Issues and Approaches to Management of Sensor Networks

Kang G. Shin
Real-Time Computing Laboratory
Department of EECS
University of Michigan

Outline of Talk

- Generics, issues, and approaches of sensor networks
  - Hardware
  - Communication
  - Software
- Examples of research on sensor networks at:
  - University of Michigan
  - UC Berkeley
  - University of Virgina
  - UCLA
  - ….

Characteristics of Sensor Networks

- Large # of small, resource-limited sensor nodes, operating *in aggregate*
- Usually battery-powered, hence *energy-constrained*
- Wide range of sensing capabilities
  - Temperature, light, sound, magnetic fields, motion, vision
- Low-power wireless networking
- Unattended, inaccessible, prolonged deployment
- Requires *in-network processing*
- *Time-varying* functions/roles

=> Must be self-organized, self-maintaining and programmed *in situ* to operate at very low duty cycle

Uses of Sensor Networks

- Commercial
  - Manufacturing plant monitoring, integrated robotics, vehicle/object tracking, security/safety monitoring, inventory control and manuals/instructions (RFIDs), etc.
- Research
  - Environmental monitoring (habitat and agricultural studies)
- Military
  - Tracking, intrusion detection
- Homeland security
  - Surveillance of public/critical infrastructures such as buildings, bridges, utility distribution and water supply systems

