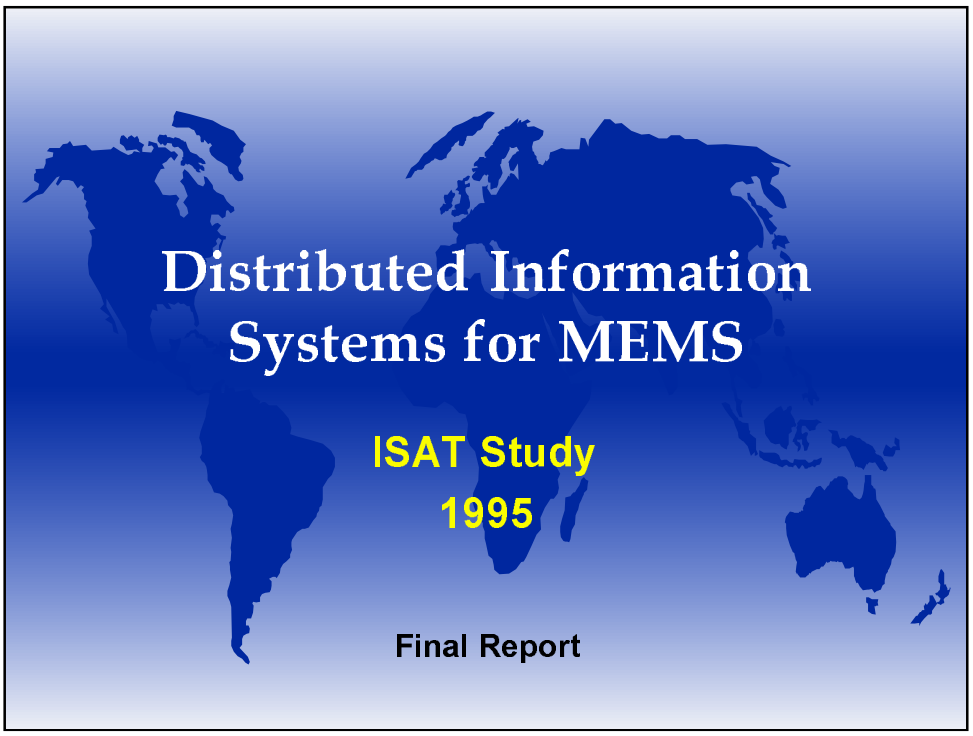

Mission Statement

The purpose of this study is to examine the long-term benefits of using distributed MEMS technology to couple computation to the physical world, and to recommend research in information technology required to achieve these benefits.

Study authors and their affiliations:

Andrew A. Berlin (Chair)	Xerox Palo Alto Research Center
Harold Abelson	Massachusetts Institute of Technology
Natalya Cohen	Massachusetts Institute of Technology
Larry Fogel	Natural Selection, Inc.
Chih-Ming Ho	UCLA
Mark Horowitz	Stanford University
Jonathan How	Stanford University
Thomas F. Knight	Massachusetts Institute of Technology
Richard Newton	University of California at Berkeley
Kris Pister	UCLA

Study Participants

Distributed Systems:
Hal Abelson (MIT)
Bernardo A. Huberman (Xerox)
Larry Fogel (NSI)
Natalya Cohen (MIT)

MEMS:
Chih-Ming Ho (UCLA)
Kris Pister (UCLA)
Andrew A. Berlin (Xerox)
Peter Will (USC/ISI)

Computer Architecture:
Tom Knight (MIT)
Mark Horowitz (Stanford)

Microelectronics/CAD:
Richard Newton (UCB)
Patrick Scaglia (Cadence)

Aeronautics/Astronautics (Active Control):
Jonathan How (Stanford)

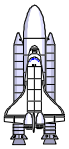

Invited Speakers

- Bernardo A. Huberman
  - » Distributed Computation
- Ken Yip
  - » Computer Vision Applied to Analysis of Fluid Flow Data
- Antony Jameson
  - » Issues in Fluid Dynamics and Aircraft Design
- Kris Pister
  - » MEMS and Microrobots
- Chih-Ming Ho
  - » Turbulent boundary layer control
- Peter Will
  - » Active Surfaces for Manufacturing
- Steve Jacobsen
  - » MEMS in Robotics
- Mark Horowitz
  - » Future Directions in Computing
- Jonathan How
  - » The Middeck Active Control Experiment

For further information about this study, please contact:
Dr. Andrew A. Berlin
Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304
E-mail: berlin@parc.xerox.com
Phone: (415) 812-4372

