

Investigating Integrated Access and Backhaul on the Aether 5G Testbed

Jack Brassil
Dept. of Computer Science
Princeton University
Princeton, NJ, USA
jbrassil@princeton.edu

Abstract—The envisioned dense deployment of millimeter wave small cells threatens to require expensive and potentially disruptive installation of fiber backhaul from each cell location to a Mobile Network Operator’s (MNO) nearest point-of-presence. Integrated Access and Backhaul (IAB) uses part of the wireless spectrum in lieu of fiber for the backhaul connection, promising to reduce the cost of fiber deployments. But IAB architectures also introduce design challenges and open research questions at multiple protocol layers, ranging from wireless self-interference to multi-hop topology management. To explore these issues we construct a replicable research testbed to investigate systems level architectural and performance issues in multi-hop IAB settings. We show how the Open Networking Foundation’s (ONF) *Aether*TM platform can be used to emulate alternative wireless backhaul architectures across a range of topologies and radio technologies. Using commercially available Citizens Broadband Radio Service (CBRS) radios and small cells, we evaluate the performance of a UE communicating over a multi-hop topology.¹

Index Terms—experimental testbeds, hybrid cloud computing, 5G, wireless backhaul, edge computing, Aether, OpenRAN, CBRS, UPF, wireless backhaul

I. INTRODUCTION

Many proposed 5G applications will rely on the higher bandwidths only achievable through the use of higher frequency bands than used in previous cellular generations. But high-band (e.g., millimeter wave) frequency deployments are coverage-limited and will require dense placement of base stations. Depending on the setting, cell densification may incur expensive site acquisition costs and/or infrastructure attachment, rights-of-way access, or application and review fees payable by providers to local governments.

On outdoor properties such as a university campus, the infrastructure required to support a single cell can be formidable, including a physical structure (e.g., post), and underground cable conduits with capacity to support multiple fiber pairs, and electric connections. In *neutral host* deployments – where a structure is shared by cells from different MNOs – the required connectivity and associated costs increase further. The bulk of these small cell installation costs can usually be attributed to trenching and installation [1]. Estimating

installation costs in either greenfield or brownfield settings is difficult, as are cell siting decisions. Additional complicating factors on private properties including some campus settings involve infrastructure ownership and the sharing of installation costs between property owners and MNOs. Further, given the unpredictable evolution of technology, it is unclear how current physical infrastructure investments will be reusable in next generation cellular systems.

Given the potential for high small cell deployment costs it is unsurprising that wireless backhaul – where part of the wireless spectrum is allocated for the backhaul connection of base stations instead of fiber – is a potentially cost-effective deployment solution. Wireless backhaul can use separate spectrum from the access network (out-of-band), or shared spectrum (in-band). To address this need for current cellular systems, the 3rd Generation Partnership Project (3GPP) has advanced a standard multi-hop IAB network architecture in the Release 16 standard for 5G Phase 2 [2].

While our understanding of IAB has been aided by simulation, analysis, bench experiments, and limited field trials, experimental research on architectures continues to be hampered by a lack of components, technologies, open source software implementations, and experimental testbeds. In this paper we describe an extensible testbed we constructed to create repeatable 5G experiments. Section II reviews the design and operation of IAB. We describe the testbed architecture and implementation in Section III. Based on the *Aether* system, the collaborative university-industry testbed can be replicated by other researchers to explore 5G RAN, edge and core technologies. The next section presents preliminary results on the performance of an IAB approach on an emulated multihop wireless topology. We conclude by discussing some objectives




Fig. 1: An illustration of a simple multi-hop IAB topology.

¹This material is based upon work supported by the National Science Foundation under Grant No. CNS-1923692 and DARPA under Dispersed Computing HR0011-17-C-0047.

for future work, and arguing that the testbed will be a powerful tool for exploring 5G systems and services well beyond IAB.

II. IAB BACKGROUNDER

Wireless backhaul is a well-established technology; its application based on LTE relaying was introduced by 3GPP in LTE Rel. 10. However, relaying was little used; there were few dense LTE small cell deployments and a desire to preserve limited 4G spectrum by avoiding wireless backhaul. Yet various studies [3] demonstrated the potential benefits of IAB. 5G New Radio (NR) is expected to be a more suitable setting for IAB, where limited range mmWave small cells will drive cell densification, with correspondingly larger backhaul infrastructure requirements.