Typical Sensor Node: X-Bow Mica Mote

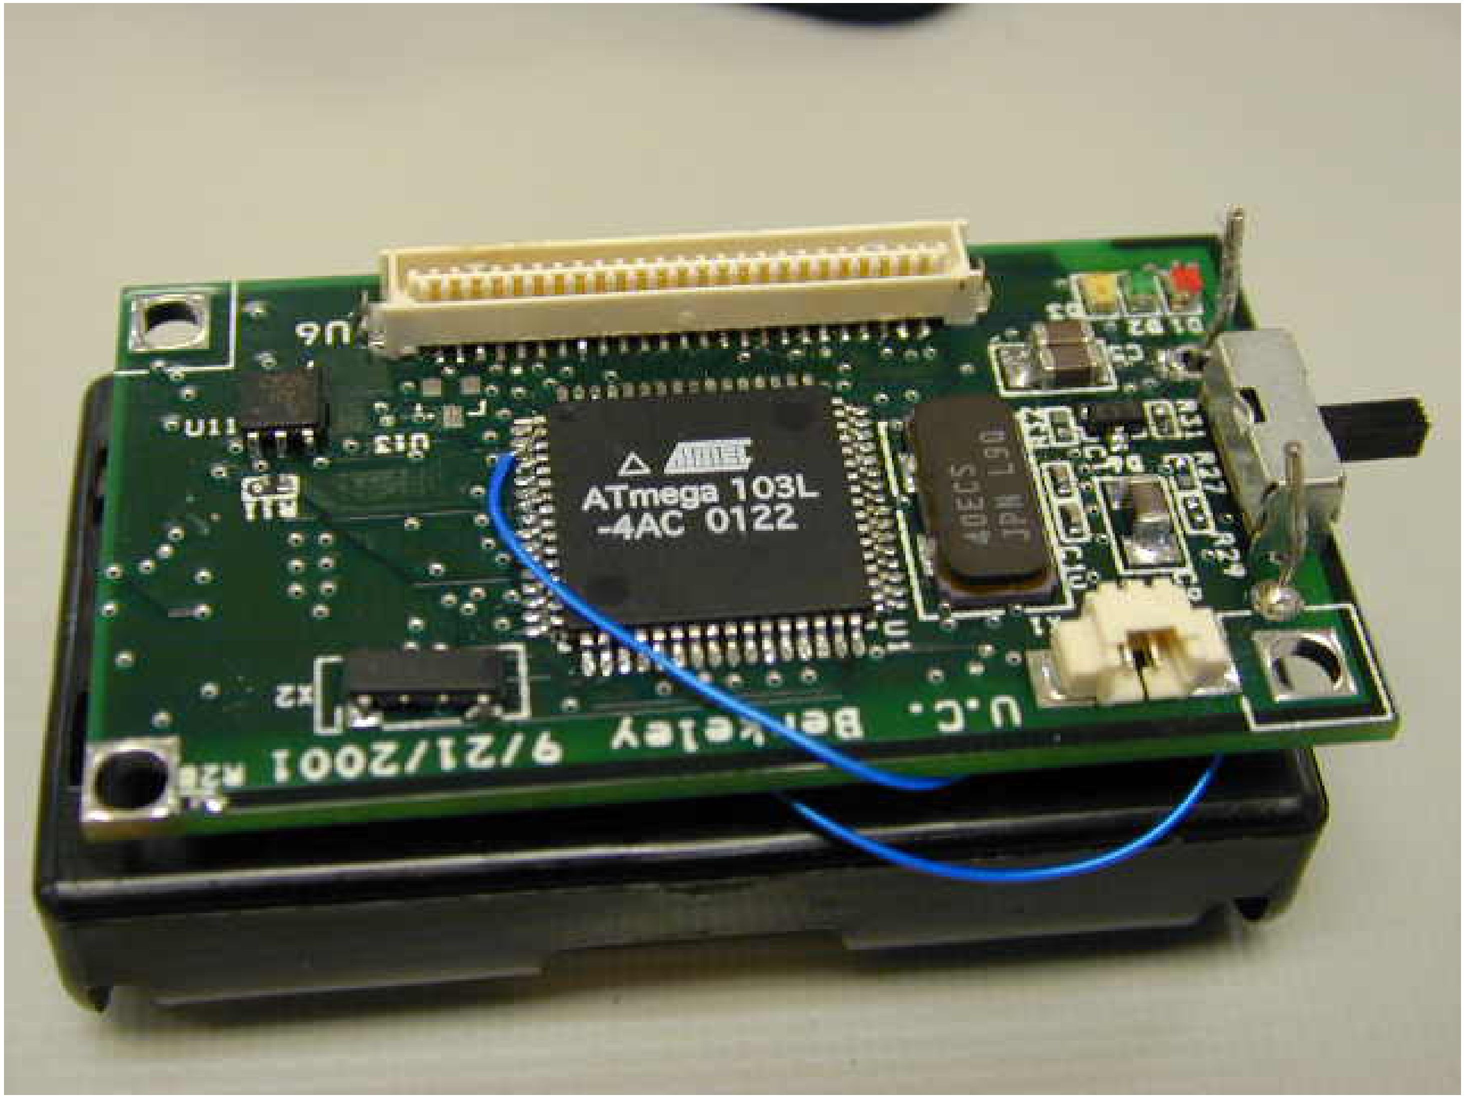

MICA Architecture

- Atmel ATMEGA103
  - 4 MHz 8-bit CPU
  - 128KB Instruction Memory
  - 4KB RAM
- 512 KB flash (`AT45DB041B`)
  - SPI interface
  - 1-4 μj/bit r/w
- RFM TR1000 radio
  - 50 kb/s – ASK
  - Focused hardware acceleration
- 6 ADC channels
- Unique serial IDs
- Network programming
- 51-pin connector
  - Analog compare + interrupts
- TinyOS tool chain