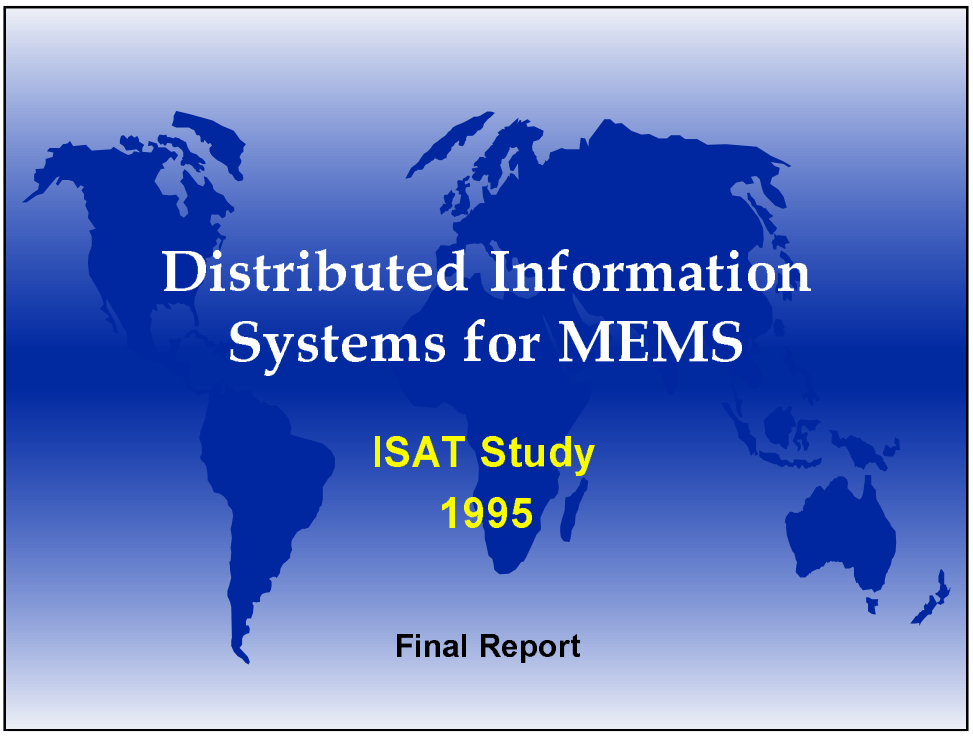

Mission Statement

The purpose of this study is to examine the long-term benefits of using distributed MEMS technology to couple computation to the physical world, and to recommend research in information technology required to achieve these benefits.

