

Cyclops: An FSO-based Wireless Link for VR Headsets

Himanshu Gupta
Stony Brook University
Stony Brook, NY, USA

Max Curran
Stony Brook University
Stony Brook, NY, USA

Jon Longtin
Stony Brook University
Stony Brook, NY, USA

Torin Rockwell
Stony Brook University
Stony Brook, NY, USA

Kai Zheng
Stony Brook University
Stony Brook, NY, USA

Mallesham Dasari
Stony Brook University
Stony Brook, NY, USA

ABSTRACT

The ultimate goal of virtual reality (VR) is to create an experience indistinguishable from actual reality. To provide such a "life-like" experience, (i) the VR headset (VRH) should be *wireless* so that the user can move around freely, and (ii) the wireless link, connecting the VRH to a high-performance renderer, should support high data rates (tens to hundreds of Gbps). Industry is already pushing towards such wireless VRHs; however, these wireless links can only support a few Gbps rates. In general, current radio-frequency (RF) links (including mmWave) are not able to provide desired data rates. In this paper, we build a system, we call Cyclops, which uses free-space optical (FSO) technology to create a high-bandwidth VR wireless link. FSO links are capable of very high data rates (up to Tbps) due to the high frequencies of light waves and narrow beams. The main challenges in developing an effective FSO link are: (i) designing a link with sufficient movement tolerance, and (ii) developing a viable tracking and pointing (TP) mechanism which maintains the link while the VRH moves. As traditional TP approaches seem infeasible in our context, we develop a novel TP approach based on learning techniques, leveraging the VRH's in-built tracking system. We build robust 10 Gbps and 25 Gbps link prototypes from commodity components, demonstrate their viability for expected movement speeds of a VRH, and show that, with certain custom-built components, we can support much higher movement speeds and bandwidths.

CCS CONCEPTS

• **Networks** → **Physical links; Home networks; Traffic engineering algorithms; Network control algorithms.**

KEYWORDS

Free-space optics, virtual reality, wireless link, tracking and pointing.

Work done by all the authors while at Stony Brook University. Current affiliations: (i) Max Curran, Google Inc., WA, USA; (ii) Torin Rockwell, Hoshizaki America, GA, USA; (iii) Kai Zheng, Apple Inc., CA, USA; (iv) Mallesham Dasari, Carnegie Mellon University, PA, USA..

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.

SIGCOMM '22, August 22–26, 2022, Amsterdam, Netherlands

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9420-8/22/08...\$15.00

<https://doi.org/10.1145/3544216.3544255>

ACM Reference Format:

Himanshu Gupta, Max Curran, Jon Longtin, Torin Rockwell, Kai Zheng, and Mallesham Dasari. 2022. Cyclops: An FSO-based Wireless Link for VR Headsets. In *ACM SIGCOMM 2022 Conference (SIGCOMM '22)*, August 22–26, 2022, Amsterdam, Netherlands. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3544216.3544255>

1 INTRODUCTION

Virtual Reality (VR) is an exciting and rapidly emerging technology that creates a simulated environment within which the user becomes fully immersed. The ultimate goal of a VR systems is to provide an experience that is indistinguishable from physical reality—deemed *Life-Like VR* [31]. Standard high-performance VR systems use a powerful computer that renders the virtual scene which is then displayed in a VR headset (VRH) worn by the user. (Self-contained VR systems such as Oculus Quest compromise on performance; see §2.1). Typically, a wired connection connects the computer to the VRH, due to the high data rates required. However, to provide a true life-like and immersive experience, it is highly desirable that the VRH-renderer connection be wireless to allow the user to experience VR freely with “no strings attached.” Moreover, to match the limits of human visual perception, the VRH-renderer link should be able to support very high data rates (tens to hundreds of Gbps). The goal of our work is to design such a high-bandwidth VRH wireless link.

Recently, there has been strong interest to create wireless VRHs, in the industry as well as the research community; however, they are unable to provide the desired data rates, mainly due to the fundamental limitations of the wireless technologies used. E.g., self-contained Oculus Quest 2 uses a WiFi link, and HTC-Vive uses mmWave link that can provide up to a few Gbps. Similarly, recent works [18, 22, 60] have proposed or built solutions based on mmWave IEEE 802.11ad standard that can provide up to 7 Gbps.

To circumvent the bandwidth limitation of above technologies, we use free-space optical (FSO) technology to build a VRH link. FSO links can provide high data rates, up to Tbps [51], due to high frequencies of the light waves and narrow beams; commercially available optical transceivers can provide up to 400 Gbps [3, 54]. However, maintaining a FSO link can be challenging when the terminal(s) may be subjected to large and/or rapid motions as is the case of a head-worn VRH viewing a video. In addition, FSO links have line of sight (LOS) requirements and eye-safety concerns. In this work, we address the above challenges and develop a working prototype of a system that maintains a robust high-bandwidth link in presence of VRH movements. The core components of our approach and Cyclops system are: (i) an FSO link with sufficient

movement tolerance, and (ii) a tracking and pointing (TP) mechanism that maintains the link continuously by tracking the VRH position and pointing/realigning the beam as needed. The unique aspects of our setting makes the direct application of traditional TP approaches very challenging and likely even infeasible (see §3). Thus, we develop a novel TP mechanism based on function-approximation learning techniques that steers the beam as needed using steering mirrors at both ends.

Contributions. In the above context of developing an FSO link for VRH, we make the following contributions: **(a)** We design and develop a novel TP mechanism based on learning techniques leveraging the VRH’s built-in tracking system; the designed TP is able to achieve high accuracy and is highly effective in keeping the beam realigned as the VRH moves. **(b)** We design a steerable FSO link, using commodity optical components, for the VRH; the link uses a diverging (rather than collimated) beam with an appropriately chosen divergence angle to provide sufficient link movement tolerance. **(c)** We build a working prototype for Cyclops, demonstrate its robustness using 10G and 25G FSO links for a range of VRH motions (with bounded speeds) and publicly-available user traces. Our work does not raise any ethical issues.

2 PROBLEM AND BACKGROUND

In this section, we provide some background on VR headsets and FSO links, and motivate the problem and approach.

2.1 Virtual Reality



Figure 1: Oculus Rift S.

Virtual Reality (VR) is a simulated experience that, using a variety of sensors, display, and specialized equipment, allows a user to move around and interact with a virtual world. VR is used for many purposes, including entertainment, games, and education. The ultimate goal of VR is to create a life-

like experience indiscernible from reality. High-performance VR systems use a powerful GPU-equipped desktop (*renderer*) to render a high-definition video and then stream the video via a high-bandwidth wired link to a head-mounted display called *VR headset (VRH)*. For a truly immersive VR experience, it is imperative that the connection between the renderer and VRH be wireless so that the user can move around freely [22, 31].

Bandwidth Requirements for a Life-Like VR Experience.

There have been several attempts [45, 46, 50] to optimize for high quality VR content by using various schemes, but these entail decoding burden on the VRH which can introduce high motion-to-photon latency, consequently motion sickness, and decreased quality of user experience. Thus, ideally, we would like to stream “raw” data from the renderer to the VRH. The raw data rate for high-precision content such as 4K/8K-360° and volumetric 6-DoF videos comprising of point clouds and meshes [53], light field video with RGB+Alpha+depth channels [26] can range from a few 10s to 100s of Gbps. E.g., even a 2D uncompressed 8K RGB video at 30 frames

per second (fps) requires ≈ 24 Gbps; adding the Alpha+depth channels for transparency and 3D world would increase the required data rates to as high as 200Gbps [26, 53]. A recent work [31] estimates the bandwidth requirements for a life-like rendered video to be as high as 2.7 to 27 Tbps based on 1800 frames/sec, and expects such content to be available in commercial VRHs by 2028.

Related Works; State-of-the-Art Wireless VRHs. The goal of our work is to design a high-bandwidth wireless link between the renderer and VRH, to facilitate a life-like VR experience. Self-contained VRH systems such as Oculus Quest 2 [20] connect wirelessly to a remote renderer on internet, but these cannot provide a rich content due to the low-bandwidth WiFi link. Only available VRH wireless adapter [18] for HTC-Vive is based on mmWave/WiGiG technology (IEEE 802.11ad) and thus limited to a few Gbps throughput. Prior research works on developing wireless VRH links are unable to provide the desired bandwidths. In particular, [44] proposes use of visible-light communication technology connecting multiple transmitters on the ceiling with multiple receivers on each VRH. They estimate that multi-Gigabit bandwidth can be provided via such a system, but do not build a prototype. Some works [22, 60] have proposed using mmWave-based links that can potentially provide up to 7 Gbps: [60] demonstrates a 850 Mbps prototype, while [22] proposes a mirror-based alignment solution to circumvent the line-of-sight (LOS) issue. The mmWave and MIMO based IEEE 802.11ay standard in the making aims to provide a peak theoretical bandwidth of 176 Gbps [21, 37], but it is in early stages of development [56]: recent ns3-based simulation and evaluation of the 802.11ay standard with four bonded channels achieved a maximum throughput of only 30 Gbps [25]. TeraHertz is a promising wireless technology that can potentially provide data rates of upto 100 Gbps [27, 28], but is currently at a very nascent stage with prototypes developed mostly in research settings. Our preliminary work [55] explored feasibility of an FSO wireless link using a simple TP approach under simplifying assumptions (e.g., no mixed motions; see §3), by modeling the link in an optical simulator without developing any prototypes. In this paper, we develop a novel TP mechanism with a working link design and prototype.