IAB enables the flexible design of dense cell deployments with only a sublinear increase in the fiber backhaul support. In addition to lowering costs, a diverse set of novel use cases and deployment scenarios can be envisioned, including rapidly constructed emergency responder networks for disaster support. In more conventional settings, network density and coverage can be added incrementally, with IAB nodes temporarily deployed for evaluation purposes; wired backhaul can be installed subsequently if needed.

A. Evolving 5G Access and Edge Systems

The Rel. 16 IAB architecture aligns with current architectural initiatives opening 5G edge and access systems to innovation. In the 5G NR system architecture a gNodeB is a base station providing User Plane (UP) and Control Plane (CP) protocol terminations towards the User Equipment (UE). Aether implements a decentralized, hybrid cloud-based 5G core network (5GC) by combining a Mobile Edge Computing (MEC) platform implementing disaggregated edge services, and a Google Cloud Platform (GCP) based connectivity manager implementing centralized core services.

The User Plane Function (UPF) performs data handling functions such as packet routing, QoS handling, and serving as an anchor point for mobility. A *split* gNodeB comprises a Central Unit (CU) serving one or more Distributed Units (DUs). The IAB architecture was designed to support the DU/CU split option(s).

The standardized IAB architecture crucially supports multi-hop wireless backhauling to increase *both* capacity and coverage, and backward operating compatibility of legacy UEs with new IAB nodes. 3GPP adopted a forwarding-based IAB architecture, where each downstream IAB *child* base station is assigned an IP address that is routable from a parent *IAB-donor* base station with fiber backhaul. Intermediate IAB nodes forward the packets based on these destination addresses and optional route identifiers. Fig. 2 depicts a conventional single wireless hop between a UE and small cell labeled IAB_1 . UP packets are tunneled from the gNodeB to the 5GC using the GPRS Tunneling Protocol (GTP) [4], decapsulated at the UPF, and forwarded to the Data Network (DN).

Fig. 3 depicts a two-hop IAB forwarding network. An end-to-end GTP tunnel from the UE-associated small cell to




Fig. 2: A single wireless hop small cell network architecture.




Fig. 3: A depiction of a 2-hop IAB architecture as defined by the 3GPP Rel. 16 standard.

the 5GC is forwarded through the IAB-Donor cell. This forwarding approach introduces lower overhead than alternative architectures by 1) avoiding potential tunneling-over-tunneling or hop-by-hop tunneling protocol overhead; 2) eliminating the need for UP/CP gateway or proxy functionality at each child IAB node; and 3) reducing the complexity of corresponding core functions supporting these IAB approaches.

Each IAB node implements DU functions, and the IAB-donor implements CU functions. A compliant IAB node's protocol stack contains two sides. The Mobile Termination (MT) part exposes a UE-like interface, and is used to communicate upstream with a parent cell. The DU part communicates with either the downstream MT child node or a UE.

Efficient dynamic multi-hop topology management is supported by the Backhaul Adaptation Protocol (BAP). Since wireless backhaul links are susceptible to outages (e.g., due to obstructions) and topology configuration changes, BAP provides network resilience through autonomous topology adaptation if an active backhaul path is impaired or disconnected. While outside the scope of this work, we anticipate BAP will play a significant role in our future work.

IAB nodes may use in-band or out-of-band communications for backhaul. Simplicity favors in-band operation, where the same carrier frequencies are used for access and backhaul. But in the case of in-band operation, the MT (or DU) part of an IAB node typically cannot receive when the DU (or MT) is transmitting. This IAB *half-duplex constraint* is a performance limitation not present in out-of-band (or wired) backhaul, and is being actively addressed by modern full-duplex radio implementations [5].

Two final observations regarding wireless backhaul are worth noting. First, most settings will involve fixed cell installations such that signal directionality can be engineered. Second, in some settings (e.g., academic campus) out-of-band spectrum capacity might be readily available for backhaul

applications, including unlicensed (e.g., 2.4 Ghz, 5 Ghz for Wi-Fi) and lightly licensed bands (e.g., 3.5 Ghz for CBRS).

