Architecture overview and data summary of a 5.4 km free-space laser communication experiment

Architecture Overview and Data Summary of a 5.4 km Free-Space Laser Communications Experiment

John D. Moores, Frederick G. Walther, Joseph A. Greco, Steven Michael, William E. Wilcox Jr., Alicia M. Volpicelli, Richard J. Magliocco, and Scott R. Henion
Massachusetts Institute of Technology Lincoln Laboratory, Lexington, MA 02420-9108

ABSTRACT

MIT Lincoln Laboratory designed and built two free-space laser communications terminals, and successfully demonstrated error-free communication between two ground sites separated by 5.4 km in September, 2008. The primary goal of this work was to emulate a low elevation angle air-to-ground link capable of supporting standard OTU1 (2.667 Gb/s) data formatting with standard client interfaces. Mitigation of turbulence-induced scintillation effects was accomplished through the use of multiple small-aperture receivers and novel encoding and interleaver hardware. Data from both the field and laboratory experiments were used to assess link performance as a function of system parameters such as transmitted power, degree of spatial diversity, and interleaver span, with and without forward error correction.

This work was sponsored by the Department of Defense, RRCO DDR&E, under Air Force Contract FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the authors and are not necessarily endorsed by the United States Government.

Keywords: free-space optical communications, laser communications, bit error ratio, optical scintillation, pointing and tracking, interleaving, diversity, field demonstration

1. INTRODUCTION

In a five-month rapid development effort, the Advanced Lasercom Systems and Operations Group at MIT Lincoln Laboratory (MIT/LL) modeled, designed, built, integrated, and fielded two custom 2.7 gigabaud (2.7 billion symbols per second), 1550 nm-band lasercom terminals at field sites in the western suburbs of Boston, separated by 5.4 km of wooded countryside. The communications transmitter was situated in a shed attached to a fire tower, while the receiver was housed in the dome of MIT/LL's Firepond facility: see Figure 1. The fire tower terminal represented an aircraft terminal in this set of experiments.

Free-Space Laser Communications IX, edited by Arun K. Majumdar, Christopher C. Davis,
Proc. of SPIE Vol. 7464, 746404 · © 2009 SPIE · CCC code: 0277-786X/09/$18 · doi: 10.1117/12.825591

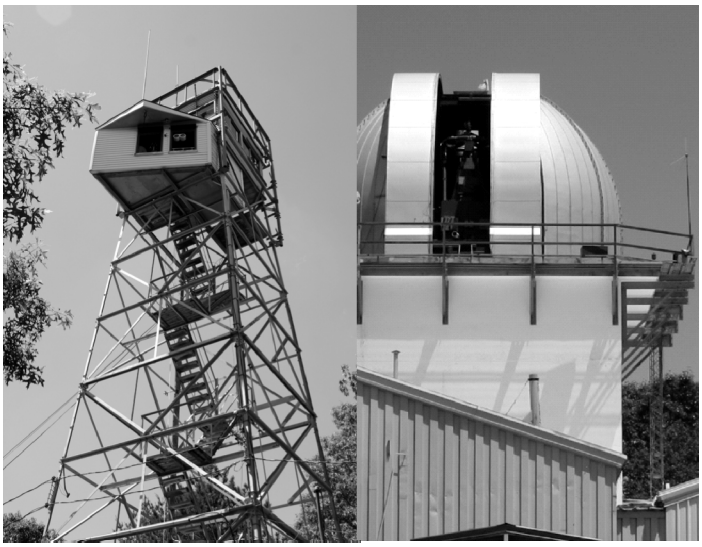

Fig. 1. Fire tower housing the transmitter and telescope dome housing the receiver. These platforms were separated by a distance of 5.4 km.