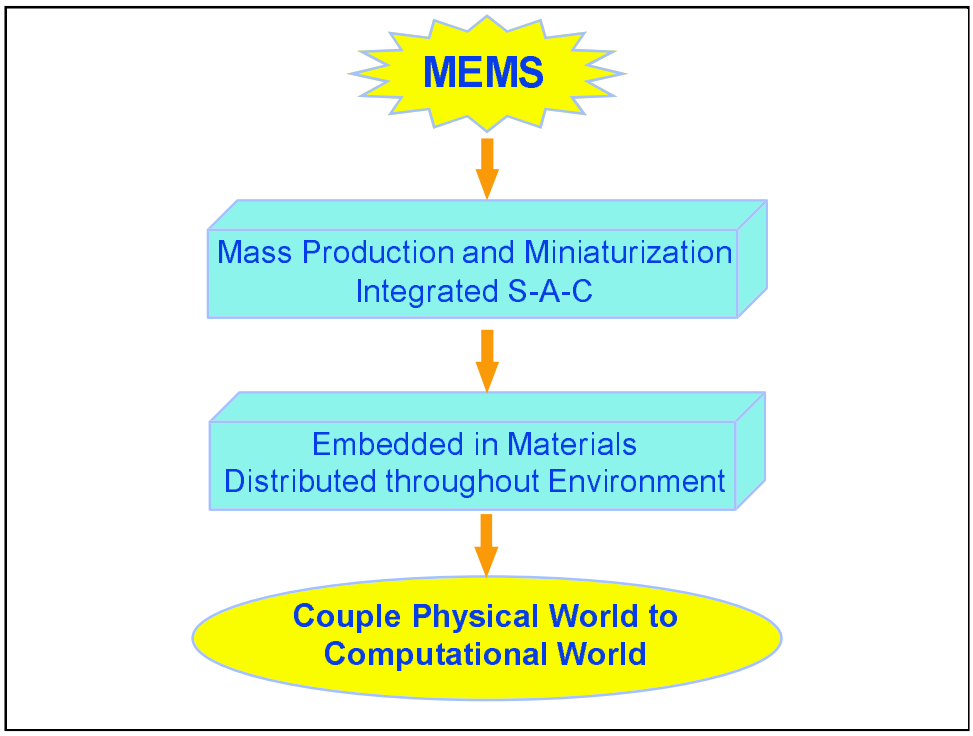

We view MEMS as a set of technologies that make it possible to miniaturize and mass-produce integrated sensors, actuators, and computers. We use the term MEMS broadly to include a range of technologies that are required to construct these integrated devices, including micromachining as well as things like thin-film battery technology and smart materials.

In this study, we are looking 5-10 years down the road to a time when this mass production and miniaturization will make it possible to embed large numbers of MEMS devices within materials, to coat them on top of surfaces, and to distribute them throughout the environment, effectively coupling computation to the physical world on a scale that has never been possible before. The purpose of this study is to explore what some of the benefits of this coupling are, and to recommend research in information technology required to exploit those benefits.

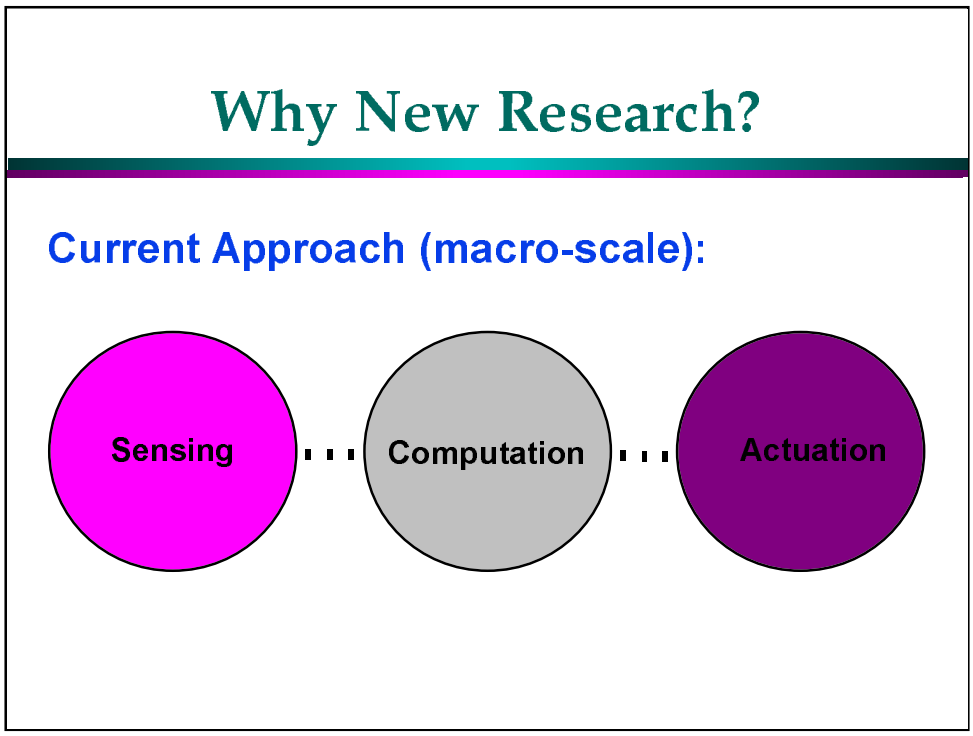

For the past 20 years there has been an ongoing trend towards increased reliance on active control technology, enabled by the electronics revolution that allowed computers to become small enough to be embedded within everyday products. Today's active products typically consist of a handful of discrete sensors and discrete actuators that are physically wired together to some control electronics.

A rich variety of computational tools and control strategies have been developed for these systems. However, for the most part efforts on active control have focused on systems with many degrees of mechanical freedom, rather than on systems with many different sensing and actuation sites.

