

Video Semaphore Decoding for Free-Space Optical Communication

M Last, B Fisher, C Ezekwe, S Hubert, S Patel, S Hollar, B Leibowitz, KSJ Pister

Berkeley Sensor and Actuator Center
497 Cory Hall
University of California, Berkeley
Berkeley, CA 94720-1774

ABSTRACT

Using real-time image processing we have demonstrated a low bit-rate free-space optical communication system at a range of more than 20km with an average optical transmission power of less than 2mW. The transmitter is an autonomous one cubic inch microprocessor-controlled sensor node with a laser diode output. The receiver is a standard CCD camera with a 1-inch aperture lens, and both hardware and software implementations of the video semaphore decoding (VSD) algorithm. With this system sensor data can be reliably transmitted 21 km from San Francisco to Berkeley.

Intelligent encoding and video processing algorithms are used to reject noise and ensure that only valid message packets are received. Dozens of independent signals have been successfully received simultaneously.

A software implementation of the VSD algorithm on a Pentium computer with a frame grabber was able to achieve an effective frame rate of 20 fps, and a corresponding bit rate of 4bps. This bit rate is adequate to transmit real-time weather sensor information.

The VSD algorithm has also been implemented in a custom hardware board consisting of a video ADC, RAM, Xilinx FPGA, and a serial output to PC. This system successfully operated at 60 fps with a bit rate of 15 bps.

Keywords: Free Space, Laser, Wireless, Optical, Communication, Long Range, San Francisco, Low Power

1. INTRODUCTION

1.1 Free Space Laser Communication

For applications where line of sight is available between a transmitter and a receiver, free-space optical communications systems can present tremendous advantages over their RF counterparts. Perhaps most important, optical power can be collimated in tight beams even from small apertures. Diffraction enforces a fundamental on the divergence of a beam, whether it be from an antenna or a lens. Laser pointers are cheap examples demonstrating milliradian collimation from a millimeter aperture. To get similar collimation for a 1 GHz RF signal would require an antenna 100 meters across, due to the difference in wavelength of the two transmissions. As a result, optical transmitters of millimeter size can get antenna gains of one million or more, while similarly sized RF antennas are doomed by physics to be mostly isotropic. With this kind of transmitter gain, microwatt signals can be sent over multi-kilometer distances with a strong SNR. Micro Electro-Mechanical Systems (MEMS) technology has the potential to integrate a laser diode, collimating optics and a beamsteering mirror into a package with a volume of only a few cubic millimeters¹. Using a MEMS steered laser transmitter, the next generation of infrared ports may have ranges of kilometers instead of the current centimeters, or data rates of a few gigabits per second instead of only a few megabits per second.

The high antenna gains possible using optical frequencies have advantages for the receiver end of the link as well. Since each pixel in an imaging optics system has a very narrow field of view, and high levels of integration allow many optical transducers on a chip, an imaging receiver such as a CCD or photodiode array can function as a collection of independent optical receivers. In this scenario, each pixel in an imaging array stares at a very narrow field of view, limiting background noise and focusing most of the received optical power onto a very small receiver area. If done properly, this can result in an extremely high signal-to-noise ratio limited only by the noise in the receiver electronics. Additionally, since each pixel in the array receives light from a

Correspondence: M Last; Email: mattlast@eecs.berkeley.edu; Telephone: (510)643-2236.
KSJ Pister; Email: pister@eecs.berkeley.edu; Telephone: (510)643-9268

different portion of the field of view of the receiver, spatially separated transmitters image onto different pixels in the array, forming a type of spatial division multiple access (SDMA), and can therefore transmit freely without interfering with one another.

In the next section of this paper, the feasibility of low-power, long-range optical communication links is discussed. To experimentally verify these conclusions, three optical communication systems were built and tested (Section 3). The results of these tests are described and discussed in Section 4.

2. THEORY

2.1 Signal Power

The single largest factor determining how much signal power will arrive at a receiver in a long-range optical link is the divergence angle of the transmitted beam. For a diffraction-limited optical system, the divergence full-angle of the output beam in an ideal system is given by:

$$\theta_{div} = \frac{1.22\lambda}{d} \quad [1]$$

where λ is the center wavelength of the laser and d is the limiting aperture of the system. The power density D in Watts/Steradian is given by:

$$D = \frac{P_o}{2\pi \left(1 - \cos\left(\frac{\theta_{div}}{2}\right)\right)} \quad [2]$$

The intensity of signal at range R is the power/area: $I_s = \frac{D}{R^2}$. This describes a spherical cap of constant power. The receiver aperture defines the portion of the transmitted power that is focused onto the photodetector. The power at the photodiode is given by $P_{pd} = I_s A_r$, where A_r is the area of the aperture of the receiver.

