On the Limits and Applications of MEMS Sensor Networks

Kris Pister, UC Berkeley

Abstract

Wireless sensor nodes will become inexpensive and common over the next decade. Some of the physical limits to the underlying technology are discussed. Predictions of future performance are made. Three militarily relevant applications examples are given.

Introduction

We are surrounded by sensor networks. We drive in cars (which have seat occupation and belt sensors) on roads (that have car presence sensors), to work in buildings (that have temperature and motion sensors), which are all part of the tremendous infrastructure that we take for granted, in part because of the sensor networks that help to make it maintainable.

We are increasingly surrounded by wireless communication networks. The cell phone and pager networks are the most obvious and recent examples. Microwave towers and satellite links have become so common that they are no longer noticed.

We are surrounded by computation. Most of us carry at least one (admittedly simple) computer on our person all day long - wristwatch, cell phone, hearing aid, etc. In the latter two cases, the signal processing capabilities of the silicon we wear exceeds the capabilities of the most powerful computers just a few decades ago, yet we complain that the batteries run low too quickly!

No one seriously questions the exponential improvement in computing technology. This work explores a few of the military implications of exponential improvement in all three of the above capabilities: sensing, computation, and communication. Where are the limits, and what are some of the applications?

Technology Roadmap

Existing technology

There are many groups currently working under DARPA funding on wireless sensor networks using MEMS technology. Initially the DARPA effort focussed on developing the sensor technology itself. As sensor capabilities improved, the emphasis shifted to developing sensor systems. (One of?) the first of these was an effort led by Ken Wise at the University of Michigan[i], in which the goal was to produce a wrist-watch sized, battery powered sensor system (Figure 1.). Shortly thereafter, Bill Kaiser and the humble author at UCLA launched the LWIM project (Low Power Wireless Microsensors)[ii] with the goal of putting a completely autonomous sensor node with power, processing, and communication into a cubic centimeter volume. LWIM has been quite successful, with many technology demonstrations in military exercises. The success of the project has spawned many follow-on contracts at UCLA including WINS and AWAIRS. In 1998 the humble author, now at UC Berkeley, was funded to build autonomous sensors in a cubic millimeter volume, and the term Smart Dust[iii] was launched. Early motivation and concept development for Smart Dust was a result of a RAND workshop[iv] and two DARPA ISAT meetings[v].

With several wireless sensor network projects making progress, it became clear that one of the major roadblocks in sensor networks was power. This was one of the reasons why the most recent round of DARPA MEMS program funding was in the area of MEMS power generation, focussing mostly on the conversion of hydrocarbon fuels to electric power.

