

Design and validation of a portable modular free-space quantum optical terminal towards mobile deployment over 20 km

Hamish Johnson^a, Kimia Mohammadi^b, Wilson Wu^a, Howard Schimmer^a, Indira Sonea^a,
Sungeun Oh^a, Sasan V. Grayli^{ac}, and Thomas Jennewein^{ab}

^aDepartment of Physics, Simon Fraser University, Burnaby, Canada

^bInstitute for Quantum Computing, University of Waterloo, Waterloo, ON, Canada

^cInstitute for Quantum Computing and Department of Electrical and Computer Engineering,
University of Waterloo, Waterloo, ON, Canada

ABSTRACT

We report the design and coarse-pointing field validation of a portable, SWaP-constrained free-space optical terminal for inter-regional quantum networking. The terminal is designed for a 20.44 km mixed urban-maritime link between Simon Fraser University and Cypress Mountain, British Columbia, Canada. This horizontal path lies entirely within the atmospheric boundary layer, providing a stringent terrestrial test for hardware intended for subsequent mobile and high-altitude-platform deployment. A mean power-in-bucket loss model, informed by preliminary turbulence measurements over the path, is used to derive terminal requirements. The payload is designed within a 200 mm, sub-7 kg envelope, compatible with an off-the-shelf drone gimbal, and operates across 780–900 nm.

We describe the optical design, embedded control, two-stage acquisition, pointing and tracking (APT) architecture, and a fixed pupil-plane interface allowing fine-steering hardware to be exchanged without redefining the surrounding optics. The implemented coarse stage is characterised in the laboratory and over the full link, obtaining 99th-percentile residual pointing errors of 70 μ rad and 80–250 μ rad, respectively. Preliminary atmospheric-coherence measurements indicate that higher-order wavefront correction will be required for the proposed ground receiver. The terminal provides a modular testbed for terrestrial, mobile, and airborne free-space quantum links within the Vancouver Quantum Network.

Keywords: free-space quantum communications, quantum networking, mobile optical terminals, high-altitude platforms, acquisition pointing and tracking, horizontal free-space links, atmospheric turbulence, drone-based quantum communications

1. INTRODUCTION

Quantum technologies have progressed from isolated experiments towards systems in which physically separated processors, sensors, and users must exchange quantum states. Distributing entanglement between distant nodes provides the non-classical correlation required for quantum teleportation, entanglement-based quantum key distribution (QKD), distributed quantum sensing, and delocalised quantum processing.¹ QKD provides one near-term networking application, using quantum states encoded in photons to establish cryptographic keys whose security can be bounded by the measured channel statistics.² More generally, entanglement distribution provides the physical resource required to connect remote quantum processors and sensors.³ A useful quantum network must therefore provide an optical transport layer capable of sending these quantum states between nodes with low loss in potentially high-noise environments, while preserving spatial and temporal photonic modes.⁴

The resulting network vision is multiscale, as illustrated in Fig. 1. Fibre links are well suited to local and metropolitan networks, and enable increasingly sophisticated multi-node demonstrations.^{5,6} Recent demonstrations have extended towards dynamically reconfigurable networks, including distribution among multiple users

Further author information: (Send correspondence to H.J.)

H.J.: E-mail: hamishj@sfu.ca

over metropolitan fibre.^{7–9} However, even state-of-the-art low-loss fibre accumulates approximately 0.1 dB/km loss,¹⁰ attenuating signal exponentially with distance, while an unknown quantum state being amplified is prohibited by the no-cloning theorem.¹¹ Quantum repeaters offer one solution, but memory-based repeater networks remain an experimental challenge.^{6,12} Free-space optics (FSO) links provide a complementary approach, connecting terrestrial, airborne, and orbital nodes over progressively larger spatial scales.^{13–15}




Figure 1. Representative scales of free-space quantum network links. (A) Multiscale optical networking between terrestrial, airborne, and satellite nodes; the solid red line represents this work, and the dashed lines represent a vision for reconfigurable networking. (B) Line of sight of the 20.44 km SFU–Cypress free-space link studied in this work.

Free-space quantum communications have progressed from the laboratory to increasingly demanding field demonstrations. Long-distance horizontal terrestrial channels have demonstrated QKD over tens of kilometres, including a 53 km daylight link using coarse and fine tracking over a high-altitude plateau,¹³ while metropolitan experiments have demonstrated wavefront correction under daytime atmospheric conditions.¹⁶ Airborne QKD has been demonstrated with moving aircraft over kilometre-scale paths;¹⁴ more recent vehicle and drone-mounted terminals have demonstrated two-stage acquisition, pointing, and tracking (APT),^{17,18} and are actively under development. At larger scales, satellite experiments have distributed entanglement between ground stations over distances exceeding 1000 km¹⁵ and quantum keys over 10 000 km.¹⁹

Together, these demonstrations establish many of the technologies required for free-space quantum networking, but leave an important systems gap:

portable optical terminals that combine atmospheric reach over tens of kilometres (inter-regional), mobile-platform size, weight, and power (SWaP) constraints, and robust APT.

This gap is relevant to mobile ground terminals and high-altitude platforms (HAPs), which offer persistent or reconfigurable free-space connections between metropolitan and satellite-scale networks. Recent work proposes stratospheric balloon relays as components of long-range quantum network architectures²⁰ and hybrid satellite–HAP links.²¹ Compared with a horizontal terrestrial path, airborne slant paths can reduce the path-integrated turbulence from the atmospheric boundary layer (ABL). Unlike conventional FSO terminals, a quantum terminal cannot compensate channel loss through optical amplification, must recover single photons against background, and may require single-mode spatial coupling at the receiver.^{13,22}

Within the Vancouver Quantum Network (QVaN), we are developing a free-space communications layer to connect metropolitan quantum nodes with future mobile, airborne, and satellite infrastructure.²³ The present work targets a 20.44 km mixed urban–maritime link between Simon Fraser University (SFU) and Cypress Mountain, British Columbia, Canada. Because the path remains within the ABL, it provides a severe turbulence stress test²⁴ for the development of hardware intended for mobile and airborne deployment. The transmitter is therefore designed as a portable, SWaP-constrained optical terminal rather than a fixed observatory-scale system. It operates across a 780–900 nm band and provides a fixed, modular interface for interchangeable fine-steering hardware.

