Demonstration of a high-efficiency free-space optical communications link

Kevin Birnbaum, William Farr, Jonathan Gin, Bruce Moision, Kevin Quirk, Malcolm Wright

Jet Propulsion Laboratory

California Institute of Technology

4800 Oak Grove Drive

Pasadena, CA 91109

ABSTRACT

In this paper we discuss recent progress on the implementation of a hardware free-space optical communications test-bed. The test-bed implements an end-to-end communications system comprising a data encoder, modulator, laser-transmitter, telescope, detector, receiver and error-correction-code decoder. Implementation of each of the component systems is discussed, with an emphasis on 'real-world' system performance degradation and limitations. We have demonstrated real-time data rates of 44 Mbps and photon efficiencies of approximately 1.8 bits/photon over a 100m free-space optical link.

Keywords: Optical Communications, Pulse-Position-Modulation

1. INTRODUCTION

Although optical communications has been successfully demonstrated in near-Earth scenarios (SILEX, OICETS, NFIRE), links are vastly more difficult at planetary ranges due to large inverse range squared losses. Unlike thermal noise dominated RF communications, paradigm shifts may be required at certain range limits due to shot noise limitations at very low received powers. Due to spacecraft power limitations, low-mass and power-efficient terminal implementations are required for deep space operations where links are photon starved, relative to terminals sufficient for near-Earth operations. A state of the art paradigm for hundreds of megabits per second (Mbps) data transmission from Mars planetary ranges consists of a spacecraft terminal comprising:

- a fiber laser transmitter with high peak to average power to implement power efficient modulation;
- a single optical telescope with near-sun pointing capability that functions as both transmit and receive aperture;
- a pointing-acquisition-tracking system that includes a point-ahead-mechanism to compensate for transmit-receive angular offsets and relies on a dim beacon from Earth for low frequency (to DC) pointing;
- a vibration isolation system to eliminate high frequency pointing disturbances from the host spacecraft,

and an Earth terminal comprising:

- a multi-meter diameter receive antenna with near-sun pointing capability;
- a photon counting detector system that can support flux rates of hundreds of mega-photons per second;
- a digital processor system with tens of gigahertz (GHz) bandwidth to recover the data stream;

E-mail:{Kevin.M.Birnbaum, William.H.Farr, Jonathan.W.Gin, Bruce.E.Moision, Kevin.J.Quirk, Malcolm.W.Wright}@jpl.nasa.gov

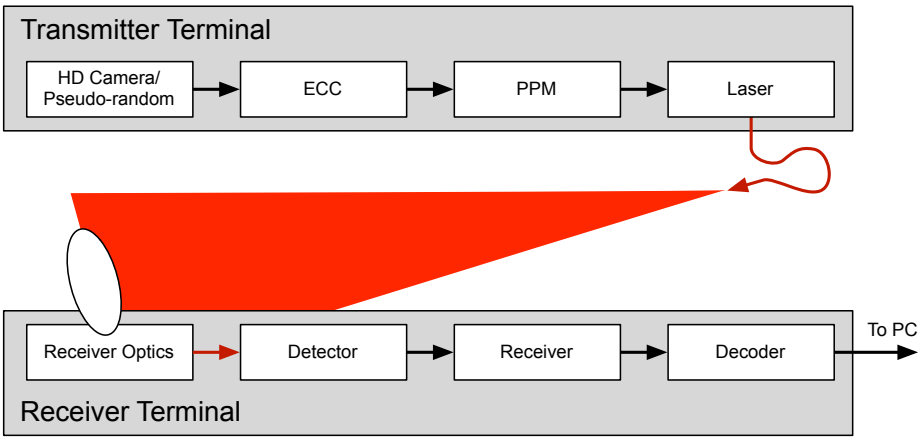

Figure 1. Block diagram of transmitter and receiver terminals

- a power efficient error-correction code;
- and a high average power beacon laser transmitter that includes forward data transmission capability.

At the present time, with the possible exception of the spacecraft vibration isolation system, none of these subsystems has had any flight qualified deep-space operations. In order to reduce mission risk prior to flight operations, JPL has aligned elements of its deep space optical communications technology program into a series of emulated deep space demonstrations in order to validate subsystem models and operations. These demonstrations have evolved from initial in-fiber lab validations to current day-and-night operational free space links that functionally validate transmitter and receiver systems in real time at data rates over 44 Mbps with efficiencies on the order of two bits per photon. Furthermore, as we have found that testing with pseudo-random data is not always sufficient to validate robust system operations (for example, temporal acquisition and tracking issues due to repetitive data sequences may be masked), we have included a channel for live transmission of high-definition television (HDTV) at a compressed nominal data rate of 30 Mbps.

In this paper, we describe the components of the real-time operational optical communications test-bed, illustrated in the block diagram of Figure 1. The test-bed implemented a building-top to building 100 m range link with the receiver placed behind a window. The receiver field of view allowed varying levels of sky background light on the detector, depending on the sun angle. The link was typically operated under day-time conditions with suitable bandpass, spatial, and neutral density filters to provide the desired detector signal and background levels. Alignment was accomplished by manual adjustment and facilitated by a co-aligned visible beam from the transmitter.

The paper is organized as follows. In Sections 2, 3, 4, 5, and 6 we describe the major functional elements: transmitter, receiver optics, photon counting detector, photon counting receiver, and error control code, respectively. In Section 7 we illustrate sample data from the test bed, and discuss estimating the photon efficiency of the system.

