

Amir Ashtari Gargari, Andrea Ortiz, Matteo Pagin, Anja Klein, Matthias Hollick, Michele Zorzi, Arash Asadi “Safehaul: Risk-Averse Learning for Reliable mmWave Self-Backhauling in 6G Networks”, in *IEEE International Conference on Computer Communications INFO-COM*. 2023

©2023 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this works must be obtained from the IEEE.

Safehaul: Risk-Averse Learning for Reliable mmWave Self-Backhauling in 6G Networks

Amir Ashtari Gargari*, Andrea Ortiz[†], Matteo Pagin*, Anja Klein[†], Matthias Hollick[‡], Michele Zorzi*, Arash Asadi[§]

*Department of Information Engineering, University of Padova, Italy, {amirashtari, paginmatte, zorzi}@dei.unipd.it

[†]Communications Engineering Lab, TU Darmstadt, Germany, {a.ortiz, a.klein}@nt.tu-darmstadt.de

[‡]Secure Mobile Networking Lab (SEEMOO), TU Darmstadt, Germany, mhollick@seemoo.tu-darmstadt.de

[§]Wireless Communications and Sensing Lab (WISE), TU Darmstadt, Germany, aasadi@wise.tu-darmstadt.de

Abstract—Wireless backhauling at millimeter-wave frequencies (mmWave) in static scenarios is a well-established practice in cellular networks. However, highly directional and adaptive beamforming in today’s mmWave systems have opened new possibilities for self-backhauling. Tapping into this potential, 3GPP has standardized Integrated Access and Backhaul (IAB) allowing the same base station to serve both access and backhaul traffic. Although much more cost-effective and flexible, resource allocation and path selection in IAB mmWave networks is a formidable task. To date, prior works have addressed this challenge through a plethora of classic optimization and learning methods, generally optimizing a Key Performance Indicator (KPI) such as throughput, latency, and fairness, and little attention has been paid to the reliability of the KPI. We propose Safehaul, a risk-averse learning-based solution for IAB mmWave networks. In addition to optimizing average performance, Safehaul ensures reliability by minimizing the losses in the tail of the performance distribution. We develop a novel simulator and show via extensive simulations that Safehaul not only reduces the latency by up to 43.2% compared to the benchmarks, but also exhibits significantly more reliable performance, e.g., 71.4% less variance in achieved latency.

Index Terms—Millimeter-Wave Communication, Integrated Access and Backhaul (IAB), Self-backhauling, Wireless Backhaul

I. INTRODUCTION

The emergence of mmWave cellular systems created a unique opportunity for cellular operators to leverage a scalable and cost-effective approach to deal with network densification. The fact that mmWave base stations can support fiber-like data rates facilitates the use of the same base station for both access and backhaul traffic, a solution which in 3GPP parlance is referred to as Integrated Access and Backhaul (IAB). Consequently, 3GPP has included IAB in the standard [1], [2] covering the details on architecture, higher layer protocols, and the radio. Although Release 17 of 5G-NR defines the interfaces, architectures, and certain system parameters, the actual configuration and resource allocation is left to operators.

Traditional self-backhauled networks featured fixed-wireless links decoupled from access networks with static configurations. In contrast, IAB should account for the dynamic nature of the backhaul links (particularly in dense mmWave deployments) and their integration with the access network. Further, IAB allows the traffic to traverse several hops (i.e., base stations) to reach its destination, adding a new dimension

to the problem’s complexity. *In addition to the scheduling problem, an IAB network should: (i) solve the problem of path selection and link activation at the backhaul while considering inter-cell interference and (ii) decide on serving access or backhaul traffic depending on the access load and the ingress backhaul traffic from neighboring base stations.*

Prior work. Methodologically, the majority of the existing works [3]–[11] focus on classic optimization techniques to solve the above-mentioned problem. However, given the large number of parameters involved, such formulations often result in non-convex problems that are too complex for real-time operations, but are nonetheless valuable indicators as performance upper bounds. Recently, some works focus on more practical solutions which can be deployed in real networks [12]–[14]. Specifically, these works leverage Reinforcement Learning (RL) to tackle both resource allocation and/or path selection in IAB mmWave networks and demonstrate that RL-based solutions achieve real-time performance.

Regardless of the methodology, prior works mostly aim at maximizing the network capacity [3]–[10], minimizing latency [15], [16] and improving throughput fairness [4], [17]. Although optimizing capacity and latency is a challenging task by itself, network operators are often more *concerned about the reliability of such approaches*. For this reason many commercial products rely on *simplified* but reliable algorithms for resource allocation, despite their sub-optimal performance.

In this article, we propose Safehaul, a reinforcement learning-based solution for reliable scheduling and path selection in IAB mmWave systems under network dynamics. We use the concept of risk aversion, commonly used in economics [18], [19], to measure and enhance the reliability of Safehaul. The following summarizes our contributions:

- We model the scheduling and path selection problem in IAB mmWave networks as a multi-agent multi-armed bandit problem (Section III). We consider multiple fiber base stations simultaneously supporting many self-backhauled mmWave base stations. In our model, the self-backhauled base stations independently decide the links to activate. The consensus among the base stations is reached via standard-defined procedures (Section IV-C).
- We present the first solution to provide reliable performance in IAB-enabled networks (Section IV). Specifi-

cally, we investigate the joint minimization of the average end-to-end latency and its expected tail loss. To this aim, we propose Safehaul, a learning approach that leverages the coherent risk measure Conditional Value at Risk (CVaR) [18]. CVaR measures the tail average of the end-to-end latency distribution that exceeds the maximum permitted latency, thus ensuring the network's reliability.

- We provide a new means of simulating multi-hop IAB networks by extending NVIDIA's GPU-accelerated simulator Sionna [20] (Section V). Specifically, we add codebook-based analog beamforming capabilities for both uplink and downlink communications. Further, we extend Sionna by implementing system-level components such as layer-2 schedulers and buffers and Backhaul Adaptation Protocol (BAP)-like routing across the IAB network. We believe our IAB extensions will be instrumental for the open-source evaluation of future research on self-backhauled mmWave networks.
- Exploiting the above simulator, we evaluate and benchmark Safehaul against two state-of-the-art algorithms [16], [21]. The results confirm that Safehaul is significantly more reliable than the considered benchmarks, as it exhibits much tighter variance in terms of both latency (up to 71.4%) and packet drop rate (at least 39.1%). Further, Safehaul achieves up to 43.2% lower average latency and 11.7% higher average throughput than the reference schemes.

II. SYSTEM MODEL

We consider a cellular system with N base stations capable of self-backhauling and D base station with a fiber connection to the core network. Following 3GPP terminology, we refer to self-backhauled base stations as IAB-nodes (BS-nodes) and the fiber base stations as IAB-donors (BS-donors). Each IAB-node connects to the core network via a (multi-hop) wireless link to an IAB-donor. The sets of all BS-nodes and BS-donors are denoted by $\mathcal{N} = \{1, \dots, N\}$ and $\mathcal{D} = \{1, \dots, D\}$, respectively. The system works in a time-slotted fashion starting from time slot $i = 1$ until a finite time horizon I . All the time slots $i = 1, \dots, I$ have the same duration. The IAB-nodes are equipped with two RF chains. One RF chain is used exclusively for the communication with cellular users (access network), while the second RF chain is used for self-backhauling. In line with the 3GPP specification [22], we assume half-duplex self-backhauling, i.e., in each time slot i an IAB-node can either transmit, receive or remain idle.