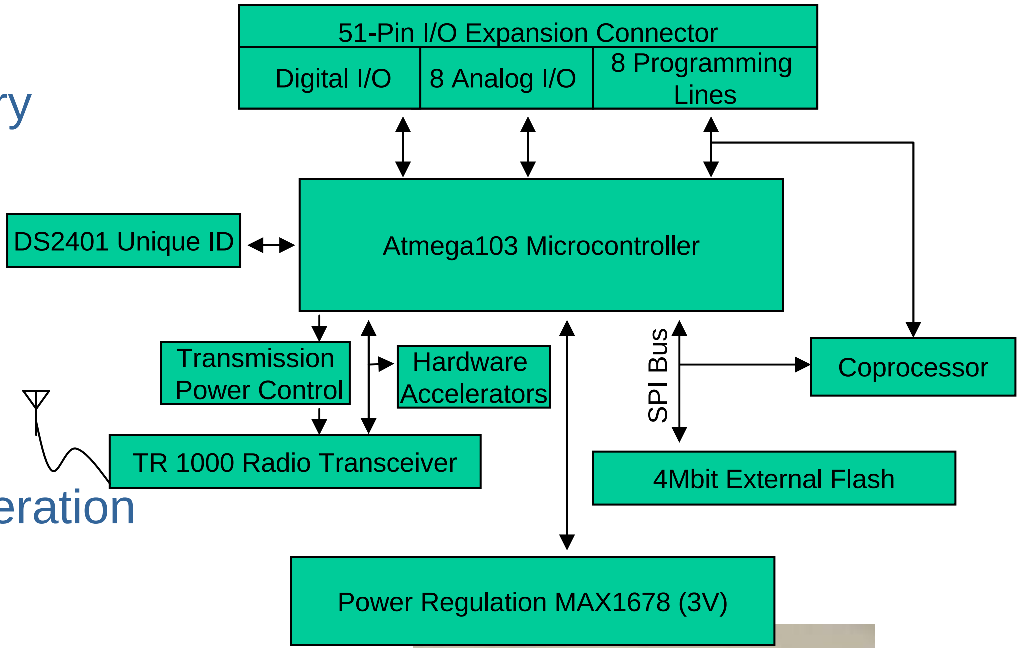

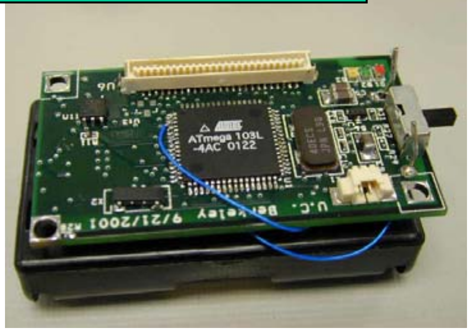

2xAA form factor

Cost-effective power source

Sensor Board Device Placement

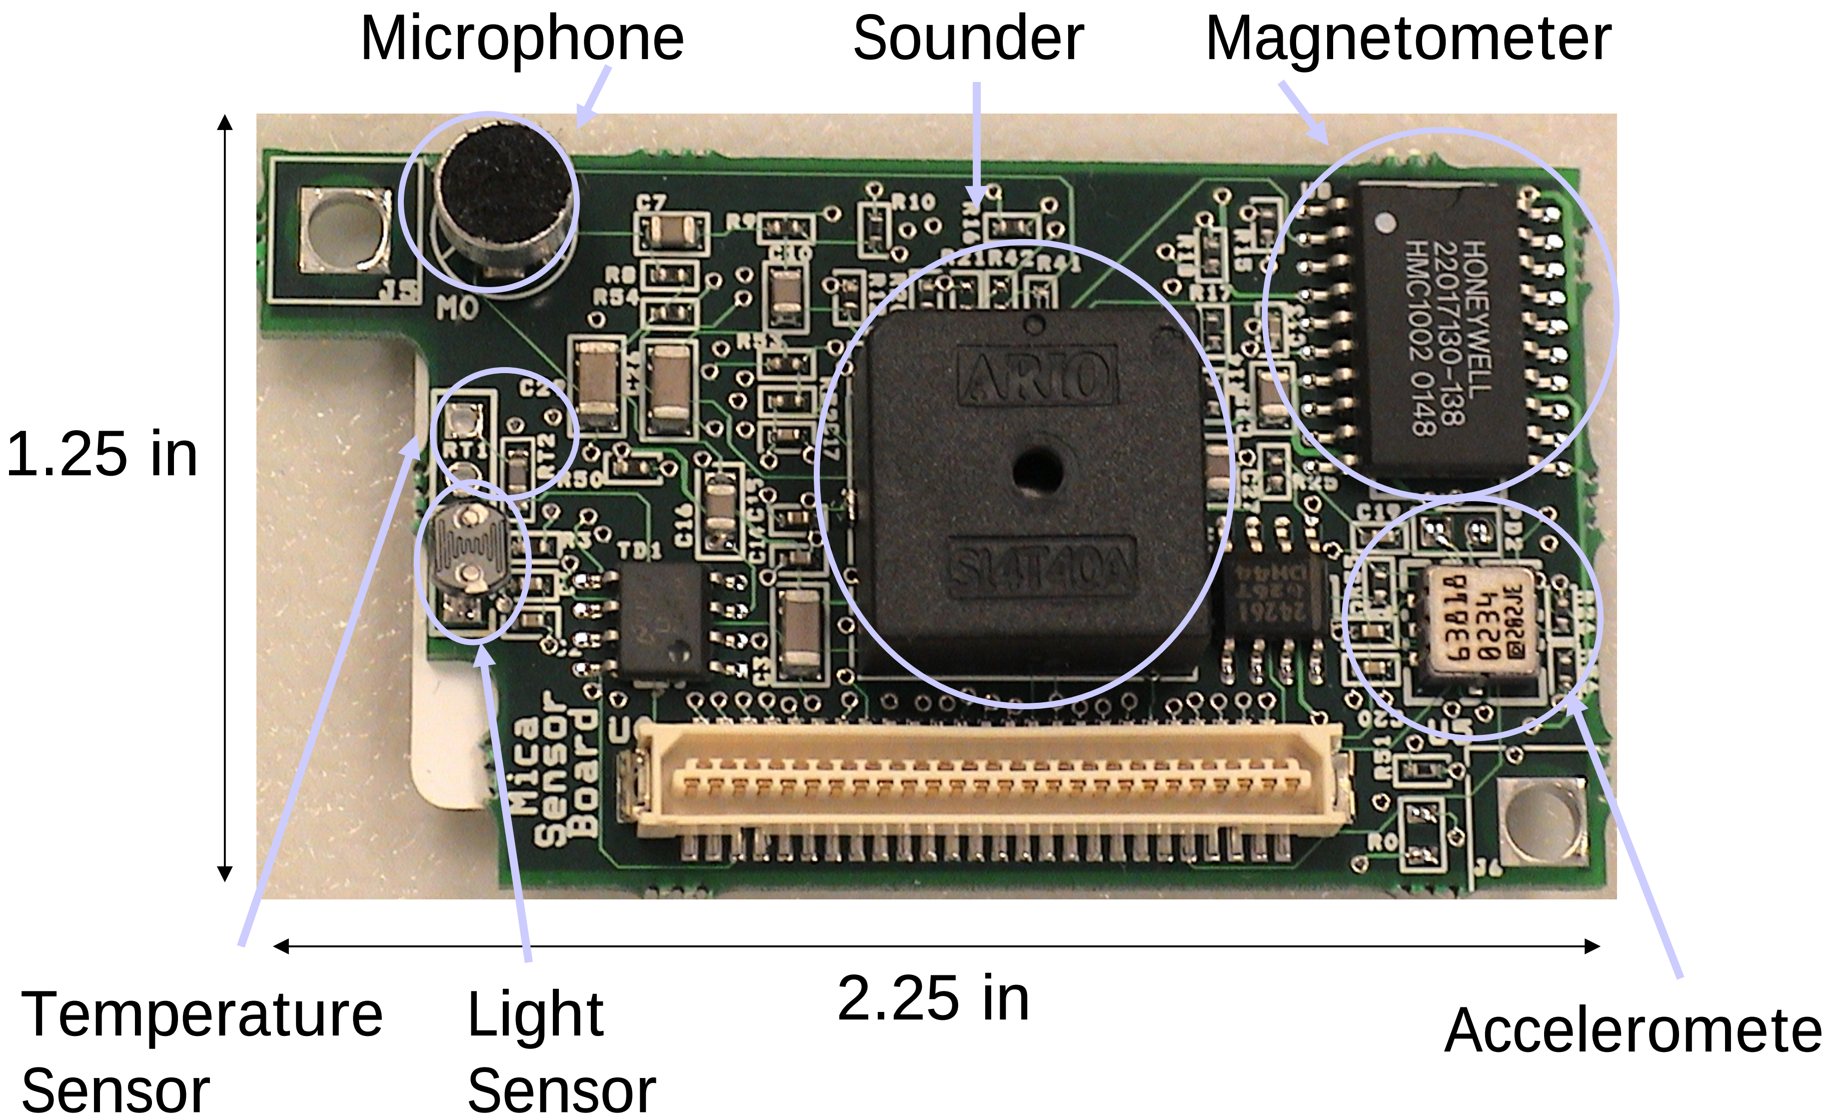

Mica Mote with Sensor Board

Research Areas/Issues

- Sensing and architecture
  - Sensor hardware design (MEMS)
  - Signal/Data processing
  - Rich interfaces and simple primitives allowing cross-layer optimization
  - Low-power processor, ADC, radio, communication, encryption
- Resource management (operating system)
  - Limited computational power, memory, code space, electrical power
  - Node computation & communication, and their scheduling
- Networking and distributed services
  - Medium Access Control & routing
  - Clock synchronization, localization, and data aggregation
- Programming
  - Software component models and middleware
  - Describe global behavior, synthesize local rules that have correct, predictable global behavior
- Applications
  - Long-lived, self-maintaining, dense instrumentation of previously unobservable phenomena
  - interaction with a computational environment

Networking

- Medium Access Control (MAC)
  - Main issues with wireless communication
    - Collisions
    - Limited range
      - Hidden node/terminal problem
    - Transmission errors
  - Motes use CSMA (Carrier Sense Multiple Access)
    - Cannot send and receive at the same time
    - Cannot detect collision
  - Work is being done to create inherently collision-free MAC protocols
    - TDMA in a region; may be closely-coupled with applications
  - … or to reduce the probability of collision
    - Implicit acknowledgements
    - S-MAC – coordinates sleep cycles to *save energy* and *avoid collisions*
  - Non-Mote systems (esp. simulations and more powerful sensors) use 802.11 MAC or its variations: Stargate and Stareast

