ABSTRACT

Advances in hardware technology has enabled very compact, autonomous and mobile nodes each having one or more sensors, computation and communication capabilities ,and a power supply. The Smart Dust project is exploring whether an autonomous sensing, computing, and communication system can be packed into a cubic-millimeter mote to form the basis of integrated, massively distributed sensor networks. It focuses on reduction of power consumption, size and cost. To build these small sensors, processors, communication devices, and power supply , designers have used the MEMS (Micro electro mechanical Systems) technology.

Smart Dust nodes otherwise known as “motes” are usually of the size of a grain of sand and each mote consists of :

1. sensors
2. transmitter & receiver enabling bidirectional wireless communication.
3. processors and control circuitory
4. power supply unit

Using smart dust nodes, the energy to acquire and process a sample and then transmit some data about it could be as small as a few nanoJoules.

These dust motes enable a lot of applications, because at these small dimensions ,these motes can be scattered from aircraft for battle field monitoring or can be stirred into house paint to create the ultimate home sensor network.

CONTENTS

WHAT IS A SMART DUST?

Autonomous sensing and communication in a cubic millimeter

Berkeley's *Smart Dust* project, led by Professors Pister and Kahn, explores the limits on size and power consumption in autonomous sensor nodes. Size reduction is paramount, to make the nodes as inexpensive and easy-to-deploy as possible. The research team is confident that they can incorporate the requisite sensing, communication, and computing hardware, along with a power supply, in a volume no more than a few cubic millimeters, while still achieving impressive performance in terms of sensor functionality and communications capability. These millimeter-scale nodes are called "Smart Dust." It is certainly within the realm of possibility that future prototypes of Smart Dust could be small enough to remain suspended in air, buoyed by air currents, sensing and communicating for hours or days on end.

'Smart dust' — sensor-laden networked computer nodes that are just cubic millimetres in volume. The smart dust project envisions a complete sensor network node, including power supply, processor, sensor and communications mechanisms, in a single cubic millimetre. .Smart dust motes could run for years , given that a cubic millimetre battery can store 1J and could be backed up with a solar cell or vibrational energy source

The goal of the Smart Dust project is to build a millimeter-scale sensing and communication platform for a massively distributed sensor network. This device will be around the size of a grain of sand and will contain sensors, computational ability, bi-directional wireless communications, and a power supply. Smart dust consists of series of circuit and micro-electro-mechanical systems (MEMS) designs to cast those functions into custom silicon. Microelectromechanical systems (MEMS) consist of extremely tiny mechanical elements, often integrated together with electronic circuitry.

THE MEMS TECHNOLOGY IN SMART DUST

Smart dust requires mainly revolutionary advances in miniaturization, integration & energy management. Hence designers have used MEMS technology to build small sensors, optical communication components, and power supplies. Microelectro mechanical systems consists of extremely tiny mechanical elements, often integrated together with electronic circuitory. They are measured in micrometers, that is millions of a meter. They are made in a similar fashion as computer chips. The advantage of this manufacturing process is not simply that small structures can be achieved but also that thousands or even millions of system elements can be fabricated simultaneously. This allows systems to be both highly complex and extremely low-cost.

Micro-Electro-Mechanical Systems (MEMS) is the integration of mechanical elements, sensors, actuators, and electronics on a common silicon substrate through microfabrication technology. While the electronics are fabricated using integrated circuit (IC) process sequences (e.g., CMOS, Bipolar processes), the micromechanical components are fabricated using compatible "micromachining" processes that selectively etch away parts of the silicon wafer or add new structural layers to form the mechanical and electromechanical devices. MEMS realizes a complete System On chip technology.

Microelectronic integrated circuits can be thought of as the "brains" of a system and allow microsystems to sense and control the environment. Sensors gather information from the environment through measuring mechanical, thermal, biological, chemical, optical, and magnetic phenomena. The electronics then process the information derived from the sensors and through some decision making capability direct the actuators to respond by moving, positioning, regulating, and filtering, thereby controlling the environment for some desired purpose. Because MEMS devices are manufactured using batch fabrication techniques similar to those used for integrated circuits, unprecedented levels of functionality, reliability, and sophistication can be placed on a small silicon chip at a relatively low cost. The deep insight of MEMS is as a new manufacturing technology, a way of making complex electromechanical systems using batch fabrication techniques similar to those used for integrated circuits, and uniting these electromechanical elements together with electronics.Historically, sensors and actuators are the most costly and unreliable part of a sensor-actuator-electronics system. MEMS technology allows these complex electromechanical systems to be manufactured using batch fabrication techniques, increasing the reliability of the sensors and actuators to equal that of integrated circuits. The performance of MEMS devices and systems is expected to be superior to macroscale components and systems, the price is predicted to be much lower.