2.2 Noise Sources

The main source of noise in this type of link derives from the presence of reflected sunlight in the field of view of the receiver. This ambient light power causes two problems. First, it adds a large DC offset. Ideally this does not degrade the signal, but in any practical receiver with finite output swing it reduces the usable dynamic range of the receiver. Second, since the photons arrive with a Poisson distribution, this DC offset contributes shot noise power directly proportional to the ambient light power. When this optical input is integrated as collected charge over some exposure period, the result is a noise variance that is equal to the total amount of charge collected. Therefore the optical SNR is the square of the received signal charge divided by the received charge due to ambient light. Although the use of narrow-band filters alleviates this ambient light problem, even with a 15nm full-width, half-maximum interference filter the noise power from the sun may still be unacceptable. The use of an imaging receiver can further alleviate this problem.

Imaging receivers dramatically decrease the ratio of ambient optical power to received signal power. Ideally, the imaging receiver will focus all of the received power from a single transmission onto a single photodetector. If the receiver has an $N \times N$ array of pixels, then the ambient light that each pixel receives is reduced by a factor N^2 compared with a non-imaging receiver. Typically, using a value for N between 8 and 32 makes the shot noise power due to ambient light negligible compared with the electronic noise in the analog electronics. The remaining DC offset from the ambient light may still be comparable to the signal strength without degrading the receiver performance appreciably.

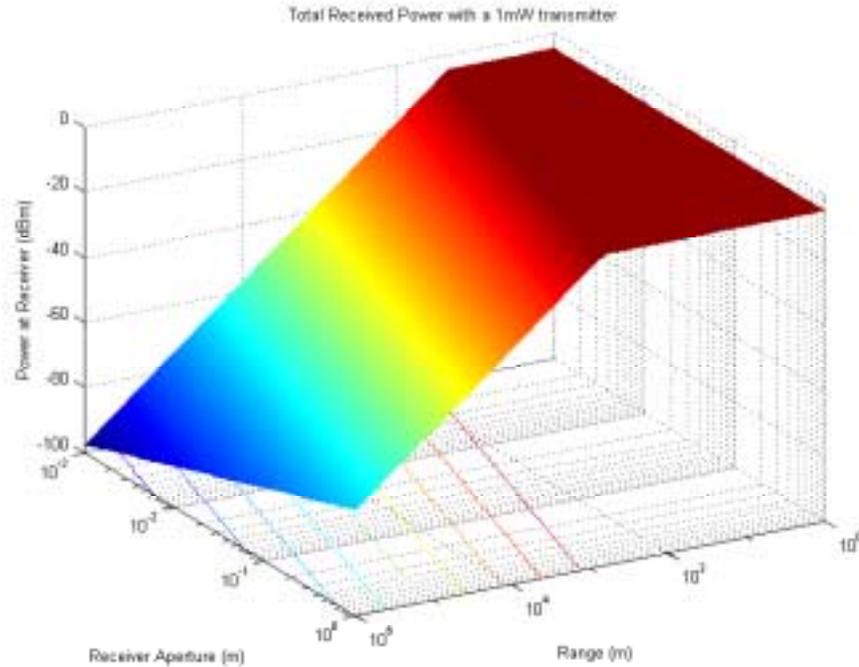


Figure 1: Graph of received power versus range and receiver aperture. The graph flattens out for short ranges or large receiver apertures when the beam spot is smaller than the receiver aperture. This was calculated assuming that the transmitter broadcasts using 1mW (0dBm) from a 0.5mm diameter aperture.

Electronic noise is generally dominated by a few noise sources at the front end of the analog electronics. In a CMOS process, these noise sources are the thermal channel noise and flicker noise of the input transistor and its load device as well as noise sources from any feedback devices, such as thermal noise in the feedback resistor if a transimpedance amplifier is used at the front end. In addition, the analog-to-digital (ADC) converter that digitizes the video signal adds quantization noise to the signal.

3. SYSTEM DESCRIPTION

Three optical communication systems developed in our lab are presented in this section. Each system consists of a laser transmitter subsystem and a receiver subsystem. The following two sections describe the differences between the systems.

3.1 Laser transmitter

A miniature weather station incorporating light intensity, humidity, temperature, and barometric pressure sensors was built^{2,3}. A microcontroller chip on this board controlled sensor data acquisition, filtering and packetization of the acquired data, and the low-level control of the laser diode. The sensor data was acquired once per second. Each packet consisted of a training sequence, followed by four eight-bit sensor values, followed by a sixteen-bit cyclic redundancy check (CRC) that was added to the end of each packet to ensure the validity of received data. In order to ensure a nearly constant average optical power level, 3 bit transitions are introduced into each byte by inverting the 2nd, 4th, and 6th bit as shown in Figure 2.

