the fastest links outgoing from s, using up to K interference-free links concurrently, gives the following lower bound.

Fact 1. *Given a communication system with an MBS s and a set R of μBSs, and a collection of data to deliver from s to the μBSs within the model from Section 4, in the presence of interference, consider a schedule $\mathcal{S}$ that solves the MMWBS problem. Let*

$$C = \max_{\substack{R' \subseteq R: |R'| \leq K: \\ \forall a, b \in R': I_{(s,a),(s,b)}=1}} \sum_{i \in R'} c_{(s,i)}.$$

Then, the length of $\mathcal{S}$ (i.e., the makespan) is as follows:

$$|\mathcal{S}| \geq \left\lceil \alpha + \frac{\sum_{j \in R} d_j}{C} \right\rceil. \tag{14}$$

Another lower bound, based on link activation, is proved for a given schedule in Lemma 1, and a universal lower bound (for any schedule) is given in Theorem 4. These bounds are relevant when the amount of data to send is small and the makespan is dominated by link activations.

Lemma 1. *Given a communication system with an MBS s, a set R of n μBSs, and a collection of data to deliver from s to $n' \leq n$ μBSs within the model from Section 4, consider a schedule $\mathcal{S}$ that solves the MMWBS problem. If the set of μBSs that receive data directly from s is $D \subseteq R$ then, even without interference,*

$$|\mathcal{S}| \geq d + \max \left\{ 0, \ \left\lceil \frac{n' - |D| - (d(d-1)/2)\,K}{|D|} \right\rceil \right\},$$
$$\text{with } d = \lceil |D|/K \rceil.$$

Proof. Let the data to be delivered be called simply data. Let a μBS that has data, for itself or for other nodes, be called informed. Consider the sequence of time slots $t_1, t_2, \ldots, t_d$ when the μBSs in D are informed (possibly interleaved with other time slots when no μBS in D is informed). Recall that D is defined to be the set of μBSs that receive directly from the MBS. Then, in each time slot, at most K new μBSs in D may be informed, and it is $d \geq \lceil |D|/K \rceil$. For $1 \leq i \leq d$, let $D(t_i)$ be the subset of μBSs in D that have been informed by time t_i. Then, $|D(t_i)| \leq iK$ for $1 \leq i < d$, and $|D(t_d)| = |D|$.

Let $\mathcal{I}$ be the set of μBSs that do not receive directly from s in the schedule $\mathcal{S}$. For any time slot t, let $\mathcal{I}(t)$ be the subset of μBSs in $\mathcal{I}$ informed during time slot t. Then, given that μBSs in $\mathcal{I}$ are only informed by the μBSs in D, we have that

$$|\mathcal{I}(t)| \leq \begin{cases} 0 & \text{for } t \leq t_1; \\ |D(t_i)| & \text{for } t_i < t \leq t_{i+1} \text{ and } 1 \leq i < d; \\ |D| & \text{for } t > t_d. \end{cases}$$

Since $|D(t_i)| \leq |D|$ for all $1 \leq i \leq d$, to prove the lower bound we assume as a worst case that all μBSs in D are informed in the first $d = \lceil |D|/K \rceil$ time slots. Then, the sequence of numbers of μBSs informed along time slots $1, 2, \ldots, d$, is $|D(t_1)|, |D(t_2)| + |\mathcal{I}(t_2)|, \ldots, |D| + \sum_{i=2}^{d} |\mathcal{I}(t_i)|$. So, at the end of slot d, the number of informed μBSs is

$$|D| + \sum_{i=2}^{d} |\mathcal{I}(t_i)| \leq |D| + \sum_{i=2}^{d} (i-1)K = |D| + \frac{d(d-1)}{2}K. \tag{15}$$

From slot $t_d + 1$, at most $|D|$ new μBSs are informed in each slot. Thus, to inform remaining μBSs (if any), we need at least $\left\lceil \left(n' - |D| - \frac{Kd}{2}(d-1) \right) / |D| \right\rceil$ additional slots. Adding this time to the previous d time slots, the claim follows. □

Theorem 4. *Given a communication system with an MBS s, a set R of n μBSs, and a collection of data to deliver from s to*

$n' \leq n$ *μBSs within the model from Section 4, consider a schedule $\mathcal{S}$ that solves the MMWBS problem. Then,*

$$|\mathcal{S}| \geq \left\lceil \sqrt{1/4 + 2\,(n'/K + 1)} - 1/2 \right\rceil.$$

Proof. Consider the maximum number of μBSs that may be informed in slots $1, 2, \ldots$. In the first time slot at most K μBSs may be informed. In the second slot at most K μBSs may be informed directly and another K may be informed by relaying. We continue the same analysis to compute the first time slot t when $tK + \sum_{i=2}^{t-1} iK \geq n'$, i.e., when

$$t^2 + t - 2(n'/K + 1) \geq 0.$$

Solving the quadratic equation the claim holds. □

Another lower bound on the makespan is given by the LP of Fig. 4 (see Section 6.3 for an explanation of the constraints). The objective function here is simply to minimize the makespan, whereas the variables indicate the amount of data (flow) that has to be sent through each path (of at most two hops), and the amount of time that each link is used.

Notice that the formulation does not restrict the temporal order in which the links must be used (as opposed to the MILP of Section 5). For instance, data that is being relayed can be delivered to a given destination only after reaching the relay. This LP does not restrict such temporal order. Therefore, the makespan obtained from the solution of this LP is only a lower bound on the optimal makespan. The reason being that, anyway, flow and time-period restrictions have to be observed, but the optimal makespan for MMWBS could be even larger after additionally restricting the order in which links are used.

In the experimental evaluation, we compare our algorithms with the maximum of the lower bounds obtained in Fact 1, Theorem 4, and the solution of the LP in Fig. 4.

7 HEURISTICS

7.1 Greedy heuristic

We present first a simple greedy heuristic to be compared with the bounds derived in Section 6. The heuristic greedily uses the links that are faster as soon as they are available.

The schedule of transmissions with Greedy is built as follows. At any time slot t in which an RF chain k in the MBS s is available (e.g., initially or when it completes sending data to a μBS), it schedules a new transmission across the fastest link (s, i) from s to any available μBS $i \in R^R$. The data $d_j > 0$ sent in this transmission will be for the available μBS j that has the fastest link (i, j), among the μBSs whose data is still at s. Such scheduling begins in the first time slot later or same than t where the download of d_j by link (s, i) and the relaying of d_j by link (i, j) does not interfere with other allocated transmissions. Nodes s and i are then unavailable for a period of $\lceil \alpha + d_j/c_{si} \rceil$ time slots in case link (s, i) was not used through RF chain k in the time slot prior to this link allocation, or for a period of $\lceil d_j/c_{si} \rceil$ otherwise. When the data d_j is completely received at i, the MBS becomes available again in the next time slot and at the same time the data is forwarded to j via the link (i, j). μBSs i and j are then unavailable for a period of length $\lceil \alpha + d_j/c_{ij} \rceil$ time slots. Here, the activation cost α is considered since link (i, j) cannot be active in any previous time slot. In case

at some point an RF chain in the MBS becomes available and there is only one available μBS i (because all the others already have their data or because they are relaying/busy), the MBS schedules a direct download of data d_i for μBS i in the earliest time slot that does not interfere.

In summary, at each iteration, while there is data to be served, Greedy takes the fastest available μBS and sends it data d_j for the fastest available neighbour μBS that still does not have its data. Hence, we enable parallel transmission as soon as possible, as intended in our mmWave backhaul with relaying. In case this is not possible, Greedy schedules a direct download. Since checking interference issues only consists into logical checks, it takes at most n iterations, one per μBS waiting for a file, to decide a final schedule. Thus, the computational complexity of this algorithm is $\mathcal{O}(n)$.

7.2 Resched heuristic

The second heuristic is based on reschededuling an initial communication assignment in order to iteratively improve the makespan at each step until no further improvement is possible. We call this heuristic Resched.

We consider an initial feasible schedule provided by Direct Download (cf. Algorithm 1), which consists of sequential direct downloads without relaying from the MBS to the μBSs under MBS-coverage that have a file to be served, while using all K RF chains and avoiding interfered connections, and sequential relayed downloads for those μBSs out of MBS-coverage. The set of μBSs that have to receive a file is sorted from lowest to highest delivery time. At this point, every μBS $r \in R$ such that $d_r > 0$ is scheduled through an RF chain $k_r \le K$ and at a time slot in which it begins its download, which we call its initiation instant: t_{sr}. For those $r' \in R^R$ such that $d_{r'} = 0$ we let $t_{sr'}$ unset. Then, we iteratively modify such assignment by rescheduling the transmissions, taking advantage of relaying.

