
Wireless Communications for Smart Dust

V.S. Hsu, J.M. Kahn, and K.S.J. Pister

Department of Electrical Engineering and Computer Sciences
514 Cory Hall
University of California
Berkeley, CA 94720

Phone: (510) 643-8848

Facsimile: (510) 642-2739

Internet: {jmk, vhsu, pister}@eecs.berkeley.edu

January 30, 1998

Abstract:

The goal of this research is to design a wireless communication system for sending and receiving data from distributed sensor networks, or Smart Dust systems. The devices for communication on the dust motes are subject to size and power constraints, and are required to communicate over several hundred meters at a minimum data rate of 1 kbps. Optical links with active and passive transmitters and RF links are considered. A detailed analysis of the link from a passive transmitter, or corner-cube retroreflector (CCR), on the dust mote to an imaging system, either CMOS or CCD, shows the feasibility of this link. The technique used for modeling and analyzing the performance of imperfect CCRs is discussed, and some results from this analysis are presented. The design of the imaging system for high data rate reception and optimized performance is investigated. Factors governing the choice of an appropriate communication protocol are discussed.

CONTENTS

- I. Introduction 1**
- II. System Design Options 1**
 - A. Uplink using Corner-Cube Retroreflectors and Imaging Receiver 3
- III. Link Components 8**
 - A. Modeling and Optimization of Corner-Cube Retroreflectors 8
 - B. Design of Imaging Receiver 10
 - C. Eye-Safe Interrogation Laser for CCR-Based Uplink 12
- IV. Communication Protocols 13**
- V. Conclusions 14**
- VI. References. 15**
- VII. Tables and Figures 16**

I. Introduction

Smart Dust are millimeter scale sensing and communication platforms. Distributed sensor network systems can consist of hundreds to thousands of dust motes and one or more interrogating transceivers. Each dust mote consists of a power supply, a sensor or sensors, analog and digital circuitry, and a system for receiving and transmitting data. Depending on the power source, solar cells, thick film batteries, or commercially available batteries, the dust mote can vary in size from 1 mm^3 to as large as a sugar cube.

There are both military and commercial applications for the dust motes. The military could use dust motes containing acoustic, vibration, and magnetic field sensors distributed across many square miles of territory to monitor the passage of vehicles. The sensors could be delivered to the area by unmanned air vehicles (UAV), artillery, or distributed like seeds from moving vehicles. They could be interrogated by manned air vehicles (MAV) or soldiers with modified binoculars. In the future, chemical and biological sensors could be incorporated into the dust motes to detect the use of chemical or biological agents in combat.

Both the military and industry could use dust motes to monitor the performance of critical parts of aircraft, vehicles, and manufacturing equipment. This could dramatically reduce the cost of maintenance. A specific application is monitoring high cycle fatigue (HCF) of rotating compression blades. Currently, engine parts need to be inspected regularly for fatigue. The ability to attach sensors would reduce the amount of maintenance required. A major limitation in using sensors to monitor HCF has been the difficulty in wiring sensors to rotating devices. There have also been problems with vibration and acceleration levels that prevent the use of conventional hardware and connectors. Dust motes would provide a solution to these problems.

Clearly, wireless communication is essential for each of these applications. The design of the communication system is subject to several constraints due to the size and low power of the dust motes. Several possible systems will be introduced including optical links with active and passive transmitters and RF links. Based on the requirements for implementing these systems, the optical link with a passive transmitter on the dust mote is selected for further performance analysis and design.

II. System Design Options

A viable wireless communication architecture for smart dust must satisfy a number of requirements. It must support half- or full-duplex, bidirectional communication between a central transceiver and up to 1000 dust motes. The downlink (central transceiver to dust motes) must broadcast to all of the dust motes at a bit rate of several kbps. The uplink (dust motes to central transceiver) must permit each of 1000 dust motes to convey about 1 kb of data within 1 s, an aggregate throughput of 1 Mbps. Options for uplink multiplexing include

time-, frequency-, code- and space-division multiplexing, as described below. The central transceiver must be able to resolve the position of each dust mote with an angular resolution of the order of 1/100 of the field of view. The link should operate over a range of at least several hundred meters. The dust mote transceiver must occupy a volume of the order of 1 mm^3 , and consume an average power not exceeding $1 \text{ }\mu\text{W}$. If possible, the uplink and downlink should afford a low probability of interception. It should be noted that we do not envision performing direct communication between dust motes.

We will perform a detailed comparison of several alternative system designs based on radio frequency (RF) and free-space optical transmission. This study will reveal the best downlink and uplink designs, and permit us to optimize the design of the relevant components. Optimum downlink and uplink designs are likely to be rather different. Indeed, optimized designs of both should exploit the system's asymmetry, i.e., the larger power consumption, complexity, size and cost permissible in the central transceiver unit, as compared to a dust mote.

RF transmission can be employed for the uplink and/or the downlink. In principle, uplink multiplexing can be achieved using time-, frequency- or code-division multiplexing (TDMA, FDMA or CDMA) [2], but each technique may be problematic in this application. TDMA requires each dust mote to transmit at an excessively high bit rate, as high as the aggregate uplink capacity in the absence of other multiplexing techniques. Moreover, TDMA requires each dust mote to coordinate its transmission with all the other dust motes. FDMA requires accurate control of the dust-mote oscillator frequency, while CDMA requires high-speed digital circuitry to operate for a relatively extended time interval, potentially consuming excessive power. In order to avoid coordination between dust motes, both FDMA and CDMA require individual dust motes to be preprogrammed with unique frequencies or codes. If RF transmission is to be used, it is probably necessary to employ an judicious combination of two or more of these multiplexing techniques. In principle, space-division multiplexing (SDMA) represents another means to multiplex uplink RF transmissions. In SDMA, the central transceiver employs an antenna array to separate transmissions from different dust motes. Given the limited size of the central transceiver, however, it would be difficult for SDMA to achieve the required spatial resolution. For similar reasons, it would be difficult for a RF uplink to provide information on the location of the dust motes. The resolution of RF-based location determination can also be degraded by multipath propagation.

Free-space optical transmission at visible or near-infrared wavelengths (400-1600 nm) represents an attractive alternative for the downlink and uplink. On the downlink, a single laser transmitter can broadcast an on-off-keyed signal to the collection of dust motes. Each dust mote would be equipped with a very simple receiver consisting of a bandpass optical filter, a photodiode, a preamplifier and a slicer. This receiver would involve only low-speed baseband electronics, making it far simpler than its RF counterpart.

