Information Processing in Sensor Networks

Feng Zhao
PARC (formerly Xerox PARC)
www.parc.com/zhao

parc
Palo Alto Research Center

Acknowledgement

- PARC CoSense team:
  - Patrick Cheung, Maurice Chu, Leo Guibas, Qingfeng Huang, Jie Liu, Julia Liu, Jim Reich
  - Student Interns: Jaewon Shin, Qing Fang, Judy Liebman, Elaine Cheong, Soham Mazumdar, Dragan Petrovic
- External collaborations:
  - Diffusion: Deborah Estrin, John Heidemann, Fabio Silva (UCLA/USC/ISI)
  - DOA Estimation: Kung Yao (UCLA)
  - TinyGALS: Berkeley
- Funded in part by the DARPA SensIT Program under contract F30602-00-C-0175

parc
Palo Alto Research Center

PARC Smart Matter Research

Smart Matter Research since 1994

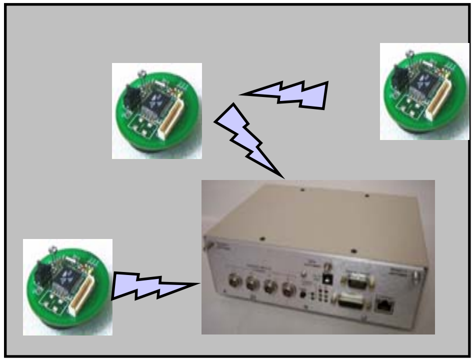

Sensor networks

- IDSQ/CADR
- Location services
- Group management
- TinyGALS/PIECES
- Distributed attention
- Applications

- Collaborative processing
- Embedded software design
- Scalable Information architecture

- Modular reconfiguration
- Constraint based control
- Adaptive optimization

- MEMS signal processing
- Energy harvesting
- Large-area sensor/actuator arrays

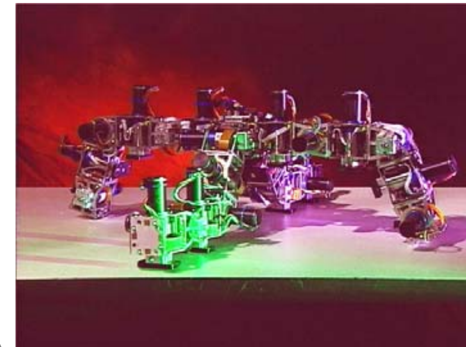

Actuator networks

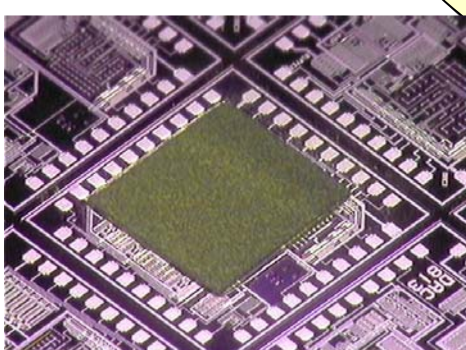

MEMS devices

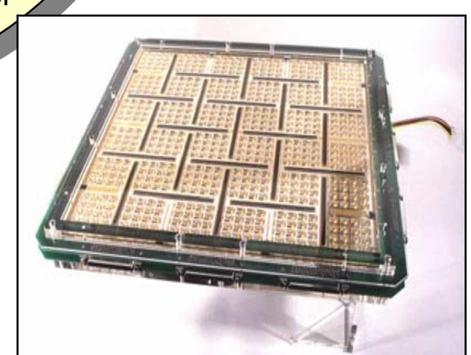

Smart, Networked Sensors

- Of 9.6 billion uP to be shipped in 2005, 98% will be embedded processors!
- Intel plans to put a radio on every uP

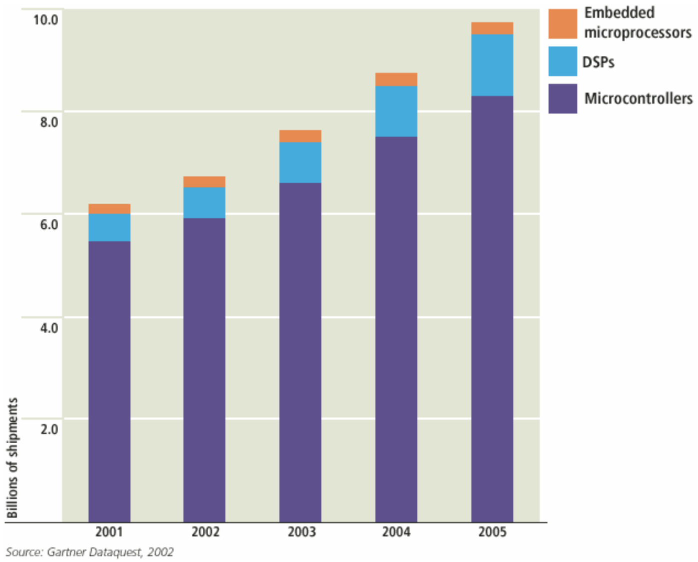

- Riding on Moore's law, smart sensors get

More powerful

Sensoria WINSNG 2.0
CPU: 300 MIPS
1.1 GFLOP FPU
32MB Flash
32MB 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

Super-cheap & tiny

Smart Dust (in progress)
CPU, Memory: TBD (LESS!)
Sensors: integrated

parc Palo Alto Research Center

Ubiquitous Sensing will Change the Way People Live, Work, and Play

Healthcare

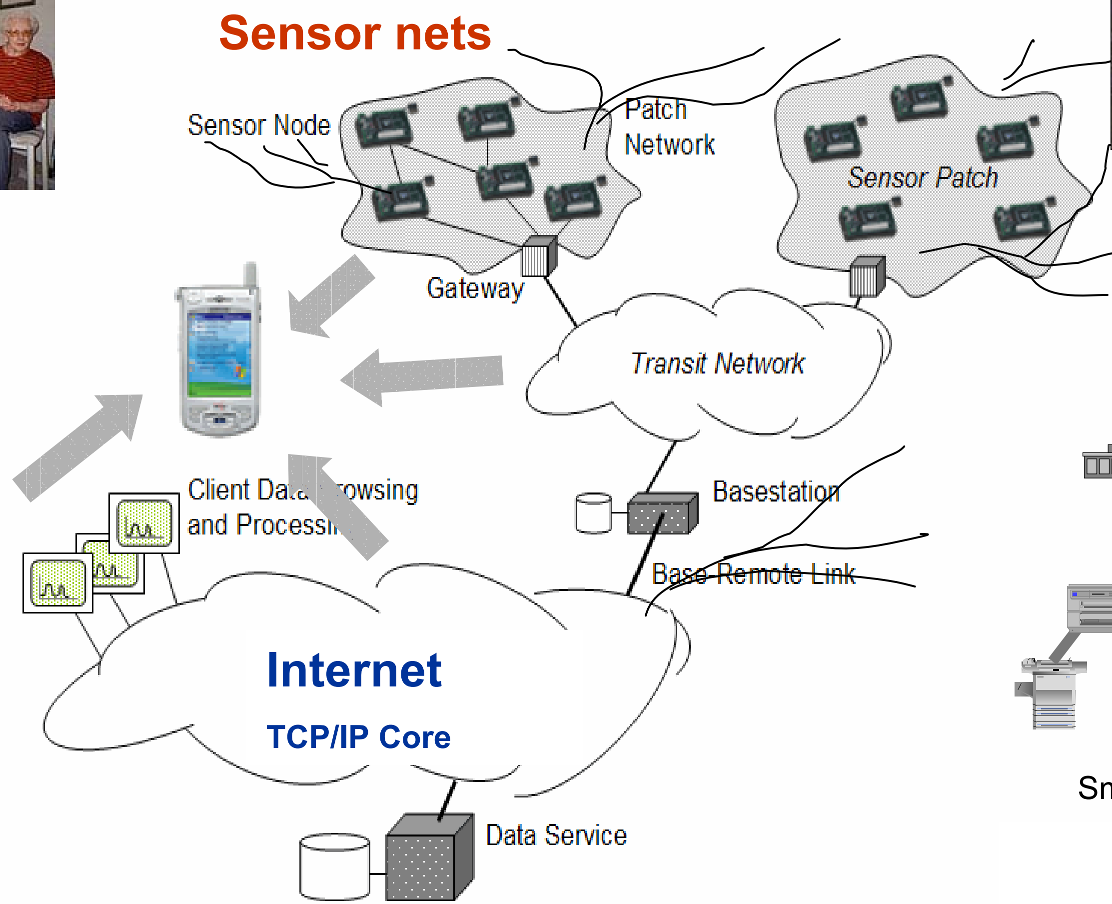

Networked transportation

Ubiquitous appliances

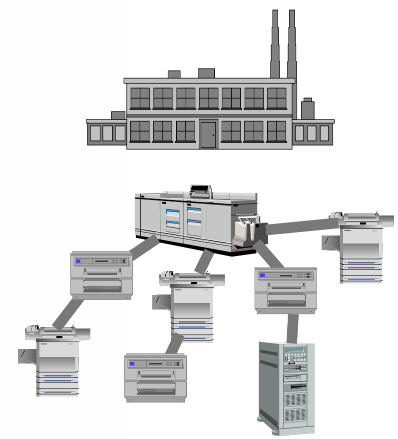

Smart Factories

"Networks of tiny sensors will track everything from weather to inventory." BusinessWeek, Aug 18, 2003

Challenges

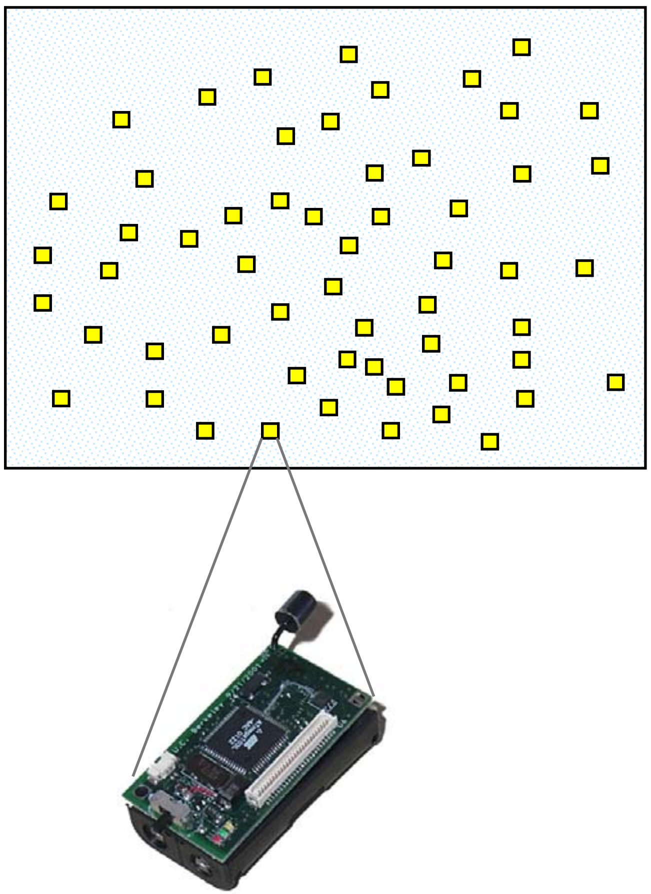

Small programs on tiny devices

Hardware challenges

- Limited capabilities
  - Processing, storage, comm
- Limited resources
  - Power, bandwidth