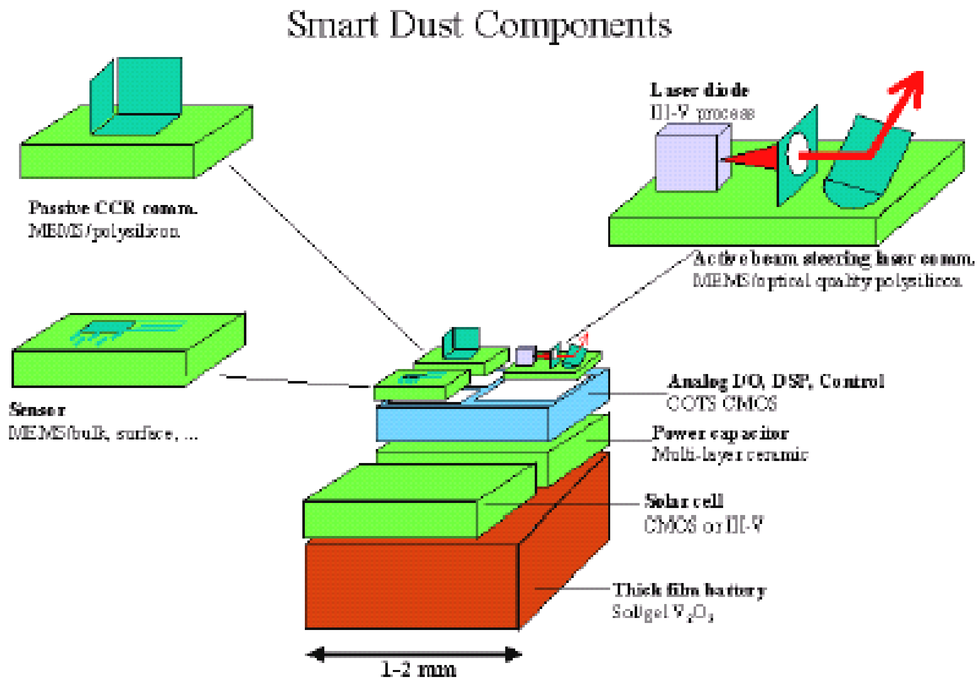

Figure 1: *Conceptual Diagram of the Smart Dust Mote for $1mm^3$ autonomous distributed sensing and communication*

SMART DUST TECHNOLOGY

Integrated into a single package are :-

1. MEMS sensors
2. MEMS beam steering mirror for active optical transmission
3. MEMS corner cube retroreflector for passive optical transmission
4. An optical receiver
5. Signal processing and control circuitory
6. A power source based on thick film batteries and solar cells

This remarkable package has the ability to sense and communicate and is self powered. A major challenge is to incorporate all these functions while maintaining very low power consumption.

- Sensors collect information from the environment such as light , sound, temperature ,chemical composition etc
- Smart dust employs 2 types of transmission schemes:-passive transmission using corner cube retroreflector to transmit to base stations and active transmission using a laser diode & steerable mirrors for mote to mote communication.
- The photo diode allows optical data reception
- Signal processing & control circuitory consists of analog I/O ,DSPs to control &process the incoming data
- The power system consists of a thick film battery,a solar cell with a charge integrating capacitor for a period of darkness.

OPERATION OF THE MOTE

The Smart Dust mote is run by a microcontroller that not only determines the tasks performed by the mote, but controls power to the various components of the system to conserve energy. Periodically the microcontroller gets a reading from one of the sensors, which measure one of a number of physical or chemical stimuli such as temperature, ambient light, vibration, acceleration, or air pressure, processes the data, and stores it in memory. It also occasionally turns on the optical receiver to see if anyone is trying to communicate with it. This communication may include new programs or messages from other motes. In response to a message or upon its own initiative the microcontroller will use the corner cube retroreflector or laser to transmit sensor data or a message to a base station or another mote.

The primary constraint in the design of the Smart Dust motes is volume, which in turn puts a severe constraint on energy since we do not have much room for batteries or large solar cells. Thus, the motes must operate efficiently and conserve energy whenever possible. Most of the time, the majority of the mote is powered off with only a clock and a few timers running. When a timer expires, it powers up a part of the mote to carry out a job, then powers off. A few of the timers control the sensors that measure one of a number of physical or chemical stimuli such as temperature, ambient light, vibration, acceleration, or air pressure. When one of these timers expires, it powers up the corresponding sensor, takes a sample, and converts it to a digital word. If the data is interesting, it may either be stored directly in the SRAM or the microcontroller is powered up to perform more complex operations with it. When this task is complete, everything is again powered down and the timer begins counting again.