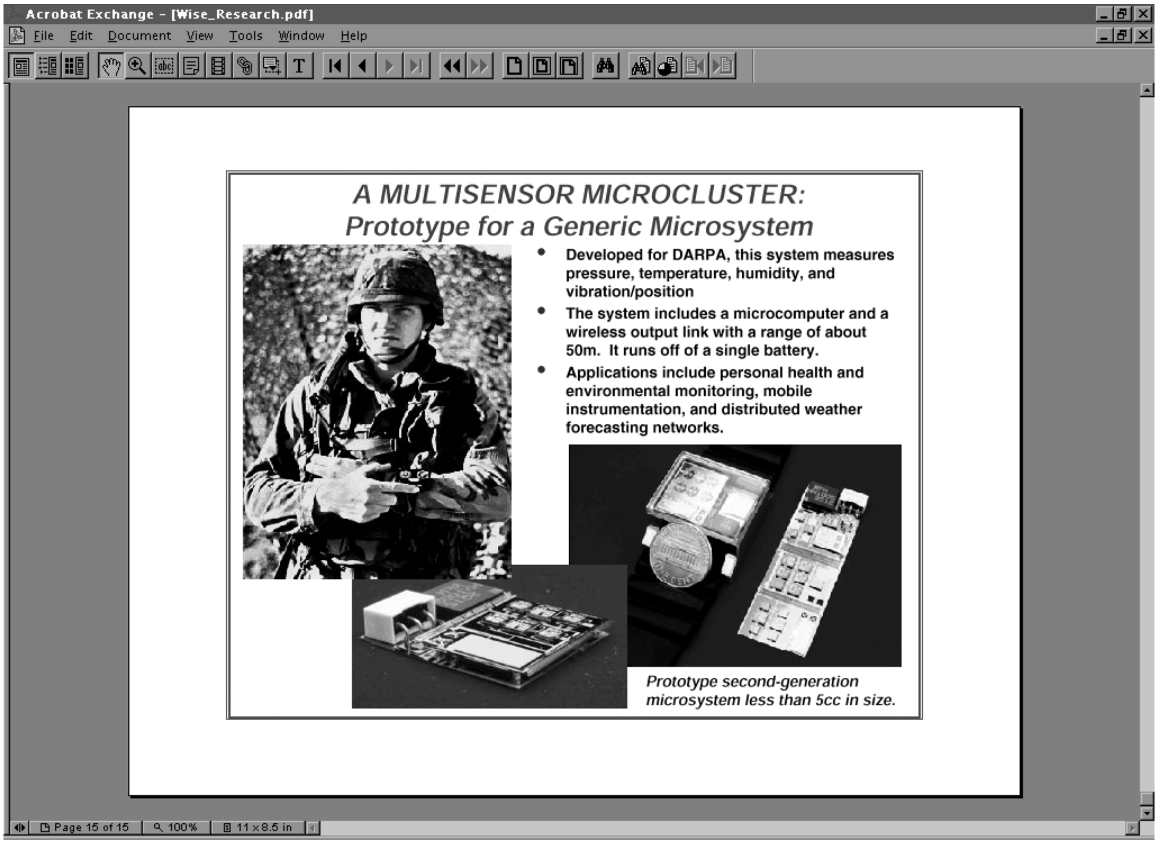

A MULTISENSOR MICROCLUSTER:
Prototype for a Generic Microsystem

- Developed for DARPA, this system measures pressure, temperature, humidity, and vibration/position
- The system includes a microcomputer and a wireless output link with a range of about 50m. It runs off of a single battery.
- Applications include personal health and environmental monitoring, mobile instrumentation, and distributed weather forecasting networks.

Prototype second-generation microsystem less than 5cc in size.

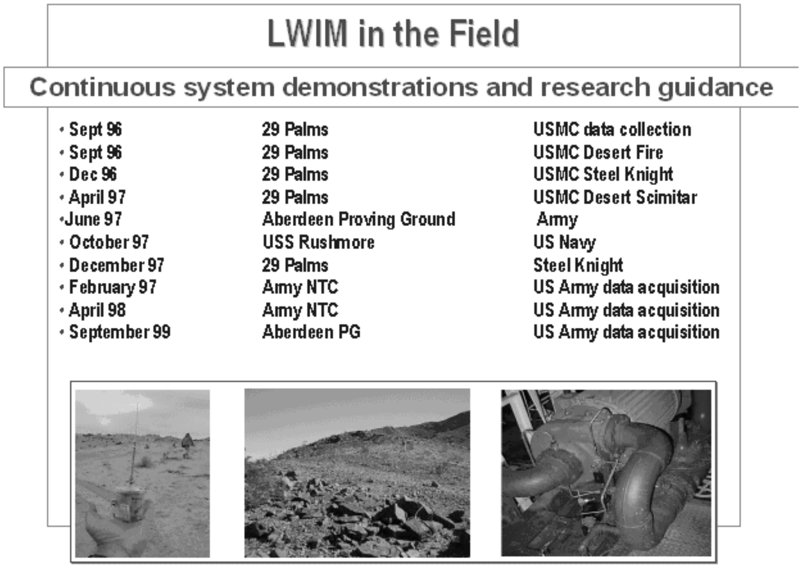

LWIM in the Field

Continuous system demonstrations and research guidance

• Sept 96	29 Palms	USMC data collection
• Sept 96	29 Palms	USMC Desert Fire
• Dec 96	29 Palms	USMC Steel Knight
• April 97	29 Palms	USMC Desert Scimitar
•June 97	Aberdeen Proving Ground	Army
• October 97	USS Rushmore	US Navy
• December 97	29 Palms	Steel Knight
• February 97	Army NTC	US Army data acquisition
• April 98	Army NTC	US Army data acquisition
• September 99	Aberdeen PG	US Army data acquisition

Figure 2 The UCLA LWIM project demonstrations.

