FSONet: A Wireless Backhaul for Multi-Gigabit Picocells Using Steerable Free Space Optics†

Md. Shaifur Rahman, Max Curran, Himanshu Gupta, Kai Zheng, Jon Longtin, Samir R. Das, Thanvir Mohamed

***Abstract*—Expected increase in cellular demand has pushed recent interest in picocell networks which have reduced cell sizes (100-200m or less). For ease of deployment of such networks, a wireless backhaul network is highly desired. Since RF-based technologies are unlikely to provide the desired multi-gigabit data rates, we motivate and explore use of free space optics (FSO) for picocell backhaul. In particular, we present a novel network architecture based on steerable links and sufficiently many robust short-range links, to help circumvent the key challenge of outdoor effects in reliable operation of outdoor FSO links. Our architecture is motivated by the fact that, due to the high density of picocells, many short-range links will occur naturally in a picocell backhaul. Moreover, use of steerable FSO links facilitates networks with sufficient redundancy while using only a small number of interfaces per node. We address the key problems that arise in the context of such a backhaul architecture, viz., an FSO link design with desired characteristics, and related network design and management problems. We develop and evaluate a robust 100m FSO link prototype, and simulate the proposed architecture in many metro US cities while show its viability via evaluation of key performance metrics.**

I. Introduction

With the advent of 5G cellular network, it is expected that there will a drastic increase in mobile broadband data consumption. Some projections estimate a 1000-fold increase in capacity demand within a few years [2]. Researchers have been addressing this impending 'capacity crunch' using various mechanisms i.e, improving spectral efficiency, developing shared spectrum regimes (e.g., white spaces), and utilizing very high frequency spectrum bands (e.g., 30-300 GHz). Despite all these efforts, *spatial reuse* of wireless spectrum is the key to address the explosive increase in wireless broadband demands. Due to spatial reuse, cell radius have fallen steadily from tens of kms in the 1980s to just a few hundred meters currently [2].

Picocells' Backhaul Network. The above trend has led to significant interest in *picocells*. These are the cells that have a coverage radius of 100–200 m or even less. Recent research anticipates multi-Gbps capacity of such picocells [3]. One significant concern in deployment of such networks is backhauling of the base stations to a close-by "gateway" for further connectivity. Given the high capacity requirement of the backhaul network, optical fiber communication is often preferred as it can deliver very high data rates. However, deploying optical fibers for such dense outdoor picocells can be significantly expensive; thus, a wireless solution is desired. Unfortunately, the existing or upcoming RF based solutions are unlikely to be able to provide the required data rates of 10Gbps or more (see §II) at desired ranges (100–200 m and more).

† A preliminary version of this work appeared in ACM MobiCom 2017 [1]. This work was supported by NSF award #1514017.

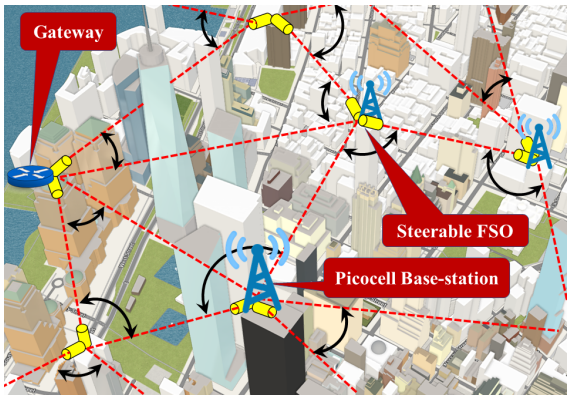

Fig. 1: FSO-based backhaul network for picocells. Each base station has a network switch connecting steerable FSO devices.

FSO-based Backhaul Solution. Free Space Optics (FSO) communication is similar in principle to (wired) optical fiber communication, except that, the laser beam travels in free space instead in an optical fiber. In general, FSO links can provide data rates—even up to ≈100 Gbps—over long distances (≈kms) [4], due to high frequencies of lightwaves and absence of regulatory restrictions. Also, the spectrum is free of any licensing fee while most RF spectrum is auctioned to the tune of billions of dollars. However, reliability of FSO outdoor links is challenging due to outdoor effects, especially at long ranges. For short ranges (≈100 m or below) some of the challenges can be overcome easily. Thus, picocells offer a unique setting for FSO-based network as short links arise naturally in densely deployed picocells.

In this work, we propose the use of *steerable* FSO devices which create a "dynamic" network with sufficient network redundancy yet minimal number of FSO interfaces per base-station. See Fig. 1. Our overall approach (called FSONet) [1] is based on the following multi-pronged strategy: (i) use of steerable FSO links to create a "dynamic" network with sufficient network redundancy, (ii) having a sufficient number of robust short links that can handle most weather effects, and additional longer links for improved hop-count and network capacity in favorable conditions.

Our Contributions. In context of the above proposed vision, our work makes the following contributions:[1]

1) We propose a novel FSO network architecture, based on robust short links and steerable FSO links, that can offer sufficient network availability in face of outdoor effects (§II).

[1]The new contributions with respect to our preliminary work [1] include analysis of bi-directional links (§III), tailored FSO networks (§IV), and more robust and realistic simulation results in §VI (over 3D buildings, comparison with hybrid FSO/RF architectures, and weather simulation).

2) We design and evaluate a 100m link prototype with active tracking and pointing (§III) with desired link margin and tolerance requirements, using commodity hardware.
3) We address the key problem of base station placement and dynamic network design that arises in our context. Since the problem is intractable, we design an efficient multi-step heuristic that delivers good quality network solutions in real settings (§IV).
4) For the runtime operation of FSONet, we address the network reconfiguration problem which selects a network topology and flow-routes at runtime, in response to changing traffic and/or link failures (§V).

Finally, we conduct extensive simulations to demonstrate the effectiveness of our techniques (§VI).

II. Motivation and Proposed Architecture

We start with making a case of an FSO-based network for picocell backhaul.

Why not an RF-based backhaul? In a wireless mesh network, since the individual picocell base stations are expected to provide multi-gigabit capacity through mmWave bands [3], we need backhaul links capable of supporting up to 10 Gbps, perhaps even 40 Gbps, depending on the network topology. However, no known RF technology can provide such data rates at ranges of ≈100–200 m.[2]

Why FSO? FSO links use visible or infrared (IR) laser beams, in the free space. Laser beams are very narrow, thus eliminating crosstalk interference. As optical spectrum is also unregulated, FSO links can easily offer Gbps-Tbps of bitrates at long distances. However, at long ranges, FSO links can suffer transient failure due to various outdoor effects (see Table I). Fortunately, short-range links with sufficient link "margin" can withstand most outdoor effects. Due to the high density of picocell networks, short links will arise naturally in a backhaul network connecting these base stations. In addition, to create a mesh network with sufficient redundancy with minimal number of interfaces per node, we propose to use steerable FSO devices. A steerable FSO device can essentially steer the FSO beam to any one of many intended receivers—thus, creating a "dynamic" network, where the network topology can be controlled at runtime. Such flexibility also helps in efficient runtime handling of moving hotspots [1].