Another timer controls the receiver. When that timer expires, the receiver powers up and looks for an incoming packet. If it doesn't see one after a certain length of time, it is powered down again. The mote can receive several types of packets, including ones that are new program code that is stored in the program memory. This allows the user to change the behavior of the mote remotely. Packets may also include messages from the base station or other motes. When one of these is received, the microcontroller is powered up and used to interpret the contents of the message. The message may tell the mote to do something in particular, or it may be a message that is just being passed from one mote to another on its way to a particular destination. In response to a message or to another timer expiring, the microcontroller will assemble a packet containing sensor data or a message and transmit it using either the corner cube retroreflector or the laser diode, depending on which it has. The laser diode contains the onboard laser which sends signals to the base station by blinking on and off. The corner cube retroreflector , transmits information just by moving a mirror and thus changing the reflection of a laser beam from the base station.

This technique is substantially more energy efficient than actually generating some radiation. With the laser diode and a set of beam scanning mirrors, we can transmit data in any direction desired, allowing the mote to communicate with other Smart Dust motes.

COMMUNICATING WITH A SMART DUST

COMMUNICATING FROM A GRAIN OF SAND

Smart Dust's full potential can only be attained when the sensor nodes communicate with one another or with a central base station. Wireless communication facilitates simultaneous data collection from thousands of sensors. There are several options for communicating to and from a cubic-millimeter computer.

Radio-frequency and optical communications each have their strengths and weaknesses. Radio-frequency communication is well under-stood, but currently requires minimum power levels in the multiple milliwatt range due to analog mixers, filters, and oscillators. If whisker-thin antennas of centimeter length can be accepted as a part of a dust mote, then reasonably efficient antennas can be made for radio-frequency communication. While the smallest complete radios are still on the order of a few hundred cubic millimeters, there is active work in the industry to produce cubic-millmeter radios.

Moreover RF techniques cannot be used because of the following disadvantages :-

1. Dust motes offer very limited space for antennas, thereby demanding extremely short wavelength (high frequency transmission). Communication in this regime is not currently compatible with low power operation of the smart dust.
2. Furthermore radio transceivers are relatively complex circuits making it difficult to reduce their power consumption to required microwatt levels.
3. They require modulation, band pass filtering and demodulation circuitory.

So an attractive alternative is to employ free space optical transmission. Studies have shown that when a line of sight path is available , well defined free space optical links require significantly lower energy per bit than their RF counterpaths.

There are several reasons for power advantage of optical links.

1. Optical transceivers require only simple baseband analog and digital circuitory .
2. No modulators,active band pass filters or demodulators are needed.
3. The short wavelength of visible or near infra red light (of the order of 1 micron) makes it possible for a millimeter scale device to emit a narrow beam (ie, 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 from 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 dust mote's small size.

Semiconductor lasers and diode receivers are intrinsically small, and the corresponding transmission and detection circuitry for on/off keyed optical communication is more amenable to low-power operation than most radio schema. Perhaps most important, optical power can be collimated in tight beams even from small apertures. Diffraction enforces a fundamental limit on the divergence of a beam, whether it comes from an antenna or a lens. Laser pointers are cheap examples of milliradian collimation from a millimeter aperture. To get similar collimation for a 1-GHz radio-frequency signal would require an antenna 100 meters across, due to the difference in wavelength of the two transmissions. As a result, optical transmitters of millimeter size can get antenna gains of one million or more, while similarly sized radio-frequency antennas are doomed by physics to be mostly isotropic.

Collimated optical communication has two major drawbacks. Line of sight is required for all but the shortest distances, and narrow beams imply the need for accurate pointing. Of these, the pointing accuracy can be solved by MEMS technology and clever algorithms, but an optical transmitter under a leaf or in a shirt pocket is of little use to anyone. We have chosen to explore optical communication in some depth due to the potential for extreme low-power communication.

OPTICAL COMMUNICATIONS

We have explored two approaches to optical communications: passive reflective systems and active-steered laser systems. In a passive communication system, the dust mote does not require an onboard light source. Instead, a special configuration of mirrors can either reflect or not reflect light to a remote source.

Passive reflective systems

The passive reflective communication is obtained by a special device called CCR (Corner cube retro reflector) consists of three mutually orthogonal mirrors. Light enters the CCR, bounces off each of the three mirrors, and is reflected back parallel to the direction it entered. In the MEMS version, the device has one mirror mounted on a spring at an angle slightly askew from perpendicularity to the other mirrors.

In this position, because the light entering the CCR does not return along the same entry path, little light returns to the source—a digital 0. Applying voltage between this mirror and an electrode beneath it causes the mirror to shift to a position perpendicular to other mirrors, thus causing the light entering the CCR to return to its source—a digital 1. The mirror's low mass allows the CCR to switch between these two states up to a thousand times per second, using less than a nanojoule per 0→1 transition. A 1□0 transition, on the other hand, is practically free because dumping the charge stored on the electrode to the ground requires almost no energy. Our latest Smart Dust device is a 63-mm3 autonomous bidirectional communication mote that receives an optical signal, generates a pseudorandom sequence based on this signal to emulate sensor data, and then optically transmits the result. The system contains a micromachined corner-cube reflector, a 0.078-mm3 CMOS chip that draws 17 microwatts, and a hearing aid battery. In addition to a battery based operation, we have also powered the device using a 2-mm2 solar cell. This mote demonstrates Smart Dust's essential concepts, such as optical data transmission, data processing, energy management, miniaturization, and system integration.

