

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

Max Curran[§], Md. Shaifur Rahman[§], Himanshu Gupta, Kai Zheng, Jon Longtin, Samir R. Das, Thanvir Mohamed
Stony Brook University
Stony Brook, NY

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.

CCS CONCEPTS

• **Networks** → **Network design and planning algorithms**; • **Hardware** → *Emerging optical and photonic technologies*;

KEYWORDS

Small Cells, Backhaul, Free Space Optics, Reconfigurable Networks

1 INTRODUCTION

It is widely expected that the explosion of mobile broadband data consumption will continue unabated. Some projections estimate a 1000-fold increase in capacity demand within a decade [32]. Researchers have been addressing this impending ‘capacity crunch’

[§] Co-first student authors.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiCom '17, October 16–20, 2017, Snowbird, UT, USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-4916-1/17/10...\$15.00

<https://doi.org/10.1145/3117811.3129239>




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

using various mechanisms. These include improving spectral efficiency, developing shared spectrum regimes (e.g., white spaces), and utilizing very high frequency spectrum bands (e.g., 30-300 GHz). While all these mechanisms are bringing dividends, it is well understood that to keep up with the anticipated capacity demand, *spatial reuse* of wireless spectrum must improve dramatically. Spatial reuse has indeed been one of the key strategies used — e.g., cell radius has fallen steadily from tens of kms in the 1980s to just a few kms currently [32].

Picocells’ Backhaul Network. The above trend of reduced cell sizes has recently led to a significant interest in *picocells* with ranges of around 100–200 m or even less, specifically using the millimeter wave bands (30–300 GHz) [37, 65, 68]. Recent research also anticipates multi-Gbps capacity of such picocells [63, 68]. 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 the technology of choice as it can deliver very high data rates over long ranges. However, deploying optical fibers for such dense outdoor picocells can be significantly expensive; thus, a wireless solution is desired. Unfortunately, the existing or upcoming RF¹ based solutions are unlikely to be able to provide the required data rates of 10Gbps or more (see §2.1) at desired ranges (100–200 m and more).

FSO-based Backhaul Solution. Use of free-space optical (FSO) communications can have a strong potential in creating a wireless backhaul for picocells. FSO communication is similar in principle

¹Including Sub 6 and traditional (6–42 GHz) microwave, licensed (70–95 GHz) and unlicensed mmWave (60 GHz).

to (wired) optical fiber communication, except that in FSO, the laser beam travels in free space instead of being enclosed in an optical fiber. In general, FSO links can provide very high data rates—even up to ≈ 100 Gbps—over long distances (\approx kms) [39, 56]. Such tremendous speeds are enabled by high frequencies of lightwaves and absence of regulatory restrictions. However, while the FSO technology itself is not new, its commercial use has been quite limited [15]. This is because significant engineering efforts must be spent in overcoming outdoor effects, especially for long ranges. While short range (≈ 100 m or below) communication can simplify some of these efforts, so far no demanding applications existed to commercialize the short-range use.

Picocells offer a unique setting for FSO-based networks: short-range links arise naturally in densely deployed picocells and robust short-range links with sufficient “link margin” can handle most outdoor effects. To provide sufficient network redundancy with minimal number of FSO interfaces per node, we propose use of *steerable* FSO devices which can be steered with low latency to enable many links and thus create a “dynamic” network. See Fig. 1. Our overall approach (called FSONet) to designing an FSO-based backhaul that can provide desired network availability is based on the following multipronged strategy: (i) Use of steerable FSO links, creating a “dynamic” network with sufficient network redundancy, (ii) Having a sufficient number of robust short-range links that can handle most weather effects, and additional longer links for improved hop-count and network capacity in favorable conditions.

Our Contributions. To implement the above proposed vision, we make the following contributions:

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 (§2).
2. We design and evaluate a 100m link prototype with active tracking and pointing (§3) 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 (§4).
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 (§5).

Finally, we conduct extensive simulations to demonstrate the effectiveness of our techniques (§6).

2 MOTIVATION, CHALLENGES AND PROPOSED ARCHITECTURE

Picocell networks are highly dense, with coverage range of about 100–200 m or less [68]. The nodes (base stations)² are connected to a small number of gateways for external connectivity, through a backhaul network. As mentioned before, a wired backhaul network

can be very expensive to deploy. Thus, a wireless backhaul network is desired. Below, we make a case for an FSO-based backhaul network.

2.1 Case for FSO-based Backhaul for Picocells

Why not a point-to-point network? Creating point-to-point wireless links from each base station to a gateway may be technically plausible via FSO links, but unlikely to be practical for multiple reasons: (i) FSO links require a clear line of sight that is harder to obtain for long links. (ii) Number of transceivers needed at each gateway would be prohibitively high, due to many more nodes than gateways. (iii) The reliability of long-range FSO links is generally not satisfactory due to outdoor effects (discussed later).

Why not an RF-wireless mesh network? The capacity of links in a mesh network must be multiples of the capacity of individual base stations due to traffic aggregation. Given that the individual picocell base stations are expected to provide multi-gigabit capacity through mmWave bands [68], we need backhaul links capable of supporting up to 10 Gbps, perhaps even 40 Gbps, depending on the network topology. No known RF technology can provide such data rates at ranges of ≈ 100 –200 m.³

Why FSO? Free-space optical (FSO) communication [51] links use modulated visible or infrared (IR) laser beams in the free space. Laser beams are very narrow, which eliminates crosstalk interference, and optical spectrum is also unregulated. Thus, FSO links can easily offer Gbps-Tbps of bitrates at long distances (several kms) using relatively low transmit power [39, 56]. However, at long ranges, FSO links can suffer transient failure due to various outdoor effects (discussed below). Fortunately, most of these attenuation problems become manageable at shorter links—in particular, 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 §2.3, we describe an FSO-based architecture based on robust short-range links and other characteristics such that it can provide high network availability.

Why Steerable FSO? For better handling of outdoor effects, we should create a mesh network with sufficient redundancy which requires each node to have more than a few links in the network. This would require many transceivers (perhaps, even ten or more) per node for a network with a few thousand nodes, since the strong directivity of FSO beams entails one transceiver per link. However, equipping a node with more than a few FSO transceivers is likely infeasible due to cost and size constraints. Thus, we propose the use of 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 (see below). Finally, the steering mechanism is also leveraged in designing the tracking and pointing system (§3.2).

Demonstrating Changing Hotspots. There is some evidence in literature that cellular traffic does not peak in many base stations at

²We use the word *node* interchangeably with *base station* throughout the description.

³While THz technology may be able to achieve such data rates, it is relatively immature at this time and has limited outdoor range due to high atmospheric attenuation [29]. Massive MIMO in the mmwave bands can hypothetically deliver 10–40 Gbps at these distances, but the needed signal processing at such frequencies is very challenging.

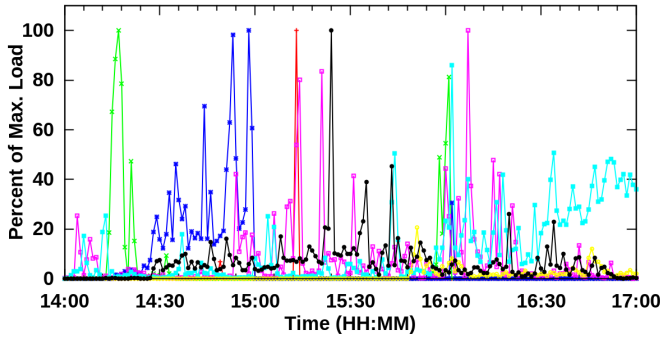


Figure 2: Time series of load seen on different sniffers (shown by different colors).