A. FSONet Backhaul Architecture

We envision a dense picocell network, with each picocell having a range of about 100–200 m and capable of providing multi-gigabit (2-5 Gbps or more) peak capacity [3]. Picocell base stations themselves may be co-located indoors or outdoors for appropriate coverage requirements. We assume that indoor picocells have a fiber extended to an outdoor location to facilitate a line-of-sight backhaul connection to other picocells. Our goal is to create an FSO mesh network that connects the picocells base stations to a small number of gateways. *Note that the availability requirement of picocells is less stringent than "5 nines" and is generally 99 to 99.9% due to the default availability of macrocells in urban areas [6].* To that end, we design FSONet, a picocell backhaul with the following features:

[2]THz technology is relatively immature and has limited outdoor range due to high atmospheric attenuation [5]. Massive MIMO requires very challenging signal processing at high frequencies.

Cause	Effect	Mitigation
Weather: fog, snow, rain	Attenuation up to 400 dB/km	short range and/or link margin
Turbulence: scintillation, beam-wander, beam-spreading.	minimal impact at ≤ 500m	short range and/or link margin
Building motions: movement of deploying platforms	misalignment	Tracking & Pointing (TP)
Blockage: by objects, e.g., birds	transient link failure	packet retransmission or rerouting

TABLE I: Challenges for Outdoor FSO Links.

1) Each base station has a network switch (see Fig. 1) connected to a small number of FSO devices. The FSO devices enable steerable bi-directional (either via two wavelengths on the same optical path [7] or via two separate optical paths) links. This creates a "dynamic" network.
2) We restrict the maximum range of links to about 500m, and classify them into two categories: short (≤ 100m) and longer (100-500m) links.
   a) All links have an optical power *budget* of about 30dB (to allow for SFP-based prototypes III).
   b) Each short (≤ 100m) link has about 15dB link margin allocated specifically to withstand weather effects; we refer to such links as *robust short links*. Such links should have an availability of more than 99.5% in most US cities (see below).
   c) All links handle building movements via the link's movement tolerance and an active tracking and pointing (TP) mechanism (§**III**).
3) The network contains a subnetwork called a *backbone* that (i) comprises solely of robust short links, and (ii) provides coverage, and connectivity from backbone nodes to gateways (§**IV**). Thus, in unfavorable weather conditions, the network availability is maintained (albeit, at a lower capacity than the full network) via a backbone whose robust short links can handle weather effects. In favorable weather conditions, the full network (including longer links) is operational. From (2b) above, this guarantees at least 99.5% network availability.
4) We can handle transient link outages due to blockages etc. by retransmission or rerouting flows by exploiting network redundancy (§**V**).

Link Budget Analysis. In short links, of the total 30dB budget, we allocate about 15dB margin for weather effects, and 10–15dB for *geometric* losses (due to beam divergence and limited receiver aperture) and link's movement tolerance. For longer links, we allocate about 25dB margin for geometric losses and movement tolerance (needed for TP mechanism) and 2-5dB for any scintillation effects.

Tolerance to Weather Effects. Based on the fog attenuation model in [8], a 100m link with XdB link margin can withstand

fog conditions that have a visibility of greater than $(1.3/X)$ km; thus, a 100m link with 15dB (13dB) can withstand fog with > 85m (100m) visibility. The visibility data of 20 US cities [9] implies that most cities have a 99.5% chance of having more than 250m visibility. As non-fog weather factors have much less attenuation, a 100m link with 13-15dB should should have an availability of much more than 99.5% in most US cities. Note that, as mentioned before, the availability requirement of picocells is generally 99 to 99.9% due to a macrocells [6].

Commercial FSO Link Solutions. Commercially available FSO links are fixed (i.e., do not steer), are bulky (up to 2 cubic feet [10], [11]) and expensive ($6K-25K for a device [12]); this is because their objective is to provide highly robust singular links for longer ranges (e.g., multiple kms). On the other hand, we envision lightweight and less expensive solutions by targeting smaller ranges (at most 500m) and embedding redundancy in the overall network rather than individual links.

III. FSO Link Design in FSONet

In this section, we describe our design of a lightweight tracking and pointing (TP) mechanism, and a 100m FSO link prototype (equipped with TP) with desired characteristics for FSONet and evaluate its performance. As in prior works [13], [1], we use Galvo Mirrors [14] as the beam steering mechanism.

Active TP based on GMs and Photodiodes. Typical TP mechanisms [15], [16] include a fast steering mirror or gimbal controlled by digital servos [17], [18], [19] for pointing, along with some tracking detectors to track the target or beam. Common tracking detectors include positioning sensing diodes [18] (e.g., CCDs [17], photodiode arrays [20]) etc.[3] we leveraged the steering mechanism GMs also as the pointing mechanism. The pointing accuracy of GMs (10μrad) is much less than 1/8 of the beam divergence (about 280μrad) of our collimator [22], and thus should guarantee a less than 0.5dB tracking error [17]. For beam tracking, we use photodiodes, due to relatively short range (< 500m) of our links which entails minimal intensity fluctuations in outdoors [23], [24].

In particular, we locate four long-wavelength photo detectors in a quadrant around the circumference of receiving lens assembly. With the beam perfectly aligned for an operational link, we record the light intensity at the detectors. When the beam moves, the intensity on each detector will change independently; this variation is then used to provide a correction signal to the GMs at the transmitter terminal. We build a table that contains four-tuple diode readings for different 2D deviations of the TX from the initial (aligned) position. Then, to correct a misaligned link, we do a "reverse" look up in the table, i.e., find the deviation that would have resulted in (close to) the current diode readings, and use it to move the beam back via the GM at TX. Tracking feedback can be transmitted across terminals over an in-band or an out-of-band (e.g., macrocellular) control channel. The above mechanism can also be used for the initial acquisition of the receiver, and thus, yielding a complete ATP (acquisition, tracking, and pointing) mechanism.

[3]For very short links, simpler techniques can be used, e.g., [21] uses photodetector at *transmitter* with a reflective film at RX for a 1m link, but incurs a latency of several hundreds of msecs due to scanning.

A. Robust 100m Link Prototype with TP

We set up our SFP-based 100m link in a long hallway. At 100m range, our indoor link prototype should exhibit similar performance as an outdoor link in clear weather; as mentioned in §II-A, atmospheric effects are negligible at these ranges. We address weather effects by demonstrating a sufficient link margin; as mentioned in §II-A, a 100m link with 13-15dB margin can withstand most weather effects. We use the 10G 1550nm SFP+ [25] with a 10GBASE-ZR interface[4], a transmission power of 0-4dBm and a receiver sensitivity of -25dBm yielding a total optical budget of 25-29dB. For simplicity, we use a uni-directional FSO link; later, we analyze design of bi-directional links using an optical simulator. We use a GM at TX, and surround the RX's collimator by four photodiodes. See Figure 2(a)-(c). For photodiodes, we use Thorlabs' InGaAs variable-gain photodiodes [26] which have a low (70ns) rise time and include an amplifier in a compact package. Each of the photodiodes is covered by a bandpass optical filter [27] to minimize errors from stray light. The diodes are read via a DAQ [28]. We address noise in the diode readings by increasing the gain of the diodes appropriately and also by taking an average of 100 samples per diode over a few milliseconds. The collimator [22] has a divergence of 0.28mrad and generates a beam of diameter 29mm at 100m which, being larger than the 25mm collimator, falls partially on the photodiodes at RX. The entire TX assembly is placed on a motorized rotational stage [29] to enable fine angular movement of the TX assembly at varying speeds, to simulate building motions.

