

Emerging Challenges: Mobile Networking for “Smart Dust”

Joseph M. Kahn, Randy Howard Katz, and Kristofer S. J. Pister

Abstract: Large-scale networks of wireless sensors are becoming increasingly tractable. Advances in hardware technology and engineering design have led to dramatic reductions in size, power consumption and cost for digital circuitry, wireless communications and Micro ElectroMechanical Systems (MEMS). This has enabled very compact, autonomous and mobile nodes, each containing one or more sensors, computation and communication capabilities, and a power supply. The missing ingredient is the networking and applications layers needed to harness this revolutionary capability into a complete system. We review the key elements of the emergent technology of “Smart Dust” and outline the research challenges they present to the mobile networking and systems community, which must provide coherent connectivity to large numbers of mobile network nodes co-located within a small volume.

I. INTRODUCTION

As the research community searches for the processing platform beyond the personal computer, networks of wireless sensors have become quite interesting as a new environment in which to seek research challenges. Many researchers have recently shown that it is possible to integrate sensing, communication, and power supply into an inch-scale device using only off-the-shelf technology. These have been enabled by the rapid convergence of three key technologies: digital circuitry, wireless communications, and Micro ElectroMechanical Systems (MEMS). In each area, advances in hardware technology and engineering design have led to reductions in size, power consumption, and cost. This has enabled remarkably compact, autonomous nodes, each containing one or more sensors, computation and communication capabilities, and a power supply.

Fig. 1 shows two examples of this off-the-shelf sensor technology. The device on the left contains a microprocessor, temperature sensor, light sensor, 900 MHz radio, and battery (hidden underneath). The radio range is about 10 m. The device on the right consists of a microprocessor and four sensors, and the radio has been replaced with a laser pointer driven by the microprocessor. Transmitting at 4 bps to a small CCD camera attached to a PCMCIA frame grabber in a laptop computer, that device is capable of sensing and communicating weather information at a distance of over 20 km. Communication was demonstrated from San Francisco to the authors’ building in Berkeley, across the San Francisco Bay. Both devices have a full-duty lifetime of about a day.

Manuscript received June 2000.

The authors are with the department of Electrical Engineering and Computer Sciences, University of California, Berkeley, USA, e-mail: jmk@eecs.berkeley.edu, randy@eecs.berkeley.edu, pister@eecs.berkeley.edu.




Fig. 1. Cubic-inch wireless sensor nodes (motes) constructed using off-the-shelf technology. On the left is a radio-frequency mote with temperature and light sensors. On the right is a laser mote with temperature, light, humidity, barometric pressure, and sensors.

With improvements in integration, packaging, circuit design, and process technology, autonomous sensor nodes like these will continue to shrink in size and power consumption while growing in capability. Berkeley’s *Smart Dust* project, led by Professors Pister and Kahn, explores the limits on size and power consumption in such autonomous sensor nodes.

Size reduction is paramount, to make the nodes as inexpensive and easy-to-deploy as possible. The research team is working to incorporate the requisite sensing, communication, and computing hardware, along with a power supply, in a volume no more than a cubic millimeter, while still achieving impressive performance in terms of sensor functionality and communications capability. These millimeter-scale nodes are called “Smart Dust.” Although mimicking the mobility of dust is not a primary goal, future prototypes of Smart Dust will be small enough to remain suspended in air, buoyed by air currents, sensing and communicating for hours or days on end. At least one popular science fiction author has articulated just such a vision [1].

In this paper, we are concerned with the networking and applications challenges presented by this radical new technology. These kinds of networking nodes must consume extremely low power, communicate at average bit rates measured in kilobits per second, and potentially need to operate in high volumetric densities. These requirements dictate the need for novel ad hoc routing and media access solutions. Smart dust will enable an unusual range of applications, from sensor-rich “Smart spaces” to self-identification and history tracking for virtually any kind of physical object.

The study of “Smart Dust systems” is very new. The main purpose of this paper is to present some of the technological opportunities and challenges, with the goal of getting more systems-level researchers interested in this critical area. The remainder of this paper is organized as follows. Section II presents an overview of the technology that underlies Smart Dust. Sec-