We report the link-analysis-derived requirements, optical and electronic implementation, firmware architecture, and two-stage APT architecture of the transmitter. We then characterise the implemented coarse pointing stage with a fixed mirror at the fine-steering pupil, first in the laboratory and then over the SFU–Cypress link. Section 2 develops the atmospheric propagation model and terminal design requirements; Sec. 3 describes the terminal implementation; and Sec. 4 presents laboratory and field validation together with the resulting fine-steering requirements.

2. LINK ANALYSIS & DESIGN REQUIREMENTS

2.1 Mean Atmospheric Propagation Efficiency

Horizontal FSO within the atmospheric boundary layer is affected by extinction, turbulence-induced beam spreading, scintillation, and beam wander.^{24,25} We first consider the ensemble-mean optical power delivered to the receiver aperture. This establishes the mean-efficiency envelope used to design the transmitter; instantaneous fading and receiver single-mode coupling are discussed separately in Secs. 2.3 and 2.2, respectively.

For a link distance L , the mean power-in-bucket (PiB) efficiency at the receiver aperture is written as

$$\bar{\eta}_{\text{PiB}} = \eta_{\text{Tx}} \eta_{\text{atm}}(V, \lambda) \eta_{\text{geo}}(C_n^2, D_{\text{Tx}}, D_{\text{Rx}}, L, \lambda), \quad (1)$$

where η_{Tx} accounts for transmitter efficiency and aperture truncation, η_{atm} describes atmospheric extinction, and η_{geo} is the fraction of the long-term propagated beam overlapped by the receiver aperture. Equation 1 describes an aperture-plane mean loss model, and it does not predict instantaneous channel transmittance or single-mode fibre (SMF) coupling at the receiver.

The transmitted field is approximated as Gaussian with $1/e^2$ intensity radius w_0 at the transmitter aperture. For a circular clear aperture $D_{\text{Tx}} = 2a_{\text{Tx}}$, the transmitted fraction is

$$\eta_{\text{Tx}} = \eta_{\text{opt,Tx}} \left[1 - \exp \left(-\frac{2a_{\text{Tx}}^2}{w_0^2} \right) \right], \quad (2)$$

where $\eta_{\text{opt,Tx}}$ is the transmission efficiency of the optical train. A nominal Gaussian fill factor $w_0/a_{\text{Tx}} \simeq 0.89$ is used as the design reference, providing a near-optimal balance between truncation and far-field divergence.²⁶

The transmitter optics use near-infrared anti-reflection coatings covering 780–900 nm. The nominal optical train throughput estimated from vendor data is $\eta_{\text{opt,Tx}} = 96.9\%$ at 785 nm, 96.5% at 810 nm, and 95.0% at 890 nm. The band accommodates the free-space wavelengths targeted by the QVaN testbed and overlaps wavelengths previously used for free-space quantum communication.^{27,28}

Atmospheric extinction is separated into visibility-dependent scattering and molecular absorption, $\eta_{\text{atm}}(V, \lambda) = \eta_{\text{vis}}(V, \lambda) \eta_{\text{abs}}(\lambda)$. The visibility term is evaluated using the Kim model,

$$\eta_{\text{vis}} = \exp \left[-\frac{3.91}{V} \left(\frac{\lambda}{550} \right)^{-q(V)} \times L \right], \quad (3)$$

where V and L are expressed in kilometres and λ in nanometres, and $q(V)$ is an empirical wavelength exponent of the model.²⁹ Hourly visibility observations from the Environment and Climate Change Canada Meteorological Service of Canada station at Vancouver International Airport (Climate ID 1108395) were analysed over 2013–2026.³⁰ The annual mean visibility was (32 ± 3) km, where uncertainty denotes year-to-year variation. The seasonal 5th percentile visibility was lowest in winter, (4 ± 3) km, and highest in summer, (19 ± 4) km. The upper limit is reporting-limited at 48.3 km (30 mi). These aeronautical observations are biased by local coastal weather, and are therefore used as an environmental proxy rather than a site-specific measurement. Molecular absorption, $\eta_{\text{abs}}(\lambda)$, is calculated from HITRAN spectroscopic data using HAPI.³¹ Atmospheric loss over the 20.44 km channel is shown in Fig. 2.

In the absence of turbulence, the Gaussian beam radius is $w(L) = w_0[1 + (L/z_R)^2]^{1/2}$, where $z_R = \pi w_0^2/\lambda$.³² For ensemble-averaged broadening under turbulence, we use the turbulence-broadened beam radius approximation,²⁵

$$w_{\text{LT}} = w(L) \sqrt{1 + 1.63 (\sigma_R^2)^{6/5} \Lambda}, \quad \text{where} \quad \Lambda = \frac{\lambda L}{\pi w^2(L)}. \quad (4)$$




Figure 2. Modelled atmospheric transmission over the SFU-Cypress link, combining HITRAN molecular absorption with the Kim visibility model (dashed). The shaded band denotes the operating band, with dashed lines for 785 nm, 810 nm, and 890 nm.

Here $\sigma_R^2 = 1.23 C_n^2 k^{7/6} L^{11/6}$ is the plane-wave Rytov variance, with $k = 2\pi/\lambda$.³² Approximating the long-term irradiance profile as Gaussian, the mean geometric collection efficiency is

$$\eta_{\text{geo}} = 1 - \exp \left[-\frac{D_{\text{Rx}}^2}{2w_{\text{LT}}^2} \right]. \quad (5)$$

Figure 3 illustrates the aperture-efficiency trade-off over the design space. Without turbulence, increasing D_{Tx} reduces diffraction until the receiver captures most of the beam. With increased turbulence, long-term beam broadening reduces the benefit of a larger transmitter aperture. A $D_{\text{Tx}} = 50.8$ mm is chosen as a SWaP-constrained baseline approaching the region of diminishing returns.




Figure 3. Mean receiver-aperture loss landscape at 810 nm over the 20.44 km horizontal link under (A) no turbulence and (B) median path-equivalent turbulence; atmospheric extinction is excluded from these panels. Selected terminal geometry is represented by a white diamond. (C) Mean geometric loss over the broader turbulence sweep for the selected geometry; the shaded region represents the preliminary 5th–95th percentile range. (D) Empirical cumulative distribution function (CDF) of the effective r_0 from 53 non-overlapping 6 s integrations.