A passive communication system suffers several limitations. Unable to communicate with each other, motes rely on a central station equipped with a light source to send and receive data from other motes. If a given mote does have a clear line of sight to the central station, that mote will be isolated from the network. Also, because the CCR reflects only a small fraction of the light emitted from the base station, this system's range cannot easily extend beyond 1 kilometer. To circumvent these limitations, dust motes must be active and have their own onboard light source.

Active-steered laser systems

For mote-to-mote communication, an active-steered laser communication system uses an onboard light source to send a tightly collimated light beam toward an intended receiver. Steered laser communication has the advantage of high power density; for example, a 1-milliwatt laser radiating into 1 milliradian (3.4 arcseconds) has a density of approximately 318 kilowatts per steradian (there are 4 π steradians in a sphere), as opposed to a 100-watt lightbulb that radiates 8 watts per steradian isotropically. A Smart Dust mote's emitted beam would have a divergence of approximately 1 milliradian, permitting communication over enormous distances using milliwatts of power. Each mote must carefully weigh the needs to sense, compute, communicate, and evaluate its energy reserve status before allocating precious nanojoules of energy to turn on its transmitter or receiver. Because these motes spend most of their time sleeping, with their receivers turned off, scheduling a common awake time across the network is difficult. If motes don't wake up in a synchronized manner, a highly dynamic network topology and large packet latency result. Using burstmode communication, in which the laser operates at up to several tens of megabits per second for a few milliseconds, provides the most energy-efficient way to schedule this network. This procedure minimizes the mote's duty cycle and better utilizes its energy reserves. The steered agile laser transmitter consists of a semiconductor diode laser coupled with a collimating lens and MEMS beam-steering optics based on a two degree-of-freedom silicon micromirror. This system integrates all optical components into an active 8-mm3 volume as the figure shows

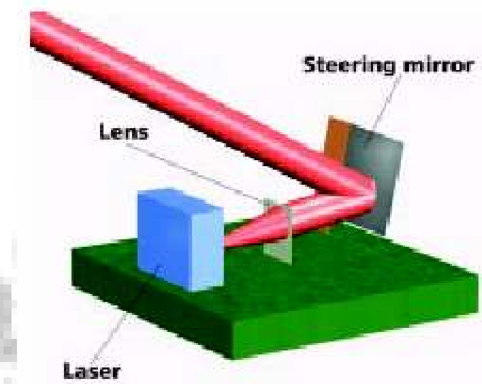

CORNER CUBE RETROREFLECTOR

These MEMS structure makes it possible for dust motes to use passive optical transmission techniques ie, to transmit modulated optical signals without supplying any optical power. It comprises of 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 contains an electrostatic actuator that can deflect one of the mirrors at kilohertz rate. 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 , passive transmitter consumes little power. Using a microfabricated CCR, data transmission at a bit rate upto 1 kilobit per second and upto a range of 150 mts ,using a 5 milliwattt illuminating laser is possible.

It should be emphasized that CCR based passive optical links require an uninterrupted line of sight. The CCR based transmitter is highly directional. A CCR can transmit to the BTS only when the CCR body diagonal happens to point directly towards 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.

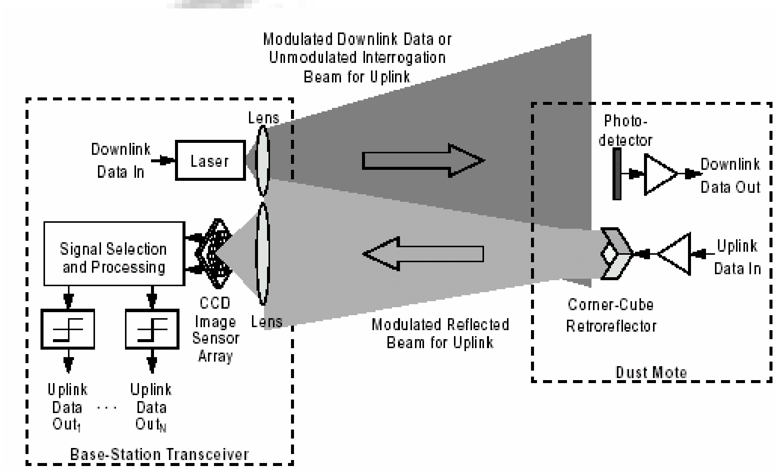

The figure illustrates free space optical network utilizing the CCR based passive uplink. The BTS contains a laser whose beam illuminates an area containing dust motes. This beam can be modulated with downlink data including commands to wake up and query the dust motes. When the illuminating beam is 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 the CCR signals as lights blinking on and off. It decodes these blinking images to yield the uplink data. Analysis show that this uplink scheme achieves several kilobits per second over hundreds of metres in full sunlight. At night ,in clear ,still air ,the range should extend to several kilometres. Because the camera uses an imaging process to separate the simultaneous transmissions from dust motes at different locations, we say it uses ‘space division multiplexing’. The ability for a video camera to resolve these transmissions is the consequence of the short wavelength of visible or near infra red light. This does not require any coordination among the dust motes.