Cause (rate)	Attenuation (db/km)
Building Movement (without tracking)	40-100 [34]
Dense Fog (40-70 m visibility)	143-250 [33]
Thick Fog (70-250 m visibility)	40-143 [33]
Moderate Fog (250-500 m visibility)	20-40 [33]
Light Fog (500-1000 m visibility)	10-20 [33]
Snow (light to whiteout)	3-30 [33, 67]
Rain (25 to 150 mm/hour)	6-35 [33]
Scintillation (clear day)	5-13 [34]

Figure 3: Link attenuations due to outdoor factors.

exactly the same moment [59, 60]. To understand this effect in WiFi traffic in densely populated areas (which is perhaps most reflective of future picocellular networks), we ran the following measurement study: we deployed 7 time-synchronized sniffers in our building in a dense and active WiFi network of 30+ access points used by 200+ faculty and PhD students and many transient users. Post-processing the captured traces reveals that (i) only 1-2 out of 7 locations carry a significant fraction of the load at any moment, and (ii) these “hotspot” locations are different at different times. See Fig. 2.

2.2 Key FSO Challenge: Outdoor Effects

One concern with use of an FSO-based network in an urban setting is the presence of many high buildings which may limit the number of FSO links as they require a clear line-of-sight; we alleviate this concern by simulating FSONet over many metro US cities (see §6). The other key concern is the attenuation of FSO signal due to various outdoor effects. We discuss these effects below by dividing them into four categories.

- **Building Movements:** Since FSO links require precise alignment of transceivers, minor movements of deployment platforms can result in significant signal attenuation [34], depending on the movement and link design. See Fig. 3. Solutions include active tracking and pointing and/or wider beams [18].
- **Weather Effects:** The signal attenuation from weather effects can be anywhere from 10 to 250 dB/km [33]. Among these,

fog causes the worst attenuating factor, since the size of fog particles is comparable to the optical wavelengths. Note that attenuation from snow is generally less severe than from fog.

- **Turbulence Effects:** These include scintillation, beam wander, and beam spreading. For ranges up to 1km [66], beam spreading is negligible while beam wander is minimal (beam deviation of at most 2-3cm at 1km). Scintillation (optical intensity changes due to temperature variations) happens only when the weather is clear and results in minimal attenuation for links smaller than 500m [34].
- **Blockage:** Finally, since FSO links require a clear line of sight, obstructions such as birds can result in transient failures; these have generally been handled by using multiple beamlets [35].

In the following subsection, we present our approaches to handle each of the above effects.

2.3 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 [68]. 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 [27].* To that end, we design FSONet, a picocell backhaul with the following features:

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 [8] or via two separate optical paths) links. This creates a “dynamic” network.
2. We restrict the maximum range of links to about 500m,⁴ and classify them into two categories: short ($\leq 100m$) and longer (100-500m) links.
 - (a) All links have an optical power *budget*⁵ of about 30dB (to allow for SFP-based prototypes 3).
 - (b) Each short ($\leq 100m$) 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 (§3).
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 (§4).

⁴Links much longer than 500m are unlikely to be useful in a dense picocell network or even feasible in an urban area due to obstructions. In addition, links much longer than 500m can suffer significant scintillation effect.

⁵Optical budget is TX power minus RX sensitivity. If the path loss is higher than the budget, the link may still remain operational albeit at a lower data rate [31]. However, for simplicity, we implicitly assume the link to be non-operational when the path loss exceeds the optical budget.

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 (§5).

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 and 2-5dB for any scintillation effects. The movement tolerance accounted above is needed to support the TP mechanism (see §3.2). We anticipate only minimal (0-2dB) coupling and optical system losses.

Tolerance to Weather Effects. Based on the fog attenuation model in [33], a 100m link with X dB 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 > 85 m (100m) visibility. The visibility data of 20 US cities [52] 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 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 [27].

Commercial FSO Link Solutions. Commercially available FSO links are fixed (i.e., do not steer), are bulky (up to 2 cubic feet [6, 15]) and expensive (\$6K-25K for a device [14]); this is because their objective is to provide highly robust singular links for longer ranges (e.g., multiple kms). They achieve this by using multiple transmitters per link, and bulky TP mechanisms [53]. 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; the combination of these factors allow us to use simpler solutions to handle outdoor effects (e.g., a lightweight TP solution).

3 FSO LINK DESIGN IN FSONET

In this section, we design and develop FSO link prototypes with desired characteristics for FSONet. In particular, we build a 3-node steerable link network, design a lightweight tracking and pointing (TP) mechanism, and develop and evaluate a prototype of a 100m link with the TP mechanism. We start with an introduction to SFPs and SFP-based FSO links.

SFPs. An SFP (small form-factor pluggable) transceiver is a small ($1/2'' \times 1/2'' \times 2''$) and compact commodity optical transceiver [9], widely used to interface optical fibers with network switches. An SFP contains a laser source and a photodetector, for transmission and receiving respectively. In a traditional (wired) optical link, the transmitter SFP modulates the incoming bits onto the laser and

launches the laser beam into the fiber; the beam is captured by the photodetector of the receiver SFP where it is demodulated.

SFPs are available with a variety of laser sources, varying in the wavelength (typically, between 800nm to 1550nm) of the emitted beam as well as the supported data rate (anywhere from 10Mbps to recent variants called CFPs with 100Gbps [1]). SFP+ refers to an enhanced version of the SFP that supports data rates up to 16Gbps. *Bidirectional* SFPs transmit two beams at different wavelengths one in each direction on the same optical fiber. In an SFP-based wireless link, wireless interference is handled by collimation, as described below.

SFP-based FSO Links. We build our FSO link prototypes using SFP transceivers, as in prior works [38, 46, 47]. Our approach of using commodity SFPs is driven by convenience, as it serves our objectives without the need to use independent laser, photodetector, and modulation/demodulation logic. Moreover, SFP-based prototypes demonstrate feasibility of a form-factor manageable and cost-effective FSO link, as well as eye-safety since SFPs are classified as Class 1 safe [13, 26]. SFP+ (i.e., 10Gbps SFPs) with 30dB budget are readily available [2–4].

To create an FSO link using SFPs, the beam emanating from the transmitter SFP is channeled into a short optical fiber which feeds into a collimator. The collimated beam is then launched into free space towards the receiving SFP, where it is captured by another collimating lens and focused back into an optical fiber connected to the receiving SFP. Wireless interference is implicitly handled via two mechanisms: (i) First, a collimated beam is very narrow with minimal divergence, and is thus received only by the intended receiver; (ii) Second, in reconfigurable networks such as ours wherein a beam may be steered to different receivers, appropriate reconfiguration logic must ensure that a receiver receives beam signal from only one transmitter. Note that, in a real system, the received beam can be focused directly to an exposed photodetector (of about $30\mu\text{m}$ diameter for 10Gbps [34]) rather than an optical fiber; we compensate for this discrepancy by using a short $50\mu\text{m}$ multimode fiber at the receiver.⁶

3.1 Steerable FSO Links

We first discuss our choice of steering mechanism.