At each iteration, we take the μBS $u \in R$ such that $d_u > 0$ that receives its file the latest and that has been tried to be reallocated less times. Then, we reallocate the path transmission of its data, d_u. For such reallocation, we search for a μBS $r \in R^R$ that may potentially relay the data d_u to u. Thus, we check each μBS $r \in R^R$ in the order of the sorted list of μBSs and take the one that reduces most the current makespan and has no interference issues (this is, the new retransmission of the data d_u cannot be scheduled through links that interfere with the already allocated transmissions of the other files of the system). The way in which the retransmission of data d_u is allocated is the following:

We take the initiation instant of r, t_{sr}. Let τ_r^δ be the number of time slots the μBS r is downloading files and let τ_r^ρ be the number of time slots the μBS r is relaying files, according to the current schedule. In case r is already downloading data from the MBS, i.e. $\tau_r^\delta > 0$, we take the RF chain k_r through which such download is scheduled, and define a binary indicator $\xi_r = 1$. Otherwise, i.e. if $\tau_r^\delta = 0$, we select an RF chain $k_r \le K$ in the MBS such that k_r stops being used in the earliest time slot t and define the binary indicator $\xi_r = 0$. Then, we take $t_{sr} = t$. The μBS r begins to download in time slot t_{sr} the data d_u through the RF chain k_r. Since data d_u is now downloaded by r, instead of by u, the RF chain k_u gets free for those times slots corresponding

to the direct download, $[t_{su}, t_{su} + \lceil \alpha + d_u/c_{(s,u)} \rceil - 1]$, so that other transmissions can be allocated in such time slots later in the heuristic. The download of d_u will take place through RF chain k_r in the time slots

$$\left[t_{sr} + \tau_r^\delta - 1, \quad t_{sr} + \tau_r^\delta + \lceil \alpha \cdot \xi_r + d_u/c_{(s,r)} \rceil - 1 \right].$$

Then, we update τ_r^δ to $\tau_r^\delta = \tau_r^\delta + \lceil \alpha \cdot \xi_r + d_u/c_{(s,r)} \rceil$. Please note that, in case r was already downloading data before allocating to it the download of data d_u, we do not consider the activation cost α, while in the opposite case we do[7]. Once r downloads all the files currently allocated to it, r relays the data d_u to u in the time slots

$$\left[t_{sr} + \tau_r^\delta + \tau_r^\rho - 1, \quad t_{sr} + \tau_r^\delta + \tau_r^\rho + \lceil \alpha + d_u/c_{(r,u)} \rceil - 1 \right].$$

Then, we update τ_r^ρ to $\tau_r^\rho = \tau^\rho + \lceil \alpha + d_u/c_{(r,u)} \rceil$. Regardless the case, α has to be considered when updating τ_r^ρ because link (r, u) could not have been activated before.

The reallocation described for data d_u clearly affects the scheduling of those files downloaded through RF chain k_r. Since now r has to use more time slots to download one file more, the transmissions beginning later than t_{sr} through RF chain k_r must be delayed, as well as the relaying of such files. Thus, for all $r' \in R$ such that $t_{sr'} > t_{sr}$, the initiation instant of r' is updated to:

$$t_{sr'} = t_{sr'} + \lceil \alpha \cdot \xi_r + d_u/c_{(s,r)} \rceil.$$

Thus, all the time slot intervals used for relaying are delayed as well, according to the new value of $t_{sr'}$, and the current makespan is modified.

If now the makespan is shorter than before and there are not interference issues, we keep the new schedule. Otherwise, we discard such reallocation of μBS u and try to relay data d_u through a different relay $r \in R^R$ not selected for u yet. If all relays in R^R have already been selected for u without success, another iteration starts for the relaying of data from the next μBS in the sorted list of μBSs, and u is replaced the last in the sorted list.

Resched ends when all non-reallocated nodes waiting for a file have been tried to be rescheduled without success.

In terms of complexity, note that Resched begins with an initial feasible allocation and then, it checks μBS by μBS with a file in the MBS if its download can be rescheduled to reduce the makespan. For each μBS, the heuristic checks up to $|R^R| \le n$ possibilities for its rescheduling. Thus, since checking interference issues only consists into logical checks, the computational complexity is $\mathcal{O}(n^2)$.

8 EXPERIMENTS

We perform experiments in scenarios that accurately reproduce real mmWave communication systems. We consider an MBS and a set of μBSs, as described in Section 4. We perform experiments with different choices for sets R^R and R^D where we analyze the behaviour of our algorithms, as well as different numbers of RF chains in the MBS. Also, we model interference in the network based on a model for mmWave antenna patterns that measures the radiating

---

7. We consider that every relay $r \in R^R$ first downloads all the data allocated to it, so that the activation cost with the MBS is considered only once. After all the downloads by r end, r begins to relay the files with the other μBSs allocated to r.

beamwidths, as detailed later. To shed light on scenarios based on real measured rates and beam-patterns, we perform simulations in which link SNR, link rates and interference are obtained from real experiments, as detailed later.

To collect consistent statistics, we simulate each scenario 1000 times. We show in each plot the lower bounds detailed in Section 6.5. Nevertheless, we compact them onto the maximum lower bound that we get on each experiment, i.e., we select the lower bound with better guarantees for the makespan. Regarding upper bounds, here we mention that `Constant Approximation Schedule` (cf. Algorithm 2) provides theoretical guarantees for the makespan in absence of interference, as proved in Theorem 3. However, the `Direct Download` schedule described in Algorithm 1 practically provides better upper bounds in any case (either in absence or presence of interference), although it does not give theoretic guarantees. Hence, here we show the makespan provided by `Direct Download`.

The labels of the figures' legends refer to the makespan obtained with the following schemes:

- **Upper Bound:** Makespan resulting from `Direct Download` (cf. Algorithm 1) and Observation 1.
- **Lower Bound:** Longest makespan from Fact 1, Theorem 4 and Fig. 4.
- **Resched:** Makespan obtained with `Resched` heuristic.
- **Greedy:** Makespan obtained with `Greedy` heuristic.
- **Optimum:** Optimum makespan obtained by solving the MILP for MMWBS problem derived in Section 5.

8.1 Experimental setup

In the experiments, we consider that the MBS has data of random length for each μBS, drawn from a truncated exponential distribution ranging from 1 MB to 80 MB, with an average of 10 MB. We deploy a circular cell of radius R_C centered in the MBS and place uniformly at random n μBSs inside the cell. We consider a fixed transmit power of $P_t = 30$ dBm, and fixed antenna gains of $G_t = 25$ dB and $G_r = 25$ dB for transmitter and receiver respectively. According to the Friis equation, the power received in dB, P_r is $P_r[dB] = P_t + G_t + G_r + 10\eta \log_{10}\left(\frac{\lambda}{4\pi\delta}\right) - 5$, where $\eta = 2$ is the path loss exponent in free space, λ is the wavelength in meters for 60 GHz carrier frequency, and δ is the distance between transmitter and receiver. Besides, we subtract an implementation loss of 5 dB. The thermal noise power of Johnson-Nyquist in dBm is: $P_N[dBm] = -174 + 10\log_{10}(W)$, where $W = 2.16 \cdot 10^9$ is the bandwidth in Hz used for transmission. We amplify this noise by the receiver noise factor of 40 dB, so the actual noise in dBm in the receiver is $N[dBm] = P_N + 40 = -174 + 10\log_{10}(W) + 40$. The achieved signal-to-noise ratio is $SNR = 10^{\frac{P_r - N + 30}{10}}$, and the electronic sensitivity S in dBm at the receiver is $S[dBm] = 10\log_{10}(k(T_a + T_{Rx})W \cdot SNR) + 30$, where k is the Boltzmann constant, T_a is the noise temperature in Kelvin of the antenna at the input of the receiver, T_{Rx} is the noise temperature in Kelvin of the receiver referred to its input, and W is again bandwidth. We use $T_a = T_{Rx} = 290$ K.

We use the link rates resulting from the above computation and corresponding to the SCPHY modulation and coding schemes of [53], which have been implemented in commercial mmWave devices. SCPHY rates range from 385 Mb/s to 4620 Mb/s. We further use an EIRP of 55 dBm. This makes possible to find theoretically feasible links with $R_C = 35$ m, as we adopt.