ACTIVE OPTICAL TRANSMITTERS

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. Active transmitters make possible peer to peer communication between dust motes, provided there exists a line of path of sight between them. Power consumption imposes a trade off between bandwidth and range. The dust motes can communicate over long distances at low data rates or higher bit rates over shorter distances. The relatively higher power consumption of semiconductor lasers dictates that these active transmitters be used for short duration burst mode communication only. Sensor network using active dust mote transmitters will require some protocol for dust motes to aim their beams towards the receiving parties.

LISTENING TO A DUST FIELD

Many Smart Dust applications rely on direct optical communication from an entire field of dust motes to one or more base stations. These base stations must therefore be able to receive a volume of simultaneous optical transmissions. Further, communication must be possible outdoors in bright sunlight which has an intensity of approximately 1 kilowatt per square meter, although the dust motes each transmit information with a few milliwatts of power. Using a narrow-band optical filter to eliminate all sunlight except the portion near the light frequency used for communication can partially solve this second problem, but the ambient optical power often remains much stronger than the received signal power.

Advantages of imaging receivers

As with the transmitter, the short wavelength of optical transmissions compared with radio frequency overcomes both challenges. Light from a large field of view field can be focused into an image, as in our eyes or in a camera. Imaging receivers utilize this to analyze different portions of the image separately to process simultaneous transmissions from different angles. This method of distinguishing transmissions based on their originating location is referred to as space division multiple access (SDMA). In contrast, most radio-frequency antennas receive all incident radio power in a single signal, which requires using additional tactics, such as frequency tuning or code division multiple access (CDMA), to separate simultaneous transmissions.

Imaging receivers also offer the advantage of dramatically decreasing the ratio of ambient optical power to received signal power.Ideally, the imaging receiver will focus all of the received power from a single transmission onto a single photodetector. If the receiver has an $n \times n$ array of pixels, then the ambient light that each pixel receives is reduced by a factor n2 compared with a nonimaging receiver.

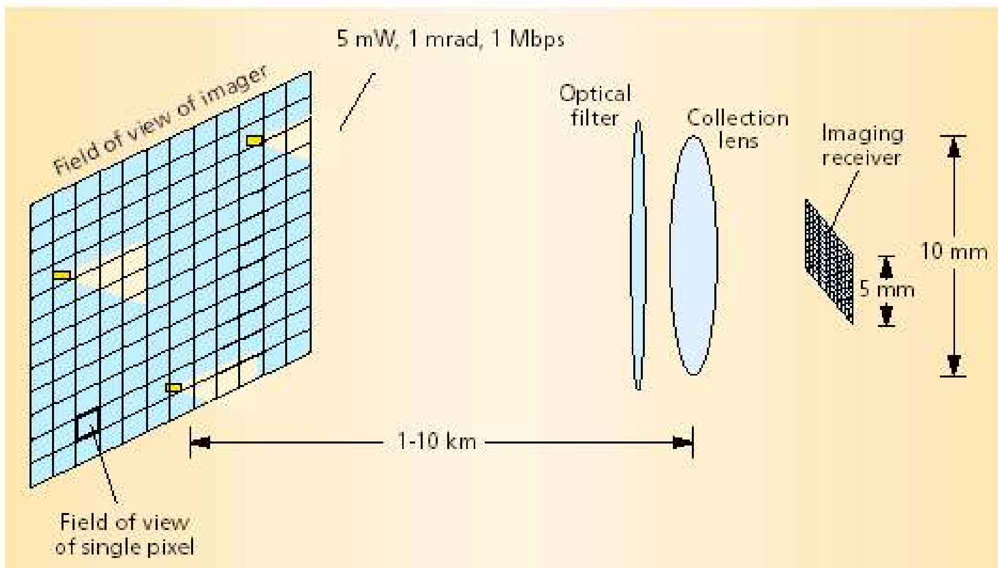

Figure 6. Pictorial representation of an integrated imaging receiver in action with predicted specifications. The receiver leverages the power of integrated circuits and CMOS imaging sensors to create a microchip with a complete asynchronous receiver circuit integrated into every pixel in the imaging array. Each pixel autonomously monitors its own signal as it searches for a transmission, decodes it, and alerts the network when it receives a data packet.

Video camera.