Galvo Mirrors as Steering Mechanism. We desire a steering mechanism that can steer a beam to an intended target with minimal latency. Several steering solutions [57] exist, but most are not commodity yet. Prior works [46, 47] on reconfigurable FSO networks have explored switchable mirrors, Galvo mirrors (GMs), and digital micromirror devices (DMDs). Switchable mirrors have high latency (10-20 msecs) and DMDs have a very low angular range ($\approx 6^\circ$) unless intricate mirror assemblies are used [46]. On the other hand, GMs [23] have high angular range (up to 40°) with latency (small angle) of a fraction of a msec. Thus, we use GMs as our choice of steering mechanism. GMs also support much higher beam widths (up to 50mm [11]) than DMDs, and have an angular accuracy of $10\mu\text{rad}$ or less. A typical GM essentially consists of one or two mirrors at right angles, each of which can be independently rotated using an electrical signal. A laser beam can be steered using

⁶Intermodal dispersion in the fiber is not an issue [7], as we can use a very short fiber (a few cms).

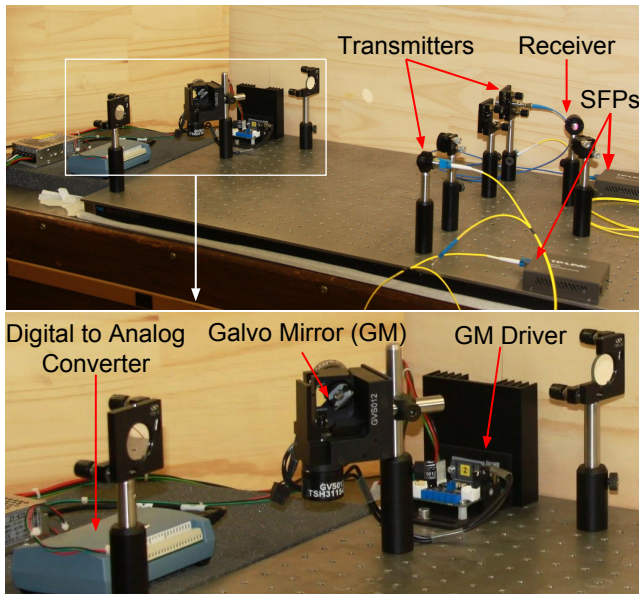


Figure 4: (Top) Three node experimental setup on an optical bench demonstrating FSO steering, (Bottom) Zoomed in to show the Galvo Mirror related set-up.

a GM by reflecting off its mirrors; in effect, a GM with two mirrors can steer a beam towards any target within a predetermined rectangular cone (called *coverage cone*). MEMS [41] are conceptually similar to GMs, with reduced latency at the cost of limited angular range and/or smaller mirrors [17]; small mirrors limit link range (due to small beam-width and higher divergence).

Steerable Link Prototype.

Though steerable FSO links have been proposed before [41, 46, 47], but no prior work has actually built and evaluated a steerable link prototype, to the best of our knowledge. We have recently built a three-node prototype to evaluate GM’s performance in steering an FSO link. See Fig. 4. Each node in the prototype is an FSO end-point using a bidirectional SFP+ connected to a computer. One node acts as a receiver (RX), while the other two act as senders (TX1 and TX2). To evaluate any throughput loss due to steering, we send 16KB UDP packets as fast as possible from TX1 and TX2 to the single receiver RX, while using the GM (GVS012 [23]) to switch the FSO link between TX1-RX and TX2-RX at varying rates (1-50 Hz). *Results:* We observe that the aggregate throughput for a steerable

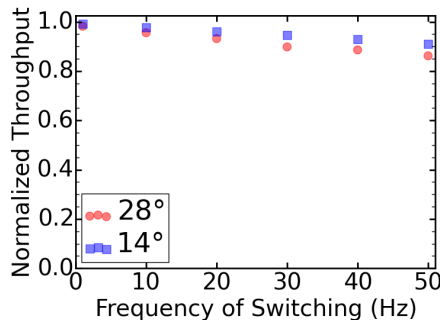


Figure 5: Normalized throughput vs. switching frequency at different steering angles.

link is only 1-2% less (vs. a fixed link) for a steering rate of 1 Hz, and about 8% less at 50 Hz. Large part of throughput loss is due to the recovery of the optical physical/link layers in the SFPs [46], and should not manifest in a real system not based on SFPs. In a non-SFP real system, the throughput loss should be closer to 2% at 50Hz, since our GM’s switching latency is only about $400\mu\text{m}$ [23].

3.2 Active Tracking and Pointing (TP)

To keep an FSO link aligned in an event of building/tower (deployment platform) motion, we need to have a active tracking and pointing (TP) mechanism to *track* the movement of terminal(s) and a *pointing* mechanism to realign the link in real-time. We start with discussing the requirements of such a system in FSONet.

Angular Speed Requirements due to Building Motions. The fundamental reason for link misalignment is the movement of the beam at the receiver plane (due to the movement of the TX and/or RX assemblies). The tracking and pointing (TP) mechanism at the TX corrects any misalignment by moving the TX beam appropriately. Since any correction mechanism incurs a non-zero latency, a TP system can only be effective up to a certain *speed* of RX beam movement. To relate the RX beam movement to building motions that cause the beam movement, we map the movement speed of the RX beam to the most dominant⁷ building motion in our context, viz., the angular movement speed of the TX assembly. Accordingly, we measure effectiveness of the TP mechanism in terms of the angular speed (i.e., rotation rate) of TX assembly (or building motion) that the TP mechanism is able to handle.

Prior studies suggest that building motions rarely have an angular speed of more than a few mrad/sec. In particular, based on the motion frequencies, building motions can be roughly divided into three categories [35]: (i) Vibrations can have a high frequency (1-10Hz), but rarely exceed 0.5mrad [35] in amplitude; (ii) Moderate frequency (fraction of a Hz) motions due to wind, etc. can have an amplitude of up to a few milliradians [34, 58]; (iii) Motion due to thermal variations have very low frequency (periods from minutes to months). Maximum angular speed (frequency times amplitude) across the above categories is at most a few mrad/sec. Below, we demonstrate that our link prototype is able to handle angular speeds of up to 15 mrad/sec, which is more than sufficient to handle any expected building motions.

Link Tolerance Requirements. To ensure link’s operation at all times, the link should have an “angular tolerance” of more than the building motion that is possible within the TP system’s latency; here, the link’s *angular tolerance* is defined as the maximum angular motion of TX that can be tolerated (*absence of a TP system*) without losing the link. Due to our choice of GM as the pointing mechanism, we expect the TP latency to be of the order of few msec. Thus, to handle building motion speeds of say 5 mrad/sec using a TP system with latency (to correct the link “sufficiently”) of say 5 msec, we need the link’s angular tolerance to be at least 0.025 mrad. To be conservative and to account for other factors (e.g., noise), we aim for link’s tolerance of 0.05–0.1 mrad.

⁷RX beam movement due to angular movement at TX depends on the link range, and thus can be significant (few 10s of cms) at 100m range. In contrast, RX beam movement due to linear movement of RX or TX does not depend on the link range, and is expected to be much less. Rotation of the RX assembly causes only slight misalignment.

