Information Processing in Microsensor Networks: an introduction

Feng Zhao

Palo Alto Research Center

(formerly Xerox PARC)

http://www.parc.com/zhao

parc
Palo Alto Research Center

Smart Sensors and Sensor Net

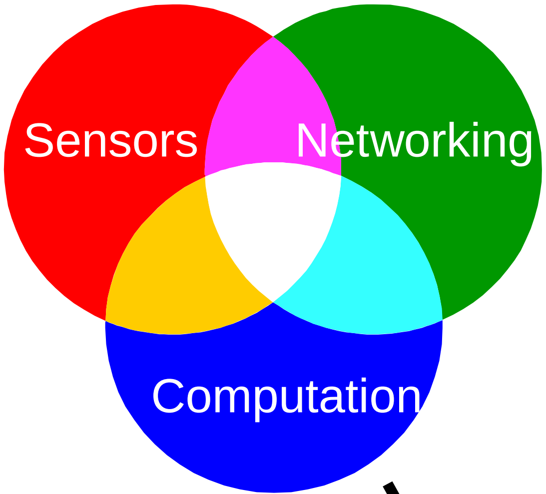

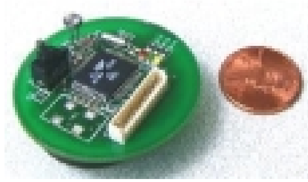

Berkeley Motes

Rockwell HiDRA

UCLA WINS

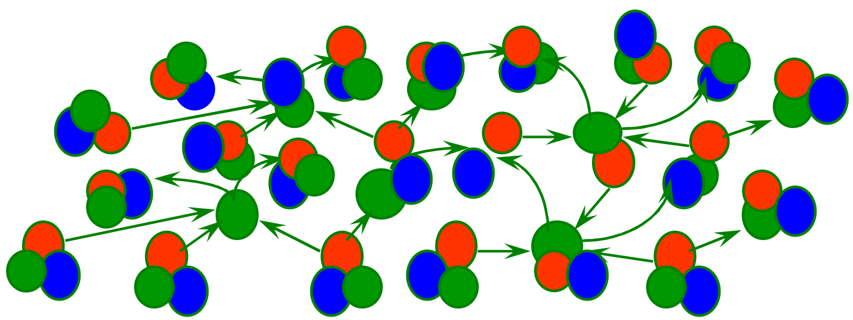

- **Environmental sensing**
  - –Traffic, habitat, hazards, security, anti-terrorism
- **Industrial sensing**
  - –Machine monitoring and diagnostics
  - –Power/telecom grid monitoring
- **Human-centered computing**
  - –Context-aware environment

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.

-- Horst Stormer in *Business Week, 8/23-30, 1999*.

parc
Palo Alto Research Center

Wireless sensor trend

- Of 9.6 billion uP to be shipped in 2005, 98% will be embedded processors!

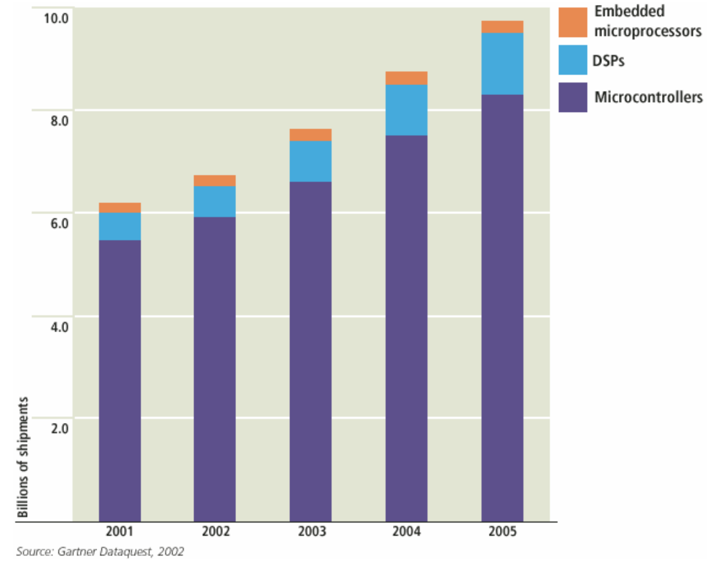

- Riding on Moore's law, smart sensors get

More powerful

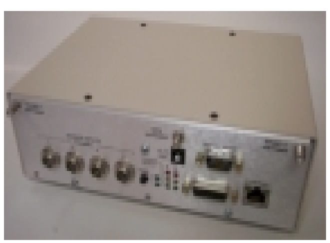

Sensoria WINSNG 2.0
CPU: 300 MIPS
1.1 GFLOP FPU
32MB Flash
32MB RAM
Sensors: external

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

Sources of constraints:

- Energy and physics of sensing, processing, and communication
- Task requirements from application domains

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

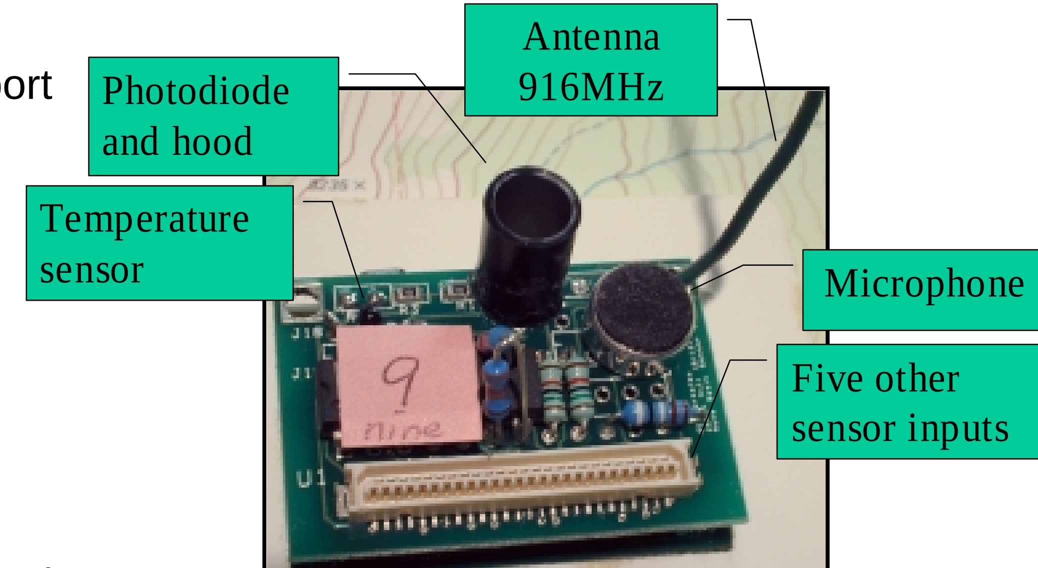

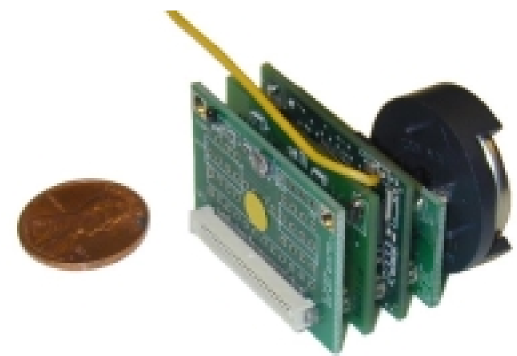

Hardware

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

Panasonic CR2354
560 mAh

Computation/communication 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

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

Distributed sensor net: multi-hop RF advantages

RF power attenuation near ground:

$$P_{receive} \propto \frac{P_{send}}{r^{\alpha}}, \qquad \alpha : 3-5$$

Or equivalently, $P_{send} \propto r^{\alpha} P_{receive}$

Power advantage:

$$\frac{P_{send\,(Nr)}}{N \cdot P_{send\,(r)}} = \frac{(Nr)^{\alpha} P_{receive}}{N \cdot r^{\alpha} P_{receive}} = N^{\alpha-1}$$