Sample Sensor Hardware: Berkeley motes

- CPU:
  - 8-bit, 4 MHz Atmel processor
  - No floating-point arithmetic support
- Radio:
  - 916 MHz RFM @10Kbps
  - Distance 30-100ft
  - Adjustable strength for RF transmission & reception
- Storage:
  - 8 KB instruction flash
  - 512 bytes data RAM
  - 512 bytes EEPROM (on processor)
- OS:
  - TinyOS, event driven (3.5KB code space)
- Sensors:
  - 10-bit ADC mux'd over 7 analog input channels
  - Sensing: light, sound, temperature, acceleration, magnetic field, pressure, humidity, RF signal strength

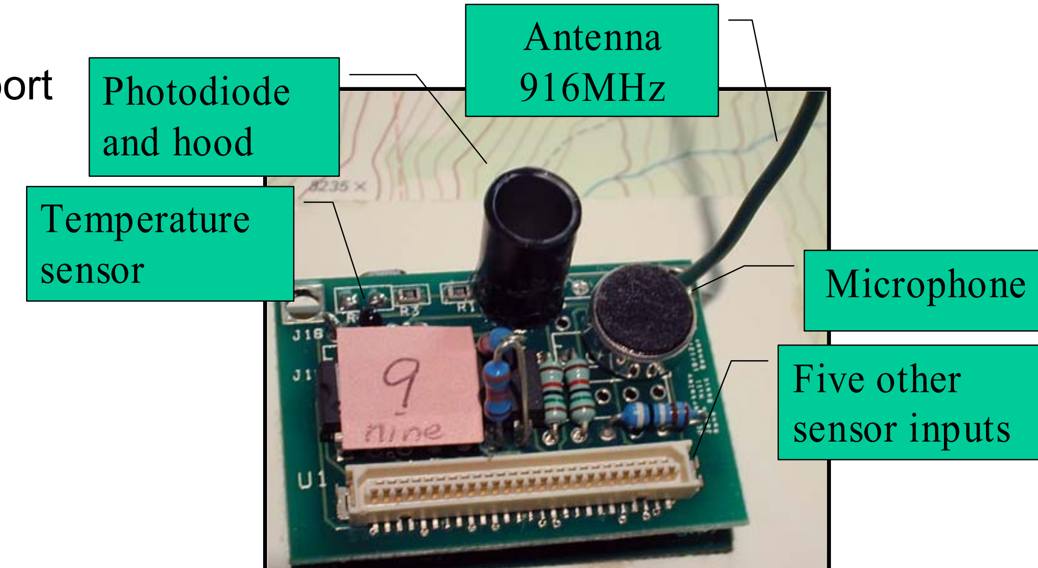

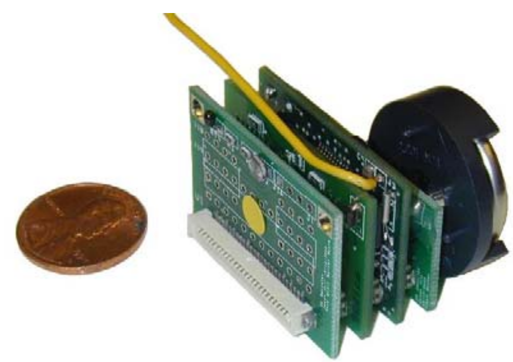

Hardware

parc
Palo Alto Research Center

Power Breakdown…

	Active	Idle	Sleep
CPU	5 mA	2 mA	5 μA
Radio	7 mA (TX)	4.5 mA (RX)	5 μA
EE-Prom	3 mA	0	0
LED's	4 mA	0	0
Photo Diode	200 μA	0	0
Temperature	200 μA	0	0

Rene motes data, Jason Hill

Communication/computation ratio:

- Rene motes:
  - Comm: (7mA*3V/10e3)*8=16.8μJ per 8bit
  - Comp: 5mA*3V/4e6=3.8 nJ per instruction
  - Ratio: 4,400 instruction/hop
- Sensoria nodes:
  - Comm: (100mW/56e3)*32=58μJ per 32bit
  - Comp: 750mW/1.1e9=0.7nJ per instruction
  - Ratio: 82,000 instruction/hop

Panasonic CR2354
560 mAh

This means

– Lithium Battery runs for 35 hours at peak load and years at minimum load, a three orders of magnitude difference!

parc Palo Alto Research Center

Challenges

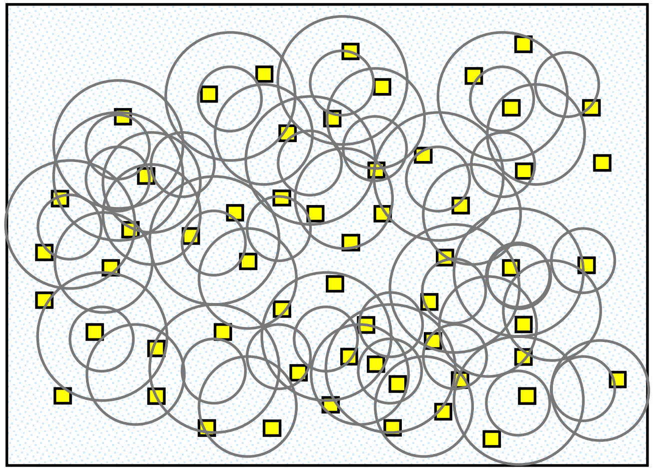

Hardware challenges

- Limited capabilities
  - Processing, storage, comm
- Limited resources
  - Power, bandwidth

Networking challenges

- Limited support
  - Peer-to-peer, mesh topology
  - Dynamic, mobile, unreliable connectivity
  - No universal routing protocols
  - No central name and registry services
- Both router and application host

Limited infrastructure support

parc Palo Alto Research Center

An application of wireless sensor network: fire monitoring

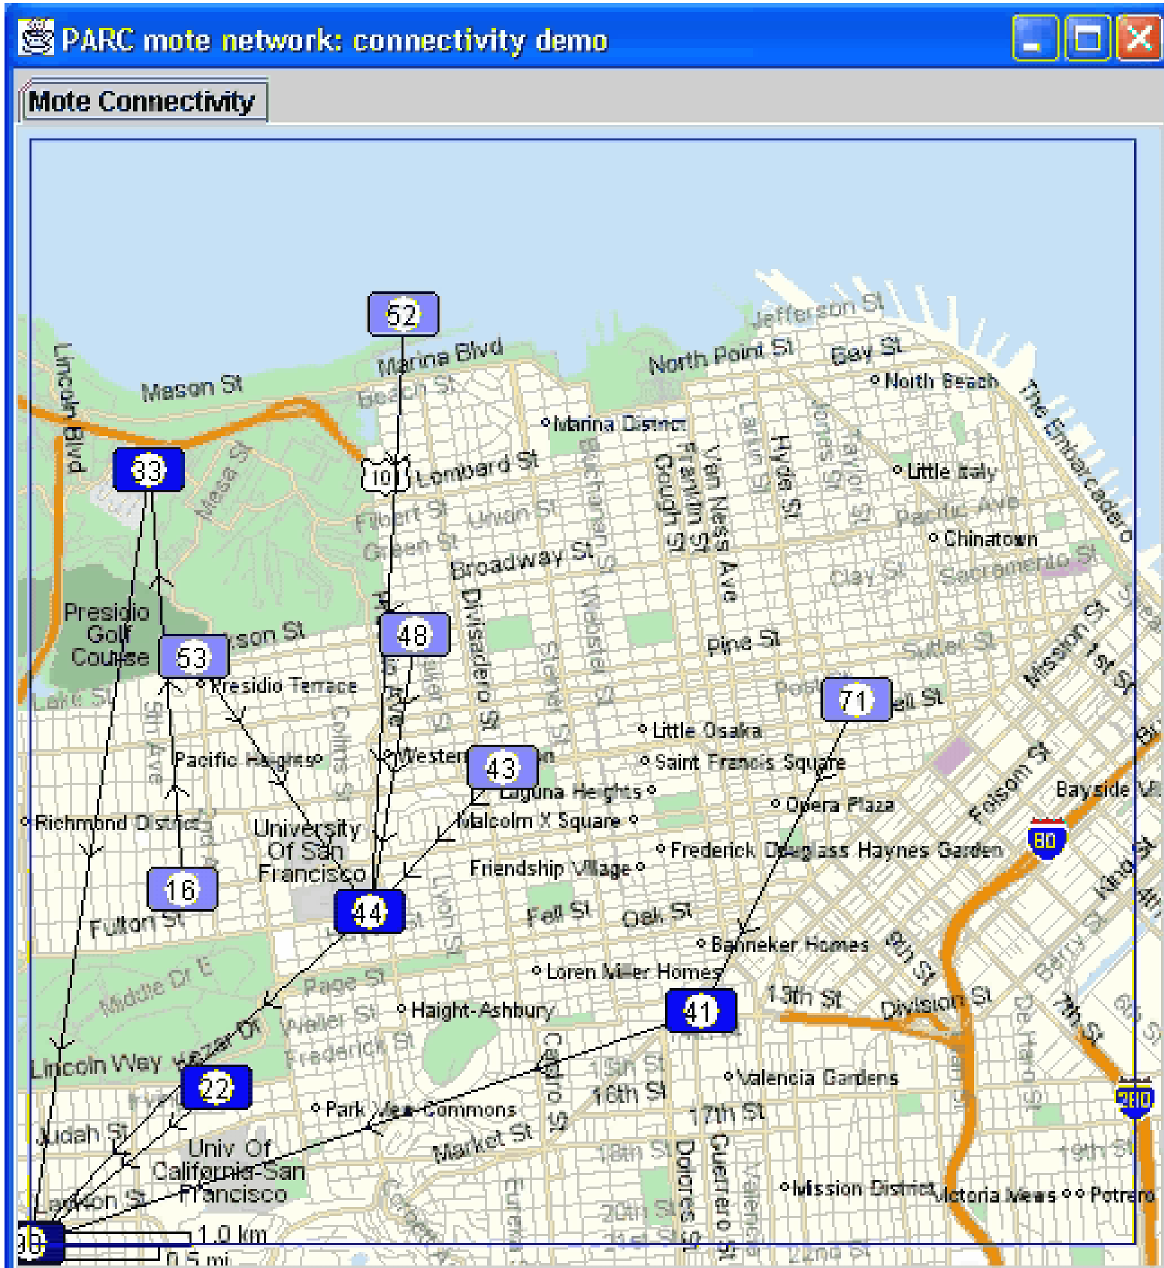

parc
Palo Alto Research Center

Challenges

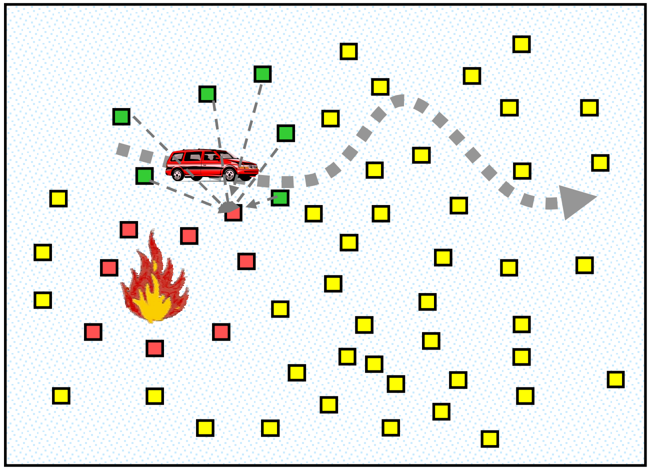

Hardware challenges

- Limited capabilities
  - Processing, storage, comm
- Limited resources
  - Power, bandwidth