2. TRANSMITTER

The transmitter takes the input signal, encodes it into the appropriate data format and amplitude modulates a laser source. The transmitter laser signal is produced by a directly modulated single mode fiber-coupled semiconductor laser diode. The laser is integrated with a multi-GHz high speed driver and pulse-position-modulation(PPM) data encoder. An integrated fiber Bragg grating (FBG) in the output fiber ensured narrow linewidth operation during modulation of the diode laser. The PPM data encoder[1] utilizes a PC interface.[2] The performance of the directly modulated seed source was tested as part of a master oscillator power amplifier (MOPA) system.[2] The device could provide a nominal output power of tens to hundreds of mWs for CW

operation, however, only mWs of average power were typically transmitted with the PPM modulated signal. The laser was biased at threshold with the modulation current added above that level, so the laser is never turned completely off.

A fused fiber tap coupler was connected to the output so that the transmitted power or outgoing pulse pattern could be monitored. No transmit optics were used for beam propagation. The output was allowed to freely expand from the fiber across the link range. By not using any beam collimation optics, the beam divergence is on the order of 15 degrees full angle and the collected power at the receiver aperture is significantly reduced. This emulates a larger link range with a small received power. A co-aligned collimated visible laser source was mounted next to the transmitter to facilitate initial alignment and a beam tube was used to avoid any eye hazards at the output aperture of the fiber coupled transmitter.

The transmitted data source could be selected to be either a pseudo-random data stream or a real time HDTV signal compressed and sent through a FireWire interface from a HD camera. A control computer outputs the data stream through a interface to a PPM encoder that is integrated with a laser driver board in a compact module. The PPM encoder divides the incoming stream into blocks of 7536 bits. These blocks have a cyclic-redundancy-check (CRC) and two termination bits appended, and are then encoded with a rate 1/2 serially-concatenated-convolutionally-coded-PPM[3] (SCPPM) code, producing a digital signal composed of PPM symbols to the laser driver. Although the input signal could be clocked directly into the PPM encoder board, the internal clock (the encoder board clock) was not stable enough during typical operation. To decrease the timing jitter and facilitate receiver synchronization, an external clock reference was input to the PPM encoder board as well. This allowed the clock reference to be adjusted for different data rates.

3. RECEIVER OPTICS

The receiver optics collect the signal light coming from the transmitter and relay it to the photon counting detector while filtering out background light. The receiver optics also provide monitoring of the power and spatial alignment of the light. A diagram of the receiver optics is shown in Figure 2.

Light enters the receiver and is focused by the primary mirror, an aluminum-coated parabola . A spatial filter (field stop) is placed at the prime focus to reduce the amount of background light sent to the rest of the optical system. After this focus, the light is re-collimated by a achromatic doublet and folded off axis by a flat mirror. The light then passes through a series of filters to reduce the background and wave plates to control the polarization. The light then enters a sealed enclosure, where a second spatial filter (an iris in an image plane) is used to control the field of view of the subsequent detectors in order to adjust the received signal and noise power. Next in the optical path is a polarizing beam splitter (PBS) cube. Depending on the setting of the wave plates, the light is either transmitted to a camera for imaging or reflected on to the rest of the system. This PBS cube is also used to inject test signals, by-passing the telescope. These test signals enter the enclosure on a polarization maintaining single mode fiber, are collimated, and then pass through the PBS cube. Continuing on the optical path, there is a half-wave plate to rotate the polarization of the light before it strikes another PBS cube. The combination of wave plate and PBS are used to make a variable splitter. The light is then sent to the photon counting detector and a reference power meter, with a split ratio dependent on the angle of the half-wave plate. Light that goes along the path to the photon counting detector also passes through a quarter-wave plate. Any light that reflects off the detector face passes through the quarter-wave plate twice and returns back to the PBS cube with the opposite polarization, passes through the cube, and is directed to another camera. The image of the detector face from this camera is used to help align the detector.

4. PHOTON COUNTING DETECTOR

The photon counting detector physically translates signal encoding in the photonic domain to the electronic domain. Ideally, this detector would produce a well defined electrical pulse for every incident photon at the laser transmitter wavelength, produce no false output pulses, and have an invariant time delay between incident photon and pulse output. A single anode hybrid photodiode[4] (HPD) was selected as the photon counting detector for end to end testbed validation for data rates up to 60 Mbps. The HPD comprises an InGaAsP vacuum emission photocathode and a GaAs avalanche diode anode. Much of the device gain, on the order of 10^3, comes from