2.2 FSO: Motivation and Background

We now motivate our approach and provide FSO background.

Why FSO? The main motivation for using FSO for a wireless VRH link is its ability to provide 100s of Gbps to Tbps bitrates [49]; e.g., recently, [51] demonstrated a 1.72 Tbps FSO link over 10.45 kms. The advantages of FSO links comes from the very high frequencies and collimating properties of light waves which yield narrow beams with low divergence, and unregulated spectrum. FSO links are also intrinsically secure, as (i) they are hard to intercept without detection, and (ii) they are naturally blocked by walls. Nevertheless, FSO beams, being narrow, need to be precisely aligned which can be challenging when the transceiver(s) are mobile. In addition, FSO communication is also limited by LOS, and can have eye safety concerns. In our work, we are able to circumvent the above challenges and design a viable link.

SFP-Based FSO Link Design. We build our prototype using commodity SFP transceivers [4] as in [32, 38–40], as they demonstrate feasibility of a small form-factor link [40] and eye-safety (SFPs are




Figure 2: SFP-based FSO link [40].

Class 1 safe [10, 19]). SFPs are available in a variety of wavelengths and data rates (up to 400 Gbps [13, 54]). To create an FSO link using SFPs, a collimated beam is launched into the air from a transmitter SFP via an optical fiber connected to a collimator; the beam is then captured at the receiver using a collimator into a fiber connected to an SFP. See Fig. 2. For the beam to be optimally captured into the fiber at the receiver, the beam must strike the receiver collimator lens at its center at a 90° incidence angle.

Tracking and Pointing (TP) Mechanisms. As FSO beams are very narrow, any movement of a transceiver can potentially disconnect the link. To keep an FSO link continuously connected, we need to have an active tracking and pointing (TP) mechanism consisting of a *tracking* mechanism to track/localize the transceiver(s) and a *pointing* mechanism to realign the beam in real-time. Typical TP mechanisms [41, 43] include detectors (e.g., position-ensing diodes [24, 35, 48], accelerometers, cameras, GPS [59]) for tracking, and steering mirrors [32] or gimbal [24, 48, 59] for pointing. In Cyclops, we use a novel TP mechanism based on function-approximation/learning techniques, while leveraging the VRH’s own built-in tracking system.

TP Metrics. The key performance metrics of a TP system are latency and accuracy, which dictate the maximum movement speeds of the terminals that can be tolerated without disconnecting the link. The TP *latency* is mainly composed of: (a) the tracking latency, since any tracking mechanism has a finite update frequency; (b) the latency incurred to realign the link’s beam. The TP *accuracy* signifies how accurately is the beam realigned.




Figure 4: Schematic of a Galvo Mirror (GM).

Galvo Mirrors (GMs). Our designed TP mechanism uses Galvo Mirrors (GMs) as steering mechanisms to steer the beam (as in [32, 40]). A typical GM consists of two orthogonal mirrors (called *first*




Figure 3: CDFs of VRH linear and angular speeds for VR applications.

and *second* mirrors), each of which can be independently rotated with extreme precision by applying an appropriate voltage to its motor. See Figure. Such a GM can thus steer a beam towards any target within a predetermined rectangular cone (called *coverage cone*). GMs can be **extremely reliable**; some have been life-test proven to tens of billions of cycles of operation [6].

Characterizing VRH Movement. A TP mechanism’s viability depends on the movement speeds of the link terminals. One way to characterize movement speeds is to measure linear and angular speeds of movements. Recently, we [55] conducted such a study and observed that during normal use, the angular and linear speeds of a VRH were at most 19 deg/s and 14 cm/s respectively. See Fig. 3. We wish to design a wireless VRH link that can handle such VRH movement speeds.

3 CYCLOPS’S HIGH LEVEL STRATEGY

Overall Cyclops System. Our overarching goal is to build a robust FSO link from a renderer (or transmitter, TX) to a VRH (or receiver, RX). The VRH is worn on the user’s head. To maintain clear LOS, we envision affixing the TX on the ceiling. See Fig. 5. For simplicity, we focus only on a uni-directional link from the renderer to the VRH. To create a viable FSO link that maintains the link connectivity while the user moves her head, we need to design:

1. A link with sufficient movement tolerance. This is essentially achieved by an appropriate optical design to generate a wide-enough beam and capture it efficiently at the receiver. We discuss this in §5.1.
2. A robust tracking and pointing (TP) mechanism, which tracks the RX movement and a realigns the beam as RX moves. The effectiveness of a TP mechanism depends on the link’s movement tolerance. The TP design is the key optics contribution of our work. We discuss the overall TP and its tracking mechanism below, while the pointing mechanism is described in §4.¹

Occlusions, Coverage, Cost, Size, Power, and Eye Safety. To circumvent occasional occlusions and/or limited field-of-view coverage of the GMs, we can use multiple TXs on the ceiling with appropriate handover techniques. As analyzed and shown in [40], steerable SFP-based links can indeed be designed with a small size, cost and power² footprint of terminals, especially when commercialized to

¹Note that the TP design is independent of the link design; the only link/beam property implicitly assumed in designing the TP (in particular, the Lemma 1) is that, for maximum received power, the beam must strike the lens at 90° degrees incidence angle.

²Total power usage of our system (with two SFPs and two GMs) should be at most a few watts, resulting in minimal (\$1-10/year) electricity usage cost.

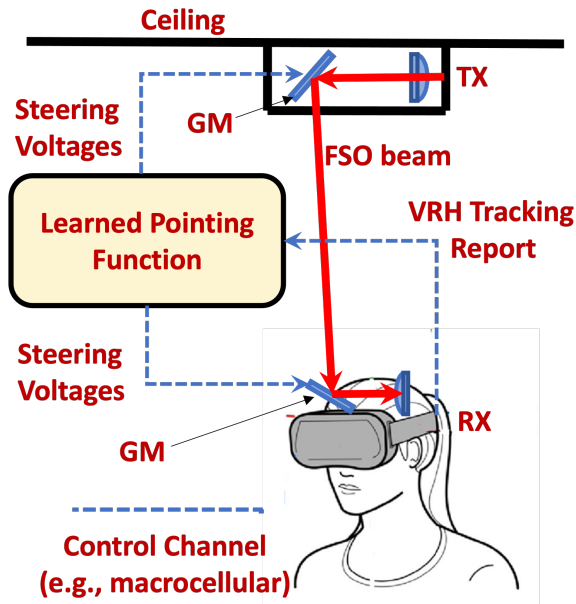


Figure 5: Cyclops's Tracking and Pointing approach.

scale; thus, there are no serious size or cost concerns with our FSO-based solution. Our prototypes use Class I lasers, with amplifiers used only to compensate for signal attenuation (see §5); thus there are no eye-safety concerns either.

Cyclops's Tracking and Pointing (TP) Approach. The main challenge in creating a robust TP system for Cyclops is that the RX may have relatively significant lateral and/or angular movement from its aligned position; this means that, to realign the beam perfectly, just steering the beam at the TX to point to the RX collimator's center is not sufficient and we need to also steer the beam at RX to make the incident angle 90 degrees at RX (see §2.2). The above is particularly important for short-range links such as ours as these links are very sensitive to RX angular movement which can be significant and frequent in our context. The need to steer beam at RX as well as TX makes a direct application of traditional TP approaches to our context very challenging and likely even infeasible. E.g., photodiode-[32] or probe-based [55] tracking is challenging to adapt here due to RX's angular movement, and would need to be deployed at both ends; more importantly, the associated pointing technique will incur prohibitively high latency due to the need to *jointly* optimize the TX and RX steering parameters. Thus, for Cyclops, we develop a novel TP mechanism based on learning (function approximation) techniques, while leveraging the VRH's inbuilt tracking system. At a high-level, Cyclops's TP approach is (see Fig. 5):

1. Use VRH's built-in tracking, described below, to track the *RX assembly* (which includes the VRH).
2. Use steering mirrors (we use GMs) at both the TX and RX to steer the beam, as needed.
3. At real-time, use a pointing function/mechanism \mathcal{P} to determine the voltages to be applied to GMs to steer and realign the beam, based on VRH tracking report. We learn the function \mathcal{P} using

pre-deployment and at-deployment training, as discussed in §4. A control channel (e.g., macro-cellular) may be used to transmit the VRH tracking reports and the computed GM voltages.