On the uplink, optics offers two alternatives for transmission. Each dust mote can be equipped with an active laser-diode-based transmitter. This has the disadvantage of requiring excessive power. To minimize the power consumption, the dust mote should employ a directional beam and an active beam-steering mechanism, but this complicates the dust-mote design, and adds the need for beam-tracking protocols. Alternatively, each dust mote can employ an optically passive transmitter consisting of a corner-cube retroreflector (CCR). By displacing one of the mirrors, the CCR reflectivity can be modulated at bit rates up to 10 kbps. The collection of dust motes can be illuminated by a single interrogating laser. Because of its small power consumption (of the order of $1 \mu\text{W}$) and because it avoids the need for active beam-steering, the CCR-based transmitter seems much more promising than the laser-diode-based design. Because they involve only low-speed baseband digital electronics, both transmitter designs are much simpler than their RF-based counterparts.

With either uplink optical transmitter design, the uplink receiver should employ an imaging receiver, consisting of a lens and a CCD or CMOS image sensor array. With either transmitter design, periodic pulsation of the downlink or interrogating laser beam can be used to synchronize the transmissions from all dust motes to the frame clock of the imaging array. Such an uplink architecture offers several advantages over a RF solution. The short optical wavelength enables a reasonably sized camera to achieve high spatial resolution. Not only does this easily yield information on the dust-mote locations, but it enables the uplink to rely solely upon SDMA for multiplexing [4]. The 1000 dust motes can transmit simultaneously and without any mutual coordination, greatly simplifying their design and the system protocols.

Based on the above considerations, we currently favor free-space optical transmission for both downlink and uplink, with the uplink relying on a CCR-based transmitter. This system design is shown in Fig. 1.

A. Uplink using Corner-Cube Retroreflectors and Imaging Receiver

In this section, we analyze the signal-to-noise ratio of a free-space optical link using a corner-cube transmitter and a CMOS image sensor receiver, as illustrated in Fig. 1. This analysis will prove quantitatively the feasibility of this uplink design and help identify critical design parameters. A similar analysis has also been carried out for the downlink shown in Fig. 1, and shows that the downlink can rather easily achieve the required signal-to-noise ratio.

The symbols used in our analysis are summarized in Table 1. The interrogating laser emits a CW beam of power P_t , which uniformly illuminates a field of semiangle θ_f . At the CCR range r , the illuminated field has a diameter $d_f = 2r \tan \theta_f$. The irradiance incident upon the CCR is thus:

$$I_c = \frac{P_t}{\pi r^2 \tan^2 \theta_f} . \quad (1)$$

The CCR has an effective diameter d_c when not tilted with respect to the link axis. When the CCR is tilted, its effective capture area is reduced by a factor f_{cap} . The CCR captures a power $I_c \pi d_c^2 \cos \theta_c / 4 = P_t d_c^2 \cos \theta_c / 4 r^2 \tan^2 \theta_f$. Including the effect of multiple internal bounces, the CCR has effective reflectivity R_c . One CCR mirror is modulated so that the device emits an on-off-keyed signal with non-return-to-zero pulses. Assuming that zeros and ones are equiprobable, the average power reflected by the CCR is:

$$P_c = \frac{P_t R_c d_c^2 \cos \theta_c}{8 r^2 \tan^2 \theta_f} . \quad (2)$$

Because of diffraction, the signal emitted by the CCR spreads as it propagates back toward the receiving camera. When the CCR is not tilted with respect to the link axis, we model it as a circular aperture having diameter d_c , and compute the diffracted irradiance at the lens using Fraunhofer diffraction theory [1]. When the CCR is tilted, the diffracted irradiance at the lens is assumed to be reduced by a factor f_{dif} . Thus, the average irradiance at the lens is given by:

$$I_l = \frac{P_c \pi d_c^2 f_{dif}}{4 \lambda^2 r^2} = \frac{\pi P_t R_c d_c^4 f_{cap} f_{dif}}{32 \lambda^2 r^4 \tan^2 \theta_f} , \quad (3)$$

where λ is the interrogation wavelength. In this analysis, we have neglected any imperfections in the CCR and any atmospheric attenuation (e.g., due to rain). We assume that to reject ambient light, the camera employs an optical bandpass filter having a bandwidth $\Delta \lambda$ and a transmission T_f . The lens has an effective entrance diameter d_l and a transmission T_l . A fraction f_{act} of each camera pixel is sensitive to light, converting it to current with a responsivity R . Thus, the average received photocurrent is given by:

$$i_{sig} = \frac{I_l \pi d_l^2 T_l T_f f_{act} R}{4} = \frac{\pi^2 P_t d_c^4 d_l^2 \cos^2 \theta_c R_c T_l T_f f_{act} R}{128 \lambda^2 r^4 \tan^2 \theta_f} . \quad (4)$$

We note that the received signal may be divided among several pixels, and (4) represents the total over all of these pixels.

We now calculate the ambient light noise received by the camera. Assume that the region surrounding the CCR is illuminated by ambient light (e.g., sunlight) having a spectral irradiance p_{bg} and reflects this ambient light with reflectivity R_{bg} . The camera employs a total of N pixels, and its field-of-view lies within a semi-angle θ_f , which is matched to the interrogating beam. In the plane of the CCR, each camera pixel views an area $\pi r^2 \tan^2 \theta_f / 4N$. Within the bandwidth of the optical bandpass filter, the total power reflected from that area is:

$$P_{bg} = \frac{\pi p_{bg} R_{bg} \Delta \lambda r^2 \tan^2 \theta_f}{N}. \quad (5)$$

Assuming that the region surrounding the CCR is a Lambertian reflector [1], the reflected irradiance at the lens (within the filter bandwidth) is:

$$I_{bg} = \frac{P_{bg}}{\pi r^2} = \frac{p_{bg} R_{bg} \Delta \lambda \tan^2 \theta_f}{N}. \quad (6)$$

The d.c. photocurrent per pixel due to ambient light is:

$$i_{bg} = \frac{I_{bg} \pi d_l^2 T_f T_{f_{act}} R}{4} = \frac{\pi p_{bg} R_{bg} \Delta \lambda \tan^2 \theta_f d_l^2 T_f T_{f_{act}} R}{4N}. \quad (7)$$

This d.c. photocurrent induces a white shot noise having (one-sided) power spectral density per pixel:

$$S_{bg} = 2q i_{bg}, \quad (8)$$

where q is the electronic charge.