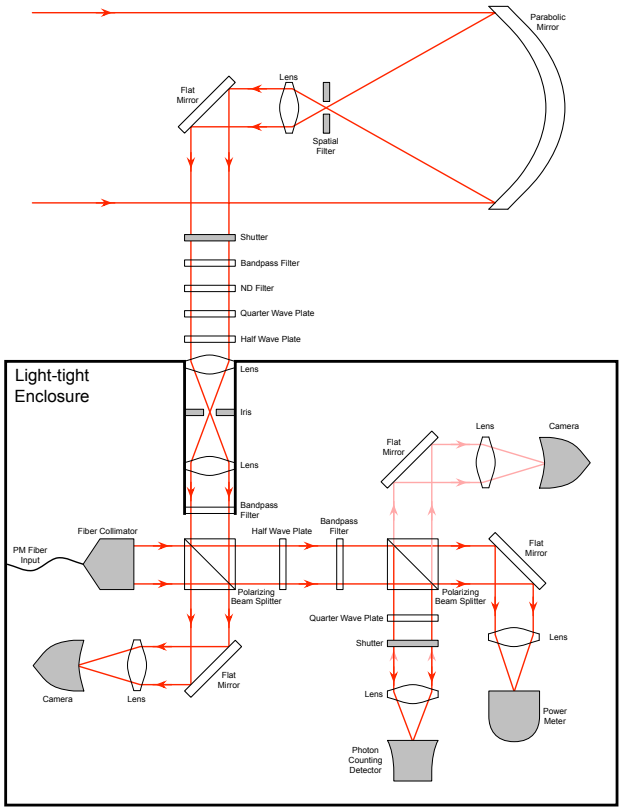

Figure 2. Diagram of the receiver optics used in the end-to-end testbed. Test signals enter the receiver either via the 60 cm telescope show at the top or the single mode polarization maintaining fiber shown at the left.

kinetic energy imparted to a primary photoelectron while traversing an accelerating field region between the photocathode and the anode. The resulting high energy electron bombardment of the anode creates a shower of secondary electrons with a reduced variance (Fano factor) compared to a Poisson distribution. Secondary gain is provided by avalanche multiplication on the order of about 10 within the anode diode, which contributes excess noise to the output signal according to McIntyre avalanche gain statistics. The total internal device gain is on the order of 10^4 with an excess noise factor of 1.1 or less.

An active transferred electron photocathode is used to maximize the quantum efficiency (QE),[4] defined here as the ratio of the number of emitted photoelectrons from the photocathode to the number of photons incident on the active photocathode area. The detection efficiency (DE), defined here as the ratio of the number of countable pulses*at the anode output to the number of photons incident on the active photocathode area, is typically no more than 70% of the QE. Much of the difference between the QE and DE is expected from the basis of atomic backscatter from Ga and As in the anode structure.[5]

Figure 3 depicts the circuit used to measure QE and Figure 4 illustrates typical results for a well performing tube operating at room temperature. Due to thermoelectric offsets at device connections, we have not been able to make reliable cooled quantum efficiency measurements. Figure 5 depicts a similar circuit for measuring detection efficiency. Detection efficiency and dark rate for the same tube illustrated in Figure 4 are shown in Figure 6.

The HPD pulse has a full-width-half-maximum of less than 500 picoseconds and no recovery time. Figure 7

*The rate of countable pulses is a function of electrical waveform at the detector output and the counting algorithm. We measure the efficiency at a sufficiently low flux rate such that a simple threshold crossing decision rule is sufficient to measure nearly all pulses, to render dependence on the counting algorithm negligible. The resulting DE is used as a reference, even though it is known that the DE changes with the incident rate.

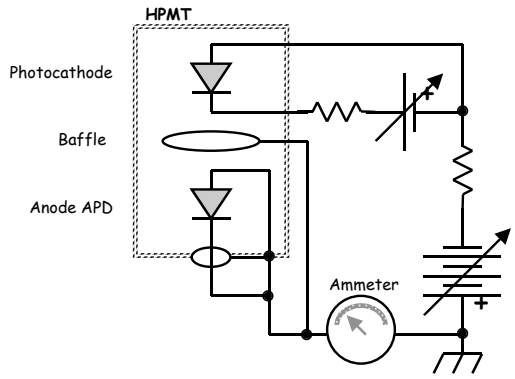

Figure 3. Biasing topology for HPD quantum efficiency measurement.

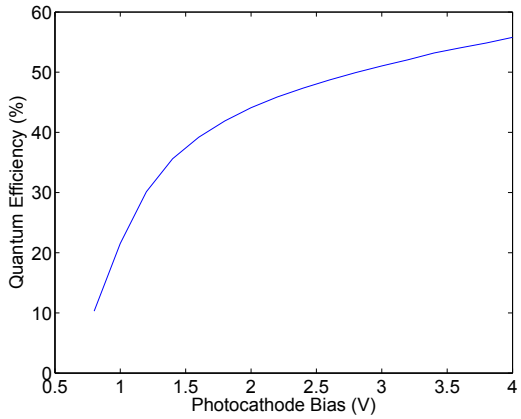

Figure 4. Measured HPD quantum efficiency.

illustrates photon rates as a function of incident rate at the output of three measurement devices: a theoretical counter that registers a count for any pulses with negligible overlap, a measured rate from a commercial counter with a 1 GHz maximum count rate, and a measured rate from the receiver described in Section 5 configured as a flux rate counter with a maximum count rate of 3 GHz.

For the end-to-end testbed, the circuit of Figure 5 was used to generate the input signal to the photon counting receiver. An HPD with a 1 mm active photocathode diameter was selected and operated at room temperature (for convenience) under biasing conditions that yielded a nominal 12% detection efficiency with a 200 KHz dark rate.