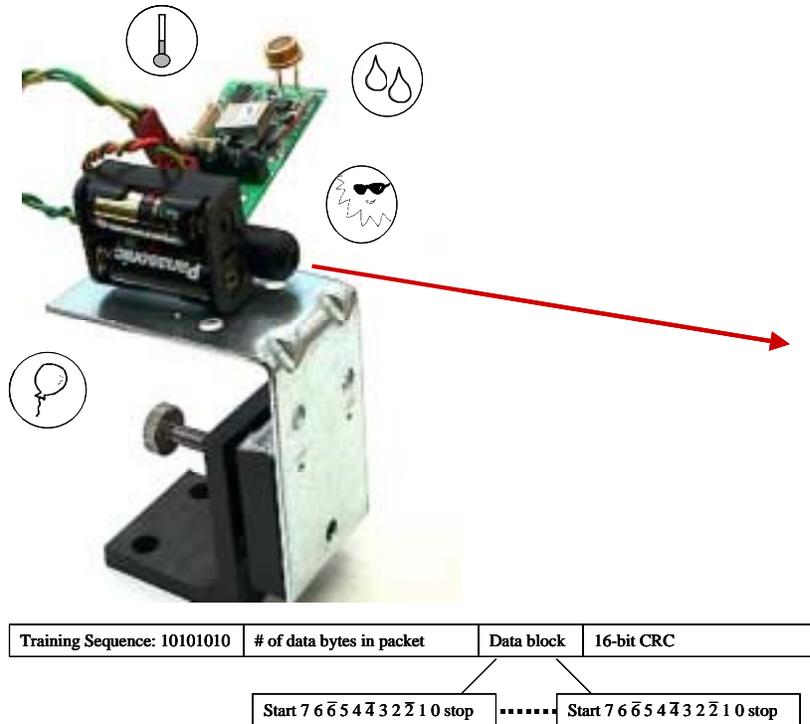


Figure 2: Miniature weather station with integrated laser transmitter. Sensors include light intensity, humidity, temperature, and barometric pressure. The station measures 1x3x0.5 inches. Each packet of data sent over the laser consists of an 8-bit training sequence, a byte for the number of data packets to follow, a byte of data with the current value for each sensor, and a 16-bit cyclic redundancy check (CRC) to ensure data consistency.

A standard laser pointer had its battery case removed and was wired up to the microcontroller. This laser pointer drew 30mA at 3.0V and radiated 3mW of optical power while on. The best method of alignment of the transmitter for long-range operation was to super-glue the laser to a riflescope and align the laser spot to the crosshairs at a range of at least 15m. In order to hold the transmitter steady for long periods of time and to avoid alarming the police during the experiment, the laser and riflescope combination was not attached to a rifle, instead, a small optical breadboard and large gimbal-mounted mirror were used to mount and aim the riflescope and laser.

This system is sufficient to provide accurate laser pointing at a range up to 21 km. However, manual alignment of the laser requires skill and patience. Open-loop alignment, where the operator simply lines up the receiver in the crosshairs of the rifle scope, is sufficient for accurate targeting of the laser at ranges up to approximately 15 km. For longer links, feedback from the receiver station regarding signal strength is necessary. With this system, this is accomplished using cellular phones and a technician at each end of the link. An automatic alignment system, where the laser transmitter raster scans its beam and a matching transmitter at the other end of the link provides feedback on received signal strength, was developed as an alternative. This is described in section 3.4.

3.2 PC-based Video Semaphore Decoder

The receiver in the PC-based VSD consists of a CCD camera and a laptop computer equipped with a PCMCIA video framegrabber card. The CCD camera is equipped with a variable zoom lens and a 15nm Full-Width, Half-Maximum (FWHM) interference filter to minimize extraneous light from other light sources in the field of view. The video framegrabber digitizes images from the video stream as fast as the images can be processed. Due to limitations in the speed of the computer, this video capture rate rarely exceeds 20 frames per second. This low image capture rate provides an upper bound on the speed of the communication link and provides the impetus for the development of the system described in section 3.3.

