

Proactive Resilience in 1-2-1 Networks

Mine Gokce Dogan^{*}, Martina Cardone[†], Christina Fragouli^{*}

^{*} University of California, Los Angeles, CA 90095, USA, Email: {minedogan96, christina.fragouli}@ucla.edu

[†] University of Minnesota, Minneapolis, MN 55455, USA, Email: mcardone@umn.edu

Abstract—Millimeter Wave (mmWave) (and beyond) is expected to play an increasingly important role in our wireless infrastructure by expanding the available spectrum and enabling multi-gigabit services. Despite the promising aspects of mmWave communication, mmWave links are highly sensitive to blockage. In this paper, we develop *proactive* transmission mechanisms that suitably distribute the traffic across multiple paths in the mmWave network, with the two-fold objective of ensuring resilience against link blockages and achieve high end-to-end packet delivery rate. We present examples of resilience-capacity trade-off curves and show that there exist network topologies for which the worst-case and average approximate capacities are achieved by activating overlapping paths. We also show that this can provide additional benefits, such as decreasing the variance of the achieved rate.

I. INTRODUCTION

Millimeter Wave (mmWave) (and beyond) is an enabling technology that is playing an increasingly important role in our wireless infrastructure by expanding the available spectrum and enabling multi-gigabit services [1]–[6]. A number of use cases are currently built around multi-hop mmWave networks, such as the Facebook’s Terragraph network [7] that uses flexible mmWave backbones to connect clusters of base stations. Other examples of scenarios include private networks, such as in shopping centers, airports and enterprises; mmWave mesh networks that use mmWave links as backhaul in dense urban scenarios; military applications employing mobile hot spots; and mmWave based vehicle-to-everything (V2X) services, such as cooperative perception [3]–[5], [8]–[12].

Despite the promising aspects of mmWave communication, mmWave links are highly sensitive to blockage, channels may abruptly change and paths get disrupted [13]–[16]. Moreover, as in traditional wireless networks, mmWave network nodes are susceptible to component failures due to natural hazards and resource depletion. It becomes therefore of fundamental importance to deploy transmission mechanisms that are *resilient* against such disruptions.

In this paper, we aim to develop resilient transmission mechanisms to suitably distribute the traffic across multiple paths in the mmWave network by building on the so-called 1-2-1 network model that offers a simple yet informative model for mmWave networks [17]–[19]. In particular, [17] proved that the capacity of a 1-2-1 network can be approximated to

within a constant (i.e., which only depends on the number of network nodes) additive gap and its optimal beam schedule can be computed in polynomial-time, *without considering resilience to link failures*.

We leverage the results in [17] to develop *proactive* transmission mechanisms that build resilience against link failures in advance, with the only knowledge of the probability of link failures, while achieving a high end-to-end packet rate (fraction of packets delivered¹). The probability of link failures can be highly asymmetric in mmWave networks, yet known in advance through accurate models [13], which can be leveraged to achieve a superior performance². Our goal is to identify which paths to use and how to suitably schedule them, so that the packet rate is as large as possible in the presence of link blockages. A challenging aspect is beam scheduling in mmWave networks, where nodes communicate with each other by using beamforming and scheduling. Which nodes should communicate and for how long, is a non-trivial optimization problem. Our main contributions are summarized as follows.

- For a mmWave network with arbitrary topology, we characterize both the worst-case and the average approximate capacities through Linear Programs (LPs), and present examples of resilience-capacity trade-off curves. For the worst-case approximate capacity, our method is polynomial in the network size, but exponential in the maximum number of blocked links.
- We analyze the structure of the LPs and show that, out of an exponential number of paths (in the number of relay nodes N) that potentially connect a source to a destination, in 1-2-1 networks we only need to utilize at most $2N + 2$ paths to characterize the average approximate capacity.
- We show that, when the link capacities are all *equal*, the worst-case and the average approximate capacities can always be achieved by activating only edge-disjoint paths. However, when the link capacities are *unequal*, there exist network topologies for which the worst-case and average approximate capacities are achieved by activating overlapping paths. We also show that operating overlapping paths can provide additional benefits, such as decreasing the variance of the achieved rate.

¹We note that coupling our approach with erasure correcting codes allows to translate packet rates to information throughput.

²The link blockage probability depends on explicitly known factors such as physical distances, and indoor or outdoor propagation.

The work of M.G. Dogan and C. Fragouli was supported in part by the Army Research Laboratory under Co-Operative Agreement W911NF-17-2-0196 and by the U.S. National Science Foundation under Grant CNS-2146838. The work of M. Cardone was supported in part by the U.S. National Science Foundation under Grants CCF-2045237 and CNS-2146838.