We model the connections between the base stations in slot i as a graph $\mathcal{G}_i = \{\mathcal{V}, \mathcal{E}_i\}$, see Fig. 1. The set $\mathcal{V} = \mathcal{N} \cup \mathcal{D}$ of vertices is formed by all the BS-nodes and BS-donors in the system. The set \mathcal{E}_i of edges is composed of the available wireless links (n, l) between a BS-node $n \in \mathcal{N}$ and any BS (BS-donor or BS-node) $l \in \mathcal{V}$ in time slot i . Note that \mathcal{G}_i is not static. In a given time slot i , some links may be unavailable due to failure, blockage, or interference. Thus, only feasible wireless links are considered in the set \mathcal{E}_i . The path $X_{n,d}$ from BS-node n to any BS-donor d is a sequence of intermediate




Fig. 1: Example of a graph \mathcal{G}_i

links (n, l) . $X_{n,d}$ changes over time according to the traffic loads of the intermediate BS-node and the channel conditions. We model the activation of link (n, l) with the binary variable $x_{n,l,i}$. When $x_{n,l,i} = 1$, the link is activated and BS-node n transmits to BS $l \in \mathcal{V}$ in time slot i . $x_{n,l,i} = 0$ indicates that the link is deactivated and no transmission occurs.

Each BS-node n has a finite data buffer with capacity B_n^{\max} to store the backhaul data to be transmitted to any of the BS-donors. In each time slot i , BS-node n is characterized by its load and average queuing time. The load, denoted by $B_{n,i} \in \mathbb{N}$, indicates the number of data packets stored in the buffer at the beginning of time slot i . The average queuing time $t_{n,i}^q \in \mathbb{R}^+$ is the average number of time slots the current packets in the data buffer have been stored. Additionally, we denote by $t_{n,l,i}^{\text{tx}} \in \mathbb{R}^+$ and $M_{n,l,i} \in \mathbb{R}^+$ the transmission time from BS-node n to BS l in time slot i , and the amount of successfully transmitted data on the link, respectively. We define the maximum tolerable latency T_{\max} as the maximum time a packet can take from its source BS-node to any BS-donor. Any packet that is not delivered before T_{\max} milliseconds will be dropped. The average maximum end-to-end latency $\bar{T}_{n,d}$ from BS-node n to BS-donor d is the average, over the complete time horizon I , of the maximum delay a packet originating from BS-node n takes to reach any BS-donor d in time slot i . This is calculated as:

$$\bar{T}_{n,d} = \frac{1}{I} \sum_{i=1}^I T_{n,d,i}, \quad (1)$$

where $T_{n,d,i}$ is the maximum end-to-end latency among all the packets originating in BS-node n which reach BS-donor d in time slot i . $T_{n,d,i}$ is a sample of the random variable $T_{n,d}$ drawn from an unknown stationary probability distribution P that depends on the activated links $x_{n,l,i}$, the cellular user's mobility, the location of the BS-node n , the interference in the system, and the queue dynamics. Considering (1), the average end-to-end latency in the system \bar{T} is defined as

$$\bar{T} = \frac{1}{ND} \sum_{n=1}^N \sum_{d=1}^D \bar{T}_{n,d}. \quad (2)$$

III. PROBLEM FORMULATION

The joint minimization of the average end-to-end latency and the expected value of its tail loss in IAB-enabled networks is formulated in this section. We first introduce CVaR, the risk metric accounting for minimizing the events in which the end-to-end latency is higher than T_{\max} . Next, we formulate the optimization problem in the complete network.

A. Preliminaries on CVaR

Traditionally, latency minimization in IAB-enabled networks has focused on optimizing the expected value of a latency function [15], [16]. However, such an approach fails to capture the time variability of the latency distribution, thus leading to unreliable systems in which $T_{n,d,i} > T_{\max}$, for any $i = 1, \dots, I$, $n \in \mathcal{N}$ and $d \in \mathcal{D}$. For this purpose, we consider not only the average end-to-end latency \bar{T} in the system, but also its expected tail loss based on the CVaR [18], [23].

Having in mind that $T_{n,d}$ is a random variable, we assume it has a bounded mean on a probability space (Ω, \mathcal{F}, P) , with Ω and \mathcal{F} being the sample and event space, respectively. Using a risk level $\alpha \in (0, 1]$, the $\text{CVaR}_\alpha(T_{n,d})$ of $T_{n,d}$ at risk level α quantifies the losses that might be encountered in the α -tail. More specifically, it is the expected value of $T_{n,d}$ in its α -tail distribution [23]. Formally, $\text{CVaR}_\alpha(T_{n,d})$ is defined as [18]

$$\text{CVaR}_\alpha(T_{n,d}) = \min_{q \in \mathbb{R}} \left\{ q + \frac{1}{\alpha} \mathbb{E} [\max\{T_{n,d} - q, 0\}] \right\}, \quad (3)$$

where the expectation in (3) is taken over the probability distribution P . Note that lower $\text{CVaR}_\alpha(T_{n,d})$ results in higher reliability in the system because the expected end-to-end latency in the α -worst cases is low. Moreover, note that α is a risk aversion parameter. For $\alpha = 1$, $\text{CVaR}_\alpha(T_{n,d}) = \mathbb{E}[T_{n,d}]$ which corresponds to the traditional risk-neutral case. Conversely, for $\alpha = 0$, $\lim_{\alpha \rightarrow 0} \text{CVaR}_\alpha(T_{n,d}) = \sup\{T_{n,d}\}$. CVaR has been shown to be a coherent risk measure, i.e., it fulfills monotonicity, subadditivity, translation invariance, and positive homogeneity properties [24].

B. Optimization Problem

We jointly minimize the average end-to-end latency and its expected tail loss for each BS-node. For this purpose, we decide which of the (n, l) links to activate in each time slot i during the finite time horizon I . In the following, we formulate the optimization problem from the network perspective and consider the sum over all BS-nodes in the system. The latency minimization problem should consider three different aspects: (i) link activation is constrained by the half-duplex nature of self-backhauling, (ii) only data stored in the data buffers can be transmitted, and (iii) packet drop due to buffer overflow should be avoided. Formally, the problem is written as:

$$\begin{aligned} & \underset{\{x_{n,l,i}, i \in \{1, \dots, I\}\}}{\text{minimize}} && \sum_{n \in \mathcal{N}} \sum_{d \in \mathcal{D}} \bar{T}_{n,d} + \eta \text{CVaR}_\alpha(T_{n,f}) \\ & \text{subject to} && \end{aligned} \quad (4a)$$

$$\sum_{\substack{l \in \mathcal{V} \\ l \neq n}} x_{n,l,i} + \sum_{l \in \mathcal{N}} x_{l,n,i} = 1, \quad n \in \mathcal{N}, \quad (4b)$$

$$\sum_{j=1}^i B_{n,i} \geq \sum_{j=1}^i M_{n,l,i}, \quad n \in \mathcal{N}, l \in \mathcal{V}, \quad (4c)$$

$$\sum_{j=1}^i B_{l,j} + \sum_{j=1}^i M_{n,l,j} \leq B_l^{\max}, \quad n \in \mathcal{N}, l \in \mathcal{V}, \quad (4d)$$

$$x_{n,l,i} \in \{0, 1\}, \quad n \in \mathcal{N}, l \in \mathcal{V}. \quad (4e)$$

In (4a), $\eta \in [0, 1]$ is a weighting parameter to trade between minimizing the average end-to-end latency $\bar{T}_{n,d}$ and the expected loss of its tail. As the considered scenario is not static, solving (4) requires complete non-causal knowledge of the system dynamics during the complete time horizon I . However, in practical scenarios, knowledge about the underlying random processes is not available in advance. For example, the IAB-node's loads $B_{n,i}$ depend not only on the transmitted and received backhaul data, but also on the randomly arriving data from its users. Similarly, the amounts of transmitted data $M_{n,l,i}$ depend on the varying channel conditions of both BS n and l . As a result, the exact values of $T_{n,l,i}$, $B_{n,i}$ and $M_{n,l,i}$ are not known beforehand. For this reason, we present in Sec. IV Safehaul, a multi-agent learning approach to minimize in each BS-node the average end-to-end latency and the expected value of the tail of its loss.

