




Wireless Sensor Networks: The Planet's New Skin

Leonidas Guibas


Computer Science Department
Stanford University



Wireless Sensor Networks

- Distributed systems consisting of small nodes, capable of sensing, processing, and wireless communication
- Often nodes operate untethered, under serious power constraints
- Communication is far more expensive than computation
- Applications in scientific, industrial, commercial, and military settings

notes



RFID



DOT



MICA



MS Spot Watch




PDA



Sensiora Node

Wireless Sensor/Embedded Processor Trends

- Of 9.6 billion μ P's to be shipped in 2005, 98% will be embedded processors!



- Riding on Moore's law, smart sensor nodes get:

More powerful



Sensoria WINSNG 2.0

CPU: 300 MIPS
 0.1 GFLOP FPU
 2MB Flash
 2MB RAM
 Sensors: external

Easy to use



HP iPAQ w/802.11

CPU: 240 MIPS
 32MB Flash
 64MB RAM
 Both integrated and off-board sensors

Inexpensive & simple



Crossbow MICA mote

4 MIPS CPU (integer only)
 8KB Flash
 512B RAM
 Sensors: on board stack

Supercheap & tiny

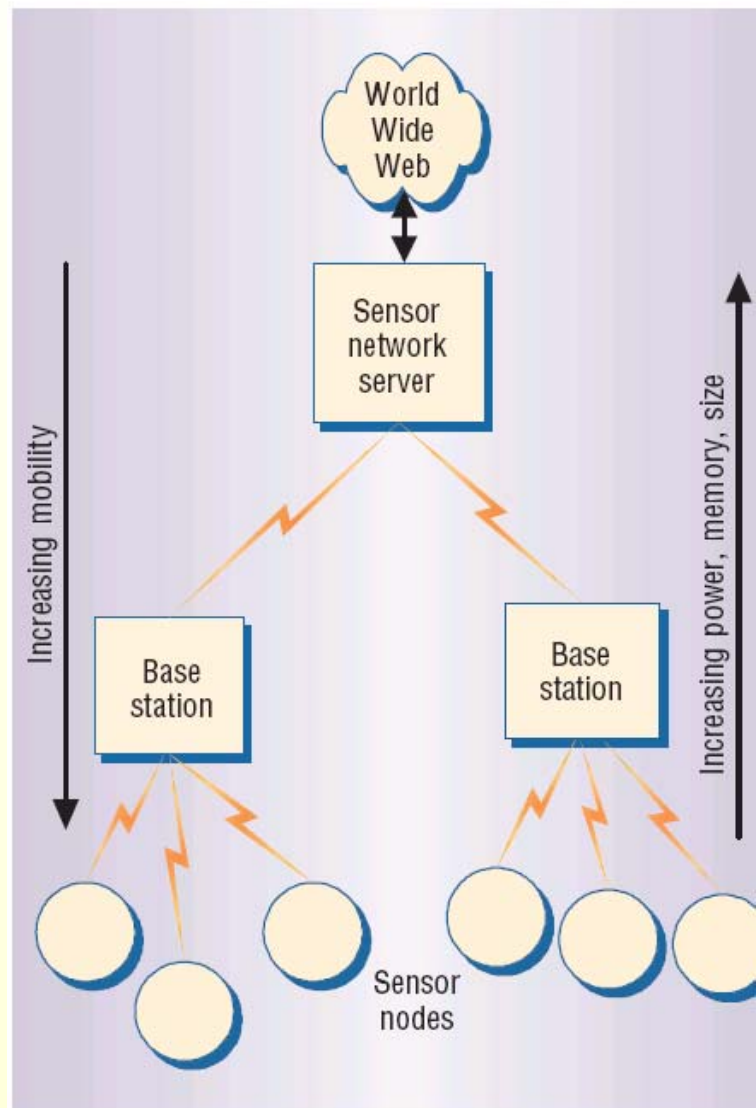


Smart Dust (in progress)

CPU, Memory: TBD (LESS!)
 Sensors: integrated

Monitoring the World

- Monitoring the **environment** and other spaces
- Monitoring **objects**
- Monitoring **interactions** between objects, or between objects and their environment



Sensor Network Deployment

Advantages:


- sensors can be close to signal sources, yielding high SNR
- monitoring of phenomena widely distributed across space and time
- scalable, robust and self-repairing systems
- significant savings on cabling, etc

British Columbia winery with networked temperature sensors



Other monitoring: temperature in data centers, oil tanker vibrations, soil contaminants, etc.