Networking challenges

- Limited support
  - Peer-to-peer, mesh topology
  - Dynamic, mobile, unreliable connectivity
  - No universal routing protocols
  - No central name and registry services
- Both router and application host

Application challenges

- Dynamic collaboration among nodes
- Global property from local execution
- Competing events/tasks
- Real-time missions

Massively distributed multitasking

parc Palo Alto Research Center

Collaborative processing: monitoring multiple events

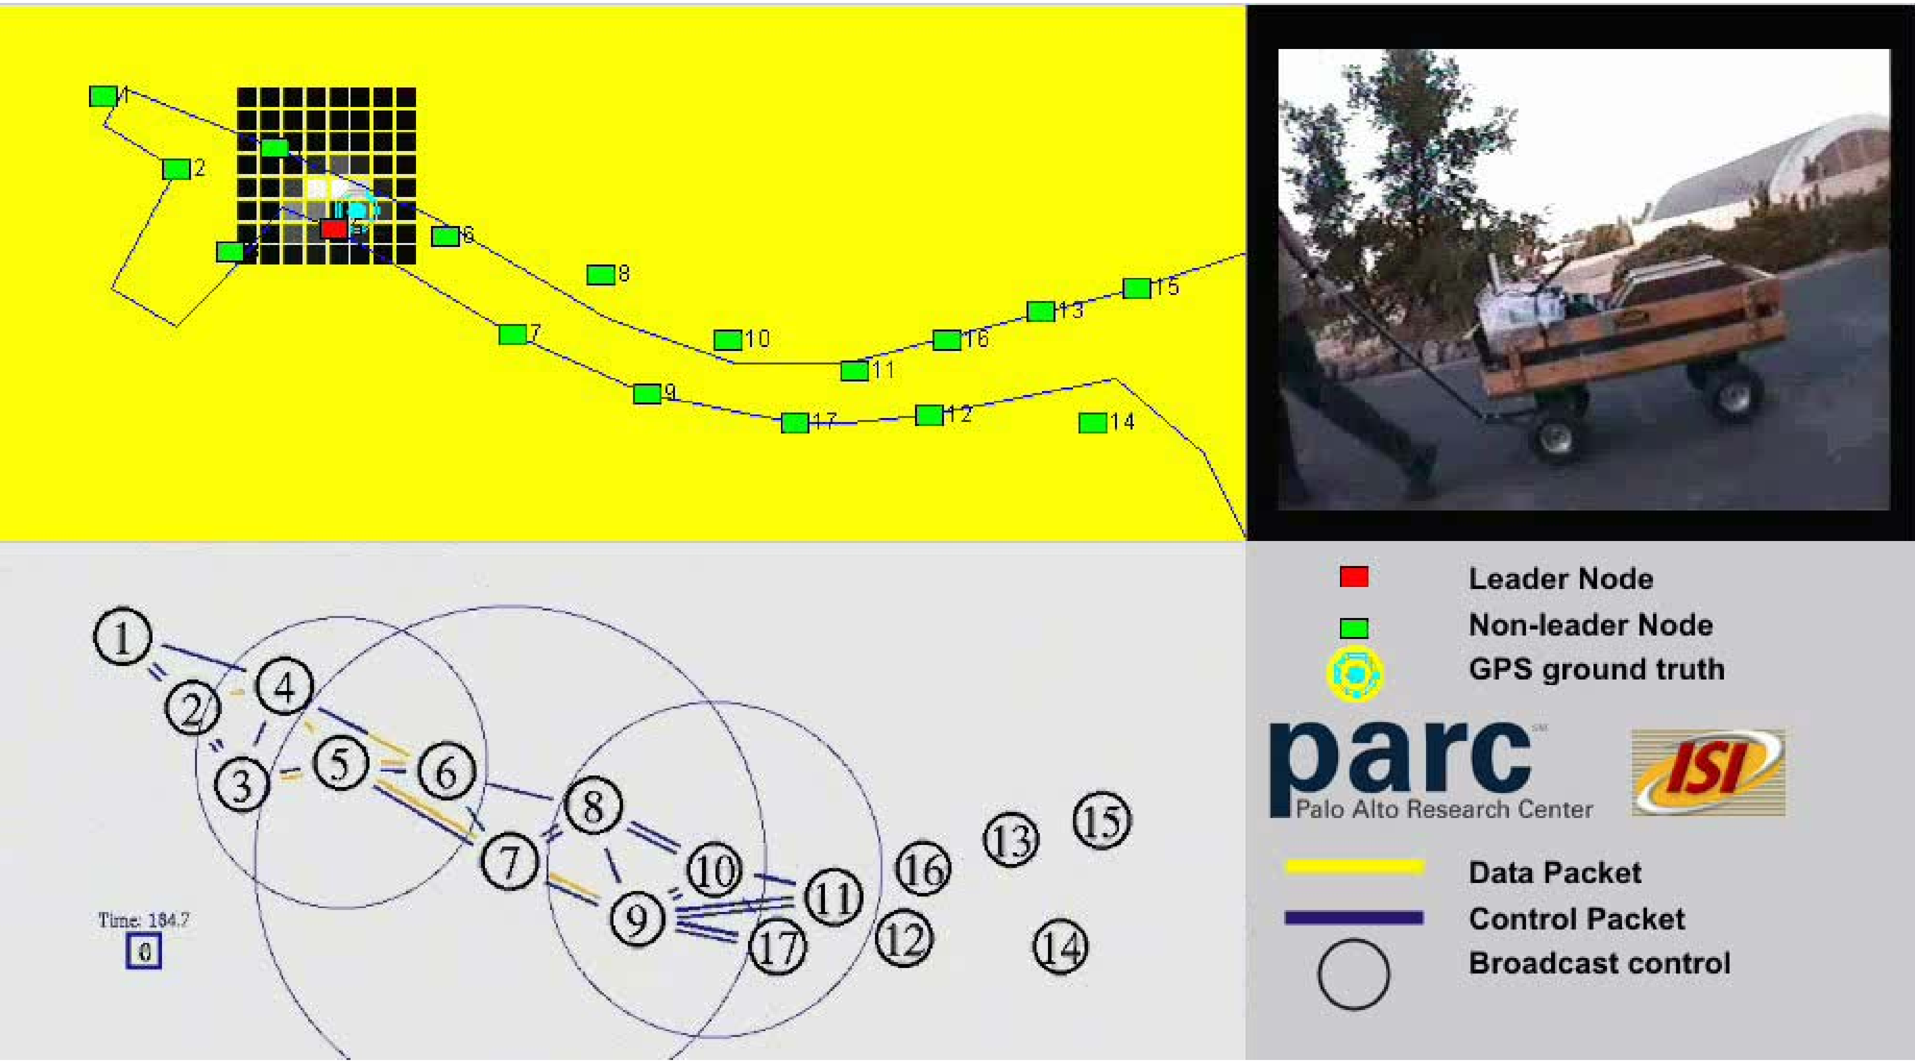

Rethinking the network infrastructure

In wireless ad hoc networks, networking is intimately coupled with sensing, interaction, and control needs and hence application semantics

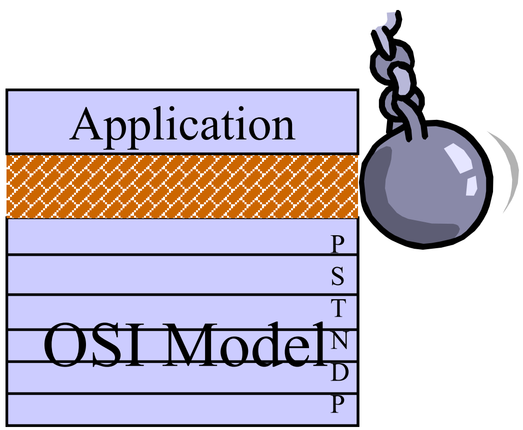

- Break down traditional barriers of OSI model
  - Consider both communication cost *and* application requirements to plan routes and task sensrors
- Data-centric and ad-hoc
  - Address nodes based on geography and capability, not by name
- Group management vital to scalability
  - Limit data propagation to sensors relevant to measurement at hand

As untethered sensors, actuators, embedded processors become ubiquitous, we need new ways to program and organize them

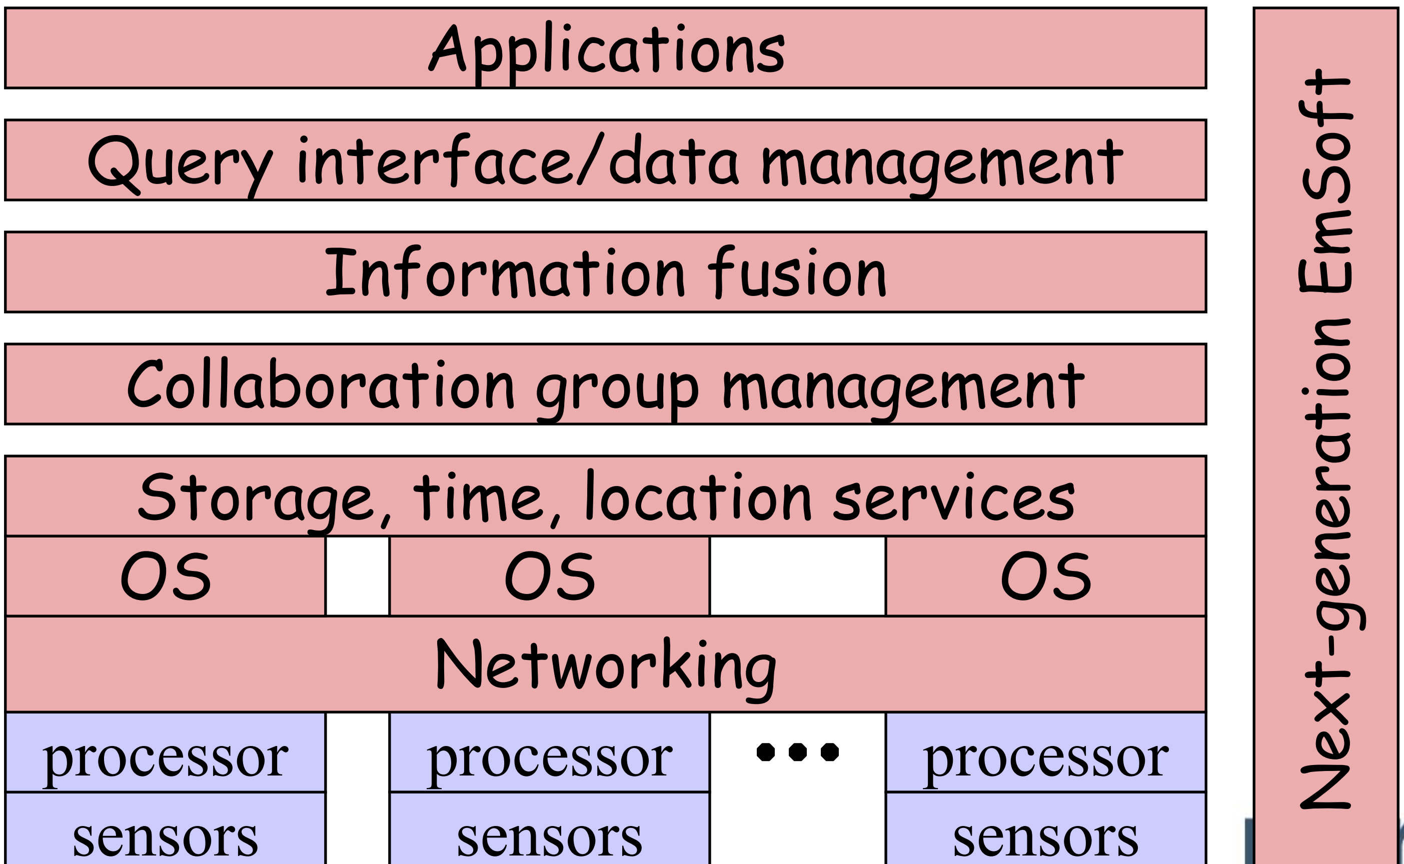

Palo Alto Research Center

A central problem: define and manage collaborative sensor groups dynamically based on their relevance to the current task and available network resources

Where is the data and how to move it to where it will be needed?