2.2 Receiver Coupling Requirements

Efficient SMF coupling additionally requires spatial overlap and sufficient phase coherence across the receiver pupil.³³ We therefore treat $\bar{\eta}_{\text{pIB}}$ as a receiver-aperture quantity and consider receiver coupling separately.

The atmospheric coherence is characterised by the Fried parameter r_0 . For a constant C_n^2 channel under the spherical-wave model used here,³²

$$r_0 = (0.16k^2C_n^2L)^{-3/5}. \quad (6)$$

Preliminary beacon measurements over the SFU–Cypress link provide a quantitative estimate of atmospheric coherence. A 680 nm beacon LED was observed through an 80 mm aperture telescope with a 40 nm bandpass filter at 670 nm. Eleven post-sunset measurements acquired on 12 August 2026 provided 53 non-overlapping 6 s integrations. Effective r_0 was inferred from the long-exposure beacon width using the Kolmogorov seeing approximation, $\text{FWHM}_{\text{LE}} \approx 0.98\lambda/r_0$.³⁴ At 680 nm, the 5th, median, and 95th percentiles were 25.5 mm, 32.4 mm, and 38.4 mm, shown in Fig. 3(D). These correspond under Eq. 6 to a path-equivalent C_n^2 range of 8.2×10^{-16} – $1.62 \times 10^{-15} \text{ m}^{-2/3}$, with a median of $1.09 \times 10^{-15} \text{ m}^{-2/3}$. The measurements span one evening; static aberration and residual image motion were not independently calibrated, potentially biasing the inferred r_0 low. Scaling $r_0 \propto \lambda^{6/5}$ gives $D_{\text{Rx}}/r_0 \approx 4.2$ – 6.4 at 810 nm for the proposed $D_{\text{Rx}} = 200$ mm receiver. We therefore identify higher-order receiver adaptive optics (AO) as a design requirement for this link. For the coupling efficiency upper bound, we use $\eta_{c,0} \approx 0.78$,³⁵ the maximum overlap of a circular pupil with a Gaussian fundamental mode.

2.3 Fading Channel Effects

Atmospheric turbulence produces a time-dependent transmittance through wander, scintillation, and aberration. For QKD, a mean-loss budget alone is insufficient to determine channel performance.⁴ We use centroid dropout as the operational measured quantity for APT robustness to intermittent target loss, including channel fading. During on-site tests, centroid dropouts had a median duration of 5 ms, with 99th-percentile dropout duration varying from 35 ms to 40 ms between trials. These measurements were obtained from a 75 mm, $f/2.8$ lens assembly and represent an APT robustness requirement.

2.4 Mean Link Budget and Derived Terminal Requirements

For the selected terminal geometry, the mean aperture-plane loss,

$$\bar{L}_{\text{PiB}} = -10 \log_{10} \bar{\eta}_{\text{PiB}}, \quad (7)$$

spans 8–43 dB across the design sweep ($C_n^2 \in [10^{-16}, 10^{-14}] \text{ m}^{-2/3}$) and Vancouver visibility envelope ($V \in [19, 100] \text{ km}$), represented in Fig. 3. This budget excludes stochastic fading, receiver AO residuals, and other receiver losses. It is therefore interpreted as a lower bound on the end-to-end system loss for the purpose of establishing optical terminal requirements.

Table 1. Design parameters and derived requirements for optics and fine-steering mirror (FSM) from link modelling and coarse tracking results.

Requirement	Value	Driver
Tx aperture (D_{Tx})	50.8 mm	Mean PiB–SWaP trade-off.
Beam $1/e^2$ diameter	43.5 mm	Aperture-fill target.
Operating band	780–900 nm	Atmospheric window, source heritage. ²⁸
Rx aperture	200 mm	Assumed receiver baseline.
Rx AO	High-order	Measured $D_{\text{Rx}}/r_0 \approx 4.2$ – 6.4 .
Coarse residual 99 th percentile	$(250 \pm 40) \text{ } \mu\text{rad}$	Fine-steering capture.
FSM mechanical half-range (S_{FSM})	$\geq (0.56 \pm 0.09) \text{ mrad}$	Coarse residual and FSM gain.
Fine-steering residual target	$\lesssim 12 \text{ } \mu\text{rad}$	Long-term beam quarter-angle (Sec. 3.2.1)
Fine-steering bandwidth (f_{FSM})	50 Hz	On-site coarse reverse-cumulative RMS.
Dropout tolerance (τ_{drop})	40 ms	Centroid-dropout 99 th percentile.

3. SYSTEM DESIGN & METHODS

3.1 Platform Integration and SWaP Requirements

The FSO transmitter is integrated on a Freefly MōVI Pro three-axis gimbal. Optical tracking provides coarse pointing in pan and tilt, while the gimbal’s internal stabilisation loop rejects higher-frequency disturbances. The gimbal additionally provides the mechanical mounting and electrical interfaces for the optical payload. The tilt stage unit (TSU) carries the optical transmitter, coarse-tracking camera, beacon, and the control electronics. As summarised in Fig. 4, image acquisition and memory-intensive processing are performed by the single-board computer (SBC), while deterministic closed-loop control is implemented on the microcontroller (MCU). The quantum signal and AO reference light are combined in fibre upstream of the transmitter.

The gimbal platform constrains the transmitter to a compact SWaP envelope. The TSU provides a $176 \times 190 \times 200$ mm volume, a payload limit of 6.8 kg, and supplies up to 90 W from a 172.8 Wh battery. These values are treated as platform design constraints rather than measured SWaP values.




Figure 4. System architecture of the FSO transmitter. The gimbal provides coarse pan-tilt actuation and stabilisation, as well as the mechanical and electrical interface to the TSU. The TSU holds the optical transmitter, MCU, SBC, tracking camera and beacon. The dashed fine-steering loop denotes the reserved actuator and position-sensor interface, not enabled in this work. Here PSD refers to a position-sensitive detector.