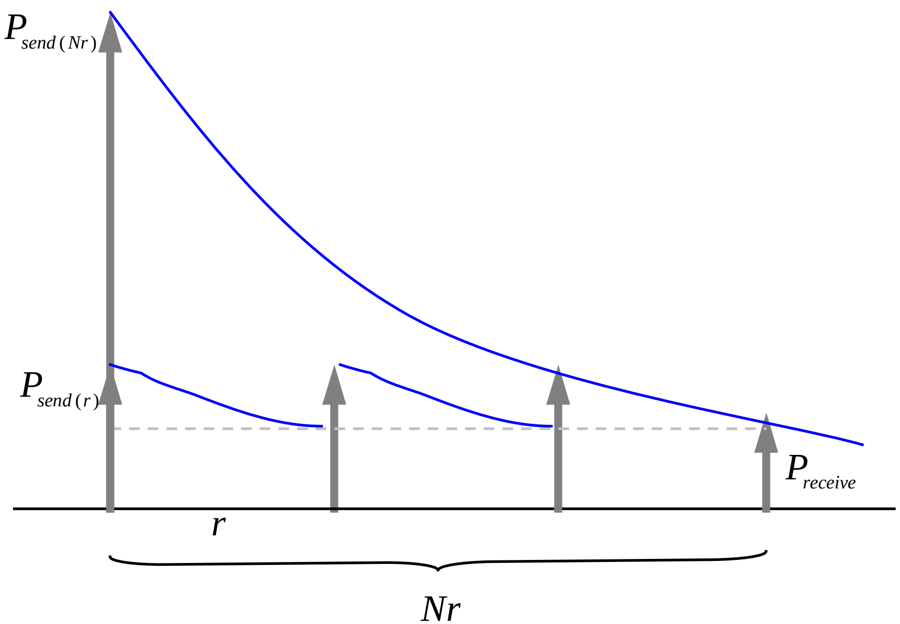

parc
Palo Alto Research Center

Distributed sensor net: detection and SNR advantages

Sensors have a finite sensing range. A denser sensor field improves the odds of detecting a target within the range. Once inside the range, further increasing sensor density by *N* improves the SNR by 10log*N* db (in 2D). Consider the acoustic sensing case:

Acoustic power received at distance *r*: $P_{receive} \propto \frac{P_{source}}{r^2}$

Signal-noise ratio (SNR):

$$SNR_r = 10\log P_{source} - 10\log P_{noise} - 20\log r$$

Increasing the sensor density by a factor of *N* gives a SNR advantage of:

$$SNR_{\frac{r}{\sqrt{N}}} - SNR_r = 20\log\frac{r}{\frac{r}{\sqrt{N}}} = 10\log N$$

parc
Palo Alto Research Center

Tracking as a canonical problem for studying information processing in sensor networks