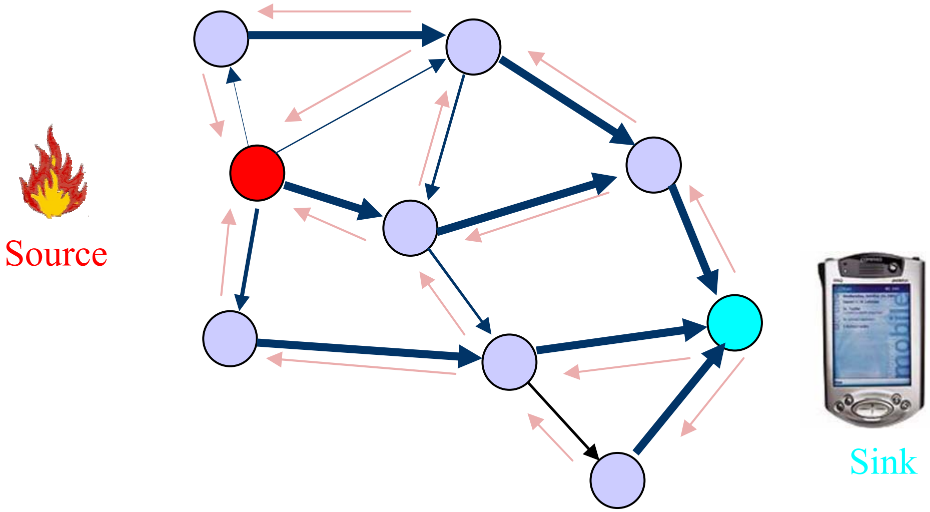

For example, use directed diffusion routing (Estrin et al)

- Publish and subscribe
  - Interest from user/data attribute from source => gradient
  - Finding shortest paths in graph

But we must also consider the content of data …

parc
Palo Alto Research Center

Moving Center of Aggregation Protocol

A leader node (blue square) carries belief state

- Choose sensor in the neighborhood with good information
- Hand off current belief to chosen sensor (new leader) and update

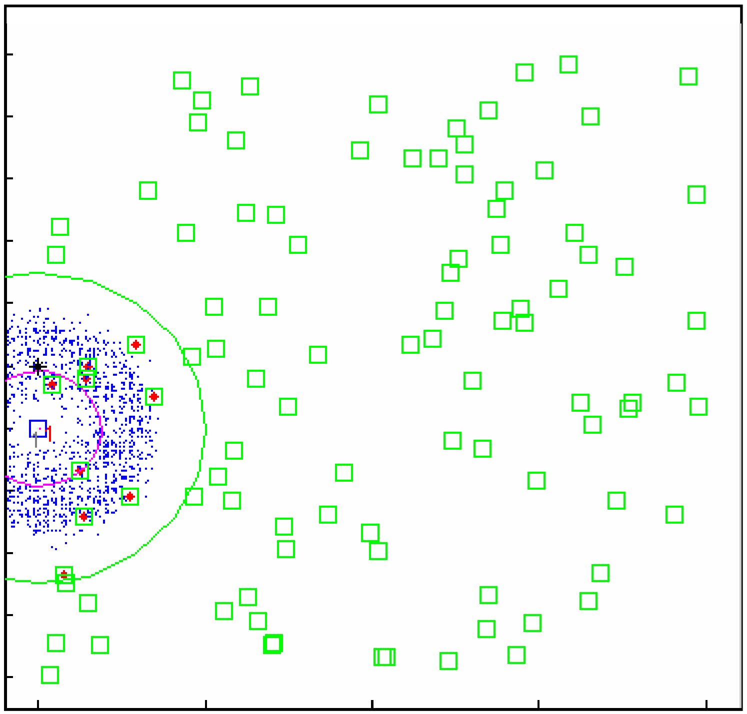

Target moving in straight line; Tracking using particle filter

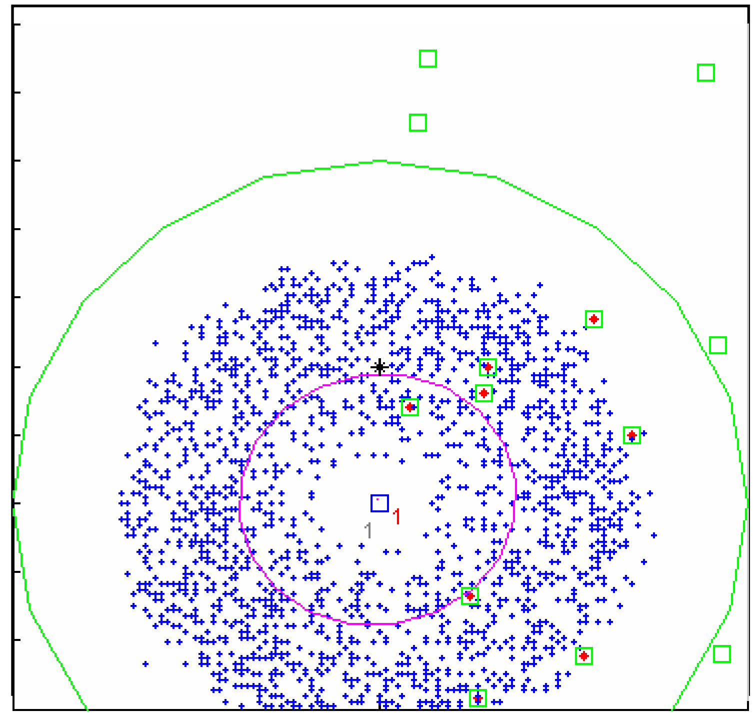

Close-up of target (particles show velocity vectors)