A New Scientific Instrument



- A sensor network can sense the environment and provide information that was impossible to obtain before
- At the same time, a sensor network operates under many processing, memory, communication, and power constraints



biology


astronomy



untethered micro sensors will go anywhere and measure anything -- traffic flow, water level, number of people walking by, temperature. This is developing into something like **a nervous system for the earth.** -- first Stormer in *Business Week*, 8/23-30, 1999.

More Demanding Sensor Network Applications

- Beyond simple data collection and aggregation
 - simultaneous tracking of multiple objects
 - distributed, wide-area phenomena
 - many detractors
 - integration of geographically scattered evidence
 - distributed attention: focus and context
- Network must adapt to highly dynamic foci of activity
- Resources must be apportioned between detection, tracking, etc.



The Challenge of Distributed Sensing

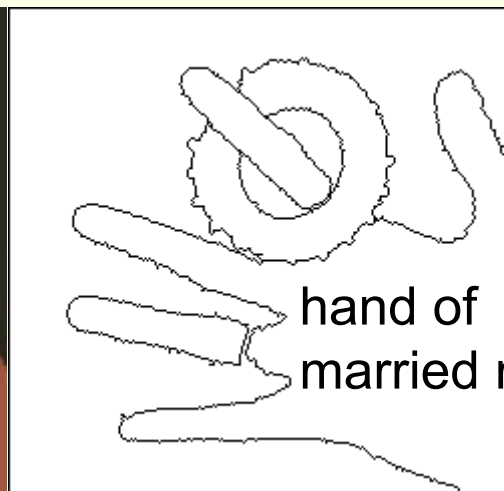
- Data is not all available at once or at the same place
- Data has to be collected only as needed – information acquisition has a cost
- Data may contain errors



sensor data




processing: segmentation



recognition

The Challenge of Networking Sensor Networks

- Network support for **a small number of collaborative tasks**.
- **Data-centric**, (as opposed to a node-centric) view of the world.
- **Monitoring processes may migrate from node to node**, as the phenomena of interest move or evolve.
- Communication flow and structure is dictated by the **geography of signal landscapes and the overall network task**.