Performance Metrics and Requirements. We measure effectiveness of the link prototype, equipped with the TP mechanism, in terms of: (i) the angular speed (i.e., rotation rate) of TX assembly (or building motion) that the link is able to tolerate, and (ii) the link margin, which dictates its ability to handle weather effects as mentioned in §II. Prior studies [30] suggest that building motions rarely have an angular speed of more than a few mrad/sec, and below we show demonstrate that our link prototype is able to handle angular speeds of up to 15 mrad/sec. In addition to the above metrics, we also measure TP latency and link's angular tolerance, which is defined as the maximum angular motion of TX that can be tolerated (*absence of a TP system*) without losing the link. In fact, for a fixed TP latency, the angular speed of movement a link is able to tolerate must be proportional to link's angular tolerance.

Link Margin and Tolerance. At best alignment, the link's received power was around -12dBm, yielding a link margin of 13dB. In Fig. 3(a), we present reduction in RX power due to angular movement (tilt) of the TX, with or without active TP. Without TP, we move the TX slowly and observe that the RX power reduces minimally at 0.1 mrad tilt, and remains

[4]The physical layer standard used by 10GBASE-ZR interfaces is proprietary and not specified under IEEE 802.3ae.

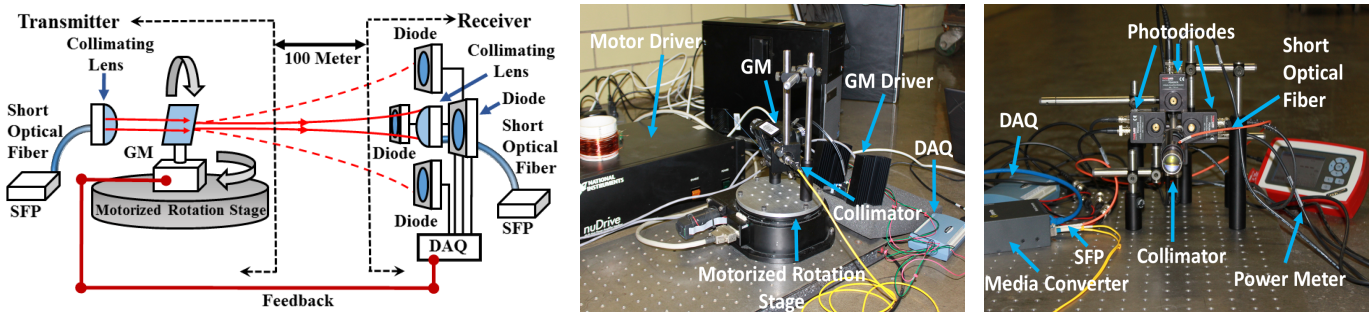

Fig. 2: (a) Schematic diagram of the 100m link prototype, (b) Transmitter assembly with GM and motorized rotational stage, and (c) Receiver assembly with 4 photodiodes around the collimator.

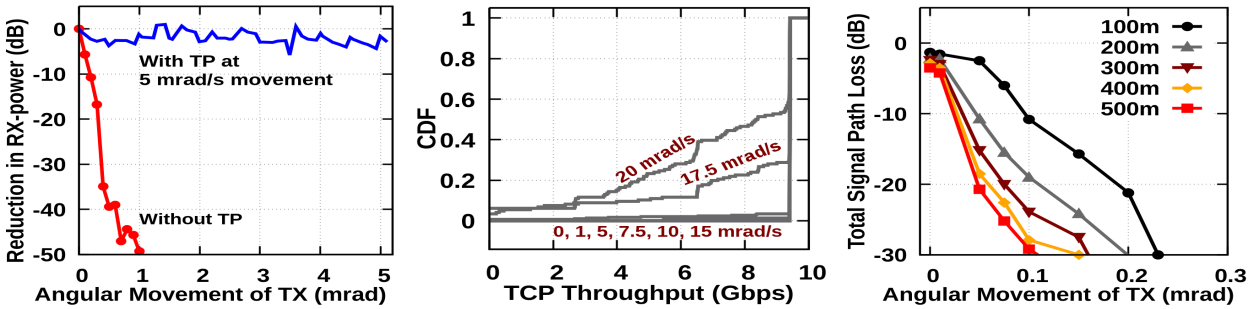

Fig. 3: (a) Reduction in received power, with and without TP, (b) CDF of the link's TCP throughput for varying angular speeds of the TX assembly, (c) Path loss estimates based on Zemax for varying angular tilts and ranges.

less than the initial link margin (13dB) up to a tilt of 0.2 mrad—implying an angular tolerance of a bit more than 0.2 mrad. With active TP, we use the motorized rotation stage to provide an angular movement at 5 mrad/sec, and observe that the RX power remains relatively stable with a reduction of a few dB over a wide range (up to 100 mrad; figure shows only up to 5mrad) of angular movement.

TCP Throughput for Varying Angular Speeds. To demonstrate the link operation during continuous terminal movement, we compute the CDF of link's TCP throughput with TP active and TX rotating at varying angular speeds: 0–20 mrad/sec, using a motorized rotational stage, with an amplitude of 100 mrad. See Figure 3(b). Here, each run is over a 10 minute period, with throughput measured over every second. We observe that the throughput CDFs for speeds $\leq$ 15 mrad/sec are near-identical to that of the fixed link, with an average throughput of about 9.4 Gbps. The throughput deteriorates somewhat for 17.5 and 20 mrad/sec. We also logged RX power every 3-5 msec using [31] and observed that it remained around -15dBm for 0-15 mrad/sec speeds; this implies a link margin of 10dB and thus ability to handle fog conditions with visibility more than 130m even *during* 15 mrad/s motions.

TP Latency. The latency incurred in one iteration of the TP mechanism is around 6 msecs. Reading 100 samples from the diodes takes about 4 msecs, while the remaining tasks (feedback transfer, correction computation, GM latency) take about 2 msecs.

Analyzing Longer (200-500m) Links. Creating longer (200-500m) link prototypes is beyond the scope of this current work. Instead, we model an FSO link in the optical simulator Zemax [32] based on given optical elements and estimate link's signal loss at various ranges. Recall that, for links longer than

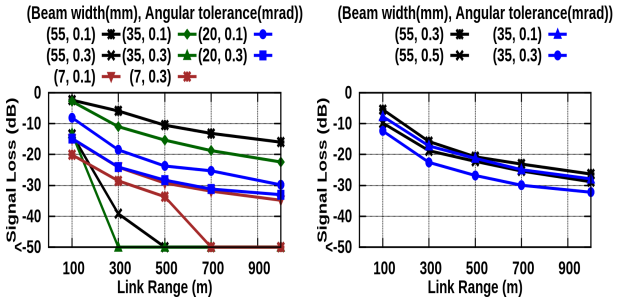

Fig. 4: Signal loss for varying link ranges for various (beam width, angular-tilt) combinations for (a) 50μm and (b) 200μm MMF connecting the the SFP+ to the collimator.

100m, we don't allocate any link margin for weather effects—so, the link should just have sufficient angular tolerance to support the active TP mechanism to handle expected building motions. Fig.3(c) shows that, assuming an optical budget of about 30dB, the angular tolerance is approximately 0.23 mrad for 100m, 0.15 mrad for 300m, 0.1 mrad for 500m, etc. Based on the performance results (Fig. 3(b)) of the 100m link prototype, the 0.1 mrad angular tolerance of a 500m link should suffice to handle motions of 5–7.5 mrad/sec angular speeds.