The time slot is fixed to $T_s = 10$ ms, and the activation time, namely A_t, of every link is 1.0349 ms, unless otherwise specified. This activation time corresponds to the antenna steering time spent to explore a finite number of sectors, as specified, e.g., in the IEEE 802.11ad amendment [53]. Hence, this time is composed of a preamble and a header of 4291 and 4654 nanoseconds each, plus the time needed to steer the antenna towards the intended receiver and the intended transmitter. Here we assume that base stations have arrays of antennas of $N_A = 32$ sectors. Since the time to transmit the sector sweep frame (SSW) is 15.76 μs, and both transmitter and receiver need to steer their beams one after the other, the final activation time will be:

$$A_t = 2N_A \cdot 15.76\ \mu s + 18.2\ \mu s + 4.29\ \mu s + 4.65\ \mu s \approx 1.0349\ \text{ms.} \quad (16)$$

The first term in Eq. (16) corresponds to the N_A SSWs from the transmitter and N_A SSWs from the receiver. The second term corresponds to one feedback frame from the transmitter, and the remaining terms are the preamble and header duration, respectively. Since in the paper we assume slot length normalized to 1, we have that every link has an activation cost of $\alpha = A_t/T_s = 0.10349$.

We discuss two cases of the model described in Section 4:

- Full Network: We consider that the full network helps to relay data, i.e., $R^R = R$ and $R^D = \emptyset$, and that all the μBS $r \in R$ have data of size $d_r > 0$ to download. This case fits the main purpose of this paper, when a mmWave backhaul network is deployed to obtain the files for the user equipments served by the μBSs as quickly as possible. Thus, relaying is enabled in the full network.
- Small Cell Network: We consider that a set of μBSs act as relays whose only purpose is to help the network by relaying data. They do not have data for themselves, so $\emptyset \neq R^R \subsetneq R$ and $d_r = 0, \forall r \in R^R$. Although this paper is intended for a mmWave wireless backhaul network, this case more generically represents and sheds light on scenarios where a mmWave access network is deployed, so that there may be mmWave devices willing to help the network through relaying but do not claim any file.

Furthermore, we also study scenarios under the effects of interference, as described in the model of Section 4. In order to capture such effects, we consider an interference model based on antenna patterns with one main lobe of a given half power beamwidth (HPBW), and sidelobes with a power lower than such HPBW. This model represents real directional antennas that may radiate and interfere only within an angle given by the HPBW. Such assumption is based on a simple analytical model [54] which characterizes the electric field of an antenna as a function of the beamwidth. As depicted in Fig. 2, actual beam-patterns may radiate non-negligible power in sidelobe directions. The beam patterns depicted in the figure correspond to actual patterns experimentally obtained in a recent work [39], [40]. In our experiments, we use HPBW values large enough to account for sidelobes. In any case, the interference matrix I (i.e., parameter I_{ℓ_1, ℓ_2}) is assumed as an input of the problem, so that the performance of our algorithms is not affected by the exact shape of the antenna patterns. Moreover,
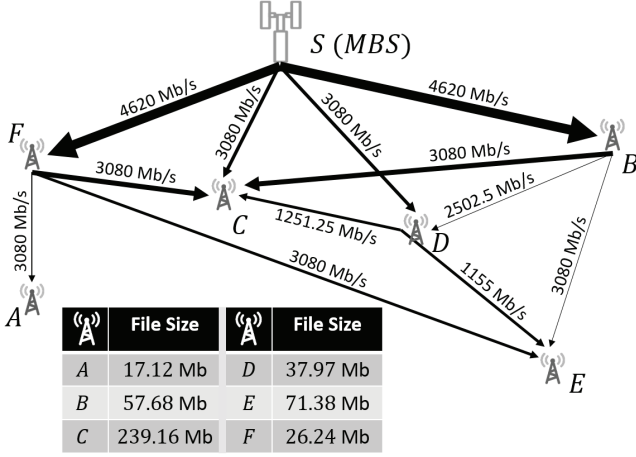
Fig. 5. Example of optimal scheduling with 6 μBSs and one MBS with single RF chain ($K = 1$). The figure shows the logical topology, the set of links used and their utilization, the scheduling and the makespan with its bounds.
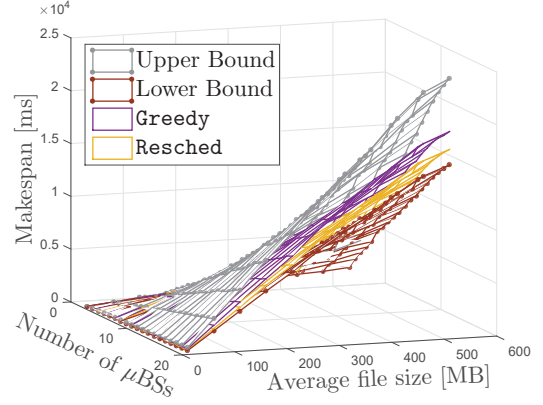

Fig. 6. Makespan in the full network case, with $K = 1$ RF chain and without interference: Impact of the size of the network, n, and the average file size.

since the deployment of μBSs is uniformly random, the resulting average results are the same as if an arbitrary number of sidelobes were considered. In [39], [40], authors have observed that a typical aggregate radiating width for commercial antennas is around 22.5 degrees, which approximately corresponds to $\pi/8$ radians, as we mainly adopt in our numerical evaluation. In addition, we show at the end of the section results in which we use measured data for the exact shape of beam-patterns and link rates obtained from a commercial mmWave device: the TP-Link Talon AD7200 router. While this mmWave device is for indoor use and mmWave backhaul has somewhat different characteristics (more refined beam-patterns, higher rates, ...), the underlying RF technology and specifically the phased arrays are similar enough for these measurements to give meaningful results. The exact shape of these beam-patterns is available at [55] and the data can be downloaded at [56]. In order to build the real binary interference map we need to know which beam-pattern each μBS will use with each of its neighbors. Hence, we simulate the beam-training of links based on the link SNR: when a link is trained, each of the μBSs involved tests all of its beam patterns and chooses the one that provides the highest SNR. The achievable link rates of these mmWave devices have been investigated in [57], from where we obtain the link data rate based on the distance between two devices. This allows to study our framework for realistic relaying and spatial reuse scenarios. We further provide a performance evaluation based on synthetic and modeled beam-patterns that represent future backhaul applications using better hardware.

8.2 Numerical results

Here we present multiple results of numerical experiments. We use R2018a version of MATLAB in order to simulate

channel conditions, packet sizes, interference topology and positions of μBSs. We use the CPLEX optimizer to find optimum values for small instances of the MMWBS problem. In fact, it is hard to obtain optimal solutions from the MILP formulation presented in Section 4, due to the NP-hardness proven in Section 6.2. Thus, optimum values are only available for small numbers of μBSs, which have been obtained through exhaustive search methods as Branch & Bound [58]. In the figures shown in this section, error bars represent 95% confidence intervals centered on the average extracted from 1000 simulations.

We show in Fig. 5 a simple example of an optimal schedule. Here, we consider $n = 6$ μBSs, labeled from A to F, and an MBS with only one RF chain ($K = 1$). The link rates shown are extracted from the SCPHY modulation and coding schemes defined in the standard. The rate corresponds to the SNR experienced by the receiver according to the power and noise computed as described in Section 8.1. Therefore, rates depend on the distance between transmitter and receiver, and on the adopted channel model. The upper part of the figure shows the logical topology (not the physical one), the links that are used, the data rate of links (label next to the link), the utilization of each link (thickness of the lines), and the size of downloaded data (in the table on the left). The bottom part of the figure shows the optimal schedule, which corresponds to a makespan of 12 time slots (*TS* in the figure) for this example, and is very close to the lower bound of 10 slots indicated above the scheduling table. This makespan is a 30% lower than the upper bound of 17 slots of the schedule with no relaying. Hence, end-user demands can be served much faster at the serving μBS. We also show at each time slot which links have been scheduled (marked with an arrow), which μBSs are intended to receive the data sent over such links (marked with the μBS label), and how much data in Mb is sent for the intended destination. In what follows, we show that this makespan improvement is the general behavior, thus proving the importance of studying relay and spatial reuse featured in wireless mmWave backhaul networks.