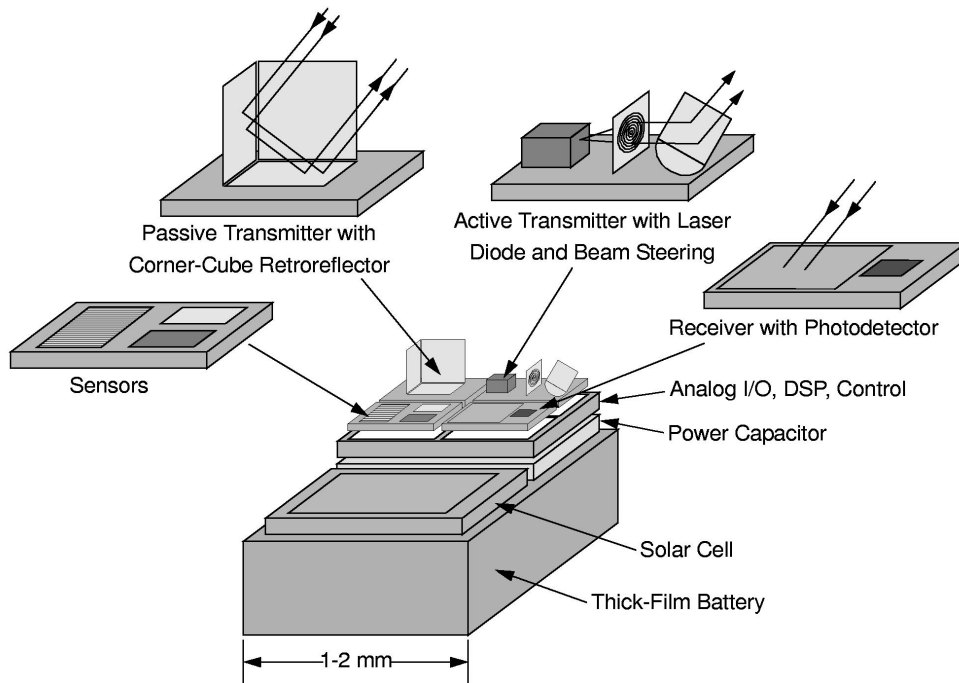


Fig. 2. Smart dust mote, containing microfabricated sensors, optical receiver, passive and active optical transmitters, signal-processing and control circuitry, and power sources.

tion III outlines the key networking challenges presented by this technology. In Section IV, we describe some of the potential applications of Smart Dust and the challenges they pose. Section V discusses related projects from the research community. Section VI presents our summary and conclusions.

II. SMART DUST TECHNOLOGY

The goal of the project, a Smart Dust *mote* is illustrated in Fig. 2. Integrated into a single package are MEMS sensors, a semiconductor laser diode and MEMS beam-steering mirror for active optical transmission, a MEMS corner-cube retroreflector for passive optical transmission, an optical receiver, signal-processing and control circuitry, and a power source based on thick-film batteries and solar cells. This remarkable package will have the ability to sense and communicate, and be self-powered!

A major challenge is to incorporate all these functions while maintaining very low power consumption, thereby maximizing operating life given the limited volume available for energy storage. Within the design goal of a cubic millimeter volume, using the best available battery technology, the total stored energy is on the order of 1 J. If this energy is consumed continuously over a day, the dust mote power consumption cannot exceed roughly $10 \mu\text{W}$. For comparison, this is roughly the “Shutdown” power of individual low power ICs found in today’s laptop computers. The functionality envisioned for Smart Dust can be achieved only if the total power consumption of a dust mote is limited to microwatt levels, and if careful power management strategies are utilized (i.e., the various parts of the dust mote are powered on only when necessary). To enable dust motes to function over the span of days, solar cells could be employed to scavenge as

much energy as possible when the sun shines (roughly 1 J per day) or when room lights are turned on (about 1 mJ per day).

With energy as the most precious resource, and time not as likely to be critical, elementary operations and ultimately algorithms are likely to be judged in terms of their energy cost, rather than power consumption. Energy-optimized microprocessors currently use roughly 1 nJ per 32-bit instruction. Commercially available data acquisition chips approach 1 nJ per sample. Bluetooth radio-frequency (RF) communication chips will burn about 100 nJ per bit transmitted. There is much excitement in the RF circuit community right now over picoradios, which target 1 nJ/bit. Dramatic improvements will be made in some of these categories, but for every Joule that a mote stores, it will have the ability to perform some billions of operations. The networking and information theory challenge is to determine how to allocate the energy: Sense? Compute? Transmit?

Techniques for performing sensing and computation at low power are reasonably well understood. Developing a communications architecture for ultra-low-power represents a more critical challenge. The primary candidate communication technologies are based on RF or optical transmission techniques. Each technique has its advantages and disadvantages. RF presents a problem because dust motes offer very limited space for antennas, thereby demanding extremely short-wavelength (i.e., high-frequency) transmission. Communication in this regime is not currently compatible with low power operation. Furthermore, radio transceivers are relatively complex circuits, making it difficult to reduce their power consumption to the required microwatt levels. They require modulation, bandpass filtering and demodulation circuitry, and additional circuitry is required if the transmissions of a large number of dust motes are to be multiplexed using time-, frequency- or code-division multiple access

[2].

An attractive alternative is to employ free-space optical transmission. Studies have shown that when a line-of-sight path is available, well-designed free-space optical links require significantly lower energy per bit than their RF counterparts [2]. There are several reasons for the power advantage of optical links. Optical transceivers require only simple baseband analog and digital circuitry; no modulators, active bandpass filters or demodulators are needed. The short wavelength of visible or near-infrared light (of the order of $1\ \mu\text{m}$) makes it possible for a millimeter-scale device to emit a narrow beam (i.e., high antenna gain can be achieved). As another consequence of this short wavelength, a base-station transceiver (BTS) equipped with a compact imaging receiver can decode the simultaneous transmissions from a large number of dust motes at different locations within the receiver field of view, which is a form of space-division multiplexing.

Successful decoding of these simultaneous transmissions requires that dust motes not block one another's line of sight to the BTS. Such blockage is unlikely, in view of the dust motes' small size. A second requirement for decoding of simultaneous transmission is that the images of different dust motes be formed on different pixels in the BTS imaging receiver. To get a feeling for the required receiver resolution, consider the following example. Suppose that the BTS views a $17\ \text{m} \times 17\ \text{m}$ area containing Smart Dust, and that it uses a high-speed video camera with a 256×256 pixel imaging array. Each pixel views an area about $6.6\ \text{cm}^2$. Hence, simultaneous transmissions can be decoded as long as the dust motes are separated by a distance roughly the size of a pack of cigarettes.