Related Work. A multitude of works in the literature study multi-path diversity to achieve reliability in wireless networks [20]–[24]. However, these works focus on traditional wireless networks and they do not consider the scheduling constraints of mmWave communication. Some studies on mmWave communication have focused on profiling the distribution of the Signal-to-Interference-plus-Noise Ratio (SINR) in random environments both in cellular and ad-hoc network settings [25], [26]. However, these works consider communication over a single-hop either between ad-hoc nodes or in a cellular system between a base station and a user equipment, and they do not characterize the worst-case or the average approximate capacity as we do. There exist studies that aim to reduce link outages in mmWave networks by taking reactive approaches [27]–[33]. However, such a reactive mechanism adds the complexity of identification and adaptation, as well as feedback latency. In [34], to achieve resilience to link and node failures, the authors explored a state-of-the-art Soft Actor-Critic (SAC) deep reinforcement learning algorithm, which adapts the information flow through the mmWave network, without using knowledge of the link capacities or network topology. However, this proposed approach also reactively adapts to network disruptions. Several works proposed proactive approaches that constantly track users using side-channel information [35], [36] or external sensors [37]–[39]. These solutions have limited accuracy, and possibly require sensitive information, such as user location [38], [39]. They are different from our work, which leverages scheduling properties of mmWave links, as well as the blockage asymmetry, to design schemes that achieve the average and the worst-case approximate capacities and proactively offer high resiliency.

Paper Organization. Section II provides background on the 1-2-1 network model for mmWave networks. Section III presents our main results. Section IV concludes the paper.

II. SYSTEM MODEL AND BACKGROUND

We consider the Full-Duplex (FD) Gaussian 1-2-1 network model that was introduced in [17] to study the information-theoretic capacity of multi-hop mmWave networks. In an N -relay Gaussian FD 1-2-1 network model, N relays assist the communication between the source node (node 0) and the destination node (node $N + 1$). Each node in the network can simultaneously transmit and receive by using a single transmit beam and a single receive beam. At any particular instance, a node can transmit to at most one node and it can receive from at most one node. In order for two nodes to communicate, they need to steer their beams towards each other so as to activate a link (edge) that connects them. We next discuss some known capacity results for Gaussian FD 1-2-1 networks.

Capacity of Gaussian FD 1-2-1 networks. In [17], it was shown that the unicast capacity of an N -relay Gaussian FD 1-2-1 network can be approximated to within an additive gap that only depends on the number of nodes in the network. In particular, the following LP was proposed to compute




Fig. 1: MmWave network example with 4 nodes.

the unicast approximate capacity and its optimal schedule in polynomial-time,

$$\begin{aligned}
 \text{P1 : } \bar{C} &= \max_{x_p, p \in \mathcal{P}} \sum_{p \in \mathcal{P}} x_p C_p \\
 (\text{P1a}) \quad &x_p \geq 0, \quad \forall p \in \mathcal{P}, \\
 (\text{P1b}) \quad &\sum_{p \in \mathcal{P}_i} x_p f_{p, nx(i)}^p \leq 1, \quad \forall i \in [0 : N], \\
 (\text{P1c}) \quad &\sum_{p \in \mathcal{P}_i} x_p f_{p, pr(i)}^p \leq 1, \quad \forall i \in [1 : N + 1],
 \end{aligned} \tag{1}$$

where: (i) \bar{C} is the approximate capacity; (ii) \mathcal{P} is the collection of all paths connecting the source to the destination; (iii) C_p is the capacity of path p ; (iv) $\mathcal{P}_i \subseteq \mathcal{P}$ is the set of paths that pass through node i where $i \in \{0, 1, \dots, N + 1\} := [0 : N + 1]$; (v) $p.nx(i)$ (respectively, $p.pr(i)$) is the node that follows (respectively, precedes) node i in path p ; (vi) x_p is the fraction of time path p is used; and (vii) $f_{j,i}^p$ is the optimal activation time for the link of capacity $\ell_{j,i}$ when path p is operated, i.e., $f_{j,i}^p = C_p / \ell_{j,i}$. Here, $\ell_{j,i}$ denotes the capacity of the link going from node i to node j where $(i, j) \in [0 : N] \times [1 : N + 1]$.

Although the number of variables in LP P1 (particularly, the number of paths) can be exponential in the number of nodes, this LP can indeed be solved in polynomial-time through an equivalent LP as proved in [17]. We refer readers to [17] for a more detailed description.

Remark 1. In LP P1, the beam scheduling allows to share traffic across multiple paths, both over space and time without considering resilience to link and node failures. Our goal is to identify which paths to use and how to schedule them with the two-fold objective of ensuring resilience against link blockages and achieve a high end-to-end packet delivery rate.