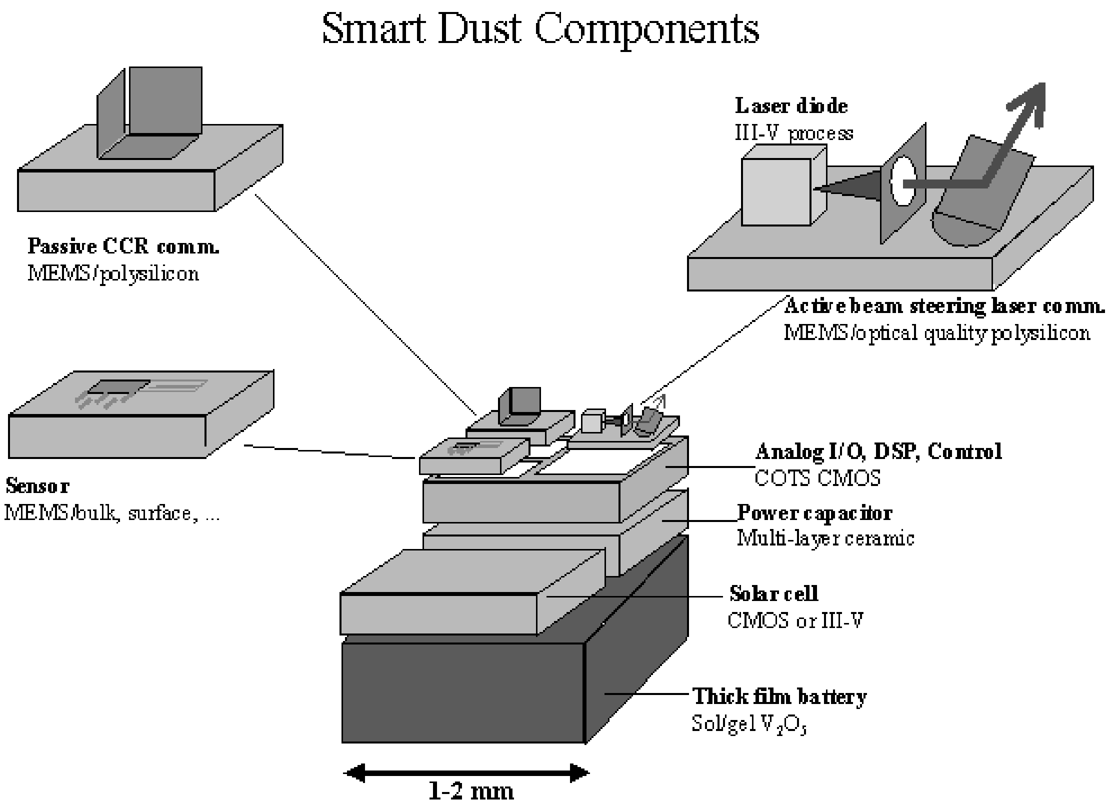

Figure 3 The UCB Smart Dust project goal.

The parameters of greatest interest in most wireless sensor networks are sensor performance, power, and cost. The issue of size is typically *not* a constraint for most applications once the move to MEMS sensors is accomplished. For most applications, the difference between a cubic inch sensor node and a cubic millimeter sensor node is relatively unimportant. However, size is indirectly important because of it's relationship to cost. If we assume an integrated solution to the autonomous sensor problem, then size and cost are mostly likely strongly correlated.

Sensor Performance

Sensor performance is often *inversely* related to size, despite what most MEMS researchers would like you to believe. Certainly the raw sensitivity of pressure sensors, accelerometers, gyroscopes, and microphones all degrade substantially as the size of the sensor decreases. On the plus side, frequency response of most sensors does improve with decreasing size.

The fundamental limit in most MEMS sensor systems is thermal noise. In a nutshell, the vibration of molecules (which is the very definition of temperature), causes all mechanical and electrical devices to jitter around as well, with an average kinetic energy of kT/2 (a few thousandths of a billionth of a billionth of a Joule). While your desk and the power coming out of the wall are relatively unaffected by this amount of energy, MEMS components (and the electronics that interface to them) are small enough that this amount of energy is important. In particular, the proof mass of a MEMS accelerometer is not much bigger than the pollen grains that Robert Brown saw through his microscope in 1827. What use is an accelerometer if random collisions with air molecules cause it to bounce around with Brownian motion? Not much. This then provides a lower limit on the size to which we can miniaturize our sensors - they must be either massive enough or stiff enough to not be unduly influenced by the air itself[1].

For example, a device like the ADXL202, a two-axis, +/-2g full-scale accelerometer is within spitting distance of the thermal limit to sensitivity. More performance can only be achieved in this device by either increasing the size of the proof mass, decreasing the bandwidth, or increasing the power dissipated in the excitation and sensing electronics. Similarly, for hearing aid microphones (both MEMS based and non-MEMS), the noise in the microphone signal is not very much larger than the fundamental thermal limit, the amount of noise caused by the thermal vibration of the microphone membrane itself.

---

[1] One partial solution to this problem is to run the sensors in vacuum. This can lead to dramatic improvements in thermal noise performance for some sensors, but other damping mechanisms will still impose a thermal-noise limit.

Power consumption in sensors

Sensor excitation and sensor electronics power requirements are intimately related to the thermal noise in the sensor itself. This is one area where most MEMS products have not made much progress, because sensor power constraints from the system level are generally mild. One exception is the latest few accelerometer products from Analog Devices which burn dramatically less power than the original ones did, even though their performance is better. Typically, there are several orders of magnitude available between the hundreds of milliWatts currently used by most of these sensor systems and the theoretical limits of the sensor and electronics, which are typically in the microWatt to milliWatt range.

Power consumption in computation

Currently, power consumption in a power-optimized microprocessor[vi] is roughly 1nJ/instruction[2]. This corresponds to a general-purpose 32 bit microprocessor. For specific tasks, application specific integrated circuits (ASICs) typically outperform general purpose processors by a factor of 100 to 1000 in the area of power consumption, so we can look forward to power consumption in the 1-10 pJ/instruction with dedicated silicon.

Power consumption in RF communication

It is difficult to make generalizations about power consumption in communication systems, because there are so many variables that come into play in evaluating performance of these systems. However, the fundamental limits are again related to thermal noise. For a receiver with a noise bandwidth B (roughly the bit rate), the thermal noise power from the antenna is kTB. The quality of the electronics in the receiver determines how close the actual noise performance is to this theoretical limit, and is represented by the noise figure of the receiver, N_f. N_f is the ratio of the actual noise to the thermal limit. The strength of the radio signal received needs to be greater than the noise by an amount determined by the down-stream signal processing of the signal, and is given by SNR_{min}. This means that overall, the signal power received by the antenna must be greater than $kTB\ N_f\ SNR_{min}$

---

[2] The unit nJ/inst is actually the reciprocal of MIPS/mW. The StrongARM family of 32bit microprocessors, for example, achieve more than two hundred MIPS when burning two hundred mW.