The field-site effort included data collection throughout much of September 2008, repeatedly verifying high-availability communications over the link under widely varying turbulence conditions, using modest, eye-safe levels of optical power. Error-free transmission using less than 50 mW of optical power through severe turbulence was made feasible through the use of spatial and temporal diversity techniques. Four receive apertures separated by more than the coherence diameter were used to capture a single transmit beam, achieving spatial diversity. Temporal diversity was achieved using symbol interleaving combined with forward error correction coding (FEC), with interleaver spans of up to 2 s. Small, 12 mm, apertures were used, which facilitated pointing, tracking, and coupling into single-mode fiber. The high availability, predominantly error-free performance achieved using diversity techniques enabled the use of standard terrestrial network protocol interface — in our case Ethernet over SONET. Our claims of error-free performance and high availability are based on rigorous testing with protocol analyzers, rather than qualitative assessment of received video. This is an important distinction because video tends to conceal errors: a video stream may not utilize the full channel bandwidth, and temporary corruption of isolated pixels tends not to render images unusable. The protocol tester identifies every error.

Potential applications for lasercom terminals similar to those demonstrated include networking for forces on the move and transmission of aggregated tactical data in theater. Because the characteristics of the ground-to-ground link are similar to those of a much longer air-to-ground link, our demonstration results can be extrapolated to applications such as the offload of remote sensing data from aircraft to the ground.

2. ARCHITECTURE

The system constructed for this experiment consists of a communications data plane; a control plane; a set of instruments; and a pointing, acquisition, and tracking (PAT) system. A conceptual diagram showing the data flow from transmitter to receiver through the optical channel is illustrated in Figure 2.

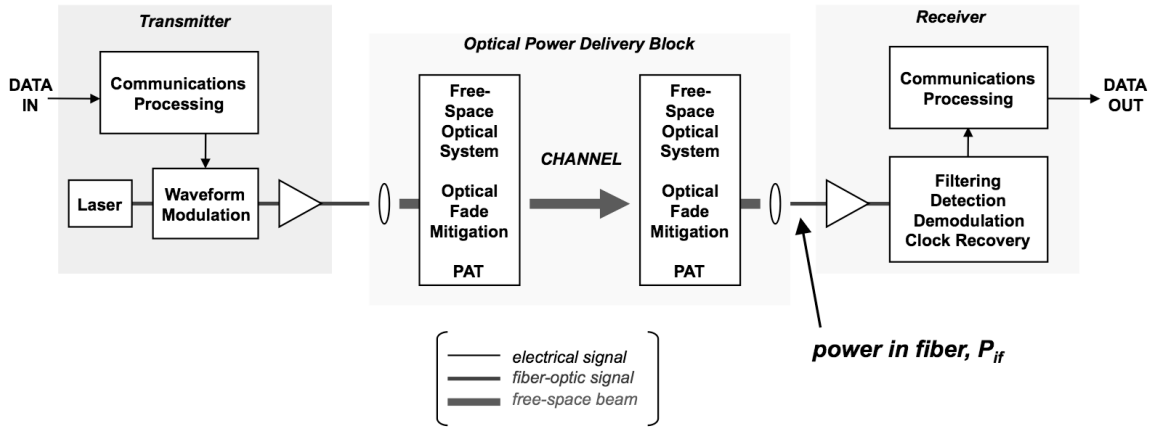

Fig. 2. Free-space laser communication system block diagram. A pointing, acquisition, and tracking (PAT) system was used at both ends to minimize pointing errors.

The communications data plane at the Fire Tower in Groton, MA includes dual Gigabit Ethernet (GbE) interfaces to the data source. A commercial Ciena 3605 multiservice platform maps these two GbE streams into an OC-48 SONET stream. The OC-48 stream feeds into the High-speed FEC and Interleaving Cards (HFICs), a set of custom FPGA-based boards developed by MIT Lincoln Laboratory, which perform OTU 1 framing, OTU 1 standard Reed-Solomon (255,239) encoding, and custom interleaving of RS(255,239) symbols with interleaver span (time interval from first to last symbol of a single codeword) as long as 2 seconds. The output of the HFICs feeds into a photonic, electro-optic return-to-zero on-off keyed (RZ-OOK) transmitter with a booster erbium-doped fiber amplifier (EDFA). The output of the EDFA is converted to a free-space beam with a fiber launch assembly and is sent through a 12 mm aperture across the 5.4 km link.