8.2.1 Full network case

Fig. 6 shows the impact of the number of nodes and of the average data size on the makespan. Here we do not consider interference and we use $K = 1$. As expected in this case,
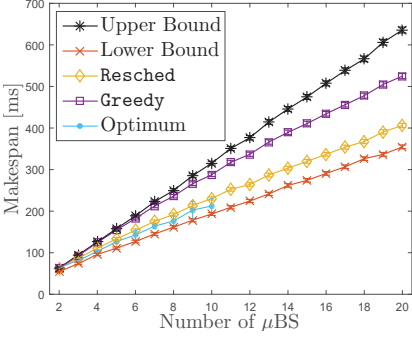
Fig. 7. Full network case with $K = 1$, no interference, average data sizes of 10 MB and comparison to optimum.
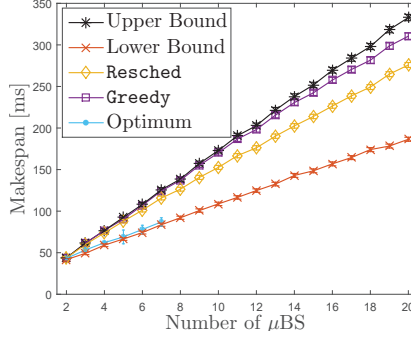

Fig. 8. Full network case with $K = 2$, no interference: Impact of network size n, with comparison to optimum.
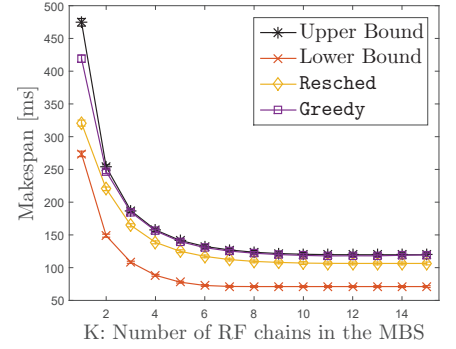

Fig. 9. Full network case, with $n = 15$ and interference of HPBW$=\pi/8$ rads: Impact of K.
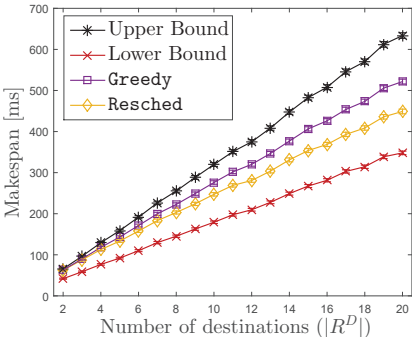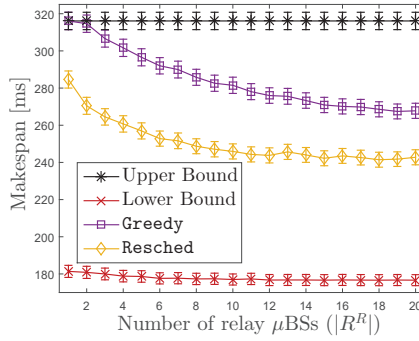

Fig. 10. Small cell network case, with $K = 1$ RF chain in the MBS and without interference: impact of the number of μBSs in R^D for a fixed number ($|R^R| = 15$) of relay μBSs in R^R.

Fig. 11. Small cell network case, with $K = 1$ RF chain in the MBS and without interference: impact of the number of relay μBSs in R^R (with no data for themselves) for $|R^D| = 10$ destinations.
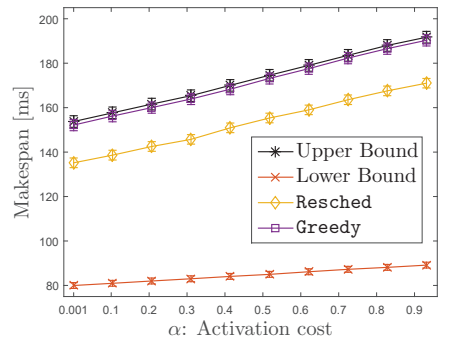
Fig. 12. Small cell network case, with $K = 4$ RF chains in the MBS and with an interference of HPBW = $\pi/8$: impact of the activation cost, α, for $|R^R| = 15$.

the makespan grows as long as the file size average grows, while heuristics reduce the makespan with respect to the direct download from the MBS. Indeed, Resched gets closer to the theoretical lower bound.

Due to the NP-hardness of the problem, optimum results are computationally unfeasible, thus we show optimality only for small networks, and with an average data size of 10 MB, in Fig. 7. Here the optimum is computed by solving the MILP formulation for MMWBS. Bounds and heuristics are reported for comparison (and they are as in the slice obtained for average file sizes of 10 MB in Fig. 6).

In Fig. 7, the makespan grows linearly with the size of the network since the average burden of data at the MBS grows linearly with the number of nodes. We observe that the optimal results perform close to the lower bound and the heuristics, in particular, the Resched heuristic operates near-optimally when the MBS has one RF chain. Greedy achieves worse results than Resched, which is not surprising given its low computational complexity.

In Fig. 8 we also show more optimal results for the makespan when the MBS disposes of $K = 2$ RF chains. Again, the makespan grows linearly with the size of the network, although the achieved values are much lower in this case. This fact is due to the possibility of using up to two simultaneous links from the MBS. In this case, Resched behaves better than Greedy, although the distance from the lower bound and from the optimum is higher than with $K = 1$. This is due to the fact that our heuristics give priority to direct download, when it can be fast, so that the MBS can transmit more often when K increases.

Instead, we show the impact of K for the case in which we consider interference, in Fig. 9. In there, we fix the size of the network to $n = 15$ μBSs using antenna patterns radiating and interfering within $\pi/8$ rads. The figure shows that the makespan tends to decrease considerably as long as K increases. Such decrease follows an interesting shape that slows down the decrease until converging to an almost constant makespan. The reason behind this behaviour is that the more RF chains we have, the more parallel links can be active from the MBS and the less relaying takes place in the network. Still, Resched is able to reduce the makespan and take advantage of relaying in all the cases, despite the small gap of 50 ms (5 slots) between upper and lower bounds with high values of K. Indeed, the figure shows how tight our bounds are (although the lower bound is not necessarily a feasible schedule) and how Resched can achieve makespan reductions of 30% to 75% with respect to direct download (i.e., the upper bound) and a Greedy heuristic, using reasonable values of K. Note also that a greedy approach to relay yields no practical benefit as soon as the MBS can use two or three RF chains.

In general, Figs. 7-9 show that Resched provides good results, but never converges to the lower bound. Since heuristics always provide a makespan larger than the optimal, the optimal is not necessarily the lower bound. Heuristics aim to be as close as possible to optimal schedules, not to the lower bound. However, for those cases where obtaining the optimum schedule is not computationally feasible, Resched should approach as much as possible the lower bound, since it cannot be compared to other benchmarks.
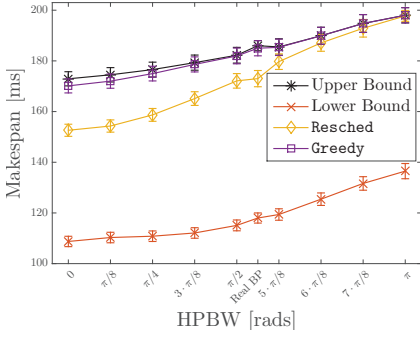
Fig. 13. Full network case, with $n = 10$ μBSs and $K = 2$ RF chains in the MBS: impact of non-ideal beamwidths causing an interference.
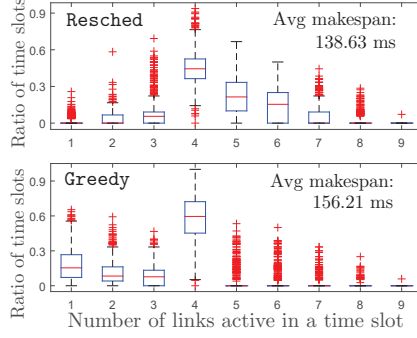

Fig. 14. Full network case, with $n = 15$ μBSs and $K = 4$ RF chains in the MBS: reuse of links in the presence of interference caused by a HPBW of $\pi/8$ rads.
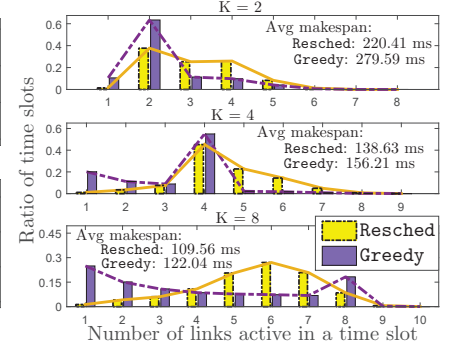