5. PHOTON COUNTING RECEIVER

The photon-counting receiver takes as input the voltage output of the photo-detector assembly and produces estimates of the slot and codeword boundaries, the number of photon arrivals in each slot, and the mean signal and noise flux rates. The slot counts, signal, and noise-rates are used to compute symbol log likelihoods, which are transmitted to the decoder. The 16 PPM prototype receiver is composed of a custom front end Photon-Discriminator-Deserializer (PDD) board and a custom FPGA processing board . The receiver may be partitioned functionally into three areas: detection, synchronization, and estimation.

5.1 Detection

The input to the receiver from the photo-detector assembly may be modeled as

$$x(t) = \sum_k \alpha_k p(t - \tau_k) + n(t),$$

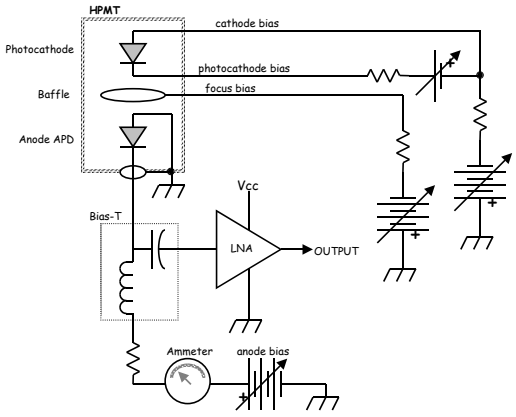

Figure 5. Photon counting biasing topology for the HPD detector.

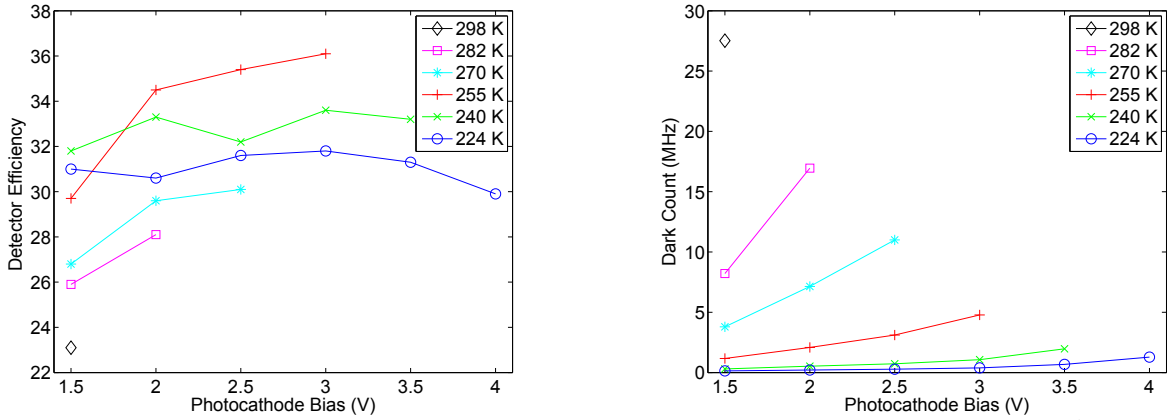

Figure 6. Dark count rate and detection efficiency as a function of bias and temperature for an InGaAsP HPD.

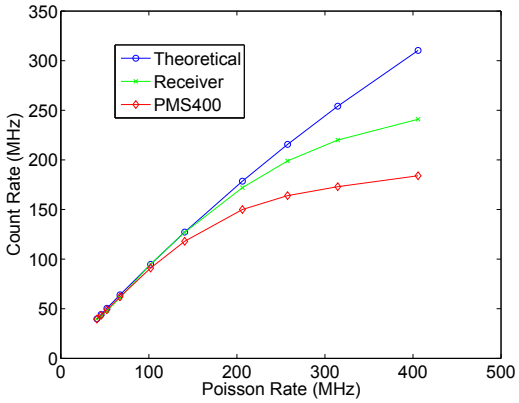

Figure 7. HPD Count rate as a function of incident rate. 'Theoretical' illustrates a theoretical counter which records a count for pulses with negligible overlap, 'Receiver' is the observed rate using the Receiver described in Section 5, and 'PMS400' is from a commercially available counter.

a superposition of bandlimited pulses, $p(t)$, each corresponding to an absorbed photon or dark event with an unknown arrival time τ_k, a gain term α_k that is randomly distributed, and additive thermal noise $n(t)$. In order to estimate the number of events within a slot, the receiver must first reliably detect the locations of the pulses. To accomplish this, the signal $x(t)$ is sampled at a fixed rate of $T_{slot}/8$. Each sample is compared to a selectable threshold, producing a binary sequence x_n, where x_n is 1, or 'high' when the threshold is exceeded and 0, or 'low', otherwise. Since the detector pulse is bandlimited with a duration greater than the sample time, over counting from multiple threshold crossings produced by a single detector pulse can degrade performance.[6] To emulate edge-detection and avoid over-counting, a photon arrival is declared only on a low-to-high transition. Letting y_i be the estimated photon count corresponding to x_i, we have

$$y_i = \begin{cases} 1 & x_i = 1 \cap x_{i-1} = 0 \\ 0 & otherwise \end{cases},$$

At high flux rates, this leads to saturation of the receiver, since photon arrivals resulting in overlapping pulses are not counted [†]. Given the sample decision sequence and knowledge of the slot boundaries the eight sample decisions that represent each slot are summed to form a slot count k_i. These slot counts are passed on to the decoder assembly and are used by the receiver to synchronize the receiver slots and symbol boundaries and to estimate the receiver operating point.