The communications data plane at the receiver, at Firepond in Westford, MA is similar to that at the Fire Tower. There are four receive apertures, two 12 mm in diameter and two 11 mm in diameter, oriented in a square 10 cm per side. Figure 3 is a photograph of the setup illustrating the ability to adjust the locations of the apertures, which enabled us to test different spacings to ensure sufficient decorrelation. The 10 cm aperture spacing was sufficiently wide to largely, but not entirely, decorrelate the fading of the four received beams under the strongest measured turbulence. The four received beams are coupled into single mode optical fiber under the control of the PAT system. Each beam is detected in a separate photodiode. The photocurrents are electrically summed in the receiver module. This incoherent sum of uncorrelated or nearly uncorrelated power time series greatly reduces the received power variance due to fading. The sum feeds a clock recovery circuit and a decision circuit. The output of the decision circuit feeds into the HFICs, which deinterleave, decode and deframe, sending an OC 48 stream to a Ciena 3605 multiservice platform. The Ciena provides two GbE interfaces to the data sink.

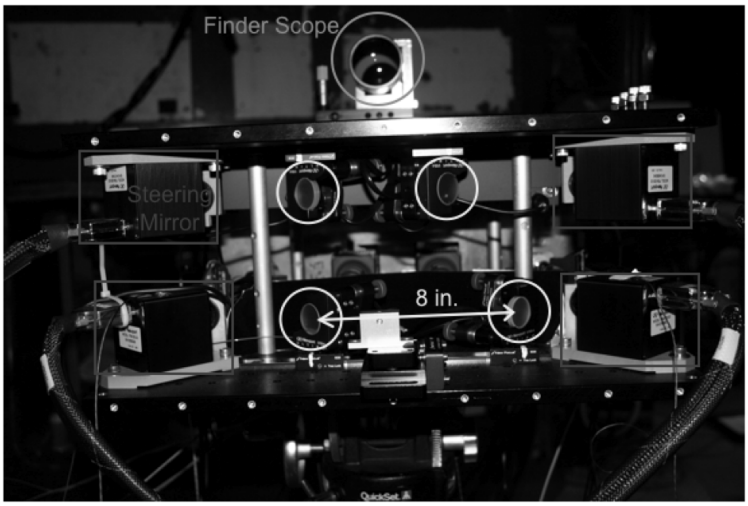

Fig. 3. Photograph of the link receiver showing the locations of the four receiver apertures (inside the four smaller circles) and the locations of the steering mirrors (inside the rectangles). Finder Scope on top (inside large circle) is used for initial alignment. The aperture spacing was adjustable.

The pointing, acquisition, and tracking system is designed to compensate for platform vibrations at both sites and for atmospheric turbulence-induced phase aberrations. This ensures accurate pointing from transmitter to receiver and ensures efficient coupling from free-space into single mode fiber at the receiver. The choice of aperture diameters that are smaller than the spatial coherence diameter r_0 of the path simplifies the PAT system considerably because the phase aberrations are effectively linear across the aperture. The tip/tilt can be corrected using a fast steering mirror (FSM), without the need for adaptive optics to correct higher-order phase distortions. From the receive side, four beacon beams were transmitted out the same apertures through which the communications beams were received. Wavelength (de)multiplexing and isolation enabled simultaneous operation of beacons and communications beams. Each beacon was at a different wavelength, with sufficient wavelength separation for incoherent summing at the Fire Tower. The light received at the Fire Tower was captured on a focal plane array (FPA) and in an optical power meter. The integration time of the FPA was adjusted so as to eliminate reflections off the Firepond telescope dome, ensuring that the beacon beams were guiding the pointing. At the Firepond side, from the free space channel, each receive optical path included a piezo-driven slow steering flat mirror (analogous to a gimbal), fast steering mirror, and a beam splitter to send received light to the FPA for PAT and to the optical fiber for communications (and PAT). A hollow retro-reflector (HRR) could be flipped in between the FSM and beam splitter for initial boresighting the fiber and FPA. The tracking bandwidth of the primary tracking loop was 50 Hz. Two additional tracking loops were used. The first involved nutation of the FSM and detection of power in fiber to minimize FPA-fiber co-alignment error. This is a slow correction ~0.2 Hz. The second loop, also ~0.2 Hz, used the FSM position sensors to detect FSM angular bias and adjust the steering flats to minimize FSM bias (analogous to a gimbal).