Bi-Directional FSO Links. We now analyze design of bi-directional links in the optical simulator Zemax. Bidirectional links use two beams at different wavelengths to carry data in both directions, using bidirectional SFP+s [7]. This reduces the hardware cost compared to two independent unidirectional links. Bidirectional links must have a symmetric design—in particular, the same optical fiber launches as well as captures the beam. Single-mode fiber (9μm core diameter) is not a

viable for bidirectional link [13]. Thus, we explore use of multi-mode fibers (MMF) (core diameter 50–200μm). [5] However, a beam launched via an MMF has a higher divergence, which can be reduced by using a wider beam via a "beam expander" assembly. To explore various design choices, we model a link in the Zemax simulator, and analyze link's signal loss (from geometric, optical system, and coupling) at various ranges and angular tilts for various choices of launch fibers and beam widths. Recall that, for a viable FSO link for our purposes, it should have an angular tolerance of 0.1-0.2 mrad, and in particular, as per our link budget allocation in §II-A, should incur signal loss of at most 10-15dB at 100m and 25dB at 500m, at angular tilts of 0.1-0.2 mrad. In Figure 4(a)-(b), we observe that our above link requirements are met by various design combinations: (i) 200μm fiber provides an angular tolerance of almost 0.3 mrad with 35mm beam width; (ii) 50μm fiber can provides up to 0.1 mrad angular tolerance with a beam width of 20mm or more.

IV. Dynamic Backhaul-Network Design (DBND)

In our context, the network design problem entails placement of picocells and creating "potential" FSO links over the located picocells covering an area. Creation of potential FSO links in turn entails orientation of GMs associated with the FSO devices. In this section, we define the problem after providing relevant definition. We term this as "Dynamic Backhaul Network Design" (DBND) problem. Since, the problem is NP-hard, we provide a heuristic solution to it.

Candidate Links, Dynamic Network, Realizable Topologies. Each picocell/node has a coverage area associated with it. To provide backhaul connectivity to gateways, each node is connected to a small number of steerable FSO devices via a network switch. For steerability, each FSO is equipped with a GM which is preconfigured to a coverage cone within which the FSO can steer its beam. See Figure 1. A link can be potentially established between two FSOs if there is a clear line of sight between them and each is contained in the coverage cone of the others' GMs. We refer to such links as *candidate link*. The set of all candidate links forms what is called a *dynamic network*. At runtime, the network is "reconfigured" to activate a subset of candidate links (by steering the beams appropriately) such that each only one candidate link per FSO is active; this network reconfiguration is done based on the prevailing traffic and state, and is discussed in the following section. A subset of candidate links that can be active simultaneously forms a *realizable topology* of the given dynamic network.

DBND Problem. Given a geographical area, which needs picocell coverage, potential locations where picocells can be placed, and locations of a few fixed gateways connected to the outside network, the DBND problem is to design a dynamic network $\mathcal{D}$ with maximum "average-hotspot-flow" (as defined below) under the following constraints: (i) $\mathcal{D}$ must have a realizable topology that is a backbone (as defined below), (ii) the total number of nodes in $\mathcal{D}$ is less than a given constant k, and (iii) the number of FSOs per node in $\mathcal{D}$is bounded by a given constant k'. Solving the DBND problem entails determining (a) locations to place picocell nodes, and (b) candidate links (with clear line-of-sight) over these placed nodes by orienting the GMs associated with the nodes' FSOs.

DBND's Constraints. The created dynamic network $\mathcal{D}$ must have a realizable topology τ called *backbone* such that (i) τ consists solely of short links (i.e., ≤ 100 m), (ii) the set of nodes of τ provides full coverage to the given area, and (iii) each node i in τ is connected to a gateway via a path in τ. See Fig. 5(c)-(d). The purpose of the backbone is to maintain network availability (i.e., area coverage) in most adverse outdoor conditions, by ensuring its short links have sufficient link margin. To constrain infrastructure cost, we also impose constraints on the total number of nodes as well as the number of FSOs per node in $\mathcal{D}$.

DBND's Optimization Objective. At a high-level, we wish to create a "flexible" dynamic network $\mathcal{D}$ that has a rich set of realizable topologies sufficient to "handle" most of the anticipated traffic patterns. As observed in §II, we are particularly interested in traffic patterns with only a small number of hotspot nodes. Thus, a reasonable optimization objective is as follows. Let X be the fixed number of expected hotspots, and let $\mathcal{P}$ be the set of all possible X-size subsets of nodes in $\mathcal{D}$. Then, the *average-hotspot-flow (AHF)* optimization objective is to *maximize* the following:

$$\boxed{\text{avg}_{p\in\mathcal{P}} \max_{\tau} \text{ (max-flow from } p \text{ to gateways in a realizable topology } \tau \text{ of } \mathcal{D})} \quad (1)$$

Essentially, the above objective is to maximize the average (across all possible sets of hotspots) network-flow from the hotspots to the gateways; note that the objective implicitly also depends on the path lengths from hotspots to the gateways.

Intractability, and Four-Step Heuristic (FSH). It is easy to see that the DBND problem is NP-hard [1] via a straightforward reduction from the set cover or Steiner tree problem. The closest known problem is the *Group Steiner Tree* (GST) problem [34], which, in our context, represents the "partial" problem of constructing the smallest backbone without the node degree constraint. The known approximation algorithms [34], [35] for GST, though polynomial, are very inefficient [34] to be useful here. The key challenges in solving the complete DBND problem arise from the node degree constraint and the network dynamicity which requires an intricate objective function. We propose the following multi-step heuristic (referred as the *Four-Step Heuristic (FSH)*) that has no theoretical worst-case performance guarantees but delivers good quality network solutions in real settings as shown in §VI. The basic idea (see Fig. 5) is to first construct a backbone with minimal number of nodes, and then, augment it with additional nodes and links to optimize the objective.

1) Select a minimal number of nodes n_c to cover the given area, using a set-cover like greedy approach.
2) Connect each of the nodes in n_c to a gateway using only short-range links and possibly, additional nodes; this

[5] Intermodal dispersion due to MMFs will be negligible, as we will use a very short fiber (a few cms). MMFs can easily transmit a 10Gbps signal over few hundred meters [33].

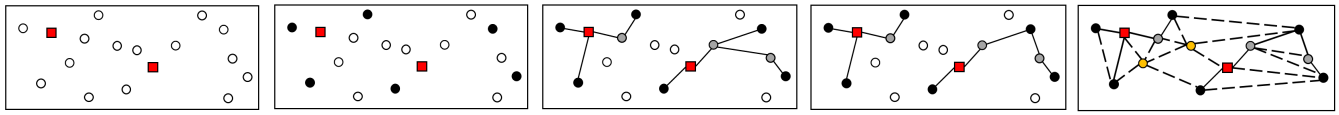

(a) (b) (c) (d) (e)

Fig. 5: Four-Step Heuristic. (a) Given a geographic area (the rectangle), set of potential locations, and gateways (red). (b) Set of nodes n_c chosen to cover the area. (c) Steiner forest connecting the nodes n_c to the gateways. (d) Reduction of the maximum node (non-gateway) degree, (e) Augmenting the backbone with additional nodes (yellow) and links (dashed).

"Steiner-forest" problem is a slight variation of the well-known Steiner-tree problem [36]. This yields a backbone subnetwork with the desired properties.

3) Minimize the maximum degree of a node, by an iterative "exchange" of links (using a slight variation of the Furer-Raghavchari's trick [37]), following by a reconstruction phase as described below. This step addresses the constraint on number of FSOs per node.
4) Finally, we augment the above backbone with additional nodes and links to introduce dynamicity and improve its objective value. We do this by iteratively adding a "path" of nodes that most improves the objective.