Tracking (the RX assembly). As mentioned above, Cyclops uses the VRH's built-in tracking system, henceforth referred to as VRH-T, to track the RX assembly's position. Here, by *position*, we mean location (i.e., x, y, z coordinates) and orientation (represented by, e.g., three angles), and the *RX assembly* includes the VRH and the optical elements including GM at the RX. The key challenge in using the VRH-T is that the position reported by VRH-T is the **position of some unknown point within VRH in an unknown coordinate space** (henceforth, referred to as *VR-space*). We address this challenge in §4.

VRH Tracking System (VRH-T). VRH-T uses an inertial motion unit (IMU) to compute the position. To compensate for error over time, VRH-T also utilizes independent cameras to localize and reduce the overall error. The precise camera localization system used in the VRHs varies across models. The update frequency and accuracy of VRH-T on current VRHs, while adequate for typical VR applications, are however not sufficiently ideal for our purposes (§5.2). However, revising the VRH-T was not our work's focus; we instead design a robust link, leveraging the given VRH-T.

4 CYCLOPS'S POINTING MECHANISM

In this section, we discuss the pointing mechanism which, on RX movement, realigns the beam by steering the GMs at TX and RX by applying appropriate voltages to them. This is done based on the VRH position as reported by the VRH-T. Thus, the *pointing function* \mathcal{P} can be written as:

$$\mathcal{P}(\text{VRH Position}) = \langle v_1^{tx}, v_2^{tx}, v_1^{rx}, v_2^{rx} \rangle$$

where v_1^{tx} (v_1^{rx}) and v_2^{tx} (v_2^{rx}) are the voltages to be applied to the TX (RX) GMs for aligning the beam. It is easy to see that \mathcal{P} is a well-defined function.

Challenges. As the \mathcal{P} function's form or expression is unknown, we learn it via function-approximation techniques. The key challenge in learning \mathcal{P} is that we require a very high accuracy since the link's movement tolerance is expected to be at most a few mms. Learning \mathcal{P} *directly* is infeasible, due to the prohibitive training cost: it would take many years to collect the training data, as per our preliminary analysis.³ Thus, we learn \mathcal{P} indirectly as described below.

Our Approach to Learning \mathcal{P} . One of the fundamental "components" of the \mathcal{P} function is the GM model, which relates the GM voltages to the output beam direction. Thus, our approach to learning \mathcal{P} involves first learning the GM model, and has the following high-level steps (see Fig. 6).

1. *Learn the GMA Model/Function*⁴ \mathcal{G} . The GMA (GM assembly; formally defined shortly) function \mathcal{G} maps a pair of GM voltages

³Even for a small space of 1 m^3 of RX movement, the needed domain size of \mathcal{P} is of the order of 10^{18} for mm-level accuracy. For a complex function \mathcal{P} , we stipulate that we need tens of thousands or many magnitudes more samples to learn \mathcal{P} with mm-accuracy; this can take years, as gathering each sample takes several minutes as determining the (four) voltages that align the link takes a few minutes of exhaustive search [32]. In our experiments, we tried to learn the much simpler function \mathcal{G}' (see §4.3) directly, but even several hundred training samples yielded an error of a few cms.

⁴We use the words model and function interchangeably.

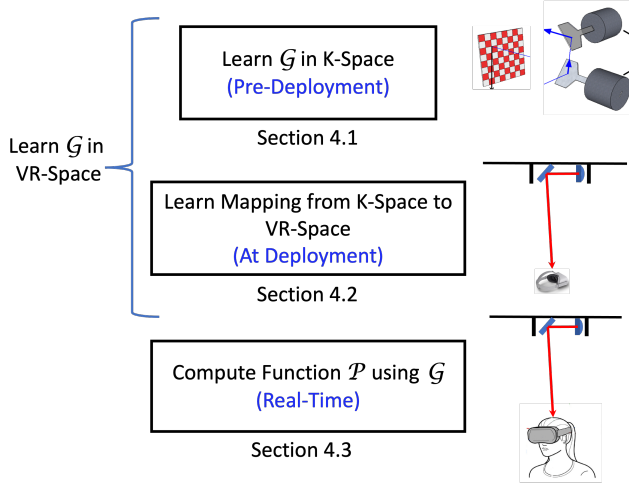


Figure 6: Three stages of learning Cyclops's Pointing Mechanism.

onto the output beam specifications. We must learn \mathcal{G} in the VR-Space. We do so by first learning it in a *known* coordinate space K-space (§4.1), and then learning the mapping from K-space to VR-space (§4.2).

2. *Pointing Function \mathcal{P}* . Once we have learned the TX- and RX-GMA functions in VR-space, we can use them iteratively to compute the pointing function (§4.3).

Offline vs. Online Training, Re-training, and Deployment Steps. We note that learning of \mathcal{G} happens offline, with the step of learning \mathcal{G} in K-space (§4.1) done at pre-deployment by the GM manufacturer and the mapping-step (§4.2), which is deployment/location specific, is done one-time at deployment. Computation of \mathcal{P} happens in real-time for each VRH position.⁵ Thus, in case of re-deployment or VRH-T drift, the only re-training (calibration) that needs to be re-done is the mapping step. Overall, deployment of the Cyclops system at home will involve the following steps: (i) affixing the TX to the ceiling, and (ii) conducting the training/calibration of the system, by moving the VRH around.

4.1 Learning the GMA Function in K-space

We describe our technique to learn the GMA function in a known co-ordinate space, say K-space. We start with formalizing a few terms.

- *GM Assembly (GMA)*. GMA is responsible for launching (or capturing) and steering the optical beam. In particular, the GMA consists of: (i) SFP, which generates or captures the beam, (ii) optical assembly, i.e., collimator, the coupling optical fiber, etc., and (iii) a GM, used to steer the beam. The TX-GMA (i.e., the GMA at the TX) essentially launches a beam within the GM's cone, based on the two voltages applied to the two GM mirrors. Similarly, the RX-GMA steers the received beam, based on the voltages applied, on to the collimator at a desired point or incident angle.

⁵We could potentially also compute \mathcal{P} offline for all possible VRH locations and merely do a look-up at real-time, but such a scheme is not feasible due to the large number (e.g., $\approx 10^{18}$ in a m^3 space) of VRH positions required for mm-level accuracy.

- *GMA Model \mathcal{G}* . Our goal, in this subsection, is to learn the GMA function \mathcal{G} which map the two voltages applied to a GM's mirrors and outputs the originating point and direction of the output beam. More formally, for a *given position* of a GMA, we represent the GMA function as:

$$\mathcal{G}(v_1, v_2) = (p, \vec{x}),$$

where v_1 and v_2 are the voltage inputs to the GM, and (p, \vec{x}) represents the emanating beam with \vec{x} being its direction/vector and p being the originating point on the second GM mirror.⁶ See Fig. 7. The GMA model \mathcal{G} essentially represents the optical and physical configuration of the GMA's internal components.

High-Level Strategy to Learn \mathcal{G} in K-space. First, we note that learning \mathcal{G} directly is infeasible as it is nearly impossible to accurately determine the output beam's originating point p (see Fig. 7), which is needed for any training sample.⁷ Thus, we *learn* the GMA model \mathcal{G} efficiently in two steps: (a) Determine a closed-form parameterized expression for \mathcal{G} ; (b) Learn the parameter values in the parameterized expression, in a *known* coordinate space K-space. The above steps, discussed below, obviates the need to determine the output beam's specifications. We note that the above steps are same for TX-GMA and RX-GMA, so we just focus on TX-GMA and simply use GMA to refer to TX-GMA. In a commercial system, TX-GMA and RX-GMA may use the same components and design, and hence, may have the same known model. However, in a prototype development, TX-GMA and RX-GMA need to be manually assembled using a GM and a collimator, and hence, will likely have different values for p_0 and \vec{x}_0 parameters.

(A). Parameterized Expression for \mathcal{G} . To determine a parameterized expression for the GMA function \mathcal{G} , we analyze the beam path for a given configuration of mirrors and related optical components. Formally, the parameterized expression for \mathcal{G} is derived over the following GMA parameters. See Figure 7.

1. *Input Beam*. The input beam is represented by: (i) p_0 , the originating point, and (ii) \vec{x}_0 , the beam's direction vector.
2. *Mirror Specifications*. The first (i.e., the mirror that the input beam hits first) mirror in the GM is represented by: (i) \vec{n}_1 , the normal of the mirror at zero voltage; (ii) q_1 , a point on the mirror's plane and its rotation axis; (iii) \vec{r}_1 , the direction vector of the rotation axis. Similarly, the second GM mirror is represented by \vec{n}_2 , q_2 , and \vec{r}_2 .
3. *Voltage to Angle Function*. For each mirror, we need to also specify how the mirror's normal changes with a change in voltage. In experiments, we observed that the change in mirror's normal angle is proportional to the change in voltage Δv and thus can be estimated by $\theta_1 \Delta v$ for some constant θ_1 . We assume θ_1 to be same for both mirrors.

Using the above parameters, we can now derive a parameterized expression for \mathcal{G} where $\mathcal{G}(v_1, v_2) = (p, \vec{x})$ as follows.

⁶In simpler applications with limited range of motions, p may be assumed to be a constant as in [32, 33], but in reality it depends on the voltages—this dependence results in distortion [58] and needs to be considered for high accuracy.