Another advantage of free-space optical transmission is that a special MEMS structure makes it possible for dust motes to use passive optical transmission techniques, i.e., to transmit modulated optical signals without supplying any optical power. This structure is a corner-cube retroreflector, or CCR (see Fig. 3). It comprises three mutually perpendicular mirrors of gold-coated polysilicon. The CCR has the property that any incident ray of light is reflected back to the source (provided that it is incident within a certain range of angles centered about the cube's body diagonal). If one of the mirrors is misaligned, this retroreflection property is spoiled. The microfabricated CCR includes an electrostatic actuator that can deflect one of the mirrors at kilohertz rates. It has been demonstrated that a CCR illuminated by an external light source can transmit back a modulated signal at kilobits per second. Since the dust mote itself does not emit light, the passive transmitter consumes little power. Using a microfabricated CCR, Chu *et al.* have demonstrated data transmission at a bit rate up to 1 kbps, and over a range up to 150 m, using a 5-mW illuminating laser [3].

It should be emphasized that CCR-based passive optical links require an uninterrupted line-of-sight path. Moreover, a CCR-based passive transmitter is inherently directional; a CCR can transmit to the BTS only when the CCR body diagonal happens to point directly toward the BTS, within a few tens of degrees. A passive transmitter can be made more omnidirectional by employing several CCRs oriented in different directions, at the expense of increased dust mote size. If a dust mote employs only one or a few CCRs, the lack of omnidirectional transmission




Fig. 3. Microfabricated corner-cube retroreflector, consisting of three gold-coated polysilicon mirrors. The base mirror can be deflected electrostatically, modulating the optical signal reflected from the device (taken from [3]).

has important implications for feasible network routing strategies (see Section III-A.2).

Fig. 4 illustrates a free-space optical network utilizing the CCR-based passive uplink. The BTS contains a laser whose beam illuminates an area containing dust motes. This beam can be modulated with downlink data, including commands to wake up and query the dust motes. When the illuminating beam is not modulated, the dust motes can use their CCRs to transmit uplink data back to the base station. A high-frame-rate CCD video camera at the BTS “sees” these CCR signals as lights blinking on and off. It decodes these blinking images to yield the uplink data. Kahn and Pister’s analysis show that this uplink scheme achieves several kilobits per second over hundreds of meters in full sunlight [2]. At night, in clear, still air, the range should extend to at least a kilometer. Because the camera uses an imaging process to separate the simultaneous transmissions from dust motes at different locations, we say that it uses *space-division multiplexing*. The ability for a video camera to resolve these transmissions is a consequence of the short wavelength of visible or near-infrared light. This does not require any coordination among the dust motes, and thus, it does not complicate their design.

When the application requires dust motes to use active optical transmitters, MEMS technology can be used to assemble a semiconductor laser, a collimating lens and a beam-steering micro-mirror, as shown in Fig. 2. Active transmitters make possible peer-to-peer communication between dust motes, provided there exists a line-of-sight path between them. Power consumption imposes a trade-off between bandwidth and range. The dust motes can communicate over longer ranges (tens of kilometers) at low data rates or higher bit rates (megabits per second) over shorter distances. The relatively high power consumption of semiconductor lasers (of the order of 1 mW) dictates that these active transmitters be used for short-duration burst-mode communication only. Sensor networks using active dust mote transmitters will require some protocol for dust motes to aim their beams toward the receiving parties.

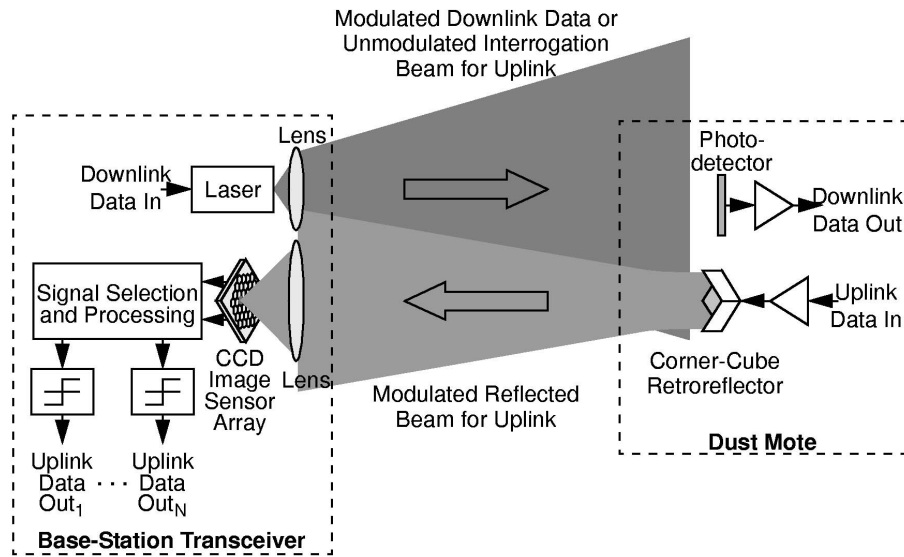


Fig. 4. Design of a free-space optical network in which a base-station transceiver communicates simultaneously with a collection of many dust motes (only one dust mote is shown). A single laser at the base station supplies optical power for the downlink and the uplink.