Tailored General FSO Networks. In general, we would like to design optimal FSO networks tailored to a given geography or weather statistics. That is, given a geographical area and associated weather statistics, cost constraints, and network availability requirements, our goal is to design a dynamic FSO network with optimal *expected* performance (i.e., expectation of AHF over expected weather conditions). We can generalize our FSH heuristic to solve the above general problem as follows. First, we determine the highest attenuation rate ρ (dB/km) that the FSO network must be able to handle; this is done by mapping each weather condition to signal attenuation rates, and considering the "worst" weather condition that must be handled (note that the network availability requirement need not be 100%). Then, we can generalize the above described FSH to create a dynamic network as follows: (i) To create a backbone subnetwork, we use "short" links of range l kms that have a a minimum link-margin of ρ. If there are multiple feasible values of l, then we design a dynamic network for each l, and pick the one with optimal objective value. (ii) In the final step of augmenting the created backbone, the optimization objective is an *expectation* over varying weather conditions; this is handled by considering the availability of a path for each weather condition.

V. Runtime Network Management

We now address the runtime reconfiguration of the network, which requires us to select a realizable (runtime) topology from the given dynamic network and compute routes for the flows over the selected topology, based on the prevailing traffic and network state. Optimization objective could be to maximize the total traffic demand satisfied. Since reconfiguration happens at runtime, we must design very fast algorithms.

Network Management Overview. We leverage software-defined networking (SDN) [38], [39], [40] in designing the network management layer, that essentially coordinates link and network reconfiguration. The network switches, which connect the FSO devices at each node, are envisioned to be SDN-capable receiving configurations (i.e., steering commands, routing table entries, etc.) from a logically centralized controller. The controller continually receives the status of the entire network status (e.g., link status, traffic patterns) from the SDN-enabled switches [39] via a control channel. The key component of the controller is the *reconfiguration engine*, which is described below. We preserve network connectivity and consistency during reconfigurations [41], [13] via techniques from [13]. except that we execute reconfigurations one at a time rather than concurrently as they are expected to involve more links but occur less frequently in our context.

FSONet's Reconfiguration Engine. We reconfigure FSONet in response to certain network events of interest. In the context of a picocell network, the main event of interest is mostly arrival of a flow. We model a *flow* as arriving at a certain node with a *bandwidth-demand*. An effective reconfiguration strategy should be infrequent, fast (run in a few tens of milliseconds), and incremental (incur minimal topology and routing-table changes). Note that, in FSONet, the network is relatively sparse and all flows are to and from a small number of gateway nodes; thus, we employ a simple and fast algorithm which is somewhat localized and done only when beneficial. See Algorithm 1. Here,

Algorithm 1: Runtime Network Management

Data: A flow f with bandwidth-demand b_f.
Parameters. β (large-flow threshold), τ (high satisfaction threshold)
Result: **Accept** and **Route** f, or **Reject** f
if *$b_f \le \beta$ or if τ% of b_f is satisfiable* **then**
  Accept and **Route** f using current routing table without reconfiguration;
else if *a multi-path route P (without rerouting existing flows) is found in the current topology such that P meets $\ge \tau$% of b_f* **then**
  Accept and **Route** f using path P;
else
  Find a multipath route P in the given dynamic network such that:
  i) P involves at least one currently-inactive link
  ii) Activation of links in P does not require deactivation of any "heavily-used" links nor disconnection of any node from the gateways
  if *P found* **then**
    Accept and Route f using P
  else
    Reject flow f.

Link outages can be handled based on their duration: (i) short outage: retransmit dropped packets (e.g., due to bird

blockage), (ii) longer outage: reroute flows, by reconfiguring links, if needed, (iii) semi-permanent outage: reconfigure network. Note that the above is oblivious of short vs. long links, though in the case of semi-permanent outage (e.g., due to unfavorable weather) the network may be reconfigured to a topology with only small links.

VI. Evaluation

In this section, we evaluate the quality of FSONet networks created for multiple cities in the world. In particular, we first evaluate the network quality in terms of #-of nodes (and other metrics) needed to provide coverage and connectivity via our proposed architecture for different cities in the world. Then, via extensive traffic simulations, we evaluate the online performance of the FSONet. In particular, we compare the performance of FSONet with that of a hybrid FSO/RF architecture, which is another known architecture to address outdoor challenges faced by FSO links.

3D Building Data of Large Cities. In our simulation study, we create FSONet networks to provide *outdoor* coverage in 8 large metro cities around the world. We use topological data from [42] which contains the locations and 3D geometry of the buildings in the cities. We use building corners as the *potential locations* for picocell base stations. We determine line of sight between each pair of corners of the buildings.

Parameter Values. We vary the (ground) coverage radius of a node from 100 to 200 meters, with 100 m as default; We allow 3 to 6 FSOs per base station, with 4 as the default; note that only a low number is practical. We randomly choose one node as gateway for every 1km × 1km area of each map. To constrain the total number of nodes: note that the first three steps of FSH attempt to create a small size backbone — an *essential* part of a FSONet network. Thus, we constrain only the number of nodes added in the fourth step of FSH; we allow at most $n_b/10$ nodes to be added in the fourth step, where n_b is the created backbone size.

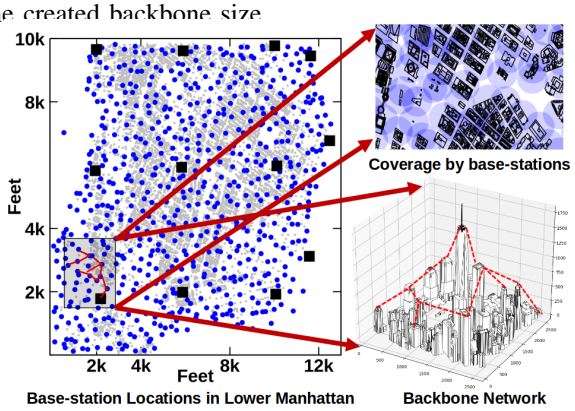

Fig. 6: Basestation locations computed by FSH algorithm for NYC; here, *blue dots*: locations chosen at step 1; *gray dots*: available corners; *red lines*: backbone links.

A. FSONet Network Quality

Given the above set-up, we create FSONet networks in 8 cities, and focus on evaluating (i) the number of nodes n_b in the backbones and (ii) the average-hotspot-flow (AHF) objective, in the created FSONet networks.

Coverage	Backbone (n_b)	Coverage Factor $\left(\frac{n_c}{A/\pi r^2}\right)$	Conn. Factor (n_b/n_c)	AHF Ratio	
				FSONet / Static	FSONet /Upper-bound
NYC	848	1.85	1.25	3.82	0.85
Boston	328	1.75	1.15	3.16	0.82
Austin	329	1.90	1.24	3.12	0.83
Philly	325	1.88	1.16	3.25	0.81
Berlin	454	1.86	1.18	3.58	0.86
Montreal	393	1.84	1.26	3.72	0.85
Vienna	358	1.92	1.30	3.22	0.88
Adelaide	338	1.85	1.28	3.45	0.81

TABLE II: Statistics for FSONet networks for 8 cities.