To characterize the worst-case and average approximate capacities, we consider a permanent blockage (failure) model, where each link $\ell_{j,i}$ is blocked with probability $p_{\ell_{j,i}}$ (which can be learned in advance [13]). This model is different from the erasure channel model where a packet is blocked with probability $p_{\ell_{j,i}}$ at every channel use. As we show through the following example, the optimal solution of the erasure channel model does not necessarily give a feasible solution for the permanent blockage model.

Example 1. Consider the network with 4 nodes in Fig. 1, where the link capacities are $\ell_{3,1} = 6$, $\ell_{3,2} = 12$, $\ell_{4,3} = 6$ and the link blockage probabilities are zero except for $p_{\ell_{3,2}} = 1/2$. In Fig. 1, for the erasure channel model, we can simply replace the link capacities $\ell_{j,i}$ with the average link capacities $(1 - p_{\ell_{j,i}})\ell_{j,i}$ and solve LP P1 in (1) to find the average approximate capacity. The solution of LP P1 would use equal time sharing across the links $1 \rightarrow 3$ and $2 \rightarrow 3$ and the average

flow arriving at node 4 would be equal to 6. However, for the permanent blockage model, two scenarios can happen, namely: (1) the link with capacity $\ell_{3,2}$ is blocked and hence, with an equal time sharing schedule, the flow arriving at node 4 is equal to 3; and (2) none of the links is blocked for which case an equal time sharing schedule between the links $1 \rightarrow 3$ and $2 \rightarrow 3$ becomes infeasible since the link with capacity $\ell_{4,3}$ cannot support a flow equal to 9. \square

III. MAIN RESULTS

In this section, we aim to build scheduling mechanisms that are resilient against link blockages/failures in a mmWave network with arbitrary topology. In particular, we are interested in characterizing both the *worst-case* and the *average* approximate capacities.

In the *worst case*, we assume that $k_\ell \in [0 : |\mathcal{E}|]$ links fail, where \mathcal{E} is the set of network links. We highlight that the number k_ℓ and the set of the k_ℓ blocked links are not known, i.e., only the link blockage probabilities are known. Under such assumptions, we can find the *worst-case* approximate capacity by solving P1 in (1) with the objective function modified as follows (the constraints are the same as those in LP P1):

- *Worst-Case Approximate Capacity:*

$$\max_{x_p, p \in \mathcal{P}} \min_{a \in \mathcal{A}} \sum_{p \in \mathcal{P}^{(a)}} x_p C_p, \quad (2)$$

where: (i) \mathcal{A} is the set of all combinations of k_ℓ links from the $|\mathcal{E}|$ links; and (ii) $\mathcal{P}^{(a)}$ is the set of unblocked paths when the k_ℓ links $a \in \mathcal{A}$ are blocked.

In the *average case*, we remove the assumption of k_ℓ -size failure patterns, and find the average approximate capacity over *all* failure patterns by solving the LP P1 in (1), where the objective function is now modified as follows (the constraints are the same as those in LP P1):

- *Average Approximate Capacity:*

$$\max_{x_p, p \in \mathcal{P}} \sum_{p \in \mathcal{P}} x_p C_p \left(\prod_{(j,i) \in \mathcal{E}_p} (1 - p_{\ell_{j,i}}) \right), \quad (3)$$

where \mathcal{E}_p is the set of links in path $p \in \mathcal{P}$.

We highlight that solving P1 in (1) with one of the above two objective functions leads to a schedule that achieves either the worst-case or the average approximate capacity.

Remark 2. LP P1 in (1) with the objective function in (2) can be equivalently formulated with a number of variables polynomial in N by taking the same steps as in [17]. Thus, if k_ℓ does not depend on any network parameters (i.e., k_ℓ is a constant and hence $|\mathcal{A}|$ in (2) is also constant), the worst-case approximate capacity and an optimal schedule for it can be computed in polynomial-time in N . For the average approximate capacity, we have an exponential number of variables (the number of paths). Finding an alternative formulation with a polynomial number of variables in the network size is an interesting open research direction, which is currently under




Fig. 2: MmWave network example with $N = 5$ relay nodes.

investigation. As pointed out in Example 1, simply replacing each link capacity with the average link capacity and solving LP P1 would not find the average approximate capacity.

As we show through the following simple example, LP P1 in (1) without the modifications in (2) and (3) does not ensure resilience against link failures/blockages.