The suite of instruments and sensors included a COTS Scintec BLS900 scintillometer (transmitter at Fire Tower, receiver at Firepond), video cameras at each site, GPS receivers, NTP servers, a weather station, and a near infrared camera for capturing 8”x8” scintillation patterns. A representative image measured in the aperture plane of the receiver is displayed in Figure 4.

Instruments that provided system performance metrics included optical power meters with 70 dB dynamic range and Innocor TestPoint protocol generator/analyzers. The HFICs provided key communications performance metrics. The HFIC counts the number of uncorrectable subrows (USRs), which is the number of RS(255,239) codewords that could not be properly decoded by the decoder. The HFIC also provides the number of corrected errors, providing a measure of the uncoded, raw channel performance.

The control plane included command, control, and telemetry of nearly every element in the system, using standard Ethernet. Serial port concentrators and media access units were used to interface to devices lacking built-in Ethernet ports. At the heart of the control plane were three key pieces of hardware: (1) a server providing basic services such as NFS, TFTP, DNS, as well as the data logging service of the telemetry system, (2) an ESX server supporting multiple virtual machines, and (3) a RAID array for capturing telemetry. Thin clients were distributed throughout working areas

enabling access to nearly every element of the system from any location. Key to the ability to support the kHz-rate telemetry and control processing was aggregation of messages, rather than sending each small message in a separate Ethernet frame.

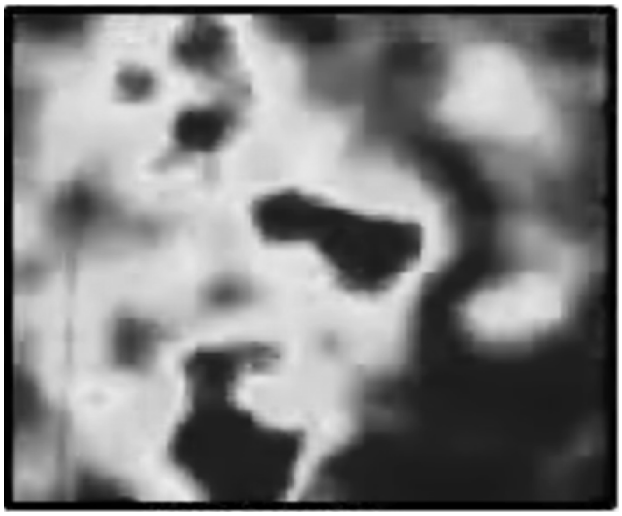

Fig. 4. Measured intensity pattern over an 8"x 8" area. The characteristic size of the scintillation features is comparable to the separation between the receiver apertures.

The primary goal of field testing was to measure optical power received in single-mode fiber, sampled at 1 kHz. This measurement characterizes the fading of the channel, the PAT system's ability to track, and the effectiveness of receive-side spatial diversity. The PAT system telemetry consistently indicated that tracking was excellent and did not degrade communications performance. Data were captured throughout much of the month of September, 2008, on weekdays, typically from mid-morning hours through sunset, and sometimes into the night. Because of concerns for the equipment, power in fiber data were not captured during intervals of precipitation. Scintillometer and weather station measurements could be captured continually.

In the field, mitigation of the channel effects was achieved using spatial diversity with four spatially separated receive apertures. The 10 cm separation of the apertures provided good decorrelation of the received time series in the four apertures. The received optical power in each aperture was measured, enabling analysis of the correlation between the time series of the four apertures. Figure 5 shows probability distribution functions (PDFs) for received optical power in single mode fiber vs. received optical power for three different turbulence severities. The broader PDFs (with lower mean received optical power in fiber) are for a single aperture, while the narrower PDFs (with higher mean received optical power) represent the incoherent superposition of all four apertures. The mean received power is 6 dB higher with four apertures, and the variance decreases by a factor of between 2 and the ideal 4 depending upon the strength of turbulence.