5.2 Synchronization

Due to differences in the reference oscillators of the transmitter and receiver, clock references at the transmitter and receiver will differ and change with time. For a link to a spacecraft, Doppler will also contribute to clock drift. In order to demodulate the PPM symbols and decode the SCPPM codewords, the receiver must obtain knowledge of the slot, symbol, and code-word boundaries. To acquire and synchronize the slot and symbol boundaries in the received PPM signal an inter-symbol guard-time is employed.[7] This synchronization technique introduces an empty slot between each transmitted symbol, or after a given number of slots. It is utilized because it does not expend transmitter power, it allows for a low complexity receiver design, and it may be implemented in a manner that is independent of the PPM order. An estimate of the delay offset of the PPM symbol boundary is formed by performing a circular correlation of the slot count sequence with the inter-symbol guard time basis for each of the possible slot locations of the empty slot.[7] For our system, we utilized PPM order $M = 16$, and a single guard-time slot per symbol.

During acquisition, consecutive estimates of the delay offset are used to generate an estimate of the frequency offset between the reference symbol rate and the received symbol rate. These estimates are used in a first order feedback loop to vary the sample clock frequency to reduce this frequency offset to an amount appropriate for subsequent tracking of the delay offset. After acquisition, the delay offset estimate is fed into a second order type II feedback control loop that tracks the location of the slot and symbol boundaries. The bandwidth of this loop is set by varying the length of the correlation, the update rate of the loop, and the bandwidth of the loop filter.

Each codeword consists of 3780 PPM symbols. Once symbol synchronization is established, the receiver performs codeword synchronization. To determine the codeword boundary, a search of the 3780 possible codeword boundary offsets is performed. For each codeword offset the CRC, incorporated into each codeword for error detection, is checked. In the case where no codewords are passing the check the codeword offset is incremented and the process continues until the codewords begin passing the CRC check indicating the correct code word boundary has been found. It is important that a modified CRC[8] be used, otherwise trailing zeros in the codeword may lead to multiple valid synchronization locations. By requiring multiple successive CRC checks to pass before declaring sync, one may drive the probability of false sync to effectively zero. Once the codeword boundary is known it does not change given the slot boundaries remain synchronized.

[†]Saturation loss may be mitigated by applying a decision directed threshold.[6]

5.3 Estimation

Let p_0 and p_1 denote the probability mass functions of the counts for nonpulsed (noise) and signal slots, respectively. In the receiver, these mass functions are modeled as Poisson, and the photon counts in slots are presumed to be conditionally independent given the transmitted symbol. Given these assumptions, the PPM symbol log-likelihood-ratio (LLR) for a symbol with a pulse in slot i is given by

$$\text{LLR} = \ln\left(1 + \frac{n_s}{n_b}\right) \cdot k_i + \text{ constant}, \tag{1}$$

where k_i is the photon count for slot i, n_s is the average number of signal counts in a pulsed slot, and n_b is the average number of background counts per slot. The constant is common to all symbols with the same collection of slots, and may be neglected, since the receiver may offset any collection of log-likelihoods by a constant with no change to the performance (only relative values are relevant). To evaluate (1), estimates of n_s and n_b are needed. Estimates are obtained by performing a circular correlations of the received slot counts with the inter-symbol guard-time basis. The circular correlation of the received signal slot counts, k_i, with the inter-symbol guard time basis for an offset of iT_s is[7]

$$\mathbf{z}_i = \sum_{n=0}^{N\cdot(M+1)-1} k_n \delta_{(n \bmod (M+1)),i} \qquad i = 0, 1, \cdots, M,$$

where $\delta_{n,m}$ is a Kronecker delta function, N is the number of symbols in the correlation, and M is the PPM order, and we have assumed there is one guard time slot per symbol. The minimum slot correlation

$$\hat{\mathbf{k}} = \arg\min_i \{z_i, i = 0, 1, \ldots, M\},$$

is used as an estimate of the location of the inter-symbol guard time slot. Slot synchronization at this point is not required. It is sufficient that the inter-symbol-guard-time slot be partially contained in the slot pointed to by $\hat{k}$. As the intersymbol guard time slot contains only background counts and the other slots contain both signal and background counts, by adding and subtracting combinations of these, slots estimates of n_s and n_b can be obtained. For example, with $M = 16$, estimates $\hat{n}_s$ and $\hat{n}_b$ may be formed with

$$\begin{aligned}
\hat{n}_s =& \frac{M}{N}\left(\left(\mathbf{z}_{(\hat{\mathbf{k}}+7) \bmod M+1} + \mathbf{z}_{(\hat{\mathbf{k}}+8) \bmod M+1} + \mathbf{z}_{(\hat{\mathbf{k}}+9) \bmod M+1}\right)\right. \\
& \left. - \left(\mathbf{z}_{(\hat{\mathbf{k}}-1) \bmod M+1} + \mathbf{z}_{\hat{\mathbf{k}}} + \mathbf{z}_{(\hat{\mathbf{k}}+1) \bmod M+1}\right)\right), \\
\hat{n}_b =& \frac{1}{4N}\left(\mathbf{z}_{(\hat{\mathbf{k}}+7) \bmod M+1} + \mathbf{z}_{(\hat{\mathbf{k}}+8) \bmod M+1} + \mathbf{z}_{(\hat{\mathbf{k}}+9) \bmod M+1} + \mathbf{z}_{(\hat{\mathbf{k}}+10) \bmod M+1} - 4 \cdot \hat{n}_s\right).
\end{aligned}$$

Assuming that the inter-symbol-guard time slot is partially contained in the slot pointed to by $\hat{k}$, one can use the first moments of z_i[7] to show the estimates of n_s and n_b are unbiased.