Each pixel is assumed to have a d.c. leakage current i_{leak} , which induces a white shot noise having power spectral density per pixel:

$$S_{leak} = 2q i_{leak}. \quad (9)$$

In calculating the receiver SNR, we first assume that the signal is incident upon only one pixel, which is the best case. The bit rate is B . The receiver is synchronized to the received signal, and integrates the received photocurrent for a fraction g_{int} of each bit interval, i.e., for an interval g_{int}/B per bit. For analytical convenience, the integrator is assumed to have a gain of B . At the end of this interval, the integrator output is sampled, and this sample is compared to a threshold set midway between the “zero” and “one” levels. The average signal component of this sample is $i_{sig} g_{int}$. The noise component of this sample has variance:

$$\sigma_{tot}^2 = \frac{(S_{bg} + S_{leak})Bg_{int}}{2}. \quad (10)$$

The average SNR is:

$$SNR_{best} = \frac{(i_{sig}g_{int})^2}{\sigma_{tot}^2} = \frac{2i_{sig}^2g_{int}}{(S_{bg} + S_{leak})B}. \quad (11)$$

Depending on the size of the CCR image spot, the pixel size, and the position of the corner-cube, the received signal may overlap several pixels. The number of pixels among which the signal is divided is denoted here by N_o . When $N_o > 1$, the optimum way to process these signals is to perform maximal-ratio combining (MRC), i.e., to form a linear combination of the N_o photocurrents, weighting each in proportional to its signal-amplitude-to-noise-variance ratio [2][3]. A common suboptimal technique involves selecting the best of the N_o pixels (SB). For either technique, for a given N_o , the SNR is most degraded when the signal is equally split among all N_o pixels. In this case, MRC and SB achieve SNRs given by:

$$SNR_{worst}^{MRC} = \frac{SNR_{best}}{N_o}, \quad (12)$$

and

$$SNR_{worst}^{SB} = \frac{SNR_{best}}{N_o^2}, \quad (13)$$

respectively. Assuming that all noise sources are Gaussian-distributed, the bit-error probability of the receiver is given by:

$$P_b = Q(\sqrt{SNR}), \quad (14)$$

where $Q(x)$ is the Gaussian Q function [2]. For example, to achieve $P_b = 10^{-6}$ requires a SNR of about 13.5 dB.

In the image sensor, there are gaps between pixels, and each pixel has regions that are not light-sensitive. This leads to two alternative design strategies. In the first strategy, we hope to successfully receive the CCR signal whenever the CCR is pointing toward the receiver, regardless of where the image spot lies on the pixel array. This implies that we should make the pixel size comparable to the spot size, so that a large fraction of the image spot always falls upon light-sensitive regions of the pixel array. In a typical implementation, the image spot and pixels might each be about 10 μm across. In order to minimize the maximum value of N_o , one should

employ hexagonal pixels, so that N_o ranges from one to three. If the receiver employs MRC, when SNR is expressed in decibels, $SNR_{worst}^{MRC} = SNR_{best} - 4.8$ dB. If SB is employed, $SNR_{worst}^{SB} = SNR_{best} - 9.5$ dB. Thus, it is highly preferable to use MRC instead of SB. Even with MRC, however, there is a large penalty incurred when the signal spot overlaps three pixels. Perhaps an even bigger disadvantage of this approach is the need to combine three analog signals with variable weights prior to the slicer, and to perform SNR estimation to determine those weights.

The second design strategy is to make the pixel size much larger than the spot size. For example, the spot size might be less than 10 μm across, while the pixels would be 30 μm squares. In this case, it is very likely that the spot will fully overlap only one pixel, in which case $N_o = 1$. Thus, the receiver need not employ MRC or SB, and there is no corresponding SNR penalty. This greatly simplifies the receiver design, since there is no need for SNR estimation in order to perform MRC or SB. It should be noted that if the image spot falls on the gap between two pixels or on the non-light-sensitive region of a pixel, no signal will be received from that CCR. In most applications, this is acceptable, because the large number of dust motes within the receiver field of view makes it unnecessary to receive the signal from any one particular dust mote. Thus, the second design strategy is preferred over the first in most applications.

To determine the effect of various design parameters on the link performance, we assume that ambient-induced shot noise is dominant, and use (11), (4) and (7). We find that to achieve a given target SNR, the required interrogation power depends strongly on several parameters. The required value of P_t is proportional to r^4 , $\tan^3 \theta_f$ and d_c^{-4} . In order to determine how the required interrogation power depends on wavelength, we must take account of the wavelength dependence of the responsivity R and of the ambient light spectral irradiance p_{bg} . The wavelength dependence of the responsivity is described by $R(\lambda) = e\lambda\eta(\lambda)/hc$, where $\eta(\lambda)$, the detector quantum efficiency, depends on the wavelength and on the depth of the detector active region. When the ambient light is dominated by sunlight, the wavelength dependence of $p_{bg}(\lambda)$ is well-approximated by a blackbody spectrum, i.e., $p_{bg}(\lambda) \propto \lambda^{-5} (e^{hc/\lambda kT} - 1)^{-1}$, where $T \cong 5000$ K. For present purposes, it is sufficient to note that over the wavelength range of interest (400-900 nm), $p_{bg}(\lambda)$ does not vary by more than about a factor of two. Taking account of these wavelength dependences, the required value of P_t is proportional to $\lambda^{3/2} [p_{bg}(\lambda)]^{1/2} [\eta(\lambda)]^{-1/2}$. The required interrogation power P_t depends weakly on some other parameters, being proportional to $B^{1/2}$, $\Delta\lambda^{1/2}$, $f_{act}^{-1/2}$, and $g_{int}^{-1/2}$.

It is important to choose N , the number of pixels, carefully. We restrict our attention to the second design strategy described above, in which $N_o = 1$, and the SNR is always given by SNR_{best} . Examining (11) and (7), we see that $SNR_{best} \propto N$, because i_{sig} is independent of N , $i_{bg} \propto 1/N$, and for a fixed overall array area,

typically $i_{leak} \propto 1/N$. We should choose N to be as large as possible such that the pixel size is sufficiently larger than the image spot size. For example, if the interrogation wavelength is $\lambda = 830$ nm and the lens diameter is $d_l = 0.1$ m, the minimum spot diameter lies between 5 and 10 μm (as limited by diffraction). If the image sensor die size cannot exceed 1 cm square, then it is reasonable to employ 10^5 pixels, each 30 μm square.