III. MOBILE NETWORKING CHALLENGES

A. Overview

Development of mobile networking protocols for Smart Dust represents a significant challenge. Some critical limitations are: 1) The free-space optical links requires uninterrupted line-of-sight paths, 2) the passive and active dust mote transmitters have directional characteristics that must be considered in system design, and 3) there are severe trade-offs between bit rate, energy per bit, distance and directionality in these energy-limited free-space optical links. These limitations are described in more detail in the following subsections.

A.1 Line-of-Sight Requirement

An unbroken line-of-sight path is normally required for operation of free-space optical links for Smart Dust. These links cannot operate reliably using non-line-of-sight propagation, which would rely on reflections from one or more objects between the transmitter and receiver. As shown in Section III-A.3, the transmitted beam should have a small angular spread in order to achieve a high signal-to-noise ratio with acceptably small transmitter power. Specular reflection may not significantly increase a beam's angular spread, but the existence of a properly aligned specular reflector would be a rare event. Diffuse reflection scatters a beam's energy over a wide range of angles, making alignment less critical, but usually scatters insufficient energy toward the receiver. Hence, diffuse, non-line-of-sight transmission is likely to be feasible only when active transmitters are used over very short distances (probably under 1 meter). It is probably impossible to use diffuse, non-line-of-sight transmission with passive transmitters (based on CCRs), because both the interrogating beam and the reflected beam would be subject to scattering over a wide range of angles.

A fixed dust mote without a line-of-sight path to the BTS can communicate with the BTS via multihop routing, provided that a suitable multihop path exists. The existence of such a path

is more likely when the dust mote density is higher. Multihop routing increases latency, and requires dust motes to be equipped with active optical transmitters. Constraints on size and power consumption of the dust mote digital circuitry dictate the need for low-complexity ad hoc multihop routing algorithms.

When dust motes are floating in the air or otherwise not fixed, a line-of-sight path to the BTS may become intermittently available. In such cases, the BTS can continuously interrogate the dust motes. When a line-of-sight path to a mote becomes available, the mote can transmit a packet to the BTS. When the average time between occurrence of viable line-of-sight paths is much longer than the packet duration, latency will probably be minimized by using multihop routing instead.

A.2 Link Directionality

In most Smart Dust systems, the BTS interrogating beam angular spread should be matched to the field of view of the BTS imaging receiver. These two should be matched in all systems using passive dust mote transmitters, and in systems using active dust mote transmitters when the application involves frequent bi-directional transmission between the BTS and dust motes. Intuitively, it makes little sense for the BTS to interrogate dust motes from which it cannot receive, and vice versa. In these systems, the interrogating beam and imaging receiver will be mounted rigidly together in the BTS, and will be aimed together as a unit. For example, the BTS may reside in a hand-held unit resembling a pair of binoculars, which is aimed by a human operator.

In certain applications using active dust mote transmitters, it may be desirable to use a BTS transmitter beam whose angular spread is smaller than the BTS receiver field of view. In these applications, the interrogating beam will be aimed at various locations within the receiver field of view.

Because of limited available space, the dust mote's optical receiver probably cannot employ an imaging or non-imaging op-

tical concentrator in front of the photodetector. As a result, the dust mote receiver will be fairly omnidirectional, i.e., it will be able to receive from most of the hemisphere located in front of the dust mote. In most applications, it should not be necessary to aim the dust mote receiver.

The dust mote's transmitter will exhibit markedly different directional characteristics than its receiver. A passive dust mote transmitter is based on the CCR. This device reflects light directly back to the source within a narrow beam¹, provided that it is illuminated from a direction that lies within a few tens of degrees of the cube body diagonal. If dust motes use only one CCR each, then any given dust mote, if fixed in a random, upright orientation, has only about a 10% probability of being able to transmit to the BTS. This probability can be increased significantly by equipping each dust mote with several CCRs, each oriented along a different direction. As an alternative, a single CCR may be mounted on a MEMS aiming mechanism. This mechanism need only aim the CCR with an accuracy of the order of 10 or 20°

Still other solutions exist for coping with the CCR's directionality. It may be possible to distribute randomly an excess number of dust motes, with the goal of communicating only with those whose CCRs happen to point toward the BTS. If the dust motes are not fixed, it may be best for a dust mote to simply delay transmitting until it moves into an orientation that enables transmission to the BTS.

An active dust mote transmitter is based on a laser diode. It should employ a narrow beamwidth, typically less than a degree (see Section III-A.3). This necessitates equipping the dust mote with an active beam-steering mechanism. Pister and his students are working on a MEMS-based mechanism capable of steering a beam to any position within a hemisphere. Beam-steering algorithms for systems with active dust mote transmitters represent a current research challenge. It would be desirable for each dust mote to autonomously steer its beam toward the desired direction. One approach would be to make the dust mote receiver directional, and to mount the receiver and transmitter on the same aiming mechanism. Accordingly, by aiming its receiver so as to maximize the signal received from the BTS or another mote, the dust mote would be aiming its transmitter at that node. The need for active dust mote transmitters to determine the direction to other nodes slows down connection set up, but if nodes remain fixed then the directions of various nodes, once determined, can be stored in the dust mote for future use.