⁷It can be simpler to instead learn (directly) the reverse function \mathcal{G}' which maps a target-point to voltages, as it doesn't involve measuring output beam's specifications. However, we still need a closed-form expression for \mathcal{G}' to learn the mapping K-space to VR-space (§4.2), and unfortunately, deriving an expression for \mathcal{G}' seems very challenging.




Figure 7: GM Assembly (GMA) parameters (second mirror parameters not shown). Also shown (in red) are the output beam specification parameters (p, \vec{x}) .




Figure 8: Learning \mathcal{G} 's parameter values: For each grid point (x, y) on the board, we record the pair of voltages needed to hit it.

1. *New Normals of the Mirrors.* Let $R(\vec{r}, \theta)$ denote the rotation matrix that rotates a vector by angle θ about the rotation axis \vec{r} . Then, the new normals of the two mirrors are given by:

$$\begin{aligned}\vec{n}_1 &= R(\vec{r}_1, \theta_1 v_1) \vec{n}_1 \\ \vec{n}_2 &= R(\vec{r}_2, \theta_1 v_2) \vec{n}_2\end{aligned}$$

2. *Reflection from the Mirrors.* Let \mathcal{R} be the reflection function for a mirror that maps an input beam parameters to the output beam parameters, given the mirror position. Then, the output (p, \vec{x}) of \mathcal{G} function can be derived as below.

$$\begin{aligned}(p_{mid}, \vec{x}_{mid}) &= \mathcal{R}(p_0, \vec{x}_0, \vec{n}_1, q_1) \\ (p, \vec{x}) &= \mathcal{R}(p_{mid}, \vec{x}_{mid}, \vec{n}_2, q_2)\end{aligned}$$

Above, (p_{mid}, \vec{x}_{mid}) is the specification of the beam after hitting the first but before the second mirror; note that the points q_1 and q_2 on the mirrors did not change with the rotation in (1), as they are on the rotation axis.

(B). Learning \mathcal{G} 's Parameter Values. We now discuss the methodology used to learn/estimate the values of the parameters in the \mathcal{G} function expression, viz., $p_0, \vec{x}_0, \vec{r}_1, \vec{n}_1, q_1, \vec{r}_2, \vec{n}_2, q_2$, and θ_1 . At a high-level, we *learn* the parameter values of function \mathcal{G} by first

collecting a set of appropriate training samples, and then employing learning techniques to estimate the parameter values that minimize the total "error". In particular, we set up a system consisting of a planar board with grid lines, and the GMA placed (fixed) in front of it at a certain distance. See Figure 8. For simplicity, we assume K-space coordinate system's x - y plane to be the board. In this set-up, we have the GMA launch a laser beam, and vary the voltages supplied to the GM so that the beam hits each of the grid (*target*) points on the board. More formally, for each grid point (x, y) on the board, we determine the voltage pairs (v_1, v_2) that must be supplied to the GM for the beam to hit (x, y) . We gather about 266 (see below) such 4-attribute training samples (x, y, v_1, v_2) . Then, we use non-linear optimization techniques [57] to estimate/learn the parameter values of \mathcal{G} such that the error $\sum_{(x,y)} d((x, y), f(\mathcal{G}(v_1, v_2)))$ is minimized; here, $d(\cdot)$ is the distance function, \mathcal{G} is the learned GM function, and $f(\cdot)$ maps the output beam specification to the point on the board hit by the beam. Essentially, we make a good initial guess, and then use the solver [57] to iteratively improve the parameters to minimize the error; in the current context, the initial guess came from the available CAD design of the GM from its manufacturer and manual measurement of GM's position in the K-space. Even though the above training is based on target points on a 2D plane, we are able to learn a *general* \mathcal{G} function for arbitrary target points in space; this is likely because the set of training samples, though on a 2D plane, are not really "specialized" with respect to the GM, due to the distortion effect [58].⁶

Training Sample Details. In our prototype development, the planar board had a grid of 20×15 cells with each cell being 1 inch by 1 inch; the board was placed at a distance of 1.5 meters from the GM. The board having 300 cells, gives us an ability to take 336 samples (one at each intersection/boundary point). but for high accuracy we only used the internal points on the board, i.e., 266 points.

4.2 Learning the Mapping to VR-space

Note that the function \mathcal{G} learned in the previous subsection yields the output beam specifications in K-space. To get the beam specifications in the VR-space (of the deployed set-up), we need to learn the mapping from K-space to VR-space. TX-GMA and RX-GMA have their own independent K-space, and we need to learn the mapping for each⁸—for a total of 12 parameters, as each mapping can be represented by 6 parameters [30]. We refer to these 12 values as *mapping parameters*. We learn these 12 parameters jointly by gathering training data using aligned links, as described below.

Joint-Learning of the 12 Mapping Parameters. As accurately determining the GMA positions is very challenging, we learn the above 12 mapping parameters jointly. We do so by leveraging the obvious precision of an automated-exhaustive search to optimally align a beam; the exhaustive search finds the optimal combination of the

⁸Since the RX-GMA may move in real-time, its K-space does not remain fixed. However, the *relative* position of RX-GMA within RX-assembly (and hence, wrt VRH) remains fixed; thus, the mapping parameters for RX-GMA can be looked upon as determining its (or, its K-space) relative-position wrt to the point X in VRH whose position VRH-T reports. Now, in a commercial system, even though RX-GMA's relative-position is fixed at manufacturing time, the VR-space (and in particular, the point X) is determined only at deployment; thus, the RX-GMA's relative-position wrt to X needs to be learned at deployment.




Figure 9: Originating and Target points.


four voltages that maximizes the received power at the RX.⁹ a and control the GMs Since the parameter estimation is done one-time, the time taken (1-2 mins) by the search is tolerable. Overall, to estimate the above described 12 mapping parameters, we undertake the following three steps:

1. First, we extend the parameterized expression of \mathcal{G} from §4.1, which is in K-space, to expressions in VR-space for TX-GMA and RX-GMA; these expressions will naturally include their respective mapping parameters (six each).
2. Then, for various positions of the VRH in the deployed set-up, we collect 5-tuple training samples $(v_1, v_2, v_3, v_4, \Psi)$ where Ψ is the VRH position reported by the VRH-T and v_1 to v_4 are the voltages to the GMs' mirrors that align the beam as perfectly as possible via exhaustive search.
3. We estimate the 12 mapping parameters, using non-linear optimization [57], to minimize the below error function.

Error Function. Consider estimated \mathcal{G} functions, viz., \mathcal{G}_T and \mathcal{G}_R for TX-GMA and RX-GMA respectively, based on the current estimate of the 12 mapping parameters. For a given 5-tuple $(v_1, v_2, v_3, v_4, \Psi)$ for a fully-aligned beam as defined above, we use \mathcal{G}_T over (v_1, v_2) to estimate p_t (originating point of the beam) and τ_t (the target point that the beam emanating from TX-GMA hits on the RX-GM's second mirror). See Figure 9. The point p_t is a direct output of \mathcal{G}_T , while the point τ_t can be computed from the beam specification output of \mathcal{G}_T , 12 mapping parameters, and the VRH position (we skip the tedious details). Similarly, we use \mathcal{G}_R to compute p_r and τ_r . We now define the *error function* as: $\sum d(p_t, \tau_r) + d(p_r, \tau_t)$, where $d(\cdot)$ is the distance function and the summation is over all the 5-tuple training samples. The rationale for the above error function is that, at perfect beam alignment, each of the pairs of points (p_t, τ_r) and (p_r, τ_t) should coincide if the \mathcal{G} functions are perfectly accurate. See below Lemma.

LEMMA 1. *Given a FSO link set-up in Cyclops with each terminal (TX and RX) equipped with a collimator and a GM. Consider two optical paths: (i) Optical path of the beam emanating from TX, and (ii) the optical path of an imaginary beam emanating from RX. See Figure 9. We claim that the configuration (i.e., four voltages) of the two GMs that maximizes the received power at RX (from the actual beam emanating from TX) is the same as the configuration that ensures*

⁹To monitor the receiver power, we surround the RX's collimator by four photodiodes connected to a DAQ, as in our earlier work [32]

Figure 10: \mathcal{G}' Iteration.

that (i) p_t and τ_r coincide, and (ii) p_r and τ_t coincide. We assume that GMs can be steered with unbounded resolution.¹⁰ ■

We note that the above lemma is fundamental to Cyclops's design; in particular, it is fundamental to: (i) learning \mathcal{G} accurately in the VR-space which has many unknowns, and (ii) designing an accurate pointing function (§4.3).

4.3 Designing a Pointing Mechanism

The above yields TX-GMA and RX-GMA models in the common VR-space. We now show how to use these models to design the pointing mechanism \mathcal{P} which, in real-time, maps a given VRH position to the four GM voltages required to align the beam. The **high-level idea** is as follows. For a given VRH position, we can compute (from the parameters learned in previous subsections) an approximate point τ on RX-GMA's second mirror that the beam from TX should hit; then, we can determine TX-GMA voltages that would yield a beam hitting τ , by using the TX-GMA's \mathcal{G} function iteratively. To do the latter, we design an algorithm for a reverse GMA function \mathcal{G}' that maps a given target point τ to the two GM voltages that yield a beam passing through τ . And, to determine the precise (instead of just approximate) hitting point(s), we use Lemma 1 and \mathcal{G} and \mathcal{G}' functions for both GMAs iteratively. We discuss these in below paragraphs.