III. TESTBED ARCHITECTURE

We next describe an experimental testbed allowing us to emulate many RAN and edge innovations, specifically including IAB system architectures and protocols. We seek to emulate a variety of wireless multi-hop network topologies and approaches, and not simply construct and report on the performance of a single fixed topology or standardized solution. We also seek to create a testbed that will enable investigation of wireless backhaul service over multiple, heterogeneous radio channels. This will allow us to explore both current and future radios, as well as both in-band and out-of-band backhaul approaches. Finally, we strive to construct testbed infrastructure requiring minimal hardware investment, and hence implement a framework that can be easily replicated by other researchers, permitting repeatable research.

We focus on the systems and architectural aspects of IAB rather than issues at the link level and below. We make this decision in part because there has been extensive research and bench testing of 5G wireless channels and components. Further, emerging testbeds (e.g., PAWR systems including COSMOS, POWDER, ARA and AERPAW [6]) are well suited to support repeatable field testing of advanced wireless technologies. While a 2 wireless hop IAB theoretical model has been analyzed [7] and analytical studies are increasing steadily [8], the majority of published studies of IAB to date have been simulation based.

Simulations can provide a high fidelity representation of a complete implementation of an IAB system and its protocols, and permit modeling of arbitrary multi-hop wireless network topologies [9]. But given the growing availability of both 5G and CBRS radios and private LTE and 5G cores, we argue that an emulation of IAB can provide increasingly valuable insights, grounding learning in application of commercial components. But the first challenge in emulating an IAB system is that no standard, open, commercial IAB small cell technology is currently easily and widely available to academic researchers. Second, most field trials with commercial IAB systems have reported on small scale topologies constructed on an opportunistic or ad hoc basis [10]. While these studies have been revealing about operational issues (e.g., outdoor coverage, weather affects), they are not readily repeatable by other research groups.

Achieving our goals forces us in the short term to make various testbed implementation decisions that depart from the Rel. 16 standard IAB implementation; we will return to our approach to overcoming these limitations in Sect. III-B. Our first challenge was not having certain key open source components, in particular a small cell with a standard-compliant stack (including components such as MT). We discuss a partial workaround in the next subsection. To simplify our initial testbed, we elected to focus first on data plane operation and performance. We defer immediate consideration of other CP

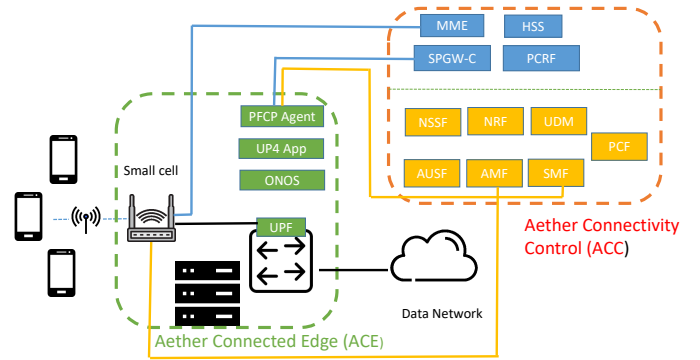


Fig. 4: The *Aether* software-defined, hybrid cloud Non-Standalone (NSA) core provides local breakout for data traffic with a line-rate, programmable switch-based UPF.

protocols, and investigating lower layer issues associated with the wireless channel.

For future study we leave multiple intriguing research topics, including half-duplex operations, topology management and associated control plane protocols (e.g., BAP), channel behavior and measurement, etc. Instead we strive to create an extendable, replicable design which can first help us identify lower bounds on the throughput of a standardized relaying architecture. As we will discuss, we anticipate that over time the initial differences between our emulated testbed implementation and standardized implementations will narrow as new commercial equipment becomes available and can be deployed in the testbed.

A. *Aether* System

To address our goal of constructing a replicable testbed we turned to the *Aether* system [11] to form the core of our wireless testbed. *Aether* is an open source, geographically decentralized 5G Connected Edge platform that provides mobile connectivity and edge cloud services. This unique system is a hybrid cloud instantiation of a mobile core network. *Aether* seeks to combine the complementary benefits of both a locally-controlled, private enterprise core as well as convenient, centralized cloud-based OAM services [12].