Under most of the scenarios discussed above, the dust mote's transmitter and receiver have different angular spreads. This leads to non-reciprocal link characteristics, wherein a dust mote may receive from another node, but be unable to transmit to it, or vice versa. As a consequence, a dust mote may receive queries from other nodes, and may attempt to answer them, unaware that its transmissions are in vain. When dust motes are fixed, in order to conserve dust mote power, the other nodes should acknowledge this dust mote's transmissions, and this dust mote should not answer further queries from nodes that do not acknowledge its transmissions.

¹In a well-designed CCR, the angular spread of the reflected beam is limited by diffraction to the order of $\theta \sim \lambda/a$, where λ is the optical wavelength and a is the effective diameter of the CCR.

It is known that in free-space optical networks, non-reciprocity can lead to "hidden nodes" which can cause collisions during medium access. For example, this effect is observed in networks having a shared-bus physical topology, and using MAC protocols based on random time-division multiplexing, such as CSMA-CA with RTS/CTS [4]. In Smart Dust networks, the uplink (dust mote to BTS) uses space-division multiplexing. As discussed in Section II, uplink collisions will not occur as long as the dust motes are sufficiently separated that their transmissions are detected by different pixels in the BTS imaging receiver. Collisions during active peer-to-peer communications are a potential problem in Smart Dust networks. A peer-to-peer collision avoidance scheme must cope with a dynamic network configuration, while not introducing excessive complexity or latency.

A.3 Trade-Offs Between Bit Rate, Distance and Energy per Bit

Free-space optical links are subject to trade-offs between several design parameters. For simplicity, we consider the case of links employing active laser transmitters. The receiver signal-to-noise ratio (SNR) is given by

$$SNR = C \cdot \frac{E_b^2 R_b A^2}{N_0 d^4 \Phi^4}. \quad (1)$$

Here, C is a constant, E_b is the average transmitted energy per bit, R_b is the bit rate, A is the receiver light collection area², N_0 is the receiver noise power spectral density, d is the link transmission distance, and Φ is the transmitter beam angular spread. This expression assumes that Φ is small, and that the transmitter beam is well-aimed at the receiver. The SNR governs the probability of bit error, and must be maintained at a suitably high value to insure reliable link operation. From (1), we see that in order to achieve a given SNR with all other parameters fixed, the required value of E_b is proportional to $R_b^{-1/2}$, i.e., the energy per bit is minimized if packets are transmitted in short bursts at a high bit rate.

The average transmitter power (during transmission of a packet) is $P_t = E_b/R_b$. Hence, transmission at a high bit rate requires a high-power transmitter. In practice, P_t should be chosen to be as high as possible, within constraints posed by eye safety and by dust mote current-drive limitations. Rewriting (1) in terms of P_t , we obtain

$$SNR = C \cdot \frac{P_t^2 A^2}{N_0 R_b d^4 \Phi^4}. \quad (2)$$

Given a limit on P_t , to maximize the bit rate R_b and the distance d , we should maximize the receiver area A and minimize Φ , i.e., use a highly directional transmitter.

Once all other parameters have been fixed, to maintain a required SNR, the permissible bit rate and distance are related by $R_b \propto d^{-4}$. Hence, it is possible to extend the link distance by drastically lowering the bit rate. If a multihop route is available, overall latency may be minimized by transmitting at a higher bit rate over several hops.

²On a link from BTS to dust mote or from dust mote to dust mote, A corresponds to the dust mote photodetector area. On a link from dust mote to BTS, A corresponds to the BTS camera's entrance aperture area.

To give a concrete example, the inch-scale laser node shown in Fig. 1 has an average optical power of 1.5 mW and a beam divergence of roughly 1 mrad. The receiver used in the 21-km trans-bay demonstration had a 1-inch diameter lens aperture, and the signal to noise ratio at 4 bps was more than 30 dB (the upper bound was not measurable because of pixel saturation and blooming due to excessive signal strength).

B. Mobile Networking Opportunities

B.1 Overview

The optical free-space communication method presents many opportunities beyond low-power, passive communications. Since the application of interest in sensor networks is primarily sensor read-out, the key protocol issues are to perform read-out from a large volume of sensors co-located within a potentially small area. Random access to the medium is both energy-consuming and bandwidth inefficient. So it is extremely useful to exploit passive and broadcast-oriented techniques when possible. Fortunately the free-space approach supports multiple simultaneous read-out of sensors, mixes active and passive approaches using demand access techniques, and provides efficient and low-latency response to areas of a sensor network that are undergoing frequent changes. These are described in more detail in the following subsections, with emphasis on passive dust mote transmitters.

B.2 Parallel Read-Out

A single wide beam from the BTS can simultaneously probe many dust motes. The imaging receiver at the BTS receives multiple reflected beams from the motes, as long as they are sufficiently separated in space to be resolved by the receiver's pixel array. The probe beam sweeps the three dimensional space covered by the base station on a regular basis, most likely determined by the nature of the application and its need for moment-by-moment sensor readings.

B.3 Demand Access

To save transmit power, if the mote must use active communications, then it is best to use the active transmitter in a high-bit-rate, short-burst mode. Familiar demand access methods can be used to combine the low latency advantages of active communications with the low-power advantages of the passive approach.