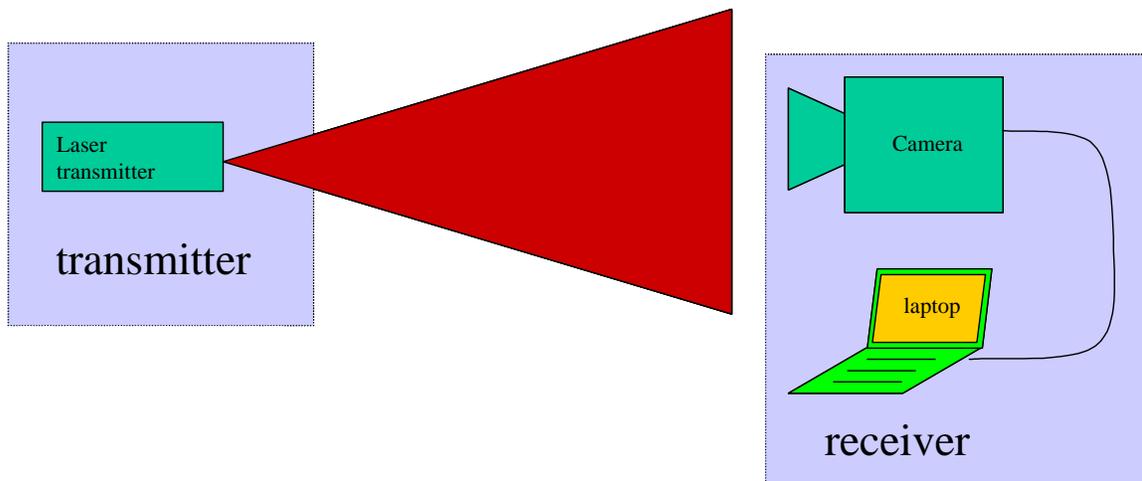


Figure 3: Diagram of system described in sections 3.1 and 3.2 of this paper. A laser transmitter sends data to a CCD camera. This camera is connected to a video framegrabber in a notebook computer. Software in this notebook computer decodes the data and saves it to disk.

Custom software developed in our lab controls the video framegrabber and performs the image manipulation steps needed to locate, filter, and decode data streams present in the field of view of the camera. As each new video frame is digitized, the previous frame is subtracted from it. This high-pass filters the image stream in time, highlighting bit transitions. This subtracted image is then binarized using an adaptive threshold. A size filter removes spots smaller than a threshold value input by the user, and all remaining spots are considered transmitter candidates and are labeled. Each spot is assumed to be the result of one laser transmitter, and therefore the pixel intensity values are correlated. Using the binary thresholded image as a pixel mask, the pixels in each spot are combined using Maximal Ratio Combining to come up with an intensity value for each transmitter. A decision is then made about whether each transmitter is on or off based on a threshold computed separately for each transmitter candidate. This threshold value can be determined using the known bit sequence present at the beginning of each packet of data. Each packet of data is examined to ensure that a valid CRC is present. In the absence of this CRC, the packet is discarded. This results in a highly reliable link, as the chances of an invalid data packet being misinterpreted as valid data is only 1 in $2^{16}=65,536$, corresponding to an invalid packet being accepted as valid on average once every 473 hours of continuous transmission of invalid data. However, since most noise sources in natural scenes are time-varying and highly random (e.g. specular reflections from sunlight, trees moving in the wind), noise sources tend to disappear long before the mean time to error. This receiver has proven to be fairly robust, and is able to handle multiple transmitters (up to several dozen simultaneous transmissions have been successfully decoded), slowly moving transmitters (less than one spot width per frame of motion), and variations of each transmitter's intensity over time.

3.3 Xilinx-based Video Semaphore Decoder

The major shortcoming in the software-based receiver described in section 3.2 is the limited data link bitrate. The sampling rate in a video receiver is the frame rate of the rate-limiting step of the receiver. In our case, the two main rate-limiting steps are the rate at which the data can be written to main memory by the framegrabber card and the rate at which the data is processed. If the receiver were not rate limited by the framegrabber or the image processing algorithms, the video camera used has a frame rate of 30Hz (60Hz interlaced). This results in an increase in the maximum data rate from 10 to 30 bits per second. However, if we use a high-speed camera, we can increase this maximum rate up to several hundred bits per second.

By implementing the image processing algorithms in hardware instead of software, the image processing tasks can be pipelined and each pixel can be processed at the pixel scan rate, up to the clock frequency of the hardware used. There are several approaches to implementing this kind of pipelined video processor; a DSP chip or a Field-Programmable Gate Array (FPGA) can be used. The FPGA can be made into a pipelined pixel processor, where all eight bits of a pixel are read in and the results of a previous calculation are output each clock cycle. A DSP however would require multiple clock cycles to implement all required signal-processing steps. A workaround for this would be to have a separate DSP chip for each step. This is essentially the same

thing as implementing the pipeline in firmware on an FPGA – the processing tasks are each given dedicated hardware. However, a single chip solution results in savings in power, cost, and board space.

It was decided that the Xilinx-based solution was better for our needs. A board was designed in-house, fabricated by Advanced Circuits⁴, and assembled in our lab. The board consists of an eight-bit analog-to-digital converter (ADC), a Xilinx FPGA, 256 KB of SRAM, a 24 MHz clock for the Xilinx, a video synch separator, and a level shifter to implement an RS-232 serial port connection. A system diagram is shown in Figure 4.