In Table 1, we list a set of typical design parameters. These show that it is feasible to achieve a reliable 10-kbps link over a 500-m range using a CCR diameter of 0.5 mm and an interrogating power of 10 W, and an array of 10^5 pixels. In this example, the dominant noise arises from ambient-induced shot noise.

III. Link Components

Based on detailed system-level analysis, we will design critical components to optimize the link performance. In this section, we discuss some of the factors governing the optimized design of uplink transmitter and receiver components.

A. Modeling and Optimization of Corner-Cube Retroreflectors

In order to optimize the CCR optical performance, we will perform detailed modeling that takes account of several effects. The mirror faces can buckle slightly as the device cools down from its fabrication temperature, degrading its reflection efficiency. Similarly, non-orthogonal alignment of the faces can degrade reflection efficiency. Diffractive spreading of the reflected beam puts a fundamental limit on the performance of a long-range optical link using this device. In order to model nonflatness, misalignment and diffraction on an equal footing, we will perform a finite-element analysis of the device.

Each face of the CCR may be subject to some curvature, spherical or otherwise, which is represented by an equation describing the surface. The CCR faces are bounded by the planes $x = 0$, $x = l$, $y = 0$, $y = l$, $z = 0$ and $z = l$, where l is the length that the sides would have if they were perfectly flat and mutually perpendicular. The surfaces are expected to be close to flat and orthogonal, so the error introduced by not adjusting the boundary conditions is negligible. Each surface is divided into a specified number of discrete elements that are bounded by equally spaced planes parallel to the x - y , y - z , and x - z planes. For perfectly flat surfaces, these discrete elements would be squares. For non-flat surfaces, the discrete elements have non-uniform surface area and shape. The number of discrete elements should be large enough such that each element can be assumed to be essentially flat.

Fig. 2 depicts the geometry used to model the CCR. The CCR is illuminated along the direction $-\hat{n}_i$ by a uniform plane wave having an irradiance I_i . For each discrete element m , a ray trace is performed to determine the direction of the ray leaving the CCR. The ray starts from a reference plane where all of the rays that strike

the CCR are in phase. This reference plane is required to be normal to the incident light and to pass through an arbitrary point near the CCR. The axes of the plane are chosen to be an arbitrary pair of unit vectors normal to the incident vector and each other. It is important to specify the axes for future calculations requiring change of basis. The ray first strikes the specified discrete element. The normal vector of the discrete element is determined by computing the Jacobian of the equation describing the surface, evaluated at the center of the element. The direction of the reflected ray can be determined from the incident vector and the normal vector. The next surface that the ray strikes, if any, is then determined, and the ray trace continues. Finally, the ray leaves the CCR and is terminated on a second reference plane. This plane is required to be normal to the ray leaving the CCR and to pass through an arbitrary point near the CCR. Unlike the first reference plane, this second reference plane is not common to each discrete element unless the surfaces are perfectly flat. The axes of this reference plane can be arbitrarily chosen to be any two unit vectors, u and v , normal to the ray from the CCR and each other. It is necessary to specify these axes to perform the required change of basis to calculate the Fraunhofer diffraction integral. It is also important to keep track of the distance each ray propagates through the ray trace d and the number of reflections that occur N_{ref} . These values will be needed to specify the phase change and the transverse extent of the wave reflected from the discrete element.

The wave propagating along the path of the ray trace is a plane wave, and is nonzero only within a region given by the projection of the discrete element onto the first reference plane. When the CCR is broken into a sufficiently large number of discrete elements, this shape can be approximated as a parallelogram. The boundaries of the parallelogram are determined by projecting two adjacent sides of the discrete element onto the first reference plane. Once the boundaries of the parallelogram are defined, the image can be projected to the second reference plane either inverted, for an even number of reflections, or not inverted, for an odd number of reflections. Fig. 2 shows the ray trace for a particular discrete element and its corresponding parallelogram for a CCR with four discrete elements on each surface.

It is important to note that the entire parallelogram propagates to the location of the termination on the second reference plane. The number of discrete elements must be large enough such that this method is valid. For example, if a surface has only one discrete element, the entire surface is expected to propagate to the location on a single ray trace. For certain incidence vectors, this may mean the entire surface appears to be reflected back to the source when, in actuality, only a fraction of the surface is part of the effective area.

In order to treat diffraction effects, we employ Fraunhofer diffraction theory [1], which is valid when the receiver lies in the far-field. More precisely, this requires that $R > a^2/\lambda$, where R is the distance from the aper-

ture to the receiver, a is the largest dimension of the aperture, and λ is the wavelength of the light. The complex amplitude of the electric field at the observation point contributed by discrete element m is given by:

$$E_m(\hat{\mathbf{n}}_i, \hat{\mathbf{n}}_o, R) = \frac{\sqrt{2I_i} e^{-i[k(R+d) + \pi N_{ref}]}}{\lambda R} \iint_{\text{parallelogram}} e^{ik(Uu + Vv)/R} dS, \quad (15)$$

where $k = 2\pi/\lambda$. The sum over all discrete elements gives the complex amplitude of the total electric field at the receiver:¹

$$E_o(\hat{\mathbf{n}}_i, \hat{\mathbf{n}}_o, R) = \sum_m E_m(\hat{\mathbf{n}}_i, \hat{\mathbf{n}}_o, R). \quad (16)$$

The irradiance at the receiver can be calculated from this total electric field using:

$$I_o(\hat{\mathbf{n}}_i, \hat{\mathbf{n}}_o, R) = \frac{1}{2} |E_o(\hat{\mathbf{n}}_i, \hat{\mathbf{n}}_o, R)|^2. \quad (17)$$

The differential scattering cross section $d\sigma(\hat{\mathbf{n}}_i, \hat{\mathbf{n}}_o)/d\Omega_o$ (m^2/sr) is the scattered power per unit solid angle of the observation per unit illumination irradiance. This can be calculated from the irradiance using:

$$\frac{d\sigma(\hat{\mathbf{n}}_i, \hat{\mathbf{n}}_o)}{d\Omega_o} = \frac{I_o R^2}{I_i}. \quad (18)$$

The integral of the differential scattering cross section over observation angle equals the effective area of the CCR, or total scattering cross section, which has units of m^2 :

$$\int_{4\pi} \frac{d\sigma(\hat{\mathbf{n}}_i, \hat{\mathbf{n}}_o)}{d\Omega_o} d\Omega_o = \sigma(\hat{\mathbf{n}}_i) \quad (19)$$