Coverage	Backbone (n_b)	Coverage Factor $\left(\frac{n_c}{A/\pi r^2}\right)$	Conn. Factor (n_b/n_c)	FSO per Node	AHF Ratio	
					FSONet /Static	FSONet /Upper-bound
100 m	847	1.85	1.25	3	4.18	0.82
				4	3.82	0.85
				5	3.25	0.88
				6	3.68	0.91
150 m	378	1.90	1.21	3	4.25	0.83
				4	3.91	0.84
				5	3.35	0.90
				6	2.70	0.92
200 m	212	1.93	1.18	3	4.31	0.84
				4	3.90	0.86
				5	3.36	0.92
				6	2.86	0.93

TABLE III: Statistics for FSONet network in NYC for different values of coverage radius and #-of FSOs/node.

Backbone Size (n_b) and Related Ratios. Recall that the backbone subnetwork is an essential part of FSONet network, and hence represents the smallest FSONet network possible using our techniques. First, we observe that the backbone size n_b varies from about 325 to 848 nodes across the cities, for default coverage radius of 100m. See Table II. Delving in deeper, we observe that n_c, the number of nodes selected in FSH's first step to provide coverage, is only about 1.75 to 1.92 times more than the loose lower bound (= total coverage area divided by $\pi 100^2$) which ignores both the coverage area's irregularity and the limited locations available to place nodes. Moreover, we observe that ratio of n_b to n_c is 1.15 to 1.3; this essentially represents the cost of enforcing the backbone condition (the key aspect of FSONet).

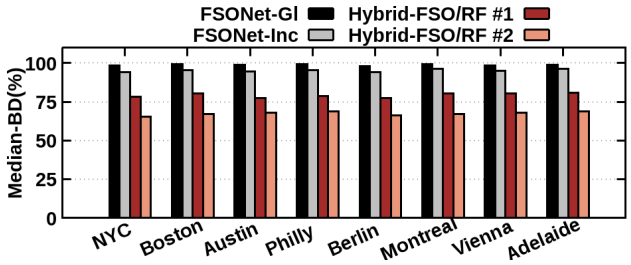

Fig. 7: Comparison of `Median-BD` values for various cities for default parameters.

Average-Hotspot-Flow (AHF) Comparison. To evaluate the quality of the created FSONet networks, we *estimate* their AHF values and compare them to that of a representative "static" (i.e., with no steering FSO capability) network as well

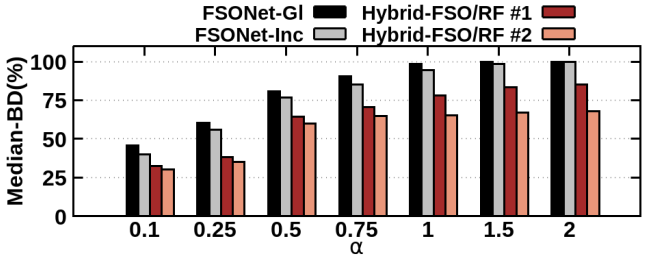

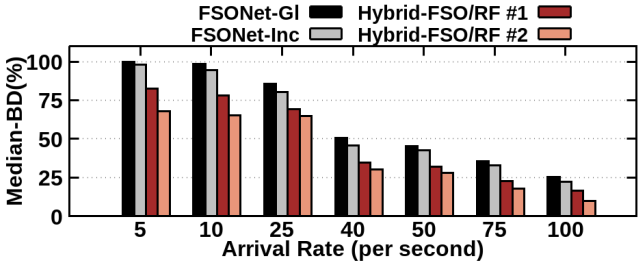

Fig. 8: `Median-BD` values (a) for varying α with flow arrival rate of 10K/sec. (b) for varying flow arrival rates with $\alpha = 1$.

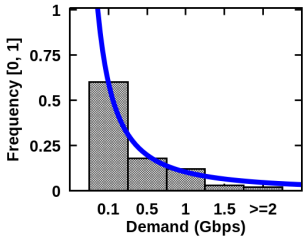

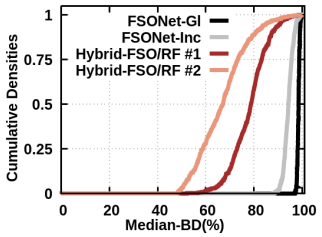

Fig. 9: (a) Distribution of bandwidth demands, and (b) CDF of `Median-BD` values for NYC.

as a loose upper-bound. In particular, for a FSONet network $\mathcal{D}$, we compare its AHF value with the static network that contains $\mathcal{D}$'s backbone subnetwork, has the same number of nodes as $\mathcal{D}$, and wherein each node's degree is at most the number of FSOs per node in $\mathcal{D}$. To estimate the AHF value, we use Eqn.(1) with $X = 5$ and consider about 200 random subsets p of size 5. See Table II. We make two observations about the AHF estimate of FSONet networks: (i) it is about 3.12 to 3.82 times the AHF estimate of the corresponding static graph—confirming the benefit of steerability/dynamicity in a network, (ii) it is about 0.81 to 0.86 of the trivial upper bound of 200 Gbps (= number of hotspots *times* the maximum number of FSOs per node *times* 10Gbps link capacity)—suggesting the effectiveness of our techniques.

Table III shows range of statistics for NYC for other values of coverage radius and # of FSOs/node, and we observe similar statistics.

B. Online Performance

Traffic Model. To evaluate the online performance of FSONet networks (for default radius 100m and 4 FSOs/node), we run traffic simulation and measure appropriate performance metric. In particular, we consider a traffic model wherein the arrival of flows in the network is modeled as a Poisson process with rate varying from 5k to 100k (default being 10k) per second; each flow is randomly assigned to arrive at a network node. The flow's size is modeled as a Pareto random variable α in the range 10 KB to 1 TB; the default value of α is 1 so that 90% of the flows are less than 100KB. We assign the bandwidth-demand as follows: with a 10% probability we pick a random rate between 100Mbps and 6Gbps, and with the remaining 90% probability, we assign the bandwidth-demand increasingly based on the size (e.g., 10-10,000KB flows get 100Mbps, 10-100MB get 500Mbps, and so on). The latter can be implemented by increasing the bandwidth-demand as more and more of the flow size is routed and discovered. The maximum values of size and bandwidth-demand roughly correspond to high-resolution multi-hour movies. Figure 9(a) shows the CDF of the bandwidth demands for a simulation run for the default setting.

Reconfiguration Schemes. We consider the following reconfiguration schemes: (i) `Inc`: Algorithm 1, the incremental approach, executed for every large flow with β = 100Mbps [39] and $\tau = 90$, (ii) `Gl`: the global periodic-reconfiguration of [1] executed at the arrival and completion of every large flow. We term the global and incremental update version of FSONet as FSONet-Gl and FSONet-Inc respectively. Note that the `Gl` approach is impractical due to high computation time and global changes to the network but serves as an apt comparison benchmark (upper-bound).