Super Node I: Stargate Board

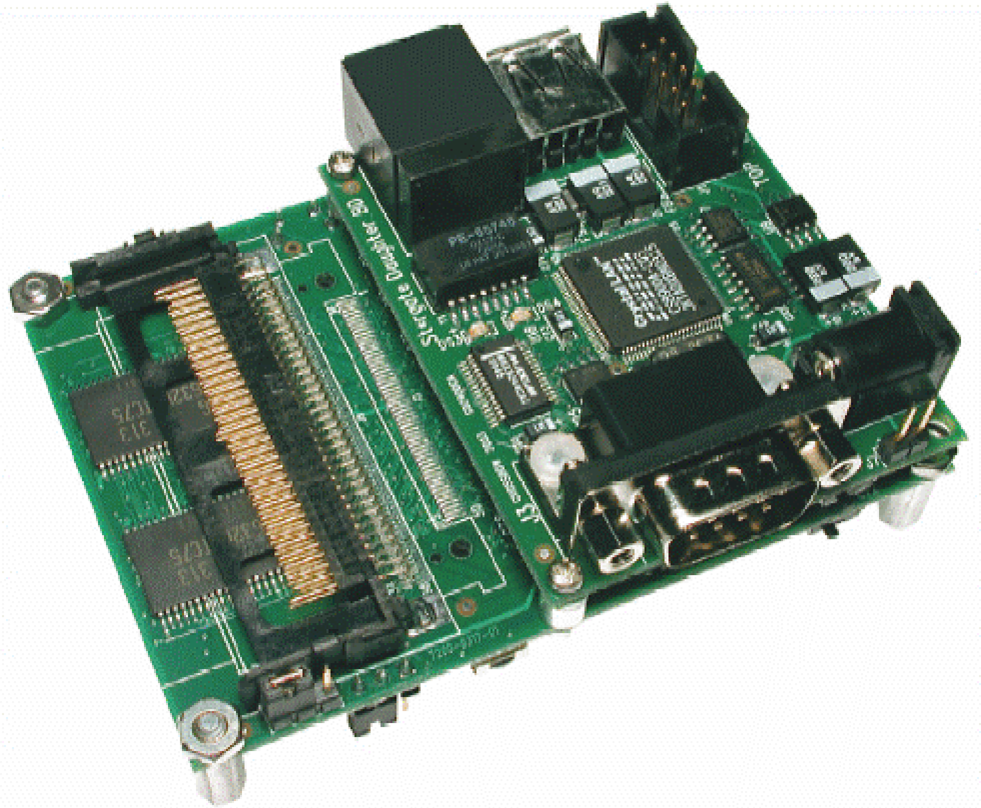

Super node II: Stareast Boards

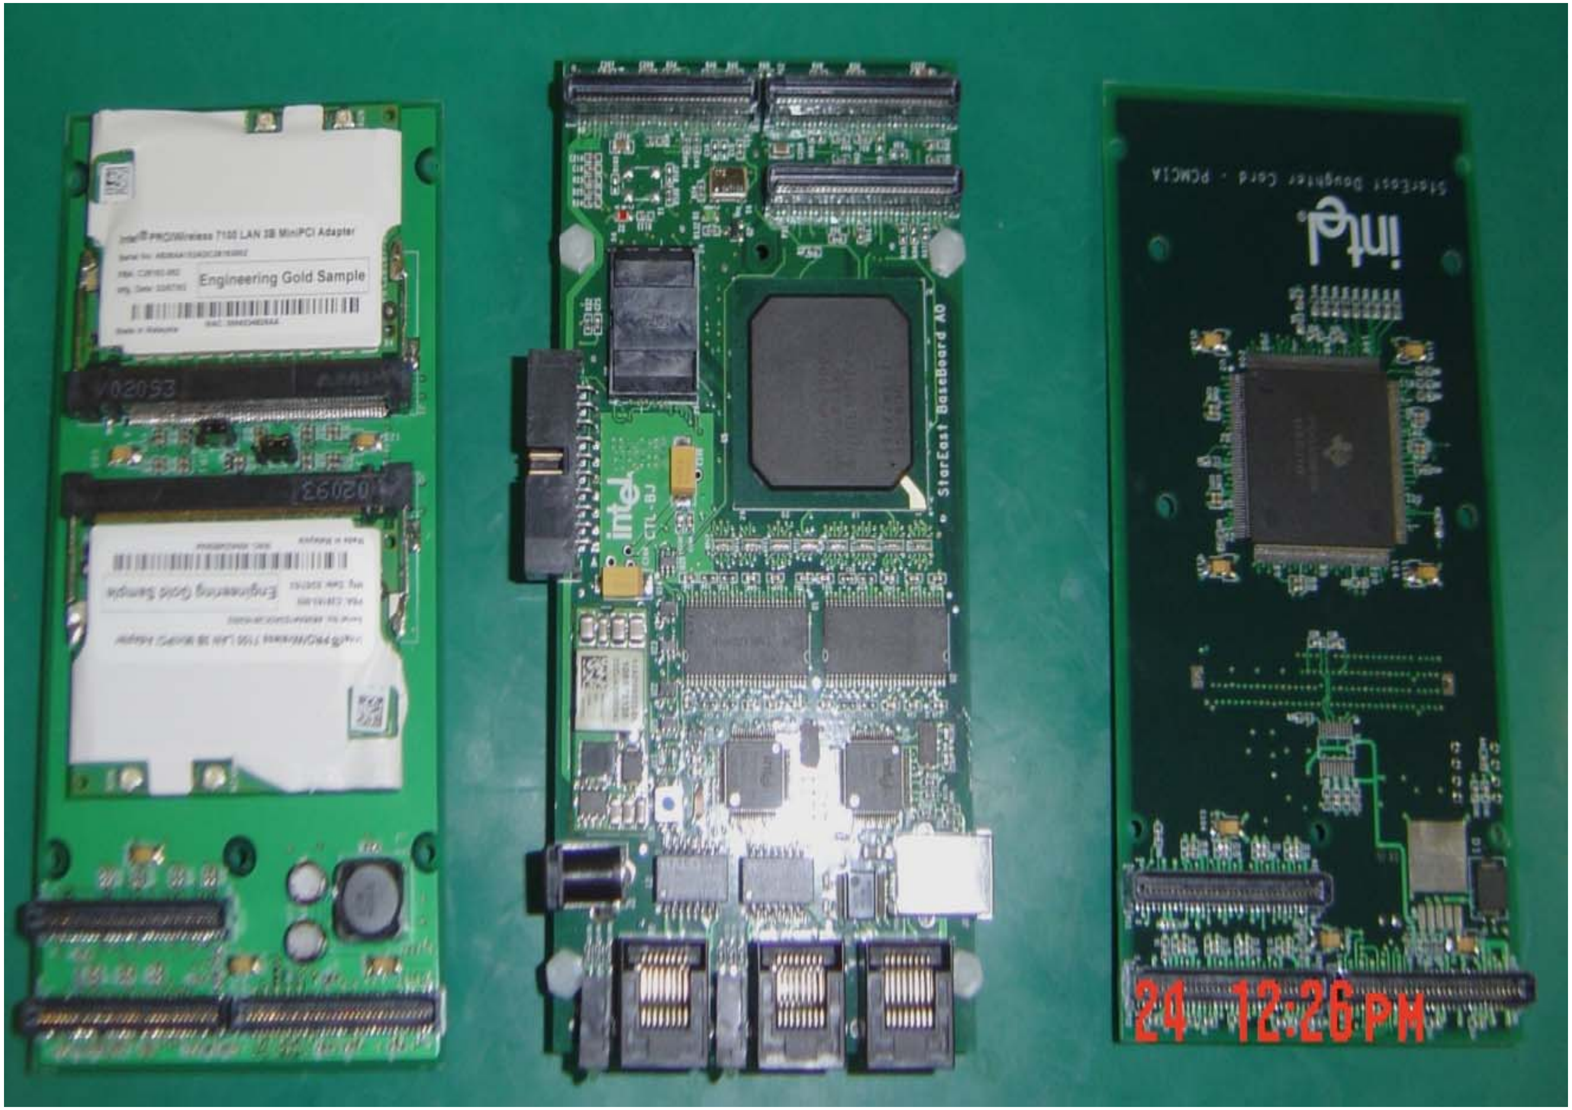

Routing Protocols

- Spanning tree within a cluster/region
- Geographic routing
  - Route messages to a specific location
  - Each node knows its location
  - No routing tables maintained
- Cluster-based routing
  - Use simple table-based routing protocols to route to cluster head (e.g., dynamic source routing, ad-hoc on demand distance vector routing)
  - Use higher-level protocol (e.g., geographic) to route between cluster heads
- Landmark routing
  - Similar to cluster-based routing, but without the cluster formation overhead
  - Messages are routed to known landmarks, from which they are routed to their final destination

Routing Protocols, cont'd

- Gradient Routing
  - Requires only local information at each node
  - An "interest" is propagated outward by a sink node
  - Each node receiving the interest remembers it and passes it along
  - Different topologies arise due to forwarding policies
  - Data from a source traces back links to the sink
  - Preferred data paths may be reinforced
    - Lowest energy
    - Shortest path
    - Least latency

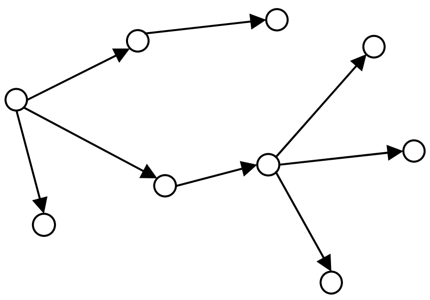

Spanning Tree

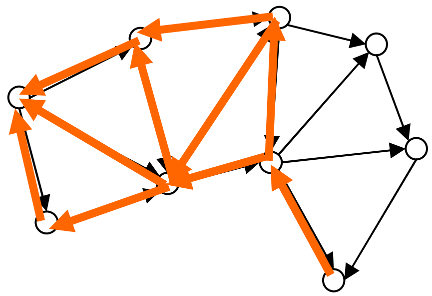

Directed Acyclic Graph

Sensor Network Programming

- Embedded systems
  - Lightweight OS, e.g., tinyOS, EMERALDS
  - OS and application software are compiled and linked together, then downloaded to the node
  - Programmed once and deployed
  - Some work is being done on network reprogramming
    - Expensive in terms of energy
    - Takes a node out of service while reprogramming
    - Scalability issues
- Software structured using component models
  - Support modularity
  - Only essential components are compiled into the system
  - Easy to upgrade/replace components *during development*

Example University Research Efforts

- University of Michigan