Fig. 15. Full network case, with $n = 15$ μBSs and $K = 2, 4, 8$ RF chains in the MBS. Distribution of spatial reuse with interference caused by HPBW $= \pi/8$ rads.
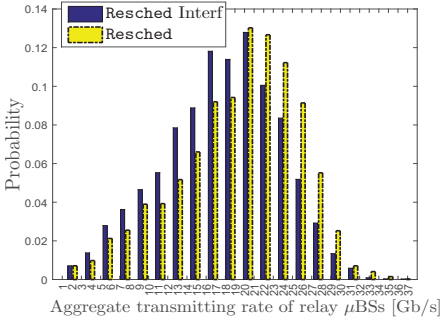

Fig. 16. Full network case, with $n = 15$ μBSs and $K = 8$ RF chains in the MBS: distribution of the aggregate transmitting rate of μBSs with Resched without and with interference of $\pi/8$ rads HPBW.
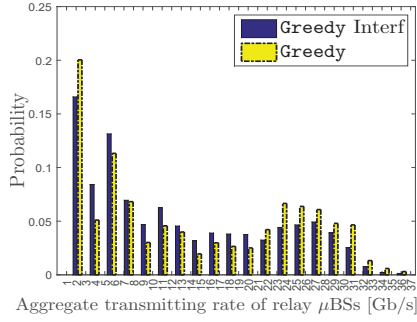
Fig. 17. Full network case, with $n = 15$ μBSs and $K = 8$ RF chains in the MBS: distribution of the aggregate transmitting rate of μBSs with Greedy without and with interference of $\pi/8$ rads HPBW.
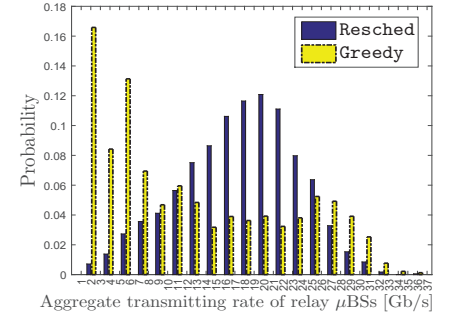
Fig. 18. Full network case, with $n = 15$ μBSs, $K = 8$ RF chains in the MBS and interference of a HPBW of $\pi/8$ rads: distribution of the aggregate transmitting rate of μBSs with Resched and Greedy.

8.2.2 Small cell network

We next take a set $\emptyset \neq R^R \subsetneq R$ of μBSs with no files ($d_r = 0, \forall r \in R^R$) whose only task is to help the network through relaying as detailed in Section 8.1. In this scenario, all other μBSs i in R^D do not relay traffic, which is implemented by setting to 0 the rate of any link (i, r). For a given small cell network, we consider a uniformly random placement of relays in R^R, as well as for other nodes.

In Fig. 10 we show the impact of fixing the number of relay μBSs in R^R to 15 and increasing the destination nodes in R^D in the network. Here, without interference and with $K = 1$, the behaviour is expected to be a linear increase since we increase the traffic burden in the MBS as long as we add destination nodes. Resched behaves as a good scheduler for the makespan, not very far from the lower bound and reducing more than 50% the gap between both bounds, while Greedy suffers from the simplicity in its design.

More interestingly, to complement the results of Fig. 10, in Fig. 11 we show the impact of growing the small cell network through increasing the number of relays in R^R. We fix the number of end-nodes in R^D to $|R^D| = 10$. Again, we consider no interference and assume $K = 1$. We place randomly the destination nodes of R^D in each instance and keep increasing the number of relays. Such increase brings more chances for the end-nodes in R^D to be helped and then the makespan keeps reducing with both Greedy and Resched, although the gain flattens for high numbers of relays. The gain of Resched is remarkable, since it can practically save between 15% to 50% of time to complete

the download of all data. Results with interference and with $K > 1$ yield very similar behaviours.

In Fig. 12, we study the impact of the link activation cost, α. Since the behaviour is similar in all cases, we only show an example with $K = 4$ and with interference. We find that the makespan of heuristics and of bounds increases linearly with the activation cost. The slope of the lower bound is, however, smaller than for the other curves. This result is due to the fact that our heuristics and the upper bound do not allow to split a data download in multiple separate chunks. Instead, when a download is scheduled, the entire file is transmitted. The increase of the lower bound is less remarkable here because such lower bound is computed with the LP of Fig. 4, which does not use time slots, so that it can be more efficient (although the resulting scheduling is not necessarily feasible). In practice, while increasing the link activation cost may provoke the increase of an integer number of time slot for heuristics and upper bound, the LP used for the lower bound only increases file transmission times by α, without using any discretized schedule. Here we have also tested an activation cost of 1 μs, for those cases in which one can assume that the beam-training is saved an not repeated when links activate.

8.2.3 Interference and spectrum reuse

In Fig. 13 we study the impact on the makespan in the presence of interference, as a function of the HPBW, for a full network of 10 μBSs and $K = 2$. We observe how beamwidths below $\pi/4$ rads barely affect network performance, and so
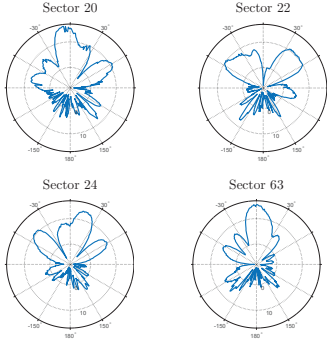
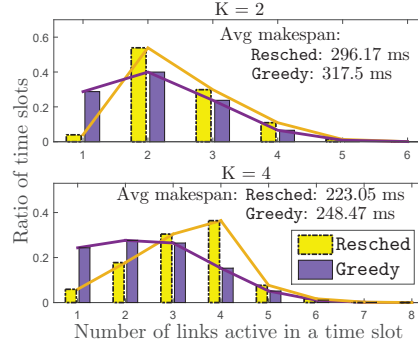Fig. 19. Measured beam-patterns in [55] from commercial off-the-shelf mmWave devices.

Fig. 20. Full network case, with $n = 15$ μBSs and $K = 2, 4$ RF chains in the MBS. Distribution of spatial reuse with interference caused by real beam-patterns measured in [55].
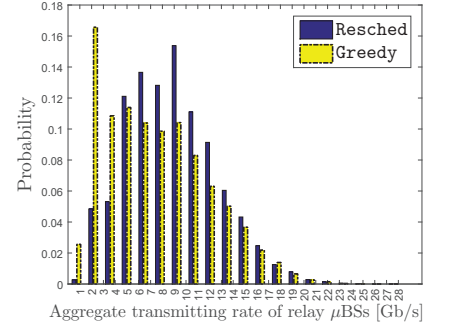

Fig. 21. Full network case, with $n = 15$ μBSs, $K = 8$ RF chains and interference caused by real beam-patterns [55]: distribution of the aggregate transmitting rate of μBSs.

`Resched` behaves in a similar way in presence or absence of interference. However, as interference increases because of larger beamwidths, `Resched` becomes less impaired than `Greedy`. Still, the overall makespan increases because the interference limits network capacity. For beamwidths of π rads the performance of the makespan is already like the direct downloads without spectrum reuse. We also compare the behavior of `Resched` and `Greedy` by considering the reuse of links in the presence of interference, as shown in Fig. 14 and Fig. 15. Clearly, `Resched` achieves higher reuse factors, even though the degree of interference considered in the figure is low (with beamwidths of $\pi/8$ rads). Specifically, in Fig. 14 we use *box-and-whisker* plots to observe the distribution for the ratio of time slots in which a number of links are active when we have $K = 4$ RF chains in the MBS. Fig. 15 reports just average values for such ratio, and compares different values of K. Since there can be up to K simultaneous transmissions from the MBS, this is often the most common number of simultaneous links active in the network. However, we can see from the figures that `Resched` is able to take advantage of relaying and can use even more than K active links per time slots with high frequency, while `Greedy` suffers from its simplicity and often uses K active links. Nonetheless, we note from the bottom part of Fig. 15 that `Resched` uses less than K links in parallel with high frequency when K is high. The resulting makespan is however shorter than for `Greedy`. The reason behind this counter-intuitive example is that high spatial reuse does not necessarily lead to faster downloads when interference can build up, which takes us to the next set of results, in which we consider the rates actually used.