Hybrid FSO/RF Setting. In addition to above, we also compare the performance of our proposed architecture with the hybrid FSO/RF architecture [43], wherein the FSO outdoor challenges are handled by equipping each link with a back-up RF transceiver and using the RF link (which is largely unaffected by weather conditions) whenever FSO link is inoperable due to inclement weather conditions. To the best of our knowledge, hybrid FSO/RF is the only other technology known to circumvent the outdoor challenges in FSO, while maintaining an all-wireless architecture. In particular, we consider two hybrid FSO/RF architectures: (i) In `Hybrid-FSO/RF #1`, we select a random *static* topology T from the created FSONet and replace each FSO link in T with a hybrid FSO/RF link. (ii) In `Hybrid-FSO/RF #2`, we relax the short-links restriction in FSONet, as it is not pertinent to the hybrid FSO/RF architecture; in particular, for `Hybrid-FSO/RF #2`, we design a network of hybrid FSO/RF links by using 1-3 steps of FSH (§IV) but allowing up to 500m links in the second step. In both the hybrid architectures, we model the RF link as a 60GHz mmWave link [44][6]. The 60GHz spectrum is unlicensed, and the links can potentially support data rates of a few Gbps for short distances— in particular, the RF capacity varies from 8Gbps at few meters to 2Gbps at 500 meter (see Figure 11). Note that the cost

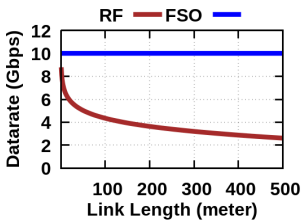

Fig. 11: Throughput of 60GHz RF link [44]

[6]Conservatively, we assume that, via spatial reuse, the same mmWave band may be used for cellular coverage.

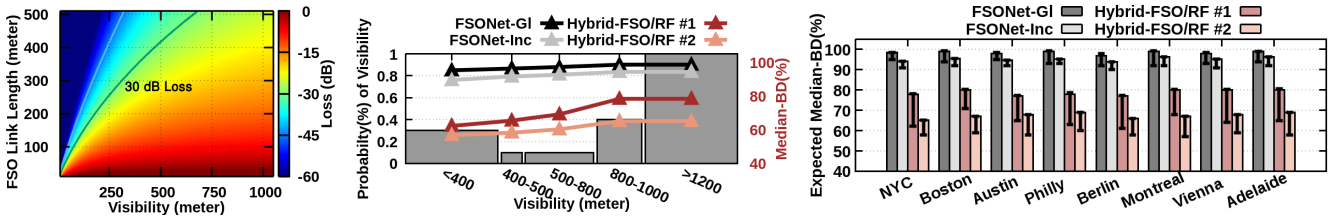

Fig. 10: (a) Heatmap of loss (dB) for varying visibility and FSO-link lengths (b) Probability distribution (left-Y axis) of visibility in NYC according to [9] and corresponding `Median-BD` values (right-Y axis) (c) Comparison of `Expected Median-BD` for different cities (the candlesticks show the difference between best and worst values).

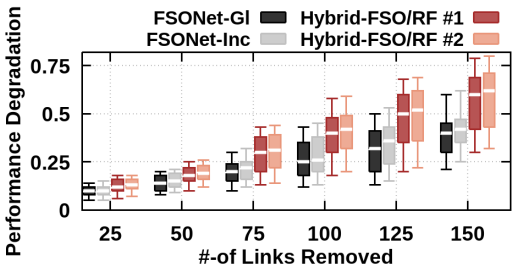

Fig. 12: Performance degradation for varying # of link failures for FSONet and Hybrid FSO/RF in NYC

of Hybrid FSO/RF architectures is only higher than that of the corresponding FSONet architecture, as the hybrid links need an additional transceiver for each link;[7] the steering mechanism needed for reconfigurability in FSONet is needed anyway for TP mechanism to handle building motions (§III).

Performance Metric and Comparison. We run the simulation for sufficient time to reach a steady state and measure, at each time instant, the total bandwidth-demand satisfied as a percentage of the total requested by all the flows at that time, and then we take the median of all the values (since unlike mean, median is resistant to outliers). We use `Median-BD` to refer to this average metric, and use it for comparison.

In Figure 7, we compared `Median-BD` metric values between FSONet and Hybrid FSO/RF networks for all the cities. We observe that i) FSONet outperforms Hybrid networks by a margin of at least 25% ii) `Inc` performs very close to `Gl` differing by only up to 2% margin. Figure 8(a) and (b) show the metric values for varying size parameter α of Pareto distribution and Poisson arrival rate of flows respectively for NYC. In all cases FSONet outperforms the rest. The CDF plots for default parameters in Figure 9(b) show that- 90% of the `Median-BD` values are > 90% for FSONet.

Network Redundancy. To evaluate and compare network redundancy for FSONet and Hybrid FSO/RF networks, we randomly deactivate the same set of links in FSONet and hybrid networks and measure the above defined `Median-BD` metric for each, and compare the "performance deterioration". We define the *performance deterioration* as $(a - b)/b$ where b and a are the `Median-BD` values of the network before and after deletion of links. See Figure 12. We observe that the NYC FSONet network exhibit much less deterioration in performance than the hybrid networks.

Weather Simulation. Figure 10(a) plots the signal attenuation (dB/km loss) of an FSO link for various visibility and link lengths, as per [45]. Based on the optical budget of 30dB of our FSO links, we observe that for FSO links of all ranges upto 500m to be operational, the fog conditions should allow for at least 750m visibility. On the other hand, short links ($\leq$ 100m) are operational as long as the visibility is higher than 50m. Figure 10(b) shows the visibility distribution for New York City (histogram bars) along with corresponding metric values for both FSONet, and Hybrid FSO/RF networks. In general, based on the distribution of varying visibility conditions of cities from [9], we estimate the `Expected Median-BD` values for all cities as shown in Figure 10(c). We observe that the expected values are almost equal to the ones without weather effect, because the probability of low visibility is very small i.e, 0.2% or less ($\sim$ 1 out of 365 days). However, the candlesticks in the plot show the performance degradation due to inclement weather; we note that the performance degradation for Hybrid FSO/RF network is 4 times or more than that of FSONet.

VII. Related Work

Backhaul of picocell networks is a key challenge, and recent works have explored various solutions. Among wireless backhaul solutions, [46] addresses resource (power and scheduling) allocation in backhaul solutions with shared wireless channels, [47] presents beamforming techniques to overcome outdoor challenges in an mm wave based solution, and [48] explores a microwave-based backhaul solution. However, low frequency radio (below 6GHz) is limited in terms of data rates due to interference problems, and high frequency radio (mm and microwave from 6 to 300 GHz) are limited in range and data rates (e.g., mmwave can only deliver upto 1-2Gbps over 100-200m). Recently, free space optics technology (FSO) has emerged as an attractive alternative. E.g., [49] explores use of FSO flying platforms to create a backhaul network with vertical point-to-point FSO links; their motivation for vertical links is to circumvent the line-of-sight issue in dense urban areas. In addition, to handle outdoor effects, some works have explored hybrid RF/FSO [43], [50] or wired/FSO [50] solutions. Instead, we alleviate the line-of-sight concern by a simulation study over many metropolitan US cities, and circumvent the outdoor challenges in an FSO-based backhaul by using many short links and embedding redundancy in the network via steerable links.

[7]The cost of `Hybrid-FSO/RF #2`architecture is lowered due to lesser number of nodes (about 90% of the nodes in FSONet) but this is easily offset by the cost of additional RF transceiver needed.

Our work is inspired by recent works on use of steerable FSO links for dynamic data center networks [13], [51]. In our context, the outdoor setting bring new challenges in FSO link engineering and network design. E.g., in a data center as the FSO devices are placed on top of racks, there is no placement problem. Moreover, in [13], a simple random network worked well as the dynamic network [13] as the racks were treated uniformly and the line of sight issue was circumvented by placing a full ceiling mirror. Similarly, in [51], use of a large number of intricate mirror assemblies allowed a complete dynamic network. In our context, the relative sparsity of network and specialized traffic pattern due to gateways calls for different reconfiguration strategies as discussed in §V.