Distributed MEMS-Based Approach

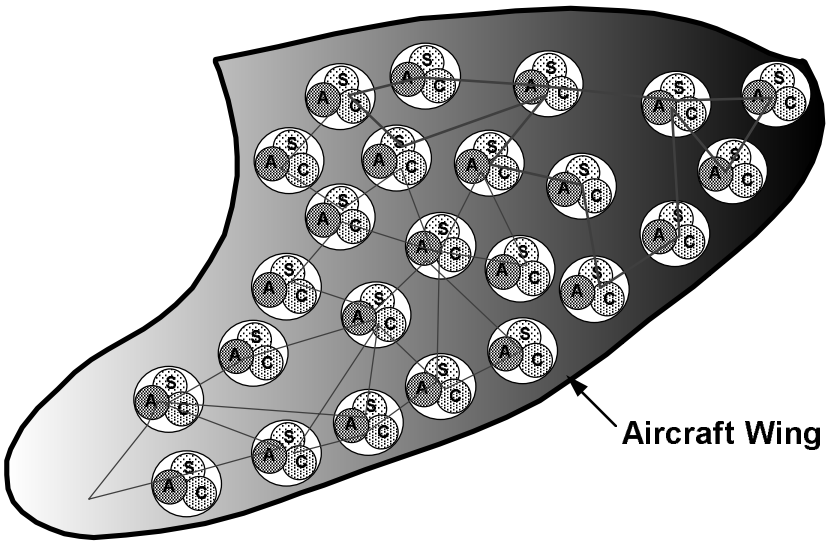

So when we start thinking about controlling systems that consist of large numbers of sensing and actuation sites distributed over a wide area, we find a new set of challenges for information technology having to do with how we are going to control and coordinate these large numbers of devices to achieve global goals.

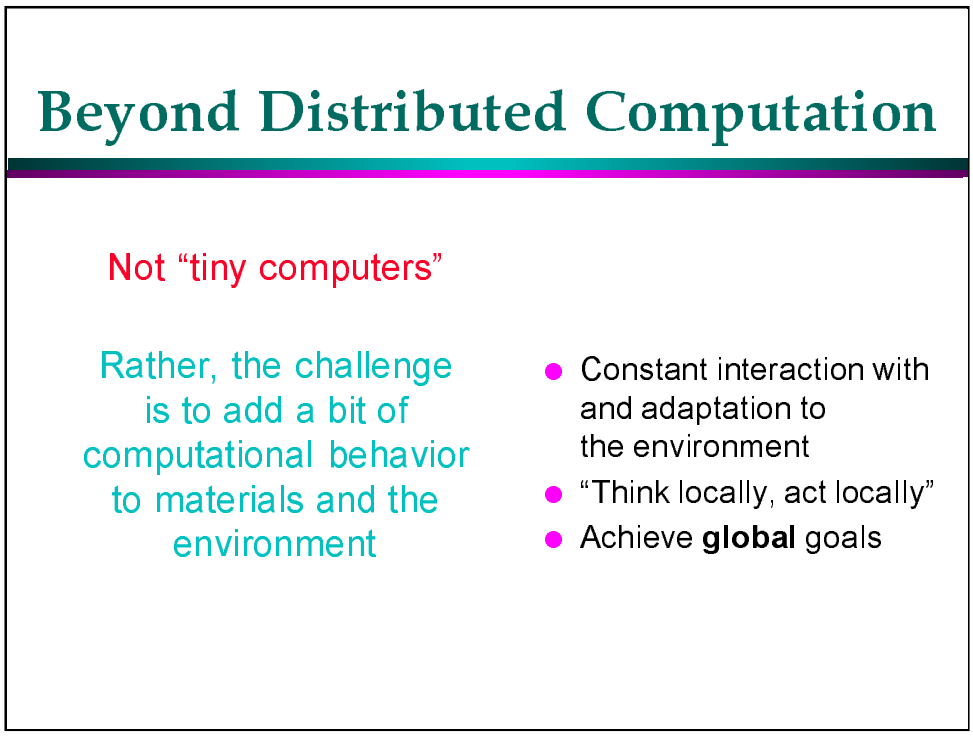

What is it about information technology for MEMS that's so different? After all, we have been working on the problem of coordinating large numbers of devices for a long time - in research on distributed computation. One of the things that makes it different is the coupling to the physical world, with constant interaction and adaptation to the environment. For instance, if we are going to couple computation to the physical world, then whether we like it or not, the differential equations that govern the physical world are going to become part of our computer programs. So the challenge that MEMS places on information technology is not merely to coordinate lots of tiny computers, but rather to add a bit of computational behavior to materials and the environment.

Challenges for Information Technology I

If computation is part of materials, what is the materials science?

Even extremely simple physical properties on the micro-scale can lead to extremely complex global properties

How does a rope "compute" a catenary?