6. ERROR-CONTROL-CODE

The data is encoded with an error-correction-code (ECC) prior to transmission. The ECC significantly improves the power efficiency of the channel, allowing performance close to channel capacity, the least upper bound on the achievable performance. This ECC, the serially-concatenated pulse-position-modulation[3] (SCPPM) code, is specifically tailored to function with PPM, by incorporating the modulation as part of the code. The code consists of a convolutional code, whose output is bit-interleaved and then encoded with a bit-accumulator, and then mapped to PPM symbols.

On reception of codeword boundaries and symbol LLRs from the receiver, a hardware decoder approximates maximum-likelihood decoding of the ECC with an iterative decoding algorithm.[9,10] The decoder is implemented through a distributed decoding architecture, where each element contains a single SCPPM decoder that when busy passes the LLRs corresponding to a single codeword on to a subsequent decoder. Each decoder element

is implemented on a custom FPGA processing board and . the throughput may be increased by adding more elements.

The fixed-point hardware decoder has a 0.2 dB implementation loss relative to a floating point implementation at a background noise rate of 0.2 detected photons/slot. The implementation loss varies with the background noise, with smaller losses at lower background rates.

7. DEMONSTRATED PHOTON EFFICIENCIES

All data presented in this section is from an operating point with a slotwidth of 2.67 nsec, $M = 16$ PPM, a single guard-time slot per symbol for synchronization, and a noise level of $n_b \approx 4 \times 10^{-3}$ mean noise photons per slot. This yields a data throughput of 44 Mbps. Figure 9 illustrates the codeword error rate as a function of the mean signal photon count at the output of the receiver. However, reporting photon efficiency relative to the receiver photon count can be misleading, since it can mask losses in the receiver algorithm. A more meaningful measure is performance as a function of the incident photon rate. Figure 8 illustrates a block diagram of the test-bed configuration used to extract an estimate of the incident photon flux rate. Photon flux rates were measured at point A, utilizing the receiver counting algorithm, and at B, utilizing a power meter. We would like to factor out the photon-detection-efficiency (PDE) of a particular device and to focus on the efficiency of the receiver and decoder algorithms, given the detector dark rate, any multiplicative noise, impulse response, and any system thermal noise. To that end, we define a point C, which reflects the detected photon rate (Figure 9 reports performance relative to the flux at A). To model the losses, we first develop a model of the nonlinear receiver counter losses, which leads to saturation at high flux.

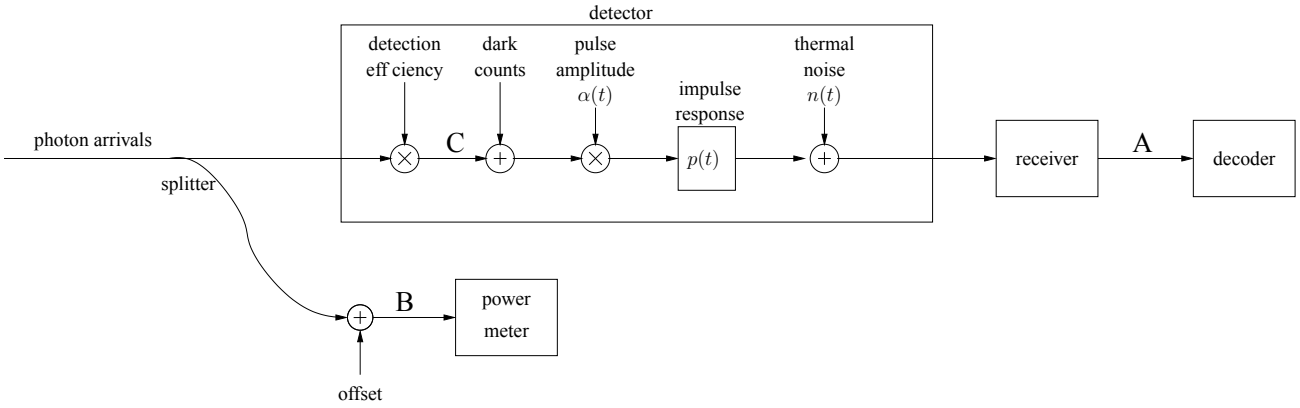

Figure 8. Model illustrating points of measured (A,B) and derived (C) flux rates

7.1 Nonlinear Receiver Losses

Recall the receiver counter maps samples to 'high' or 'low' values relative to a threshold. A photon count is registered on a low-to-high transition. This edge detection prevents over-counting, but leads to saturation at high flux rates. The receiver counter may be modeled as the K-state Markov chain illustrated in Figure 10. Each state-transition corresponds to a duration of one sample. State 0 is the reset state and K is the pulse width in samples–the number of 'high' samples per pulse. The counter in state $i > 0$ has a minimum of $K - i$ remaining samples to reset. If a pulse arrives prior to reset, the counter state returns to $i = 1$. To simplify analysis, photon arrivals are presumed to occur at the beginning of a slot, and multiplicative-pulse-amplitude and thermal noise are presumed negligible, insofar as they do not alter the outcome of comparisons to the threshold.