IV. OUR PROPOSED SOLUTION: SAFEHAUL

In this section, we describe Safehaul, a multi-agent learning approach for the joint minimization of the average end-to-end latency and its expected tail loss in IAB mmWave networks. In Safehaul, each BS-node independently decides which links (n, l) to activate in every time slot i by leveraging a multi-armed bandit formulation. The consensus among the BS-nodes is reached by exploiting the centralized resource coordination and topology management role of IAB-donors [1, Sec. 4.7.1].

A. Multi-Armed Bandit Formulation

Multi-armed bandit is a tool well suited to problems in which an agent makes sequential decisions in an unknown environment [25]. In our scenario, each BS-node n decides, in each time slot i , which of the links (n, l) to activate without requiring prior knowledge about the system dynamics. The multi-armed bandit problem at BS-node n can be characterized by a set \mathcal{A}_n of actions and a set \mathcal{R}_n of possible rewards. The rewards $r_{n,i} \in \mathcal{R}_n$ are obtained in each time slot i as a response to the selected action $a_{n,i} \in \mathcal{A}_n$. The actions are the links that BS-node n can activate, and the rewards are a function of the observed latency. We define \mathcal{A}_n as

$$\mathcal{A}_n = \{(n, l), (m, n) | n, m \in \mathcal{N}, l \in \mathcal{V}\}, \quad (5)$$

where link (n, n) indicates that BS-node n remains idle. As blockages, overloads, or failures might render certain links (n, l) temporarily unavailable, we define the set $\mathcal{A}_{n,i} \subseteq \mathcal{A}_n$ of available actions in time slot i as

$$\mathcal{A}_{n,i} = \{(n, l), (l, n) | (n, l), (l, n) \in \mathcal{E}_i\}. \quad (6)$$

Selecting action $a_i = (n, l)$ in time slot i implies $x_{n,l,i} = 1$.

The rewards $r_{n,i}$ are a function of the end-to-end latencies $T_{n,d,i}$ and depend on whether at BS-node n a link (n, l) or (l, n) is activated. BS-node n is connected to the BS-donor via multi-hop wireless links. Consequently, $T_{n,d,i}$ cannot be

immediately observed when a link (n, l) , with $l \notin \mathcal{D}$ is activated. In fact, the destination BS-donor d might not even be known to BS-node n in time slot i . To overcome this limitation, we define the rewards $r_{n,i}$ as a function of the next-hop's estimated end-to-end latency $\hat{T}_{l,d,i}$ as

$$r_{n,i} = \begin{cases} t_{l,i}^q + t_{n,l,i}^{\text{tx}} + \hat{T}_{l,d,i}, & \text{for link } (n, l) \\ t_{n,i}^q + \hat{T}_{n,d,i}, & \text{for link } (l, n), \end{cases} \quad (7)$$

where $\hat{T}_{l,d,i}$ is calculated as

$$\hat{T}_{l,d,i} = \min_{(l,m) \in \mathcal{E}_i} \hat{T}_{l,m,i}. \quad (8)$$

B. Latency and CVaR Estimation

As given in (7) and (8), BS-node n learns which links (n, l) to activate by building estimates of the expected latency $\hat{T}_{n,l}$ associated to each of them. Let $K_{n,l,i} = \sum_{j=1}^i x_{n,l,i}$ be the number of times link (n, l) has been activated up to time slot i . $\hat{T}_{n,l}$ is updated using the sample mean as

$$\hat{T}_{n,l,i+1} = \frac{K_{n,l,i} \hat{T}_{n,l,i} + r_{n,i}}{K_{n,l,i} + 1}, \quad (9)$$

where the subindex i is introduced to emphasize that the estimate is built over time. The CVaR definition given in (3) requires $T_{n,d}$ which, as discussed before, is known a priori. Hence, we leverage the non-parametric estimator derived in [26] to estimate the CVaR of a link (n, l) . To this aim, let $\tilde{r}_n^1, \dots, \tilde{r}_n^{K_{n,l,i}}$ be all the rewards received up to time i ordered in a descending fashion, i.e., $\tilde{r}_n^1 \geq \dots \geq \tilde{r}_n^{K_{n,l,i}}$. The estimated $\widehat{\text{CVaR}}_i(n, l)$ at time slot i is calculated as [26]

$$\widehat{\text{CVaR}}_i(n, l) = \frac{1}{\lceil \alpha K_{n,l,i} \rceil} \sum_{k=1}^{\lceil \alpha K_{n,l,i} \rceil} \tilde{r}_n^k. \quad (10)$$

Using the estimates in (9) and (10), BS-node computes the value $Q_n(a_{n,i} = (n, l))$ associated to the selected action $a_n \in \mathcal{A}_n$, and defined as

$$Q_n(a_{n,i}) = \hat{T}_{n,l,i} + \eta \widehat{\text{CVaR}}_i(n, l). \quad (11)$$

Note that (11) is aligned with the objective function in (4a). Actions with an associated low value $Q_n(a_{n,i})$ lead to lower end-to-end latency and a low expected value on its tail.

C. Consensus

All the BS-nodes independently decide which links to activate based on their estimates of the end-to-end latency. As a consequence, conflicting actions may be encountered. A conflict occurs when two or more BS-nodes n and m aim at activating a link to a common BS l , $l \in \mathcal{V}$, i.e., $x_{n,l,i} = x_{m,l,i} = 1$. We reach consensus by first retrieving the buffer and congestion status of the various IAB-nodes, leveraging the related BAP layer functionality [1, Sec. 4.7.3]. With this information at hand, conflicts are resolved by prioritizing the transmission of the BS-node with the larger queuing times $t_{n,i}^q$ and loads $B_{n,i}$. Then, we let the IAB-donor mark as *unavailable* the time resources of the remaining base stations

Algorithm 1 Safehaul algorithm at each BS-node

Input: $\alpha, \eta, \mathcal{A}_n$

- 1: Initialize $\hat{T}_{n,l}$, $\widehat{\text{CVaR}}(n, l)$, and Q_n for all $(n, l) \in \mathcal{E}_1$
- 2: Set counters $K_{n,l} = 0$ and initial action $a_{n,1} = (n, n)$
- 3: **for** every time slot $i = 1, \dots, I$ **do**
- 4: perform action $a_{n,i}$ and observe reward $r_{n,i}$ ▷ Eq. (7)
- 5: increase counter $K_{n,l}$ by one
- 6: update latency estimate $\hat{T}_{n,l}$ ▷ Eq. (9)
- 7: update CVaR estimate $\widehat{\text{CVaR}}(n, l)$ ▷ Eq. (10)
- 8: update $Q_n(a_{n,i})$ ▷ Eq. (11)
- 9: select next action $a_{n,i+1}$ using ϵ -greedy ▷ Eq. (12)
- 10: share $a_{n,i+1}$, $t_{n,i}^q$ and $B_{n,i}$ with the other BS-nodes
- 11: if required, update $a_{n,i+1}$ to reach consensus ▷ Sec. IV-C
- 12: **end for**

with conflicting scheduling decisions [1, Sec. 10.9]. Note that as the learning is performed at each BS-node, only the link activation decision and the weighted sum of $t_{n,i}^q$ and $B_{n,i}$ are transmitted. Thus, low communication overhead is achieved.