Sensor Network Tracking

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

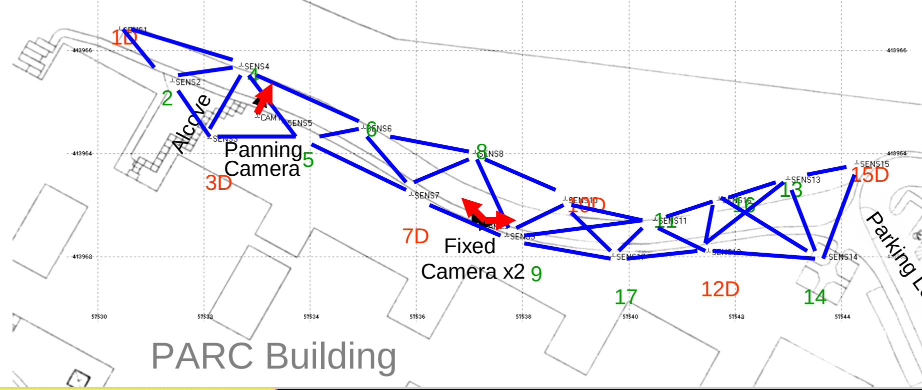

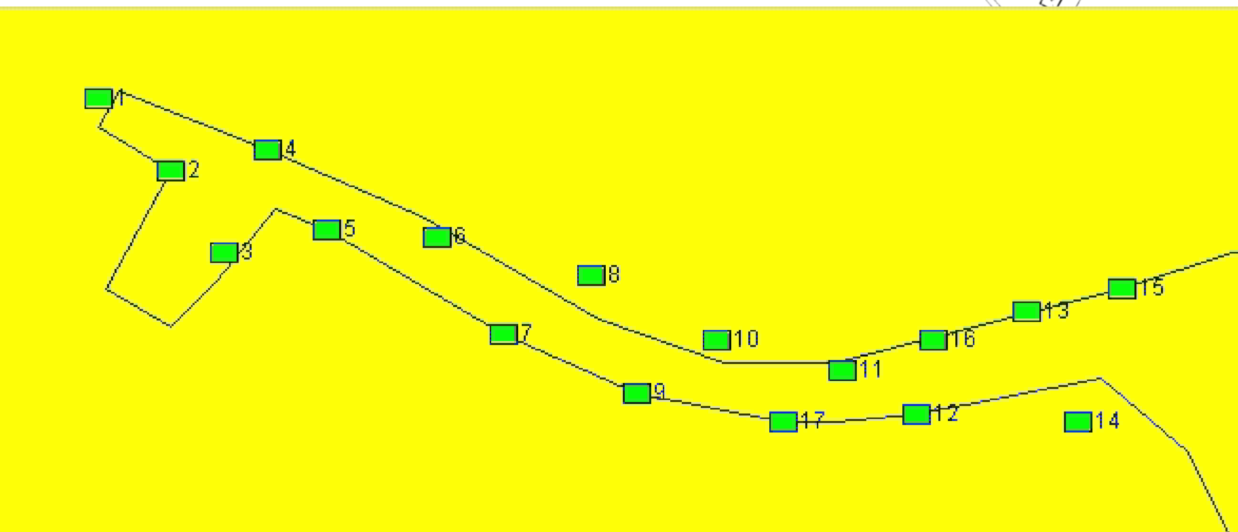

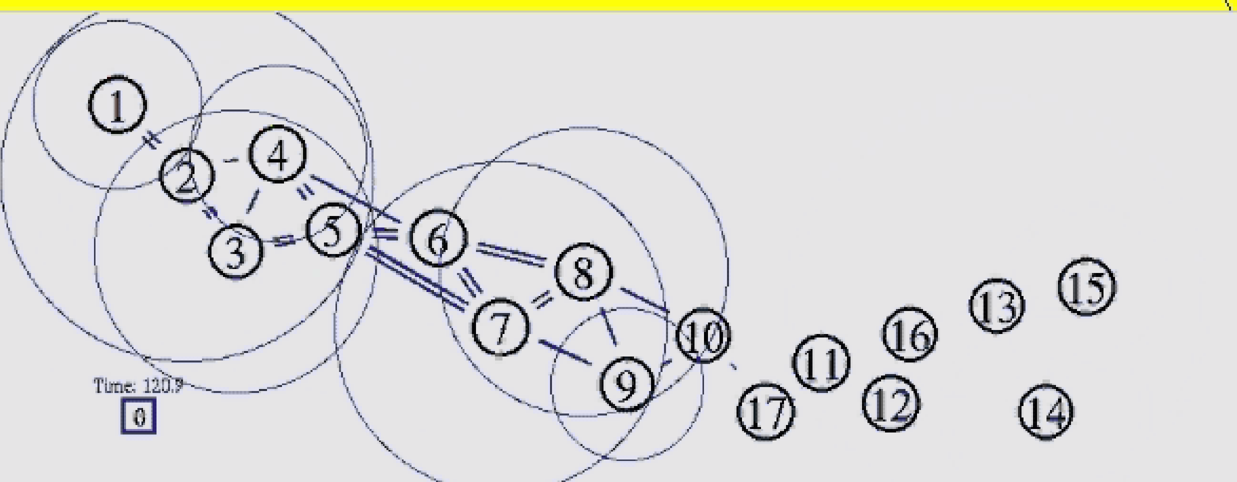

Leader Node

Non-leader Node

GPS ground truth

Palo Alto Research Center

Data Packet

Control Packet

Broadcast control

A tracking scenario

Constraints:
- Node power reserve
- RF path loss
- Packet loss
- Initialization cost
- …

1. **Discovery**: Node *a* detects the target and initializes tracking
2. **Query processing**: User query *Q* enters the net and is routed towards regions of interest
3. **Collaborative Processing**: Node *a* estimates target location, with help from neighboring nodes
4. **Communication protocol**: Node *a* may hand data off to node *b*, *b* to *c*, …
5. **Reporting**: Node *d* or *f* summarizes track data and send it back to the querying node

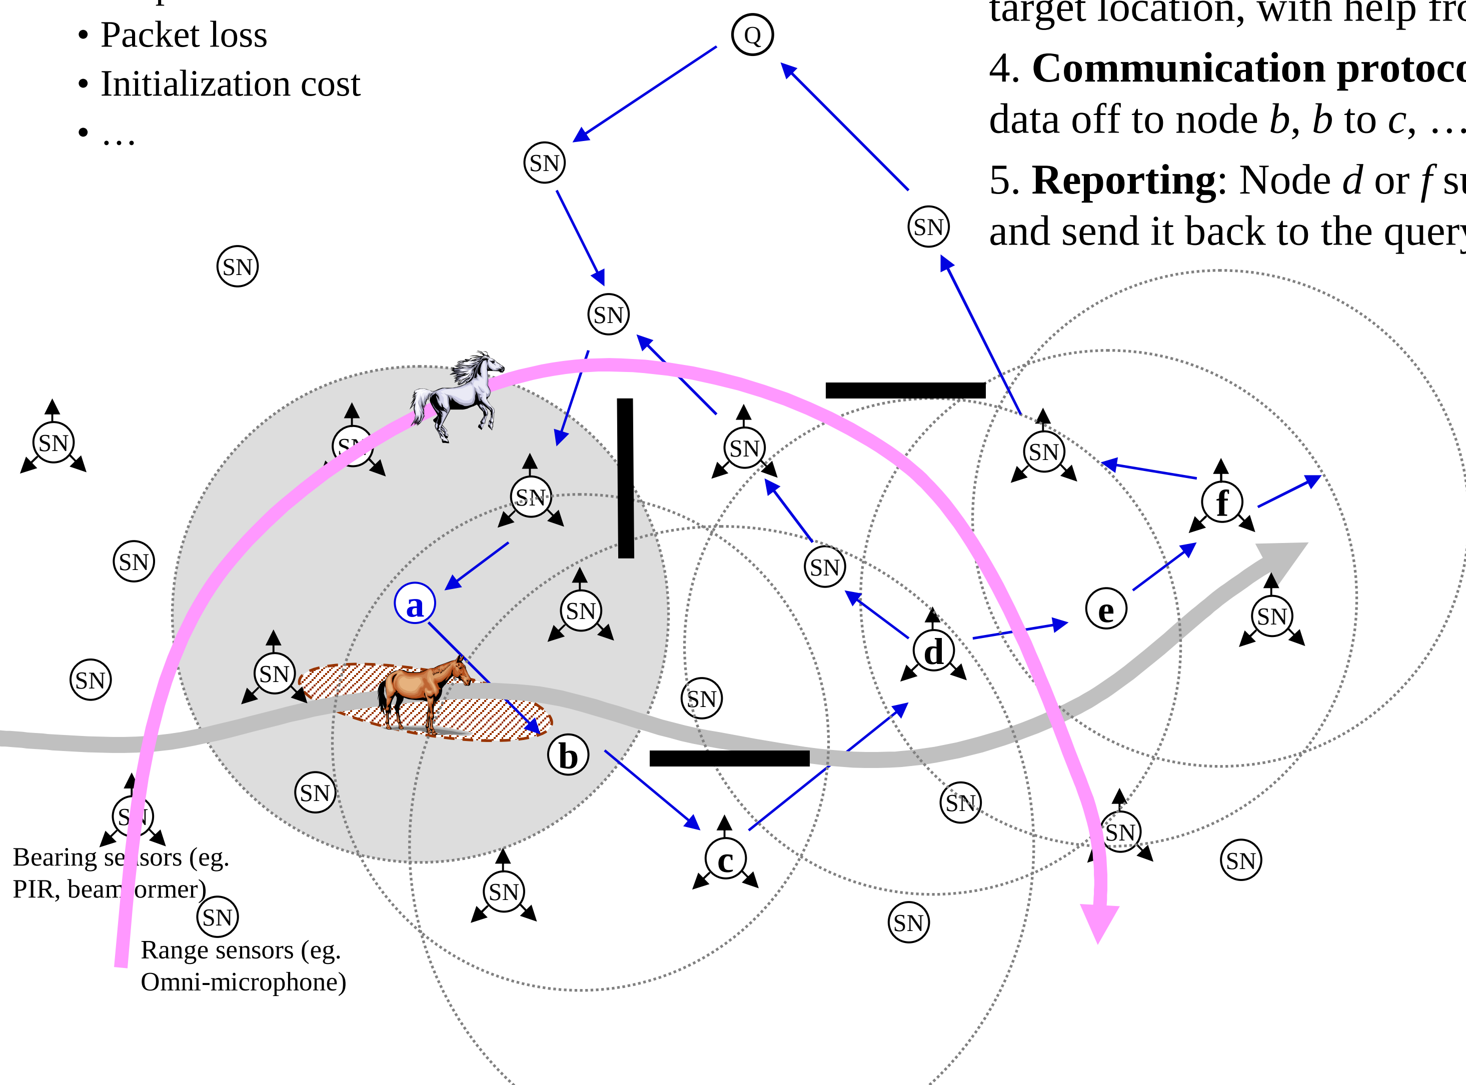

What if there are other (possibly) interfering targets?