At high flux rates, this counter will saturate. The effect of counter saturation is illustrated in Figure 11, which plots detected power at A as a function of the measured power at B. The attenuation along the detected path may be attributed to the splitter, detector losses (primarily PDE) and missed detections in the receiver counter. The loss at low flux rates is approximately linear, allowing a linear fit at low flux, which we denote as an 'ideal receiver'. The low flux fit allows us to remove any bias in the power meter. The dark count rate of the detector was negligible in our operating regime.

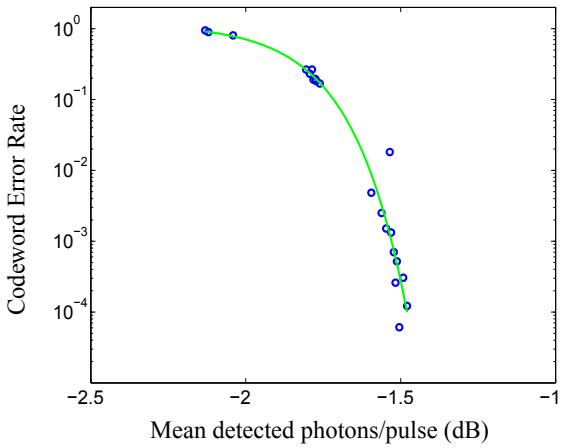

Figure 9. Measured codeword error rate as a function of detected photons per pulse reported at receiver. $M = 16$, slotwidth 2.67 nsec, rate 1/2 code.

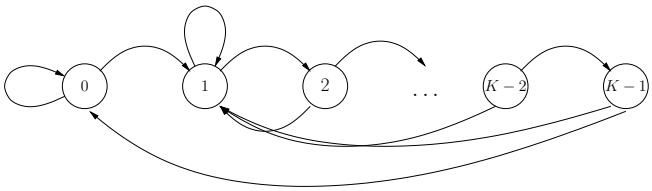

Figure 10. Markov model of receiver edge-detector

Curves illustrating the detected power predicted by the receiver model of Figure 10 are overlaid on the measured power for various detector pulsewidths. We see a good agreement with the model for a pulsewidth of $K = 5$ (corresponding to a pulsewidth of 1.67 nsec, in agreement with observations). Using the Markov model with $K = 5$ allows us to back out the detected power (A) as a function of the incident power (C), illustrated in Figure 12 . This curve is well fit as a fifth degree polynomial. We use this polynomial to convert detected rates (A) to incident rates (C). This allows us to plot performance versus incident rates, illustrated in Figure 13, demonstrating a photon efficiency of ≈ 1.8 bits/photon. Also illustrated in Figure 13 is the capacity[‡] for an ideal channel (ideal translation of incident photons to photon counts), the corresponding code performance on an ideal channel, and the performance of the channel model of Figure 10 with $K = 5$. We see a good fit between the measured performance and model prediction. Losses relative to the ideal channel may be attributed to finite-quantization, jitter, thermal noise and multiplicative noise, which are not accounted for in the ideal channel. At these low rates there are only small losses due to counter saturation.

8. CONCLUSIONS

In this paper we have described the components of a real-time operational end-to-end free-space optical communications link. The link was operated in day and night conditions over a 100 m range, demonstrating real-time data rates of 44 Mbps and photon efficiencies of 1.8 bits/photon. These tests have validated many of the subsystems required for a photon-starved deep-space or near-earth optical communications link.

ACKNOWLEDGMENTS

The research described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

[‡]The block error rate bound is from the weak converse to the channel coding theorem.[11]

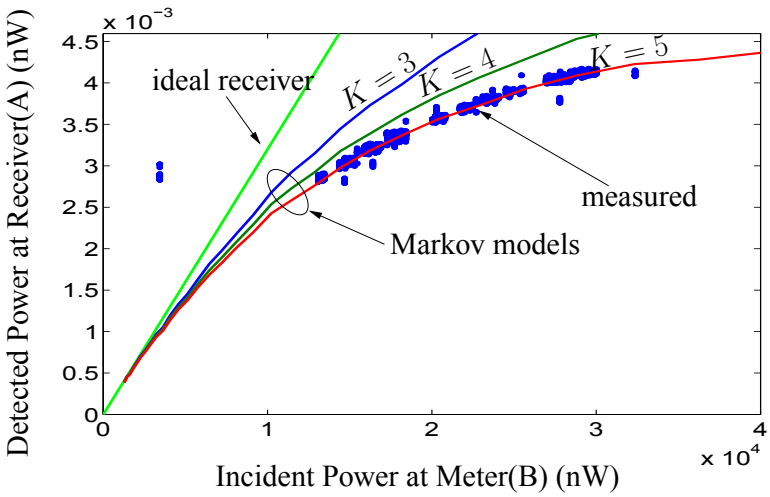

Figure 11. Measured detected power (A) as a function of measured incident power (B). Comparisons with predictions from Markov-model and ideal receiver (with linear losses only).