Example 2. Consider the network with $N = 5$ relay nodes in Fig. 2 when $k_\ell = 1$ link is blocked. There exist 5 paths connecting the source (node 0) to the destination (node 6), particularly $p_1 : 0 \rightarrow 1 \rightarrow 6$, $p_2 : 0 \rightarrow 2 \rightarrow 6$, $p_3 : 0 \rightarrow 3 \rightarrow 6$, $p_4 : 0 \rightarrow 4 \rightarrow 6$, and $p_5 : 0 \rightarrow 5 \rightarrow 6$ where the path capacities are $c^2, 4c^2, 9c^2, 16c^2$ and $25c^2$, respectively. Here, $c \geq 1$ is a constant and the capacities of the links which are on the same path are equal. The optimal solution of LP P1 activates only path p_5 to achieve the approximate capacity $25c^2$. Clearly, this solution does not ensure resilience as the worst-case rate for $k_\ell = 1$ is 0 (i.e., when any link on p_5 is blocked). However, resilience is ensured by the schedule that maximizes the objective function in (2), which activates the strongest two paths such that equal rate is sent through them. In particular, it activates p_4 and p_5 for $25/41$ and $16/41$ fractions of time (inversely proportional to the path capacities), respectively. The worst-case approximate capacity is equal to $9.76c^2$. As we increase the value of c , the gain we obtain by solving for the objective function in (2) increases. A similar result is obtained for the average approximate capacity. We assume that the path capacities are $c, 2c, 3c, 4c$ and $5c$ where $c \geq 1$, and the capacities of the links on the same path are equal. The link blockage probabilities are all equal to $1/10$ except for the links in p_5 for which the blockage probabilities are equal to $2/3$. The average rate achieved by LP P1 is $0.56c$ and the average approximate capacity in (3) is equal to $3.24c$. As we increase the value of c , the gain we obtain by solving for the objective function in (3) increases. \square

As Example 2 illustrates, over a given topology, the worst-case packet rate can significantly depend on the level of resilience (captured by the values of k_ℓ) that we want to offer. In Fig. 3, we present the resilience-capacity trade-off curve for the network in Fig. 2 for $c = 1$. The resilience-capacity trade-off for a mmWave network is the information theoretic Pareto-optimal vector of rates that can be achieved despite any $k_\ell \in [0 : |\mathcal{E}|]$ link failures. In Fig. 3, for each $k_\ell^* \in [0 : 5]$, we find the optimal schedule that achieves the worst-case approximate capacity in (2) when we assume that any k_ℓ^*




Fig. 3: Resilience-capacity trade-off curve.

network links fail (the worst-case approximate capacity for $k_l^* \geq 6$ link failures is equal to 0). Under this optimal schedule, we then find the achieved worst-case rate as the number of link failures increases from 0 to 5. For example, the optimal schedule for $k_l^* = 2$ activates the three strongest paths p_3, p_4 and p_5 ; hence, for $k_l^* = 2$, as k_l increases in Fig. 3, the achieved worst-case rate decreases and becomes 0 for $k_l \geq 3$. Fig. 3 shows that there is a trade-off between packet-rate and resilience, i.e., guaranteeing a certain amount of (worst case) resilience to say $k_l^* = 1$ link blockage, may come at the cost of a lower rate when no link is blocked (i.e., $k_l = 0$).

Remark 3. In Fig. 3, the optimal schedule for k_l^* link failures activates $k_l^* + 1$ paths. However, different strategies might be required for different values of k_l^* depending on the number of paths and path capacities. For example, if we add 5 more edge-disjoint paths in the network in Fig. 2 with capacities $c^2 i^2$ for $i \in [6 : 10]$, the optimal schedule for $k_l^* = 1$ link failure activates the strongest 3 paths, and the optimal schedule for $k_l^* = 2$ activates the strongest 5 paths. Moreover, more paths are activated if the path capacities are closer to each other: for equal path capacities, the optimal solution always (i.e., independently of k_l^*) activates all the paths. In particular, the optimal solution finds a schedule that allocates an equal rate through each path. This leads to a trade-off: if a higher number of paths are activated, every link failure in the worst case decreases the achieved rate less. However, activating a higher number of paths results in operating lower capacity paths for a longer time, and higher capacity paths for a shorter time. The optimal strategy is determined based on this trade-off.

We next leverage the following lemma, which shows that the average approximate capacity can always be achieved by using only a linear number (in N) of paths. In particular, the following lemma is a generalization of [17, Lemma 9].

Lemma 1. For any N -relay Gaussian FD 1-2-1 network, the average approximate capacity can always be achieved by activating at most $2N + 2$ paths in the network.