Figs. 16 and 17 depict the distribution of aggregate network rates for `Resched` and `Greedy`, respectively. In each of the two figures we report the distribution with and without interference, for $K = 8$ and a full network with $n = 15$. We observe that the presence of interference tends to reduce the use of better links. Finally, Fig. 18 provides a direct comparison between the two heuristics in the presence of interference. Concretely, with beam-patterns having a HPBW of $\pi/8$ rads, 15.6% of the pairs of links interfere, on average. Here, although `Greedy` selects always the fastest available links among relays, it ends up providing worse aggregate rates and longer makespans than `Resched` because it forces the use of direct downloads in absence of good inter-μBS links. Instead, `Resched` avoids scheduling too many links

when interference builds up. Therefore `Resched` achieves high spatial reuse without penalizing speeds.

8.2.4 Real-measured interference and link rates

Finally, we provide some results from scenarios with realistic measured beam-patterns and data rates from mmWave devices. Here, the interference map is based on real shapes of beam-patterns for cheap commercial antennas [55], [56] and the makespan directly depends on real rates observed in 802.11-based mmWave devices [57]. In Fig. 19 we show four of the 35 available beam-patterns integrated on the TP-Link Talon AD7200. Here we can observe that although communication is directional, there are relevant sidelobes that incur strong interference in the system. Hence, specific algorithms that carefully manage interference in order to provide high spatial reuse are needed, as `Resched` or `Greedy`. In Fig. 20 we show the spatial reuse when interference is caused by real beam-patterns. As observed before, `Resched` provides higher spatial reuse than `Greedy` and hence provides a lower average makespan. In comparison with Fig. 15, we observe a lower spatial reuse in Fig. 20 because there is more presence of non-negligible interference that affects the possibilities of spatial reuse. This fact also leads to lower aggregate transmit rate values in Fig. 21, where we observe better spatial reuse in which aggregate rates are higher, but lower than the ones observed in Fig. 18. Moreover, real SNR measures obtained in [57] are based on the imperfect beam-pattern shapes shown and hence provide lower link rates. In this case, 27.4% of the pairs of links interfere, on average.

9 LESSONS LEARNT AND DISCUSSION

Our analysis has shown that, in general, minimizing the delivery time (i.e., the makespan) of a collection of files in a mmWave backhaul network is not doable in polynomial time, unless $P = NP$. Even in simple scenarios where mmWave beams are fine-grained enough so that interference is neglected and the MBS only has one RF chain to transmit, the problem is NP-hard. So, time-efficient heuristics as `Resched` and `Greedy` are necessary to find feasible solutions. Our results show that our heuristics, specially `Resched`, are able to provide near-optimal solutions that are, with respect to the distance between upper and lower bounds, 40-80% closer to the lower bound. Hence, since
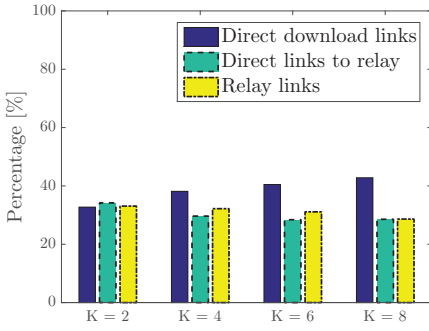
Fig. 22. Full network case, with $n = 15$ μBSs and interference of $\pi/8$ rads HPBW. Usage of direct and relay links.
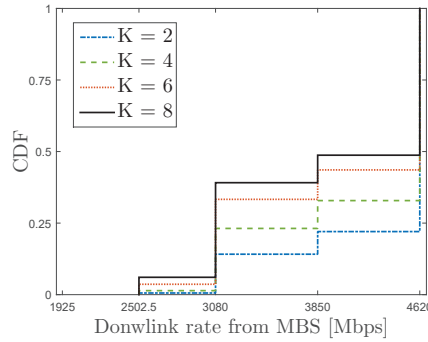
Fig. 23. Full network case, with $n = 15$ μBSs and interference of $\pi/8$ rads HPBW. CDF of downlink rates from the MBS whose links are used to relay traffic.
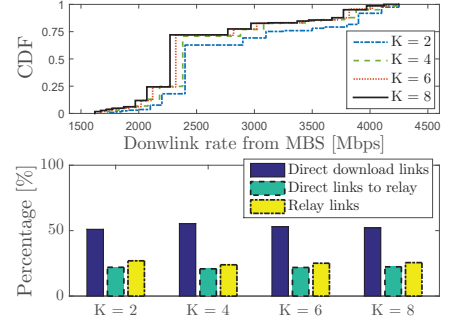
Fig. 24. Full network case, with $n = 15$ μBSs. Interference is as measured in [55] and rates depend on distance [57]. CDF of downlink rates from the MBS whose links are used to relay traffic (top); and usage of direct and relay links (bottom).

optimal scheduling cannot be implemented due to time constraints, one should always implement the `Resched` heuristic, which approximates the optimal better than `Greedy`. However, the complexity of `Resched` is quadratic with the network size, while `Greedy`'s is linear. Hence, only in those cases in which the time of the decision-making process is really tight, one would implement `Greedy`. Furthermore, for the few cases in which optimal solutions are computationally feasible, we have observed that relay reduces the makespan by a significant 35%. The results shown in Figs. 7-12 illustrate how enabling relay considerably mitigates the transmission bottleneck at the MBS. Indeed, as shown in Figs. 13 to 21, using relay is convenient with ideal and realistic antenna patterns, since it allows efficient spatial reuse with high probability to achieve high aggregate rates.

As we show in Fig. 22 for a typical case, direct download from the MBS only accounts for less than half of the aggregate utilization of links in the network. File download from the MBS to μBS relays occupies basically the same as μBS-to-μBS links, which indicates not only that relaying is convenient but also that our schemes do not blindly offload files to relays with inadequate connectivity quality to the destination μBS. `Resched` in fact tends to use relays that require a similar time to receive and retransmit the selected files. To provide some insights on which μBSs are usually selected and why, we show in Fig. 23 the CDF of downlink rates of those μBSs that act as relays. The CDF has a staircase shape because only discrete values of rates are possible, each corresponding to a given MCS. We observe in the figure that only μBSs with the best three MCS values are selected as relays, and more than 50% of relays use the maximum MCS.

The gain shown for our optimization can be impaired by non-ideal beam patterns. For instance, for the extreme case of real beam-patterns of cheap antennas as the ones studied in Figs. 19 to 21, we show in Fig. 24 that the CDF of relay rates is more spread and exhibit poorer statistics. For instance, the median for the full network case with 15 μBSs and $K = 8$ is 2.8 Gbps, against the 3.8 Gbps of the case of ideal antennas with HPBW=$\pi/8$ rads. However, the gain remains considerable as well as the fact that our optimization leads to dedicate about half of the link usage to relay, and to balance relay's in and out traffic, like in the case of ideal antennas. As a remark, while the use of cheap antennas could be common for 802.11-based inexpensive indoor devices, it is less reasonable to mount them on towers and outdoor deployments covering relatively large areas, which is where relay might be needed the more.

As a result of the lessons learnt on this research, we answer the questions raised at the introduction about the high complexity of the whole framework, the convenience of relaying and spatial reuse, and the selection of proper relays to know which μBSs relay traffic to which μBSs.

10 CONCLUSIONS AND OPEN PROBLEMS

In this paper we have defined the MMWBS problem for the compact and efficient scheduling of mmWave backhauls using an MILP formulation. We have proved that solving the problem is NP-hard because mmWave links incur a non-zero activation cost, due to antenna steering and signalling messages exchange typical of mmWave. We have derived and studied tight bounds for the length of the scheduling for a given set of download jobs, i.e., the *makespan*. We have shown that MMWBS cannot be approximated unless the interference can be neglected. In such case, based on linear programming, we are able to characterize the MMWBS problem analytically with constant-approximation guarantees using a scheduler running in polynomial time. We have also proposed practical heuristics, namely `Resched` and `Greedy`. We have evaluated the heuristics and validated the analysis by means of numerical simulations and, for small instances of the problem, compared heuristics to the optimal schedule computed with a Branch&Bound solver.

Our results show that simple heuristics show notable gains in small testable systems. On average, these heuristics find near-optimal solutions under different network topologies and base station settings, both with and without the effect of interference between transmitting mmWave links.