A video camera is a straightforward implementation of an imaging receiver. If each member in a colony of Smart Dust motes flashes its own signal at a rate of a few bits per second, then each transmitter will appear in the video stream at a different location in the image. Using a high-speed camera and a dedicated digital signal processor to process the video signal achieves higher data rates. With modern cameras and DSPs, processing video at about 1,000 frames per second should be feasible. This would allow communication at a few hundred bits per second, which is acceptable for many applications. An alternative receiver architecture provides a more elegant solution at much higher data rates, avoiding the need for computationally intensive video processing and very high speed cameras. Integrating an imaging receiver onto a single microchip imposes severe constraints in silicon area and power consumption per pixel. Only recently have continuing reductions in transistor size allowed for sufficient reductions in circuit area and power consumption to achieve this level of integration.

CORE FUNCTIONALITY SPECIFICATION

Choose the case of military base monitoring wherein on the order of a thousand Smart Dust motes are deployed outside a base by a micro air vehicle to monitor vehicle movement. The motes can be used to determine when vehicles were moving, what type of vehicle it was, and possibly how fast it was travelling. The motes may contain sensors for vibration, sound, light, IR, temperature, and magnetization. CCRs will be used for transmission, so communication will only be between a base station and the motes, not between motes. A typical operation for this scenario would be to acquire data, store it for a day or two, then upload the data after being interrogated with a laser. However, to really see what functionality the architecture needed to provide and how much reconfigurability would be necessary, an exhaustive list of the potential activities in this scenario was made. The operations that the mote must perform can be broken down into two categories: those that provoke an immediate action and those that reconfigure the mote to affect future behavior.

Proposed Architecture

Looking through the functional specifications for the core, we realized that each operation is regulated by a timed event; hence a bank of timers forms the basis of the architecture. For minimum energy, a direct mapping of a particular function into hardware is generally best, but from the list of specifications it was clear that a certain amount of reconfigurability would be necessary. Thus, the timers enable setup memories that configure functional blocks into data paths that provide only the capabilities necessary for that event. These paths are data-driven so that functional blocks are only powered up when their inputs are ready, minimizing standby power and glitching. A block diagram of this new architecture is shown in the figure

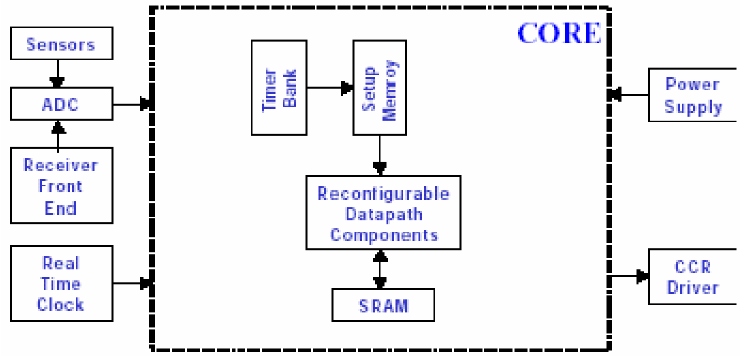

Figure 2: *Top-level diagram of the mote and the proposed core architecture architecture*

The next figure details a section of the timer bank and setup memory. The timer is loaded from the timer value memory, setting its period. When the timer expires, it enables *setupmemory 1*, which configures the data path to perform the desired function. When the data path has finished its operation, *setup memory 1* will release its configuration and **y**either the timer value can be loaded into the timer and thecountdown restarted or *setup memory 2* can be enabled.

Setup memory 2 will then configure the data path for another operation, thus facilitating multiple operations per timer event. Additional setup memory can be added for more involved sequences. Memory holds certain timer-independent configuration bits, such as timer enables. The sensor registers are used to store previous sensor readings to use in computing data changes. Various computation blocks can be included in the data path, such as an adder, comparator, and FFT unit.

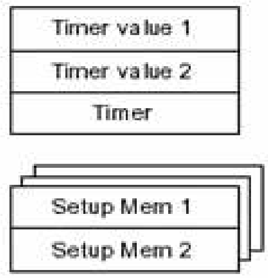

Figure 3: *Timer and associated setup memories: The timer allows multiple timer periods. The different setup memories allow multiple data path configurations per event. Energy-driven operation isfacilitatedby multiple banks of setup memory facilitated*

Multiple timer periods are desirable for several situations. For example, one might want to sample a sensor at a slow rate until an interesting signal is detected. At that point, the sampling rate should increase. In addition, the motes might be deployed without anyone coming back to talk to them for a day, so it would be desirable to be able to set the receiver wake-up timer to not wake-up for 24 hours, but then it should decrease the period dramatically to 10's of seconds in case one doesn't make it back to talk to the mote at exactly the right time. The proposed architecture facilitates this by providing multiple timer values that can be loaded into the timer depending on the results of the data path computation.