3.2 Optical Design and Modular Optomechanical Interface

3.2.1 Optical Design

The payload contains two optical tracks: the quantum signal transmitter and the coarse tracking receiver. The transmitter comprises a fibre-collimation stage, a fine-steering pupil plane (PP), and a beam expander. For the field tests reported here, the fine-steering actuator was replaced by a fixed protected-silver mirror to qualify terminal design and establish fine-steering requirements. Coarse tracking uses a 75 mm focal-length Ricoh (FL-BC7528-9M) lens and a Basler Ace 2 monochrome camera (a2A2464-77umPRO). The prototype optical assembly and optical path are shown in Fig. 5(A).

The 780–900 nm quantum signal is launched from a 780-HP SMF by a 45 mm focal-length achromatic doublet. The implemented prototype uses a 50/180 mm refractive telescope. A Zemax model using the vendor lenses and a $5 \mu\text{m}$ source mode-field diameter (MFD) at 810 nm gives a $1/e^2$ collimated beam diameter of approximately 9.3 mm, and an effective magnification of $3.59\times$, producing an output beam of approximately 33.3 mm. NIR-coated achromatic doublets support performance across the quantum band, though wavelength-dependent focus adjustments are required using axial translation stages. The prototype is assembled using a commercial off-the-shelf cage system and kinematic optomechanics on a 150×150 mm breadboard.

The requirement-driven machined design, Fig. 5(B), instead uses a 45/200 mm telescope, targeting a 43.5 mm output beam, and includes a fine-steering mirror as well as a separate optical track behind a dichroic for the




Figure 5. (A) Prototype optical assembly with tracking camera, collimated LED beacon, transmitter telescope, and fibre patch cable input coupler; corresponding optical layout in Zemax. (B) Proposed version two optical transmitter assembly developed under the design requirements presented in this work. The pupil plane (PP), beam splitter (BS), band-pass (BP) filter, and long-pass (LP) filter are annotated.

fine-steering sensor. This output beam corresponds to $w_0/a_{Tx} \simeq 0.856$, close to the target noted in Sec. 2.1. For 810 nm at the preliminary median turbulence, $w_{LT} \approx 0.47$ m, corresponding to a long-term $1/e^2$ beam half-angle $\theta_{LT} = w_{LT}/L \approx 23$ μ rad.

3.2.2 Modular Optomechanical Interface

Modular optical-bench architectures for laser terminals have previously reserved fold-mirror sites at a telescope pupil, so that a fine-steering mirror can be substituted without redesigning the surrounding optics.³⁶ To decouple fine-steering development from the transmitter’s optical layout, the fine-steering stage is defined by a fixed optical and mechanical interface. The interface defines a 9.3 mm $1/e^2$ beam diameter, pupil position, and input and output optical axes. Alternative fine-steering actuators and sensors can therefore be evaluated without redefining the surrounding transmitter optics. The steering element is located at the internal pupil plane noted in Fig. 5(B), minimising lateral beam walk across the telescope during steering. The prototype uses a fixed mirror at this location. This machined version two provides an interchangeable actuator stage located by machined dowel pins.

3.3 Electronic and Control Hardware

The control architecture separates on-arrival image processing from real-time control. A Raspberry Pi 5 SBC acquires frames from the tracking camera and performs target segmentation and centroid estimation. Centroid and status messages are transmitted over TTL UART to an Infineon PSoC 5 LP MCU, which executes the closed-loop controller and sends pan-tilt rate commands to the gimbal control unit (GCU). The GCU closes the manufacturer-provided inner stabilisation loops for each gimbal motor. The prototype electronics were assembled from stacked prototyping boards in a custom acrylonitrile butadiene styrene (ABS) frame, shown in Fig. 6; an integrated printed circuit board is in development.

This separation allows the computationally intensive image processing and logging to remain on the SBC, while the MCU maintains deterministic actuator updates. It also lends itself to a hardware abstraction layer for the MCU to interoperate with different sensors and actuators. The MCU provides analogue acquisition channels reserved for fine-steering position sensors. System telemetry from the image processing pipelines, MCU, and GCU sensors is timestamped and logged by the SBC for the characterisation described in Sec. 3.6.

3.4 Firmware Architecture

System operation is coordinated by the MCU as a hierarchical finite-state machine (HFSM), shown in Fig. 6(B). The HFSM separates standby, acquisition, tracking, loss, and error states, and supports user-triggered or autonomous operation. Following initial coarse acquisition, valid centroid packets received from the SBC enable closed-loop tracking. A stale-data timer is accumulated when centroid packets are interrupted, slowing the gimbal to a stationary hold. If valid centroid packets resume, tracking continues; otherwise, the system transitions

to the loss state. Error conditions cause a transition to a safe stationary state. The HFSM reserves an additional fine-tracking step at a separate clock rate from that of the closed-loop control to support sensing and PID control for subsequent fine-steering integration. This stage is not enabled in the prototype outlined in this work.




Figure 6. (A) Prototype electronics integration. (B) MCU hierarchical finite-state machine controlling standby, acquisition, tracking, loss, and errors. User-triggered and automatic transitions are distinguished by blue and orange arrows, respectively.

3.5 Acquisition, Pointing, and Tracking Control Architecture

The terminal uses two-stage APT. The coarse stage uses the gimbal to acquire the beacon, reject platform motion, and correct low-frequency pointing error. The future fine-steering stage will reject higher-frequency residuals. The measured coarse residual magnitude and spectrum therefore constrain the required stroke and bandwidth of the fine-steering stage.

Coarse tracking uses a collimated beacon at 680 nm imaged onto the camera. The beacon has a 48 mm clear aperture and a measured half-angle divergence of $(0.84 \pm 0.08)^\circ$. Image coordinates are converted to angular pointing error using the plate scale of the tracking optics. Image processing is divided into acquisition and tracking substates. During acquisition, frames are spatially binned and background subtracted before segmentation into targets. Once the selected target is centred and stable, the system transitions to tracking. The tracking mode uses an unbinned region of interest and the previous target state to seed sub-pixel centroid estimation, supporting frame rates up to 900 Hz, though typically limited by exposure time.

Within the MCU, pan and tilt are controlled by independent, discrete-time PID controllers. Angular centroid error is the controller input, and angular rate is the command output. Commands are issued to the GCU at a fixed 20 Hz. Controller gains were selected experimentally in the laboratory.