The on-premises *Aether* Connected Edge (ACE) allows an enterprise or campus full control of local data, and permits site-specific security policies through various mechanisms including local oversight of SIM management. There are currently 16 ACE locations in the United States, including both university and industry locations. The scalable edge architecture supports both enterprise-scale cellular communications and rapid deployment of containerized, low latency edge computing applications. ACE supports heterogeneous wireless connectivity over licensed and unlicensed spectrum. Mobile core components not requiring local placement or low latency access (e.g., billing functions) are implemented in the *Aether* Connectivity Control (ACC), a shared GCP backend managed by ONF.




Fig. 5: A depiction of the logical organization of the IAB testbed.

The number of university hosted ACE sites remains modest but is growing quickly, as it provides a platform for university-industry collaboration on 5G and beyond [13]. Our IAB testbed is implemented on the Princeton ACE site, which along with Stanford’s and Cornell’s sites are integral to the DARPA supported *PRONTO Project* [14]. PRONTO advances research on network verification, closed-loop control, and programmable networking, and leverages ONF’s community-building role and DevOps capabilities to provide a rapid path to production deployments of academic research.

Fig. 4 depicts the functional split between the ACE and ACC in Aether’s software-defined core. In the ACE an expandable set of 1U (edge) compute servers and a leaf-spine programmable switch fabric supports enterprise scale networking and low latency compute services. The ACC hosts functions that can be centralized across ACE sites, including billing, connection control, and charging. ACC implements NFV-based Open Mobile Evolved Core (OMEC) functions [15] with 5G NSA extensions (Fig. 4 ACC top half in blue), as well as *free5GC* core functions [16] (Fig. 4 ACC bottom half in yellow).

To implement IAB functionality we leverage 2 crucial and related ACE capabilities. The first capability is a local UPF, implemented as a disaggregated application collectively called *UP4* [17], [18]. The UPF is split between network controller applications and a P4 program executing on a programmable switch to provide a scalable, line-rate termination of GTP-U tunnels from a potentially huge number of served UEs. The second key ACE capability we leverage is *local breakout*, enabling UE-sourced packets destined to campus and public data networks to egress from the testbed at the edge, rather than travel to the ACC core.

B. Testbed Implementation

As noted above implementing an IAB service on the Aether platform required making numerous decisions affecting the fidelity of the emulation. Our preference was to create a standards-compliant relaying service. But no commercial MT/DU-partitioned small cell was available. Implementing forwarding was considered, but taking this path would have made the testbed less replicable to other researchers. Instead, we decided to implement multi-hop tunneling using a combination of linux *iptables* Network Address Translations (NAT) and hop-by-hop GTP-U tunnels.

Fig. 5 shows the logical operation of our testbed emulating a 2-hop network. A client compute node (left) with a dedicated radio emulates (logical) UE₁ and associates with a (logical) child IAB node labeled L-IAB₁. Rather than GTP-encapsulating packets and relaying them to an upstream IAB-donor, the GTP encapsulated packets are sent directly to ACE, and decapsulated by the UPF. Using iptables NAT rules at an intermediary edge server (shown in Fig. 6), the packets are looped back over a wired network to a second, logical UE₂, also with a dedicated radio. UE₂ also implements a NAT function, forwarding arriving packets via *its* dedicated radio channel to a second small cell, L-IAB₂, in the daisy chain, which similarly tunnels packets to the ACE UPF. From there decapsulated packets are forwarded to their destinations on the Data Network.

Hop-by-hop GTP tunnels incur more overhead than relaying and a single end-to-end tunnel shown in Fig. 3. We note that a similar proxying function to emulate forwarding was previously introduced in the LTE context in [3]. But the situation is somewhat different here for two reasons: 1) the proximity of the UPF avoids the latency associated with a round-trip to a centralized core; and 2) in steady state decapsulation is performed at line rate on a programmable switch, incurring negligible latency ($< 1 \mu s.$) compared to the radio hops (e.g., 10-20 ms. roundtrip times). Our testbed approach incurs the latency of a software-based encapsulation for each hop, as opposed to a single software-based encapsulation for an end-to-end GTP implementation. But it incurs negligible latency for the 2 hardware-based decapsulations – performed by the programmable switch at line rate – compared to 1 software-based decapsulation for the end-to-end tunnel implementation in a non-Aether based conventional private core implementation.

We also took steps to ensure the testbed implementation would scale to permit study of the behavior of many UEs operating in an IAB setting. Fig. 5 shows that we consolidated hardware by using a single compute node to emulate both UE₁ and UE₂, and a single small cell to emulate logical IAB₁ and IAB₂. Each logical UE had a separate radio; we limited the number of radios per physical CPU to 4. Here we chose the efficiency of a scalable testbed implementation over strictly independent channels, and found no ill effects in our early testing.