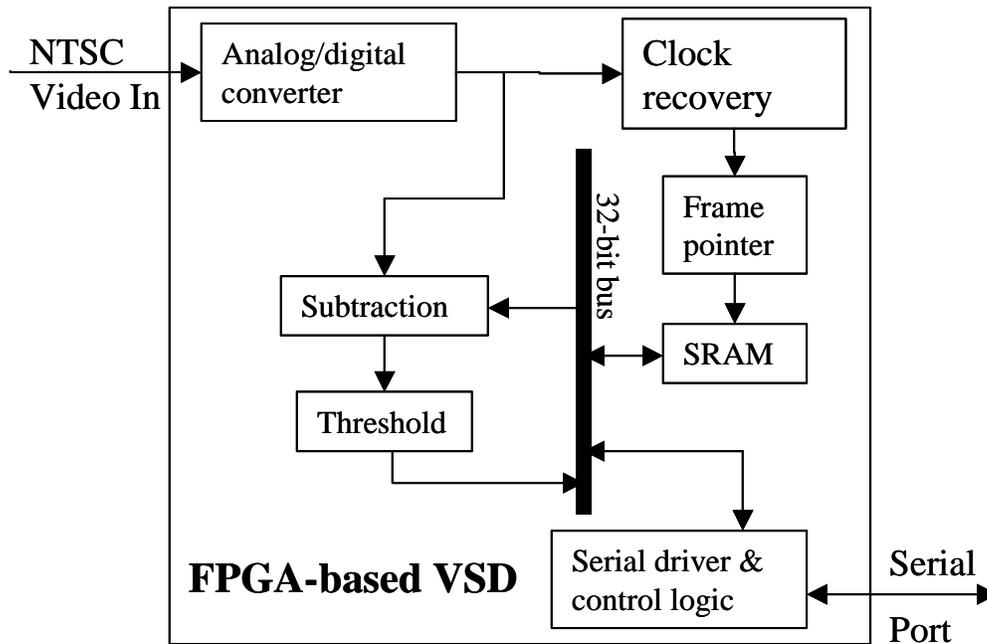


Figure 4: System architecture of FPGA-based VSD

The algorithm used by the FPGA-based system is not as fully developed as the software-based VSD. This is mainly due to the inherently slower pace of concurrent hardware/firmware development. Although capable of decoding transmitters at known pixel locations, the FPGA system relies on the user to locate transmitters in the image. It accomplishes this by sending the captured video image over the serial port at a highly reduced frame rate (limited by the speed of the serial port) and allowing the user to send back pixel locations from which he/she is interested in receiving decoded data. This information is stored in a 32 bit deep frame buffer along with 8 bits for the threshold value, 8 bits for the previous frame value, 4 bits to store the state of the finite state machine used to process each pixel, 3 bits for flags and the 12 previous bits of data.

Each pixel is processed using the following algorithm at the pixel clock rate. The previous value for each pixel location is subtracted from the current value. That subtracted value is compared to the threshold value stored in the frame buffer, and a decision is made whether the bit is a one or zero. A running average between one and zero bits is used to set the threshold value for that pixel. The bit value is stored in memory, and if the user selects that pixel location, a packet is downloaded over the serial line as soon as the CRC for that packet is validated.

Another piece of the FPGA is used to synchronize the clock to the NTSC frame and line rate. The frame rate is determined by counting clock cycles between beam blanking periods. The number of lines corresponds to the number of lines in the frame buffer. Currently, the end of line signal is not detected – instead, each frame is evenly divided into lines, and each line is evenly divided into pixels, whose number is again determined by the size of the frame buffer.

Although it may be possible to push the serial data processing strategy presented in sections 3.2 and 3.3 out into the low thousands of bits per second, this strategy does not continue to scale well past that point. CCD cameras usually use a serial readout strategy, which means that any image processing that occurs must keep up with the pixel readout rate from the camera. As the camera speed increases into the thousands of frames per second, the image processing engine will need to chew through billions of bits per second even at moderate image resolution. It becomes expensive in terms of power requirements and cost to continue to scale this technique to even higher frame rates. The upper limit of the processing engine is the maximum clock frequency that the image processing modules can run at. This is currently around 200MHz for Xilinx FPGAs, yielding a frame rate of 1,000fps at a maximum resolution of approximately 450x450 pixels. Our analysis points in the direction of parallelizing the image processing to continue to scale the speed of imaging receivers. A custom CMOS imager where each pixel is a fully independent megabit/second receiver is also presented in this conference⁵.