TP based on GMs and Photodiodes. Typical TP mechanisms [48, 50] include a fast steering mirror or gimbal controlled by digital servos [54, 64, 69] for pointing, along with some tracking detectors to track the target or beam. Common tracking detectors include positioning sensing diodes [64] (e.g., CCDs [54], photodiode arrays [42]), accelerometers, cameras, GPS [69], etc.⁸ In our context, it is natural to leverage our steering mechanism GMs as also the pointing mechanism. The pointing accuracy of GMs ($10\mu\text{rad}$) is much less than $1/8$ of the beam divergence (about $280\mu\text{rad}$) of our collimator [10], and thus should guarantee a less than 0.5dB tracking error [54]. For beam tracking, we use photodiodes, due to relatively short range ($< 500\text{m}$) of our links which entails minimal intensity fluctuations in outdoors [34, 66].

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 use the following control algorithm: First, we build a table that contains four-tuple diode readings for different 2D deviations of the TX from the initial (aligned) position; for a 100m link, 10k table entries sufficed. 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. In our prototype, we also added an integral term [22] (with a 0.1 coefficient) to the correction. The above control strategy can stabilize the link in minimal number of iterations and is fast enough, as discussed in §3.3. 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.

3.3 Robust 100m Link Prototype with TP

We now describe and evaluate our SFP-based 100m link prototype with the above described TP mechanism. The link prototype satisfies our stated requirements of sufficient link margin to handle weather effects, and the TP mechanism is effective in handling expected building motions. We also analyze longer links.

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 §2.2, atmospheric effects are negligible at these ranges. We address weather effects by demonstrating a sufficient link margin; as mentioned in §2.3, a 100m link with 13-15dB margin can withstand most weather effects. We use the 10G 1550nm SFP+ [4] with a 10GBASE-ZR interface⁹, a transmission power of 0-4dBm and a receiver sensitivity of -25dBm yielding a total optical budget of 25-29dB. We use a uni-directional FSO link for simplicity, as it is sufficient for our objectives, with the other direction of the SFPs connected by an

⁸For very short links, simpler techniques can be used, e.g., [41] 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.

⁹The physical layer standard used by 10GBASE-ZR interfaces is proprietary and not specified under IEEE 802.3ae.

optical fiber. We use a GM at TX, and surround the RX’s collimator by four photodiodes. See Figure 6(a)-(c). For photodiodes, we use Thorlabs’ InGaAs variable-gain photodiodes [21] 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 [5] to minimize errors from stray light. The diodes are read via a DAQ [19] which can sample at up to 250k/sec. 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 [10] 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 [16] to enable fine angular movement of the TX assembly at varying speeds, to simulate building motions.

Link Margin and Tolerance. At best alignment, the link’s received power was around -12dBm, yielding a link margin of 13dB. We note that most of this path loss of 12-16dB in our prototype is alignment/pointing loss from manual alignment at the RX terminal, and can be reduced significantly with a more sophisticated alignment mechanism. In Fig. 7(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 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. Here, we use a USB-interfaced power meter [24] which logs power received every 3-5 msec.

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 7(b). Here, each run is over a 10 minute period, with throughput measured over every second. We observe that the throughput CDFs for speeds ≤ 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 [24] 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. The link usually regains operational alignment within a single iteration.

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 Zemax [25] based on given optical elements and estimate link’s signal loss (assuming perfect alignment) at various ranges. Zemax is an optical design software platform

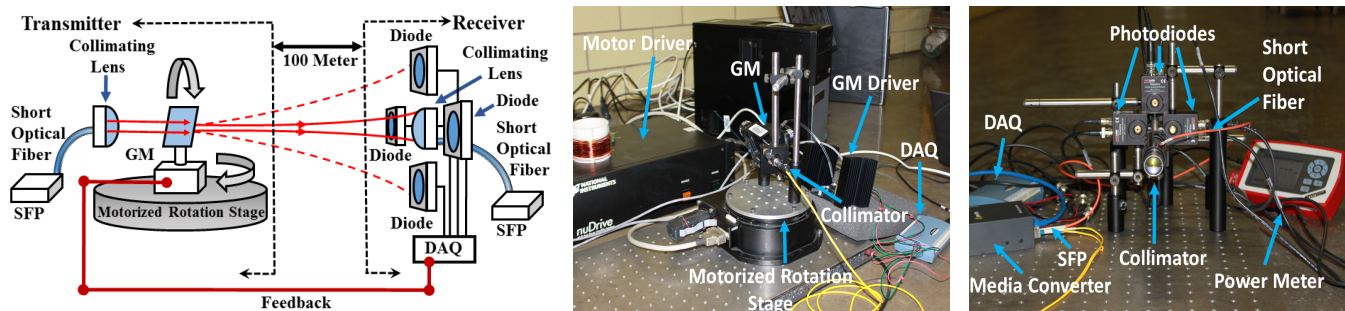


Figure 6: (a) Schematic diagram of the 100m link prototype, (b) Transmitter assembly with GM and motorized rotational stage, and (c) Receiver assembly with 4 photodiodes around the collimator.




Figure 7: (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.

widely used by optical engineers to analyze, simulate and optimize optical systems. Recall that, for links longer than 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.7(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. 7(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. Note that angular tolerance can be further improved by creating wider beams at RX [18], via a custom optical design, at the cost of slight power loss at perfect alignment. E.g., increasing the collimator divergence to 0.78 mrad improves the angular tolerance of 500m link to 0.2 mrad, while decreasing the power at perfect alignment by 1.5dB at 100m and 2dB at 500m.

4 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, given a geographical area to which we wish to provide picocell coverage to. Creation of potential FSO links in turn entails orientation of GMs associated with the FSO devices. We start with borrowing a few definitions from [47].

Candidate Links, Dynamic Network, Realizable Topologies. In our context, 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. This 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. 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. 8(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 can 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 §2.1, 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:

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

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. The above objective is somewhat related but quite different from the *dynamic bisection bandwidth* (DBW) objective [47] for dynamic networks in data centers; DBW is not a suitable metric in FSONet due to two different node categories, viz., base stations and gateways, and traffic flows only between basestations and gateways.

Intractability, and Proposed Heuristic (FSH). It is easy to see that the DBND problem is NP-hard via a straightforward reduction from the set cover or Steiner tree problem. The closest known problem is the *Group Steiner Tree* (GST) problem [12], which, in our context, represents the “partial” problem of constructing the smallest backbone without the node degree constraint. The known approximation algorithms [12, 45] for GST, though polynomial, are very inefficient [12] 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. Even an Integer Linear Program (ILP) requires an exponential number of equations for the above objective function. Even in practice, a “good” solution can be challenging to achieve due to a limited number of potential node locations (e.g., building corners, roof tops, towers) and links (limited line-of-sight). We propose the following multi-step heuristic (referred as the *Four-Step Heuristic (FSH)*) that delivers good quality network solutions in real settings as shown in §6. The basic idea (see Fig. 8) 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 “Steiner-forest” problem is a slight variation of the well-known Steiner-tree problem [44]. 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 [43]), 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.

Steps Details. For the first step, we partition the given area into “atomic” area-elements such that each area-element is either fully covered or fully uncovered by a node. Then, we greedily select nodes in the order of the number of yet-uncovered area-elements they cover.