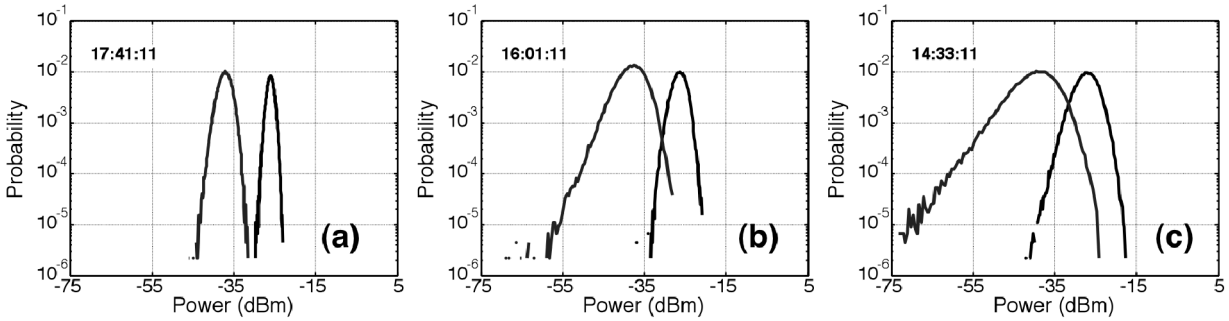

Fig. 5. Probability density functions (PDFs) vs. received power for three different turbulence strengths, increasing from (a) to (c). Each plot contains two PDFs. Within each plot, the broader PDF with lower mean received optical power is the

measured power in single mode fiber for a single Rx aperture. The narrower PDF with higher mean power in each plot is the incoherent sum of power in fiber for four Rx apertures.

Figure 6 shows the ratio of the normalized variance of a single receive aperture to the normalized variance of the incoherently summed receiver optical power vs. the measured refractive index structure parameter, C_n^2, a measure of turbulence strength. The single aperture variance is normalized to the square of the single aperture mean power, while the four aperture variance is normalized to the square of the mean of the incoherently summed power from the four apertures. Measurements of received power were made over 120 second intervals, sampled at 1 kHz. A commercial scintillometer measured C_n^2. Although there is considerable scatter in the data, there is a clear trend consistent with increased time series correlation amongst the four apertures as the turbulence becomes more severe (larger C_n^2). The implication is that with fixed aperture spacings, the spatial diversity becomes most effective when most needed.

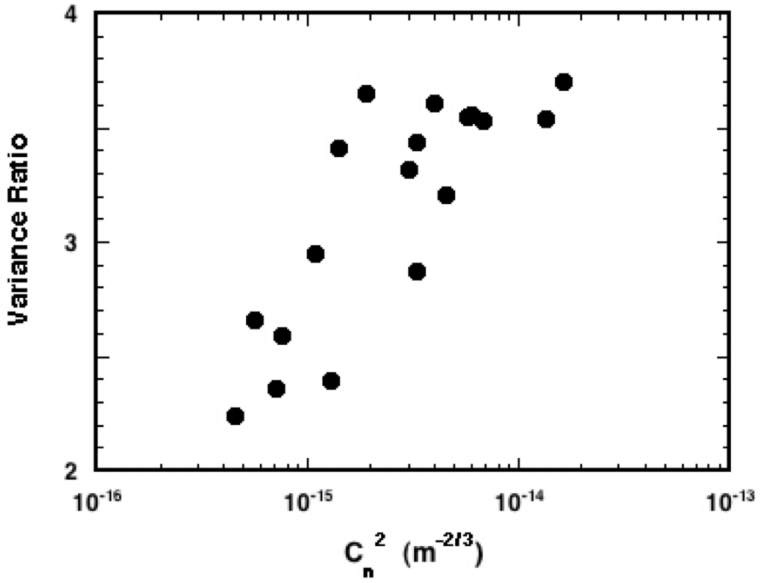

Fig. 6. Measured normalized variance reduction factor using spatial diversity vs. turbulence strength. Vertical axis - ratio of variance of PDF of incoherent sum of light from one receive aperture normalized to the square of the mean power received by that aperture to the variance of the received signal through four apertures normalized to the square of the mean power received by the four apertures. For perfectly correlated aperture power time series, the ratio should approach unity, while for perfectly uncorrelated aperture powers, the ratio should approach 4. Horizontal axis - turbulence strength, measured refractive index structure parameter.