Some of our preliminary results are shown in Fig. 3. These results show how the CCR performance is impacted by non-flat mirrors. Light is incident along $-\hat{\mathbf{n}}_i$ and is observed along $\hat{\mathbf{n}}_o$. We note that in the preliminary link performance analysis given above, the total cross section was expressed as $\sigma(\hat{\mathbf{n}}_i) = \pi d_c^2 f_{cap} f_{dif}/4$, where the factors f_{cap} and f_{dif} depend upon $\hat{\mathbf{n}}_i$. Examination of Fig. 3 shows that when the mirrors have radii of curvature of 0.3 m, the on-axis differential scattering cross section is reduced to 55% of the value achieved with flat mirrors. This modeling will allow us to establish fabrication tolerances for the devices and identify design changes that will improve performance. Accurate modeling is also required for accurate link performance analysis. Furthermore, modeling will allow us to compute the proba-

¹ The real, time-dependent electric field is related to the complex amplitude via $E_o(\hat{\mathbf{n}}_i, \hat{\mathbf{n}}_o, R; t) = \text{Re}[E_o(\hat{\mathbf{n}}_i, \hat{\mathbf{n}}_o, R) e^{i\omega t}]$, where $\omega = 2\pi c/\lambda$.

bility distribution of the scattering cross section for a randomly oriented CCR, which is required to predict the reliability of a system employing a large collection of such devices.

B. Design of Imaging Receiver

The uplink receiver lens should be chosen to have a large entrance aperture. This will increase the power captured from each CCR, improving the link SNR. It will also reduce the diffraction-limited image spot size, permitting the image sensor to employ smaller pixels, and ultimately reducing the image sensor die size. The uplink receiver should probably employ active optical image stabilization, in order to minimize the requirement for the imaging receiver to electronically track jittering dust-mote images.

The most critical component of the uplink receiver is the image sensor array. Because the design leading to optimized optical link performance differs from the conventional design optimized for imaging performance, we will design and fabricate a custom solution.

There are two main technologies available for the imaging system, CCD and CMOS. Recently, CMOS imagers with active pixel sensors (APS) have been developed that rival the performance of commercial CCD imagers. An active pixel sensor is a sensor with an active amplifier on each pixel. The advantages of using CMOS imaging systems include lower cost, increased on-chip functionality, lower power requirements, and miniaturization. The inherent ability for CMOS imagers to have on-chip/on-pixel processing makes high-data-rate reception possible. It would be nearly impossible to read out an entire CCD array and process the signals at the required rate.

As described above, if the interrogating laser is pulsed periodically in synchronization with the image sensor frame clock, all uplink transmissions can be easily synchronized to that clock. This obviates the need for the receiver to perform individual timing recovery on each signal, and permits the frame rate to be chosen no higher than the uplink bit rate, e.g., 10^4 bps. While this frame rate is higher than that of most image sensor arrays, as described below, the aggregate off-chip I/O requirement can be reduced to several tens of Mbps, making it easy to achieve the required frame rate.

As described above, because the image sensor has non-light-sensitive regions within each pixel and between pixels, there are two alternate design strategies. The first strategy seeks to successfully receive the CCR signal whenever the CCR is pointing toward the receiver, regardless of where the image spot lies on the pixel array. Thus, the pixel size is chosen to be comparable to the spot size (e.g., each would be about $10\ \mu\text{m}$ across). If possible, the pixels should be made hexagonal, to limit the maximum value of N_o to three. When the image spot overlaps more than one pixel, if possible, the receiver employs MRC, i.e., it forms a weighted sum of the analog signals from up to three pixels and makes decisions on this sum. Not only is this approach

subject to a significant SNR penalty, but it is complicated to implement. To begin with, the receiver must be capable of identifying which groups of pixels are associated with a signal spot. Then the receiver must estimate the signal-amplitude-to-noise-variance ratio of the various pixels associated with each spot. Finally, it must form a weighted sum of the corresponding analog signals, and pass this sum to a decision circuit. It would be possible to perform these functions off the image sensor chip, but this would be extremely I/O-intensive, requiring an analog signal from each pixel in the array to be brought off the chip (alternatively, a four-to-eight-bit digital representation could be employed).

In the second strategy, the pixel size is chosen to be much larger than the image spot size, and the light-sensitive fraction of each pixel, f_{act} , is chosen close to unity. As a result, an image spot will usually fall upon the light-sensitive region of one and only one pixel. The receiver identifies which pixels receive a strong signal, and simply passes the output of each such pixel to a slicer. This strategy offers two key advantages. Firstly, because it is not necessary to detect a signal spot when it is divided among several pixels, it avoids the associated SNR penalty. Secondly, and perhaps more importantly, it avoids the complications of image segmentation, SNR estimation and variable-gain analog signal combining. Of course, the signal may be lost when an image spot falls upon a non-light-sensitive region or is divided among several pixels. This is acceptable in most applications, because the large number of dust motes makes it unnecessary to receive the signal from any one particular dust mote.

We plan to adopt the second design strategy. Fig. 4 shows the image sensor architecture we propose to employ. The array has 10^5 pixels, so that if the dust motes transmit at a bit rate of 10 kbps, the potential off-chip data transfer rate requirement is 1 Gbps. The receiver is only required to receive up to 1000 signals simultaneously, so that the only information that really needs to be conveyed off-chip are the 1000 bit streams and the locations of the corresponding active pixels. This corresponds to an aggregate bit rate of several tens of Mbps, which is quite reasonable. There are several possible approaches to consolidating the data for off-chip communication. We propose to adopt the hierarchical approach shown in Fig. 4. The pixel array is subdivided into 1000 fixed clusters, each containing 100 pixels. The density of dust motes is low enough that there will very seldom be more than one or two active pixels within each cluster. Each cluster is equipped with some simple local circuitry that can identify the few pixels having the strongest receptions, decode their data, and convey this data and the locations of the corresponding pixels onto a synchronous data bus. The aggregated data and active-pixel locations are sent off-chip for further processing. For example, the off-chip processor can improve the link reliability by tracking the movement of possible received bit streams through the pixel array over time, accepting only bit streams that move along a continuous trajectory.