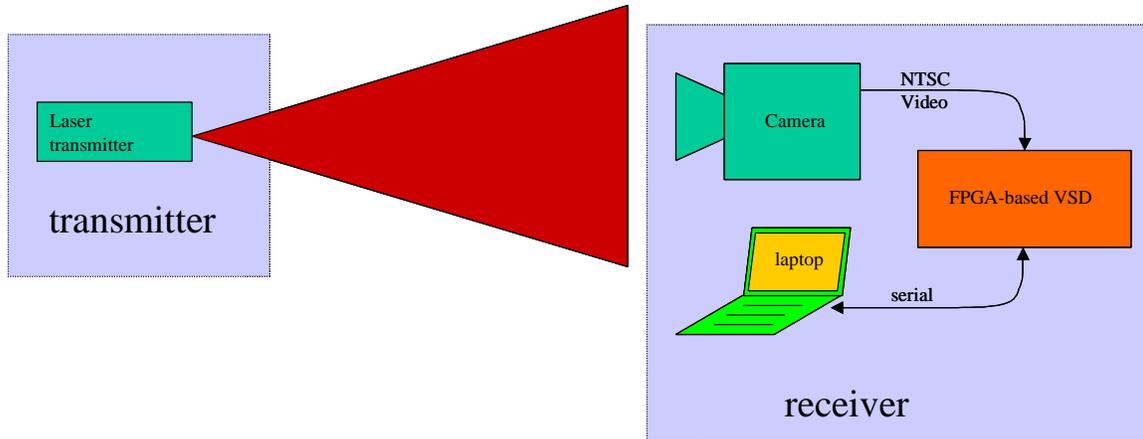


Figure 5: Laser communication system described in sections 3.1 and 3.3 of this paper. Instead of using a notebook computer to perform image processing algorithms, a custom circuit board featuring a Xilinx Field Programmable Gate Array (FPGA) digitizes and decodes the video stream in real-time using a simplified algorithm (section 3.3) and sends the decoded data to the computer over a serial port.

3.4 Integrated Laser Transceiver

In order for free-space laser transmission to become a practical method for communication over any distance, fast and robust automatic link acquisition is essential. In section 3.1, a manually aimed laser transmitter is described. The targeting process described in that section highlights the main drawback to narrow-beam communication, namely, that in order to receive any signal at all, the transmitter needs to be aligned very precisely to the receiver. This is tedious to set up initially, and is subject to long-term drift as environmental conditions change. The use of an automatic link acquisition and optimization system allows for very coarse manual alignment (+/-15 degrees) and maintenance-free operation henceforth; essentially a “fire and forget” link management strategy.

The Motorized Agile Laser Transceiver consists of a transmitter similar to the system described in section 3.1, and a receiver identical to the system described in section 3.2. There are several key differences in MALT that make it significantly more powerful than any of the other systems described so far. Computer controlled pointing of the laser beam eases manual alignment of the laser and enables automated alignment strategies to be investigated. Calibration of the camera to account for tilt and rotation of the field of view relative to the motion of the laser provides a mapping between camera pixel location and laser steering directions, so the laser can be aimed at any point visible through the camera. And finally, by linking the receiver software with the laser aiming software, MALT can automatically align itself to any transmitter whose signal it successfully decodes.

As shown in Figure 7, MALT is made of three main components assembled onto a small optical breadboard. A CCD camera is used as a receiver, a microcontroller-modulated laser pointer is used as the transmitter, and a mirror mounted on a precision-machined two degree-of-freedom gimbal is used to aim the laser beam. The gimbal is the most notable new feature in MALT. Driven by two miniature stepper motors⁶ and a custom motor controller circuit board, this gimbal provides over 15 degrees of

mechanical deflection with better than half a milliradian pointing accuracy (limited by the step size of the motors). The entire gimbal package including motors measures less than 3 cubic inches.

In order to establish a link, one or both of the transceivers (stations A and B) point their lasers at each point in their field of view. As station A scans its laser, it broadcasts its ID number to each pixel. If station B happens to see this signal, it will slew its laser around to point back towards the strongest signal it received from station A, and sends back the station A's ID number, its own ID number, and the strength of the signal that it saw. When station A sees this, it returns its laser to point towards the strongest signal it received from station B, sends station B's ID number, its own ID number, and the strength of the signal that it saw. This completes the handshaking protocol. Further optimization of the link is accomplished by performing a complete raster scan in a smaller search window and then choosing the direction that has the highest signal strength.

In cases where imaging receivers such as video cameras or photodiode arrays are available, this process is greatly simplified since only one end of a given link must perform an exhaustive search for the other end of the link.