What if there are obstacles?

parc
Palo Alto Research Center

Fundamental issues to address

- What do nodes know when they wake up
  - How do nodes discover each other, time synch, location?
- What is routing
  - Address-centric or data-centric? Is TCP/IP appropriate? What are the performance metrics?
- What is information theory for sensor nets?
  - What is the minimum number of sensors for recovering a continuous signal field? Nyquist sampling theorem for sensor net?
  - What are the compression schemes?
- What is data base for sensor net?
  - How is the data stored, organized, and retrieved? Is there a SQL interface?
- How is it connected to the Internet?
  - What are the edge network protocols? BGP?
  - How will a user browse the data? Google™ for physical world?
- How is it programmed
  - Is there an OO programming model for embedded systems?
  - What are the abstraction layers? What is the “sensor net stack”?

parc
Palo Alto Research Center

What is routing in a sensor net?

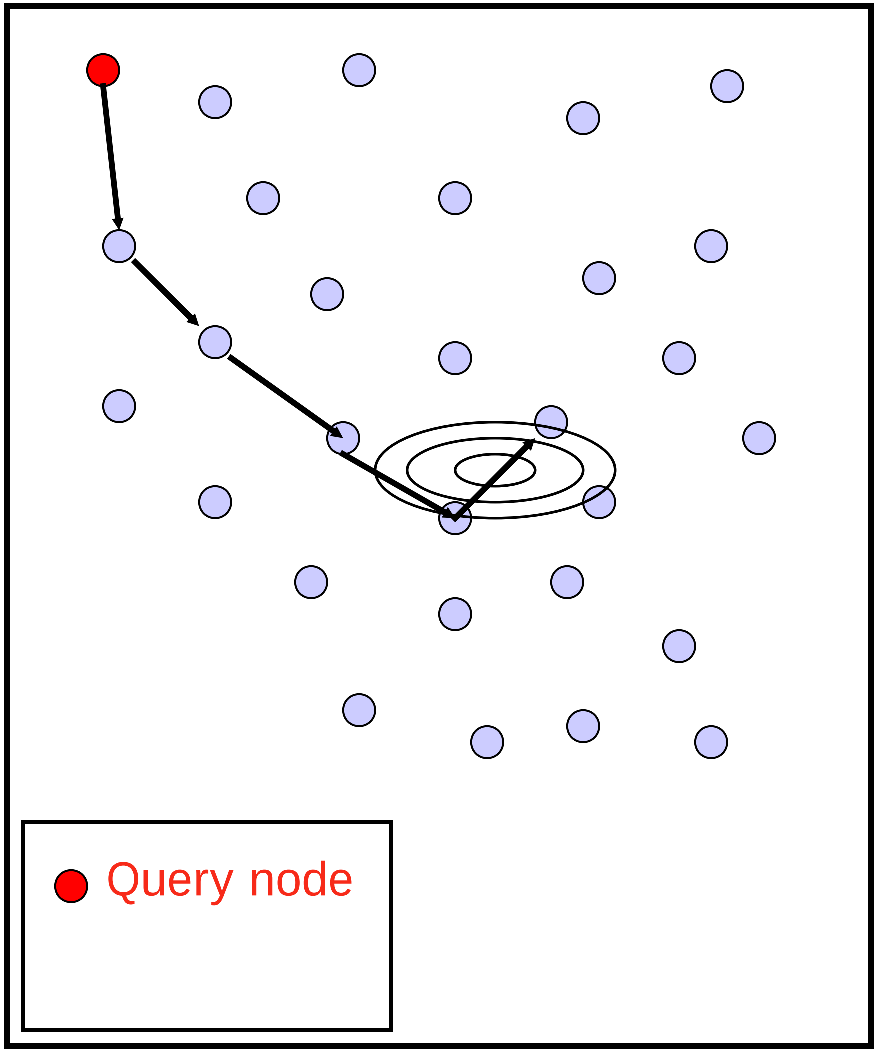

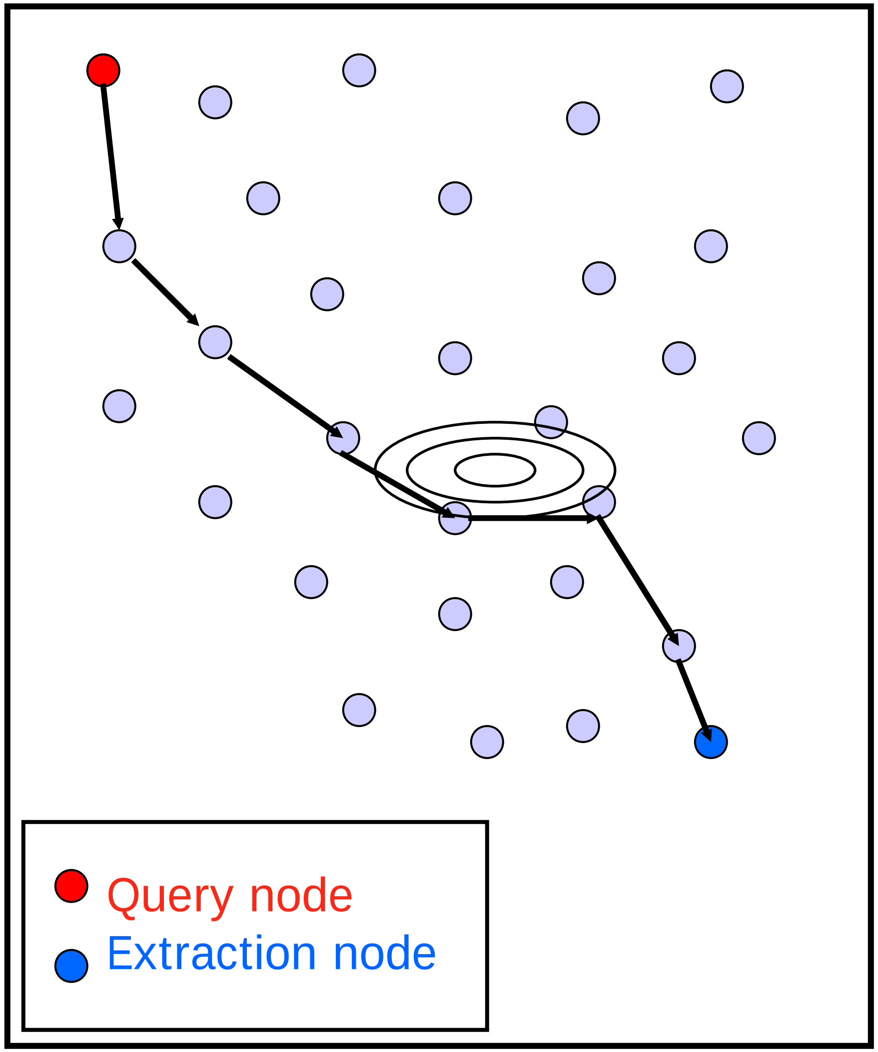

What is unique about sensor networks?