$F = k\,x$ $\qquad$ $F = k_1\,x + k_2\,x^3$

For example, consider even a very simple mechanical system such as a rope. Ropes are very simply described at the micro-scale as a linear spring, **F = k x.** We know from differential equations how to map from that local description to the overall shape of the catenary that a rope takes on under the influence of gravity. MEMS lets us modify the local properties of a system. The tricky part is that if we make even a very slight change to the local behavior, such as changing from being a linear spring to a stiffening spring (essentially adding a cubic term), then it's called **Duffing's equation**, and there are people who spend their whole lives studying it, trying to understand how to map from this simple local property to the global shape.

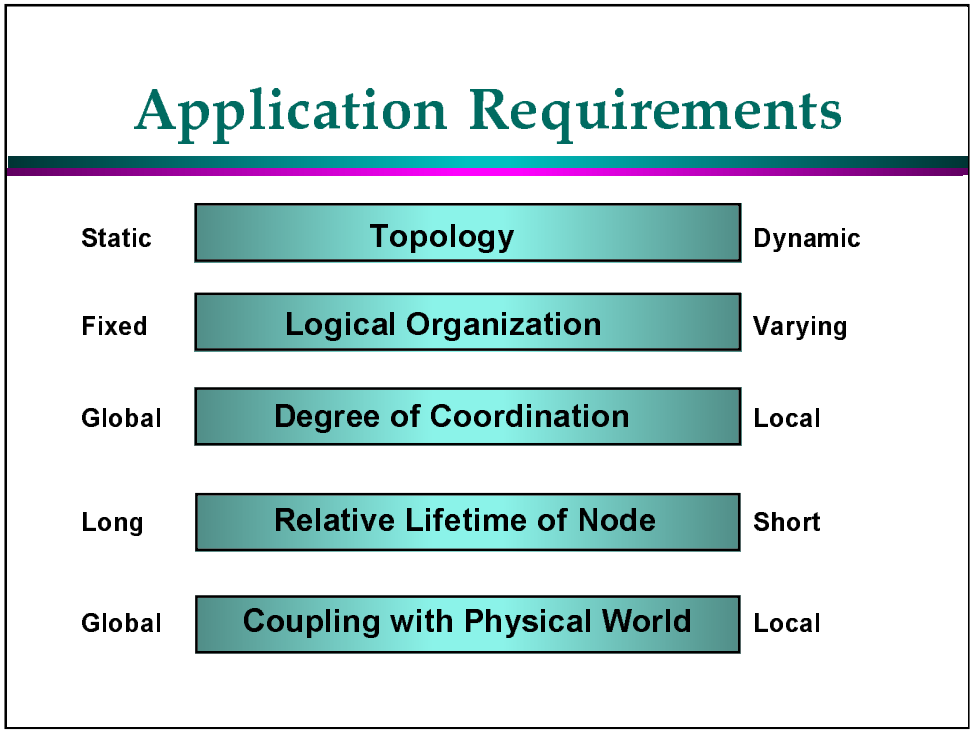

Coupling with the physical world is only one of the challenges for information technology - each MEMS application poses its own unique set of information technology requirements. We looked at a variety of applications and found that these five factors were the ones that most influenced the demands placed on information technology. For example, in the case of airborne surveillance particles, part of the challenge is a time-varying spatial configuration (e.g. dynamic topology). For active aerodynamic surfaces, control is complicated by the dynamics of the fluid being manipulated, and by the need for actuators to take on different roles depending on the location of vortices (e.g. varying logical organization and coupling with the physical world).

Challenges for Information Technology II

Coordinate the behavior of large numbers of unreliable parts with:

- Time-varying spatial configurations
- Environmental interactions
- Dynamic allocation of precious resources
  - Energy, one-shot sensors, communication, …
- Synchronization of distributed actions

Although MEMS applications have diverse information technology requirements, all of the applications that we studied had one feature in common: the need to coordinate the actions of large numbers of spatially distributed unreliable parts to achieve global goals.

Recommendations

1. Develop information technology for several different classes of applications.
2. Search for and exploit leverage points.
3. Develop new mathematics and programming paradigms.
4. Steps to take now to enable research.
5. Develop principles of control-centric design.

Recommendation I

Develop information technology
for several different
classes of applications

Different Classes of Applications

Smart Particles

Active Surfaces

Intelligent Structures

In examining a variety of applications in this study, we found 3 distinct classes of applications that place very different demands on information technology.

In this report, we present sample applications from each of these classes.