Sequential Bayesian Estimation

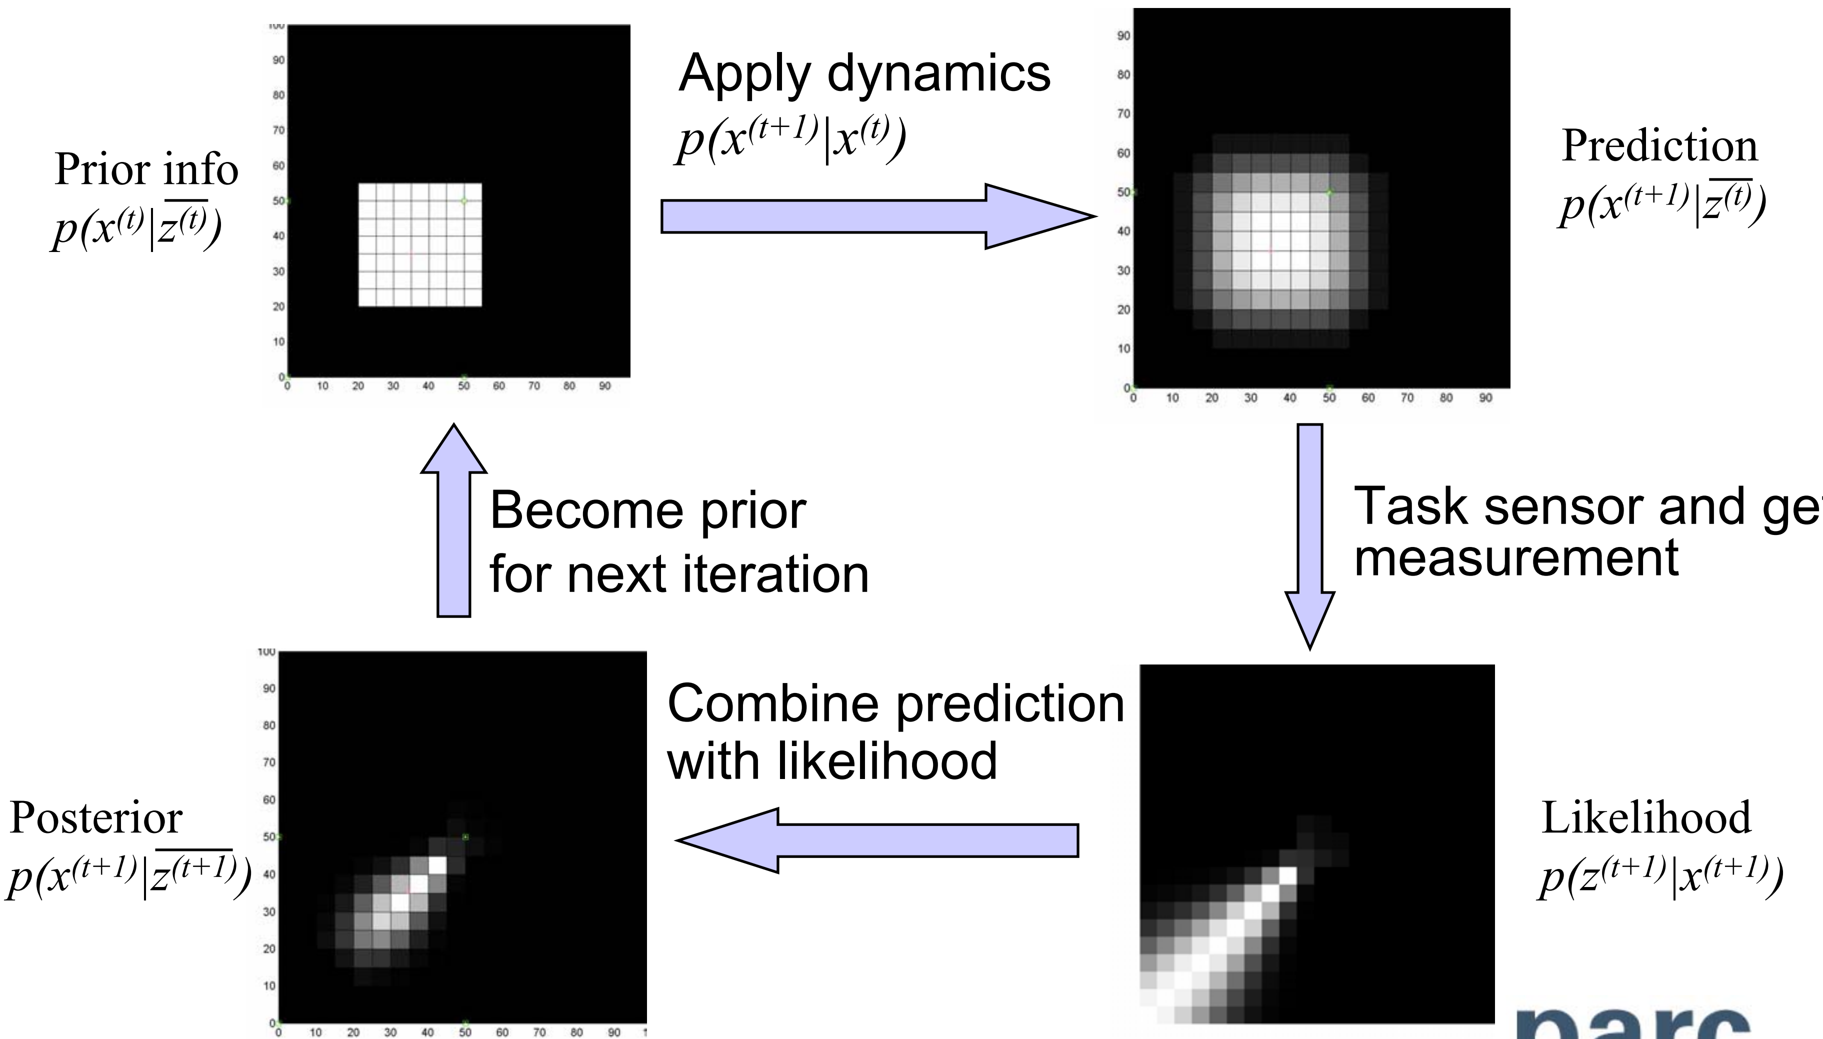

parc
Palo Alto Research Center

Information-Directed Sensor Querying (IDSQ)

- **Idea**: maximize the *predicted* information that a sensor's measurement will bring, given the current estimated distribution

- Information is measured using mutual information

$$I(U;V) = E_{p(u,v)}\left[\log \frac{p(u,v)}{p(u)p(v)}\right]$$

- IDSQ criteria: $k_{IDSQ} = \arg\max_{k \in N} I(X^{(t+1)}; Z^{(t+1)} \mid \overline{Z^{(t)}} = \overline{z^{(t)}})$

  where N is the set of candidate sensors (i.e. topological neighbors)

- This is equivalent to choosing the sensor which will give the greatest change to the current belief.

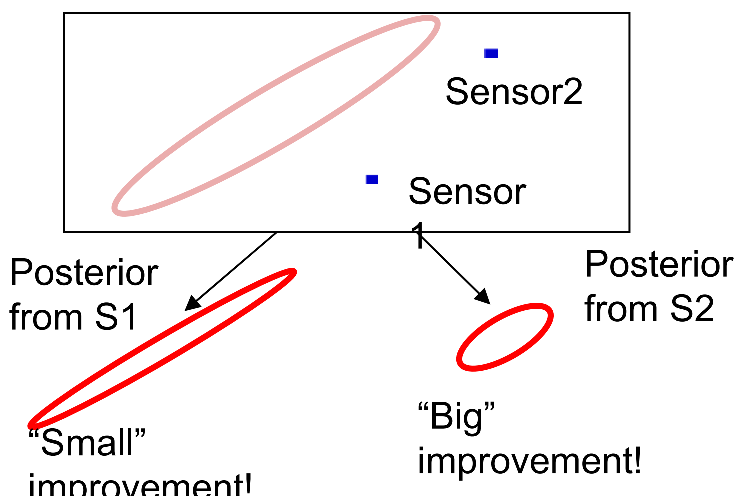

Managing Sensor Groups for Collaborative Processing Tasks

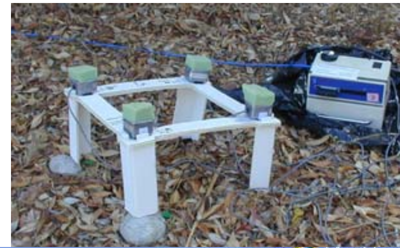

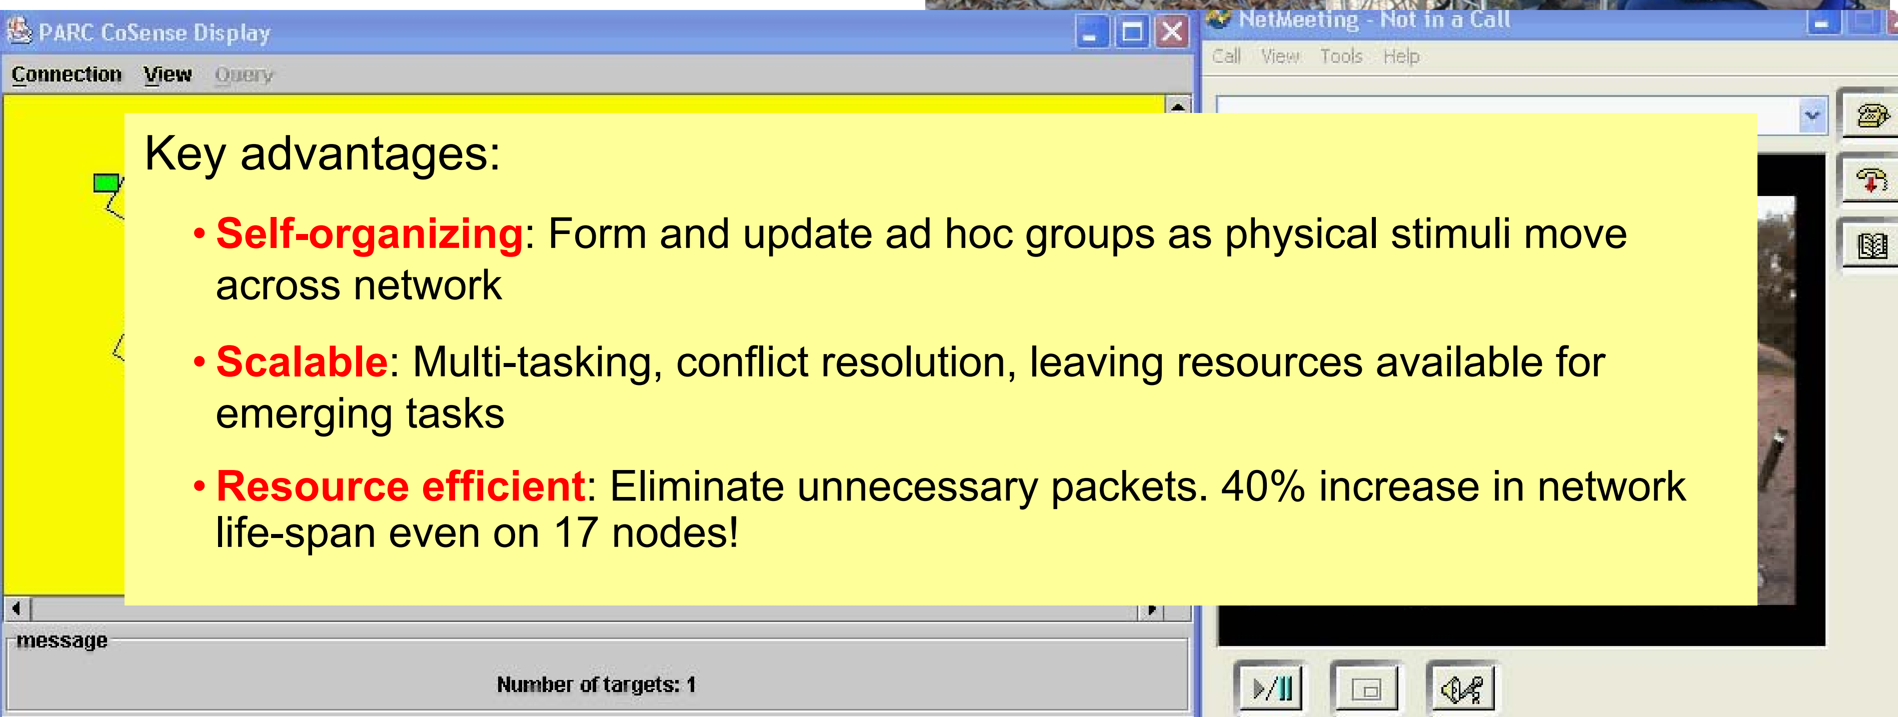

Key advantages:

- **Self-organizing**: Form and update ad hoc groups as physical stimuli move across network
- **Scalable**: Multi-tasking, conflict resolution, leaving resources available for emerging tasks
- **Resource efficient**: Eliminate unnecessary packets. 40% increase in network life-span even on 17 nodes!

- 17 nodes (6 DOA, 11 amplitude)
- Scale: 1 square=5 ft.
- 0.5 sec update interval
- Packet loss significant

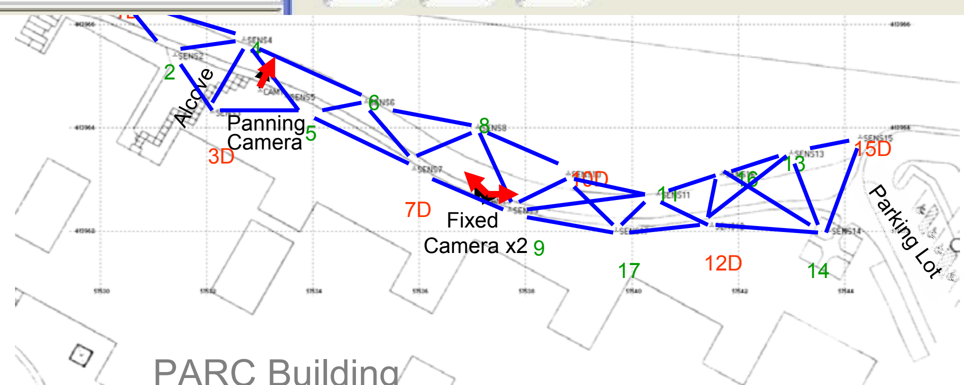

Attention is a scarce resource. How does a distributed system manage the explosion of information and attend to critical events?

Drinking from the fire hose

- **Internet**
  - Information overload
- **Sensor net**
  - Information explosion

How to manage this explosion of information?