Proof. The LP with objective function in (3) and constraints in LP P1 in (1) is bounded and thus, there always exists an optimal vertex. In particular, each vertex of this LP satisfies at least $|\mathcal{P}|$ inequality constraints with equality among (P1a), (P1b)




Fig. 4: Example network for Lemma 1.

and (P1c). Since (P1b) and (P1c) combined represent $2N + 2$ constraints, we have at least $|\mathcal{P}| - 2N - 2$ constraints in (P1a) that are satisfied with equality. Thus, at least $|\mathcal{P}| - 2N - 2$ paths are not operated. This completes the proof of Lemma 1. \square

Remark 4. A similar result as in Lemma 1 does not hold for the worst-case approximate capacity. To see this, consider the case when $k_l = 2N + 2$, which might cause $2N + 2$ paths to be blocked, hence leading to a zero rate in the worst case.

Lemma 1 shows that at most $2N + 2$ paths suffice to achieve the average approximate capacity. As we show through the next example, a much smaller number of paths might indeed be sufficient to characterize the average approximate capacity. *Example 3.* Consider the network with $N = 12$ in Fig. 4. The link capacities are unitary and the link blockage probabilities are all equal to $1/5$. There exist 2705 paths connecting node 0 to node 13. The dots from node 0 to node 2 (and similarly, the dots from node 3 to node 13 in Fig. 4 represent 50 links connecting them. From Lemma 1 we know that operating $2N + 2 = 26$ paths suffices to achieve the average approximate capacity. However, there is no need to activate all of these paths to achieve the average approximate capacity of $64/125$. Instead, it suffices to activate only 1 of these paths, for instance $0 \rightarrow 2 \rightarrow 3 \rightarrow 13$ (highlighted with dashed lines in Fig. 4). \square

A question that naturally arises in designing a schedule that is resilient to link failures, and achieves the worst-case or average approximate capacity is the following: What are the best paths to use? Or in other words, are there any intrinsic properties of the paths that should be leveraged? The next theorem provides an answer to these questions.

Theorem 1. For an N -relay Gaussian FD 1-2-1 network with arbitrary topology, we have the following properties:

- (P1) When the link capacities are all equal, the worst-case and the average approximate capacities can always be achieved by activating only edge-disjoint paths.
- (P2) When the link capacities are unequal, there exist network topologies for which the worst-case and the average approximate capacities are achieved by activating overlapping paths.

Proof. The proof of (P1) will be provided in the journal version of this paper. We here focus on proving (P2). Towards this end, we consider the network in Fig. 5. There exist 5 paths connecting the source (node 0) to the destination (node 13), particularly: $p_1 : 0 \rightarrow 1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 5 \rightarrow 13$, $p_2 : 0 \rightarrow 1 \rightarrow 2 \rightarrow 3 \rightarrow 7 \rightarrow 8 \rightarrow 13$, $p_3 : 0 \rightarrow 6 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 5 \rightarrow 13$, $p_4 : 0 \rightarrow 6 \rightarrow 2 \rightarrow 3 \rightarrow 7 \rightarrow 8 \rightarrow 13$, $p_5 : 0 \rightarrow$




Fig. 5: Example network with overlapping paths.

$9 \rightarrow 10 \rightarrow 11 \rightarrow 12 \rightarrow 13$. The link capacities are assumed to be $\ell_{1,0} = \ell_{2,1} = \ell_{4,3} = \ell_{5,4} = \ell_{13,5} = \ell_{8,7} = 1$, $\ell_{9,0} = \ell_{10,9} = \ell_{11,10} = \ell_{12,11} = \ell_{13,12} = 4$, $\ell_{6,0} = \ell_{2,6} = \ell_{3,2} = \ell_{7,3} = \ell_{13,8} = 2$.

We first consider the *worst-case* scenario for $k_\ell = 1$ (i.e., only one link is blocked). In the optimal solution of the LP with objective function in (2) and constraints in the LP P1 in (1), three paths p_3 , p_4 and p_5 are activated with activation times 0.2, 1 and 0.3, respectively. The achieved worst-case approximate capacity is 1.2. As seen in Fig. 5, p_3 and p_4 share the edges with capacities $\ell_{6,0}$, $\ell_{2,6}$ and $\ell_{3,2}$. We note that none of the feasible solutions that consist of only edge-disjoint paths reaches the same or a higher worst-case rate.