In the laboratory, more effective mitigation could be achieved with the addition of temporal diversity – the use of custom temporal interleaving. The custom interleaving is much more effective than the simple in-frame interleaving of the ITU G.709 OTN Recommendations because much greater interleaver depths are achieved (~1 s vs. ~50 μs). Measured 1/e autocorrelation widths of the received power time series ("coherence times") often exceeded 10 ms, for which the OTN interleaving is ineffective. Temporal diversity was implemented and tested in the field, but because of the suboptimal performance of the interface from the RZ-OOK receiver to the HFIC, no temporal interleaving gains could be obtained at raw uncoded channel BER > 10^{-6}. Despite using only spatial mitigation in the field, we measured very few outages, mainly during the worst-case mid-day turbulence with $C_n^2=10^{-14}$ $m^{-2/3}$. These outages were typically singular events, with automatic recovery on the order of 1-2 seconds. As an example, on the benign afternoon of 5 September 2008, from 15:00-17:00, not a single outage was measured. On a more typical day, 9/8/2008, from 12:30-14:30, during the most aggressive part of the day, only a single outage event occurred. On one of the most aggressive days measured, 10 September 2008, during the 7.5 hour interval from 10:00-17:30, not counting a brief (seconds) interval during which the beams were manually blocked, only eleven outages were measured. We believe that with an optimized receiver, no outages would have occurred.

In order to verify the effectiveness of the mitigation schemes, following field testing, we created a laboratory test bed where the measured channel conditions could be reproduced. During the interval between the end of field testing and the

beginning of laboratory testing, the interface from the optical receiver to the HFIC was corrected for better performance. This enabled measurement of coding and interleaving gains at raw, uncoded BERs as high as $3.7x10^{-3}$, approximately the breaking point for RS(255,239) with randomly distributed raw bit errors, and roughly an order of magnitude below the breaking point with highly correlated bit errors within a symbol (the relevant limit in the slowly fading channel).

At the main MIT/LL facility, a "fade generator" had been developed, consisting of a DiCon fast variable optical attenuator (VOA) and driver, and associated MIT/LL-written software to input a stream of power-in-fiber measurements from the field and reproduce the measured channel behavior. The VOA had a dynamic range of 40 dB up to 400 Hz, which was sufficient except for single channel strong turbulence cases. Truncating the power to the mean value (eliminating "surges") reduced the required dynamic range sufficiently. In the field, power in fiber measurements were made on each receive aperture individually, enabling the emulation of any sum of 1-4 channels, as desired.

The experimental setup also included a slow variable optical attenuator (SLOA), which was used to measure receiver sensitivity, with the fast VOA and HFIC (de)interleavers turned off. The HFIC decoder provided counts of raw corrected errors and uncorrectable subrows (USRs). The results of varying the received power are depicted in Figure 7. We define the sensitivity of the receiver with coding to be the minimum power level at which USRs are recorded, measured to be -40.3 dBm. The middle and right-hand curves in the figure show raw bit error ratio (BER) vs. mean received optical power in fiber. Errors were incurred at both low power levels (corresponding to fades) and at high power levels (corresponding to surges). USRs are measured at raw BERs of greater than $\sim 4x10^{-4}$. In the limit of symbols being in error by exactly a single bit (a rough approximation for the AWGN channel used in laboratory baseline testing), RS(255,239) is theoretically effective to a BER of $3.9x10^{-3}$. In the extreme bursty limit of 8 bit errors per symbol error (an approximation for the case of interleaving on a slowly fading channel), the theoretical correctable BER is $3.1x10^{-2}$.