A study of MMWBS under other interference models, such as SINR [59], link-to-link– [60] and node-to-link affectance [61], [62], and conflict graphs [41], [42], with other objective functions, such as minimizing energy consumption or optimizing beamsteering (since we consider those procedures described in current standards), and the questions of how to achieve stability, low latency, or high throughput under adversarial or stochastic packet injections, are left for future work.

ACKNOWLEDGMENTS

REFERENCES

[1] 5G PPP Architecture Working Group, "View on 5G Architecture, version 1.0," Jul. 2016 (online at https://5g-ppp.eu/wp-content/uploads/2014/02/5G-PPP-5G-Architecture-WP-July-2016.pdf).

[2] T. S. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G. N. Wong, J. K. Schulz, M. Samimi, and F. Gutierrez, "Millimeter Wave Mobile Communications for 5G Cellular: It Will Work!" IEEE Access, vol. 1, pp. 335–349, 2013.

[3] W. Feng, Y. Li, D. Jin, L. Su, and S. Chen, "Millimetre-Wave Backhaul for 5G Networks: Challenges and Solutions," Sensors, vol. 16, no. 6, Jun. 2016.

[4] R. Radhakrishnan, W. W. Edmonson, F. Afghah, R. M. Rodriguez-Osorio, F. Pinto, and S. C. Burleigh, "Survey of Inter-Satellite Communication for Small Satellite Systems: Physical Layer to Network Layer View," IEEE Communications Surveys Tutorials, vol. 18, no. 4, pp. 2442–2473, Fourthquarter 2016.

[5] P. Puri, P. Garg, and M. Aggarwal, "Bi-directional relay-assisted FSO communication systems over strong turbulence channels with pointing errors," International Journal of Communication Systems, vol. 30, no. 4.

[6] B. Hajek and G. Sasaki, "Link scheduling in polynomial time," IEEE transactions on Information Theory, vol. 34, no. 5, pp. 910–917, 1988.

[7] A. Fernández Anta, D. R. Kowalski, M. A. Mosteiro, and P. W. H. Wong, "Schedulling Dynamic Parallel Workload of Mobile Devices with Access Guarantees," ACM Transactions on Parallel Computing (TOPC), vol. 5, no. 2, p. 10, 2018.

[8] A. Halper, M. A. Mosteiro, Y. Rossikova, and P. W. H. Wong, "Station Assignment with Reallocation," Algorithmica, vol. 81, no. 3, pp. 1096–1125, 2019.

[9] H. Shachnai, A. Voloshin, and S. Zaks, "Flexible bandwidth assignment with application to optical networks," Journal of Scheduling, vol. 21, no. 3, pp. 327–336, 2018.

[10] G. B. Mertzios, M. Shalom, P. W. Wong, and S. Zaks, "Online Regenerator Placement," Theory of Computing Systems, vol. 61, no. 3, pp. 739–754, 2017.

[11] E. Arribas and V. Mancuso, "Multi-path D2D leads to satisfaction," in 2017 IEEE 18th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM). IEEE, 2017, pp. 1–7.

[12] A. Asadi, V. Mancuso, and R. Gupta, "An SDR-based experimental study of outband D2D communications," in IEEE INFOCOM 2016-The 35th Annual IEEE International Conference on Computer Communications. IEEE, 2016, pp. 1–9.

[13] G. H. Sim, A. Loch, A. Asadi, V. Mancuso, and J. Widmer, "5G Millimeter-Wave and D2D Symbiosis: 60 GHz for Proximity-based Services," IEEE Wireless Communications Magazine, 2016.

[14] J. García-Rois, F. Gómez-Cuba, M. R. Akdeniz, F. J. González-Castaño, J. C. Burguillo, S. Rangan, and B. Lorenzo, "On the analysis of scheduling in dynamic duplex multihop mmWave cellular systems," IEEE Transactions on Wireless Communications, vol. 14, no. 11, pp. 6028–6042, 2015.

[15] D. Yuan, H.-Y. Lin, J. Widmer, and M. Hollick, "Optimal joint routing and scheduling in millimeter-wave cellular networks," in IEEE INFOCOM, 2018, pp. 1205–1213.

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

[17] H. Shokri-Ghadikolaei, C. Fischione, and E. Modiano, "On the accuracy of interference models in wireless communications," in IEEE ICC, 2016, pp. 1–6.

[18] Y. Niu, C. Gao, Y. Li, L. Su, D. Jin, and A. V. Vasilakos, "Exploiting device-to-device communications in joint scheduling of access and backhaul for mmWave small cells," IEEE Journal on Selected Areas in Communications, vol. 33, no. 10, pp. 2052–2069, 2015.

[19] Y. Niu, C. Gao, Y. Li, L. Su, D. Jin, Y. Zhu, and D. O. Wu, "Energy-efficient scheduling for mmWave backhauling of small cells in heterogeneous cellular networks," IEEE Transactions on Vehicular Technology, vol. 66, no. 3, pp. 2674–2687, 2017.

[20] L. Huang, M. Rong, L. Wang, Y. Xue, and E. Schulz, "Resource Allocation for OFDMA Based Relay Enhanced Cellular Networks," in proceedings of IEEE VTC2007-Spring, April 2007.

[21] L. Wang, Y. Ji, and F. Liu, "Joint Optimization for Proportional Fairness in OFDMA Relay-Enhanced Cellular Networks," in proceedings of IEEE WCNC, April 2010.

[22] M. Al-Rawi and R. Jäntti, "A dynamic TDD inter-cell interference coordination scheme for Long Term Evolution networks," in proceedings of IEEE PIMRC, Sept 2011.

[23] N. Ruangchaijatupon and Y. Ji, "Simple Proportional Fairness Scheduling for OFDMA Frame-Based Wireless Systems," in proceedings of IEEE WCNC, March 2008.

[24] I. Ahmed and A. Mohamed, "On the joint scheduling and intra-cell interference coordination in multi-relay LTE uplink," in proceedings of IEEE Globecom Workshops, Dec 2012.

[25] X. Zhang, X. Tao, Y. Li, and J. Lu, "QoS-aware scheduling with optimization of base-station power allocation in downlink cooperative OFDMA systems," EURASIP Journal on Wireless Communications and Networking, Oct 2013.

[26] B. Yu, S. Mukherjee, H. Ishii, and L. Yang, "Dynamic TDD support in the LTE-B enhanced Local Area architecture," in proceedings of IEEE Globecom Workshops, Dec 2012.

[27] W. Huang, X. Jia, and Y. Zhang, "Interference management and traffic adaptation of femto base station based on TD-LTE," International Journal of Future Generation Communication and Networking, vol. 7, no. 1, pp. 217–224, 2014.

[28] 3GPP, "Access, Evolved Universal Terrestrial Radio: Further Enhancements to LTE Time Division Duplex (TDD) for Downlink-Uplink (DL-UL) Interference Management and Traffic Adaptation," 2012.

[29] W. Guo and T. O'Farrell, "Relay Deployment in Cellular Networks: Planning and Optimization," IEEE Journal on Selected Areas in Communications, vol. 31, no. 8, pp. 1597–1606, August 2013.

[30] R. Mudumbai, S. K. Singh, and U. Madhow, "Medium Access Control for 60 GHz Outdoor Mesh Networks with Highly Directional Links," in proceedings of IEEE INFOCOM 2009, April 2009.

[31] S. Singh, M. N. Kulkarni, A. Ghosh, and J. G. Andrews, "Tractable Model for Rate in Self-Backhauled Millimeter Wave Cellular Networks," IEEE Journal on Selected Areas in Communications, vol. 33, no. 10, pp. 2196–2211, Oct 2015.

[32] F. Gómez-Cuba, S. Rangan, and E. Erkip, "Scaling laws for Infrastructure Single and multihop wireless networks in wideband regimes," in proceedings of IEEE ISIT, June 2014.

[33] Y. Ghasempour, C. R. da Silva, C. Cordeiro, and E. W. Knightly, "IEEE 802.11 ay: Next-generation 60 GHz communication for 100 Gb/s Wi-Fi," IEEE Communications Magazine, vol. 55, no. 12, pp. 186–192, 2017.

[34] C. R. da Silva, A. Lomayev, C. Chen, and C. Cordeiro, "Analysis and Simulation of the IEEE 802.11 ay Single-Carrier PHY," in IEEE ICC, 2018, pp. 1–6.