We now consider the *average* approximate capacity over the same network in Fig. 5 with the same link capacities except for the link capacities in p_5 that are now assumed to be unitary. The link blockage probabilities are assumed to be equal to $1/5$ except for the links in p_5 for which the blockage probabilities are assumed to be equal to $1/3$. In this case, we observe a similar situation as in the worst-case scenario. In particular, the optimal solution of the LP with objective function in (3) and constraints in the LP P1 activates three overlapping paths, namely p_2 , p_3 and p_4 with equal activation times of 0.5, and the average approximate capacity is 0.39. We note that none of the feasible solutions consisting of only edge-disjoint paths reaches the same or a higher average rate. This concludes the proof of Theorem 1. \square

Intuitively, one would expect that activating edge-disjoint paths would provide higher resilience against link blockages because a higher number of paths can get blocked in the worst case as the number of shared edges increases. However, Theorem 1 shows that activating overlapping paths may provide higher resilience against link blockages in the case of unequal link capacities. One possible explanation for this is as follows. When we have two overlapping paths with equal capacity and activate only one of them, this path can stay active for at most one unit of time due to the scheduling constraints. However, if the link capacities are unequal and both of these paths are activated, the total activation time of these paths can be made larger than 1, which can result in a higher rate both in the average case and the worst case.

We now conclude this section with a couple of remarks, which showcase interesting consequences of Theorem 1.

Remark 5. *The gain between the worst-case approximate capacity (obtained by using overlapping paths) and the*

worst-case rate obtained by only using edge-disjoint paths can be arbitrarily large. As an example, consider the network in Fig. 5 with $k_\ell = 1$ and increase the link capacities in p_4 (except for $\ell_{8,7}$) and p_5 . The worst-case approximate capacity (obtained by using overlapping paths) approaches 2. The worst-case rate obtained by only using edge-disjoint paths approaches 1, thus the gain approaches 1, which is the largest gain that we can have in this network. In a similar manner, the gain between the average approximate capacity (obtained by using overlapping paths) and the average rate obtained by only using edge-disjoint paths can be arbitrarily large. In Fig. 5, as we increase the link capacities in p_4 (except for $\ell_{8,7}$) as well as the blockage probabilities of the links in p_5 , and decrease the blockage probabilities of the remaining links, the average approximate capacity (obtained by using overlapping paths) approaches 2. The average rate obtained by only using edge-disjoint paths approaches 1, thus the gain approaches 1, which is the largest gain that we can have in this network.

Remark 6. *Theorem 1 shows that if the solution of the LP with objective function in (3) and constraints as in the LP P1 activates overlapping paths for a network with equal link capacities, then there exists another solution that operates only edge-disjoint paths. However, activating overlapping paths can still provide additional benefits, such as decreasing the variance of the achieved rate as illustrated in the next example.*

Example 4. We consider the network in Fig. 5 without p_5 . The link capacities are assumed to be equal to 2 and the blockage probabilities are all equal to $1/5$. We consider the average approximate capacity given by solving the LP with objective function in (3) and constraints as in LP P1. Optimal solutions that only operate edge-disjoint paths activate only one of the paths with an activation time of 1. The average approximate capacity is $\mathbb{E}[R] = 0.52$ and the variance $\mathbb{V}(R) = 0.77$ where R is the achieved rate, $\mathbb{E}[\cdot]$ is the expected value and $\mathbb{V}(\cdot)$ is the variance. Another optimal solution activates all the four paths p_1, p_2, p_3, p_4 all with activation time equal to 0.25. In this case, $\mathbb{V}(R) = 0.38$, which is much lower than 0.77. \square

IV. CONCLUSIONS

We developed proactive transmission mechanisms that build resilience against link failures in advance, without an a priori knowledge of the failures while achieving high end-to-end packet rate. We characterized both the worst-case and average approximate capacities, and presented resilience-capacity trade-off curves. In particular, we showed that operating overlapping paths may be necessary to achieve the worst-case and the average approximate capacities and it may provide additional benefits, such as decreasing the variance of the achieved rate. A part of our future work is to develop online algorithms that will adapt to the network state changes as an alternative to centralized algorithms. Other future research directions consist of considering traffic that is not only unicast (e.g., multicast), and of including interference between concurrent beams, which was analyzed by the authors in [40] but without considering resilience against link failures.