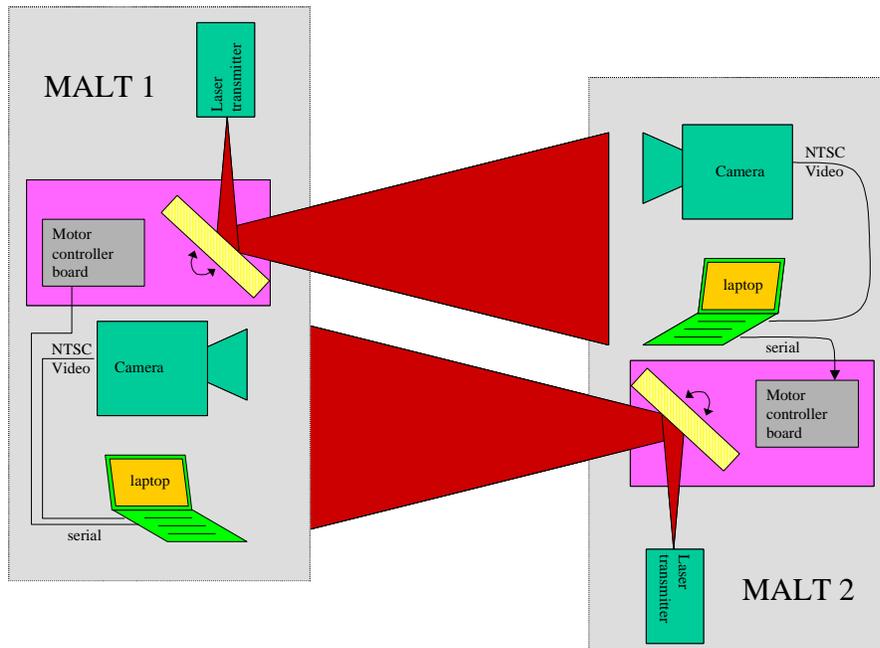


Figure 6: MALT -> MALT comm.

Performance of MALT Steered Beam Transmitter			
Pointing Resolution	480x435 μ rad	Beam Divergence	1x2 mrad
Range of Motion	314x262 mrad	Beam Optical Power	1.5mW (avg)
Settling Time	~0.5ms	Number of Spots	654x602

Table 1: Performance of MALT System

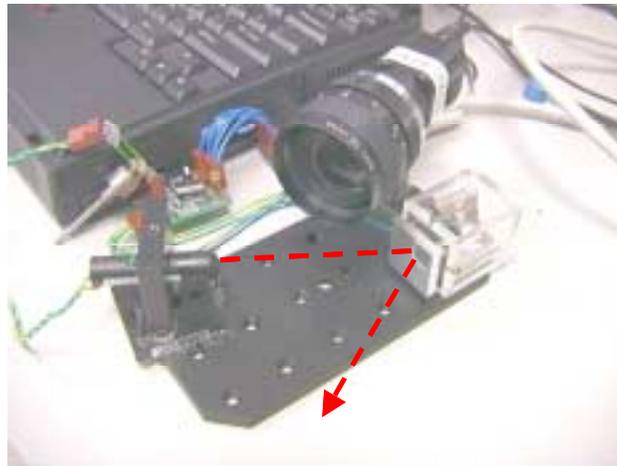


Figure 7: Motorized Agile Laser Transceiver (MALT). The laser (short cylinder on left) emits a collimated beam that bounces off the rectangular mirror glued to the aluminum plate. The plate is moved using two miniature stepper motors made by Smoovy⁶. The video camera serves as a receiver. MALT can aim the laser to reply to any transmitter in the field of view of the camera.

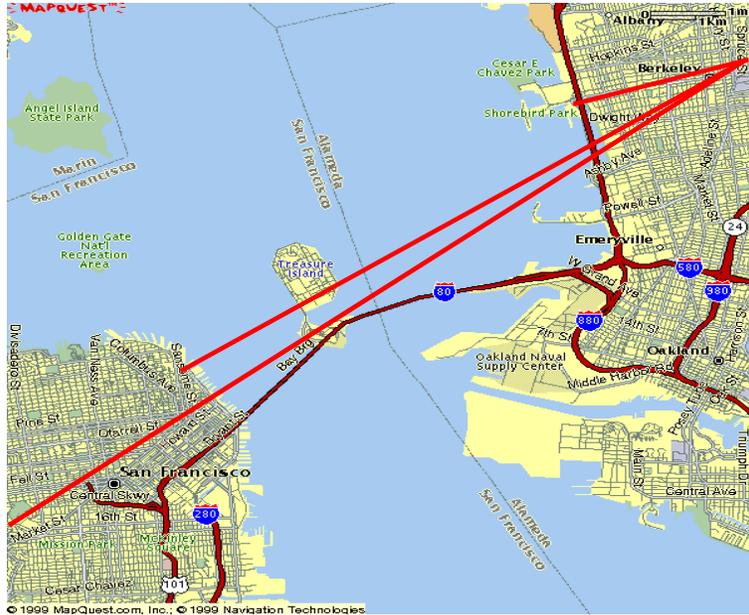


Figure 8: This is a map of 1999 communications links demonstrated using low-power laser transmitters. To date, we have successfully established links between Cory hall and (A) the Berkeley Marina (5.2km), (B) Coit Tower (15.3km), and (C) Twin Peaks (21.4km) using a 3mW optical power laser pointer with a 1mradx2mrad divergence. This results in a spot size of roughly 15 m x 30 m and an illumination of around 5 microW/m² from Coit Tower.