Distributed Attention

- Be able to search for and focus on interesting activities, while maintaining a peripheral awareness of emerging phenomena
- Optimally allocate scarce resources to possibly competing sensing tasks

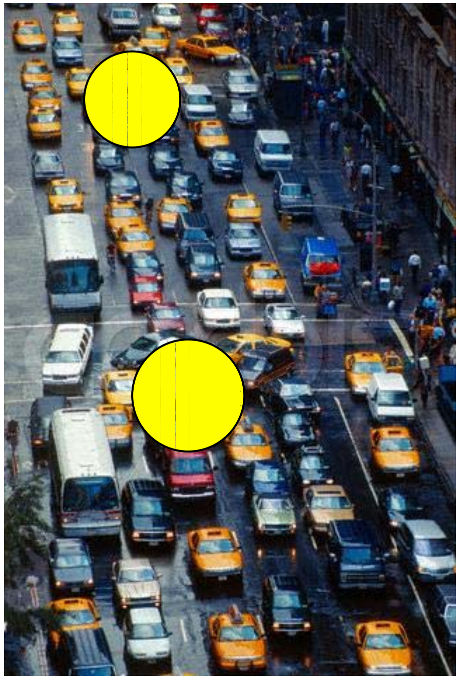

Traffic cams find trouble spots

parc Palo Alto Research Center

Think of a sensor network as a distributed market

- Sensors bid for sensing tasks
- Each sensing task has a utility
- The goal is to optimally allocate scarce resources to tasks
  - In a centralized setting, this is a classical assignment problem studied in OR and economics
- Here, it must be done in a local, peer-to-peer manner!

parc Palo Alto Research Center

Decentralized negotiation allocates resources among competing tasks

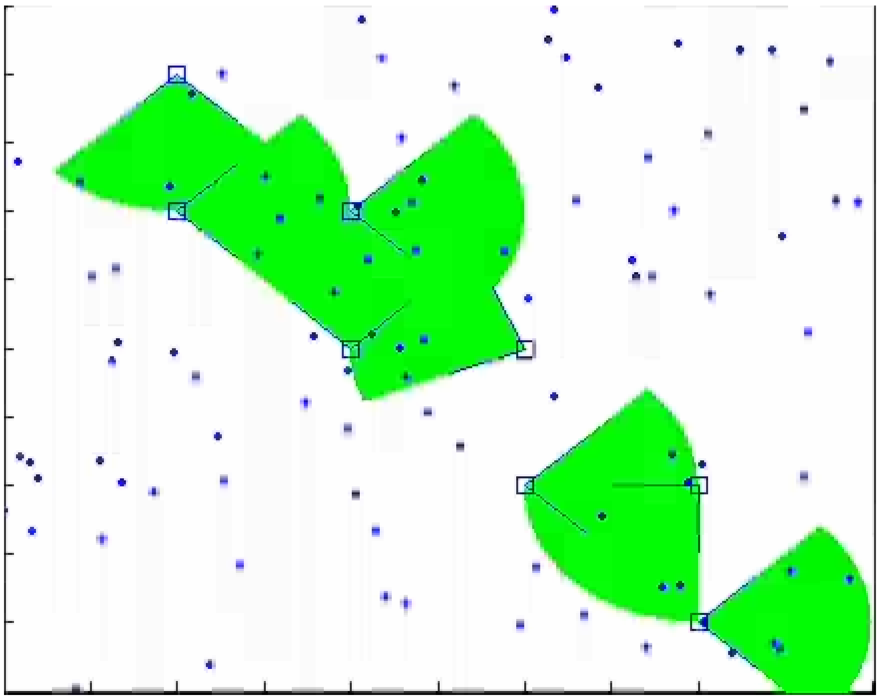

parc
Palo Alto Research Center

Recognizing and tracking "Interesting" Behaviors

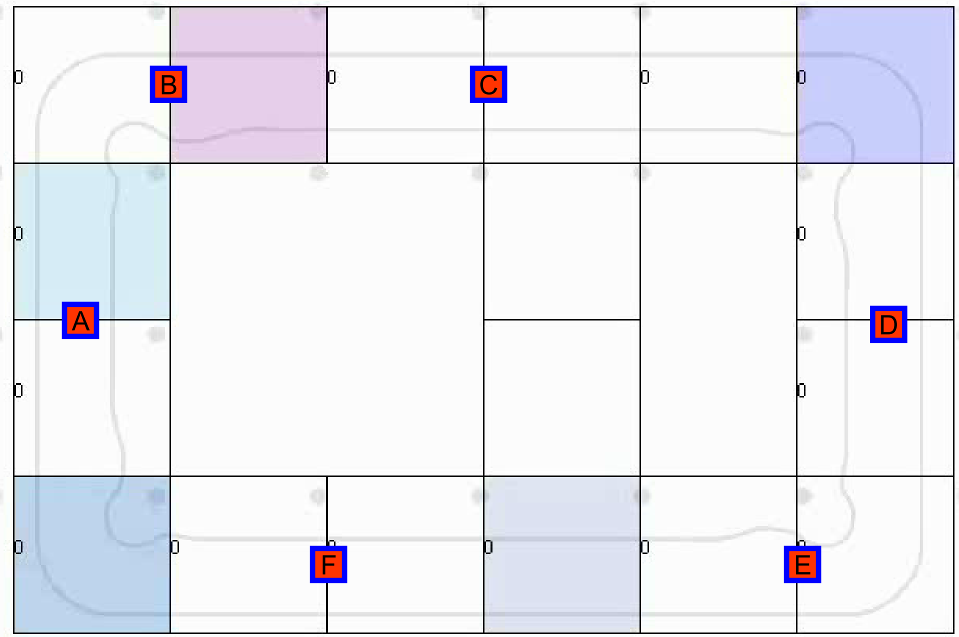

- 6 networked pan-and-tilt cameras
- FOV of cameras; each camera can see two FOV's on either side; overlapping FOVs
- Calibration points (1.5" diameter)
- Robot tracks (thickness 3/8")

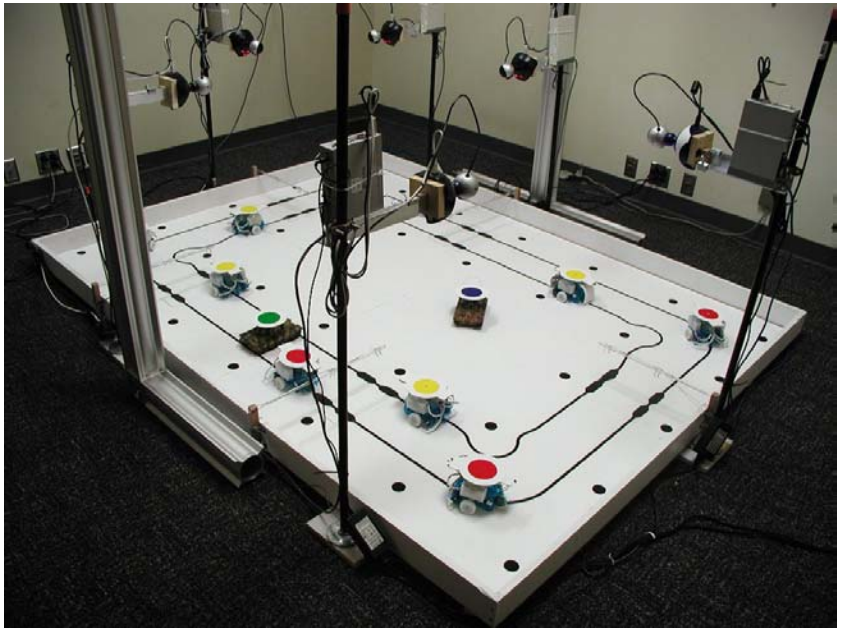

PARC CIPT testbed

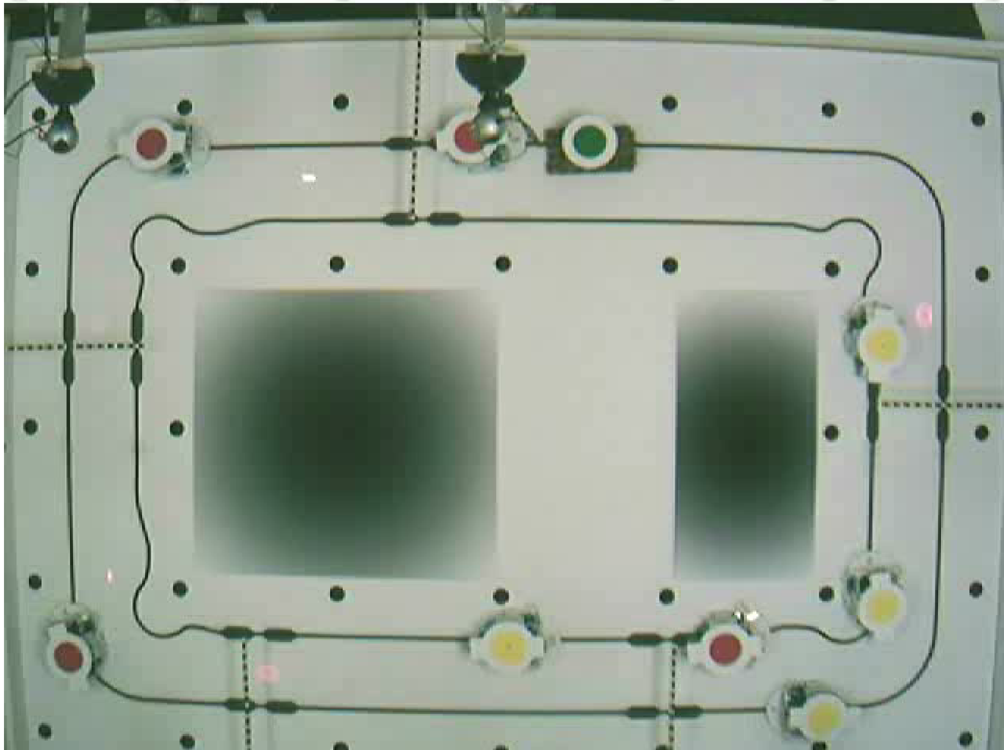

Collaborative groups raise abstraction level, to enable programming over collectives