- UC Berkeley
- University of Virginia
- UCLA
- ….

Efforts at UMICH

- DARPA:
  - SMILE: Service Models for Integration of reaL-time Embedded systems
  - Security Tradeoffs (with UMass and ASU)
- ONR, NRL, NSF, Cisco:
  - LiSP (Lightweight Security Protocol), PIV, DKMP, SyKeeper
- NSF:
  - Lightweight and Flexible Sensor Network Management
- Project personnel: 1 faculty, 1 full-time research scientist, and 9 grad students
- Project URLs: http://kabru.eecs.umich.edu/{smile,security}

Sensor Network Testbed

Sample Projects at UM

- *Adaptive Query Processing (AQP)*
- *Content-aware metadata creation in a heterogeneous mobile environment*
- *Network routing*
- *Distributed location service*
- *Sensor network security*
- Self-management

AQP Middleware

- Provides an abstraction that forms the basis for *service* & *application development* on a platform

  =>Higher-level domain services are implemented as queries and query-triggered functions

- Is based on a data-centric view of networked embedded systems
- Provides basic data access and management
- Is based on a data model that includes type, time, location, and quality parameters

Service Development on Motes

- Sensor database (SensorDB)
- Energy-aware Query Processing
  - Declarative Query Interface to
    - provide transparent adaptation and optimization
  - Energy savings in
    - communication and query processing
- Techniques proposed to increase lifetime
  - Utility/cost in *query allocation* by each coordinator
  - Energy-efficient (i.e., computationally-efficient) *query indexing* at each node

Relational Model for WSNs

- Tuples include sensor readings and associated sensor types, node ID, timestamp, energy balance, etc.
- Append-only and distributed across multiple nodes, thus supporting streamed, distributed data
- Query is *persistent* and *periodically* evaluated
- Queries themselves are treated as data upon which other queries may operate, i.e., recursive query.