When the mote needs to transmit information, it actively transmits a short-duration burst signal to the BTS. The BTS, detecting this signal, then probes in the general geographical area from which the burst was detected. Assuming that the passive transmitter (i.e., CCR) is properly oriented toward the BTS, the mote can respond by modulating the reflected probe beam with the data it needs to transmit.

Logically, the communications structure described above has much in common with familiar cellular and satellite networks [5]. The paging channel is acquired using contention access techniques. The BTS grants a channel to the node requesting attention. In a cellular network, this is accomplished by assigning a frequency, time slot, and/or code to the node. In the scheme

described for dust motes, the channel is "granted" by the incident probe beam.

Note that there are as many channels (paging or data) as there are resolvable pixels at the BTS. The BTS has no way to distinguish between simultaneously communicating dust motes if they fall within the same pixel in the imaging array. One possible way to deal with this is to introduce time slotted techniques not unlike that found in time division multiple access (TDMA) communications systems. A wide-aperture beam from the BTS could be modulated in such a fashion as to offer a common time base by which to synchronize the motes. The BTS can then signal an individual mote the particular time slot it has assigned to it for communication. The mote must await its time slot to communicate, whether it uses an active or a passive transmitter.

B.4 Probe Revisit Rates

Probe beam revisit rates could be determined in an application-specific manner. It is a well known observation from statistical data management that areas where changes are happening most rapidly should be revisited most frequently. If sensor readings are not changing much, then occasional samples are sufficient to obtain statistically significant results. So it is better to spend probe dwell time on those sensors that are experiencing the most rapid reading changes, and for which infrequent visit would lead to the greatest divergence from the current sensor values.

IV. APPLICATIONS

A. Introduction

Depending on the application, individual dust motes may be affixed to objects that one wishes to monitor, or a large collection of motes may simply be dispersed (and floating!) at random throughout an environment. The motes record sensor readings and, when queried, report these readings via the optical techniques described in Section II. In some applications, dust motes will communicate directly (and passively) with the BTS, in others, peer-to-peer active communication between dust motes will be used to relay information to the BTS. Depending on the application, the base station may be separated from the dust motes by distances ranging from tens of meters to kilometers.

For example, the BTS may actually reside in a hand-held unit, much like a pair of binoculars. This permits the user to simultaneously view a scene while displaying measured data overlaid on top of it. As another example, the BTS may reside in a small flying vehicle, which flies over an area to query the Smart Dust.

We envision numerous civilian and military applications for Smart Dust. Smart Dust may be deployed over a region to record data for meteorological, geophysical or planetary research. It may be employed to perform measurements in environments where wired sensors are unusable or lead to measurement errors. Examples include instrumentation of semiconductor processing chambers, rotating machinery, wind tunnels, and anechoic chambers. In biological research, Smart Dust may be used to monitor the movement, habits, and environment of insects or other small animals. Considering the military arena, Smart Dust may be deployed for stealthy monitoring of a hostile environ-

ment, e.g., for verification of treaty compliance. Here, acoustic, vibration or magnetic field sensors could detect the passage of vehicles and other equipment. Smart Dust could be used for perimeter surveillance, or to detect the presence of chemical or biological agents on a battlefield.

The overarching applications challenge, from a processing and communications viewpoint, is how to implement complex “ensemble” behavior from a large number of individual, relatively simple sensors. This is sometimes called “beehive,” “swarm,” or “emergent” behavior. A critical enabler is the ability for the sensors to communicate their readings with each other and with the more centralized intelligent processor residing at the base station. Proper design of the network is the key. We describe an applications scenario and some of the technology challenges to implement such a system in this section.

B. Scenario: Multi-Sensor Emergent Behavior

It is useful for sensors to operate in ensembles. Rather than implementing a broad range of sensors in a single integrated circuit, it is possible to simply deploy a mixture of different sensors in a given geographical area and allow them to self-organize.

Sensors are typically specialized to detect certain signatures. One kind detects motion, another heat, and a third sound. When one sensor detects its critical event signature, it makes other nearby sensors aware of its detection. They then orient their sensing function in a particular, signature-specific way. For example, a simple motion-detecting sensor might cue more sophisticated sensors detecting thermal or other radiation properties. The array, acting as an ensemble, not only performs the operation of detecting an intruder, but demonstrates more intelligent processing, by distinguishing between one that is a human and another that is a small animal (e.g., the former has a body heat signature spread over a larger volume than the latter).

A more complex sensor cued in this fashion may then increase its own scan rate to obtain a higher-resolution signature, or dedicate its detection energy budget into a particular narrow band or a specific direction. These operations have implications for power consumption. Maximizing detection probability and resolution while minimizing power consumption is a key optimization challenge.

C. Technology Approaches for Realizing the Scenario

There are two ways to construct such a cueing system. The first is a centralized scheme. The motion sensor communicates with the BTS, which in turn communicates with a nearby heat sensor. If passive communications techniques can be used, this may well be the most power-efficient way to propagate the detection information.

The centralized/passive schemes cannot be used if the line-of-sight path is blocked, or if the probe revisit rate is too infrequent to meet detection latency constraints. In these cases, the detecting mote must employ an active transmitter. If the line-of-sight path is blocked, then the mote will need to use ad hoc, multi-hop techniques to communicate with the BTS or nearby sensor nodes.

Detecting a blocked path between a mote and the BTS is not difficult (note that a blocked path and a disabled BTS can be

treated in the same way). We can assume some maximum duty cycle between probe visits. If sufficient time has passed since the last visit, the mote can assume that it is blocked. Weighted by the importance of what it has detected, the mote can decide to go active.