- Unique characteristics
  - Coverage
    - »Distance/area covered, number of events, number of active queries
  - Spatial diversity
    - »Dense spatial sampling, multi-aspect sensing, multiple sensing modalities
  - Survivability
    - »Robust against node/link failures
  - Ubiquity
    - »Quick/flexible deployment, ubiquitous access, info timeliness
- Particularly suited for detecting, classifying, tracking
  - Non-local spatio-temporal events/objects
    - »**Simultaneous dense spatial sampling** to identify and track large, spatially extended event
    - »**Continuous spatio-temporal sampling** to track moving objects
  - Low-observable events
    - »**Distributed information aggregation**
    - »**Non-local information validation**

parc
Palo Alto Research Center

Sensor network research is largely driven by applications …

parc
Palo Alto Research Center

Habitat monitoring: www.greatduckisland.net

Great Duck Island, 10 miles off the coast of Maine:

Remote wireless sensors are being used to find out more about birds in their natural habitat.

Petrel: Rarely seen by birdwatchers

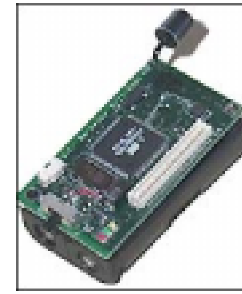

Wireless biological sensors placed in nests

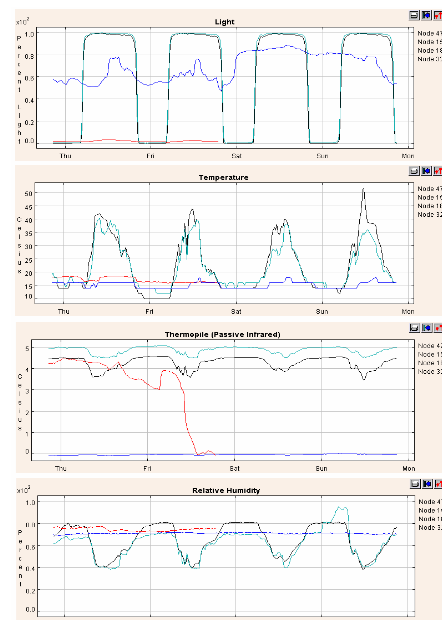

“It will enable us to study ecosystems at a level that has not been conceived.”
Steve Katona, College of the Atlantic biologist and president

BBC News, October 10, 2002

Monitoring Plants at Huntington Botanical Garden

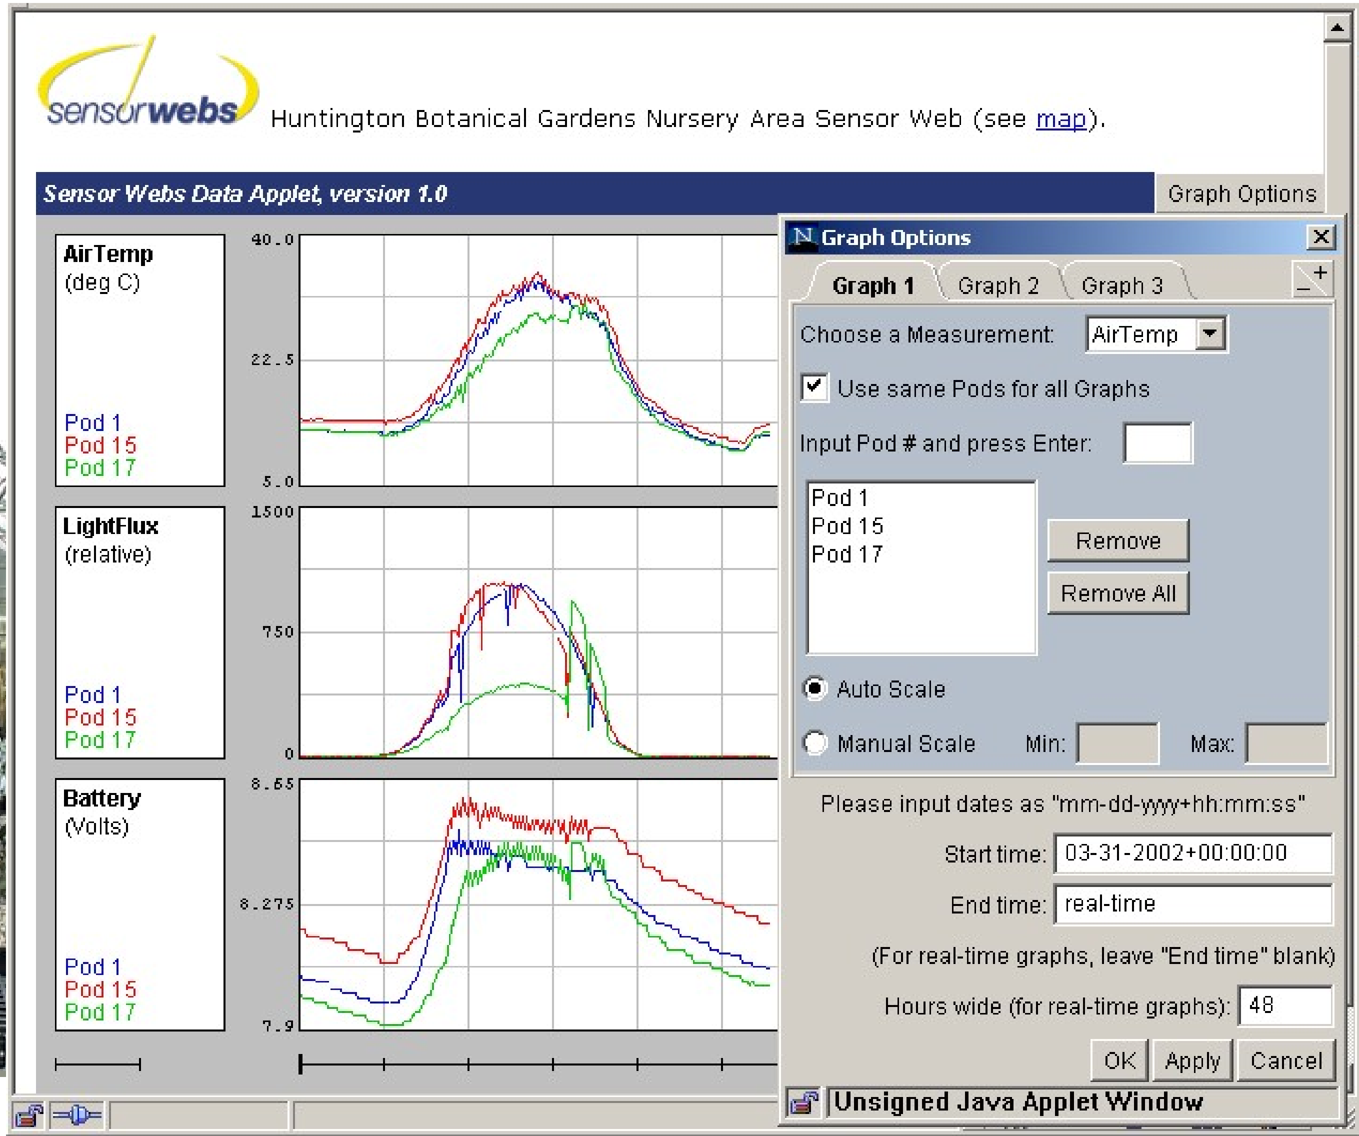

A Chemical Plume Tracking Scenario

- Ad hoc, just-in-time deployment of sensor nets for emergency response

- The Valley Authority just declared a region-wide emergency: A large-scale chemical gas leak has been detected at a plant 20 minutes ago
- National Guard has been activated to evacuate nearby towns, and to close roads/bridges
- To get a real-time situational assessment of the extent and movement of the gas release and inform the evacuation, the CS428 *SWAT Team* is called in
- Three small UAVs are immediately launched from an open field 15 miles south of the attack site, each equipped with 1000 tiny chemical sensing nodes
- Upon flying over the vicinity of the attack site, the sensor nodes are released
- The nodes self-organize into an ad hoc network, while airborne, and relay the tracking result back to the emergency response command center
  - –Where is the plume, how big, how fast, which direction?