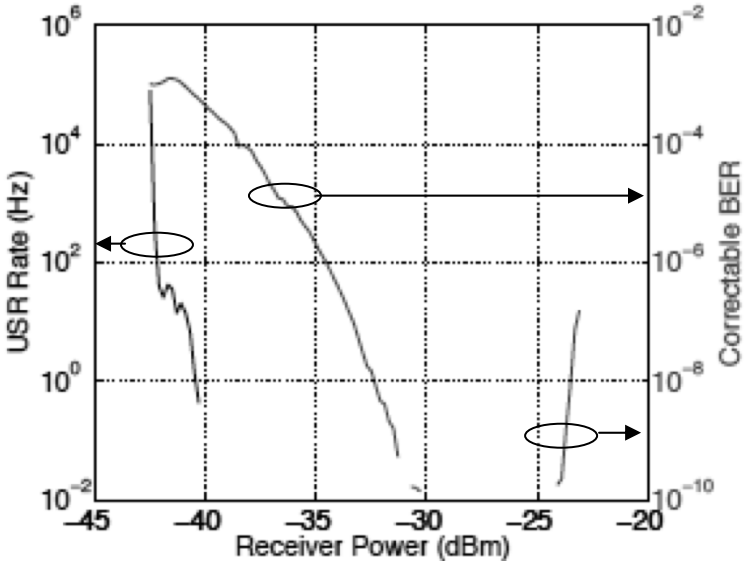

Fig. 7. Measured receiver sensitivity. Abscissa - received optical power in dBm (dB relative to 1 mW). Left ordinate and leftmost curve - uncorrectable subrow (USR) rate. Right ordinate, middle and rightmost curves - corrected bit error ratio (BER).

As anticipated, results of laboratory testing indicated that the system was able to perform error-free under any measured channel conditions including the strongest mid-day turbulence. The power transmitted out of the Fire Tower aperture was 16.7 dBm (47 mW). A one-second interleaver span was selected. Measured channel coherence times (one-sided 1/e autocorrelation widths) depended strongly upon the transverse wind on the channel and could range from 1-20+ ms. With 255-symbol Reed-Solomon codewords, a one second interleaver span is expected to fully decorrelate the symbols if the coherence time is less than 3.9 ms. We found that it for this link and power level, this interleaver span was sufficient.

In the absence of interleaving, a received power below -40.3 dBm would produce a codeword error. Therefore, for example, in order to accommodate 10 dB fades with no interleaving, the average power would need to be above -30.3 dBm for error-free communication. The interleaving and FEC provided by the HFIC allow some percentage of the

power to drop below the receiver sensitivity. Without interleaving, the duration of a codeword is far less than a typical fade duration, implying bursty errors where entire codewords are in error. The interleaver spreads the codeword symbols in time, making the errors appear to be more uniformly distributed, over more codewords, enabling effective use of RS(255,239) block coding. For a given atmosphere, the difference between the lowest average power level providing error-free performance without an interleaver and the lowest average power level for which the HFIC can correct bits in error with no decoding failures, over a fixed time interval, represents the interleaver gain.

We have measured this interleaver gain for two different atmospheres. The "moderate atmosphere" dataset represents power in fiber recorded at Firepond on 5 Sept. 2008 between 14:40 and 14:42, with a C_n^2 of $2x10^{-15}$ $m^{-2/3}$, and with a sum channel time series 1/e one-sided autocorrelation width of 2.5±1 ms. The autocorrelation width is sufficiently short that we expect the interleaved symbols on the channel to experience largely uncorrelated fades (1 s interleaver span / 254 intervals = 3.9 ms). A two minute interval with 1kHz sampling provides $1.2x10^5$ samples. The time series for the four-channel incoherent sum was used. The time series was truncated at 0 dB to prevent errors due to surges. We measured fades as low as 8.38 dB below the mean and surges as high as 5.7 dB above the mean. We ran this two minute time series through the fast VOA, also varying the SLOA in 1 dB increments so as to determine the lowest mean received power level for which no USRs were recorded. The cumulative distribution function (CDF) for this power level is depicted in figure 8a. The vertical line in the plot indicates the receiver sensitivity for a nonfading channel with coding but without interleaving. The difference between the low power end of the CDF and the vertical line is the effective interleaving gain over this two minute interval. For this case, the gain is 3 dB.