Hierarchical Architecture

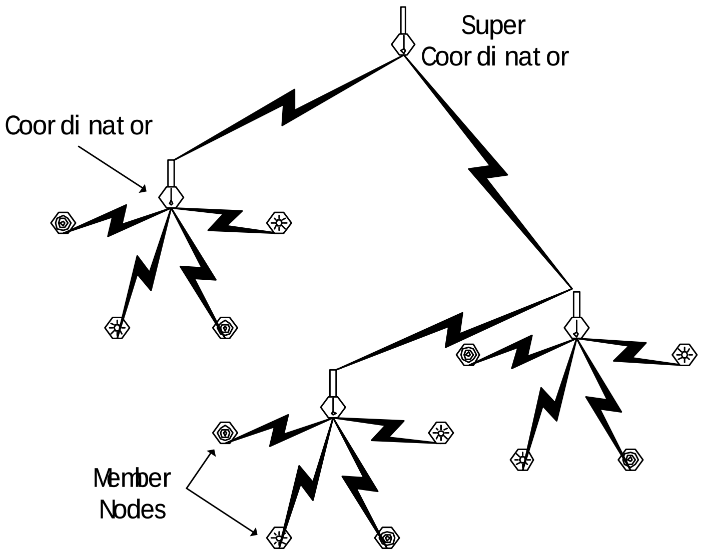

- Roles
  - Super coordinator
  - Coordinator
  - Member
- Cluster
  - Nodes in a small region
  - One-hop communication
  - Redundancy
    - Sensing
    - Communication

Filters and Aggregators

- Filters
  - Run on member nodes
  - Determine the data to be collected and sent to aggregators
- Aggregators
  - Run on coordinator
  - Collect data across space and time, and perform appropriate operations

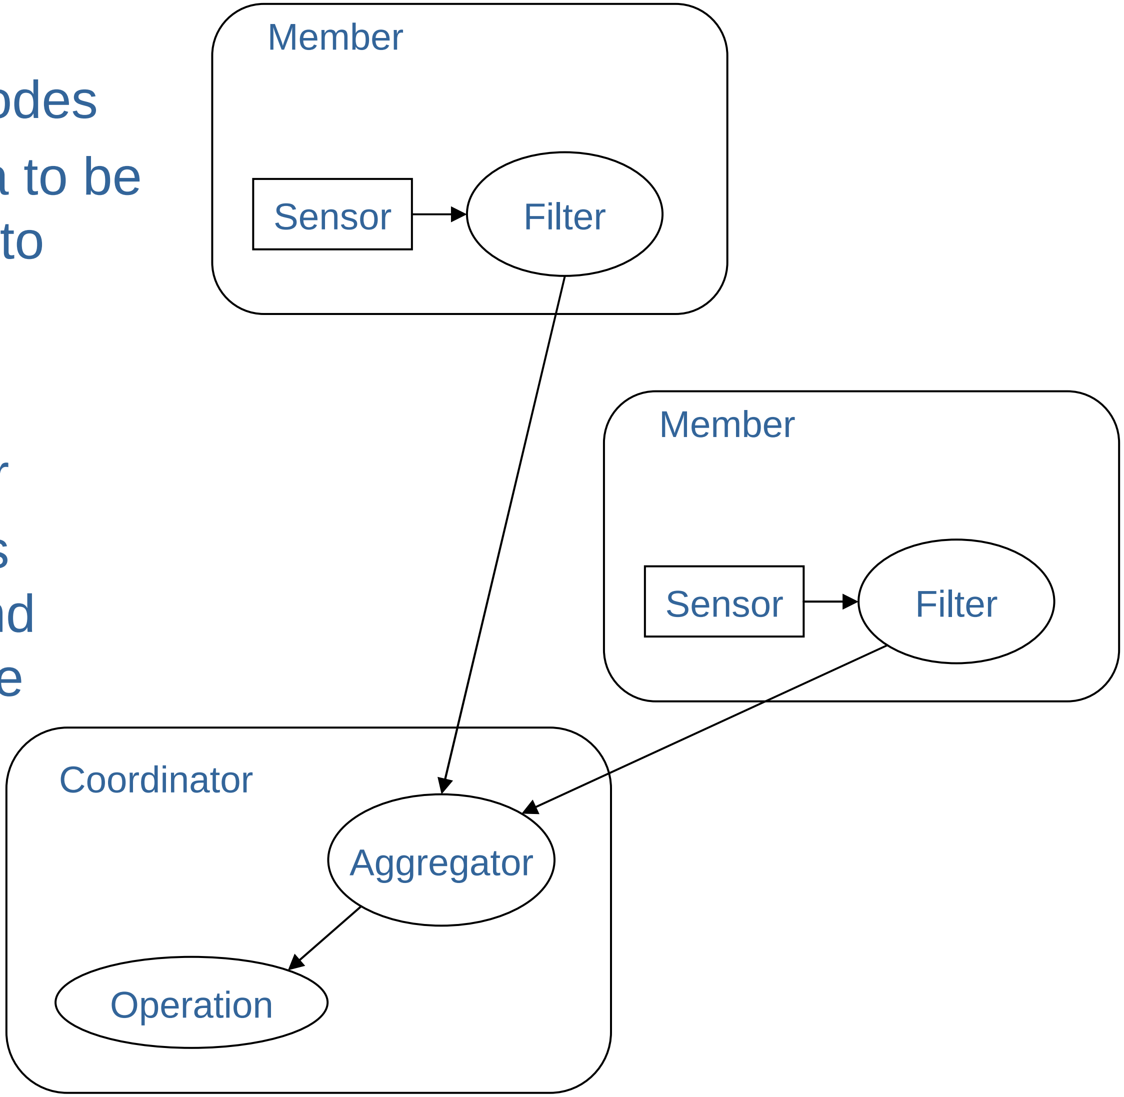

A Simple SQL-like Interface

- A sample query

```
SELECT cluster_id, AVG(mag)
FROM sensors as s
WHERE s.mag > 40
GROUP BY cluster_id
INTERVAL 1sec
DURATION 12min
```