Another feature of this architecture is energy-driven operation modes. An energy-monitoring unit selects between multiple banks of setup memory and timer values depending on the current level of the energy stores. Each bank can have different timer periods and algorithms to control energy expenditure. Two types of packets can be sent to the mote, corresponding to the two types of operations. Immediate mode operations use the packet body to configure the data path right away. Reconfiguration operations load the packet body into the setup memory for future configuration.

The following figure shows the functional blocks included in the reconfigurable data path.

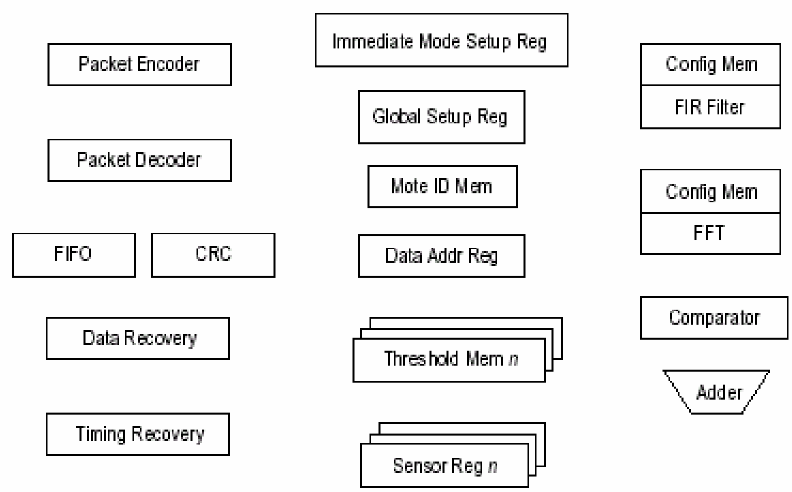

***Figure 4:** Functional blocks whose connectivity is configured by the setup memory when a timer expires. The global setup memory is always enabled.*

For the communications back end, there is a data recovery block, timing recovery block, FIR filter, packet encoder that does bits stuffing and adds the flag byte, packet decoder that does bit unstuffing, CRC block, and a FIFO. Incoming packets are stored in the FIFO until the CRC can be verified, at which point the packet body will be used as described above. The global memory holds certain timer-independent configuration bits, such as timer enables. The sensor registers are used to store previous sensor readings to use in computing data changes. Various computation blocks can be included in the data path, such as an adder, comparator, and FFT unit.

All of the functional units in the data path are data driven. The setup memory only powers up and enables the first set of units that are needed, such as the sensor and ADC. Once these units have done their job, they assert a *done* signal that is routed, based on the configuration memory, to the next unit, such as the adder, and powers it up and enables it. Likewise, when this unit has finished its job, it will power up and enable the next device in the chain. The last unit in the path will cause the timer to reload its value and cause the setup memory to stop configuring the data path. The advantages of this data driven technique include minimizing the standby power by keeping components powered down until exactly when they are needed, and ensuring that the inputs are stable before the next device is powered up, which minimizes glitches. It is significant to note that since this architecture does not use shared busses as in traditional microcontrollers, the functional components can be configured for certain parallel operations. For example, a sensor reading could be both stored in SRAM and transmitted with the CCR, although this is not necessarily a desirable capability.

PERFORMING A TASK

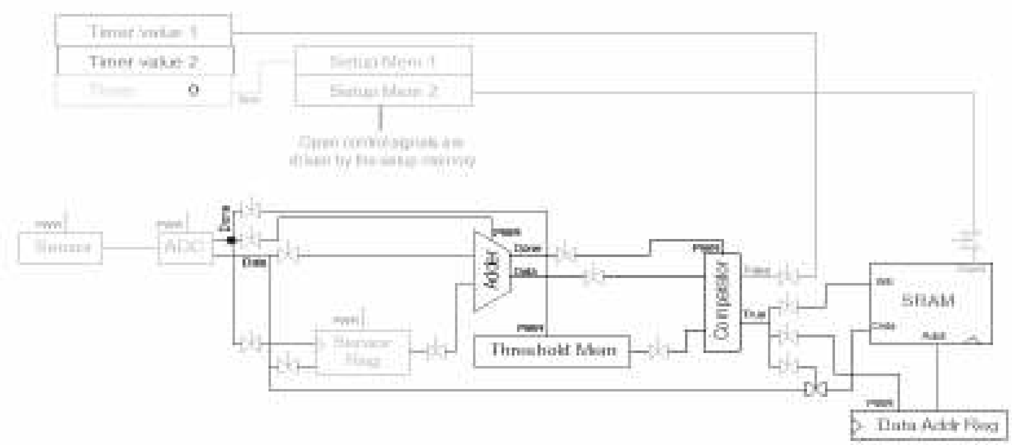

Figure 5: *Example configuration to sample data, find the change in value, and store it in the SRAM is the change exceeds a certain threshold. The colors indicate the order in which the components and signals become active, starting with the timer, proceeding to the SRAM and back to the timer.*