To put some numbers to this, consider the GSM cellular phone standard. The noise bandwdith is roughly 200kHz for a 115kbps link. The receiver has about 8 times more noise than the thermal limit, and the downstream electronics needs a signal to noise ratio of about 10 to achieve an adequately low bit error rate. In decibels relative to 1 milliWatt (dBm), that gives a sensitivity of:

-174 + 53 + 9 + 10 = -102 dBm

or just under one tenth of a picoWatt! So why does the cell phone need to transmit with several Watts of RF power if it only needs to receive 0.1pW of power to get the signal? The answer is in the path-loss. With line-of-sight and no nearby objects, the power lost between the transmitter and the receiver is proportional to the square of the distance in wavelengths. Near the ground, however, reflections off of the ground, buildings, trees, etc., cause attenuation to be proportional to the fourth power of distance (see Figure 4). For a 1GHz signal the wavelength is roughly 30cm, so transmitting a distance of 300m gives an attenuation of roughly 12 orders of magnitude! Traveling 3km gives an attenuation of 16 orders of magnitude, taking a 1Watt transmitted signal down to 0.1pW.

The power burned by the GSM receiver is roughly 200mW, and the power burned in the transmitter is roughly 4W. Given the data rate of 115kbps, this works out to a cost of around 2microJoules to receive a bit, and 50uJ to send one.

Cordless phones operate with similar data rates at less than one tenth the power, but with a range reduced to 10-100 meters. On the order of 1 uJ/bit is common.

The Bluetooth radio[vii] is designed for short range, 1Mbps communication in a household or office environment. Transmit power is 1mW, but the total radio power is still roughly one hundred mW regardless of transmit power, because of all of the radio circuit overhead. Even so, the Bluetooth standard is still the most promising for civilian sensor networks with short-range communication cost of roughly100 nJ/bit in the 2.4GHz band.