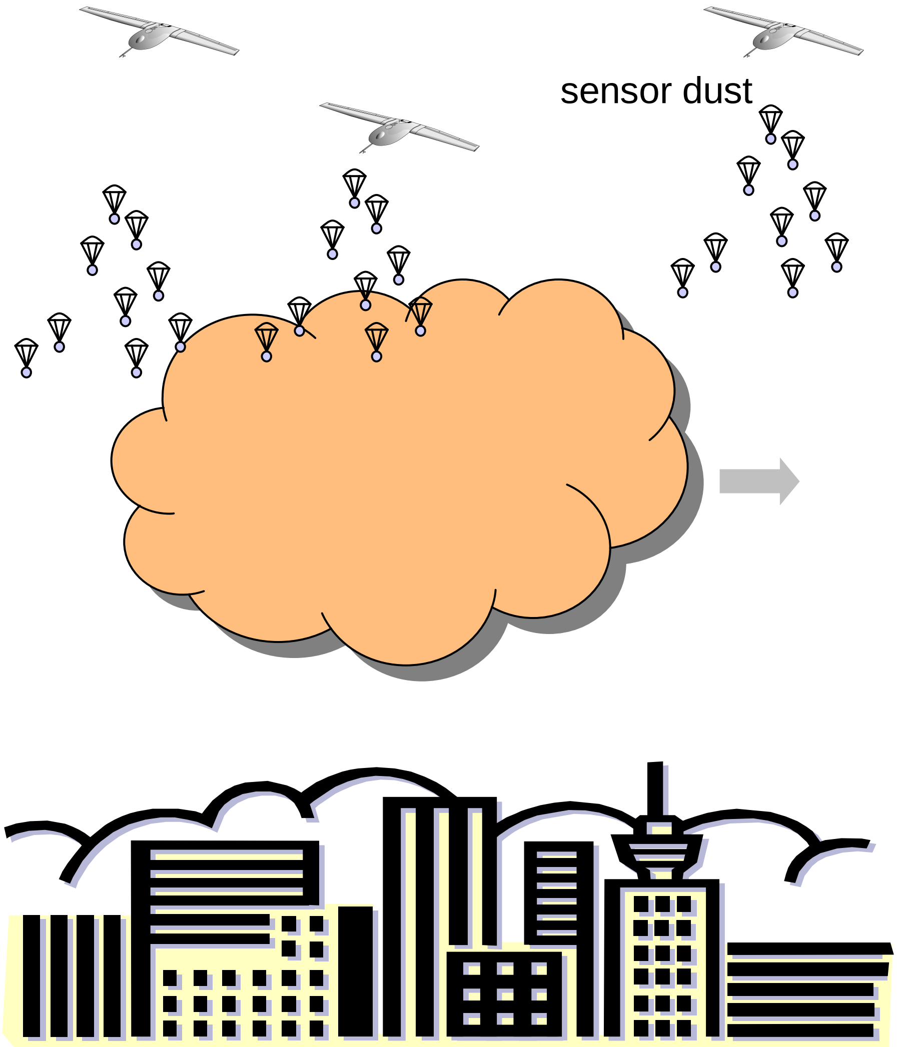

Palo Alto Research Center

Some of the possible project ideas

- Counting distinct objects
- Routing around “forbidden” zones
- Mobile sensor to improve object localization
- Query optimization for multi-object tracking
- CCW relation tracking
- “Am_I_Surrounded” tracking

Assumptions

- Each node has a number of on-board sensors (in case there is only one sensor, we use "sensor" and "node" interchangeably)
- Each node can communicate wirelessly with other nodes within a fixed radius *R* (larger than the mean inter-node distance)
- Nodes are time-synchronized to a global clock
- Targets are point sources of signals; target signals propagate isotropically in the physical space, and attenuate as a monotonically decreasing function of the distance from the source
- Each sensor has a finite sensing range, determined by a fixed minimal amplitude a sensor can sense. Signals of two targets sum at a sensor. Here we assume amplitude sensing (e.g. to infer distance to the target). DOA or restricted-view sensing (e.g., camera) may provide new constraints.
- Onboard battery power is the main limiting factor, as well as network bandwidth and latency