Without access to a 5G small cell with a standards-compliant IAB implementation, we initially validated our




Fig. 6: Details of physical testbed design.

testbed operation with a Sercomm SCE4255W LTE CBRS Cat-A Indoor Enterprise Small Cell. We experimented with 2 different x86 CPUs serving as UEs. The first was a moderately powerful desktop with a 3.00 GHz Intel Core i5-9500 CPU, and the second was a Zotac Zbox small form factor mini-PC appliance with 1.60 GHz Intel Celeron CPU N3160. We will report primarily on the latter, which is likely more representative of handheld UE processing capability. For x86 attached radios we employed Sercomm’s USB 2.0 CBRS wireless dongle (“Wingle”) [19]. In addition, we occasionally performed tests on the widely available iPhone SE 2020 smartphone.

As mentioned above, we initially considered implementing IAB forwarding on a 5G-capable small cell. But we anticipate that standards-compliant commercial devices will be available soon, and envision straightforward integration with the Aether platform. Use of available commercial devices would also support testbed replication. In the near future we expect to introduce Sub-6 GHz and 28 GHz band indoor small cells [20].

IV. PRELIMINARY RESULTS

Fig. 6 shows a detailed physical description of our testbed implementation of a simple 2-hop wireless linear topology. The order of links traversed by a packet destined to the DN is indicated. For detailed inspection of how iptables rules for port forwarding and NAT functionality implement packet recirculation, our project github site [21] provides our source code. Observe in the figure that by extending the number of wireless interfaces (and associated routing rules), the testbed can emulate various tree topologies with a variable number of hops. Use of simple mechanisms such as ‘downing’ links, and adjusting link bandwidths and latencies with standard Linux traffic control further extends the wireless network topologies and behaviors that can be emulated. If needed, the number of physical client nodes and small cells can also be increased to scale the overall emulated network size.

We employed standard network performance measurement tools including *iperf2*, *ndt7*, *speedtest-cli*, etc, to measure performance, while monitoring CPU utilization and interface traffic with *top*, *iftop*, etc. For certain HTTP transfers we elected to use a *Squid* web proxy on the server node to mitigate jitter introduced by remote data network accesses. Precise throughput and latency measurements are difficult to obtain and have been deemphasized here, as our testbed shares connectivity to public internet sites with other campus traffic.

We present two simple experiments to demonstrate the testbed’s versatility. In the first example we consider integrated access and backhaul with CBRS. The maximum downlink/uplink throughput for a single CBRS link is approximately 105/10 Mbs, as measured on an iPhone SE near the small cell. A base case (ethernet) is presented showing the maximum throughput capability of our limited capability UE (Zbox). Fig. 7 shows a few representative HTTP throughput measurements to the nearest *Measurement Lab* server for a 1, 2 and 3 hop CBRS network. The latency of each hop is significant, and it is unsurprising to see throughput decrease by nearly half

Experiment 1: Integrated Access & Backhaul

| <i>0 wireless hops: ethernet, no proxy</i> | | | |
|---|----------------|--------------------------------|------------|
| Run No. | Downlink (Mbs) | Uplink (Mbs) | |
| 1 | 426.59 | 466.40 | |
| 2 | 426.59 | 351.98 | |
| 3 | 426.59 | 498.28 | |
| <i>1 wireless hop: proxy; local iperf TCP uplink comparison</i> | | | |
| Run No. | Downlink | Uplink [iperf: client, server] | |
| 1 | 81.86 | 7.42 | 6.97, 6.66 |
| 2 | 84.19 | 5.81 | 6.05, 5.73 |
| 3 | 82.39 | 7.67 | 6.58, 6.34 |
| <i>2 wireless hops: proxy</i> | | | |
| Run No. | Downlink | Uplink | |
| 1 | 41.57 | 4.15 | |
| 2 | 43.21 | 4.04 | |
| 3 | 43.95 | 3.95 | |
| <i>3 wireless hops: proxy</i> | | | |
| Run No. | Downlink | Uplink | |
| 1 | 27.24 | 2.38 | |
| 2 | 29.17 | 2.23 | |
| 3 | 27.74 | 2.43 | |

Fig. 7: Downlink and uplink *HTTP* end-to-end throughput (Mbs) for a UE (Zbox) base case (ethernet) and 1, 2 and 3 wireless (CBRS) hops.

with each additional wireless backhaul hop. Such a throughput reduction might place a practical upper limit on the acceptable number of wireless backhaul hops.