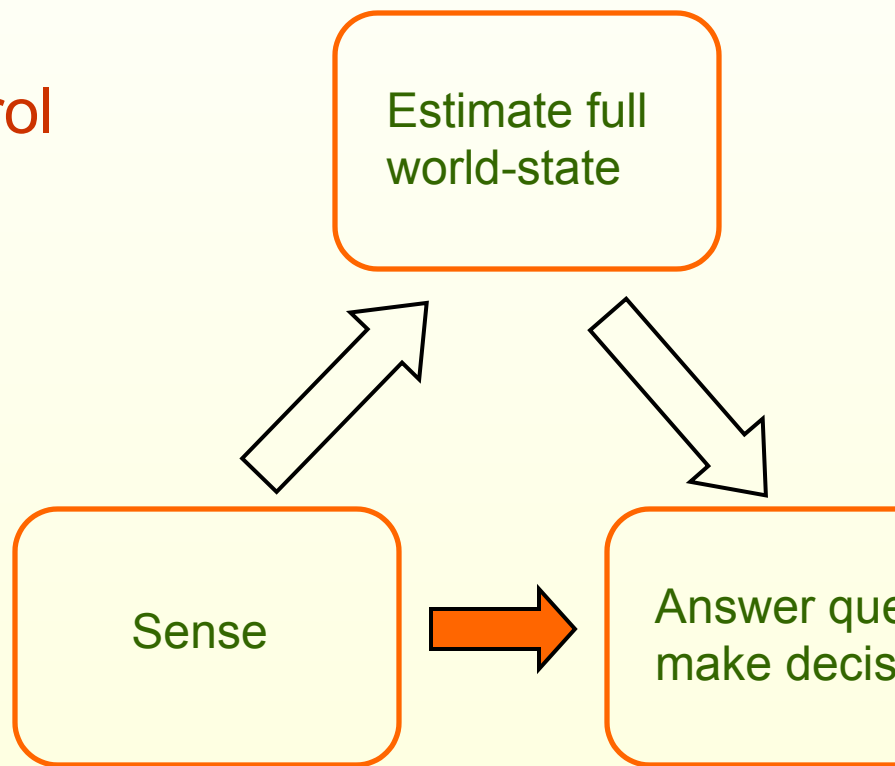


A Conundrum of Disciplines

- Low-Power Hardware Design
- Signal Processing
- Distributed Compression
- Statistical Learning and Estimation
- All layers of Networking
- Distributed Algorithms
- Reasoning under Uncertainty
- Computational Geometry and Optimization
- Data-Bases and Distributed Data Storage
- New Programming Metaphors

Sensor Network Challenges

- power awareness
- sensor tasking and control
- information brokerage
- formation of sensor collaboration groups
- in-network, distributed processing
- node management, service establishment, software layers
- coping with noise and uncertainty in the environment



A key algorithmic problem is how to sense and aggregate only the portions of the world-state relevant to the task at hand, in a lightweight, energy-efficient manner.

Image Sensor Networks

- CMOS technology enables the production of small, low-cost and low-power integrated image sensors
- Cameras (still or video) and other image sensors are becoming cheaper, smaller, and nearly ubiquitous
- However, truly distributed networked systems of image sensors are still not here



Additional Distributed Imager Challenges

- Imagers are **high data rate** sensors; therefore data must be compressed and summarized
 - compression must take into account shared data
 - goal of compression need not be reconstruction
- Vision algorithms can be **expensive** to run on low-power devices
- Visibility is **non-local** and **discontinuous** (occlusions, etc)
- Issues of privacy, etc.



The Project

- Use a camera network to obtain information about **space occupancy by people**.
- Useful for aggregate tracking, counting, etc.
- Crowd density implies multiple occlusions – **no one camera by itself can do this**.
- No image reconstruction --just **high-level distributed spatial reasoning**.



Packard 013



The Current Lab

Web cameras:

- 16 firewire webcams with 49 degree FOV
- Placed around a 22 x 19 foot room

Linux computers

- Each PC is connected to 2 webcams
- A separate process is running for each webcam to simulate an individual camera node
- All processes can communicate with each other over the network