The first application class presented is **smart particles.**

On-chip Laser and Fold-up Fresnel Lens

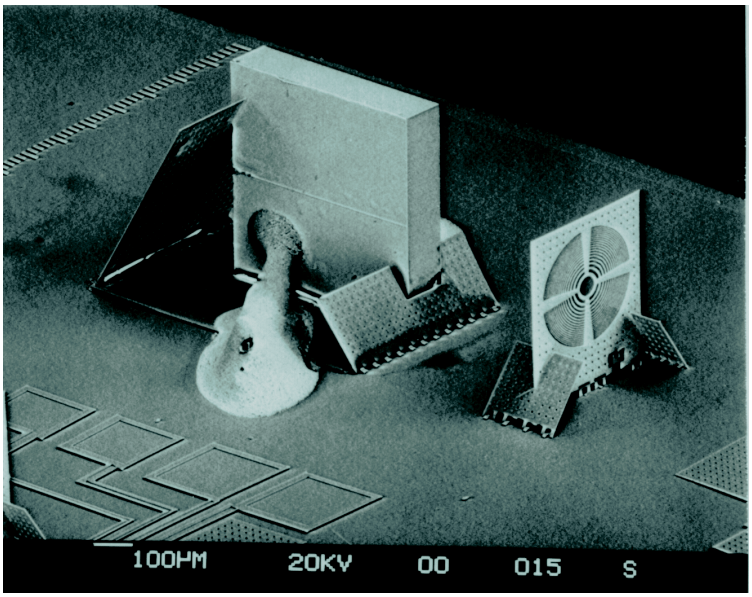

To give an idea for what is technologically possible today, here is a photograph from UCLA of a MEMS laser with optics suitable for transmitting light off-chip.

Photo: Lih-Yuan Lin and Shi-Sheng Lee, APL Vol. 66, p. 2946.

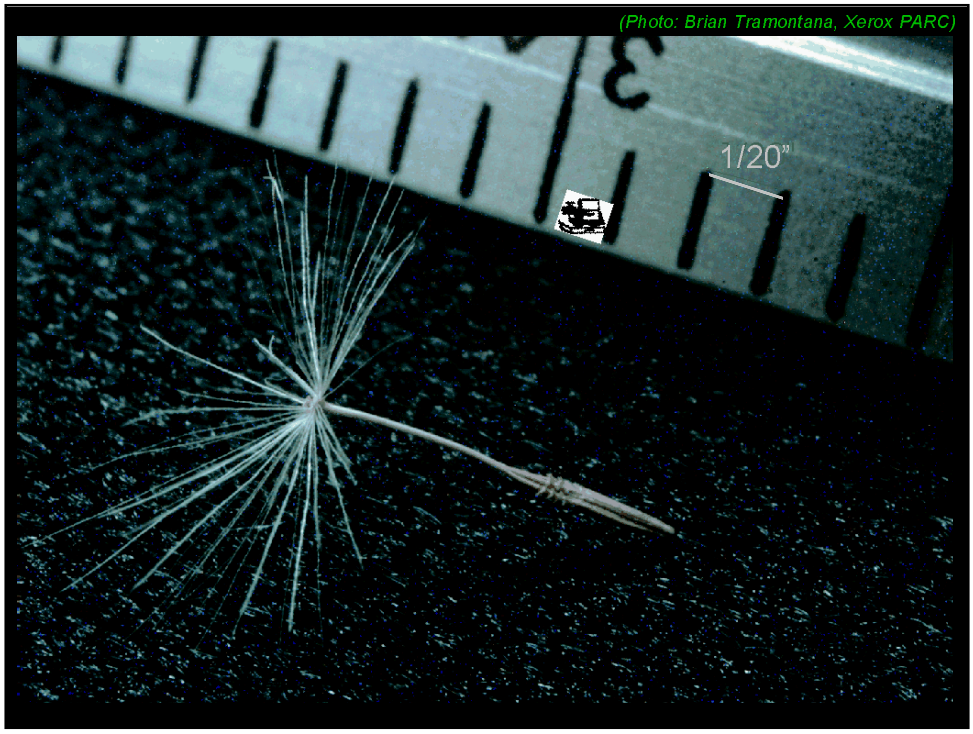

This photograph shows a single dandelion spore - something so small and light that it literally floats around in the air.

To give a feeling for size range of MEMS devices, the MEMS-based laser and lens system is shown here on the same scale as the dandelion spore.