Aggregator

Filter

- Queries that operate on queries
  - Insert, Delete, Update, Select, and Estimate

Energy-aware AQP

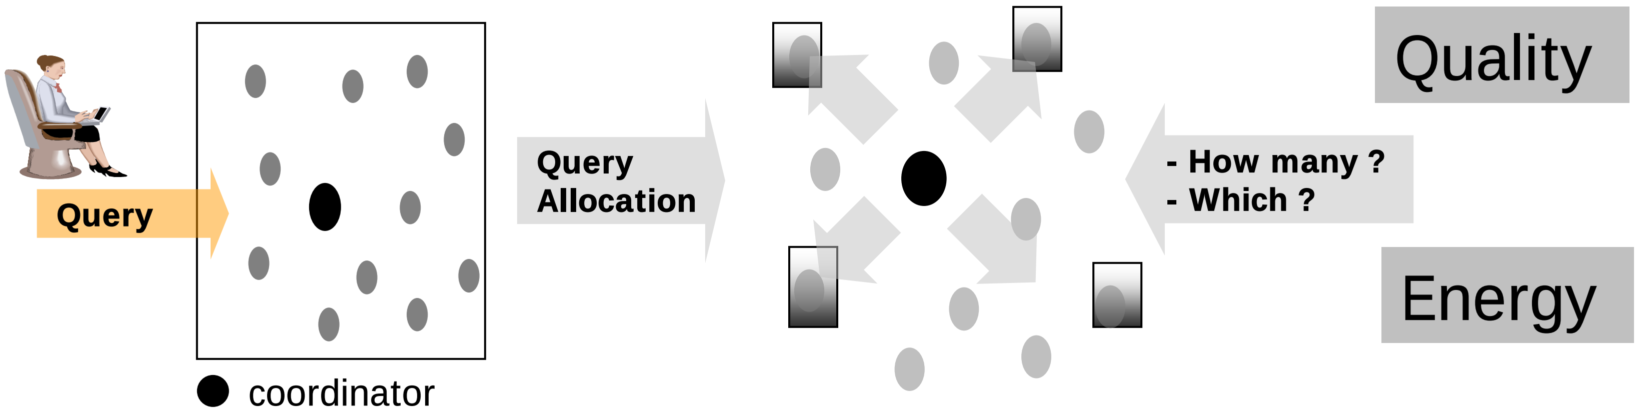

- Distribute workload using utility/cost model
- Given a local cluster of *n* substitutable nodes, adaptively distribute workload to a subset of the nodes
- Utility: accuracy of the query result
  - More nodes give better estimate of sensor value
- Cost
  - Cost associated with selecting and aggregating data
  - Models: balanced, greedy, hybrid

Comparison of Cost Models

Cost Model	Description
Balanced	Cost = 1/(Residual Energy) Balances nodes' energy consumption
Greedy	Cost = Additional Energy Consumption Minimize energy consumption by adding a new query
Hybrid	A combination of Greedy+Balanced Greedy to allocate incoming queries and Balanced to exchange existing query sets