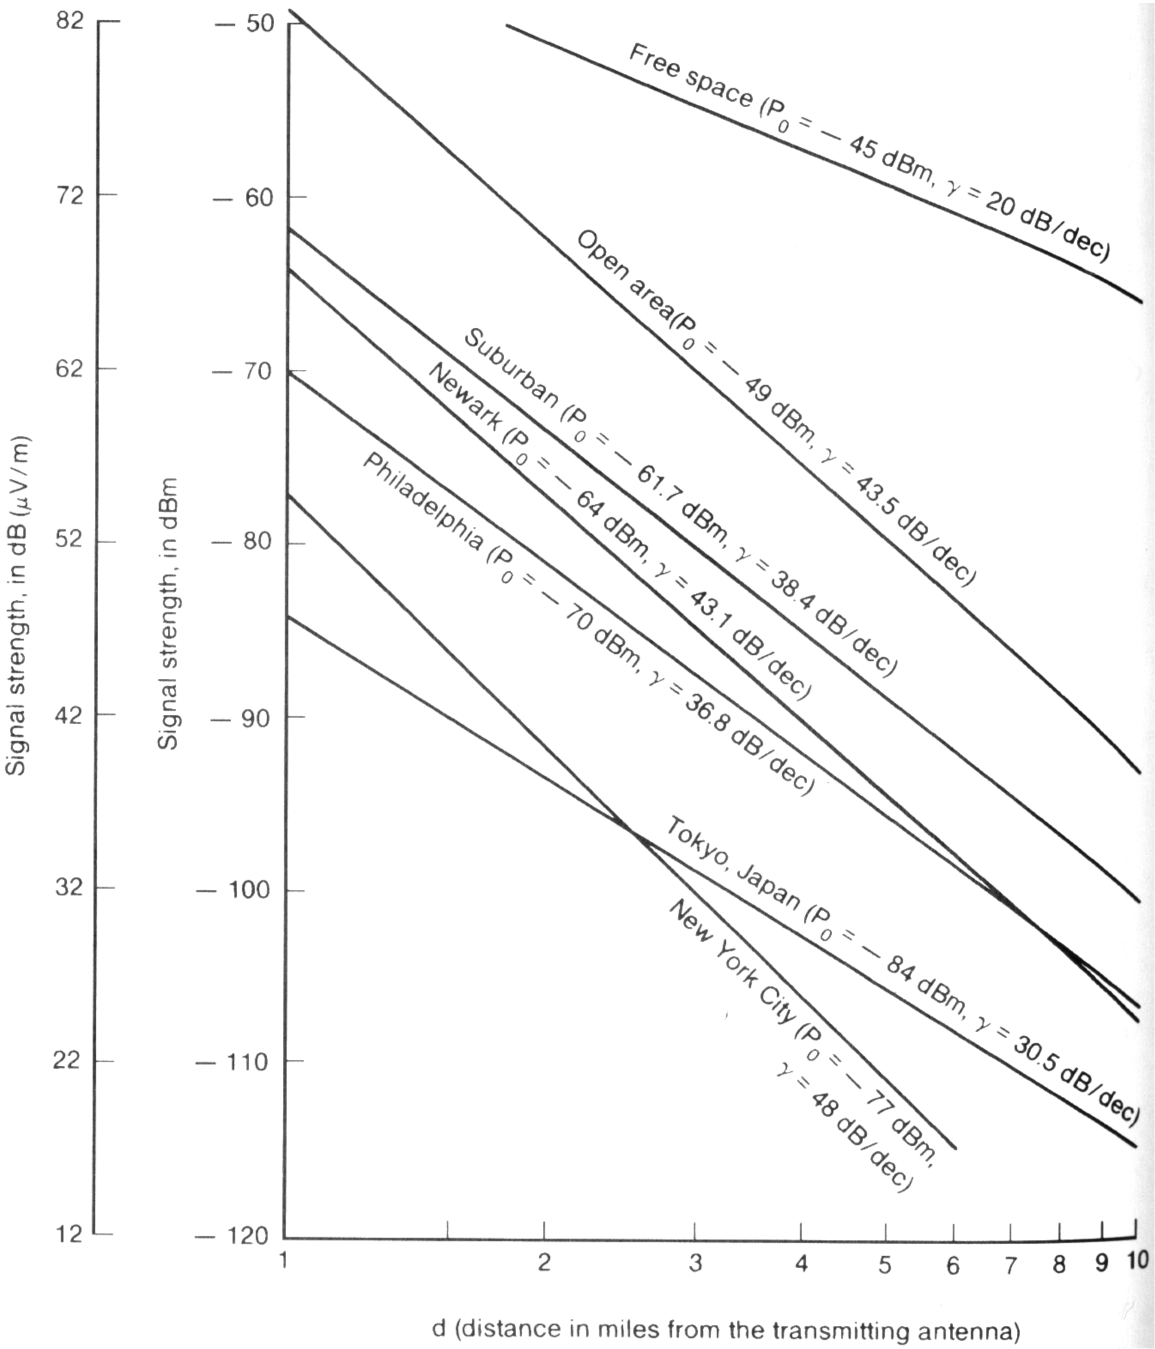

Figure 4 Signal strength vs. distance for cellular telephone. Po is the signal strength (power received, in dB relative to 1mW) at 1mile, and gamma is the attenuation exponent, giving power loss in dB per decade of increased frequency. Note that in all environments, attenuation goes as roughly the fourth power of distance (gamma=40). (From Lee [viii])

Power Generation and Storage

The most likely and simplest type of power storage for wireless sensor nodes is lithium batteries. The latest generation of lithium batteries is rechargeable, and roughly 300Watt-Hr/kg, or 2,000 J/cc. This means that you can run your 4W laptop PentiumX for about 8 minutes off a 1 cc of battery. A power optimized sensor node with duty-cycled communication might consume an average of 100uW of power, which gives a lifetime of nearly a year per cc of battery.

The latest revolution in capacitive energy storage is the Ultracapacitor[ix] which provides an energy density of nearly 10 J/cc. While this is only 1% of the energy density of a good battery, the energy from these new capacitors can be delivered in a matter of seconds, whereas most batteries can not be discharged at such high rates.

For scavenging energy from the environment it is hard to beat solar radiation as a host source. Full sunlight gives around 1mW/mm^2 and bright indoor illumination is roughly one thousandth of that. Conversion efficiency is around 30% for the best cells. For applications where duty-cycling is acceptable, solar cells or other power scavenging sources can be used to trickle-charge a capacitor (or battery), and then the stored energy can be used at much higher power rates than the charging power.

Vibration has been proposed as a scavengable energy source. Indeed, vibration spectra of office windows, copy machines, and industrial motors reveal that there is useable energy here - typically on the order of ten microWatts per gram of mass of the converter.

Existing Products

The closest existing commercial products to wireless sensor networks are home security systems and RF ID tags. During the 1990s, the home security market underwent a revolution in which all of the wired sensor nodes (window vibration, door opening, IR motion sensors, fire sensors) were converted to wireless communication. While the technology to do this was available for decades, the cost and power requirements dropped dramatically in the 1990 time frame, and so it became economically attractive to spend more money on the hardware in order to avoid the installation cost of running wires in houses.

The RF ID tag and keyless-entry system markets have existed for at two least two decades, with the Texas Instruments TIRIS system as one of the long-time leaders. Both RF ID tags and keyless-entry systems are unpowered wireless devices which absorb energy from a local RF broadcast source. In the case of the keyless-entry systems, the reflected signal from the node gives the source sufficient information to determine it's ID.

The RF ID tags, on the other hand, actually store some of the RF energy, and then transmit their own information, typically at a different frequency. These systems usually consist of a silicon chip, an inductor/antenna, and two capacitors (one of which forms a resonance with the inductor for receiving, the other for transmitting). The volume for the tag is in the hundred cubic millimeter range. The latest generation of tags has read/writeable memory on board, and a communication range of many meters, assuming an interrogator antenna with dimensions comparable to a meter.

The first fully-integrated RF ID tag was recently announced by Hitachi[x]. This system consists of 1kb of on-board memory, on-chip coil antenna, bidirectional communication circuitry at 13.56MHz, and data rates of 26kbps, this is the first complete system on a chip. The chip measures 2.3x2.3 mm^2. Addition of a MEMS sensor to this chip would presumeably be a matter of some research, but not fundamentally difficult. The drawback of the on-chip antenna is that the communication range is currently limited to less than 3 mm. For smart-card and smart-token applications, where the device will be inserted into a slot, this range is perfectly acceptable.

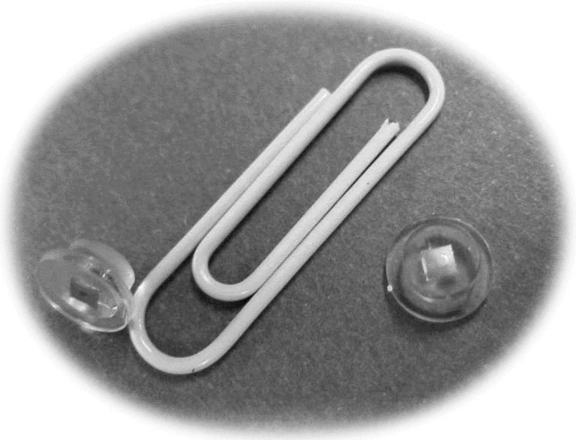

Figure 5 The Hitachi MY1007, an RF-ID system on a chip with no external components.