D. Implementation of Safehaul

Here, we describe how the above-mentioned solution can be implemented in a real system. Specifically, we elaborate on the required inputs and the interactions among the different entities as well as the pseudo-code of Safehaul, see Alg. 1.

Safehaul is executed at each BS-node n . For its implementation, the network operator provides α, η and \mathcal{A}_n as an input. α is the risk level parameter that influences the level of reliability achieved in the system. Similarly, η controls the impact the minimization of the latency in the α -worst cases has on the overall performance. Both parameters, α and η , are set by the network operator depending on its own reliability requirements. The set \mathcal{A}_n depends on the considered network topology, which is perfectly known by the network operator. The set \mathcal{A}_n includes all links (n, l) and (l, n) to and from the first-hop neighbors of BS-node n .

The execution of Safehaul begins with the initialization of the latency and CVaR estimates, and the values Q of the actions in \mathcal{A}_n . Additionally, the counters $K_{n,l}$, that support the calculations of $\hat{T}_{n,l}$ and $\widehat{\text{CVaR}}(n, l)$, are initialized for all links in \mathcal{A}_n (lines 1-2). These parameters are updated and learnt throughout the execution of Safehaul. At time slot $t = 0$, no transmission has occurred and $B_{n,0} = 0$. Hence, BS-node n remains idle for the first time slot $i = 1$, i.e., $a_{n,1} = (n, n)$ (line 2). Next, and in each of the subsequent time slots $i \in \{1, \dots, I\}$, the selected action is performed and the corresponding reward is obtained (line 4). If BS-node n transmits in time slot i , i.e., $a_{n,i} = (n, l)$, the reward $r_{n,i}$ is sent by the receiving BS l through the control channel. If $a_{n,i} = (l, n)$, the reward $r_{n,i}$ depends, as given in (7), only on the current estimates at BS-node n and the status of its buffer $B_{n,i}$. With the observed reward $r_{n,i}$, the counter for action $a_{n,i}$ is increased and the latency and CVaR estimates are updated (lines 5-7). Using the new estimates (lines 6 and 7), the value $Q(a_{n,i})$ of the performed action $a_{n,i}$ is updated (line 8). The next action $a_{n,i+1}$ is then selected according to ϵ -greedy (line 9), which is a well-known method to balance the exploitation of links with estimated low latency, and the exploration of

unknown but potentially better ones. In ϵ -greedy, a random action $a_{n,i+1}$ from the set $\mathcal{A}_{n,i}$ is selected with probability $\epsilon \in [0, 1]$. With probability $(1 - \epsilon)$, instead, the action that yields the estimated lowest value is chosen, i.e.,

$$a_{n,i+1} = \begin{cases} \text{randomly selected action from } \mathcal{A}_{n,i}, & \text{if } x \leq \epsilon \\ \operatorname{argmax}_{b_n \in \mathcal{A}_{n,i}} Q_n(b_n), & \text{if } x > \epsilon, \end{cases} \quad (12)$$

where x is a sample taken from a uniform distribution in the interval $[0, 1]$. Once the action $a_{n,i+1}$ is selected, it is shared with other BS-nodes in the network along with $t_{n,i}^a$ and $B_{n,i}$ (line 10). As described in Section IV-C, this goes through the control channel. If conflicts arise, consensus is reached by prioritizing the transmission of the BS-node with the largest loads and queuing times (line 11).

The regret of Safehaul is defined as the expected loss caused by the fact that the optimal action is not always selected [27]. For brevity, we omit the regret analysis of Safehaul. Nevertheless, a regret bound can be derived following an approach similar to the work in [28] but including the ϵ -greedy considerations [27]. Moreover, the bound should account for the fact that the probability of not selecting an optimal action also depends on the actions of the other BS-nodes.

V. SIMULATION SETUP

Given the lack of access to actual 5G (and beyond) network deployments, prior works mostly rely on *home-grown* simulators for performance evaluation. Although a valid approach, these simulators often cannot fully capture the real network dynamics, introducing strong assumptions in the physical and/or the upper layers of the protocol stack. Until very recently, the most complete simulator for IAB networks was a system-level simulator [29] developed as an extension of the ns-3 *mmWave* module [30]. However, despite accurate modeling of the IAB protocol stack, it is currently behind the latest IAB specifications¹. Moreover, the ns-3 IAB extension is unsuitable for large simulations with hundreds of nodes due to reliance on an older version of the *mmWave* module. Therefore, in our work we opt for Sionna [20], which is an open-source GPU-accelerated toolkit based on TensorFlow. The tensor-based implementation supports the native integration of neural networks and prototyping complex communication systems.

However, unlike the aforementioned ns-3 module, Sionna is a physical layer-focused simulator that does not explicitly model 5G networks, thus lacking the characterization of the 5G-NR upper-layer protocol stack. Hence, we extend Sionna by including the system-level functionalities such as MAC-level scheduling and RLC-level buffering. Furthermore, since Sionna exhibits slight differences compared to the 5G-NR physical layer, we extend Sionna's physical layer model [20] with the 5G-NR procedures. In the following, we describe the details of our extensions, which we make publicly available².

¹For instance due to the assumption of layer-3 (instead of layer-2) relaying at the IAB-nodes which was based on a draft version of the TR 38.874 [31].

²https://github.com/TUDA-wise/safehaul_infocom2023




Fig. 2: Schematic of the hierarchical beam management procedure. First, the general direction is estimated using wide beams (top). Then, the search is refined using the narrow beams codebook.

A. Extensions to Sionna's physical layer module

In this section, we describe the physical layer modification that were necessary to evaluate IAB scenarios using Sionna.

1) *Codebook-based Beamforming*: Sionna's native beamforming only supports Zero-Forcing (ZF) pre-coding in down-link. Therefore, as a first step, we extend Sionna by implementing an NR-like codebook-based analog beamforming both at the transmitter and at the receiver. Specifically, we assume that the beamforming vectors at the transmitter w_{tx} and at the receiver w_{rx} are a pair of codewords selected from a predefined codebook. The codebook is computed by defining a set of beam directions $\{\omega_{n,m}\}$ which scans a given angular sector with a fixed beamwidth. The steering vector $a_{n,m}$ corresponding to direction $\omega_{n,m}$ can be computed as:

$$a_{n,m} = \left[1, \dots, e^{j \frac{2\pi}{\lambda} d(i_H \sin \alpha_n \sin \beta_m + i_V \cos \beta_m)}, \dots, e^{j \frac{2\pi}{\lambda} d((N_H - 1) \sin \alpha_n \sin \beta_m + (N_V - 1) \cos \beta_m)} \right]^T, \quad (13)$$

where N_H and N_V are the number of horizontal and vertical antenna elements, respectively. The horizontal and vertical index of a radiating element is denoted by $i_H \in \{0, \dots, N_H - 1\}$ and $i_V \in \{0, \dots, N_V - 1\}$, respectively. α_n and β_m represent the azimuth and elevation angles of $\omega_{n,m}$. Next, we define the codebook as the set $\{(\sqrt{N_H N_V})^{-1} a_{n,m}\}$.