Programming embedded sensor networks: a comparison

OS/Network-centric view

Spend more time designing component-level abstractions

Information-centric view

Spend more time designing application-level abstractions

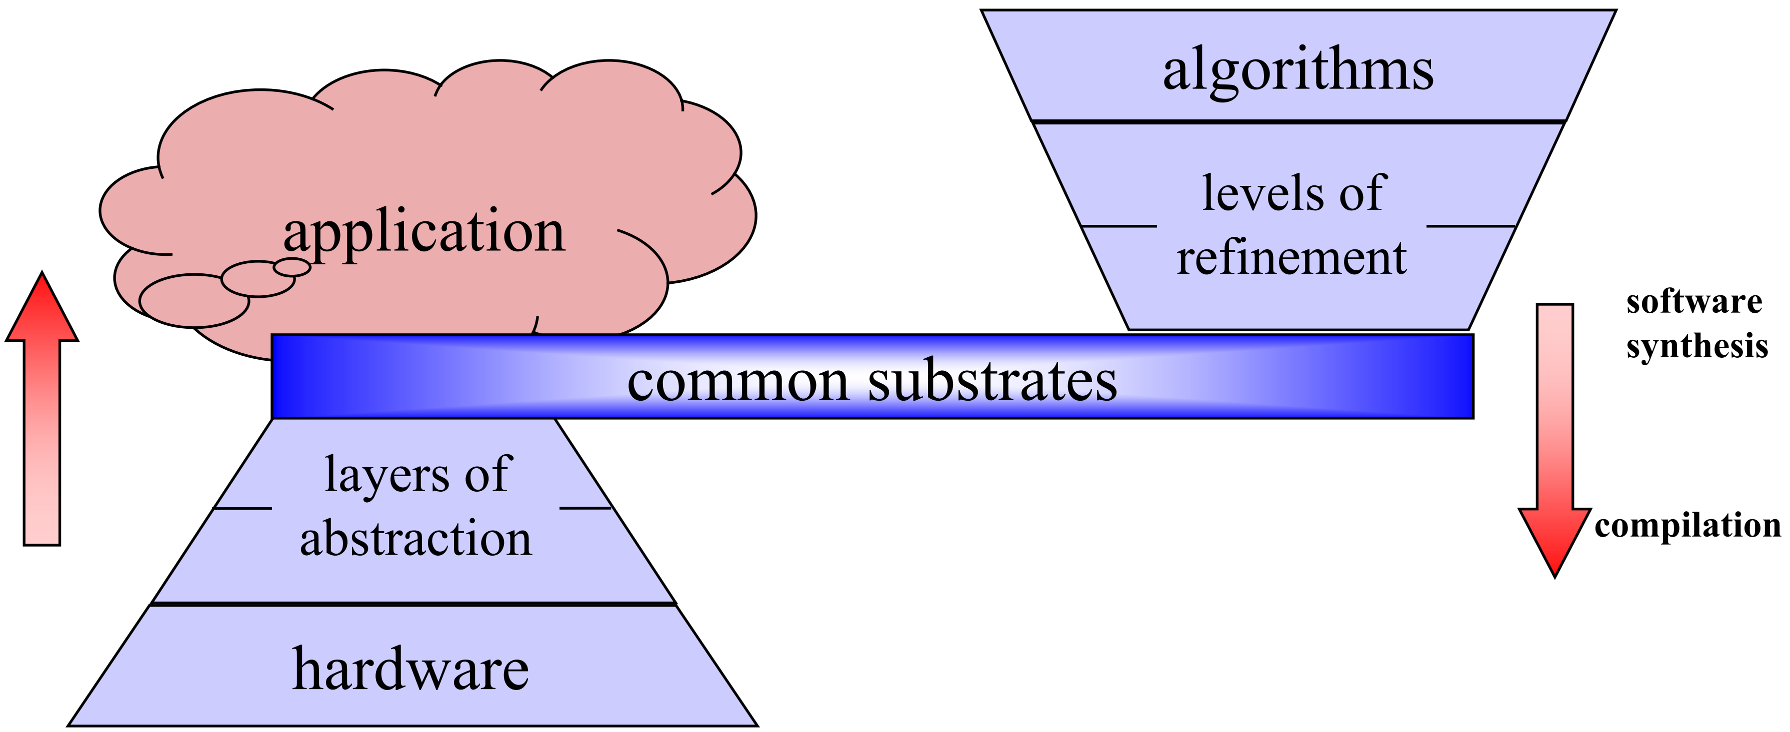

parc Palo Alto Research Center

Examples of Collaboration Groups

- N-hop neighbor group

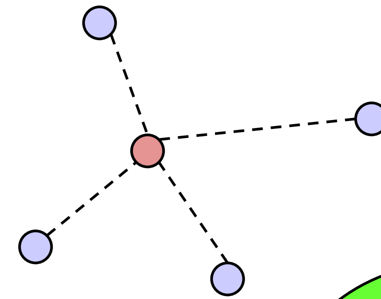

- Geographically constrained group
  - Defined by geographic extent

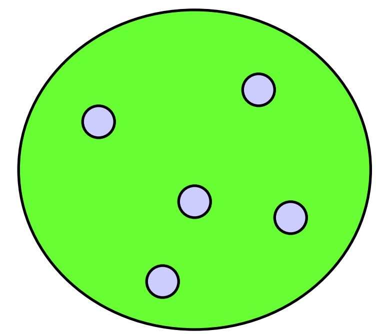

- Publish-subscribe group
  - Defined by producers and consumers of shared interests

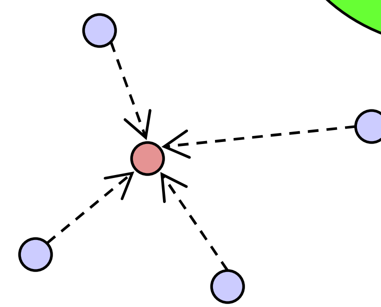

- Acquaintance group
  - Roaming members keep persistent connectivity

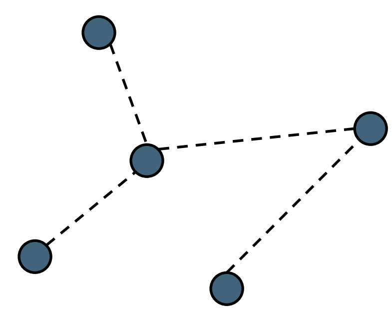

Tracking multiple, interacting events

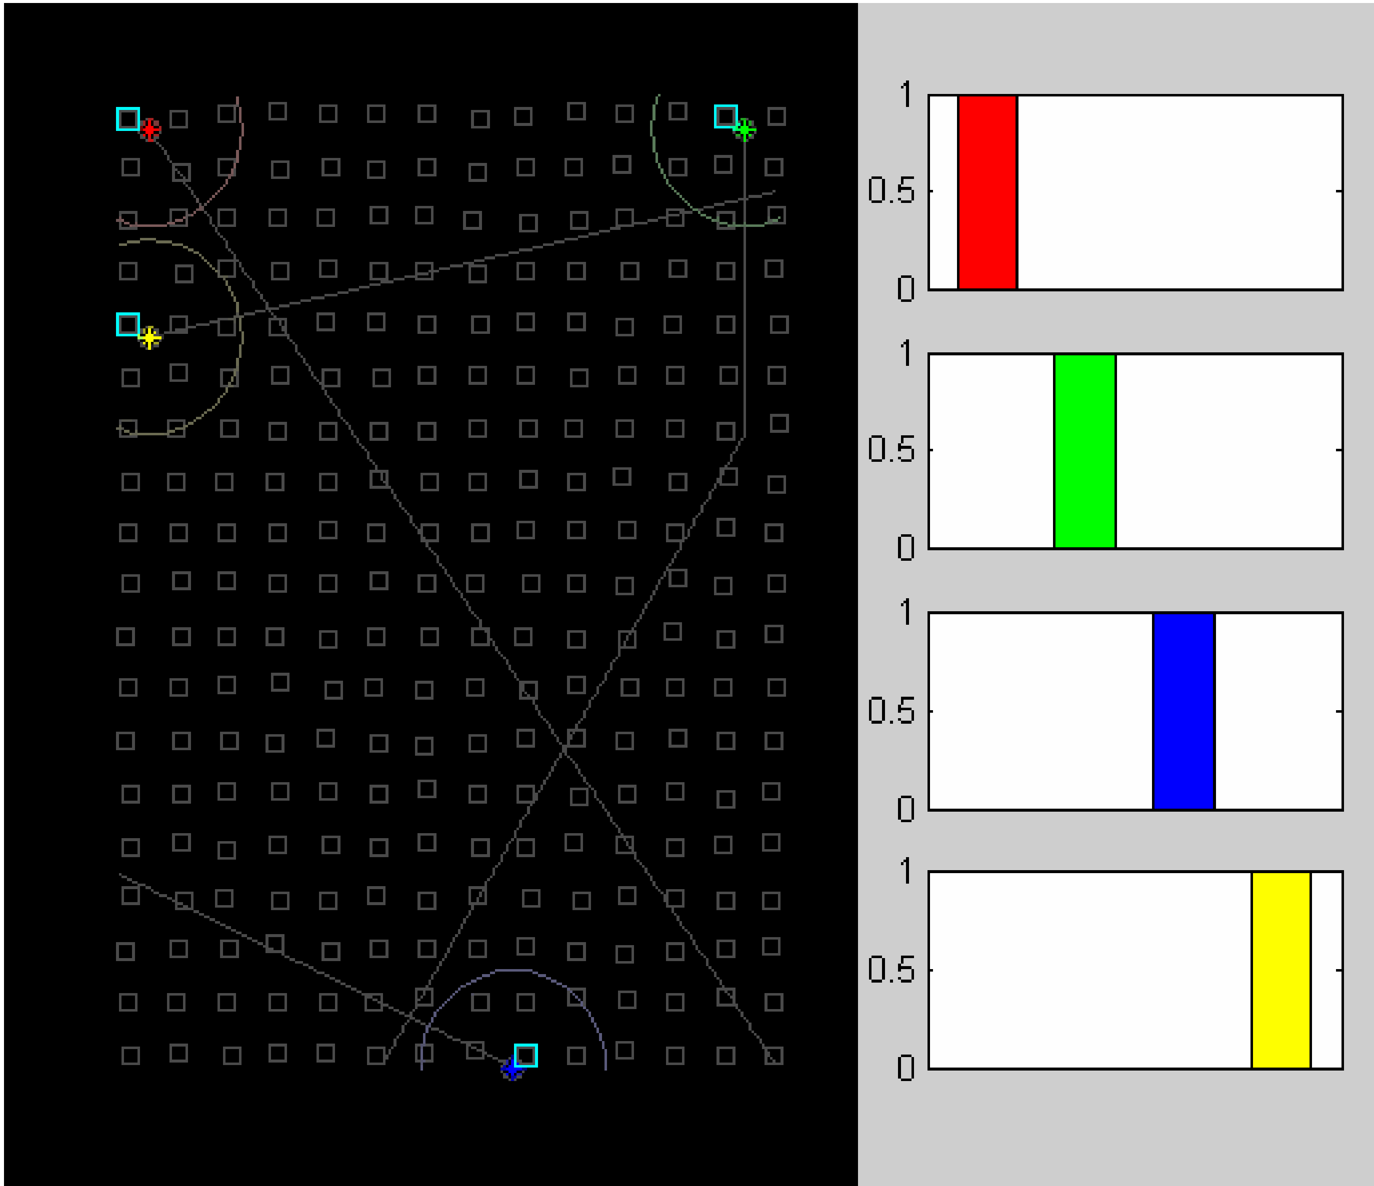

State-Centric Programming

- Raise the abstraction level beyond individual nodes
  - Provide high-level primitives that act on application “states”
  - Shield programmers from handling low-level events
- Models of collaboration
  - Abstract out common patterns of collaborative processing
  - Mix and match communication protocols from library
  - Modularize software through well-structured interfaces