Building a multihop route in this environment is quite challenging. Because of the directionality of the on-board laser, active transmission in all directions is not feasible, and we cannot assume that if a next hop node receives our transmission that we will be able to receive a transmission from it.

Determining true reachability between pairs of motes requires a full four phase handshake (“Can you see me?” “Yes, I can see you. Can you see me?” “Yes” “Good. We can communicate with each other.”). This must be executed in the context of appropriate timeouts and made robust to dynamic changes in the positions of the communicating nodes, which may be floating in the air.

Assuming a static arrangement initially, we propose the following connectivity discovery algorithm for a questing mote. We assume that the mote has a unique ID and a finite set of directions in which it can point its laser. Furthermore each direction has a unique description, e.g., (x, y) , in the mote’s coordinate system. In the first phase of discovery, the mote iterates through every direction, transmitting in each case “I am ID1, and I am pointing at location (x, y) Any mote hearing this broadcast begins its own scan, transmitting “I am ID2 pointing at location (i, j) ; and I have heard mote ID1 who was pointing at location (x, y) .” Upon hearing mote2’s message, mote1 now knows where to point to talk to mote2, and in a subsequent message can tell mote2 where to point to respond.

Routing tables can be constructed from such pairwise discovery of connectivity. However, standard routing algorithms, like RIP, OSPF, and DVRMP, assume bidirectional and symmetric links. This will not always be the case for Smart Dust. It may be possible for mote A to communicate with mode B, but not vice versa. Even if the communications is bidirectional, it need not exhibit the same bandwidth or loss characteristics in both directions.

Therefore, new routing algorithms must be developed to deal with the general case of links that are unidirectional and/or asymmetric in their performance. A strong group at INRIA in France has been leading the IETF Unidirectional Link Routing Working Group discussions on these issues [6], [7].

Unfortunately, the current efforts are focusing on supporting high-bandwidth unidirectional links where all nodes have at least low-bandwidth bidirectional links (e.g., a high-bandwidth satellite link superimposed on nodes interconnected via slow-speed telephone links). Even modifying existing algorithms will not help much, since the connectivity among floating dust motes is dynamic with short time scales. The more general case still remains to be addressed.

D. Other Applications Issues

One possible improvement is to make use of emerging MEMS technology for on-board inertial navigation circuits [8] to make sensors more aware of near neighbors even as they drift out of line-of-sight of the BTS. The BTS can determine the relative

location of dust motes within its field of view. It could then disseminate this “near neighbor information” to motes able to observe its probe beam. The on-board inertial navigation capability, combined with these periodic relative location “snapshots” could assist motes in orienting their laser and detector optics to improve their ability to establish links with nearby motes.

V. RELATED PROJECTS

Several projects have recently been initiated to investigate a variety of communications research aspects of distributed sensor networks. The following description is by no means exhaustive.

The Factoid Project [9] at the Compaq Palo Alto Western Research Laboratory (WRL) is developing a portable device small enough to be attached to a key chain. The device collects announcements from broadcasting devices in the environment, and these can be uploaded to a user’s home basestation. In its first generation, the prototype devices are much larger than smart dust motes, communications is accomplished via RF transmission, and the networking depends on short-range point-to-point links.

The Wireless Integrated Network Sensors (WINS) Project [10] at UCLA is very similar in spirit to what has been described in this paper. It is developing low power MEMS-based devices that in addition to sensing and actuating can also communicate. The essential difference is that WINS has chosen to concentrate on RF communications over short distances.

The Ultralow Power Wireless Sensor Project [11] at MIT is another project that focuses on low power sensing devices that also communicate. The primary thrust is extremely low power operation. The prototype system will transmit over a range of data rates, from 1 bps to 1 Mbps, with transmission power levels that span from 10 μ W to 10 mW. The RF communications subsystem is being developed for the project by Analog Devices. Again, optical technologies are not being investigated. Ultimately the design team will need to face the multi-hop wireless networking protocol issues outlined in this paper (e.g., see [12], [13]).

VI. SUMMARY AND CONCLUSIONS

The research community is searching for a new environments in which to generate innovative ideas and prove their effectiveness. A new paradigm beyond desktop computing is capturing the imaginations of systems designs: the so-called “post-PC” era. Wireless sensor networks is one area that promises to yield important applications and demands new approaches to traditional networking problems.

We have described *Smart Dust*, an integrated approach to networks of millimeter-scale sensing/communicating nodes. Smart Dust can transmit passively using novel optical reflector technology. This provides an inexpensive way to probe a sensor or acknowledge that information was received. Active optical transmission is also possible, but consumes more power. It will be used when passive techniques cannot be used, such as when the line-of-sight path between the dust mote and BTS is blocked.

Smart dust provides a very challenging platform in which to investigate applications that can harness the emergent behavior

of ensembles of simple nodes. Dealing with partial disconnections while establishing communications via dynamic routing over rapidly changing unidirectional links poses critical research challenges for the mobile networking community.

ACKNOWLEDGMENTS

Kahn and Pister’s research is supported in part by DARPA Contract DABT63-98-1-0018, “Smart Dust.” Katz and Pister’s research is supported in part by DARPA Contract N6601-99-2-8913, “Endeavour.”