Cost

Integrated circuit fabrication costs are on the order of $0.05/mm^2, even for high end processes[xi]. The thousand dollars that you pay for the latest Pentium is not going to pay the cost of fabrication, which are roughly $10/chip. In large volume production, manufacturing costs for a small chip like the Hitachi RF chip above will quickly reach this nickel per square millimeter level, making that chip a $0.25/copy commodity.

MEMS fabrication costs are still substantially higher than IC fabrication costs, but even so the numbers are in the tens of pennies per square millimeter. The conventional wisdom in the MEMS industry is that the cost for a MEMS part is roughly equally divided between the chip fabrication, packaging, and calibration.

Technology Forecast

Over the next five to 10 years, I believe that the following will come to pass. All of these technologies are assumed to be a part of a sub-cubic-centimeter package unless otherwise stated.

Low power Inertial Measurement

Inertial measurement units with better than micro-g acceleration sensitivity and something close to earth-rate rotation sensitivity will become widely available. Power requirements per-axis will be below 100uW.

Microphones and Pressure sensors

There will be little change in microphone performance, size, or power. Existing hearing aid microphones are close to the engineering limits of performance. More conventional pressure sensors, however, will come to achieve similar dynamic range and power levels as hearing aid microphones: ~90dB and 10uW.

Biosensors

The ability to isolate and analyze chemical species, sequence DNA, and identify pathogens in a centimeter-scale laboratory is no-doubt covered in several other DSSG papers. This field is still in its infancy, but it is clear that revolutions will occur in the coming decade.

RF communication

RF communication will improve in three ways relative to wireless sensor networks. First, low data-rate, short range radios will be built in the 100uW power range[xii] [xiii] with communication capabilities similar to a cordless phone - 10-100kbps, 10-100meters of range.

Second, extremely high data rate burst transmission radios will be implemented in the 59-64GHz oxygen-absorption band. Because these wavelengths are absorbed by oxygen, communication range is limited to under a kilometer. This is ideal for many types of distributed sensor networks. With a 5mm wavelength, MEMS relays and other tricks will be used to get reasonable antenna gain and directional communication.

These radios will require at least tens of milliWatts (possibly Watts) while on, but will have data rates above 1 Gbps.

Finally, MEMS resonators will find their way into low-power MEMS radios. Current radio designs function only because of the extremely high quality mechanical filters (SAW, FBAR, etc) which allow frequency selection to take place. MEMS filters are now starting to enter the frequency bands of interest for RF communication, for both RF and IF filtering. Integrating these filters on-chip will reduce size, power, and cost over existing radios. In addition, the power in mixers, PLLs, and other front-end radio components is inversely related to the quality of tuned LC oscillators, [xiv]. If these tuned oscillators can be implemented with high-Q, on-chip MEMS resonators, radio receiver power will drop dramatically, possibly into the microWatt range, and efficient transmitters will be possible with outputs of only tens of microWatts. Of course, the range of these radios will be quite limited, by still should be in the tens to hundreds of meters.

Optical Line of Sight Communication

Early demonstrations from the Smart Dust program make it clear that communication across tens of kilometers with less than a milliWatt from a cubic millimeter package will be achieved within the next 18 months. In fact, communication to aircraft or even satellites in low earth orbit will be possible from devices in the cubic centimeter to cubic millimeter range. 1 Gbps optical communication links are currently one of the primary focii of the DARPA/MTO Steered Agile Beams program, indicating that prototypes of these systems will be available within the next three years.

Power Generation, Storage, and Scavenging

Power generation in the tens of Watt range from cubic centimeter generators burning hydrocarbons will be demonstrated in the next few years. At present, there does not seem to be an approach to this problem that will provide low power levels for extended periods of time, but perhaps an integrated solution where a high-power hydrocarbon-burner will charge batteries will evolve.

Radioactive power sources have not received much attention as yet in the low-power community, although they are quite standard in interplanetary systems. These systems have the potential to provide tens of milliWatts of electrical power per gram for a decade or more. Application of MEMS to this area is likely to yield significant results, although environmental concerns over making many small "nuclear" devices will no doubt need to be though through.

I do not foresee any dramatic changes in energy storage technologies in the next ten years. Batteries and capacitors are close to their theoretical limits.

Integrated energy scavenging systems will be developed across the cubic millimeter to cubic centimeter size scale, combining solar cells and/or vibrational motion converters with batteries and CMOS control electronics to provide a clean interface to power using systems.

Delivery

Emplacement of sensor networks of this scale provides new options. A slingshot, for example, is sufficient to populate a good fraction of a square kilometer with marble-sized sensor nodes in a matter of minutes. More traditional delivery systems such as modified grenade or mortar shells seems a likely candidate for distributing hundreds to thousands of sensors over many square kilometers in a very short time period. There is ongoing work in the DoD MEMS community on making MEMS survive the shock associated with being fired from a gun.