In the second step, we construct a “Steiner forest” over the gateways and the nodes n_c selected in the first step. Similar to the standard 2-approx. Steiner tree algorithm [44], we start with the forest of just the gateway nodes and iteratively grow the forest by connecting a not-yet-connected node i of n_c to the current forest. To choose i , we pick the node that has the shortest path p_i of short links to the forest.

In the third step, to minimize the maximum node degree, we iteratively search for a triplet of nodes (u, v, w) such that u, v, w are nodes in the forest, (u, v) is a short-link not in the forest, degree of w is at least 2 more than the maximum degree of u and v , and w is either on the cycle created by adding (u, v) to the forest or w is an ancestor of u or v . See Fig. 9. On finding such a triplet, we replace a link of w by (u, v) — this maintains the backbone properties but reduces a node degree. We iterate until no such triplet exists. If the above still leaves certain nodes with degree higher than the node degree constraint δ , then we apply the following reconstruction step. We remove links to bound the maximum node degree by δ , connect the disconnected components to the gateways using minimal *additional* nodes, and redo the third step if needed (i.e., if a node degree is higher than δ). After sufficiently many iterations, if node-degree constraint is still not satisfied, we either reject the instance or move on to the next step with the backbone found till now with least violations.

Backbone. At the end of the third step, we have a backbone with minimal number of nodes and bounded node degree. For each node in the backbone: we equip it with d FSOs, where d is the node’s degree, and associate one FSO per link. This essentially creates a dynamic network with the backbone as its realizable topology.

Fourth Step. We now augment the dynamic network with additional nodes and links, within the given constraints, to improve its objective value. To begin with, we add any possible candidate links *between* the backbone nodes — associating these new links with already existing FSOs at the nodes. Next, we iteratively add a node or a path of nodes to improve the AHF objective. Since the AHF objective considers all hotspots uniformly, we improve the objective simply by treating the entire dynamic network as a traditional (static) network and improving the max-flow between nodes to gateways as follows. We connect all nodes to a super-source and all gateways to a super-sink, compute the max-flow, and consider the final residual network \hat{R} . Note that, in \hat{R} , no node has a path to any gateway. We then connect a node to a gateway in \hat{R} (and thus \mathcal{D}) by adding a new path. More specifically, we add the path that yields the most increase in flow per added node. We iterate on the above, until all the allowed number of nodes have been added.




Figure 8: 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).




Figure 9: Transformations used to minimize the maximum node degree in FSH’s 3rd step.

5 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) [30, 36, 40] in designing the network management layer. 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 [40] via a control channel. The key component of the controller is the *reconfiguration engine* as described below. We preserve network connectivity and consistency during reconfigurations [47, 61] via techniques from [47] 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 an arrival of a new high-demand flow. Infrequently, we also reconfigure the *entire* network periodically based on global information. We model a *flow*

as arriving at a certain node with a *bandwidth-demand*. Bandwidth-demand may be application dependent and/or may reflect the SLAs associated with the flow type. We assume it can be estimated by analyzing very early part of the flow; accurate estimation will benefit performance, but is not critical in our context. An effective reconfiguration strategy should be infrequent, fast (run in a few tens of milliseconds), and incremental (incur minimal topology and routing-table changes so that the impact to the existing traffic is minimal). We allow multi-path routing. The *triggered* reconfiguration scheme from [47], for dynamic networks in data centers, is simple and fast, but effective only in very dense networks. Likewise, the matching-based scheme of [46], though decentralized, is suited for complete networks. In FSONet, the network is much sparser and all flows are to/from a small number of gateway nodes; thus, we instead employ the following scheme which is simple/fast, somewhat localized, and done only when “beneficial.”

- If the bandwidth-demand is below a certain threshold (e.g., 100Mbps, in our simulations) *or* the current network configuration (i.e., runtime topology and the generic rules in the routing tables) is able to satisfy 90% of the bandwidth demand of the flow F , then we route F using the current generic rules.
- Else, we compute the multi-path route r (using max-flow) for the flow F , without changing the existing flow routes or topology. If r results in 90% of the bandwidth-demand of F being satisfied, then we add r as the flow-route for F in the routing tables.
- Else, we reconfigure FSONet by changing its runtime topology as well some of the existing flow-routes as follows. First, we find the shortest-path P in the dynamic network such that (a) P involves at least one inactive link, (b) Activation of links on path P entails neither deactivation of any “heavily-used” links¹⁰ nor disconnection of any node from the gateways. Then, we activate/deactivate links to add P to the current runtime topology, and route affected traffic and F appropriately on the new runtime topology. If the above still doesn’t satisfy 90% of F ’s bandwidth-demand, then we don’t reconfigure at all and just route F as best as possible in the current configuration.

To minimize the computation time, we can precompute the reconfiguration for all possible parameter (e.g., F ’s arrival node and bandwidth-demand) values for a given network. For periodic reconfiguration based on global information, we adopt a greedy approach

¹⁰Note that activation of a link may require deactivation of other links, since each FSO can only have at most one active link incident on it.

of iteratively activating a path of links from a node to a gateway, in order of nodes’ bandwidth demands.

Link/Network Outages. Link outages can be handled based on their duration: (i) retransmit dropped packets for very short outages (e.g., due to bird blockage), (ii) reroute flows for longer outages, (iii) reconfigure network (as above, since a link failure can be modelled as a flow arrival) for a semi-permanent link failure (e.g., bad weather). Finally, in case of extreme conditions where even the backbone fails, the traffic can be routed to the back-up macrocellular network.

6 EVALUATION

Having already established the performance of link prototype in §3, we now evaluate the “quality” of the FSONet networks created for many metropolitan US cities.

US Cities Data. In our simulation study, we create FSONet networks to provide *outdoor* coverage in 14 of the top US metro cities. We use topographical data from [20] which contains locations and 2D footprints of the buildings in the cities. We use building corners (at a uniform height) as the *potential locations* for picocell base stations. For determination of line-of-sight between two locations, we conservatively assume each building to be high-enough to be an obstruction.

Parameter Values. We vary the (ground) coverage radius of a node from 100 to 200 meters, with 100 m as the default value. We allow 3 to 6 FSOs per base station, with 4 as the default; note that only a low number is practical. We use $\approx 5\%$ of the expected size of FSONet network as the number of gateways, and locate them randomly over the given locations. 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.

FSONet Network Quality. Given the above set-up, we create FSONet networks in 14 of the top US 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.

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 50 to 750 nodes across the cities, for default coverage radius of 100m. See Table 1. Delving in deeper, we observe that the n_c , the number of nodes selected in FSH’s first step to provide coverage, is only about 2 to 2.8 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 2 to 3.5; this essentially represents the cost of enforcing the backbone condition (the key aspect of FSONet). Note that the n_b to n_c ratio is expected to be at least 2, as the backbone links can be at most 100m.