VIII. Conclusion

We have developed FSONet, a wireless backhaul or multi-gigabit picocells based on steerable FSO links. To the best of our knowledge, this is the first work targeting multi-gigabit wireless solutions for picocell backhauls. We have explored one viable solution and addressed design challenges therein. However, there is room for much improvements and alternative choices in architecture design and techniques, e.g., to comprehensively address key cost-performance trade-offs. Could we design efficient FSO *networks* tailored to different weather and outdoor conditions? A somewhat generalization of FSONet's approach could be to design multi-tier FSO networks with say high-bandwidth but less-reliable subnetwork backed up a low-bandwidth but very reliable subnetwork. These directions form the basis of our future research.

References

[1] M. Curran *et al.*, "FSONet: A Wireless Backhaul for Multi-Gigabit Picocells Using Steerable Free Space Optics," in *MobiCom*, 2017.

[2] J. Andrews *et al.*, "What will 5G be?" *IEEE JSAC*, 2014.

[3] Y. Zhu et al., "Demystifying 60GHz outdoor picocells," in *MobiCom*, 2014.

[4] E. Ciaramella *et al.*, "1.28-Tb/s free-space optical WDM transmission system," *IEEE Photonics Tech. Letters*, 2009.

[5] I. Akyildiz *et al.*, "Terahertz band: Next frontier for wireless communications," *Physical Communication*, 2014.

[6] "Small cell backhaul requirements," http://tinyurl.com/h8fomey, 2012, nGMN Alliance, White paper.

[7] "Bi-directional 10G SFP+," http://www.fs.com/products/31028.html.

[8] M. Awan *et al.*, "Characterization of fog and snow attenuations for free-space optical propagation," *Journal of Communications*, 2009.

[9] I. I. Kim and E. J. Korevaar, "Availability of free-space optics and hybrid FSO/RF systems," in *Proc. SPIE*, 2001.

[10] "Lightpointe FSO Links," http://www.lightpointe.com.

[11] "Artolink 10G FSO Link Specs," http://artolink.com/page/products/free_space_optics_Artolink_10Gbps/.

[12] "LightPointe FSO Link Cost," http://tinyurl.com/k86o2vh.

[13] N. Hamedazimi *et al.*, "FireFly: A reconfigurable wireless data center fabric using free-space optics," in *ACM SIGCOMM*, 2014.

[14] "Thorlabs Galvo Mirrors," http://tinyurl.com/haymezj.

[15] T.-H. Ho, "Pointing, acquisition, and tracking systems for free-space optical communication links," Ph.D. dissertation, University of Maryland, College Park, 2007.

[16] S. V. Kartalopoulos, *Free Space Optical Networks for Ultra-Broad Band Services*. John Wiley and Sons, 2001.

[17] C. Lv, H. Jiang, and S. Tong, "Implementation of FTA with high bandwidth and tracking accuracy in FSO," in *CECNet*, 2012.

[18] M. K. Al-Akkoumi, H. Refai, and J. J. Sluss, "A tracking system for mobile FSO," in *SPIE Proceedings Vol. 6877*, 2009.

[19] T. Yamashita *et al.*, "The new tracking control system for free-space optical communications," in *ICSOS*, 2011.

[20] M. S. Ferraro *et al.*, "InAlAs/InGaAs avalanche photodiode arrays for free space optical communication," *Appl. Opt.*, vol. 54, no. 31, 2015.

[21] P. Deng *et al.*, "MEMS-based beam steerable free space optical communication link for reconfigurable wireless data center," in *Proc. of SPIE*, 2017.

[22] "F810FC-1550 Collimation Package from Thorlabs," https://www.thorlabs.com/thorproduct.cfm?partnumber=F810FC-1550.

[23] H. Weichel, *Laser Beam Propagation in the Atmosphere*. SPIE Press, 1990, chapter 5.

[24] S. Bloom, "The physics of free space optics," http://tinyurl.com/yb4cyrc4, 2002, white paper.

[25] "10G SFP+ 1550nm 100km DOM Transceivers," http://www.fs.com/products/29799.html.

[26] "Photodetector," https://www.thorlabs.com/thorproduct.cfm?partnumber=PDA20CS.

[27] "1550nm Bandpass Filter," https://www.thorlabs.com/thorproduct.cfm?partnumber=FB1550-12.

[28] "Multifunction USB Data Acquisition Device," http://www.mccdaq.com/usb-data-acquisition/USB-1608G.aspx.

[29] "Micos," http://www.micosusa.com/product/subview.cfm_firstlevel=1&sublevel=17.htm.

[30] A. G. A. Muthalif *et al.*, "Active vibration isolation system to improve free space optics communication," in *Proc. of the Intl. Conf. on Info. Tech. & Comp. Intelligence*, 2013.

[31] "Thorlabs PM100D Power Meter," https://www.thorlabs.com/thorproduct.cfm?partnumber=PM100D.

[32] "Zemax," https://www.zemax.com.

[33] "BER and Eye-Diagram of 10Gpbs signal over multi-mode fiber," https://www.youtube.com/watch?v=Ab2h1oqZ-GA.

[34] "Group Steiner Tree Problem," http://theory.cs.uni-bonn.de/info5/steinerkompendium/node30.html.

[35] N. Garg, G. Konjevod, and R. Ravi, "A polylogarithmic approximation algorithm for the group steiner tree problem," in *SODA*, 1998.

[36] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, 1979.

[37] M. Fürer and B. Raghavachari, "Approximating the minimum-degree steiner tree to within one of optimal," *J. Algorithms*, vol. 17, no. 3, 1994.

[38] M. Casado *et al.*, "Ethane: Taking control of the enterprise," in *Proc. SIGCOMM*, 2007.

[39] A. Curtis *et al.*, "DevoFlow: Scaling flow management for high-performance networks," in *ACM SIGCOMM*, 2011.

[40] M. Al-Fares *et al.*, "Hedera: Dynamic flow scheduling for data center networks," in *NSDI*, 2010.

[41] M. Reitblatt, N. Foster, J. Rexford, C. Schlesinger, and D. Walker, " Abstractions for Network Update," in *Proc. SIGCOMM*, 2012.

[42] CityGML, "3D Building Model," https://www.citygml.org/3dcities/.

[43] Ahmed Douik et al., "Hybrid radio/free-space optical design for next generation backhaul systems," *CoRR*, 2015.

[44] L. Verma *et al.*, "Backhaul need for speed: 60 ghz is the solution," *IEEE Wireless Communications*, vol. 22, pp. 114–121, 2015.

[45] M. A. Esmail *et al.*, "Outdoor fso communications under fog: Attenuation modeling and performance evaluation," *IEEE Photonics Journal*, vol. 8, pp. 1–22, 2016.

[46] I. Maric et al., "Resource allocation for constrained backhaul in picocell networks," in *Info. Theory and App. Workshop*, 2011.

[47] S. Hur et al., "Millimeter wave beamforming for wireless backhaul and access in small cell networks," *IEEE Trans. on Comm.*, 2013.

[48] M. Coldrey et al., "Non-line-of-sight microwave backhaul in heterogeneous networks," in *IEEE VTC*, 2013.

[49] M. Alzenad *et al.*, "FSO-based vertical backhaul/fronthaul framework for 5G+ wireless networks," *CoRR*, vol. abs/1607.01472, 2016.

[50] Yuan Li et al., "Optimization of free space optical wireless network for cellular backhauling," *IEEE JSAC*, 2015.

[51] M. Ghobadi *et al.*, "Projector: Agile reconfigurable data center interconnect," in *ACM SIGCOMM*, 2016.