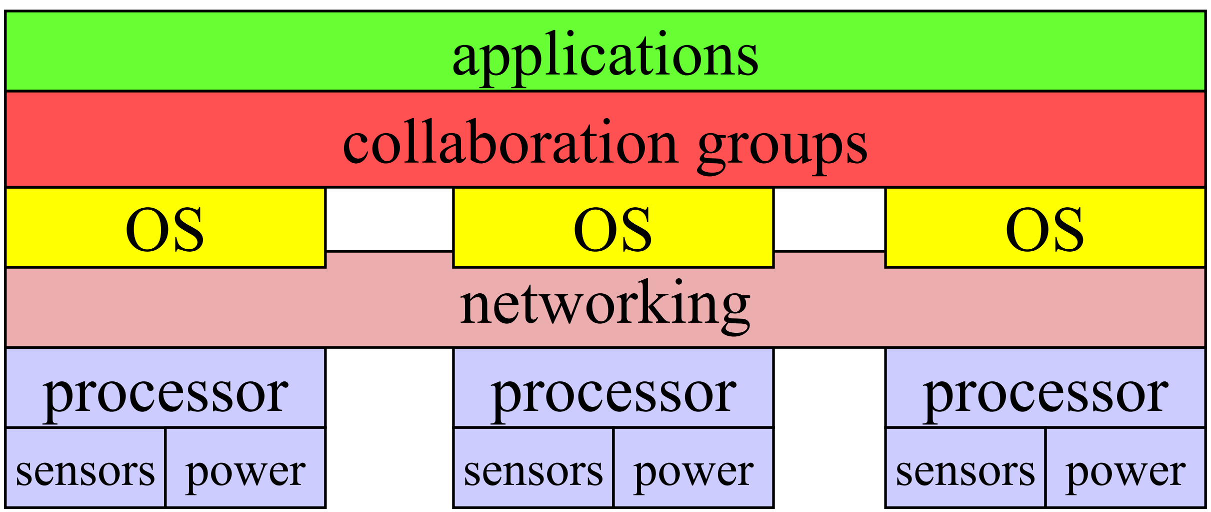

parc Palo Alto Research Center

Co-design information and software architectures

Software Technologies

- scalable software architectures
- formal methods
- software reuse
- software synthesis
- networking
- compilers
- operating systems
- …

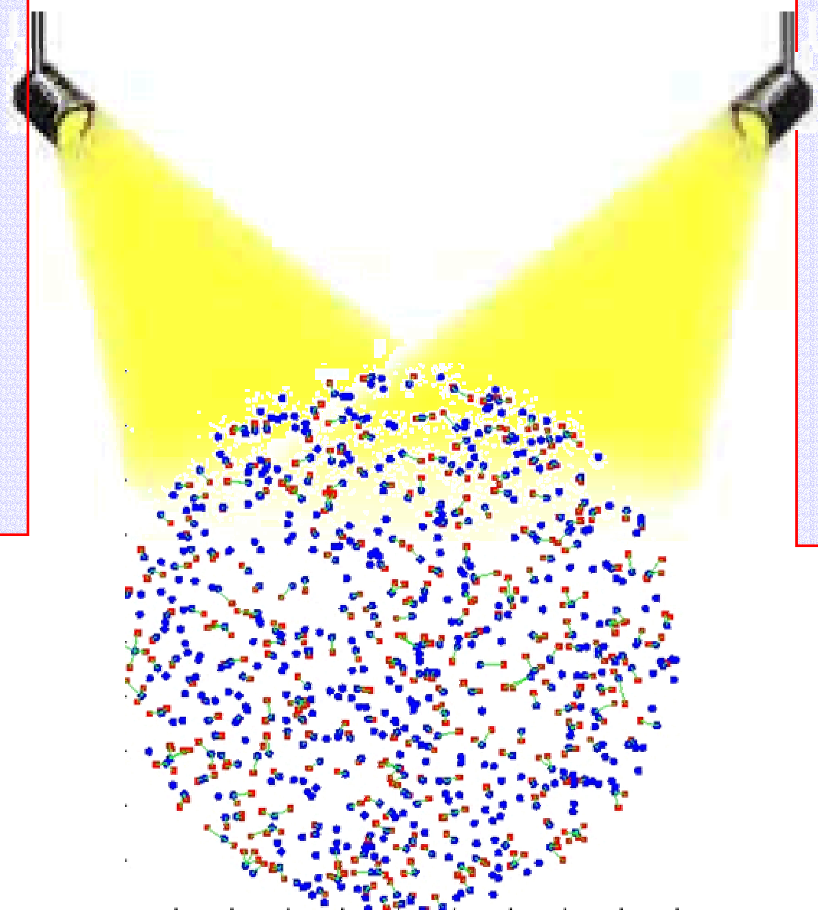

Information Technologies

- scalable information architecture
- semantics/ontology
- abstraction
- uncertainty management
- attention
- adaptation
- learning
- …

parc
Palo Alto Research Center

To Summarize …

- Couple information processing with networking
  - Sensor net requires cross-layer optimization
  - Key to resource management and scalability
- Collaborative groups as abstraction
  - Common building blocks for many applications
  - Can support programming beyond nodes

Scaling up sensor networks

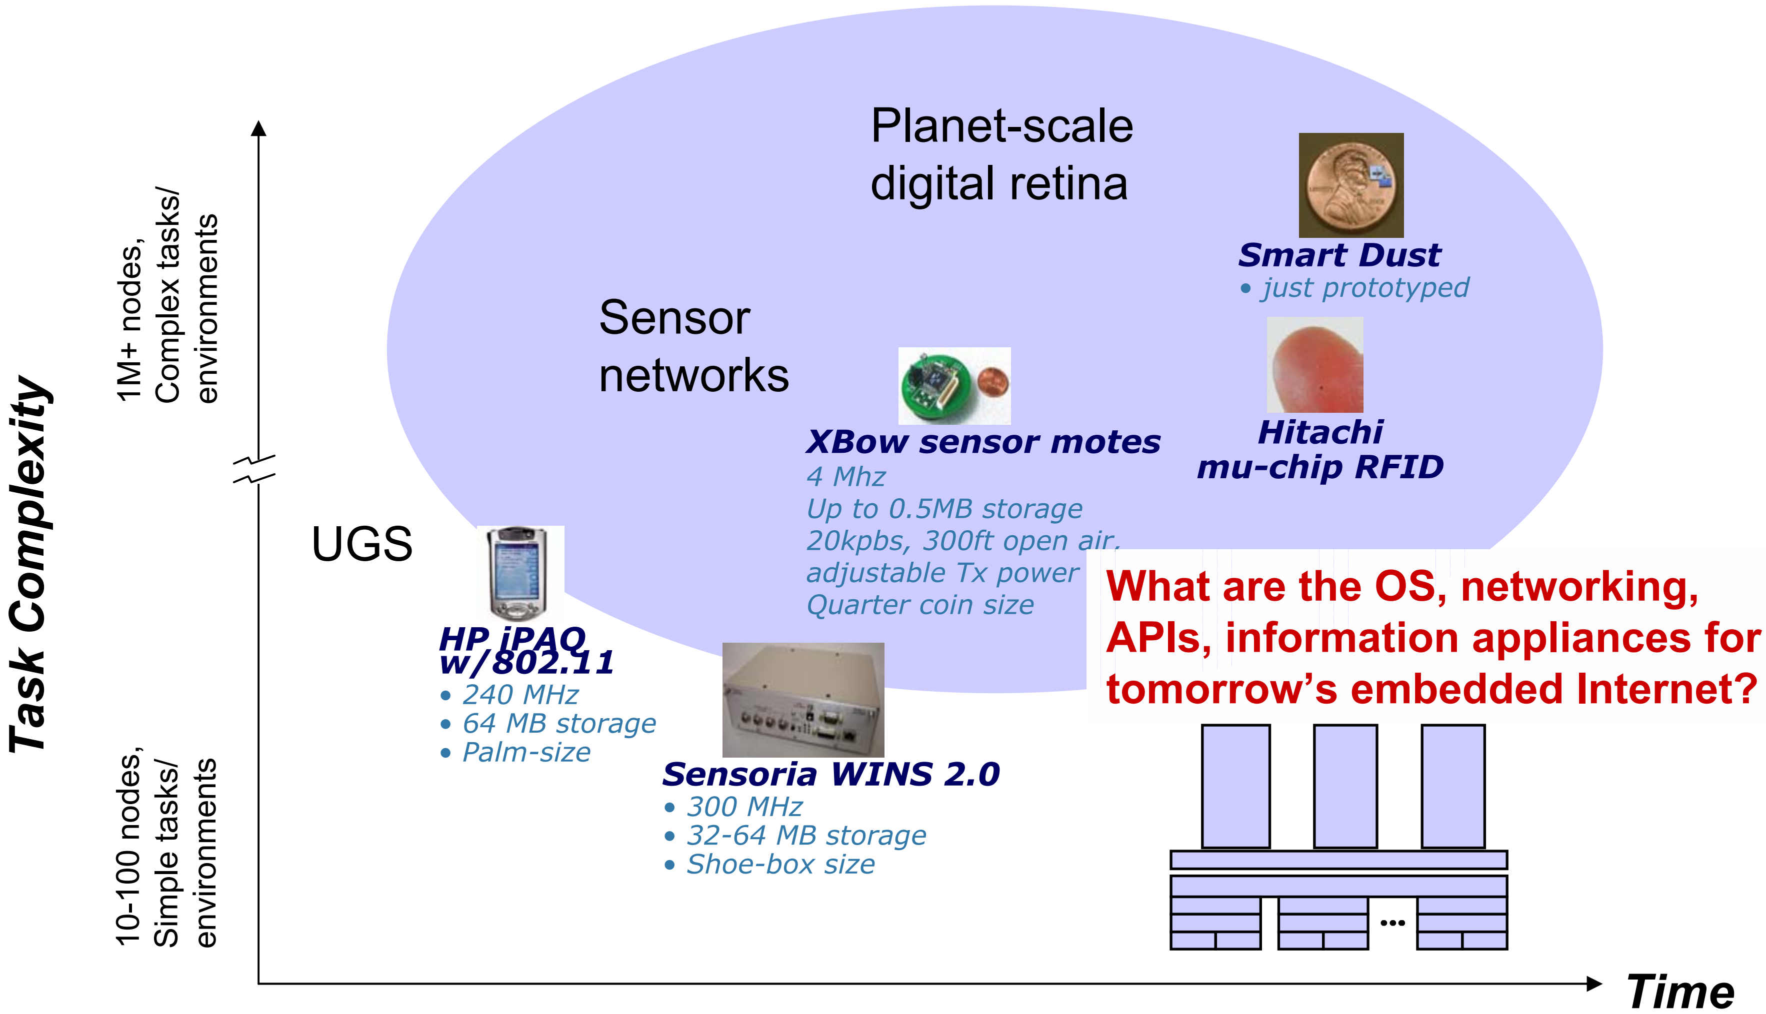

Today's systems:
Small scale, expensive, difficult to use, fixed, engineered, brittle

Tomorrow's demands:
Planet scale, cheap, simple to use, mobile, self-organized, flexible

To Probe Further

- PARC Project: www.parc.com/ecca
- Conferences:
  - IEEE/ACM IPSN04, Berkeley, April 2004
  - ACM Sensys03, Nov. 2003
  - ACM WSNA03, Sept. 2003
  - EWSN04, Berlin, Jan. 2004
  - INSS04, Tokyo, June 2004
- Journals:
  - New ACM Trans. Sensor Networks, www.acm.org/tosn
  - A dozen special issues

IPSN'04: 3rd IEEE/ACM Symposium on Information Processing in Sensor Networks

- Conference
  - April 26-27, 2004
  - Berkeley, California
- Sponsors
  - IEEE Signal Processing Soc./Communication Soc.
  - ACM SIGBED/SIGMOBILE
  - NSF, DARPA
- Web page
  - Ipsn04.cs.uiuc.edu