REFERENCES

- [1] G. W. Paper, "5G mobile communications systems for 2020 and beyond," Jul. 2016.
- [2] R. Heath, "Vehicle-to-x communication using millimeter waves," 2016.
- [3] "What role will millimeter waves play in 5g wireless systems?" <https://www.mwrf.com/systems/what-role-will-millimeter-waves-play-5g-wireless-systems>.
- [4] J. Choi, V. Va, N. Gonzalez-Prelcic, R. Daniels, C. R. Bhat, and R. W. Heath, "Millimeter-wave vehicular communication to support massive automotive sensing," *IEEE Communications Magazine*, vol. 54, no. 12, pp. 160–167, 2016.
- [5] M. Mueck, E. C. Strinati, I.-G. Kim, A. Clemente, J.-B. Dore, A. De Domenico, T. Kim, T. Choi, H. K. Chung, G. Destino, A. Parssinen, A. Pouttu, M. Latva-aho, N. Chubierre, M. Gineste, B. Vautherin, M. Monnerat, V. Frascolla, M. Fresia, W. Keusgen, T. Haustein, A. Korvala, M. Pettissalo, and O. Liinamaa, "5g champion - rolling out 5g in 2018," in *2016 IEEE Globecom Workshops (GC Wkshps)*, Dec. 2016, pp. 1–6.
- [6] S. Rangan, T. Rappaport, and E. Erkip, "Millimeter-wave cellular wireless networks: Potentials and challenges," in *Proceedings of IEEE*, vol. 102, Mar. 2014, pp. 366–385.
- [7] "Facebook connectivity," <https://terragraph.com/>.
- [8] K. Sakaguchi, T. Haustein, S. Barbarossa, E. C. Strinati, A. Clemente, G. Destino, A. Parssinen, I. Kim, H. Chung, J. Kim, W. Keusgen, R. J. Weiler, K. Takinami, E. Ceci, A. Sadri, L. Xian, A. Maltsev, G. K. Tran, H. Ogawa, K. M., and R. W. H. Jr., "Where, when, and how mmwave is used in 5G and beyond," *IEICE Transactions on Electronics*, vol. E100.C, no. 10, pp. 790–808, 2017.
- [9] "Qualcomm partners with Russian mobile industry for mmwave 5G network in Moscow," <https://www.fiercewireless.com/5g/>.
- [10] "Qualcomm introduces end-to-end over-the-air 5g mmwave test network in europe to drive 5g innovation," <https://www.qualcomm.com/news/releases/>.
- [11] 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 Transactions on Communications*, vol. 61, no. 10, pp. 4391–4403, Oct. 2013.
- [12] S. Choi, H. Chung, J. Kim, J. Ahn, and I. Kim, "Mobile hotspot network system for high-speed railway communications using millimeter waves," *ETRI Journal*, vol. 38, no. 6, pp. 1052–1063, Dec. 2016.
- [13] I. K. Jain, R. Kumar, and S. Panwar, "Driven by capacity or blockage? a millimeter wave blockage analysis," in *2018 30th International Teletraffic Congress (ITC 30)*, vol. 01, pp. 153–159.
- [14] T. Bai and R. W. Heath, "Coverage and rate analysis for millimeter-wave cellular networks," *IEEE Transactions on Wireless Communications*, vol. 14, no. 2, pp. 1100–1114, 2015.
- [15] G. R. MacCartney, T. S. Rappaport, and S. Rangan, "Rapid fading due to human blockage in pedestrian crowds at 5g millimeter-wave frequencies," in *IEEE Global Communications Conference*, 2017, pp. 1–7.
- [16] Y. Wu, J. Kokkonen, C. Han, and M. Juntti, "Interference and coverage analysis for terahertz networks with indoor blockage effects and line-of-sight access point association," *IEEE Transactions on Wireless Communications*, vol. 20, no. 3, pp. 1472–1486, 2021.
- [17] Y. H. Ezzeldin, M. Cardone, C. Fragouli, and G. Caire, "Gaussian 1-2-1 networks: Capacity results for mmwave communications," *IEEE Transactions on Information Theory*, vol. 67, no. 2, pp. 961–990, 2021.
- [18] Y. H. Ezzeldin, M. Cardone, C. Fragouli, and G. Caire, "Polynomial-time capacity calculation and scheduling for half-duplex 1-2-1 networks," in *IEEE Int. Symp. Inf. Theory (ISIT)*, 2019, pp. 460–464.
- [19] Y. H. Ezzeldin, M. Cardone, C. Fragouli, and G. Caire, "On the multicast capacity of full-duplex 1-2-1 networks," in *IEEE Int. Symp. Inf. Theory (ISIT)*, 2019, pp. 465–469.
- [20] J. P. Rohrer, A. Jabbar, and J. P. G. Sterbenz, "Path diversification: A multipath resilience mechanism," in *2009 7th International Workshop on Design of Reliable Communication Networks*, 2009, pp. 343–351.
- [21] M. Motiwala, M. Elmore, N. Feamster, and S. Vempala, "Path splicing," in *Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication*, ser. SIGCOMM '08. Association for Computing Machinery, 2008, pp. 27–38. [Online]. Available: <https://doi.org/10.1145/1402958.1402963>