REFERENCES

- [1] N. Stephenson, *The Diamond Age*, Bantam Books, New York, 1995.
- [2] V. S. Hsu, J. M. Kahn, and K. S. J. Pister, “Wireless communications for smart dust,” Electronics Research Laboratory Memorandum Number M98/2, 1998.
- [3] P. B. Chu *et al.*, “Optical communication using micro corner cube reflectors,” in *Proc. EEE MEMS Workshop*, Nagoya, Japan, Jan. 1997, pp. 350–355.
- [4] F. Gfeller and W. Hirt, “Robust wireless infrared system with channel reciprocity,” *IEEE Commun. Mag.*, vol. 36, no. 12, pp. 100–106, Dec. 1998.
- [5] D. Goodman, *Wireless Personal Communication Systems*, Addison-Wesley Longman, Reading, MA, 1997.
- [6] W. Dabbous, E. Duros, and T. Ernst, “Dynamic routing in networks with unidirectional links,” *Workshop on Satellite-Based Information Systems*, Budapest, Sept. 1997.
- [7] Unidirectional Link Routing Protocol Working Group Home Page, available at <http://www-sop.inria.fr/rodeo/udlr/>.
- [8] B. Boser, “Electronics for micromachined inertial sensors,” *Transducers’97*, Chicago, IL, pp. 1169–1172, June 1997.
- [9] Available at <http://www.research.digital.com/wrl/projects/Factoid/index.html>.
- [10] Available at <http://www.janet.ucla.edu/WINS>.
- [11] Available at http://www-mtl.mit.edu/~jimg/project_top.html.
- [12] J. Jubin and J. D. Turnow, “The DARPA packet radio network protocols,” in *Proc. IEEE*, vol. 75, no. 1, Jan. 1987, pp. 21–32.
- [13] G. S. Lauer, *Packet-radio networks*, ch. 11 in *Routing in Communications Networks*, M. Steenstrup, Ed., Englewood Cliffs, Prentice-Hall, N.J., 1995.



Joseph M. Kahn is a Professor in the Department of Electrical Engineering and Computer Sciences at U.C. Berkeley. In 2000, he co-founded StrataLight Communications, Inc., where he is currently Chief Scientist. He received the A.B., M.A., and Ph.D. degrees in Physics from U.C. Berkeley in 1981, 1983, and 1986, respectively. From 1987–90 he was a Member of Technical Staff in the Lightwave Communications Research Department of AT&T Bell Laboratories, where he performed research on multi-gigabit-per-second coherent optical fiber transmission systems, setting several world records for receiver sensitivity. He joined the faculty of U.C. Berkeley in 1990. His current research interests include wireless communication using antenna arrays, free-space optical communication, optical fiber communication, and wireless communication for networks of sensors based on micro-electromechanical systems. Professor Kahn received the National Science Foundation Presidential Young Investigator Award in 1991. He is a Fellow of the IEEE, and is a member of the IEEE Communications Society, the IEEE Information Theory Society, and the IEEE Lasers and Electro-Optics Society. From 1993–2000, he served as a technical editor of *IEEE Personal Communications Magazine*.



Randy Howard Katz received his undergraduate degree from Cornell University, and his M.S. and Ph.D. degrees from the University of California, Berkeley. He joined the faculty at Berkeley in 1983, where he is now the United Microelectronics Corporation Distinguished Professor in Electrical Engineering and Computer Science. He is a Fellow of the ACM and the IEEE, and a member of the National Academy of Engineering. He has published over 180 refereed technical papers, book chapters, and books. His hardware design textbook, *Contemporary Logic Design*,

has sold over 85,000 copies worldwide. He has supervised 32 M.S. theses and 17 Ph.D. dissertations. He has won numerous awards, including seven best paper awards, one “test of time” paper award, three best presentation awards, the Berkeley Distinguished Teaching Award, the 1999 IEEE Reynolds Johnson Information Storage Award, the 1999 ASEE Frederic E. Terman Award, and the 1999 ACM Karl V. Karlstrom Outstanding Educator Award. With colleagues at Berkeley, he developed Redundant Arrays of Inexpensive Disks (RAID), an \$18 billion per year industry sector today. While on leave for government service in 1993–1994, he established whitehouse.gov and connecting the White House to the Internet. His current research interests are Internet Services Architecture, Mobile Computing, and Computer-Telephony Integration.



Kristofer S. J. Pister received his B.A. in Applied Physics from UCSD in 1982, and his M.S. and Ph.D. in Electrical Engineering from UC Berkeley in 1989 and 1992. From 1992 to 1997 he was an Assistant Professor of Electrical Engineering at UCLA, where he developed three graduate level courses in Micro Electro Mechanical Systems: MEMS device physics and fabrication, MEMS system design, and CAD for MEMS. In 1996 he joined the faculty of Electrical Engineering and Computer Sciences at UC Berkeley as an Associate Professor. From 1996 to

1999 he was a member of the DARPA ISAT group, and in 1997 he was chosen to be a part of the Defense Science Study Group. During the last five years, his primary research interest has been the development and use of standard MEMS fabrication technologies, micro robotics, and CAD for MEMS. He is the inventor of the polysilicon hinge, now in use by many MEMS groups around the world. Most recently his work has been in wireless sensor networks and micro optics. He is an active consultant in the MEMS industry.