In line with the 5G-NR beam management procedure [32], we assume the lack of complete channel knowledge, i.e., the communication endpoints do not know the corresponding channel matrix. Accordingly, an exhaustive search is conducted to identify the best pair of codewords resulting in the highest Signal to Interference plus Noise Ratio (SINR). Specifically, we leverage a hierarchical search [33], in which the communication pairs first perform a wide-beam search in which the transmitter and receiver approximate the direction of communication, see Fig. 2. Next, the beamforming direction is fine-tuned through a beam refinement procedure going through a codebook with narrow beams. Consequently, we employ two types of codebooks, one with wide beams for sector sweep and another with narrow beams for beam refinement.

2) *SINR Computations*: Since Sionna does not natively calculate the SINR, we add this functionality to the simulator to better model the impact of interference in our simulations. We compute the SINR experienced by Transport Blocks (TBs) by combining the power of the intended signal with that of

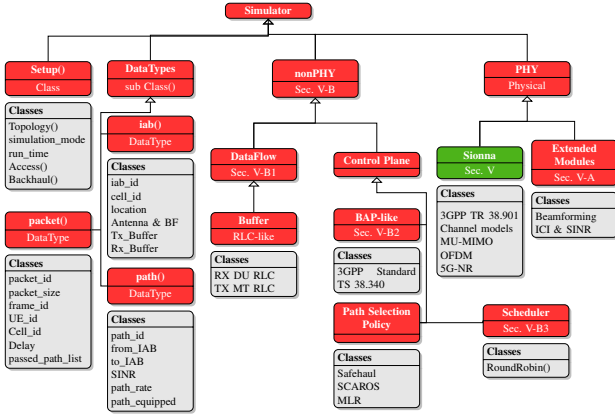


Fig. 3: Overall design of our Sionna's extension. The red blocks represent our additions to the baseline simulator, i.e., Sionna [20].

the interferers and of the thermal noise. Specifically, we first compute the power P_n of the intended signal at receiver n over frequency f and at time slot i . Then, we obtain the overall interference power by leveraging the superposition principle and summing the received power from all other interfering base stations $P_m(i, f)$ where $m \neq n$. For the purposes of this computation, we assume that each interferer employs the beamforming vector yielding the highest SINR towards its intended destination. Similarly, the transmitter and receiver use the beamforming configuration estimated via the hierarchical search procedure. Finally, the SINR is $\gamma_n(i, f) = \frac{P_n(i, f)}{\sum_{m \neq n} P_m(i, f) + \sigma(i, f)}$ where $\sigma(t, f)$ is the thermal noise at the receiver.

B. System-level extensions to Sionna

As mentioned, Sionna is mainly a physical layer simulator. However, to get closer to IAB networks as specified in Rel. 17, we have extended Sionna by implementing a selection of system-level features. To such end, we introduced a discrete-event network simulator for modeling IAB networks. This system-level extension operates on top of Sionna and provides basic functionalities such as a Medium Access Control (MAC)-level scheduler, layer-2 buffers, and data flow and path selection mechanisms. Our simulator, as depicted in Fig. 3, generates a variety of system-level KPIs such as latency, throughput, and packet drop rate.

1) *Data Flow and buffer*: 3GPP has opted for a layer-2 relaying architecture for IAB-nodes where hop-by-hop Radio Link Control (RLC) channels are established. This enables retransmissions to take place just over the afflicted hops, thus preventing the need for traversing the whole route from the IAB-donor whenever a physical layer TB is not decoded. Consequently, this design results in a more efficient recovery from transmission failures and reduces buffering at the communication endpoints [34]. To mimic this architecture, we have implemented RLC-like buffers at each base station. Specifically, each IAB-node features layer-2 buffers for both received and transmitted packets. For instance, the data flow for an uplink packet is the following. The User Equipment




Fig. 4: Locations of the 223 BS-nodes and BS-donors in Manhattan, NYC.

(UE) generates packets and sends a transmission request to the base station. Consequently, the scheduler allocates OFDM symbols for this transmission, which is eventually received and stored at the RX buffer of its Distributed Unit (DU). Next, the packet is placed into the TX buffer to be forwarded to the suitable next hop IAB-node. This procedure is repeated until the packet crosses all the wireless-backhaul hops and reaches the IAB-donor. Note that the packet can be dropped due to a violation of latency constraints or interference.

2) *Backhaul Adaptation Protocol*: To manage routing within the wireless-backhaunched network, the 3GPP introduced the BAP, i.e., an adaptation layer above RLC which is responsible for packet forwarding between the IAB-donor and the access IAB-nodes [35]. Our simulator mimics this by associating each IAB-node to a unique BAP ID. Moreover, we append a BAP routing ID to each packet at its entry point in the Radio Access Network (RAN) (i.e., the IAB-donor and the UEs for DL and UL data, respectively). Then, this identifier is used to discern the (possibly multiple) routes toward the packet's intended destination [35]. The choice of the specific route is managed by Safehaul.

3) *Scheduler*: Finally, we implement a MAC-level scheduler which operates in a Time Division Multiple Access (TDMA) mode. The scheduler periodically allocates the time resources to backhaul or access transmissions in a Round-Robin fashion. Specifically, each cell first estimates the number of OFDM symbols needed by each data flow by examining the corresponding buffer. Then, the subframe's OFDM symbols are equally allocated to the users. If a user requires fewer symbols to transmit its complete buffer, the excess symbols (the difference between the available slot length and the needed slot length) are dispersed to the other active users.

VI. PERFORMANCE EVALUATION

In our simulations, we consider a realistic cellular base station deployment in Manhattan, New York City³. Specifically, we collect the locations of $N = 223$ 5G-NR base stations in an area of 15 Km² as depicted in Fig. 4. The detailed simulation parameters are provided in Table I. We used the channel model outlined by 3GPP in TR 38.901 [36], which

³The locations correspond to the network of T-Mobile, as it has the largest deployment among the operators.




Fig. 5: Average network performance for 100 UEs and 80 Mbps per-UE source rate (Scenario 1).

TABLE I: Simulation parameters.

Parameter	Value
Carrier frequency and bandwidth	28 GHz and 400 MHz
IAB RF chains	2 (1 access + 1 backhaul)
Pathloss model	UMi-Street Canyon [36]
Number of BS-nodes N	223
Source rate	{40, 80} Mbps
IAB Backhaul and access antenna array	8H×8V and 4H×4V
UE antenna array	4H×4V
IAB and UE height	15 m and 1.5 m
IAB antenna gain	33 dB
Noise power	10 dB
Risk level α	0.1
Reliability weight factor η	1

provides a statistical channel model for 0.5-100 GHz, and analyzed the "Urban Micro (UMi)-StreetCanyon" scenario.