- [22] X. Yang and D. Wetherall, "Source selectable path diversity via routing deflections," in *Proceedings of the 2006 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ser. SIGCOMM '06. Association for Computing Machinery, 2006, pp. 159–170. [Online]. Available: <https://doi.org/10.1145/1159913.1159933>
- [23] N. Shillingford, D. C. Salyers, C. Poellabauer, and A. Striegel, "Detour: Delay- and energy-aware multi-path routing in wireless ad hoc networks," *IEEE*, Feb. 2008.
- [24] R. Yeung, *Information Theory and Network Coding*, Jan. 2008.
- [25] T. Bai, A. Alkhateeb, and R. W. Heath, "Coverage and capacity of millimeter-wave cellular networks," *IEEE Communications Magazine*, vol. 52, no. 9, pp. 70–77, 2014.
- [26] A. Thornburg, T. Bai, and R. W. Heath, "Mmwave ad hoc network coverage and capacity," in *2015 IEEE International Conference on Communications (ICC)*, 2015, pp. 1310–1315.
- [27] C. Jeong, J. Park, and H. Yu, "Random access in millimeter-wave beamforming cellular networks: issues and approaches," *IEEE Communications Magazine*, vol. 53, no. 1, pp. 180–185, 2015.
- [28] C. N. Barati, S. A. Hosseini, M. Mezzavilla, T. Korakis, S. S. Panwar, S. Rangan, and M. Zorzi, "Initial access in millimeter wave cellular systems," *IEEE Transactions on Wireless Communications*, vol. 15, no. 12, pp. 7926–7940, 2016.
- [29] S. Sur, V. Venkateswaran, X. Zhang, and P. Ramanathan, "60 ghz indoor networking through flexible beams: A link-level profiling," in *Proceedings of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, ser. SIGMETRICS '15. Association for Computing Machinery, 2015, p. 71–84. [Online]. Available: <https://doi.org/10.1145/2745844.2745858>
- [30] O. Abari, H. Hassanieh, M. Rodriguez, and D. Katabi, "Millimeter wave communications: From point-to-point links to agile network connections," in *Proceedings of the 15th ACM Workshop on Hot Topics in Networks*, 2016, p. 169–175.
- [31] H. Hassanieh, O. Abari, M. Rodriguez, M. Abdelghany, D. Katabi, and P. Indyk, "Fast millimeter wave beam alignment," in *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, ser. SIGCOMM '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 432–445. [Online]. Available: <https://doi.org/10.1145/3230543.3230581>
- [32] Y. Zhu, Z. Zhang, Z. Marzi, C. Nelson, U. Madhoo, B. Y. Zhao, and H. Zheng, "Demystifying 60ghz outdoor picocells," ser. MobiCom '14. New York, NY, USA: Association for Computing Machinery, 2014, p. 5–16. [Online]. Available: <https://doi.org/10.1145/2639108.2639121>
- [33] C. N. Barati, S. A. Hosseini, M. Mezzavilla, P. Amiri-Eliasi, S. Rangan, T. Korakis, S. S. Panwar, and M. Zorzi, "Directional initial access for millimeter wave cellular systems," in *2015 49th Asilomar Conference on Signals, Systems and Computers*, 2015, pp. 307–311.
- [34] M. G. Dogan, Y. H. Ezzeldin, C. Fragouli, and A. W. Bohannon, "A reinforcement learning approach for scheduling in mmwave networks," in *MILCOM 2021 - 2021 IEEE Military Communications Conference (MILCOM)*, 2021, pp. 771–776.
- [35] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim, "Wifi-assisted 60 ghz wireless networks," in *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '17. Association for Computing Machinery, 2017, p. 28–41. [Online]. Available: <https://doi.org/10.1145/3117811.3117817>
- [36] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer, "Steering with eyes closed: Mm-wave beam steering without in-band measurement," in *2015 IEEE Conference on Computer Communications (INFOCOM)*, 2015, pp. 2416–2424.
- [37] M. K. Haider, Y. Ghasempour, D. Koutsonikolas, and E. W. Knightly, "Listeer: Mmwave beam acquisition and steering by tracking indicator leds on wireless aps," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, 2018, p. 273–288.
- [38] T. Wei and X. Zhang, "Pose information assisted 60 ghz networks: Towards seamless coverage and mobility support," in *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, 2017, p. 42–55.
- [39] V. Va, X. Zhang, and R. W. Heath, "Beam switching for millimeter wave communication to support high speed trains," in *2015 IEEE 82nd Vehicular Technology Conference (VTC2015-Fall)*, 2015, pp. 1–5.
- [40] Y. H. Ezzeldin, M. Cardone, C. Fragouli, and G. Caire, "Gaussian 1-2-1 networks with imperfect beamforming," in *2020 IEEE International Symposium on Information Theory (ISIT)*, 2020, pp. 1599–1604.