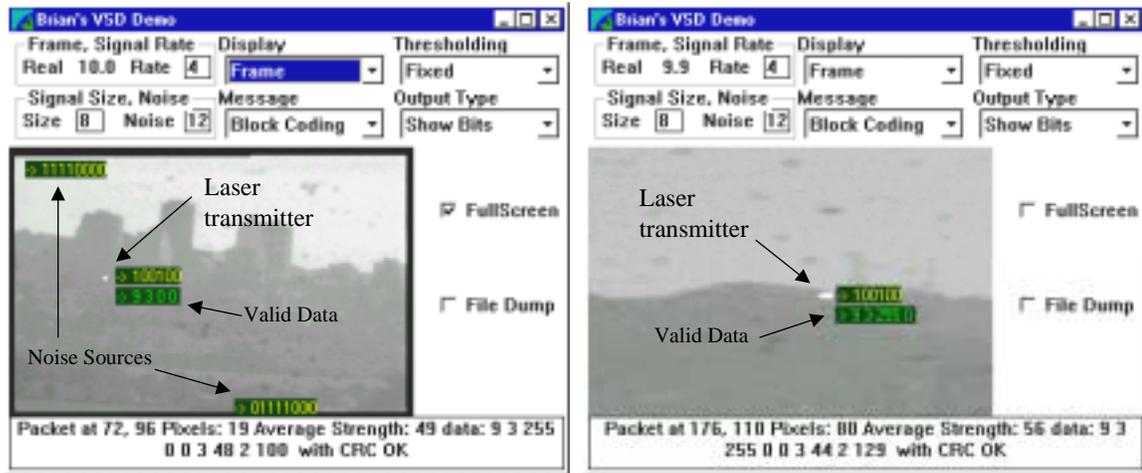


Figure 9: VSD screen captures of incoming data from Coit Tower (left) and Twin Peaks (right) experiments. Noise sources (left) that survive size and intensity filtering are rejected unless a valid Cyclic Redundancy Check (CRC) is received.

4. RESULTS AND DISCUSSION

Figure 8 shows the free-space communication links demonstrated. The setup described in section 3.1 and 3.2 was used in all of these experiments. Figure 9 shows a screen capture of the software used to decode signals during the Coit Tower → UC Berkeley (15.4km) and Twin Peaks → UC Berkeley (21.4km) experiment. MALT has been demonstrated using the shortest of these links (5.2km) only. From the strength of the received signal observed during the longest of the links, it is believed that even longer distances should be practical with the equipment described in this paper.

The communication links demonstrated so far illustrate that long-range, low-power communication using milliradian-divergence beams is possible. The communication rates are limited by the low speed of the receiver (currently 15bps with the FPGA-based VSD), and economic constraints contraindicate scaling the speed of systems such as the ones presented in this paper past a few thousand bits per second. In order to inexpensively achieve higher bandwidths, parallelizing the decoding process seems to be the best approach. In the limit of this approach, where each pixel is a fully autonomous digital receiver, speeds in the Mbps are possible⁵.

These systems demonstrate the virtues of steered narrow-beam communication. By continuing to increase levels of integration and miniaturization through the use of MEMS technology, we aim to demonstrate a complete transceiver station with a volume of just a few cubic centimeters within the next few years. We hope to advance the state of the art and to bring the benefits of steered laser communication, namely low power, long range, high speed, and intrinsic security, to applications where RF wireless just can't do the job.

REFERENCES

¹ Last, M. **An 8 mm³ digitally steered laser beam transmitter**. IEEE/LEOS International Conference on Optical MEMS, Kauai, HI, USA, 21-24 Aug. 2000. Piscataway, NJ, USA: IEEE, 2000. p.69-70.

² Macro Mote webpage – http://www.eecs.berkeley.edu/~shollar/macro_motes/macromotes.html

³ Seth's Master's Thesis – http://www.eecs.berkeley.edu/~shollar/shollar_thesis.pdf

⁴ Advanced Circuits - <http://www.advancedcircuits.com/>

⁵ BS Leibowitz, BE Boser, KSJ Pister, "CMOS 'Smart Pixel' for Free-Space Optical Communication", SPIE EI 2001 Conference 4306A, Proc. SPIE vol 4306-37.

⁶ Smoovy, APH59001 High Precision Linear Actuator. <http://www.smoovy.com/products/APH59001.htm>