Benchmarks. To provide better insights on the performance of Safehaul, we replicate two approaches from the state of the art: (i) Scalable and Robust Self-backhauling Solution (SCAROS), a learning-based approach that minimizes the average latency in the network [16], and (ii) Maximum Local Rate (MLR), a greedy approach aiming to maximize throughput by selecting the links with the highest data rate.

Our evaluation consists of four scenarios, in which we study the convergence of the algorithms to a steady state, the number of IAB-nodes, the number of IAB-donors, and the impact of risk aversion. When demonstrating the results, we show the average throughput, latency, and packet drop rate per UE. Furthermore, we show the statistical variance of the obtained results using candlesticks which include the max, min, mean, and 10 and 90 percentiles of the achieved performance.

A. Scenario 1: Average Network Performance

Analyzing the performance of the algorithms as a function of time is crucial to determine the convergence speed of the learning-based techniques, i.e., Safehaul and SCAROS. Hence, in Fig. 5 we show the average network performance over time for three metrics: latency, throughput, and packet drop rate.

In Fig. 5a, we can observe that Safehaul rapidly converges to an average latency of approximately 8.6 ms which is 12.2% and 43.4% lower than the latency of SCAROS and MLR, respectively. The high performance of Safehaul stems from

the joint minimization of the average latency and the expected value of its tail loss, which results in avoiding risky situations where latency goes beyond T_{\max} . This is not the case for SCAROS where we observe a high peak in the latency before convergence, i.e., in between zero and 1000 ms. *It is exactly the avoidance of such transients in Safehaul that leads to higher reliability in the system.* The reliability offered by Safehaul allows operators to deploy self-backhauling in an online fashion and without disrupting the network operation. Moreover, it protects the networks from the transients that may arise from changes in the network topology. The performance of MLR is constant throughout the simulation, as it is not designed as an adaptive algorithm.

Figure 5b shows that the risk-aversion capabilities of Safehaul have no negative impact on the average throughput in the network. The performance of Safehaul is comparable to that of SCAROS, approximately 79.3 Mbps, and 11.7% larger than the performance of MLR.

The performance shown in Figure 5c is consistent with the behavior observed in Figure 5a. As Safehaul additionally minimizes the α -worst latency, it achieves the lowest packet drop rate, compared to the reference schemes, namely, 16.6% and 25.0% lower than SCAROS and MLR, respectively.

B. Scenario 2: Impact of the Network Size

In Fig. 6 we evaluate the reliability of the three considered approaches for different network sizes. Specifically, we vary the number of BS-nodes starting from 25 up to 100. At the same time, we increase the load in the network by increasing the number of UEs. From the figures, we can clearly see that Safehaul consistently achieves a lower variation compared to the reference schemes. This verifies that Safehaul achieves the intended optimization goal, i.e., the joint minimization of the average end-to-end delay and its expected tail loss.

Fig. 6a shows that Safehaul is able to maintain an almost constant latency as the number of BS-nodes increases. Specifically, the variation of latency with Safehaul is 56.1% and 71.4% less than with SCAROS and MLR, respectively. Furthermore, Safehaul achieves 11.1% and 43.2% lower latency compared to SCAROS and MLR. MLR's high variance is due




Fig. 6: Network performance for $\{25, 50, 75, 100, 200\}$ IAB-nodes and 40 Mbps per-UE source rate (Scenario 2).

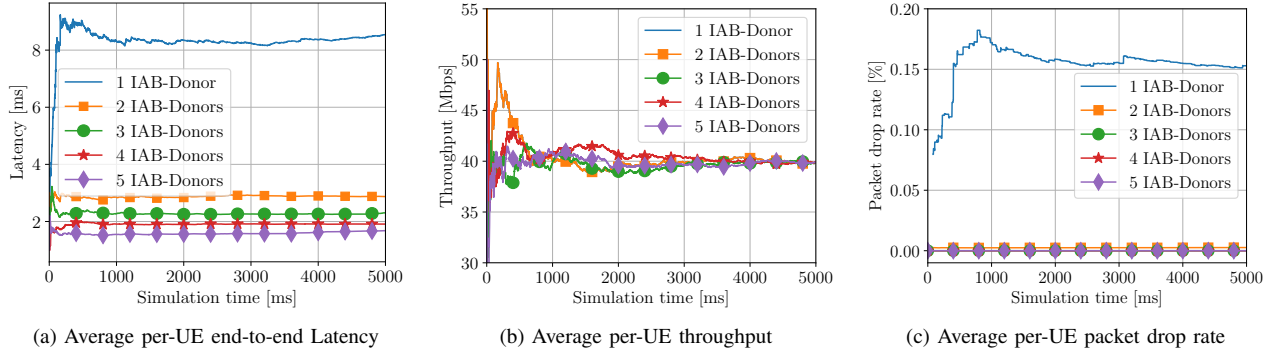


Fig. 7: Network performance for 100 UEs and 40 Mbps per-UE source rate, versus the number of IAB-donors (Scenario 3).

to a lack of adaptation capabilities, hence its latency variance is governed by the network's underlying random processes.

As shown in Fig. 6b, the average throughput of the learning-based approaches Safehaul and SCAROS remains constant for the different values of network size. However, the lowest variation in the throughput is achieved by Safehaul, i.e., only 0.40 compared to 0.51 and 0.79 in the benchmark schemes. Such behavior corroborates Safehaul's reliability capabilities.

The packet drop rate for different numbers of IAB-nodes is shown in Fig. 6c. Safehaul not only consistently outperforms the reference schemes, but also has the minimum variation in the results (by at least 39.1% compared to benchmarks). Considering the largest network size and load, i.e., 200 BS-nodes and 400 UEs, Safehaul achieves 11.2% and 24.9% lower packet drop rate compared to SCAROS and MLR, respectively.

C. Scenario 3: Impact of the number of IAB-donors

Although the benchmark schemes do not support multi-IAB-donors, Safehaul is designed to accommodate such scenarios. In Fig. 7, we investigate the impact of the number of IAB-donors on Safehaul. Specifically, the network load is constant in this scenario, i.e., the number of UEs is fixed.

We observe in Fig. 7a that the highest latency is experienced when only one IAB-donor is present in the network. This stems from the tributary effect of self-backhauling where the traffic flows towards a central entity which itself can become a bottleneck. As the number of IAB-donors increases, the traffic flow is more evenly distributed, resulting in lower

latency. Specifically, from an average latency of 8.2 ms for $D = 1$, to an average latency of 1.7 ms when $D = 5$. As mentioned, since the load is constant in this scenario, the average throughput also remains constant for all different numbers of IAB-donors, see Fig. 7b. However we should highlight that Safehaul's learning speed is maintained for the different values of D . This is an important design feature of Safehaul because having more BS-donors means that the number of paths a BS-node has to the core network increases exponentially. From a learning perspective, such increment implies a larger action set and a lower learning speed. Safehaul avoids this problem by learning the average latency based on the estimates of its neighbors and not on the complete paths to the BS-donors. Finally, Fig. 7c shows that increasing the number of BS-donors significantly reduces the packet drops, which also stems from a better distribution of traffic flows in the network, as observed in Fig. 7a.