[35] Y. Li, J. Luo, W. Xu, N. Vucic, E. Pateromichelakis, and G. Caire, "A Joint Scheduling and Resource Allocation Scheme for Millimeter Wave Heterogeneous Networks," in WCNC. IEEE, 2017, pp. 1–6.

[36] S. Sreng, B. Escrig, and M. L. Boucheret, "Exact outage probability of a hybrid satellite terrestrial cooperative system with best relay selection," in 2013 IEEE ICC, June 2013, pp. 4520–4524.

[37] M. R. Bhatnagar and M. K. Arti, "Performance Analysis of AF Based Hybrid Satellite-Terrestrial Cooperative Network over Generalized Fading Channels," IEEE Communications Letters, vol. 17, no. 10, pp. 1912–1915, October 2013.

[38] M. K. Arti and V. Jain, "Relay selection-based hybrid satellite-terrestrial communication systems," IET Communications, vol. 11, no. 17, pp. 2566–2574, 2017.

[39] T. Nitsche, G. Bielsa, I. Tejado, A. Loch, and J. Widmer, "Boon and bane of 60 GHz networks: Practical insights into beamforming, interference, and frame level operation," in ACM CoNEXT, 2015.

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

[41] M. M. Halldórsson, G. Kortsarz, P. Mitra, and T. Tonoyan, "Spanning Trees With Edge Conflicts and Wireless Connectivity," in ICALP 2018, ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 107, 2018, pp. 158:1–158:15.

[42] S. A. Jafar, "Topological interference management through index coding," Transactions on Information Theory, vol. 60, no. 1, pp. 529–568, 2014.

[43] M. Jost, J. S. Gautam, L. G. Gomes, R. Reese, E. Polat, M. Nickel, J. M. Pinheiro, A. L. Serrano, H. Maune, G. P. Rehder et al., "Miniaturized liquid crystal slow wave phase shifter based on nanowire filled membranes," IEEE Microwave and Wireless Components Letters, vol. 28, no. 8, pp. 681–683, 2018.

[44] S. F. Jilani, M. O. Munoz, Q. H. Abbasi, and A. Alomainy, "Millimeter-Wave Liquid Crystal Polymer Based Conformal Antenna Array for 5G Applications," IEEE Antennas and Wireless Propagation Letters, vol. 18, no. 1, pp. 84–88, 2019.

[45] 3GPP, "LTE; Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer procedures," 3rd Generation Partnership Project (3GPP), TS 36.213, V 13.0.0.

[46] ——, "Evolved Universal Terrestrial Radio Access (E-UTRA) Medium Access Control (MAC) protocol specification," 3rd Generation Partnership Project (3GPP), TS 36.321, V 8.0.0.

[47] D. Zuckerman, "Linear Degree Extractors and the Inapproximability of Max Clique and Chromatic Number," Theory OF Computing, vol. 3, pp. 103–128, 2007.

[48] M. Cieliebak, S. Eidenbenz, A. Pagourtzis, and K. Schlude, "On the Complexity of Variations of Equal Sum Subsets," Nord. J. Comput., vol. 14, no. 3, pp. 151–172, 2008.

[49] L. G. Valiant, "The complexity of enumeration and reliability problems," SIAM Journal on Computing, vol. 8, no. 3, pp. 410–421, 1979.

[50] N. Karmarkar, "A new polynomial-time algorithm for linear programming," Combinatorica, vol. 4, no. 4, pp. 373–396, 1984.

[51] C. E. Shannon, "A theorem on coloring the lines of a network," Journal of Mathematics and Physics, vol. 28, no. 1, pp. 148–152, 1949.

[52] P. Sanders and D. Steurer, "An asymptotic approximation scheme for multigraph edge coloring," ACM Transactions on Algorithms (TALG), vol. 4, no. 2, p. 21, 2008.

[53] IEEE, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band," IEEE Std 802.11ad-2012.

[54] M. W. Rousstia, "Switched-beam antenna array design for millimeter-wave applications," PDEng Report, SAI-ICT, Eindhoven University of Technology, 2011.

[55] D. Steinmetzer, D. Wegemer, M. Schulz, J. Widmer, and M. Hollick, "Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.1 lad Devices," in Proceedings of CoNEXT, 2017.

[56] D. Steinmetzer, D. Wegemer, and M. Hollick. (2018) Talon Tools: The Framework for Practical IEEE 802.11ad Research. [Online]. Available: https://seemoo.de/talon-tools/

[57] S. K. Saha, H. Assasa, A. Loch, N. M. Prakash, R. Shyamsunder, S. Aggarwal, D. Steinmetzer, D. Koutsonikolas, J. Widmer, and M. Hollick, "Fast and infuriating: Performance and pitfalls of 60 ghz wlans based on consumer-grade hardware," in 2018 15th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON). IEEE, 2018, pp. 1–9.

[58] E. L. Lawler and D. E. Wood, "Branch-and-bound methods: A survey," Operations research, vol. 14, no. 4, pp. 699–719, 1966.

[59] C. Scheideler, A. W. Richa, and P. Santi, "An $O(\log n)$ dominating set protocol for wireless ad-hoc networks under the physical interference model," in ACM MobiHoc. ACM, 2008, pp. 91–100.

[60] T. Kesselheim, "Dynamic packet scheduling in wireless networks," in Proceedings of ACM PODC, 2012, pp. 281–290.

[61] D. R. Kowalski, H. Kudaravalli, and M. A. Mosteiro, "Ad-hoc Affectance-selective Families for Layer Dissemination," CoRR, vol. abs/1703.01704, 2017.

[62] D. R. Kowalski, M. A. Mosteiro, and K. Zaki, "Dynamic Multiple-Message Broadcast: Bounding Throughput in the Affectance Model," CoRR, vol. abs/1512.00540, 2015.

Edgar Arribas obtained his BSc in Mathematics from Universitat de València (UV) in 2015. After completing an MSc in Telematic Engineering at Universidad Carlos III de Madrid (UC3M) in 2016, he has become a PhD student at IMDEA Networks Institute and UC3M, funded by the MECD FPU15/02051 grant. He works on optimization of dynamic relay in wireless networks.

Antonio Fernández Anta is a Research Professor at IMDEA Networks. Previously he was a Full Professor at the Universidad Rey Juan Carlos (URJC) and was on the Faculty of the Universidad Politécnica de Madrid (UPM). He was a postdoc at MIT from 1995 to 1997, and spent sabbatical years at Bell Labs Murray Hill and MIT Media Lab. He has authored more than 200 scientific publications. He received his M.Sc. and Ph.D. from the University of SW Louisiana in 1992 and 1994, respectively. He is a Senior Member of ACM and IEEE.

Dariusz R. Kowalski is a professor in the School of Computer and Cyber Sciences at Augusta University, USA and at SWPS University of Social Sciences and Humanities, Warsaw, Poland. His research is in algorithm design and analysis, with particular interests in distributed and parallel algorithms, communication and mobility, fault tolerance, and applications in other sciences.

Vincenzo Mancuso is Research Associate Professor at IMDEA Networks, Madrid, Spain, and recipient of a Ramon y Cajal research grant of the Spanish Ministry of Science and Innovation. Previously, he was with INRIA (France), Rice University (USA) and University of Palermo (Italy), from where he obtained his Ph.D. in 2005. His research focus is on analysis, design, and experimental evaluation of opportunistic wireless architectures and mobile broadband services.

Miguel A. Mosteiro is a faculty member of the Computer Science Department at Pace University. He received his PhD degree in Computer Science from Rutgers University. His current research interests include algorithms for wireless networks embedded in IoT systems and distributed computing with harshly constrained resources.

Joerg Widmer is Research Professor and Research Director of IMDEA Networks in Madrid, Spain. Before, he held positions at DOCOMO Euro-Labs in Munich (Germany) and EPFL, Switzerland. His research focuses on wireless networks, ranging from extremely high frequency millimeter-wave communication and MAC layer design to mobile network architectures. He has authored more than 150 publications and three IETF RFCs, and holds 13 patents. He is senior member of IEEE and ACM.

Prudence W. H. Wong is a Professor in the Department of Computer Science, University of Liverpool, UK. She received her PhD in Computer Science from The University of Hong Kong. She has authored more than 100 peer-reviewed publications focusing on design and analysis of algorithms and optimization with applications on energy efficient scheduling, protocols and wireless networks. She is on the Editorial Board of Information Processing Letters, The Computer Journal and Algorithms. She is currently chairing the London Mathematical Society Computer Science Committee.