C. Eye-Safe Interrogation Laser for CCR-Based Uplink

The high-power interrogation laser for the CCR-based uplink must be designed carefully in order to optimize performance and achieve eye safety. In the design example given above, for concreteness, we have assumed a laser emitting 10 W at a wavelength of 830 nm, but these parameters are likely to change when the design is further optimized. It is well-known that since the human cornea is opaque at wavelengths longer than about 1.4 μm , lasers at longer wavelengths are more easily made eye-safe than those at shorter wavelengths [5]. Unfortunately, silicon detectors are not sensitive at wavelengths beyond about 1.1 μm , precluding this approach to achieving eye safety. Instead, we will achieve eye safety by expanding the beam diameter until the emitted irradiance is eye safe. Since the interrogating camera must employ a relatively wide-aperture optical system, it would be desirable to also employ these optics for beam expansion. Duplexing of the transmitted and received signals would be achieved using a polarizing beamsplitter and a quarter-wave plate, as is done in readout systems for optical storage disks. The success of this approach depends upon achieving low back-reflection from the polarizing beamsplitter.

In our SNR analysis above, we found that the interrogation power required to achieve a desired SNR depends on the wavelength as $\lambda^{3/2} [p_{bg}(\lambda)]^{1/2} [\eta(\lambda)]^{-1/2}$. Our design example considered a near-infrared source at 830 nm because, at this time, it is easier to achieve a high-power, high-efficiency, compact source in this spectral region than in the visible range. In fact, it may be better to employ visible wavelengths in the 400-650 nm range. At these shorter wavelengths, the factor $\lambda^{3/2}$ is minimized and, in a silicon photodetector whose thickness is limited to only a few μm , the factor $[\eta(\lambda)]^{-1/2}$ is also minimized. In addition to minimizing the interrogation power requirement, visible radiation is easier to render eye-safe than infrared radiation [5], because the human eye responds to high power levels by blinking or averting its gaze. The highest-power commercial visible laser diodes we are aware of are 1200 mW devices emitting near 665 nm. It would probably be necessary to employ several such devices in the interrogation system. Visible laser-diode technology is progressing rapidly, and we expect single devices of the required power level to be available within a couple years.

IV. Communication Protocols

Because of the unique characteristics of the transmitting and receiving components, we will design simple communication protocols for wireless communication with smart dust. Likewise, optimization of the link components requires some consideration of the protocols that will be employed. The protocols will be designed to conserve power in the dust motes and avoid the need for any coordination among dust motes, by relying upon the much larger power and complexity available in the central transceiver unit. It is convenient to describe

these protocols in terms of a layered model [6], though it should be emphasized that our system is far simpler than most communication networks.

Because of the wide variations in wireless link quality and the relative unreliability of the individual dust motes, at the data-link layer, it is important to insure the integrity of received data. The limited processing power of the dust motes probably precludes the use of powerful forward error-correcting codes [2] on either the downlink or the uplink. Nonetheless, simple cyclic redundancy check (CRC) codes can be employed for error detection purposes on both downlink and uplink. Because of the large number of dust motes, when errors are detected on either the uplink or the downlink, retransmission is not possible, and the received data simply needs to be ignored.

While the proposed free-space optical communication architecture can support transmission at relatively high bit rates (e.g., 10 kbps), it is difficult to maintain a signal-to-noise ratio (SNR) sufficient to support these high bit rates under all conditions. For example, the link may not achieve the desired SNR at the full range (e.g., 500 m) during periods of rain or snow. Under these conditions, instead of incurring abrupt link failure, it is preferable to achieve graceful degradation, by lowering the bit rate until a suitably low error rate can be attained. The central transceiver unit can directly sense low-SNR conditions on the uplink, and can infer that the downlink SNR is low when a large fraction of the dust motes fail to respond. This unit can progressively lower the bit rate until reliable link operation is attained. The bit rate can be lowered either by lengthening the bit duration and increasing the receiver integration time, or by employing simple repetition coding [7].

In the proposed free-space optical system, media-access control protocols are extremely simple. The dust motes are silent until a broadcast downlink transmission instructs them to transmit. All uplink transmissions occur simultaneously via SDMA, and the imaging receiver is able to detect them all with negligible mutual interference.

A significant disadvantage of using active optical transmitters at the dust motes is the need for physical-layer beam-steering protocols. Specifically, in order to achieve sufficient link range at low transmitted power, these transmitters would almost certainly have to use directed optical beams. This would require some mechanism in the dust mote for steering the beam toward the central transceiver unit, and some protocol for aiming the beam toward that unit. A possible solution is to equip the dust mote with a directional transmitter and directional receiver that are both steered via a common mechanism; by aiming the receiver to maximize the reception, it would also be optimizing its transmission to the central transceiver.

V. Conclusions

The applications of miniature distributed sensor networks, Smart Dust, are numerous for the military as well as industry. Before these applications can be realized, a means of transmitting data to and from these sensors is required. Clearly, a wireless communication system would be advantageous. Several options for this system are being considered. Currently, optical links for both the uplink and downlink, employing a CCR for the uplink, seem favorable. Preliminary analysis of these links have demonstrated their feasibility.

The components required for these links will be investigated further. A complete analysis of the CCR will yield information on the optimal means to manufacture and assemble these devices as well as determine fabrication tolerances. Careful design of the CMOS imaging receiver will provide a means to receive data from hundreds to thousands of dust motes transmitting at nearly 10 kHz. Several design options and protocols are being considered to optimize the performance of these links.

VI. References

- [1] M. Born and E. Wolf, *Principles of Optics*, Sixth Edition, Pergamon Press, 1980.
- [2] J.G. Proakis, *Digital Communications*, Third Edition, Mc-Graw Hill, 1995.
- [3] J.B. Carruthers and J.M. Kahn, "Angle Diversity for Nondirected Wireless Infrared Communication", subm. to *IEEE Trans. on Commun.*, July 1997.
- [4] J.M. Kahn and J.R. Barry, "Wireless Infrared Communications", *Proc. of the IEEE*, pp. 265-298, February 1997 (Invited Paper).
- [5] IEC825-1 "Radiation Safety of Laser Products' Equipment Classification, Requirements, and User's Guide" (1993).
- [6] A.S. Tanenbaum, *Communication Networks*, Third Edition, Prentice-Hall, 1996.
- [7] F. Gfeller, W. Hirt, M. de Lange and B. Weiss, "Wireless Infrared Transmission: How to Reach All Office Space", *Proc. of IEEE Vehicular Technol. Conf.*, pp. 1535-1539, Atlanta, Georgia, April, 1996.