\mathcal{G}' Function: Mapping a Target Point to GM voltages. Given a GMA function \mathcal{G} and a target point τ , we wish to determine the mirror voltages v_1 and v_2 such that $G(v_1, v_2) = (p, \vec{x})$ and the beam (p, \vec{x}) passes through τ . Fortunately, we can design a purely computational procedure (i.e, requiring no additional training) to estimate \mathcal{G}' from \mathcal{G} ; note that \mathcal{G} is essential to our overall scheme anyway. In particular, first we initialize v_1 and v_2 to some reasonable values, and then update them iteratively until the output beam converges upon the target point τ . In particular, we change (v_1, v_2) at each iteration as follows. See Fig. 10.

1. Using the given GMA model \mathcal{G} , calculate $G(v_1, v_2) = (p_0, \vec{x}_0)$, $G(v_1 + \epsilon, v_2) = (p_1, \vec{x}_1)$, and $G(v_1, v_2 + \epsilon) = (p_2, \vec{x}_2)$, where ϵ is a small change to the voltages.

¹⁰When the resolutions are bounded, there may not exist pairs of points (p_t, τ_r) and (p_r, τ_t) that coincide. But, for GMs with high-enough resolutions (as is the case with our GMs), minimizing $d(p_t, \tau_r) + d(p_r, \tau_t)$ where $d(\cdot)$ is the distance function should still lead to near-optimal received power.

2. Let P be the plane perpendicular to \vec{x}_0 and containing τ , and let k_0, k_1 and k_2 be the points on P intersecting with the beams $(p_0, \vec{x}_0), (p_1, \vec{x}_1)$ and (p_2, \vec{x}_2) respectively.
3. Let \vec{u}_1 and \vec{u}_2 be the vectors on the plane P that represent the change in *direction* of the beam due to the ϵ change to v_1 and v_2 respectively. We find real numbers a and b such that $\vec{x}_0 + a\vec{u}_1 + b\vec{u}_2$ yields a vector that hits P at a point τ' closest to τ .
4. Finally, we update the voltages as $v_1 = v_1 + a\vec{u}_1$ and $v_2 = v_2 + b\vec{u}_2$. Here, we implicitly assume that change in the direction is linear to *small* changes in the voltages. We stop when a and b are smaller than ϵ .

In our evaluations, the above converged in 2-4 iterations.

Pointing Mechanism \mathcal{P} . Our overarching goal is to develop a pointing mechanism, i.e., given a VR position report, determine the 4 voltages for the 2 GMs that will align the beam perfectly. Ideally, we would like to determine the voltages based on received power at RX—however, even if we could establish a feedback based on received power, there is no efficient way to search for the 4 voltages that will maximize the received power. Neither can we configure the GMs to ensure that the beam hits the RX collimator close to its center near-perpendicularly, since there is no means to physically check that in real-time. Thus, we **leverage Lemma 1's** claim—and configure the GMs to coincide the corresponding originating and target points, which will ensure maximization of received power. Informally, we start with some initial values to the four voltages, compute originating points at both end, use these originating points as target points and use \mathcal{G}' to determine updated voltages. We repeat until the error is minimal. Below, \mathcal{G}_T (\mathcal{G}_R) and \mathcal{G}'_T (\mathcal{G}'_R) denote the TX-GMA (RX-GMA) models. Formally:

1. Initialize the four voltages v_{t1}, v_{t2}, v_{r1} , and v_{r2} to some values (e.g., zero); here, the voltages v_{t1} and v_{t2} refer to the voltages for TX-GM and v_{r1} and v_{r2} are for RX-GM.
2. Use the \mathcal{G} functions to determine the originating points p_t and p_r , i.e., $(p_t, \vec{x}_t) = \mathcal{G}_T(v_{t1}, v_{t2})$, $(p_r, \vec{x}_r) = \mathcal{G}_R(v_{r1}, v_{r2})$.
3. Use p_t as the target point for RX-GMA and p_r as the target point for TX-GMA, use \mathcal{G}' functions to determine the voltages to hit these target points, and update the voltages to these newly determined voltages. That is, set $(v_{t1}, v_{t2}) = \mathcal{G}'_T(p_r)$ and $(v_{r1}, v_{r2}) = \mathcal{G}'_R(p_t)$.
4. Repeat, until the changes to the 4 voltages is below some threshold (we use minimum GM voltage step).

In our evaluations, the above converged in 2-5 iterations.

5 CYCLOPS EVALUATION

We now design and evaluate a system prototype. In the following subsections, we discuss: (i) An FSO link design with sufficient movement tolerance, (ii) Overall prototype design, and evaluation of learned \mathcal{G} and overall TP mechanism, (iii) Throughput evaluation for varying VRH motions, (iv) Prototype simulation over publicly-available VRH user traces.

5.1 Link Design and Evaluation

In our FSO link, each terminal consists of a SFP connected via fiber to a collimator followed by a GM for beam steering. We have created

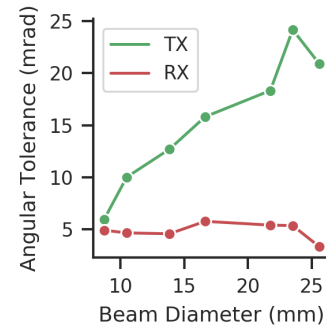


Figure 11: Angular tolerance for varying beam diameter at RX.

10Gbps and 25Gbps links of 1.5-2m length.¹¹ Our techniques would work for even higher bandwidth (40Gbps+) FSO links but would require custom-built optical elements (see §6). Some of the below discussion is specific to 10G, and we discuss the 25G link design changes in §5.3.1. For the 10G link, we use the 10G 1550nm ZR-SFPs [14] which have a transmission power of 0-4dBm and the receiver sensitivity of -25dBm, and the GVS102 [36] GM with an angular accuracy of $10\mu\text{rad}$ and small-angle latency of $300\mu\text{sec}$. As in [38], we use an amplifier [34] to compensate for the coupling losses due to using a fiber rather than an exposed photodetector as in an actual system.¹² Below, we discuss design options to maximize link movement tolerance using commodity optics.

Improving Movement Tolerance. We need to design a link with sufficient movement tolerance. Movement tolerance can be increased by using a wider beam, for which there are two options: (a) a wider collimated beam, or (b) a diverging beam that is wide enough when it reaches the RX. We discuss both options. The **first** option has minimal divergence, but the beam can also get "clipped" by the TX GM, which can defeat the whole purpose. Our GMs allow 10mm beams; using GMs that allow larger beam widths [9] also incur higher response time offsetting their advantage. The **second** option circumvents the above challenge, but it introduces additional issues: (a) variable-width received beam (due to varying link ranges), and (b) efficiently capturing a diverging beam. In our context, the TX-RX distance is small and remains relatively similar; thus, a diverging beam is a promising option.

Link Tolerance Evaluation. We thus evaluate two options: (a) a wider (20mm) collimated beam, created using a beam expander [16]; (b) a diverging beam, created using an *adjustable* aspheric lens [23] collimator at the TX which allows us to control the angle of divergence (and thus, beam diameter at RX). We use two evaluation metrics: (i) *angular movement tolerance* [32], i.e., the maximum angular movement from the aligned position for which the link remains connected; (ii) *peak received power*, i.e., the received power when the beam is aligned. Lateral tolerance is not considered here, as the design constraints due to it are subsumed by those due to angular tolerance. See Table 1 for 10G link evaluation results (see

¹¹For longer links, the TP mechanism/design remains the same, but the link design may need to be optimized via a more efficient coupling and/or higher link budget.

¹²Transmit power of SFPs (0 dBm) is much less than the safety limit [19] for 1550nm wavelength. Thus, using an amplifier retains eye safety, especially in light of our choice of diverging beam and coupling losses.

	Collimated	Diverging
TX Angular Tolerance	2.00 mrad	15.81 mrad
RX Angular Tolerance	2.28 mrad	5.77 mrad
Peak Received Power	15 dBm	-10 dBm

Table 1: Link angular movement tolerances and peak received power, using a collimated or a diverging beam (with 20mm beam diameter at RX).

§5.3.1 for 25G link), which shows the trade-off of using the diverging beam over the collimated beam: increased tolerance but lower peak received power. Since, we wish to maximize movement tolerance, the diverging beam seems a better design option. Above, we have used a diverging beam with 20mm diameter at RX; we now determine the optimal diameter. We prioritize optimizing RX (rather than TX) angular tolerance since maximum angular speed of VRH is determined only (due to geometry) by RX angular tolerance and the maximum linear speed of VRH is determined by the minimum of the two (RX and TX) angular tolerances. Figure 11 plots the TX and RX angular tolerances for varying beam diameters at RX. We observe that RX angular tolerance peaks at 5.77 mrad at the 16mm beam diameter; we thus choose this.