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

City	Backbone (n_b)	Coverage Factor (n_c /lower-bound)	Conn. Factor (n_b/n_c)	AHF Ratios	
				FSONet/Static	FSONet/upper-bound
NYC	752	2.38	3.434	2.53	0.63
Chicago	626	2.346	3.421	2.16	0.60
Seattle	533	2.278	3.25	1.99	0.62
SFO	503	2.284	3.288	1.82	0.62
DC	404	2.364	3.108	1.74	0.63
Boston	321	2.319	2.945	1.76	0.70
Portland	305	2.14	2.85	1.73	0.66
Baltimore	269	2.196	2.663	1.74	0.65
Philly	275	2.245	2.311	1.75	0.65
Houston	220	2.116	2.418	1.68	0.70
Denver	145	2.067	2.339	1.69	0.72
Detroit	139	2.179	2.279	1.62	0.73
Dallas	62	2.8	2.214	1.61	0.74
Miami	58	2.5	1.933	1.52	0.72

Table 1: Statistics for FSONet networks for 14 US cities, with default parameter values.

steering FSO capability) network as well 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 1. We make two observations about the AHF estimate of FSONet networks: (i) it is about 1.5 to 2.4 times the AHF estimate of the corresponding static graph—confirming the benefit of steerability/dynamicity in a network, (ii) it is about 0.5 to 0.7 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 2 shows range of statistics across cities for other values of coverage radius and # of FSOs/node, and we observe similar statistics except that the n_b/n_c ratio increases with the coverage radius, as expected.

Online Performance. To evaluate the online performance of the created FSONet networks (for default radius 100m and 4 FSOs/node), we run traffic simulations over the networks and measure appropriate performance metrics. In particular, we consider the traffic model wherein arrival of flows in the network is modelled 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 modelled 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 bandwidth-demand to a flow 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

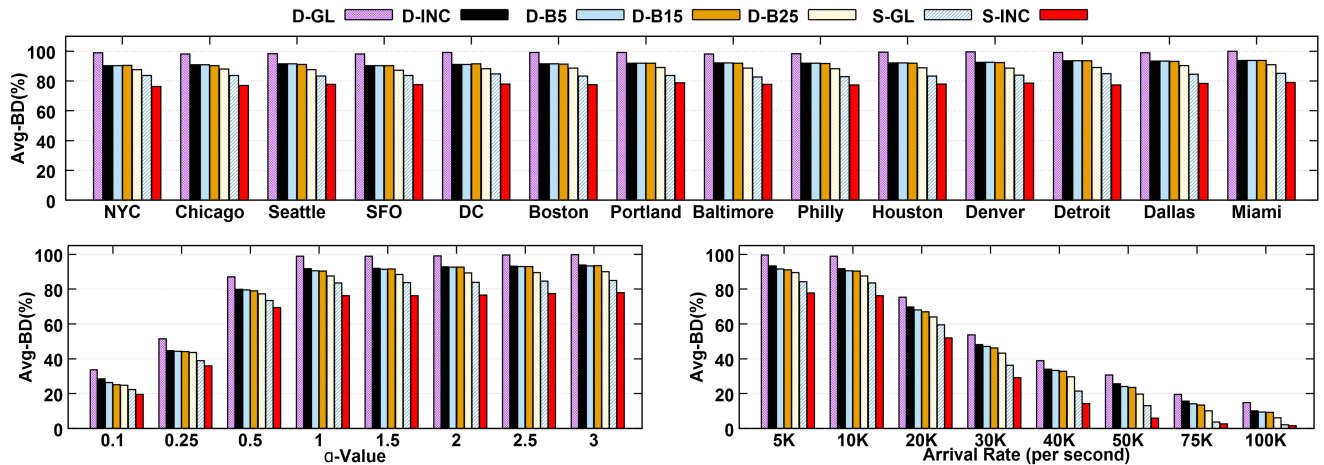


Figure 10: Performance (Avg-BD value) of various reconfiguration schemes in FSONet and its corresponding static picocell networks. (a) In FSONet networks of various cities for $\alpha = 1$ and flow arrival rate of 10,000/sec, (b) In Manhattan’s FSONet network for varying α with flow arrival rate of 10,000/sec. (c) In Manhattan’s FSONet network for varying flow arrival rates with $\alpha = 1$.

(Coverage, FSO/node)	Backbone (n_b)	Coverage Factor (n_c /lower-bound)	Conn. Factor (n_b/n_c)	AHF Ratios	
				FSONet/Static	FSONet/upper-bound
(100m, 3)	58-752	2.08-2.80	1.93-3.43	1.58-2.68	0.64-0.76
(100m, 4)	58-752	2.08-2.80	1.93-3.43	1.52-2.53	0.6-0.74
(100m, 5)	58-752	2.08-2.80	1.93-3.43	1.50-2.42	0.58-0.73
(100m, 6)	58-752	2.08-2.80	1.93-3.43	1.47-2.27	0.56-0.72
(150m, 3)	42-632	2.43-3.00	3.00-5.80	1.54-2.61	0.62-0.78
(150m, 4)	42-632	2.43-3.00	3.00-5.80	1.51-2.47	0.61-0.77
(150m, 5)	42-632	2.43-3.00	3.00-5.80	1.58-2.38	0.59-0.75
(150m, 6)	42-632	2.43-3.00	3.00-5.80	1.45-2.24	0.58-0.72
(200m, 3)	36-584	2.46-3.00	4.50-8.99	1.55-2.56	0.63-0.78
(200m, 4)	36-584	2.46-3.00	4.50-8.99	1.50-2.39	0.62-0.76
(200m, 5)	36-584	2.46-3.00	4.50-8.99	1.56-2.31	0.59-0.75
(200m, 6)	36-584	2.46-3.00	4.50-8.99	1.44-2.19	0.58-0.73

Table 2: Range of values of different statistics across FSONet networks of 14 cities, for different values of picocell coverage radius and # of FSOs per node.

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.

We consider the following reconfiguration schemes: (i) INC: the (incremental) approach from §5 executed for every large flow (i.e., > 100Mbps bandwidth-demand), (ii) GL: the global periodic-reconfiguration of §5 executed at the arrival and completion of every large flow, (iii) Bn: the slightly-modified INC approach executed for every “batch” of arrival of n large flows. Note that the GL approach is impractical

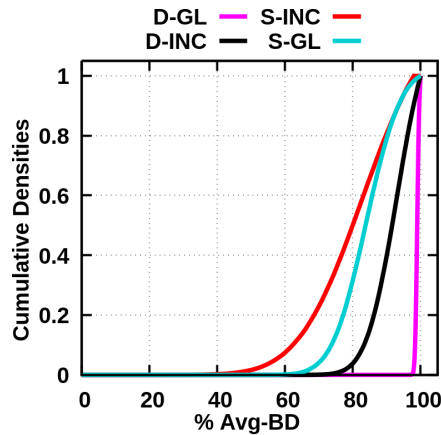


Figure 11: CDF plot for the Avg-BD metric in NYC FSONet.

due to high computation time and global changes to the network but serves as an apt comparison benchmark (upper-bound). We prefix the above acronyms with D- or S- for FSONet and static networks (as defined above) respectively.