Micro air vehicles (MAVs) provide one of the most intriguing delivery options. Existing MAVs are capable of carrying ten to tens of grams of payload, corresponding to perhaps a dozen to several thousand sensor nodes. The plane shown in Figure 6 is capable of 18 minutes of flight at 60 mph. A version of this aircraft has a magazine of 10 hearing aid batteries which can be ejected under radio control, simulating the deployment of a sensor network. Most versions of the aircraft include a color video camera with live video transmission. The plane contains a MEMS gyro to stabilize the roll axis, and is easy to fly through the video image. At 60 mph, this aircraft could dispense one Smart Dust -sized sensor every second for 1000 seconds, covering a square kilometer area with a sensor every 30 m.

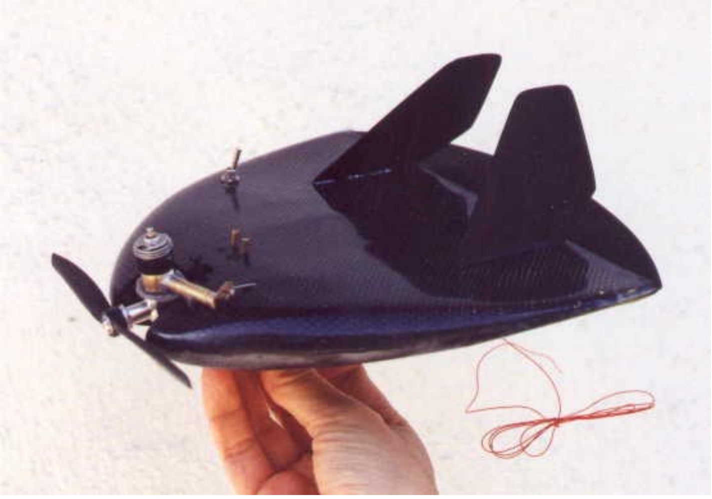

Figure 6 MLB Trochoid[xv]. 60mph top speed in level flight, gyro stabilized roll axis, 18 minute range, color video with 2.4GHz downlink, 1km communication range.

Cost

As the level of integration increases in MEMS sensor network components, the cost of the sensing, communication, and control will approach the nickel per square millimeter IC limit. Packaging for these systems will require some minor revolutions to avoid being the limiting factor in cost. Calibration will become dramatically less expensive due to the ability to communicate with the sensor system without having to mount it in a rack, and the increased level of on-board diagnostics and intelligence.

By 2010, a multi-sensor system with months to decades of life, wireless communication, and sub-cubic-centimeter volume will cost less than a dollar in large volume. Some very specialized systems with limited performance will be manufactured for under ten cents.

Applications

Here I present only three of many ideas that were generated in discussions with military personnel and DSSG mentors. Bunker mapping was chosen because it seems to address one of the "open questions" that we were given during our tours. Intrusion detection was chosen because it seems like the most immediately easy and useful demonstration of this kind of network. Stockpile Stewardship was chosen because this is an application that most people will not have thought of before.

Bunker mapping

Scenario: An underground facility is being constructed or has been constructed. The geometry of the facility is unknown, in terms of size, depth, and shape. Vehicles enter and exit the facility on a fairly regular basis.

Goal: determine the geometry of the underground facility.

Approach: Attach a small inertial measurement device to a vehicle before it enters the facility. After the vehicle leaves the facility, download the sensor data, and use it to reconstruct part of the internal structure of the facility. With multiple data sets, a comprehensive map of the internals of the facility will be constructed.

The sensor data would be downloaded by either RF or line of sight optical communication to some local retransmitter.

The mote would be placed on the vehicle by one of several methods:

- hand emplaced. This provides the best chance for hiding the mote, and guaranteeing any alignment that may be necessary. Most risky.
- ballistically delivered. The mote could be fired from a gun of some kind.
- perched MAV drop. Fly in an MAV and perch on a tree, bridge, building, or other structure that overlooks a road into the facility. This approach has the advantage that the MAV can be used as the relay station for communication after the mote has returned out of the facility.
- flying MAV drop. The MAV swoops down on the truck and delivers the mote ballistically. This method requires some level of skill in piloting, and risks discovery.

Feasibility: Assume that we have a three axis accelerometer, three axis gyro, and three axis magnetometer on the mote. Currently this combination of sensors, together with a microprocessor, bi-directional RF communication (50m), and power supply can be built in a volume of less than one cubic inch. This is already probably small enough to be used under some circumstances. Power requirements even with existing off-the-shelf components give a lifetime of days to months depending on duty cycle.

In addition to the mapping activity, the sensor nodes could be augmented with a variety of other sensors. The most useful of these might be an image sensor. Reasonable quality digital image sensor with wide angle lenses are commercially available in a few cubic centimeter volume. These could be reduced in size somewhat, but the limits imposed by optics are on the order of a few millimeters. Reasonable quality images require roughly 10kB of storage, so hundreds of images could be stored in a few MB of flash memory. Images could be programmed to be acquired on a regular timed basis, or under the control of the inertial measurement unit (e.g. every time the vehicle stops, or turns, or travels a given distance), or some combination of both. Most likely a CMOS imager would be needed, rather than a CCD image sensor, for reasons of power consumption. In addition, integration of the image compression circuitry with the imager would make for a small, lower power system. Finally, whether a custom CMOS imager or a COTS imager is used, it would be possible to use image data to augment or potentially even replace the inertial navigation data.