D. Scenario 4: Impact of the risk parameter α

The definition of losses in the tail of the latency distribution is controlled by the risk level parameter α . Its impact on the average latency is shown in Fig. 8, where an increasing behavior is observed for $\alpha \leq 0.7$. The lowest latency is achieved for $\alpha = 0.1$, which corresponds to the most risk-averse, and therefore the most reliable, case out of all the considered ones. As α grows, the performance of Safehaul tends to that of the risk-neutral case.




Fig. 8: Average latency for 100 UEs and 20 Mbps per-UE source rate, versus the risk level α (Scenario 4)

VII. RELATED WORK

Self-backhauling wireless networks have been studied in different contexts. Ranging from the so-called Heterogeneous Networks (HetNets) and IAB 5G New Radio (NR) systems, to Cloud Radio Access Networks (C-RANs), each has considered a different set of premises and optimization goals. In this section, we review the related work in the context of basic assumptions and their optimization goals.

Ideal backhaul links. Numerous works assume either an *infinite or fixed capacity backhaul link*. This is often motivated due to the presence of a wired fiber link between the Small Base Stations (SBSs) and the Macro Base Station (MBS) [3], [5]–[7]. Indeed, most of these works consider a scenario where a centralized Baseband Unit (BBU) is connected to several Remote Radio Heads (RRHs), i.e., radios which lack signal processing capabilities [3], [5], [6]. In particular, the authors of [6] consider an even more complex C-RAN scenario where RRHs feature caching and signal processing capabilities. However, in an IAB context it is fundamental to consider *limited-rate, time-varying backhaul channels* and to study the impact of such limitations on the performance of the RAN.

Constrained topologies. It is often assumed that self-backhauling networks have a *specific topology*. This assumption usually simplifies the problem and makes it tractable and/or solvable in polynomial time. For instance, the authors of [8], [9], [12] assume a single-hop network where each SBS is directly connected to the MBS. In [10], a k -ring deployment is considered, i.e., a topology where a single IAB-donor provides backhaul connectivity to k rings of IAB-nodes. Even though this topology can be used to model networks with arbitrary depth, it maintains a symmetric load for each node, an assumption which generally does not hold in real networks. In fact, the 3GPP does not impose any limits on the number of IAB-nodes which can be connected to a given IAB-donor, nor does it set an upper bound on the number of wireless hops from the latter to other wireless-backhauling base stations [22]. Accordingly, in our problem formulation we consider IAB networks with an *arbitrary number of nodes and an arbitrary maximum number of wireless hops* between MBSs and SBSs.

Simplistic traffic models. Some works either assume a *full buffer traffic model and/or impose flow conservation constraints*. In particular, the authors of [7], [37] consider systems where the capacity of each link can always be fully exploited thanks to the presence of *infinite data to transmit at*

each node. However, in actual IAB deployments the presence of packets at the MBSs and SBSs is conditioned on the *status of their RLC buffers and, in turn, on the previous scheduling decisions*. Moreover, *packets can actually be buffered at the intermediate nodes*, thus preventing the need for transmitting a given packet in consecutive time instants along the whole route from the IAB-donor to the UEs (or vice versa).

Optimization goals. The works in the literature focus on different optimization goals. Therefore, they prioritize different network metrics. For instance, the authors of [38] aim to optimize the beam alignment between MBSs and SBSs. Instead, the work of [4] aims to compute the optimal user-to-base-station association. However, they neglect backhaul associations and focus on the access only. In [4], [8], [37], [39] the objective function is a function of the users data-rate. In particular, the authors of [37] optimize the max-min user throughput, arguing that such a metric better captures the performance of the bottleneck links. In [15], the average rate of each link is maximized under bounded delay constraints. In our work, we focus on reliability by minimizing not only the average end-to-end delay, but also the expected value of the worst-case performance. The work closest to this article is SCAROS [16], a learning-based latency-aware scheme for resource allocation and path selection in self-backhauling networks. Assuming a single IAB-donor, the authors study arbitrary multi-tier IAB networks considering the impact of interference and network dynamics. In contrast, we aim at enhancing the reliability of the IAB-network by jointly minimizing the average end-to-end delay and its expected tail loss.

VIII. CONCLUSION

In this work, we proposed the first reliability-focused scheduling and path selection algorithm for IAB mmWave networks. We illustrated that our RL-based solution can cope with the network dynamics including channel, interference, and load. Furthermore, we demonstrated that Safehaul not only exhibits highly reliable performance in the presence of the above-mentioned network dynamics, but also outperforms the benchmark schemes in terms of throughput, latency and packet-drop rate. The reliability of Safehaul stems from the joint minimization of the average latency and the expected value of its tail losses. The latter is achieved by leveraging CVaR as a risk metric.

Reliability is a highly under-explored topic that definitely deserves more investigation. Some interesting research directions are the maximization of reliability under the assumption of statistical system knowledge, or the evaluation of the network's reliability when the functionality of the BAP layer is compromised. Furthermore, our system-level extension to Sionna can be further developed to support an arbitrary number of RF chains and in-band backhauling, allowing more extensive investigation of IAB protocols and architectures.

IX. ACKNOWLEDGEMENT

This research was partly funded by the DFG within the mm-Cell project and the Collaborative Research Center 1053

MAKI, by the LOEWE initiative (Hesse, Germany) within the emergenCITY center, the BMBF through the Open6GHub project and by the European Commission through Grant No. 861222 (H2020 ITN MINTS project).