In the second experiment we consider *heterogeneous* access & backhaul on a simple 2-hop network. Fig. 8 shows how HTTP performance changes when the CBRS backhaul link is replaced by a Wi-Fi (.11ac) channel with 46/25 Mbs downlink/uplink capacity. Comparing with the 2 CBRS wireless hop case of Fig. 7, we note that the downlink throughput is comparable. That is, though the Wi-Fi backhaul downlink has only approximately half the capacity as the CBRS downlink, that link still has capacity to maintain similar performance at the UE. Note too that the uplink performance has increased by approximately 50% due to the higher capacity Wi-Fi uplink (25 vs. 10 Mbs).

In our future experimentation we look forward to testing many additional topologies and radios. We also envision narrowing the current differences between our testbed capabilities and standard IAB implementation. We also anticipate additional investigation of heterogeneous wireless backhaul, which may be practical to deploy at scale in our campus setting. In other future work we anticipate exploring potentially significant coverage and performance gains when testing sub-6 GHz and mmWave small cells.

Experiment 2: Heterogeneous Access & Backhaul

2 wireless hops – CBRS with WiFi backhaul: proxy

| Run No. | Downlink | Uplink |
|---------|----------|--------|
| 1 | 36.78 | 6.46 |
| 2 | 43.35 | 6.73 |
| 3 | 44.90 | 6.28 |
| 4 | 38.51 | 6.05 |
| 5 | 44.40 | 6.67 |

Fig. 8: Measured *HTTP* end-to-end throughput (Mbs) for a 2 hop network with CBRS access and Wi-Fi backhaul. The maximum throughput of the Wi-Fi backhaul downlink/uplink channels was 46/25 Mbs.

V. CONCLUSION

We have shown that the *Aether* testbed is an available, open architecture environment for university researchers to investigate future 5G RAN and core services at scale. While we have used IAB as a demonstrator function, the testbed eases creating many other compelling applications and services. IAB offers many technical challenges that remain to be studied that will affect its deployability, and we look forward to expanding our investigation on the testbed. But many non-technical issues also remain as potential barriers to deployment. Spectrum owners consider bandwidth an exceedingly precious resource, and are naturally reluctant to use it when alternate backhaul technologies (e.g., fiber) exist at feasible cost. Infrastructure providers may consider out-of-band wireless backhaul solutions they don't own to be risky and of uncertain quality.

It remains possible that wireless backhaul will be used only opportunistically for special purpose coverage situations. In public settings MNOs initially appear likely to deploy small cells first along roadside rights-of-way, providing initial coverage in a street grid along only the most dense traffic corridors. One opportunity for wireless backhaul is to reach deeper into properties adjacent to these corridors. In some cases these properties may prioritize coverage over performance, making the reduced bandwidth achievable over multiple hops less concerning.

Many wireless challenges also remain to be studied and field tested. For example, both cross-interference between cells and self-interference between access and backhaul can potentially limit use. But new antennas with MIMO and beam steering capabilities potentially can be deployed to reduce interference and permit efficient wireless backhaul. There are also exciting opportunities and novel use cases to explore, including UAV-based IAB implementations [23] and their applications to rapidly deployed infrastructure.

Acknowledgement: Thanks to ONF and the Princeton PRONTO project team for the construction and ongoing

support of the Aether ACE site.