3.6 Pointing Characterisation Methods

The system was characterised in a climate-controlled laboratory and validated outdoors over the SFU–Cypress link under sunset and dusk conditions. Pointing performance is evaluated from timestamped centroid measurements rather than lower-frequency actuator telemetry. Laboratory measurements were acquired at approximately 500 Hz (100 μ s exposure), while the on-site validation required 200 Hz (2000 μ s exposure). In both setups, the coarse pointing controller operated at a fixed 20 Hz. Gimbal attitude, controller output, and controller telemetry were logged at the control rate.

Residual centroid error (θ_x, θ_y) is converted to angular error relative to the commanded setpoint in microradians (e_x, e_y) , defining

$$e_x(t) = \theta_x(t) - \theta_{x,0}, \quad e_y(t) = \theta_y(t) - \theta_{y,0}, \quad r(t) = \sqrt{e_x^2(t) + e_y^2(t)}, \quad (8)$$

where r is the radial pointing error. Per-axis root-mean-square (RMS) error is calculated directly from samples, quantifying jitter since the residual mean bias is negligible. Empirical radial percentiles ($R_{\%}$) quantify the dynamic capture requirement. Uncertainties on RMS pointing error are estimated from the variability of the second moment across four equal-duration subrecords, and on R_{99} using a delete-one-quarter block jackknife.

The temporal content of the coarse tracking is characterised with per-axis power spectral density (PSD), $S_{\theta}(f)$. Spectra are estimated using Welch’s method³⁷ with a Hann window and 50 % overlap. Centroid sampling is substantially faster than the coarse control rate so that residual motion above it remains observable.

To quantify residual motion as a fine-steering bandwidth requirement, the PSD is integrated from f to the Nyquist frequency f_{Nyq} ,

$$\sigma_{>f} = \left[\int_f^{f_{\text{Nyq}}} S_{\theta}(f') df' \right]^{1/2}. \quad (9)$$

This reverse-cumulative RMS represents the residual RMS contribution above f , used as an ideal disturbance-rejection benchmark. This does not include delay or sensor noise of a realised fine-steering loop.

Centroid dropout is defined as a frame in which no valid target is returned. The dropout fraction is $N_{\text{invalid}}/(N_{\text{valid}} + N_{\text{invalid}})$, and consecutive invalid frames are accumulated to estimate duration. Invalid samples are excluded from pointing-error analysis. For spectral analysis, the camera time base is retained, and missing samples are interpolated for PSD estimation. This was validated by simulating dropouts in laboratory data, where the resulting bias was less than 1 %. The centroid-pipeline sensing floor is estimated from centroid motion measured with the beacon and camera fixed mechanically to the same optical table, two metres apart.

4. RESULTS & DISCUSSION

4.1 Laboratory Coarse Tracking Performance

The coarse tracking stage was first characterised under laboratory conditions with a stationary beacon. For all measurements, the prototype transmitter contained a fixed mirror at the pupil plane, such that the measured residual represents the disturbance handed to the fine-steering stage. After acquisition and settling, the system maintained closed-loop tracking for 10–15 min trials with centroid measurements acquired at approximately 500 Hz and coarse control updates at 20 Hz.




Figure 7. Representative 12.3 min laboratory coarse tracking trial. (A) Residual pan and tilt centroid error between start and end of tracking (diamonds) with a two-second inset. (B) Joint residual-error distribution with marginal histograms and RMS ellipse. Unequal pan and tilt widths show anisotropic residuals.

Figure 7 shows a representative 12.3 min trial with acquisition, steady-state tracking, and release. The residual remained bounded about the setpoint, with per-axis RMS jitter of $(26.1 \pm 0.6) \mu\text{rad}$ and $(13.8 \pm 0.7) \mu\text{rad}$ in pan and tilt, respectively. The larger pan residual and resulting anisotropic error distribution exceed the centroid-sensing floor, consistent with the manufacturer’s internal stabilisation drift, rather than centroid estimation.

The corresponding power spectral densities are shown in Fig. 8(A). Approximately 92.4 % of the pan variance is below 5 Hz, after which the spectra converge to a similar high-frequency floor. Within 44–200 Hz, narrow spectral peaks are observed, which persist when motors are disabled, and are therefore not produced by MCU-commanded gimbal actuation. This convergence is also apparent in the reverse-cumulative RMS (Fig. 8(B)). Increasing ideal fine-steering correction bandwidth would first remove predominantly pan variance, then the remaining higher-frequency residual. These measurements establish the laboratory coarse-pointing baseline of the integrated terminal. The following section evaluates how its performance changes when tracking a beacon propagated over the full 20.44 km atmospheric channel.




Figure 8. Spectral performance of the laboratory tracking trial. (A) Per-axis PSD. Narrow features between approximately 44 and 200 Hz persist with motors disabled. Dashed curves show the PSD of the beacon centroid with the beacon and camera fixed mechanically to the same optical table, showing an increased noise floor due to a larger beacon spot size in the setup. (B) Reverse-cumulative RMS residual. Pan and tilt convergence near 6.6 Hz indicates that excess pan residual is concentrated at low frequency, while the high-frequency floor is similar between axes.

4.2 SFU–Cypress Field Validation

The coarse-pointing architecture was validated over the full 20.44 km SFU–Cypress horizontal link. On 12 August 2026, a 680 nm beacon transmitted from the SFU Bennett Library was tracked by the optical terminal at the Cypress Mountain viewpoint. The longer exposure required for the field channel reduced the centroid rate from 500 Hz in the lab to 200 Hz, while the controller was fixed at 20 Hz.

Four approximately 10 min trials were acquired after sunset, using a 40 nm bandpass filter centred at 670 nm. The beacon was detectable under daylight conditions; however, specular reflections from the urban background produced competing targets in the tracking camera, motivating narrower spectral filtering.

Figure 9(A) shows the residual pointing distribution for a representative field trial. Pan remained comparatively consistent between runs, as summarised in Table 2, with residual jitter of 24–30 μrad , while the tilt residual increased to 23–67 μrad with greater variability. Similarly, the radial 99th percentile increased to 80–250 μrad . The field distributions are visibly non-Gaussian, with elongated tails along the tilt axis.