Figure 5 delineates the operation of the architecture by showing the configuration for one of the most common tasks, acquiring sensor data, checking if it has changed more than a threshold value, then storing the result to memory. One potential hazard of this architecture is that the *done* signals can glitch as the blocks are powered up, which would provide a false trigger to the next stage. A second issue is that despite the fact that the blocks are powered down, the internal nodes do not discharge immediately. An advantage of this is that less charge will be needed when the block is powered up again. However, this stored charge will also allow the block to continue to drive its outputs despite being powered down, so the outputs will generally need tri-state buffers. These hazards will require some extra work at the circuit level to make this architecture work.

Hspice simulations were used simulations were used to determine the power and energy consumptions of some of the blocks to study the feasibility of the proposed architecture. The preliminary results of Hspice simulations show that it is possible to achieve atleast two orders of magnitude lower energy consumption, with this proposed architecture.

MAJOR CHALLENGES

1. To incorporate all these functions while maintaining a low power consumption
2. Maximising operating life given the limited volume of energy storage
3. The functionality can be achieved only if the total power consumption is limited to microwatt levels.
4. An unbroken line of sight of path should be available for free space optical links.

APPLICATIONS

1. Civil and military applications where chemical & biological agents in a battle field are detected.
2. Virtual keyboard Glue a dust mote on each of your fingernails. Accelerometers will sense the orientation and motion of each of your fingertips, and talk to the computer in your watch. Combined with a MEMS augmented-reality heads-up display, your entire computer I/O would be invisible to the people around you.
3. Inventory Control Smart office spaces The Center for the Built Environment has fabulous plans for the office of the future in which environmental conditions are tailored to the desires of every individual. Maybe soon we'll all be wearing temperature, humidity, and environmental comfort sensors sewn into our clothes, continuously talking to our workspaces which will deliver conditions tailored to our needs.
4. Individual dust motes can be attached to the objects one wishes to monitor or a large no: of dust motes may be dispersed in the environment randomly.
5. Dust motes may be used in places where wired sensors are unusable or may lead to errors. Eg:- Instrumentation of semiconductor processing chambers,wind tennels, rotating machinery etc.
6. May be used in biological research eg:- to monitor movements & internal processes of insects.

HOW FAR THEY HAVE BEEN IMPLEMENTED

1. The optical receiver for the smart dust project is being developed. The receiver senses incoming laser transmissions at up to 1Mbit/s, for a power consumption of $12\mu W$. Although this is too high for continuous use in smart dust, it is a reasonable figure for the download of small amounts of data such as a 1Kbit program.
2. For data transmission, the team is using corner cube retro-reflectors (CCRs) built using MEMS techniques. CCRs are produced by placing three mirrors at right angles to each other to form the corner of a box that has been silvered inside.

   The key property of a CCR is that light entering it is reflected back along the path it entered on. For the smart dust system, the CCR is being built on a MEMS process with the two vertical sides being assembled by hand. When a light is shone into the CCR, it reflects back to the sending position. By modulating the position of one of the mirrors, the reflected beam can be modulated, producing a low-energy passive transmission.
3. The analog-digital convertor (ADC) the 8bit ADC, has so far demonstrated with an input range of 1V, equal to the power supply, and a 70kHz sampling rate. The converter draws $1.8\mu W$ when sampling at that rate, or 27pJ for an 8bit sample.
4. The latest smart dust mote, with a volume of just 16cu mm, has been tested. It takes samples from a photo-detector, transmits their values with the CCR and runs off solar cells. So smart dust is on the way.

SUMMARY

Smart dust is made up of thousands of sand-grain-sized sensors that can measure ambient light and temperature. The sensors -- each one is called a "mote" -- have wireless communications devices attached to them, and if you put a bunch of them near each other, they'll network themselves automatically.

These sensors, which would cost pennies each if mass-produced, could be plastered all over office buildings and homes. Each room in an office building might have a hundred or even a thousand light- and temperature-sensing motes, all of which would tie into a central computer that regulates energy usage in the building.

Taken together, the motes would constitute a huge sensor network of smart dust, a network that would give engineers insight into how energy is used and how it can be conserved. In a dust-enabled building, computers would turn off lights and climate control in empty rooms. During peak energy usage times, air conditioners that cool servers -- which drain a lot of the tech world's power -- would be automatically shut off, and then turned on again if the servers get too hot. Thus it can very lead to world's energy conservation solutions.

REFERENCES

1. P.B.Chu , KSJ Pister – “optical communication using micro corner cube reflectors” 10th IEEE Int’l Workshop on Micro Electro Mechanical Systems

2. J . Kahn , R.H.Katz, KSJ Pister – “Mobile Networking for Smart Dust”

3. www.eecs.berkerly.edu