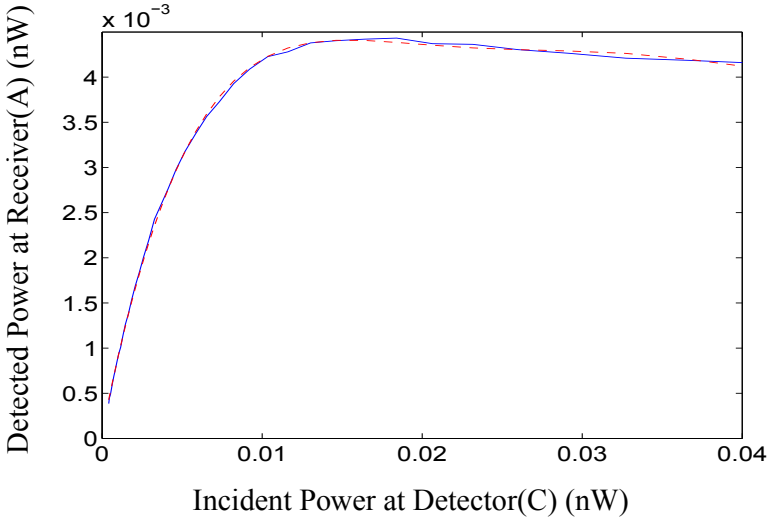

Figure 12. Modeled detected power (A) as a function of modeled incident power (C), from Markov model.

REFERENCES

[1] Wright, M., Zhu, D., and Farr, W., "Characterization of a high power fiber MOPA laser as an optical communications transmitter," *IPN Progress Report* **42–161** (2005).

[2] Wright, M. and Kovalik, J., "A fiber based MOPA laser transmitter for optical communications," *IPN Progress Report* **42–171** (2007).

[3] Moision, B. and Hamkins, J., "Coded modulation for the deep space optical channel: serially concatenated PPM," *IPN Progress Report* **42-161** (2005).

[4] Rue, R. A. L., Davis, G. A., Pudvay, D., Costello, K. A., and Aebi, V. W., "Phton counting 1060-nm hybrid photomultiplier with high quantum efficiency," *IEEE Electron Device Letters* **20**, 126–128 (1999).

[5] Klein, W., Harder, M., and Pollehn, H., "Influence of backscattered electrons on the resolution of transmission secondary electron emitters," *Applied Optics* **11**, 2337–2339 (1972).

[6] Quirk, K. J. and Milstein, L. B., "Optical PPM detection with sample decision photon counting," in [*Proceedings IEEE Global Telecommunications Conference*], (November 2005).

[7] Quirk, K. J., Gin, J. W., and Srinivasan, M., "Optical PPM synchronization for photon counting receivers," in [*Proceedings of the 2008 IEEE Military Communications Conference*], (November 2008).

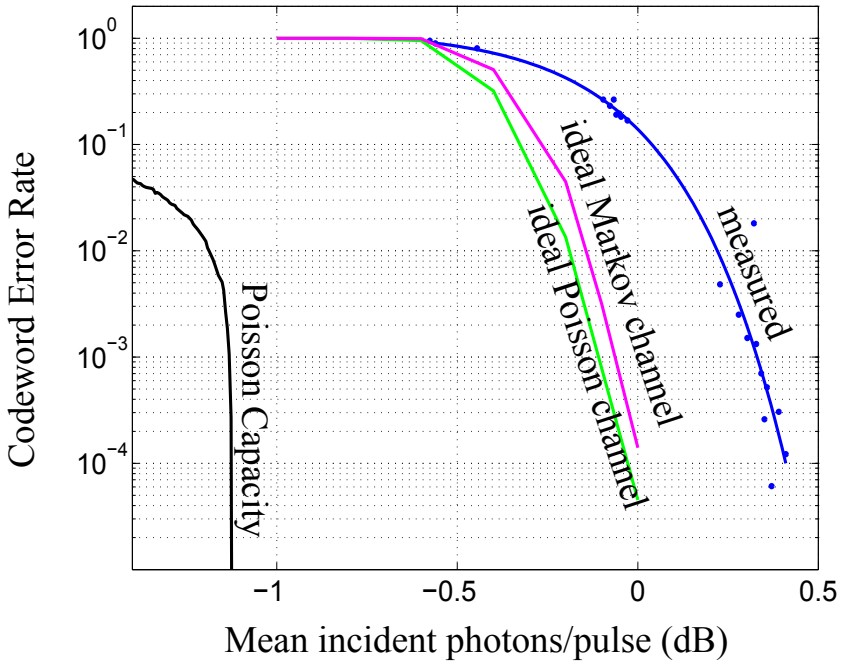

Figure 13. Performance as a function of incident photon rate: ideal Poisson, measured performance, and Markov-modeled Poisson. $M = 16$, slotwidth 2.67 nsec, rate 1/2 code.

[8] Berry, T., "A note on the modified CRC," *ACM SIGCOMM Computer Communication Review* **20**, 12–17 (October 1990).

[9] Cheng, M., Moision, B., Hamkins, J., and Nakashima, M., "Implementation of a coded modulation for deep space optical communications," in [*Proceedings IEEE Global Telecommunications Conference*], 1–5 (November 2006).

[10] Cheng, M., Nakashima, M., Moision, B., and Hamkins, J., "Optimizations of a hardware decoder for deep-space optical communications," in [*IEEE Transactions on Circuits and Systems*], **55**, 644–658 (March 2008).

[11] Cover, T. M. and Thomas, J. A., [*Elements of Information Theory*], Wiley (1991).