Dynamically Placed Intrusion Sensor Networks

Scenario: Military units clearing urban terrain must clear a building, but can not afford to leave people behind to ensure that it stays cleared.

Goal: Notify the force if anyone enters the cleared portion of the building after they have left.

Approach: Soldiers would carry something like a Pez dispenser[xvi], possibly attached to their weapons, filled with sensor nodes that could be shot or emplaced quickly by hand on a wall, stairwell, or doorway. The sensors, using some combination of acoustic, IR, visual, or vibrational cues would pass information about intruders to the appropriate person/people.

This scenario was suggested by Col. Henry Kinnison with some input by Chris Kearns[xvii]. In particular, Col. Kinnison suggested that the soldier could speak a message as he emplaced the sensor node, and that that would be the verbal message that would be relayed to the soldiers when that sensor detected an intruder. The message would typically be descriptive of where the sensor was put, e.g. "third level broken window", and would use descriptions that were relevant during the actual maneuver, rather than what might have been discussed during planning.

Feasibility: This could be done today with off the shelf components in the cubic inch size range. To be militarily useful it would certainly require substantial modification (for size, ruggedness, security of communication, etc), but Kearns and his group at the Dismounted Battlespace Battle Lab at Ft. Benning are ready to try out the off-the-shelf version as soon as someone makes it.

Stockpile Stewardship

During the DSSG visit to Los Alamos it became clear that there were many interesting applications of wireless sensor networks in Stockpile Stewardship. Many classified discussions were had with people[xviii] involved in different parts of weapons design, storage, re-manufacturing, etc. Unfortunately, virtually any interesting information about this topic is classified. The following has been cleared:

> Characterize pressure, temperature, and materials inside complicated engineered devices with as little collateral interaction as possible. This would involve insertion of sensors via hypodermic needles and catheters. Applications for Stockpile Surveillance.
>
> Characterize pressure, temperature, resistivity, and velocity differences between dissimilar materials and components subject to high accelerations and velocities. Since these ultra small sensors would be non-intrusive, they would prove useful in quantifying weapon environments.
>
> Chronic sensors for stockpile weapons. This would involve the development of ultra small feedback controllers, temperature monitors, and other devices useful for characterizing aging effects in the stockpile. General areas such as delaminating, elasticity changes, and chemical releases could be quantified using these types of devices.

This is a very promising area for the application of MEMS techniques in general, and wireless sensor networks in particular. Contacts[xix] at Sandia National Laboratory indicated that there is a large effort ongoing in this area already at SNL. While it is certainly true that SNL is doing great things in MEMS, it is also certainly true that no one at LANL had any collaboration with them in the areas discussed above. It is also certainly true that the academic community is completely unaware of the potentially great benefit of MEMS in these stockpile-related applications. While there are obviously issues of secrecy to be dealt with, this message needs to get to the university community.

Bibliography

---

[i] Michigan microsensor microcluster: http://www.engr.umich.edu/~wise

[ii] LWIM/WINS/AWAIRS projects: http://www.ee.ucla.edu/~kaiser

[iii] http://www.eecs.berkeley.edu/~pister/SmartDust

[iv] "Future Technology Driven Revolutions in Military Conflict", RAND 1992, and a follow-on study by Steeb and Brendley, "Military Applications of MEMS", RAND, 1993?

[v] DARPA/ISAT web page: http://atlas.ida.org:8500/

[vi] http://www.strongarm.com Power optimized microprocessor family with Pentium-quality speed and less than 10% the power dissipation.

[vii] http://www.bluetooth.com Local area radio.

[viii] William C.Y. Lee, Mobile Cellular Communications, 2nd. Ed., McGraw-Hill, 1995.

[ix] http://www.powercache.com A 4 Farad, 2.3 V, 2A capacitor the size of three stacked nickels.

[x] http://www.maxell.co.jp/eng/noflash/business.html Hitachi RF ID tag system on a chip.

[xi] http://www.tsmc.com Taiwan Semiconductor Manufacturing Company, a high-quality-CMOS foundry.

[xii] http://bwrc.eecs.berkeley.edu/Research/Pico_Radio

[xiii] Razieh Rofougaran Tsung-Hsien Lin, Fredric Newberg, and William J. Kaiser, *CMOS Front-End LNA-Mixer for Micropower RF Wireless Systems*, ISLPED99.

[xiv] Tsung-Hsien Lin and William J. Kaiser, *A 900MHz, 2.5mA CMOS PLL with Automatic SC Tuning Loop,* CICC2000.

[xv] http://spyplanes.com MLB Co. web site. MLB sells world-record holding hang gliders, as well as a growing family of sub-meter scale aircraft.

[xvi] http://www.pez.com

[xvii] Christopher Kearns, kearnsc@benning.army.mil, US Army Dismounted Battlespace Battle Lab, Chief, Dismounted Forces Div. Col. Henry Kinnison, KinnisonH@benning.army.mil .

[xviii] A good place to start for future discussions would be with George Hrbek, hrbek@lanl.gov . Also very interested is mbaron@lanl.gov .

[xix] Two good people to talk to are Mark Rosen, marosen@sandia.gov, and Kent Meeks, kdmeeks@sandia.gov .