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 average it over time. We use Avg-BD to refer to this average metric, and use it for comparison. See Fig. 10(a). Key observations: (i) D-INC performs quite close to D-GL (91-94% vs. 98-99%), (ii) D-INC (black bars in figure) performs significantly better than S-INC (red bars), i.e., 91-94% vs. 76-78%, and (iii) the D-Bn approaches perform almost same as the D-INC approach for $n = 5$ and 15, but worse for larger n . Since S-INC outperforms S-Bn approaches, they are not shown in Fig. 10 for clarity. Figure 10(b)-(c) plot the Avg-BD metric for varying α parameter of the Pareto

distribution and flow-arrival rate, for the NYC network. As expected, the Avg-BD values decrease due to increase in the arrival rate and decrease in the α (since smaller α makes larger flow more likely)—but the relative performances of the approaches remain similar to before. We also plot the CDF of the Avg-BD metric for the NYC FSONet for $\alpha = 1$ and flow arrival rate of 10,000/sec. See Fig. 11; for clarity, we show only the main approaches viz., D-GL, D-INC, S-GL, and S-INC. We observe that the relative performance of the schemes is similar to that observed in Fig. 10(a), and that the CDFs show minimal tail effect.

Updates, Time. In Fig. 12, we compare the number of routing table entries changed per reconfiguration across schemes for NYC network. Less changes results in less network-update time, and facilitates concurrent reconfigurations [47]. We observe that the D-Bn approaches incur less changes than the D-INC approach,

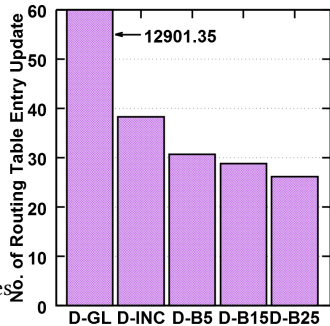


Figure 12: Number of routing table changes per reconfiguration, for various schemes.

which incurs only about 0.25% of the changes incurred by GL (the GL approach needs to change *all* the entries). Note that the number of reconfigurations in a D-Bn approach is $(1/n)^{th}$ of that in D-INC. Thus, overall, the D-B15 approach seems to be most efficient since it performs very close to D-INC, but incur much less reconfigurations routing table changes. Finally, we note that the computation times of D-INC and D-Bn approaches were of the order of 0.1 msec while that of the GL approach was 10-20 msec.

Network Redundancy. To evaluate and compare network redundancy for FSONet and corresponding static networks: We delete a number of links randomly from the static network and delete the *same* set of links from FSONet, and measure the above defined Avg-BD metric for each, and compare the “performance deterioration.” We define the *performance deterioration* as $(a - b)/b$ where b and a are the Avg-BD values of the network before and after deletion of links. See Fig. 13. We observe that the NYC FSONet network exhibit much less deterioration in performance than its static network.

7 RELATED WORK

Backhaul of picocell networks is a key challenge, and recent works have explored various solutions. Among wireless backhaul solutions, [49] addresses resource (power and scheduling) allocation in backhaul solutions with shared wireless channels, [62] presents beamforming techniques to overcome outdoor challenges in an mm wave based solution, and [55] 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

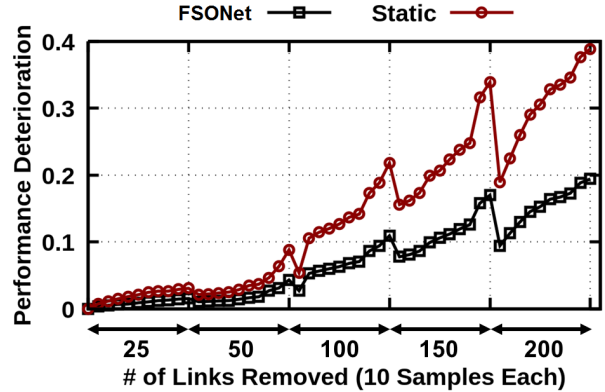


Figure 13: Performance deterioration of NYC’s FSONet and static networks for varying number of link failures.

1-2Gbps over 100-200m). Recently, free space optics technology (FSO) has emerged as an attractive alternative. E.g., [31] 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 [28, 70] or wired/FSO [70] 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.

Our work is inspired by recent works on use of steerable FSO links for dynamic data center networks [46, 47]. 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 [47], a simple random network worked well as the dynamic network [47] as the racks were treated uniformly and the line of sight issue was circumvented by placing a full ceiling mirror. Similarly, in [46], 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 §5.

8 CONCLUSIONS

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 backhuls. 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.

ACKNOWLEDGMENTS

This work was partially supported by NSF grant 1514017.