Table 2. Results of on-site SFU–Cypress coarse tracking over the proposed quantum link on 12 August 2026.

Start time (PDT)	20:51	21:03	21:13	21:24
Duration (s)	633	551	595	641
Valid centroid fraction (%)	88.5	89.1	90.8	91.3
RMS x (μrad)	(29.8 ± 1.3)	(27.3 ± 0.5)	(24.3 ± 0.5)	(25.8 ± 1.0)
RMS y (μrad)	(67 ± 10)	(66 ± 7)	(23.0 ± 1.7)	(47 ± 7)
RMS radial (μrad)	(74 ± 9)	(71 ± 6)	(33.5 ± 1.5)	(53 ± 7)
Radial R_{99} (μrad)	(250 ± 40)	(220 ± 50)	(80 ± 7)	(170 ± 30)

The spectral decomposition in Fig. 9(B) shows that the additional residual is primarily low frequency. In three of the four runs, tilt contained substantially more power below approximately 3 Hz, while pan and tilt spectra converge towards similar power at high frequency. Correspondingly, most of the excess field variance lies below the first few hertz of the reverse-cumulative RMS integral. This indicates that much of the pointing error can be compensated by low-frequency fine steering. A narrow feature was observed on-site at 46.7 Hz in all trials.




Figure 9. (A) Residual centroid error distribution for a representative field trial. (B) Per-axis PSD for four field trials, showing additional low-frequency tilt residual and a 46.7 Hz feature. (C) Reverse-cumulative RMS.

Laboratory and field radial-error distributions are compared directly in Fig. 10(A). The median radial residual increases from 23 μrad to 26–43 μrad, whereas the 99th percentile radial error increases from 70 μrad to 80–250 μrad. Relative to the laboratory, the field distributions change more strongly in the tail than at the median.




Figure 10. (A) Empirical radial CDF for the in-lab trial and four SFU-Cypress on-site closed-loop trials. (B) Reverse-cumulative centroid dropout duration distribution of the SFU-Cypress runs. In-lab measurements produced no dropouts.

Centroid dropout was observed in field data, with 8.7–11.5 % invalid centroid packets and a median dropout duration of 5 ms (one exposure). Further, 99 % of dropout events were shorter than 35–40 ms across the trials, as shown in Fig. 10(B). Centroid dropout did not result in loss of tracking lock.

4.3 Derived Fine-Steering Requirements

The residual error distributions establish the required fine-steering capture range. The laboratory measurement gives $R_{99} \simeq 70$ μrad, while the worst field run gives $R_{99} = (250 \pm 40)$ μrad. A folded Zemax model of the implemented 50/180 mm telescope gives a FSM-sky gain of -0.559 , within 0.7 % of the first-order value $-2(50/180) = -0.556$. Applying the same first-order relation to the proposed 45/200 mm telescope gives a gain

of -0.450 . The worst measured $R_{99} = (250 \pm 40) \mu\text{rad}$ therefore requires a one-sided mechanical FSM range of at least $(0.56 \pm 0.09) \text{ mrad}$, before alignment margin.

The disturbance-rejection bandwidth is obtained from the reverse-cumulative RMS. Using half the median long-term $1/e^2$ beam half-angle as the residual target, $\sigma_{>f} \lesssim \theta_{LT}/2 \approx 12 \mu\text{rad}$, our field trials indicate a minimum bandwidth of approximately 50 Hz. Together, the measured capture range, disturbance spectrum, and dropout distribution provide dynamic requirements for fine-steering implementation for the horizontal quantum link.

4.4 Outlook

The laboratory and field measurements show that, with narrow spectral filtering, the implemented coarse stage can maintain lock over the SFU–Cypress link with residuals within the capture range of candidate fine-steering actuators. The largest field degradation appears in the tails of the pointing distribution, and is dominated by low-frequency tilt motion. Measurements do not separate atmospheric fluctuations from platform disturbances, though wind gusts were qualitatively observed to coincide with the strongest tail events.

The resulting requirements are a one-sided FSM range of at least $(0.56 \pm 0.09) \text{ mrad}$, a residual target of approximately $12 \mu\text{rad}$, and a minimum disturbance-rejection bandwidth of 50 Hz. Centroid dropouts require the controller to tolerate approximately 40 ms gaps. These requirements inform the fixed pupil-plane interface, reserved fine-steering interfaces, and custom-machined optical assembly (Sec. 3.2.2) under development, allowing the fine-steering hardware to be exchanged.

The present work validated the integrated terminal architecture and coarse APT. Fine-steering, receiver AO, and end-to-end quantum coupling are beyond the scope of this manuscript. The horizontal SFU–Cypress path provides a terrestrial ABL test case for mobile platform development; subsequent mobile platform deployment requires characterisation of platform dynamics.

5. CONCLUSION

We developed a portable, SWaP-constrained free-space optical transmitter with modular optics, embedded control, and a two-stage APT architecture for inter-regional quantum links. The coarse tracking stage achieved a radial 99th-percentile residual of approximately $70 \mu\text{rad}$ in the laboratory and $80\text{--}250 \mu\text{rad}$ over the 20.44 km SFU–Cypress link. The 99th-percentile centroid-dropout duration ranged from approximately 35 to 40 ms across field trials without loss of tracking lock. Measurements constrain the fine-steering stage to a one-sided mechanical range of at least $(0.56 \pm 0.09) \text{ mrad}$, a residual target of approximately $12 \mu\text{rad}$, and a minimum disturbance-rejection bandwidth of 50 Hz. Together with the fixed pupil-plane fine-steering interface, the terminal provides a defined path for extending the demonstrated coarse stage to complete terminal APT for subsequent terrestrial, mobile, and airborne free-space quantum links.

ACKNOWLEDGMENTS

H.J. thanks T. Jennewein and S. V. Grayli for their supervision, and K. Mohammadi for advisory discussions on optical design and terminal operation. The authors further thank E. Athanasopoulos, A. Charnetski, P. N. Fard, D. Musiienko, and M. Pocrnjic for contributions to on-site field measurements. We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC) and the NSERC CREATE in Quantum Computing Program (Grant Number 543245). This project was supported in part by the NORAD IDeaS program.