Surveillance Dust

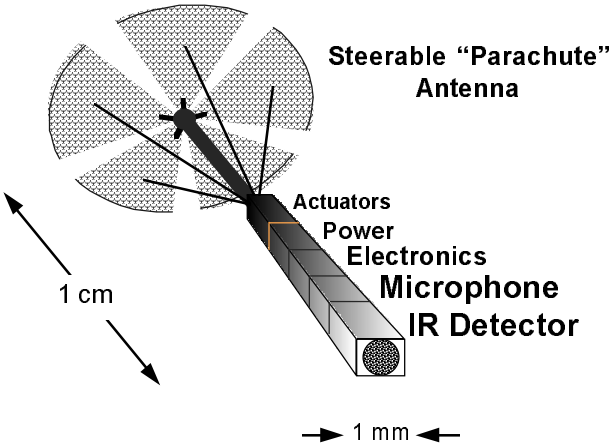

Continuing the analogy to dandelions, we envision airborne surveillance dust that looks something like this. The inside of the tube would be filled with lithium batteries, and MEMS devices would be mounted on the flat surfaces.

We have investigated technology trends and expect that it will be technologically feasible to construct devices such as this in the 5-10 year time frame.

Dust Applications

- Military:
  - » Biochem surveillance
  - » Perimeter security
  - » Access to denied areas
  - » Deception
    - – Emit decoy signals
  - » What's over the hill?
  - » "MEMS on point"
- Industrial:
  - » Precision farming
  - » Hazardous waste monitoring
  - » Smart ID tags
  - » Novel manufacturing techniques

Potential military applications of MEMS dust include tracking biochemical agents, monitoring atmospheric conditions (primarily pressure and wind velocity) for targeting artillery, and dust "bullets" that disperse large numbers of devices over an area to provide perimeter security and track troop movements (i.e. "What's over the hill?").

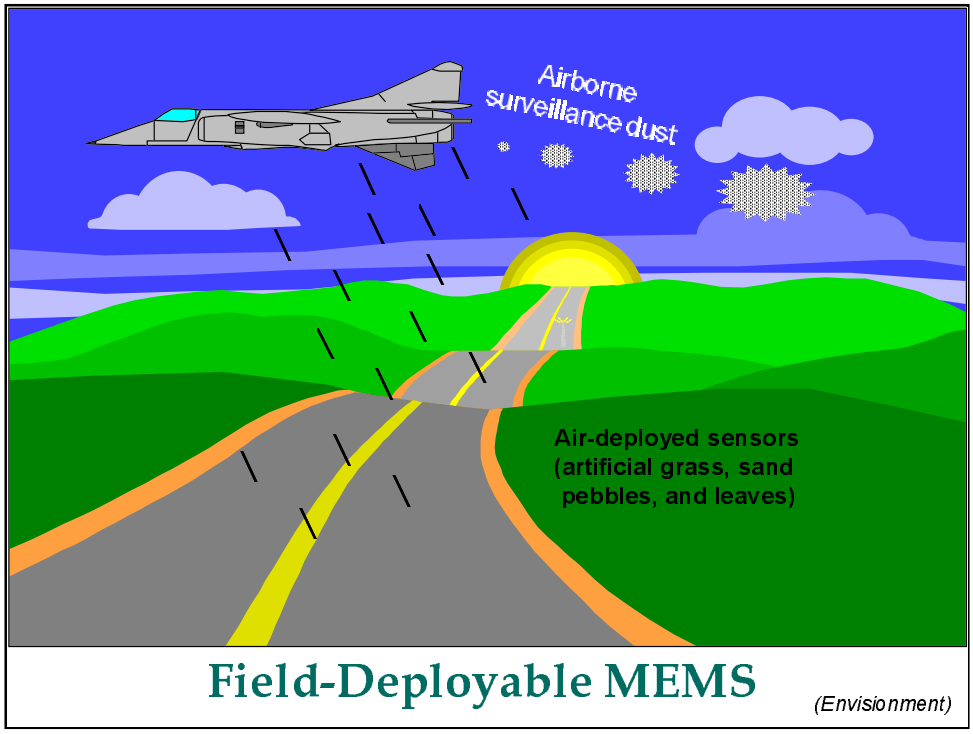

Field-Deployable MEMS *(Envisionment)*

We contacted several members of the military services who stated very clearly that small disposable surveillance devices whose focus of attention is under the control of small groups of soldiers "would save lives".[1]

In operations such as the relief mission to Somalia, civilians sneaking up on sleeping soldiers to steal equipment was a major problem. Small surveillance devices capable of providing platoon or battalion-level security to warn soldiers if ground-traffic or biochem agents are in an area would "let soldiers sleep"[1], resulting in a significant improvement in readiness.