Routing around "forbidden" zones

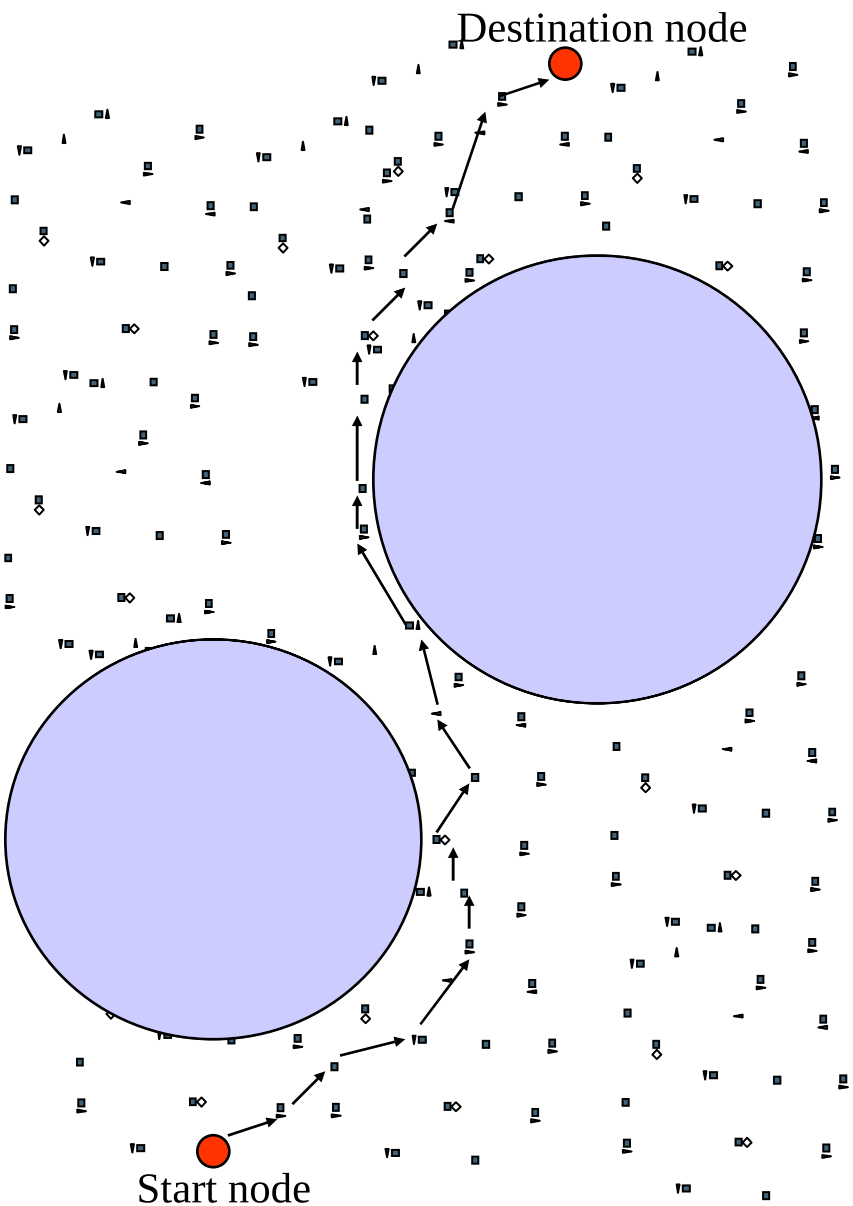

Problem spec

- Each node knows locations of nodes within RF range and that of destination node
- There are two forbidden zones where nodes are not available for sensing or comm.
  - Stationary zones (e.g., holes in the net)
  - Moving zones (e.g., quiet zones to evade enemy eavesdropping)

Task

- Discover routing path with minimal # hops
- Maintain path, when the forbidden zones move, requiring
  - minimal repair to the current path
- **Bonus point**: Suppose data along the path is aggregated for locating a stationary target in the middle of the field. Do the same as above, plus optimizing for localization accuracy as well

parc Palo Alto Research Center

Mobile sensor to improve object localization

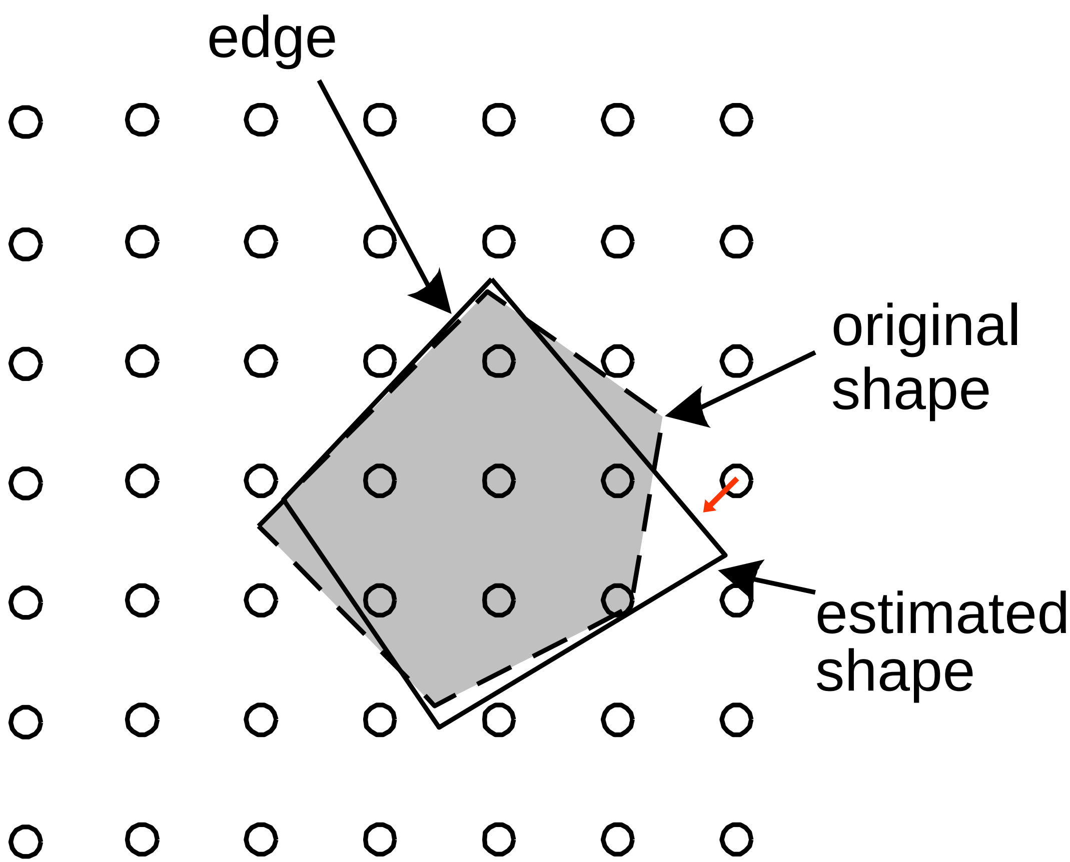

<u>Problem spec</u>

- Sensors track boundary of an polygon with *n* sides
- Each sensor can detect whether it is covered by the object or not, and can move in physical space

<u>Task</u>

- Determine which subset of sensors to move, to maximize estimation improvement while minimizing the total distance traveled
- **Bonus point**: Suppose only a fraction of the sensors can move, determine the optimal fraction for detecting a given shape