REFERENCES

- [1] Wehner, S., Elkouss, D., and Hanson, R., “Quantum internet: A vision for the road ahead,” *Science* **362**, eaam9288 (Oct. 2018).
- [2] Cao, Y., Zhao, Y., Wang, Q., Zhang, J., Ng, S. X., and Hanzo, L., “The Evolution of Quantum Key Distribution Networks: On the Road to the Qinternet,” *IEEE Communications Surveys & Tutorials* **24**(2), 839–894 (2022).

- [3] Zhu, T.-X., Liu, X., Zhou, Z.-Q., and Li, C.-F., “Remote quantum networks based on quantum memories,” *Nanophotonics* **14**, 1975–1992 (June 2025). arXiv:2508.19538 [quant-ph].
- [4] Pirandola, S., “Limits and security of free-space quantum communications,” *Physical Review Research* **3**, 013279 (Mar. 2021).
- [5] Main, D., Drmota, P., Nadlinger, D. P., Ainley, E. M., Agrawal, A., Nichol, B. C., Srinivas, R., Araneda, G., and Lucas, D. M., “Distributed quantum computing across an optical network link,” *Nature* **638**, 383–388 (Feb. 2025).
- [6] Zhu, T.-X., Zhang, C., Ou, Z.-W., Liu, X., Liang, P.-J., Hu, X.-M., Huang, Y.-F., Zhou, Z.-Q., Li, C.-F., and Guo, G.-C., “A metropolitan-scale multiplexed quantum repeater with Bell nonlocality,” (Aug. 2025). arXiv:2508.17940 [quant-ph].
- [7] Ren, Y., Wang, X., Lv, Y., Bacco, D., and Jing, J., “Distribution of Multiplexed Continuous-Variable Entanglement for Quantum Networks,” *Laser & Photonics Reviews* **16**(11), 2100586 (2022).
- [8] Wang, R., Clark, M. J., Alia, O., Bahrani, S., Aktas, D., Peranić, M., Stipčević, M., Lončarić, M., Rarity, J., Joshi, S. K., and Simeonidou, D., “Dynamic Entanglement Distribution for Multi-User and Multi-Protocol Quantum Networking,” (July 2026). arXiv:2607.15262 [quant-ph].
- [9] Vinet, S., England, D., Xu, C.-Q., and Jennewein, T., “Demonstration of a reconfigurable quantum network architecture suitable for ground-to-space communication,” *Optica Quantum* **4**, 14–21 (Feb. 2026).
- [10] Petrovich, M., Numkam Fokoua, E., Chen, Y., Sakr, H., Adamu, A. I., Hassan, R., Wu, D., Fatobene Ando, R., Papadimopoulos, A., Sandoghchi, S. R., Jasion, G., and Poletti, F., “Broadband optical fibre with an attenuation lower than 0.1 decibel per kilometre,” *Nature Photonics* **19**, 1203–1208 (Sept. 2025).
- [11] Wootters, W. K. and Zurek, W. H., “A single quantum cannot be cloned,” *Nature* **299**, 802–803 (Oct. 1982).
- [12] Azuma, K., Economou, S. E., Elkouss, D., Hilaire, P., Jiang, L., Lo, H.-K., and Tzitrin, I., “Quantum repeaters: From quantum networks to the quantum internet,” *Reviews of Modern Physics* **95**, 045006 (Dec. 2023).
- [13] Liao, S.-K., Yong, H.-L., Liu, C., Shentu, G.-L., Li, D.-D., Lin, J., Dai, H., Zhao, S.-Q., Li, B., Guan, J.-Y., Chen, W., Gong, Y.-H., Li, Y., Lin, Z.-H., Pan, G.-S., Pelc, J. S., Fejer, M. M., Zhang, W.-Z., Liu, W.-Y., Yin, J., Ren, J.-G., Wang, X.-B., Zhang, Q., Peng, C.-Z., and Pan, J.-W., “Long-distance free-space quantum key distribution in daylight towards inter-satellite communication,” *Nature Photonics* **11**, 509–513 (Aug. 2017).
- [14] Pugh, C. J., Kaiser, S., Bourgoin, J.-P., Jin, J., Sultana, N., Agne, S., Anisimova, E., Makarov, V., Choi, E., Higgins, B. L., and Jennewein, T., “Airborne demonstration of a quantum key distribution receiver payload,” *Quantum Science and Technology* **2**, 024009 (June 2017).
- [15] Yin, J., Cao, Y., Li, Y.-H., Liao, S.-K., Zhang, L., Ren, J.-G., Cai, W.-Q., Liu, W.-Y., Li, B., Dai, H., Li, G.-B., Lu, Q.-M., Gong, Y.-H., Xu, Y., Li, S.-L., Li, F.-Z., Yin, Y.-Y., Jiang, Z.-Q., Li, M., Jia, J.-J., Ren, G., He, D., Zhou, Y.-L., Zhang, X.-X., Wang, N., Chang, X., Zhu, Z.-C., Liu, N.-L., Chen, Y.-A., Lu, C.-Y., Shu, R., Peng, C.-Z., Wang, J.-Y., and Pan, J.-W., “Satellite-based entanglement distribution over 1200 kilometers,” *Science* **356**, 1140–1144 (June 2017).
- [16] Gong, Y.-H., Yang, K.-X., Yong, H.-L., Guan, J.-Y., Shentu, G.-L., Liu, C., Li, F.-Z., Cao, Y., Yin, J., Liao, S.-K., Ren, J.-G., Zhang, Q., Peng, C.-Z., and Pan, J.-W., “Free-space quantum key distribution in urban daylight with the SPGD algorithm control of a deformable mirror,” *Optics Express* **26**, 18897–18905 (July 2018).
- [17] Tian, X.-H., Yang, R., Liu, H.-Y., Fan, P., Zhang, J.-N., Gu, C., Chen, M., Hu, M., Lu, F.-Y., Zhu, C., Yin, Z.-Q., Yin, Z.-J., Yuan, M., Wang, S., Chen, W., Gong, Y.-X., Zhu, S.-N., and Xie, Z., “Experimental Demonstration of Drone-Based Quantum Key Distribution,” *Physical Review Letters* **133**, 200801 (Nov. 2024).
- [18] Conrad, A., Cochran, R., Sanchez-Rosales, D., Isaac, S., Javid, T., Rezaei, T., Schroeder, A. J., Golba, G., Gutha, A., Wilens, B., Herndon, K., Hill, A., Chapman, J., Call, I., Szabo, J., Corrigan, A., Kamin, L., Lütkenhaus, N., Gauthier, D. J., and Kwiat, P. G., “Drone- and vehicle-based quantum key distribution,” (May 2025). arXiv:2505.17587 [quant-ph].