5.2 Prototype; \mathcal{G} and TP Evaluation

Prototype. In our prototype, the TX assembly is placed in a static position. See Fig. 12(a). The RX assembly consists of link components plus an Oculus Rift S [52] VRH (for tracking). For rigidity, we place the entire RX assembly on a small breadboard. To facilitate purely linear or purely angular motion, we place the breadboard on a rotation stage mounted on a linear rail. See Fig. 12(b). To simulate arbitrary user motion, we detach the RX assembly from the linear rail and rotation stage and move it freely by hand (our current prototype is too bulky to be worn on a head). See Fig. 12(c). Most of the bulk in our prototype is only for our evaluation purposes.

GMA Model \mathcal{G} Evaluation. Recall that the GMA model \mathcal{G} is learned in two stages: (i) Learn \mathcal{G} in K-space and (ii) Map the parameters to the VR-space. Also, recall that the first stage is identical but done independently for the TX and RX GMAs, while the second stage is done jointly. We collect about 250 and 30 training samples for each stage respectively. Table 2 shows the errors for the first and the combined stages of the two GMAs; we note that it is not possible to separate the errors between the two stages once combined. As expected, the first stage errors are minimal (1-2 mm avg.) and similar for both the GMAs. The combined error for the TX-GMA is less than that of the RX-GMA; this is likely because the second stage for RX-GMA is based on its position *relative* to VRH which may not be perfectly fixed as assumed. Overall, the average combined errors are 2-4mm, which suggests that the lateral and angular tolerances of our link should be *at least* 2-4mm and 3-6 mrad respectively for our 1.5-2m length link—this somewhat matches our link evaluation results (Table 1).

TP Performance. We now evaluate the overall TP latency and accuracy, which includes the learned function \mathcal{P} 's performance. As noted in §3, traditional TP approaches are challenging or even infeasible to apply in our context.

	Avg. Error	Max. Error
First Stage (TX)	1.24 mm	5.30 mm
First Stage (RX)	1.90 mm	5.41 mm
Combined (TX)	2.18 mm	4.07 mm
Combined (RX)	4.54 mm	6.50 mm

Table 2: Errors of the first and combined (first + second) stages of estimating the TX and RX GMA models.

Tracking Frequency, and TP Latency. Recall that our TP mechanism uses the VRH-T for tracking the RX. We observed that the VRH-T reported the VRH position every 12-13ms except 0.7% of times at 14-15ms [11]; this includes the small (< 1 ms) latency due to RF control channel.¹³ Pointing latency is composed of the computation time to compute the final voltages, and the latency incurred to rotate the mirrors; the former is minimal (in μ secs), while the latter is about 1-2 msec comprised mostly of digital-to-analog conversion latency at a DAQ device. As TP's latency is much lower (1-2 ms) than the frequency at which it occurs (every 12-13 ms at VRH-T updates), **a custom VRH-T with much higher tracking frequency will improve Cyclops's performance significantly.** Note that the TP latency doesn't result in any streaming latency, as the link remains *continuously* operational as long as the movement speed is bounded.

Tracking and Overall TP Accuracy. We observed that the VRH-T is inherently noisy. E.g., over a 30 minute period, even with VRH completely stationary, the reported location and orientation varied by up to 1.79 mm and 0.41 mrad respectively. To measure overall TP accuracy (note, we can't measure pointing accuracy separately), we move the RX assembly randomly, "lock" it in place, run the TP algorithm to align the misaligned link, and finally, record the received power and average link throughput (using iperf [42]). We compare these values to an optimally (no TP) aligned FSO link. We repeat the above test 10 times. We observe that in all tests, the link achieve the optimal throughput, suggesting that the TP mechanism is accurate within the link's movement tolerance. The received power was also only slightly lower (at -13 to -14 dB) than the peak received power of -10dBm (see Table 1), implying high TP accuracy.

5.3 Throughput Evaluation

We now evaluate our prototype's throughput for various VRH motions: purely linear or angular, and arbitrary motions. All evaluations are over a link of 1.5-2m length. We start with 10G prototype evaluation, and discuss 25G prototype later.

Purely Linear or Angular Motions. As mentioned before, to simulate purely linear or angular motions, we mounted the breadboard on a rotation stage and then a linear rail (see Fig. 12(b)). Then, for purely linear motion, we move the platform along the linear rail (with the rotation stage locked), and for purely angular motion, we move the platform using the rotation stage (while being locked on the linear rail). More specifically, for purely linear motions, the RX assembly is moved continuously from one end of the rail to the other in a single smooth "stroke." The assembly, momentarily

¹³In Cyclops, the only control information transfer is the VRH position to the TX; computing \mathcal{P} is purely computational without any real-time feedback.




Figure 12: (a) TX assembly with the SFP, GM, and collimator. (b) RX assembly with the VRH, collimator, linear rail, and rotation stage. (c) RX assembly detached and moved around by hand to simulate arbitrary user motions.




Figure 13: Throughput (left) and received power (right) for linear (top row) and angular (bottom row) motions.

comes to rest to turn at one end, and is then moved in the opposite direction. This process is repeated with gradually increasing stroke speeds—until the observed throughput drops (to zero); at that point, we stop momentarily and slowly start moving again as above. For every 50ms time window, we measure the average throughput (using iperf [42]) as well as the linear speed (using VRH-T reports). Note that the link range naturally varies during this test. See Fig. 13. We observe that consistently the link throughput remains optimal [32, 40] at 9.4 Gbps for linear speeds below 33 cm/sec (and for up to 39.15 cm/sec), which is much higher than our requirements (§2.2). Note that, once the link is lost, it takes a few seconds to regain the link partly due to the SFPs taking a few seconds to report that the link is up, after receiving the light [38]. Similarly, Fig. 13 plots the results for purely angular motion, which

is done similarly as above but using just the rotation stage. We observe that the link throughput remains optimal for angular speeds below 16-18 deg/sec (and for up to 18.95 deg/sec) which is about the same as our requirements.

Received Power. We also ran *separate* tests to log received power for purely linear and angular motions, and observed that the power remained above -25 to -30dBm for the above speed bounds. See Fig 13. Interestingly, the power dropped to only about -32dBm for linear speeds of 70 cm/sec and to about -38dBm for angular of 100 deg/sec, which implies that with even a 7-13dB improvement in the coupling loss, the prototype would be able to support much higher movement speeds. Our coupling loss for the diverging beam is quite high at -30dB, as we are restricted to using commodity collimators




Figure 14: Throughput (top) and received power (bottom) for arbitrary user motions.

designed primarily to capture collimated beams; customized optics should improve the coupling loss significantly.

User Study (Arbitrary Motions). As our prototype is very bulky, it is not feasible to wear it on the head. So, to simulate realistic head movement, we **detach the RX assembly from the rotation stage and linear rail**, hold it in hands, and move it around in front of the TX (see Fig. 12(c)). Fig. 14 plots the throughput and received power, with the linear *and* angular speeds. We observe similar performance as with the pure motions; in particular, the link maintains optimal throughput for motions undergoing simultaneous linear and angular speeds of below 30 cm/sec *and* 16-18 degrees/sec respectively. We observe that the received power remains above -40dBm for angular speeds of up to 100 deg/sec with linear speeds of 30 cm/sec.

5.3.1 Higher-Bandwidth Prototype Evaluation. Though our core technique (the TP mechanism) generalizes to higher bandwidths without change, the higher-bandwidth links may require custom optical elements for an effective link design. In this section, we extend our prototype with a 25Gbps link, while illustrating some of these challenges.

25G Link Design and Challenges. To create a 25 Gbps FSO link, we use single-wavelength single-strand 25G SFPs (called SFP28s [1, 2]). However, the maximum link budget of available SFP28s is only 19-25dB, that of SFP28s ERs [2] In addition, there are no 25G network interface cards (NICs) available that are compatible with the

	Reqs. (§2)	10G		25G	
		(Pure)	(Mixed)	(Pure)	(Mixed)
Linear (cm/s)	14	33	30	25	15
Angular (deg/s)	19	16-18	16	25	15-20

Table 3: Summary of Results. Speed requirements (§2) and tolerated speeds by 10G and 25G links.

SFP28 ERs [7]; note that NICs are essential to evaluate throughput performance. Thus, we used the shorter-range SFP28 [1] (with Intel XXV710-Based NICs [8]) which have even smaller link budget of 12-18dB, which is about 13dB less than the SFPs used in our 10G prototype. We were able to achieve a small (2-3dB) improvement in the received power by use of adjustable-focus collimators [15] at TX and RX to better capture a diverging beam. The overall link yielded slightly better RX angular tolerance (0.5 deg = 8.73 mrad), but worse TX angular tolerance (also about 8-9 mrad) compared to our 10G link design; the link's TX/RX linear tolerance was about 6mm.