parc
Palo Alto Research Center

Query optimization for multi-target tracking

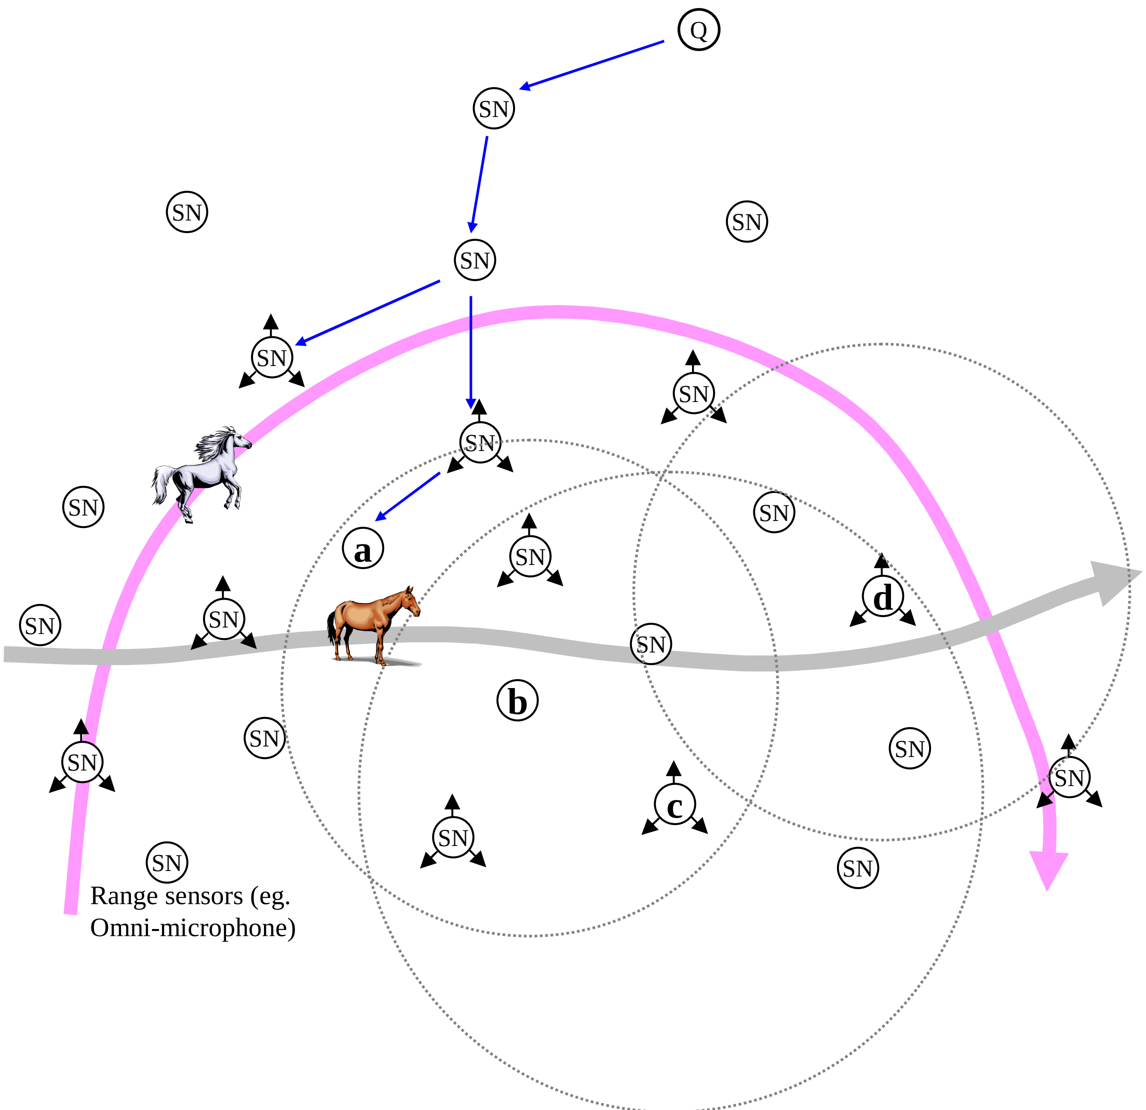

Problem spec

- Assume two or more targets
- Query how many targets are present

Task

- Route query towards each of the targets, optimizing for the total number of hops
- Determine when to split a query (targets are too far apart?), and when to merge
- Consider both stationary and dynamic target cases

Target Localization for CCW Relations

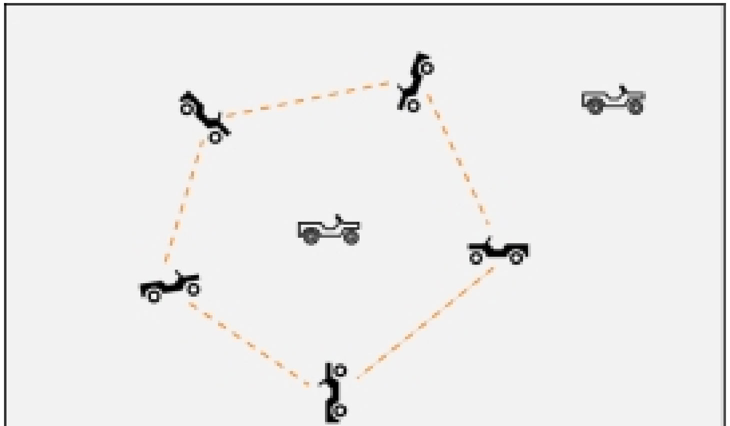

“Am_I_Surrounded?”

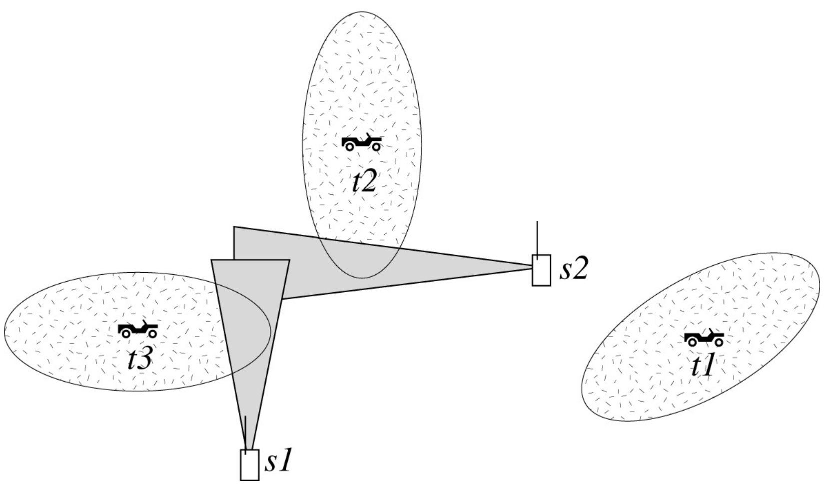

<u>Problem Spec</u>

Many complex spatial relations among targets can be understood by sensing CCW relations

<u>Task</u>

The task is to experiment with sensor selection strategies for optimally sensing CCW relations.

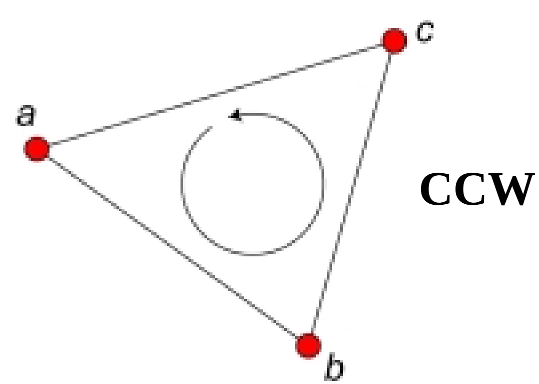

parc
Palo Alto Research Center

Tracking the “Am_I_Surrounded” Relation

Problem Spec

We want to maintain a **cage** of three black vehicles containing the white vehicle

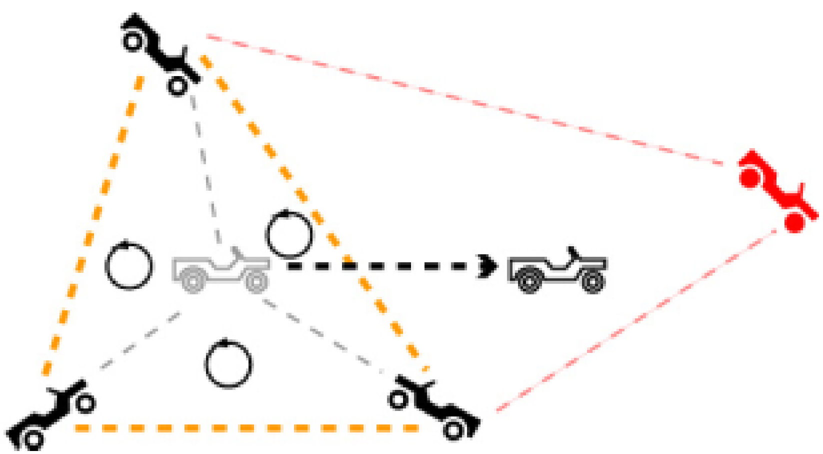

Task

When the white vehicle escapes its cage, we want to search for a new red vehicle completing the new cage.

parc Palo Alto Research Center

Preview of the classes

- Week 1: Class organization; SN introduction and applications
- Week 2: Localization and tracking
- Week 3: Networking I; class project discussion
- Week 4: Networking II
- Week 5: Network initialization and services
- Week 6: Information management I
- Week 7: Information management II
- Week 8: SW/HW architecture; resource constraints
- Week 9: Localization and tracking II
- Week 10: Applications; Final project reports