REFERENCES

- [1] CCG Consulting, "The Cost of Siting Small Cells", *POTs and PANs: Broadband for All*, March 1, 2019, <https://potsandpansbyccg.com/2019/03/01/the-cost-of-siting-small-cells/>
- [2] 3rd Generation Partnership Project, Technical Specification Group Services and Systems Aspects (Release 16), *3GPP TR 21.916*, 2020-7. <https://www.3gpp.org/release-16>
- [3] Gan Jiansong, Guo Zhiheng, Fan Rui, Liu Weihong, Wang Hai, Kristofer Sandlund, Liu Jianjun, Shen Xiaodong, Liu Guangyi, "LTE In-Band Relay Prototype and Field Measurement," *2012 IEEE 75th Vehicular Technology Conference*, 2012.
- [4] ETSI, "General Packet Radio Service (GPRS) Tunnelling Protocol (GTP) Technical Specification," *3GPP TS 29.060 version 13.7.0 Release 13*, https://www.etsi.org/deliver/etsi_ts/129000_129099/129060/13.07.00_60ts_129060v130700p.pdf, 2017-3.
- [5] J. Zhang, N. Garg, M. Holm, T. Ratnarajah, "Design of Full Duplex Millimeter-Wave Integrated Access and Backhaul Networks," *IEEE Wireless Communications*, vol. 28, no. 1, pp. 60-67, February 2021.
- [6] *Platforms for Advanced Wireless Research*, <https://advancedwireless.org/platforms/>.
- [7] C. Saha, H. S. Dhillon, "Millimeter Wave Integrated Access and Backhaul in 5G: Performance Analysis and Design Insights," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 12, pp. 2669-2684, Dec. 2019.
- [8] M. N. Islam, S. Subramanian, A. Sampath, "Integrated Access Backhaul in Millimeter Wave Networks," *2017 IEEE Wireless Communications and Networking Conference (WCNC)*, 2017, pp. 1-6.
- [9] Michele Polese, Marco Giordani, Arnab Roy, Sanjay Goyal, Douglas Castor, Michele Zorzi, "End-to-End Simulation of Integrated Access and Backhaul at mmWaves", *IEEE CAMAD*, Barcelona, Spain, 2018.
- [10] T. Tian et al., "Field Trial on Millimeter Wave Integrated Access and Backhaul," *IEEE 89th Vehicular Technology Conference*, 2019, pp. 1-5
- [11] Open Networking Foundation, *Aether Project*, 2021. <https://opennetworking.org/aether/>
- [12] Oguz Sunay, Open Networking Foundation, "Aether: An Architectural & Operational Overview (Keynote presentation)", *ONF 5G Connected Edge Cloud for Industry 4.0 Transformation*, <https://www.youtube.com/watch?v=ZzwFQCSmOJg>
- [13] Denise Barton, "Aether – Rapidly Deploying New Edge Sites Around the World," *Open Networking Foundation Blog*, March 11, 2021. <https://opennetworking.org/about-onf/onf-overview/aether-rapidly-deploying-new-edge-sites-around-the-world/>
- [14] *PRONTO Project*, 2021. <https://prontoproject.org/>
- [15] *OMEC Project*, 2021. <https://opennetworking.org/omec/>
- [16] *free5GC Project*, 2021. <https://www.free5gc.org/>
- [17] *UP4 Project*, ONF internal project presentation, 2021.
- [18] Robert MacDavid, Carmelo Cascone, Pingping Lin, Badhrinath Padmanabhan, Ajay Thakur, Larry Peterson, Jennifer Rexford, Oguz Sunay. "A P4-based 5G User Plane Function," *ACM SIGCOMM Symposium on SDN Research (SOSR '21)*, September 20–21, 2021.
- [19] Ben Lin, "Sercomm CBRS Small Cell in ONF," (presentation), 2020. <https://opennetworking.org/wp-content/uploads/2020/09/Ben-Lin-Solo-Final-Slides.pdf>
- [20] "Sercomm's 5G Enterprise mmWave Small Cell Achieves FCC Part 30 Approval," *CISION PR Newswire*, Sept. 23, 2020. <https://www.prnewswire.com/news-releases/sercomms-5g-enterprise-mmwave-small-cell-achieves-fcc-part-30-approval-301136319.html>
- [21] J. Brassil, "Investigating IAB on the Aether 5G Testbed," (PRONTO testbed backhaul implementation), 2021. <https://github.com/jackbrassil/pronto-backhaul/tree/master>
- [22] Paulo Marques et al, Federated Union of Telecommunications Research Facilities for an EU-Brazil Open Laboratory, "Experiment 1 – LSA/SAS for extended LTE capacity with E2E QOE:" D5.1: Initial experiment description and results," *FUTEBOL – H2020 688941* 2018.
- [23] "UAV-Based in-band Integrated Access and Backhaul for 5G Communications, Abdurrahman Fouda, Ahmed S. Ibrahim, Ismail Guvenc, Monisha Ghosh, 2018.