25G Prototype Throughput Evaluation. Fig. 15 plots the throughput performance achieved by the 25G prototype for varying linear and angular speeds, and mixed motions. On a close examination, we observe the following:

1. For purely linear or angular motions, the link throughput remains optimal (about 23.5 Gbps) for linear speeds below 25 cm/sec or angular speeds below 25 deg/sec; these speeds are both much higher than our requirements.
2. For mixed linear and angular motions, the link throughput remained optimal for motions undergoing simultaneous linear and angular speeds of below 15 cms/sec *and* around 15-20 deg/sec (and sometimes upto 15 cms/s and 25 deg/s), which are similar to our requirements.

Compared to the 10G link, the tolerated linear speeds are lower but surprisingly the tolerated angular speeds are slightly better (likely due to the customized link design). See Table 3.

5.4 User-Traces Based Evaluation

Above, we evaluated our prototype over simulated VRH movements, as our prototype is too bulky to be worn. Here, we evaluate our prototype over available VRH user-movement traces, by simulating it based on its performance metrics.

User Traces. We use the publicly available dataset which is collected from 50 viewers watching 1-min segments from 10 360° videos with diverse characteristics from YouTube [47]. Thus, we have 500 1-min viewing traces. The traces include head location (three co-ordinates) and orientation (three angles) recorded every 10 ms.

Simulation Methodology. To simulate our 25Gbps prototype over the above user traces, we divide time into 1 ms slots. The prototype's link starts with a perfectly aligned beam. Whenever the head/VRH position is reported (roughly every 10 ms), the TP mechanism aligns the beam in 1-2 ms with a lateral and angular error of 4.54 mm and 4.54/1.75 mrad respectively, based on the TP latency and average \mathcal{G} -function errors from Table 2 for a 1.75m link. In between two position reports r and r' , the beam (i.e., the




Figure 15: Throughput for purely linear (left top), purely angular (left bottom), and arbitrary motions (right).

RX terminal) drifts laterally (angularly) at a rate of $\frac{d(r,r')}{t(r,r')}$ per ms where $d(r,r')$ is the lateral (angular) distance/difference between the two positions and $t(r,r')$ is the time lapse in ms between the reports. In any timeslot, if the total angular or lateral drift is more than the link's angular (8.73 mrad) or lateral (6 mm) tolerance (see §5.3.1), the link is marked as disconnected in that timeslot. The above simulation essentially gives a ms-level link connectivity for the user traces.




Figure 16: Cumulative Distributive Function (CDF) of the link being disconnected for a certain percentage of timeslots.

Results and Observations. We observe that our 25Gbps link prototype is operational in 98.6% of the timeslots over all the 500 traces, with the operation percentage varying from 99.98 to 95% over the 500 traces. See Fig. 16. Note that each timeslot (being 1 ms) can transmit multiple data packets on a 25Gbps link; thus, a network protocol would be able to provide an effective bandwidth of about 23Gbps (98.6% of 23.5Gbps) for the traces.

To measure user experience, we can measure how clustered/scattered the off-timeslots are, since widely scattered off-timeslots should have minimal impact on user experience. Analyzing the distribution, we observe that most of the off-timeslots (> 60% of them) occur in frames (of 30 contiguous timeslots) with less than 10 off-timeslots.

6 CONCLUSIONS

We have demonstrated the robustness of our prototype in face of expected VRH motions. As observed in §5, the system can support much higher movement speeds with the help of: (i) a VRH tracking system with higher tracking speeds frequency, or (ii) custom-built optical elements to efficiently capture a diverging beam. For higher-bandwidth (40Gbps+) links, our designed TP mechanism remains unchanged; however, the link would likely need customized collimators that can efficiently capture a *range* of wavelengths because the high-bandwidth *single-strand* transceivers use multiple wavelengths [12, 13]. Beyond the above, the key challenges or hurdles in potentially commercializing a product based on Cyclops are cost of development, miniaturization, minimizing training effort, and the economics of the final product.

ACKNOWLEDGMENTS

We would like to acknowledge Vyas Sekar for useful comments and suggestions. The work was supported by NSF Award # 1815306: NeTS: Small: A Wireless Backhaul for Multi-Gigabit Picocells Using Steerable Free Space Optics.