- [19] Li, Y., Cai, W.-Q., Ren, J.-G., Wang, C.-Z., Yang, M., Zhang, L., Wu, H.-Y., Chang, L., Wu, J.-C., Jin, B., Xue, H.-J., Li, X.-J., Liu, H., Yu, G.-W., Tao, X.-Y., Chen, T., Liu, C.-F., Luo, W.-B., Zhou, J., Yong, H.-L., Li, Y.-H., Li, F.-Z., Jiang, C., Chen, H.-Z., Wu, C., Tong, X.-H., Xie, S.-J., Zhou, F., Liu, W.-Y., Ismail, Y., Petruccione, F., Liu, N.-L., Li, L., Xu, F., Cao, Y., Yin, J., Shu, R., Wang, X.-B., Zhang, Q., Wang, J.-Y., Liao, S.-K., Peng, C.-Z., and Pan, J.-W., “Microsatellite-based real-time quantum key distribution,” *Nature* **640**, 47–54 (Mar. 2025).
- [20] Liu, P.-X., Lin, Y.-P., Zhou, Z.-Q., Li, C.-F., and Guo, G.-C., “Global-scale quantum networking using hybrid-channel quantum repeaters with relays based on a chain of balloons,” (July 2025). arXiv:2507.15178 [quant-ph].
- [21] Mohammadi, K., Godin, P. J., and Jennewein, T., “On the viability of Transatlantic Quantum Entanglement Distribution using Combined Satellite and Stratospheric Relay Nodes,” (June 2026). arXiv:2606.09805 [quant-ph].
- [22] Kržič, A., Döll, N., Chandrashekhara, U., Spiess, C., and Steinlechner, F., “Adjustable spatial filter for optimal free-space quantum communication round the clock,” *Optics Express* **31**, 32362–32372 (Sept. 2023). arXiv:2303.09942 [quant-ph].
- [23] “Vancouver Quantum Network,” (2025). Simon Fraser University, <https://www.sfu.ca/research/vancouver-quantum-network.html>, accessed 16 August 2026.
- [24] Gruneisen, M. T., Eickhoff, M. L., Newey, S. C., Stoltenberg, K. E., Morris, J. F., Bareian, M., Harris, M. A., Oesch, D. W., Olike, M. D., Flanagan, M. B., Kay, B. T., Schiller, J. D., and Lanning, R. N., “Adaptive-optics-enabled quantum communication: A technique for daytime space-to-Earth links,” *Physical Review Applied* **16**, 014067 (July 2021). arXiv:2006.07745 [physics].
- [25] Ghalaii, M. and Pirandola, S., “Quantum communications in a moderate-to-strong turbulent space,” *Communications Physics* **5**, 38 (Feb. 2022).
- [26] Klein, B. J. and Degnan, J. J., “Optical Antenna Gain. 1: Transmitting Antennas,” *Applied Optics* **13**, 2134–2141 (Sept. 1974).
- [27] Jennewein, T., Simon, C., Fougères, A., Babin, F., Kimiaee Asadi, F., Kuntz, K. B., Maisonneuve, M., Moffat, B., Mohammadi, K., and Panneton, D., “QEYSSat 2.0—white paper on satellite-based quantum communication missions in Canada,” *Canadian Journal of Physics* **103**(4), 328–376 (2025).
- [28] Podmore, H., Cain, J., D’Souza, I., Jennewein, T., Higgins, B. L., Godin, P., Koujelev, A., Hudson, D., Taylor, M., Beica, H., and Rao, A., “Quantum and classical laser communication terminals on the QEYSSat mission,” in *[International Conference on Space Optics — ICSO 2024]*, **13699**, 1369943, SPIE (July 2025).
- [29] Kim, I. I., McArthur, B., and Korevaar, E. J., “Comparison of laser beam propagation at 785 nm and 1550 nm in fog and haze for optical wireless communications,” in *[Optical Wireless Communications III]*, **4214**, 26–37, SPIE (Feb. 2001).
- [30] Environment and Climate Change Canada, “Historical Climate Data,” (2026). Extracted from the Environment and Climate Change Canada Historical Climate Data website, https://climate.weather.gc.ca/index_e.html, accessed 11 August 2026.
- [31] Kochanov, R. V., Gordon, I. E., Rothman, L. S., Wcisło, P., Hill, C., and Wilzewski, J. S., “HITRAN Application Programming Interface (HAPI): A comprehensive approach to working with spectroscopic data,” *Journal of Quantitative Spectroscopy and Radiative Transfer* **177**, 15–30 (July 2016).
- [32] Andrews, L. C. and Phillips, R. L., *[Laser beam propagation through random media]*, SPIE Press, Bellingham, Wash, 2nd ed. (2005).
- [33] Dikmelik, Y. and Davidson, F. M., “Fiber-coupling efficiency for free-space optical communication through atmospheric turbulence,” *Applied Optics* **44**, 4946–4952 (Aug. 2005).
- [34] Tokovinin, A., “From Differential Image Motion to Seeing,” *Publications of the Astronomical Society of the Pacific* **114**, 1156–1166 (Oct. 2002).
- [35] Shaklan, S. and Roddier, F., “Coupling starlight into single-mode fiber optics,” *Applied Optics* **27**, 2334–2338 (June 1988).
- [36] Ishtiaq, S., Schmidt, C., and Moll, F., “Modular FSO optical system design for classical and quantum optical communication systems,” in *[International Conference on Space Optics — ICSO 2020]*, **11852**, 190, SPIE (June 2021).

- [37] Welch, P., “The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms,” *IEEE Transactions on Audio and Electroacoustics* **15**, 70–73 (June 1967).