System Architecture




Autonomous
Background
Subtraction,
Data Compression



Collaborative
Visual Hull
Estimation,
Camera Tasking



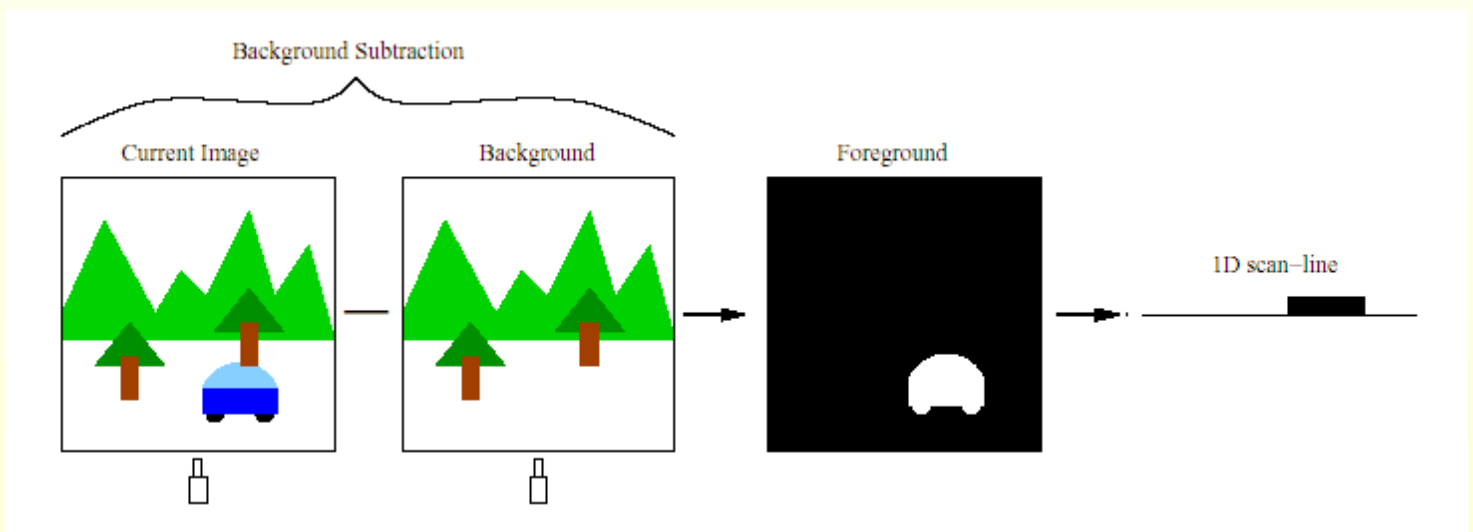
Problem
Solution




[Yang, Banos, G., ICCV'03]

Local Processing


- Perform background subtraction
- Collapse to a single scan-line
- 640x480 RGB image → 640 bit scan-line (which can be further compressed)



Occupancy Representation: The 2-D Visual Hull




a ——— 1D scan line



a ———
b ———
c ———


A Visual Hull Example

View of room
people




Scanlines
16 cameras

The Visual Hull Overestimates Occupancy




- Visual hull regions surround each object – but there is always wasted space
- Using more cameras reduces the overestimation – but it can never be fully eliminated

And Then There Are Phantom Regions ...







Visual hull regions
may also be empty;
we call these the
phantom regions



Motion can allow the pruning
of phantom regions


Pruning the Visual Hull

- Using more cameras reduces the overestimation – but it can never be fully eliminated
- Motion can allow the pruning of phantom regions





An Application: Counting People

- Given occupancy, bound the number of objects in each polygon of the visual hull
- The bounds over time can be used to constrain the count, using a tree data structure.





t





$t+1$

phantom



A Counting Example




[Video](#)

Application: Tracking




track.mov




Tasking Subsets of Cameras

To scale the system to large numbers of cameras, we must understand how to select appropriate subsets of the sensors to activate


- Using a small subset can save communication bandwidth and energy
- Often it can be done without significant estimation degradation



Optimal



Selection of 2 out of 3 cameras



Some Lessons and Issues

- A surprising amount of spatial information can be captured by cameras sharing very few bits.
- Small subsets of cameras, when appropriately tasked, can provide accurate estimates.
- The system can perform counting without tracking, thus raising no privacy issues.

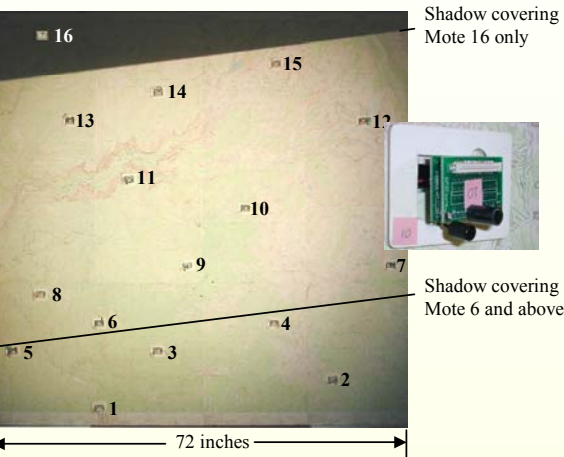
Sensor selection

Forming collaboration groups

Uncertainty and multiple hypotheses





Lightweight information integration

Other Work




Tracking large shadows over a sensor field

Forming collaboration groups and dealing with identity management



Routing using landmarks, without coordinates



Acknowledgements

Stanford Students: Qing Fang, Jie Gao, Jaewon Shin, Feng Xie, Danny Yang

External Collaborators: Patrick Cheung, Jie Liu, Julia Liu, Feng Zhao (PARC, Inc.), Hector Gonzales-Banos (Honda), Li Zhang (HP Labs), John Hershberger (Mentor Graphics)

Funding Sponsors: NSF, DARPA, ONR, Stanford NRC, Honda, Xerox



The End

Another kind of sensor network exhibiting lightweight distributed reasoning