REFERENCES

- [1] 25G SFP28 10km (LR). <https://www.fs.com/products/87017.html>.
- [2] 25G SFP28 140km (ER). <https://www.fs.com/products/100085.html>.
- [3] Available QSFP-DD. https://zebracable.com/products/400gbase-lr8-qsfp-dd-1310nm-10km-dom-transceiver-module?variant=33387770773636¤cy=USD&utm_medium=product_sync&utm_source=google&utm_content=sag_organic&utm_campaign=sag_organic.
- [4] Cisco 10GBASE SFP+ modules. http://www.cisco.com/c/en/us/products/collateral/interfaces-modules/transceiver-modules/data_sheet_c78-455693.pdf.
- [5] Data Acquisition Device (DAQ). <https://www.mccdaq.com/usb-data-acquisition/USB-1608G.aspx>.
- [6] Extremely Reliable GMs. <https://www.cambridgetechnology.com/laser-beam-technology>.
- [7] FS.com. Personal Communication. <https://www.fs.com/>.
- [8] Intel 25G NIC. <https://www.fs.com/products/75603.html>.
- [9] Large beam GMs. <https://www.cambridgetechnology.com/products/galvanometer-scanner>.
- [10] Laser safety information. https://www.lia.org/subscriptions/safety_bulletin/laser_safety_information.
- [11] Oculus Rift S Tracking Frequency. <https://vr-compare.com/headset/oculusrifts>.
- [12] QSFP+. <https://www.broadcom.com/products/fiber-optic-modules-components/networking/optical-transceivers/qsfpplus>.
- [13] QSFP28. <https://www.broadcom.com/products/fiber-optic-modules-components/networking/optical-transceivers/qsfp28>.
- [14] SFP+ Transceivers. <https://www.fs.com/products/48814.html>.
- [15] ThorLabs Adjustable-Focus Collimators. <https://www.thorlabs.com/thorproduct.cfm?partnumber=C40FC-C>.
- [16] ThorLabs Beam Expander. <https://www.thorlabs.com/thorproduct.cfm?partnumber=BE02-05-C>.
- [17] ThorLabs High-Precision Rotation Stage. <https://www.thorlabs.com/thorproduct.cfm?partnumber=PR01>.
- [18] Vive Wireless Adapter. <https://www.vive.com/us/accessory/wireless-adapter/>.
- [19] *Safety of Laser Products – Part 1: Equipment classification, requirements and users guide, International Electrotechnical Commission (IEC) Standard 60825-1*. 2007.
- [20] Standalone VR Headsets. <https://www.nytimes.com/wirecutter/reviews/best-standalone-vr-headset/>, 2020.
- [21] Ieee standard for information technology–telecommunications and information exchange between systems local and metropolitan area networks–specific requirements part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications amendment 2: Enhanced throughput for operation in license-exempt bands above 45 ghz. *IEEE Std 802.11ay-2021 (Amendment to IEEE Std 802.11-2020 as amendment by IEEE Std 802.11ax-2021)*, pages 1–768, 2021.
- [22] Omid Abari, Dinesh Bharadia, Austin Duffield, and Dina Katabi. Enabling high-quality untethered virtual reality. In *14th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 17)*, pages 531–544, 2017.
- [23] Adjustable Aspheric Lens. <https://www.thorlabs.com/thorproduct.cfm?partnumber=CFC-2X-C&pn=CFC-2X-C#3154>.
- [24] Mohammad K. Al-Akkoumi. A tracking system for mobile FSO. In *SPIE Proceedings Vol. 6877*, 2009.
- [25] Hany Assasa, Nina Grosheva, Tanguy Ropitault, Steve Blandino, Nada Golmie, and Joerg Widmer. Implementation and evaluation of a wlan ieee 802.11 ay model in network simulator ns-3. In *Proceedings of the Workshop on ns-3*, pages 9–16, 2021.
- [26] Michael Broxton, John Flynn, Ryan Overbeck, Daniel Erickson, Peter Hedman, Matthew Duvall, Jason Dourgarian, Jay Busch, Matt Whalen, and Paul Debevec. Immersive light field video with a layered mesh representation. *ACM Transactions on Graphics (TOG)*, 39(4):86–1, 2020.
- [27] Carlos Castro, Robert Elschner, Thomas Merkle, Colja Schubert, and Ronald Freund. Experimental demonstrations of high-capacity thz-wireless transmission systems for beyond 5g. *IEEE Communications Magazine*, 58(11):41–47, 2020.
- [28] Zhi Chen, Xinying Ma, Bo Zhang, Yaxin Zhang, Zhongqian Niu, Ningyuan Kuang, Wenjie Chen, Lingxiang Li, and Shaoqian Li. A survey on terahertz communications. *China Communications*, 16(2):1–35, 2019.
- [29] Collimating Lens. <https://www.thorlabs.com/thorproduct.cfm?partnumber=F810FC-1550>.
- [30] Peter Corke. *Robotics, vision and control: fundamental algorithms in MATLAB® second, completely revised*, volume 118. Springer, 2017.
- [31] Eduardo Cuervo, Krishna Chintalapudi, and Manikanta Kotaru. Creating the perfect illusion: What will it take to create life-like virtual reality headsets? In *Proceedings of the 19th International Workshop on Mobile Computing Systems & Applications*, pages 7–12, 2018.
- [32] Max Curran, Md Shaifur Rahman, Himanshu Gupta, Kai Zheng, Jon Longtin, Samir R Das, and Thanvir Mohamed. Fsonet: A wireless backhaul for multi-gigabit picocells using steerable free space optics. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, pages 154–166, 2017.
- [33] Max Curran, Kai Zheng, Himanshu Gupta, and Jon Longtin. Handling rack vibrations in fso-based data center architectures. In *2018 International Conference on Optical Network Design and Modeling (ONDM)*, pages 47–52. IEEE, 2018.
- [34] Erbium Doped Fiber Amplifier. <https://www.fs.com/products/36501.html>.
- [35] Mike S Ferraro, William R Clark, William S Rabinovich, Rita Mahon, James L Murphy, Peter G Goetz, Linda M Thomas, Harris R Burris, Christopher I Moore, William D Waters, et al. Inalas/ingaas avalanche photodiode arrays for free space optical communication. *Applied Optics*, 54(31):F182–F188, 2015.
- [36] Galvo mirrors. https://www.thorlabs.us/newgrouppage9.cfm?objectgroup_id=6057.
- [37] Yasaman Ghasempour, Claudio RCM Da Silva, Carlos Cordeiro, and Edward W Knightly. Ieee 802.11 ay: Next-generation 60 ghz communication for 100 gb/s wi-fi. *IEEE Communications Magazine*, 55(12):186–192, 2017.
- [38] Monia Ghobadi, Ratul Mahajan, Amar Phanishayee, Nikhil Devanur, Janardhan Kulkarni, Gireeja Ranade, Pierre-Alexandre Blanche, Houman Rastegarfar, Madeleine Glick, and Daniel Kilper. Projector: Agile reconfigurable data center interconnect. In *Proceedings of the 2016 ACM SIGCOMM Conference*, pages 216–229, 2016.
- [39] Navid Hamedazimi, Himanshu Gupta, Vyas Sekar, and Samir R Das. Patch panels in the sky: A case for free-space optics in data centers. In *Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks*, pages 1–7, 2013.
- [40] Navid Hamedazimi, Zafar Qazi, Himanshu Gupta, Vyas Sekar, Samir R Das, Jon P Longtin, Himanshu Shah, and Ashish Tanwer. Firefly: A reconfigurable wireless data center fabric using free-space optics. In *Proceedings of the 2014 ACM conference on SIGCOMM*, pages 319–330, 2014.
- [41] Tzung-Hsien Ho. *Pointing, Acquisition, and Tracking Systems for Free-Space Optical Communication Links*. PhD thesis, University of Maryland, College Park, 2007.
- [42] iPerf. <https://iperf.fr/>.
- [43] Stamatios V. Kartalopoulos. *Free Space Optical Networks for Ultra-Broad Band Services*. John Wiley and Sons, 2001.
- [44] Mahmudur Khan and Jacob Chakareski. Visible light communication for next generation untethered virtual reality systems. In *2019 IEEE International Conference on Communications Workshops (ICC Workshops)*, pages 1–6. IEEE, 2019.
- [45] Zeqi Lai, Y Charlie Hu, Yong Cui, Linhui Sun, Ningwei Dai, and Hung-Sheng Lee. Furion: Engineering high-quality immersive virtual reality on today’s mobile devices. *IEEE Transactions on Mobile Computing*, 19(7):1586–1602, 2019.
- [46] Luyang Liu, Ruiguang Zhong, Wuyang Zhang, Yunxin Liu, Jiansong Zhang, Lintao Zhang, and Marco Gruteser. Cutting the cord: Designing a high-quality untethered vr system with low latency remote rendering. In *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services*, pages 68–80, 2018.
- [47] Wen-Chih Lo, Ching-Ling Fan, Jean Lee, Chun-Ying Huang, Kuan-Ta Chen, and Cheng-Hsin Hsu. 360 video viewing dataset in head-mounted virtual reality. In *Proceedings of the 8th ACM on Multimedia Systems Conference*, pages 211–216, 2017.
- [48] ChunLei Lv, HuiLin Jiang, and ShouFeng Tong. Implementation of fta with high bandwidth and tracking accuracy in fso. In *2012 2nd International Conference on Consumer Electronics, Communications and Networks (CECNet)*, pages 3275–3279. IEEE, 2012.
- [49] Arun K. Majumdar. *Advanced Free Space Optics (FSO): A Systems Approach*. Springer New York, 2015.
- [50] Jiayi Meng, Sibendu Paul, and Y Charlie Hu. Coterie: Exploiting frame similarity to enable high-quality multiplayer vr on commodity mobile devices. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 923–937, 2020.
- [51] D Messier. Dlr researchers set world record in free-space optical communications. Available on: <http://www.parabolicarc.com/2016/11/05/dlrresearchers-set-world-record-freespace-optical-communications>, 2016.
- [52] Oculus Rift S. <https://www.oculus.com/rift-s/>.
- [53] Ryan S Overbeck, Daniel Erickson, Daniel Evangelakos, and Paul Debevec. Welcome to light fields. In *ACM SIGGRAPH 2018 Virtual, Augmented, and Mixed Reality*, pages 1–1. 2018.
- [54] QSFP-DD. Accelerating 400gbe adoption with qsfp-dd. <https://tinyurl.com/rnjsrx2>, March 2017.
- [55] Md Shaifur Rahman, Kai Zheng, and Himanshu Gupta. FSO-VR: steerable free space optics link for virtual reality headsets. In *Proceedings of the 4th ACM Workshop on Wearable Systems and Applications*, pages 11–15, 2018.
- [56] Swetank Kumar Saha, Yasaman Ghasempour, Muhammad Kumail Haider, Tariq Siddiqui, Paulo De Melo, Neerad Somanchi, Luke Zakrajsek, Arjun Singh, Roshan Shyamsunder, Owen Torres, et al. X60: A programmable testbed for wideband 60 ghz wlans with phased arrays. *Computer Communications*, 133:77–88, 2019.
- [57] SciPy Optimize. <https://docs.scipy.org/doc/scipy/reference/optimize.html>.
- [58] Jun Xie, Shuhuai Huang, Zhengcheng Duan, Yusheng Shi, and Shifeng Wen. Correction of the image distortion for laser galvanometric scanning system. *Optics & Laser Technology*, 37(4):305–311, 2005.
- [59] Toshiaki Yamashita, Masaki Morita, Motoaki Shimizu, Daisuke Eto, Koichi Shiratama, and Shigeru Murata. The new tracking control system for free-space

optical communications. In *2011 International Conference on Space Optical Systems and Applications (ICSOS)*, pages 122–131. IEEE, 2011.

- [60] Ruiguang Zhong, Manni Wang, Zijian Chen, Luyang Liu, Yunxin Liu, Jiansong Zhang, Lintao Zhang, and Thomas Moscibroda. On building a programmable wireless high-quality virtual reality system using commodity hardware. In *Proceedings of the 8th Asia-Pacific Workshop on Systems*, pages 1–7, 2017.

Appendices are supporting material that has not been peer-reviewed.

A HARDWARE COMPONENTS USED IN THE PROTOTYPE

Below is the list of main optical components used in our Cyclops prototype.

- (1) SFPs: We used Cisco SFP-10G-ZR100 1550nm transceivers [14] for Fiberstore Inc. The SFPs are connected to standard 10G PCI cards placed in desktops.¹⁴
- (2) On the transmitter (TX) side, the above SFP+ was connected to a single-mode optical cable, while on the receiver (RX) side the SFP+ was connected to a multi-mode optical cable with 50μ core diameter.

- (3) At TX, the optical cable from SFP connected to an amplifier [34] which connected to a collimator. At RX, the SFP connected directly to the collimator.
- (4) Collimators: For the collimated wider beam, we used beam-expanders BE02-05-C [16] from ThorLabs at both ends, and for the diverging beam, we used the adjustable collimator CFC-2X-C [23] at the TX and the standard collimator F810FC-1550 [29] at the RX.
- (5) The collimators at both ends were positioned and pointed towards the Galvo Systems. We used the 2-Axis Large Beam Diameter Scanning Galvo System GVS102 [36] from ThorLabs. Each GM is connected to its power sink unit (PSU) and servo controllers. To control voltage supplied to a GM, the GM's PSU is connected to a USB Data Acquisition Device (DAQ) [5] which is connected to a PC via USB.
- (6) The RX assembly (GM and collimator) along with the VR headset is affixed on a breadboard to ensure a rigid assembly. For the pure linear and angular motions, the breadboard is attached to an optical rotating platform (e.g., [17]) which is affixed to a linear rail.

¹⁴For the 25Gbps link, we used the 25G SFPs [2] connected to the adjustable-focus collimators C40FC-C [15] and the 25G PCI cards [8].