We repeated the above experiment with a more turbulent atmosphere. This atmosphere was recorded at Firepond on 23 Sept. 2008 between 11:52 and 11:54, with $C_n^2 = 1x10^{-14}$ $m^{-2/3}$, which is very strong turbulence for this link. The sum channel time series 1/e one-sided autocorrelation width was 11±1 ms, which is more than twice the interleaved intersymbol delay, suggesting correlation between a symbol and its two prior and two subsequent nearest neighbors. The extremes of this interval were a fade 17.2 dB below the mean and a surge 10.4 dB above the mean, after incoherent summing of four channels. Figure 8b shows the CDF of power in fiber with the lowest DC power level that produces no uncorrectable subrows. Here, the interleaving gain is 8.6 dB during this two minute interval, despite the long time series autocorrelation width.

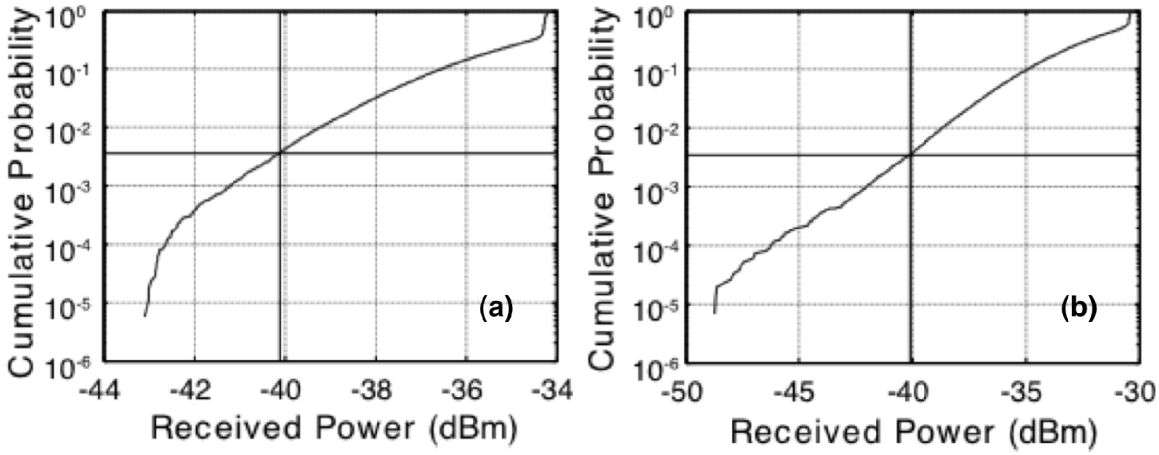

Fig. 8. Measured cumulative distribution function (CDF) of received power in fiber at the lowest average power providing error-free performance with coding and 1 s interleaving. Vertical line indicates measured power required for error-free performance for non-fading channel with RS(255,239) coding. (a) Moderate turbulence from 5 Sept. 2008 14:40-14:42 — interleaver gain is 3 dB over this interval. (b) Strong turbulence from 23 Sept. 2008 11:52-11:54 — interleaver gain is 8.6 dB over this interval.

The cumulative probability of received power less than the nonfading channel receiver sensitivity with error-free performance for both the moderate and strong turbulence cases is ~$3.7x10^{-3}$.

3. CONCLUSIONS

In summary, MIT Lincoln Laboratory performed successful field demonstrations of lasercom using MIT/LL-designed terminals located at field sites separated by 5.4 km in suburban/rural Massachusetts. Field measurements demonstrated the effectiveness of spatial diversity. The incoherent summing of the partially uncorrelated photocurrents of four spatially separated receivers reduced the fading-induced received optical power probability density function width (standard deviation normalized to mean power) by a factor of up to 2, being most effective when most needed - in the strong turbulence regime. Received optical power in fiber was measured at a sampling rate of 1 kHz. The field channel conditions were reproduced in the laboratory and the effectiveness of temporal diversity, through the use of symbol interleaving over timescales comparable to or longer than the fading time series autocorrelation width, was demonstrated in the laboratory. Interleaver gain of 8.6 dB was measured under strong turbulence for this link. We anticipate that spatial and temporal diversity techniques will be enabling to future lasercom links through aggressive atmospheric fading channels.