VII. Tables and Figures

Table 1: Symbols used in SNR analysis of CCR-based uplink.

Symbol	Description	Typical Value	Units
P_t	Peak power of interrogation beam	10	W
θ_f	Semiangle of illuminated field	1	degrees
r	Link range	500	m
d_c	Effective diameter of CCR when not tilted	5×10^{-4}	m
R_c	Effective reflectivity of CCR (incl. multiple bounces)	0.85	W/W
f_{cap}	Reduction of CCR capture cross section by CCR tilt	0.5	W/W
f_{dif}	Reduction of on-axis rec. irradiance due to CCR tilt	0.5	W/W
λ	Interrogation wavelength	830	nm
$\Delta\lambda$	Optical bandpass filter bandwidth	5	nm
T_f	Optical filter transmission	0.8	W/W
T_l	Effective transmission of camera lens	0.8	W/W
d_l	Effective diameter of camera lens entrance aperture	0.1	m
R	Pixel responsivity	0.5	A/W
f_{act}	Fraction of camera pixel area that is active	0.75	m ² /m ²
i_{sig}	Avg. detected signal photocurrent (sum over all pixels)	3.7×10^{-12}	A
p_{bg}	Ambient light spectral irradiance (bright sunlight)	0.8	W/m ² -nm
R_{bg}	Reflectivity of background behind CCR	0.3	W/W
N	Number of pixels in image sensor	10^5	-
i_{bg}	D.C. photocurrent per pixel due to ambient light	6.9×10^{-12}	A
i_{leak}	D.C. leakage current per pixel	10^{-12}	A
B	Bit rate of on-off-keyed link	10^4	bps
g_{int}	Fraction of bit interval during which camera integrates	0.95	s/s
N_o	Maximum number of pixels illuminated by signal spot	3 (hex. pix.) 4 (square pix.)	-
SNR_{best}	Best-case receiver electrical SNR	24.2 (hex. or square pix.)	dB
SNR_{worst}^{MRC}	Worst-case receiver electrical SNR with MRC	19.4 (hex. pix.) 18.2 (square pix.)	dB
SNR_{worst}^{SB}	Worst-case receiver electrical SNR with SB	14.7 (hex. pix.) 12.2 (square pix.)	dB

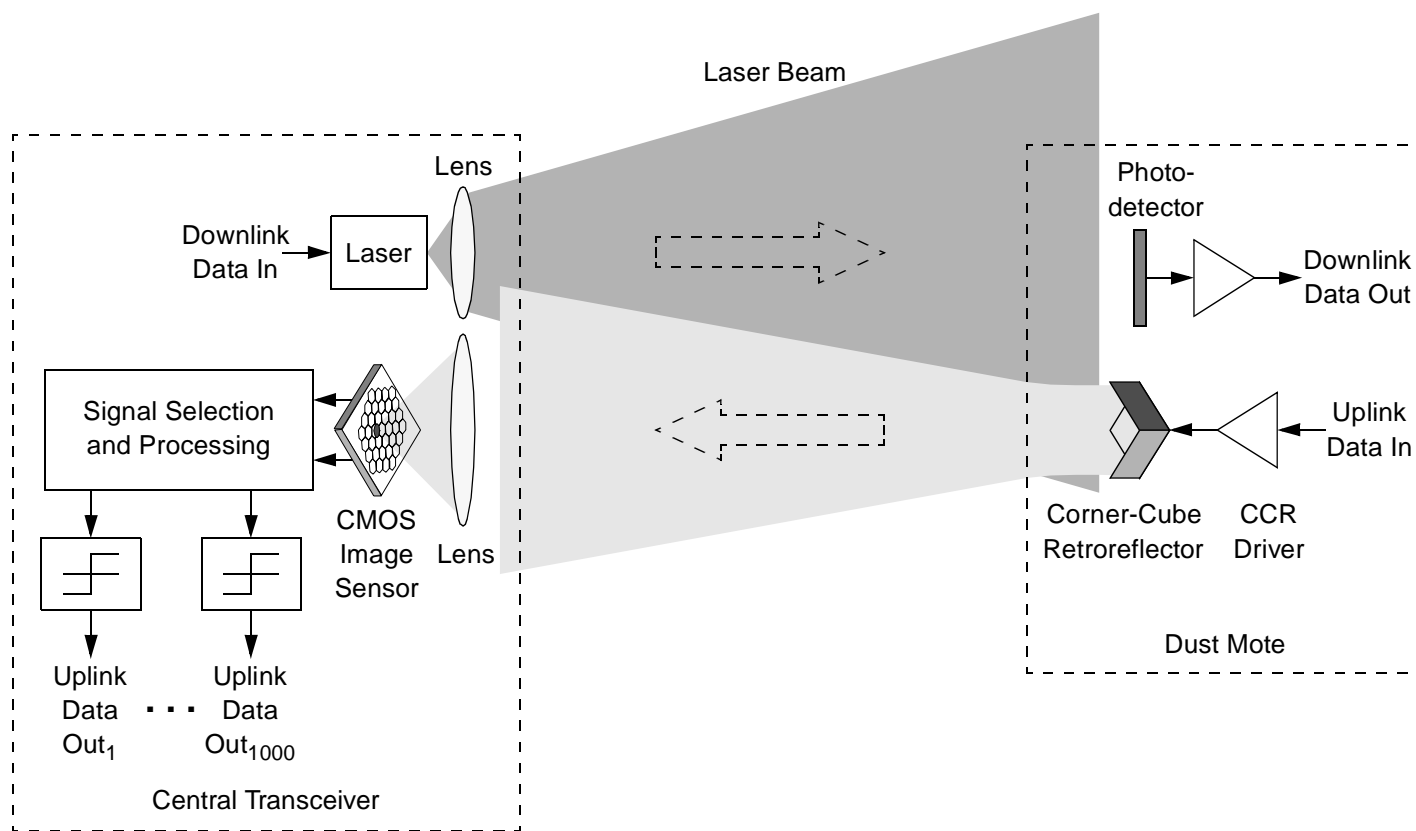


Fig. 1. Bidirectional free-space optical communication link. The laser beam is modulated with on-off-keyed data for downlink transmission, while it is CW or pulsed periodically to interrogate the CCR for uplink transmission. For clarity, only one of many dust motes is shown. Not shown are optical bandpass filters used at downlink and uplink receivers to reject ambient light noise.