REFERENCES

- [1] 3GPP, “NR; Overall description; Stage-2,” Technical Specification (TS) 38.300, Jun. 2022, version 17.1.0.
- [2] —, “NR; Integrated Access and Backhaul (IAB) radio transmission and reception,” Technical Specification (TS) 38.174, Jun. 2022, v.17.1.0.
- [3] C. Pan, H. Zhu, N. J. Gomes, and J. Wang, “Joint precoding and RRH selection for user-centric green MIMO C-RAN,” *IEEE Trans. Wireless Commun.*, vol. 16, no. 5, pp. 2891–2906, Mar. 2017.
- [4] A. Alizadeh and M. Vu, “Load balancing user association in millimeter wave MIMO networks,” *IEEE Trans. Wireless Commun.*, vol. 18, no. 6, pp. 2932–2945, Mar. 2019.
- [5] X. Huang, G. Xue, R. Yu, and S. Leng, “Joint scheduling and beamforming coordination in cloud radio access networks with QoS guarantees,” *IEEE Trans. Veh. Technol.*, vol. 65, no. 7, pp. 5449–5460, Aug. 2015.
- [6] H. T. Nguyen, H. D. Tuan, T. Q. Duong, H. V. Poor, and W.-J. Hwang, “Nonsmooth optimization algorithms for multicast beamforming in content-centric fog radio access networks,” *IEEE Trans. Signal Process.*, vol. 68, pp. 1455–1469, Jan. 2020.
- [7] M. E. Rasekh, D. Guo, and U. Madhow, “Interference-aware routing and spectrum allocation for millimeter wave backhaul in urban picocells,” in *Proc. IEEE Allerton*, Sep. 2015.
- [8] G. Kwon and H. Park, “Joint user association and beamforming design for millimeter wave UDN with wireless backhaul,” *IEEE J. Sel. Areas Commun.*, vol. 37, no. 12, pp. 2653–2668, Oct. 2019.
- [9] A. Pizzo and L. Sanguinetti, “Optimal design of energy-efficient millimeter wave hybrid transceivers for wireless backhaul,” in *15th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*. IEEE, Jun. 2017.
- [10] M. N. Kulkarni, A. Ghosh, and J. G. Andrews, “Max-min rates in self-backhauled millimeter wave cellular networks,” *arXiv preprint arXiv:1805.01040*, 2018.
- [11] L. F. Abanto-Leon, A. Asadi, A. Garcia-Saavedra, G. H. Sim, and M. Hollick, “Radiorchestra: Proactive management of millimeter-wave self-backhauled small cells via joint optimization of beamforming, user association, rate selection, and admission control,” *IEEE Transactions on Wireless Communications*, vol. 22, no. 1, Jan. 2023.
- [12] W. Lei, Y. Ye, and M. Xiao, “Deep reinforcement learning-based spectrum allocation in integrated access and backhaul networks,” *IEEE Trans. on Cogn. Commun. Netw.*, vol. 6, no. 3, pp. 970–979, May 2020.
- [13] B. Zhang, F. Devoti, I. Filippini, and D. De Donno, “Resource allocation in mmWave 5G IAB networks: A reinforcement learning approach based on column generation,” *Computer Networks*, vol. 196, p. 108248, Jun. 2021.
- [14] Q. Cheng, Z. Wei, and J. Yuan, “Deep Reinforcement Learning-based Spectrum Allocation and Power Management for IAB Networks,” in *Proc. IEEE ICC Workshops*, Jul. 2021.
- [15] T. K. Vu, C.-F. Liu, M. Bennis, M. Debbah, and M. Latva-Aho, “Path selection and rate allocation in self-backhauled mmwave networks,” in *Proc. IEEE WCNC*, Jun. 2018.
- [16] A. Ortiz, A. Asadi, G. H. Sim, D. Steinmetzer, and M. Hollick, “SCAROS: A Scalable and Robust Self-Backhauling Solution for Highly Dynamic Millimeter-Wave Networks,” *IEEE J. Sel. Areas Commun.*, vol. 37, no. 12, pp. 2685–2698, Oct. 2019.
- [17] M. Pagin, T. Zugno, M. Polese, and M. Zorzi, “Resource Management for 5G NR Integrated Access and Backhaul: A Semi-Centralized Approach,” *IEEE Trans. Wireless Commun.*, vol. 21, no. 2, pp. 753–767, Jul. 2022.
- [18] R. T. Rockafellar and S. Uryasev, “Optimization of conditional value-at-risk,” *Journal of Risk*, vol. 2, pp. 21–41, 2000.
- [19] H. Levy, *Stochastic Dominance*. Springer Science+Business Media New York, 1998.
- [20] J. Hoydis, S. Cammerer, F. A. Aoudia, A. Vem, N. Binder, G. Marcus, and A. Keller, “Sionna: An Open-Source Library for Next-Generation Physical Layer Research,” *arXiv preprint arXiv:2203.11854*, Mar. 2022.
- [21] M. Polese, M. Giordani, A. Roy, D. Castor, and M. Zorzi, “Distributed path selection strategies for integrated access and backhaul at mmWaves,” in *Proc. IEEE Globecom*, Dec. 2018.
- [22] 3GPP, “NR; Study on integrated access and backhaul,” Technical Specification (TS) 38.874, Jan. 2019, v.16.0.0.
- [23] R. Rockafellar and S. Uryasev, “Conditional value-at-risk for general loss distributions,” *Journal of Banking & Finance*, vol. 26, no. 7, pp. 1443–1471, 2002.
- [24] G. C. Pflug, *Some Remarks on the Value-at-Risk and the Conditional Value-at-Risk*. Boston, MA: Springer US, 2000, pp. 272–281. [Online]. Available: https://doi.org/10.1007/978-1-4757-3150-7_15
- [25] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. MIT Press, 2018.
- [26] S. X. Chen, “Nonparametric Estimation of Expected Shortfall,” *Journal of Financial Econometrics*, vol. 6, no. 1, pp. 87–107, Nov. 2007.
- [27] P. Auer, N. Cesa-Bianchi, and P. Fischer, “Finite-time analysis of the multiarmed bandit problem,” *Machine Learning*, vol. 47, no. 2–3, pp. 235–256, May 2002.
- [28] N. Galichet, M. Sebag, and O. Teytaud, “Exploration vs exploitation vs safety: Risk-aware multi-armed bandits,” in *Proc. 5th Asian Conference on Machine Learning*, vol. 29. Canberra: PMLR, 13–15 Nov 2013, pp. 245–260.
- [29] M. Polese, M. Giordani, A. Roy, S. Goyal, D. Castor, and M. Zorzi, “End-to-End Simulation of Integrated Access and Backhaul at mmWaves,” in *Proc. IEEE CAMAD*, Oct. 2018.
- [30] M. Mezzavilla, M. Zhang, M. Polese, R. Ford, S. Dutta, S. Rangan, and M. Zorzi, “End-to-End Simulation of 5G mmWave Networks,” *IEEE Commun. Surveys Tuts.*, vol. 20, no. 3, pp. 2237–2263, Apr. 2018.
- [31] 3GPP, “NR; Study on integrated access and backhaul,” Technical Specification (TS) 38.874, 2018, v.0.6.0.
- [32] M. Giordani, M. Polese, A. Roy, D. Castor, and M. Zorzi, “A tutorial on beam management for 3GPP NR at mmWave frequencies,” *IEEE Commun. Surveys Tuts.*, vol. 21, no. 1, pp. 173–196, Sept. 2019.
- [33] K. Hosoya, N. Prasad, K. Ramachandran, N. Orihashi, S. Kishimoto, S. Rangarajan, and K. Maruhashi, “Multiple sector ID capture (MIDC): A novel beamforming technique for 60-GHz band multi-Gbps WLAN/PAN systems,” *IEEE Trans. Antennas Propag.*, vol. 63, no. 1, pp. 81–96, Oct. 2014.
- [34] C. Madapatha, B. Makki, C. Fang, O. Teyeb, E. Dahlman, M.-S. Alouini, and T. Svensson, “On integrated access and backhaul networks: Current status and potentials,” *IEEE Open J. Commun. Soc.*, vol. 1, pp. 1374–1389, Sep. 2020.
- [35] 3GPP, “NR; Backhaul Adaptation Protocol (BAP) specification,” Technical Specification (TS) 38.340, May 2022, v.17.0.0.
- [36] —, “Study on channel model for frequencies from 0.5 to 100 GHz,” Technical Report (TR) 38.901, Jun. 2018, v.15.0.0.
- [37] D. Yuan, H.-Y. Lin, J. Widmer, and M. Hollick, “Optimal joint routing and scheduling in millimeter-wave cellular networks,” in *Proc. IEEE INFOCOM*, Oct. 2018, pp. 1205–1213.
- [38] S. Hur, T. Kim, D. J. Love, J. V. Krogmeier, T. A. Thomas, and A. Ghosh, “Millimeter wave beamforming for wireless backhaul and access in small cell networks,” *IEEE Trans. Commun.*, vol. 61, no. 10, pp. 4391–4403, Sep. 2013.
- [39] Y. Zhu, Y. Niu, J. Li, D. O. Wu, Y. Li, and D. Jin, “QoS-aware scheduling for small cell millimeter wave mesh backhaul,” in *Proc. IEEE ICC*, Jul. 2016.