[1]Conversation with C. Keers and N. Marano, U.S. Marine Corps. Background information provided by Dr. Bill Howard, Defense Airborne Reconnaissance Office.

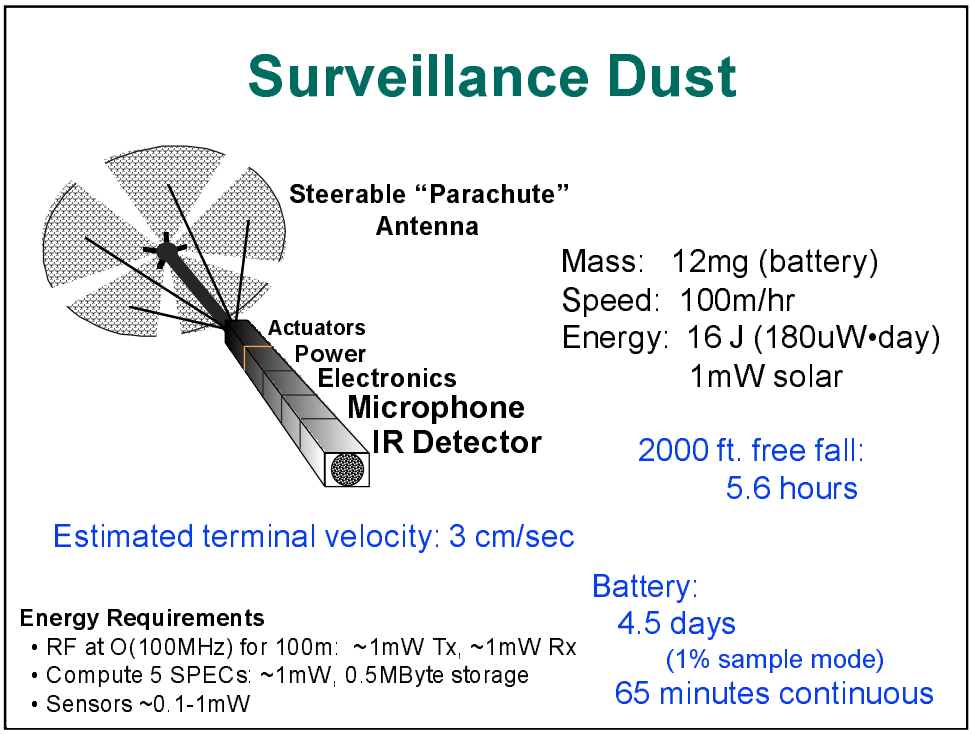

Based on our projections of technology trends we expect that a MEMS dust particle would fall at about 3 cm/sec, so that a particle released at 2000 feet would stay aloft for about **5 1/2** hours (possibly longer if it rides thermal currents).

Even without recharging by solar cells, in continuous operation of sensing, computing, and transmitting data, this device would last for about an hour. This continuous operating mode is for use when the particle finds itself in the vicinity of an interesting event. On the other hand, in sample mode, where it just wakes up roughly 1% of the time and takes a look around to see if anything interesting is happening, a MEMS dust particle would have enough power to operate for more than 4 days.

Dust Networking

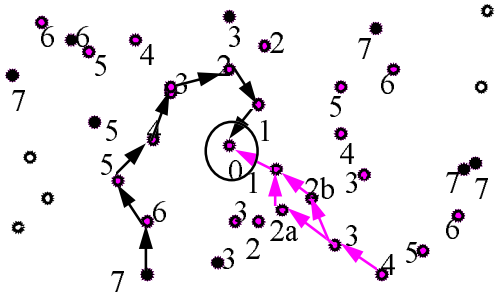

- Establish local communication links
- Establish relative (absolute) positions (with a reference)
- Establish a common timebase ('synchronize')

- Periodically verify the state of the network
- Identify and validate sensed events and communicate the required information to the outside world
- Many of the problems may already be solved in other contexts (e.g. packet radio)

When we first started thinking about how to program MEMS dust, we expected networking and communications issues to be among the most significant challenges to be overcome. We found that several of the *low-level* issues, such as how to establish local communication links and derive relative positions, have already been addressed in other contexts such as packet radio.