Fig. 2. Ray tracing in a CCR. The incident and reflected rays for a single discrete element are indicated. The Fraunhofer diffraction integral is computed for a parallelogram-shaped aperture in the u - v plane (perpendicular to the reflected ray), yielding the electric field at a point in the U - V plane.



(a)



(b)

(c)

Fig. 3. Computed performance of non-ideal corner-cube retroreflector. (a) explains the coordinate system. Light is incident along $-\hat{n}_i$ and is observed along \hat{n}_o . (b) and (c) show the differential scattering cross section $d\sigma/d\Omega_o$ when the mirrors have radii of curvature of 1000 m (nearly flat) and 0.3 m, respectively.

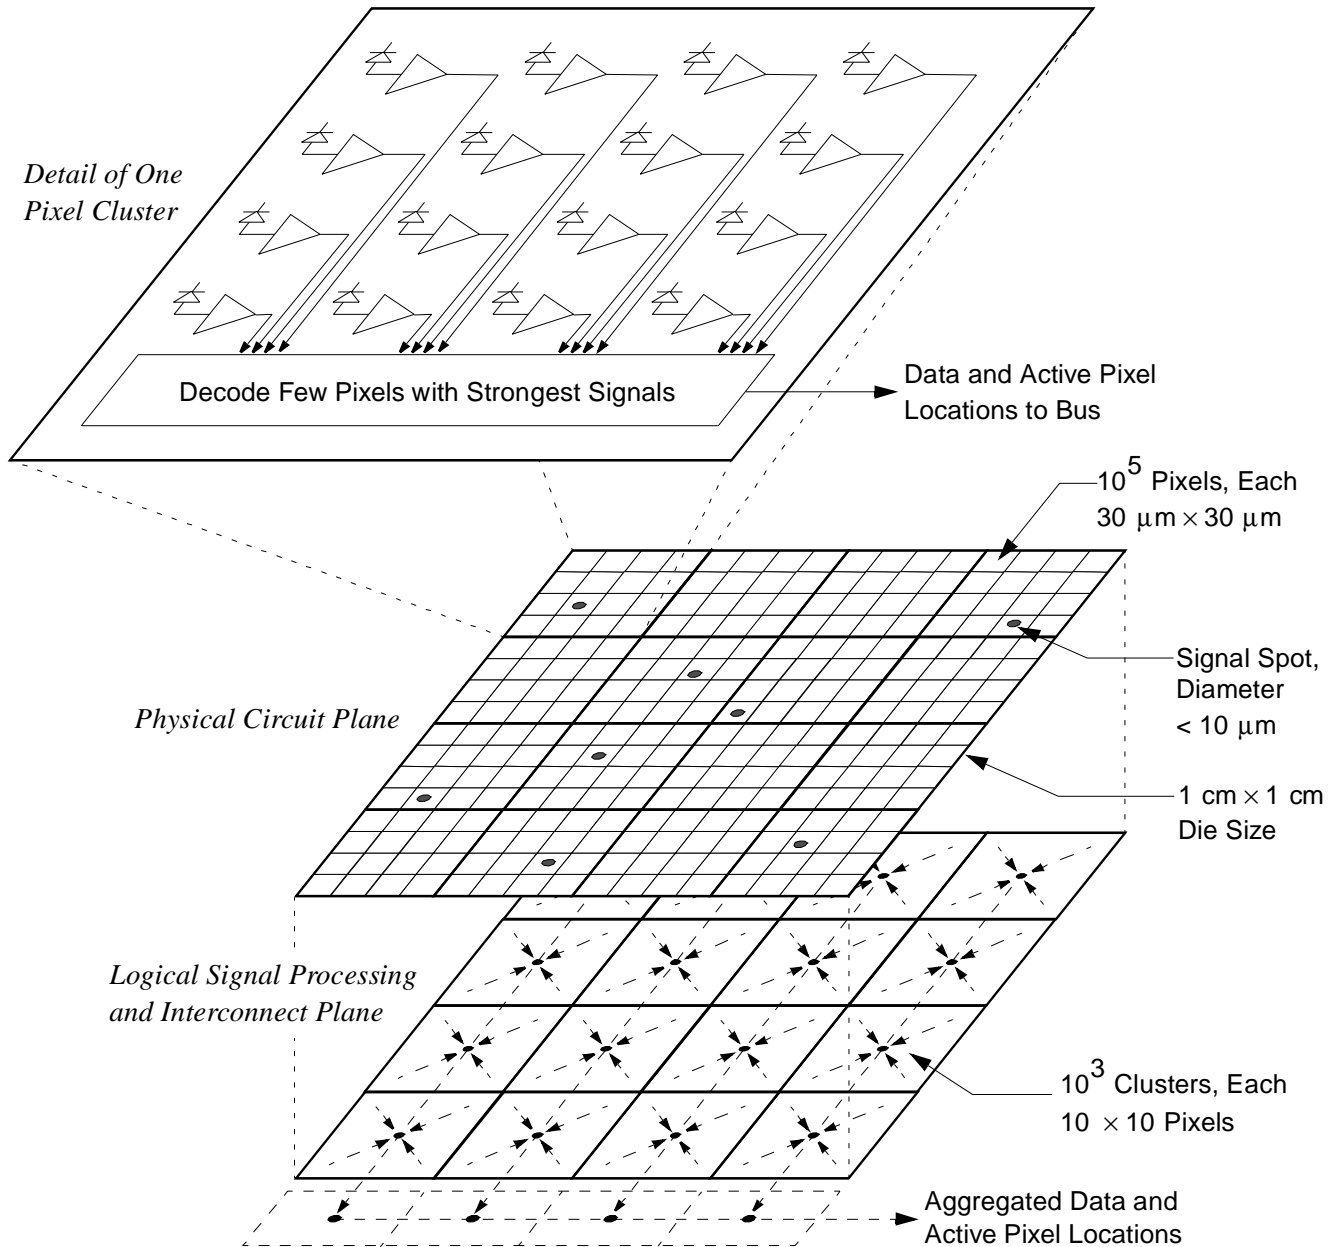


Fig. 4. Parallel optical receiver based on CMOS image sensor technology. The signal spot size is much smaller than the pixel size, so that usually most of the signal is received by a single pixel. Hierarchical signal processing and interconnection is employed. The array of 10^5 pixels is divided into 10^3 clusters, each containing 10^2 pixels. Each cluster decodes the few pixels having the strongest receptions and conveys the corresponding data and pixel locations to the bus. In this figure, the number of pixels has been made small for clarity.