REFERENCES

- [1] 100G CFP Optical Transceivers. <https://edgeoptics.com/cfp-optical-transceiver-the-basics/>.
- [2] 10G BiDi SFP+ Transceiver. <https://tinyurl.com/zqwgq9>.
- [3] 10G BiDi XFP 100km Transceiver. https://www.alibaba.com/product-detail/XFP-module-10G-Tx-1490-Rx_60536576254.html.
- [4] 10G SFP+ 1550nm 100km DOM Transceivers. <http://www.fs.com/products/29799.html>.
- [5] 1550nm Bandpass Filter. <https://www.thorlabs.com/thorproduct.cfm?partnumber=FB1550-12>.
- [6] Artolink 10G FSO Link Specs. http://artolink.com/page/products/free_space_optics_Artolink_10Gbps/.
- [7] BER and Eye-Diagram of 10Gbps signal over multi-mode fiber. <https://www.youtube.com/watch?v=Ab2h1oqZ-GA>.
- [8] Bi-directional 10G SFP+. <http://www.fs.com/products/31028.html>.
- [9] Cisco 10GBASE SFP+ modules. http://www.cisco.com/c/en/us/products/collateral/interfaces-modules/transceiver-modules/data_sheet_c78-455693.pdf.
- [10] F810FC-1550 Collimation Package from Thorlabs. <https://www.thorlabs.com/thorproduct.cfm?partnumber=F810FC-1550>.
- [11] GMs from Cambridge Technology. <http://www.cambridgetechnology.com/products/galvanometer-xy-sets>.
- [12] Group Steiner Tree Problem. <http://theory.cs.uni-bonn.de/info5/steinerkompndium/node30.html>.
- [13] Laser safety information. https://www.lia.org/subscriptions/safety_bulletin/laser_safety_information.
- [14] LightPointe FSO Link Cost. <http://tinyurl.com/k86o2vh>.
- [15] Lightpointe FSO Links. <http://www.lightpointe.com>.
- [16] Micos. http://www.micosusa.com/product/subview.cfm_firstlevel=1&sublevel=17.htm.
- [17] Mirrorcle technologies MEMS mirrors: Technical overview. <http://www.mirrorcletech.com/pdf/Mirrorcle%20Technologies%20MEMS%20Mirrors%20-%20Technical%20Overview.pdf>.
- [18] Mounting stability in fso deployments. http://www.fsona.com/tech/app_notes/fSONA-APPNOTE-Mounting_Stability.pdf.
- [19] Multifunction USB Data Acquisition Device. <http://www.mccdaq.com/usb-data-acquisition/USB-1608G.aspx>.
- [20] Openstreetmap foundation. <http://www.openstreetmap.org>.
- [21] Photodetector. <https://www.thorlabs.com/thorproduct.cfm?partnumber=PDA20CS>.
- [22] PID Controller. https://en.wikipedia.org/wiki/PID_controller.
- [23] Thorlabs galvo mirrors. <http://tinyurl.com/haymezj>.
- [24] Thorlabs PM100D Power Meter. <https://www.thorlabs.com/thorproduct.cfm?partnumber=PM100D>.
- [25] Zemax. <https://www.zemax.com>.
- [26] *Safety of Laser Products – Part 1: Equipment classification, requirements and users guide, International Electrotechnical Commission (IEC) Standard 60825-1*. 2007.
- [27] Small cell backhaul requirements. <http://tinyurl.com/h8fomey>, 2012. NGMN Alliance, White paper.
- [28] Ahmed Douik et al. Hybrid radio/free-space optical design for next generation backhaul systems. *CoRR*, 2015.
- [29] I. Akyildiz et al. Terahertz band: Next frontier for wireless communications. *Physical Communication*, 2014.
- [30] M. Al-Fares et al. Hedera: Dynamic flow scheduling for data center networks. In *NSDI*, 2010.
- [31] M. Alzenad, M. Z. Shaker, H. Yanikomeroglu, and M. Alouini. FSO-based vertical backhaul/fronthaul framework for 5G+ wireless networks. *CoRR*, abs/1607.01472, 2016.
- [32] J. Andrews et al. What will 5G be? *IEEE JSAC*, 2014.
- [33] M. Awan et al. Characterization of fog and snow attenuations for free-space optical propagation. *Journal of Communications*, 4(8), 2009.
- [34] S. Bloom. The physics of free space optics. <http://tinyurl.com/zzy3l5k>, 2002. White paper.
- [35] S. Bloom, E. Korevaar, J. Schuster, and H. Willebrand. Understanding the performance of free-space optics [invited]. *Journal of optical Networking*, 2(6):178–200, 2003.
- [36] M. Casado et al. Ethane: Taking control of the enterprise. In *Proc. SIGCOMM*, 2007.
- [37] K. Chandra et al. CogCell: Cognitive Interplay between 60GHz Picocells and 2.4/5GHz Hotspots in the 5G Era. *CoRR*, 2015.
- [38] M. Chowdhury, M. Kavehrad, W. Zhang, and P. Deng. Combined CATV and very-high-speed data transmission over a 1550-nm wavelength indoor optical wireless link. In *Proc. of SPIE*, 2014.
- [39] E. Ciaramella et al. 1.28-Tb/s free-space optical WDM transmission system. *IEEE Photonics Tech. Letters*, 2009.
- [40] A. R. Curtis, J. C. Mogul, J. Tourrilhes, P. Yalagandula, P. Sharma, and S. Banerjee. Devoflow: Scaling flow management for high-performance networks. In *SIGCOMM*, 2011.
- [41] P. Deng, M. Kavehrad, and Y. Lou. MEMS-based beam steerable free space optical communication link for reconfigurable wireless data center. In *Proc. of SPIE Vol*, volume 10128, pages 1012805–1, 2017.
- [42] M. S. Ferraro, W. R. Clark, W. S. Rabinovich, R. Mahon, J. L. Murphy, P. G. Goetz, L. M. Thomas, H. R. Burris, C. I. Moore, W. D. Waters, K. Vaccaro, and B. D. Krejca. Inalás/ingaas avalanche photodiode arrays for free space optical communication. *Appl. Opt.*, 54(31), Nov 2015.
- [43] M. Fürer and B. Raghavachari. Approximating the minimum-degree steiner tree to within one of optimal. *J. Algorithms*, 17(3), 1994.
- [44] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. 1979.
- [45] N. Garg, G. Konjevod, and R. Ravi. A polylogarithmic approximation algorithm for the group steiner tree problem. In *Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '98, 1998.
- [46] M. Ghobadi, R. Mahajan, A. Phanishayee, N. Devanur, J. Kulkarni, G. Ranade, P.-A. Blanche, H. Rastegarfar, M. Glick, and D. Kilper. Projector: Agile reconfigurable data center interconnect. In *Proceedings of the 2016 Conference on ACM SIGCOMM 2016 Conference*, SIGCOMM '16, 2016.
- [47] N. Hamedazimi et al. FireFly: A reconfigurable wireless data center fabric using free-space optics. In *SIGCOMM*, 2014.
- [48] T.-H. Ho. *Pointing, Acquisition, and Tracking Systems for Free-Space Optical Communication Links*. PhD thesis, University of Maryland, College Park, 2007.
- [49] I. Maric et al. Resource allocation for constrained backhaul in picocell networks. In *Info. Theory and App. Workshop*, 2011.
- [50] S. V. Kartalopoulos. *Free Space Optical Networks for Ultra-Broad Band Services*. John Wiley and Sons, 2001.
- [51] D. Kedar and S. Arnon. Urban optical wireless communication networks: The main challenges and possible solutions. *IEEE Communications Magazine*, 2004.
- [52] I. I. Kim and E. J. Korevaar. Availability of free-space optics and hybrid FSO/RF systems. In *Proc. SPIE*, 2001.
- [53] S. G. Lambert and W. L. Casey. *Laser communication in space*. Artech House, 1995.
- [54] C. Lv, H. Jiang, and S. Tong. Implementation of FTA with high bandwidth and tracking accuracy in FSO. In *2012 2nd International Conference on Consumer Electronics, Communications and Networks (CECNet)*, 2012.
- [55] M. Coldrey et al. Non-line-of-sight microwave backhaul in heterogeneous networks. In *IEEE VTC*, 2013.
- [56] A. K. Majumdar. *Advanced Free Space Optics (FSO): A Systems Approach*. Springer New York, 2015.
- [57] P. F. McManamon et al. A review of phased array steering for narrow-band electrooptical systems. *Proc. of the IEEE*, 2009.
- [58] A. G. A. Muthalif, K. K. Turahim, and S. A. Rahim. Active vibration isolation system to improve free space optics communication. In *Proceedings of the 2012 International Conference on Information Technology & Computing Intelligence*, pages 369–378. Springer, 2013.
- [59] A. Nika et al. Understanding and predicting data hotspots in cellular networks. *Mobile Networks and Applications*, 2015.
- [60] U. Paul et al. Understanding traffic dynamics in cellular data networks. In *INFOCOM*, 2011.
- [61] M. Reitblatt, N. Foster, J. Rexford, C. Schlesinger, and D. Walker. Abstractions for Network Update. In *Proc. SIGCOMM*, 2012.
- [62] S. Hur et al. Millimeter wave beamforming for wireless backhaul and access in small cell networks. *IEEE Trans. on Comm.*, 2013.
- [63] S. Saha et al. 60 GHz Multi-Gigabit Indoor WLANs: Dream or Reality?
- [64] M. K. A.-A. H. R. J. J. Sluss. A tracking system for mobile FSO. In *SPIE Proceedings Vol. 6877*, 2009.
- [65] T. S. Rappaport et al. Millimeter wave mobile communications for 5G cellular: It will work! *IEEE Access*, 2013.
- [66] H. Weichel. *Laser Beam Propagation in the Atmosphere*. SPIE Press, 1990. Chapter 5.
- [67] H. Willebrand and B. Ghuman. *Free Space Optics: Enabling Optical Connectivity in Today's Networks*. Sams, 2001.
- [68] Y. Zhu et al. Demystifying 60GHz outdoor picocells. In *MobiCom*, 2014.
- [69] T. Yamashita, M. Morita, M. Shimizu, D. Eto, K. Shiratama, and S. Murata. The new tracking control system for free-space optical communications. In *2011 International Conference on Space Optical Systems and Applications (ICSOS)*, 2011.
- [70] Yuan Li et al. Optimization of free space optical wireless network for cellular backhauling. *IEEE JSAC*, 2015.