Network Lifetime

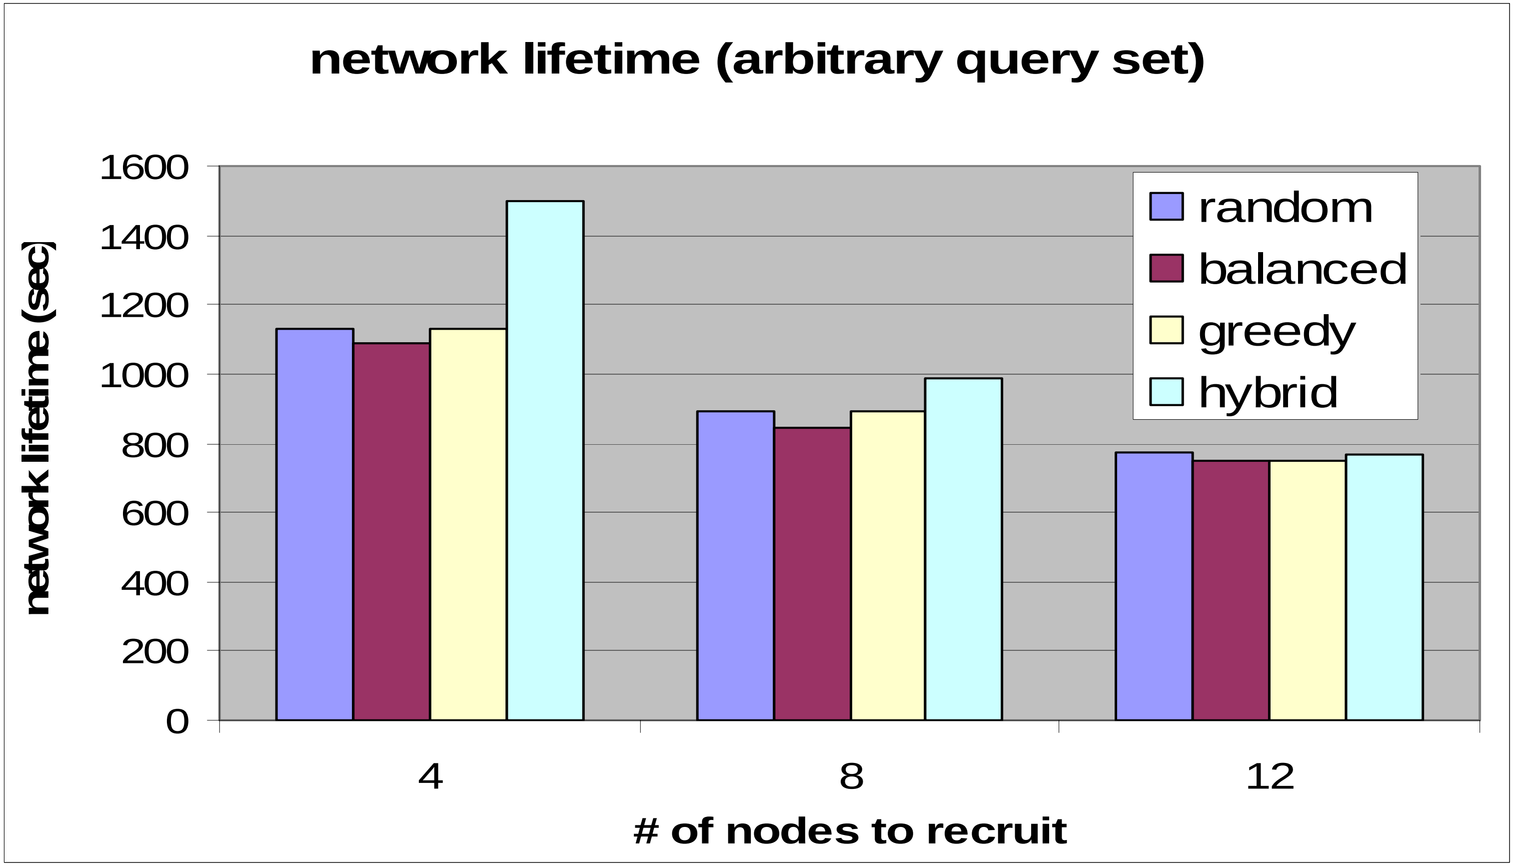

Per-node Residual Energy

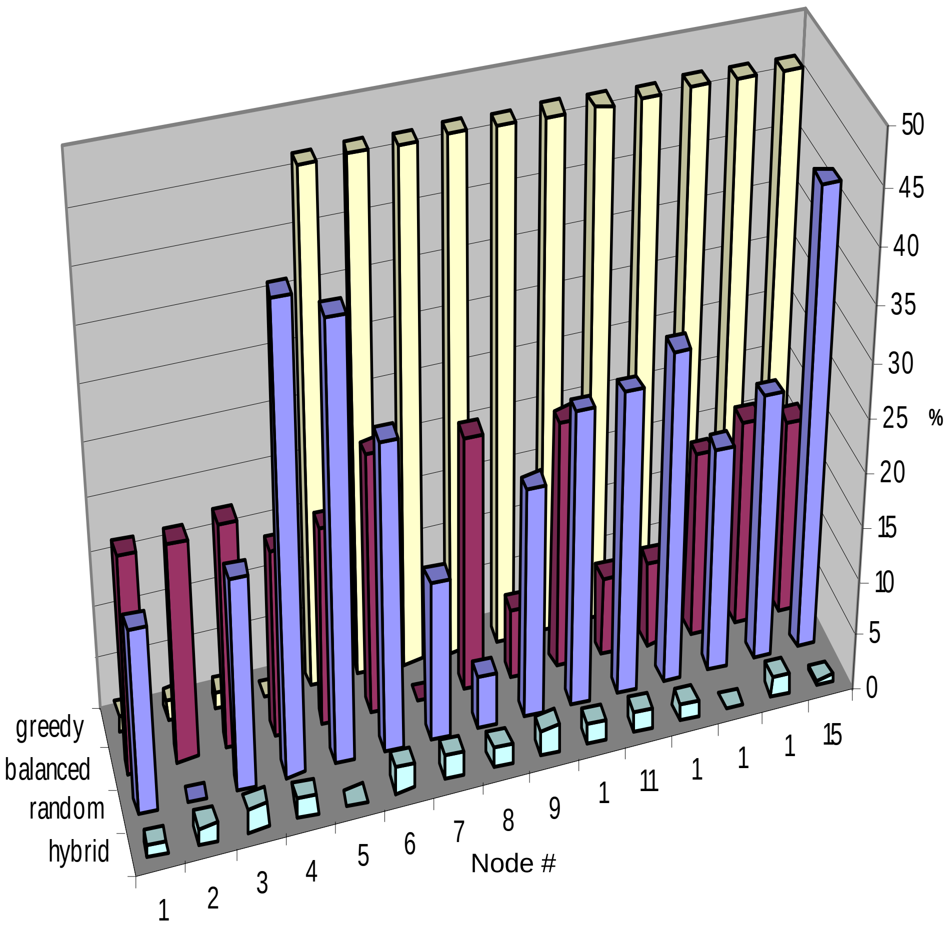

- Selects four nodes per query out of 15 possible
- Remaining energy is measured at the end of network lifetime
- Hybrid model achieves a longer lifetime by distributing power usage more evenly over available nodes

Online Query Optimization

- Why?
  - Queries may be submitted at any time
  - Availability of sensor nodes may change
- Main focus of query optimization is to save energy
  - Maximize sharing of communication and sensing costs among queries

AQP Demonstration

- Implemented support for the “Pursuer-Evader Game” scenario
  - Tracks an enemy evader through a field
  - *Location estimate* is used to pursue the evader
- Steps
  - Energy-aware Coordinator election
  - Energy-aware, geographically-distributed Sentry assignment
  - Detection and aggregation for estimation
    - Adaptive estimation
  - Re-election of Coordinators and Sentries

Content-Aware Metadata Creation and Access

- Wireless handheld devices and sensors are becoming everywhere!
- Amount of digital media data is rapidly increasing and becoming burdensome to manage

=>Difficult to find, edit, share, and reuse media because computers don't understand media content

  - Media is opaque and data-rich and lacks structured representations

So, We

Designed a framework to:

- Collect environmental information from wirelessly-enabled devices
- Associate the collected information, or "metadata," with digital media files
- Metadata facilitates easy search, categorization, and organization of files.

Communication Model

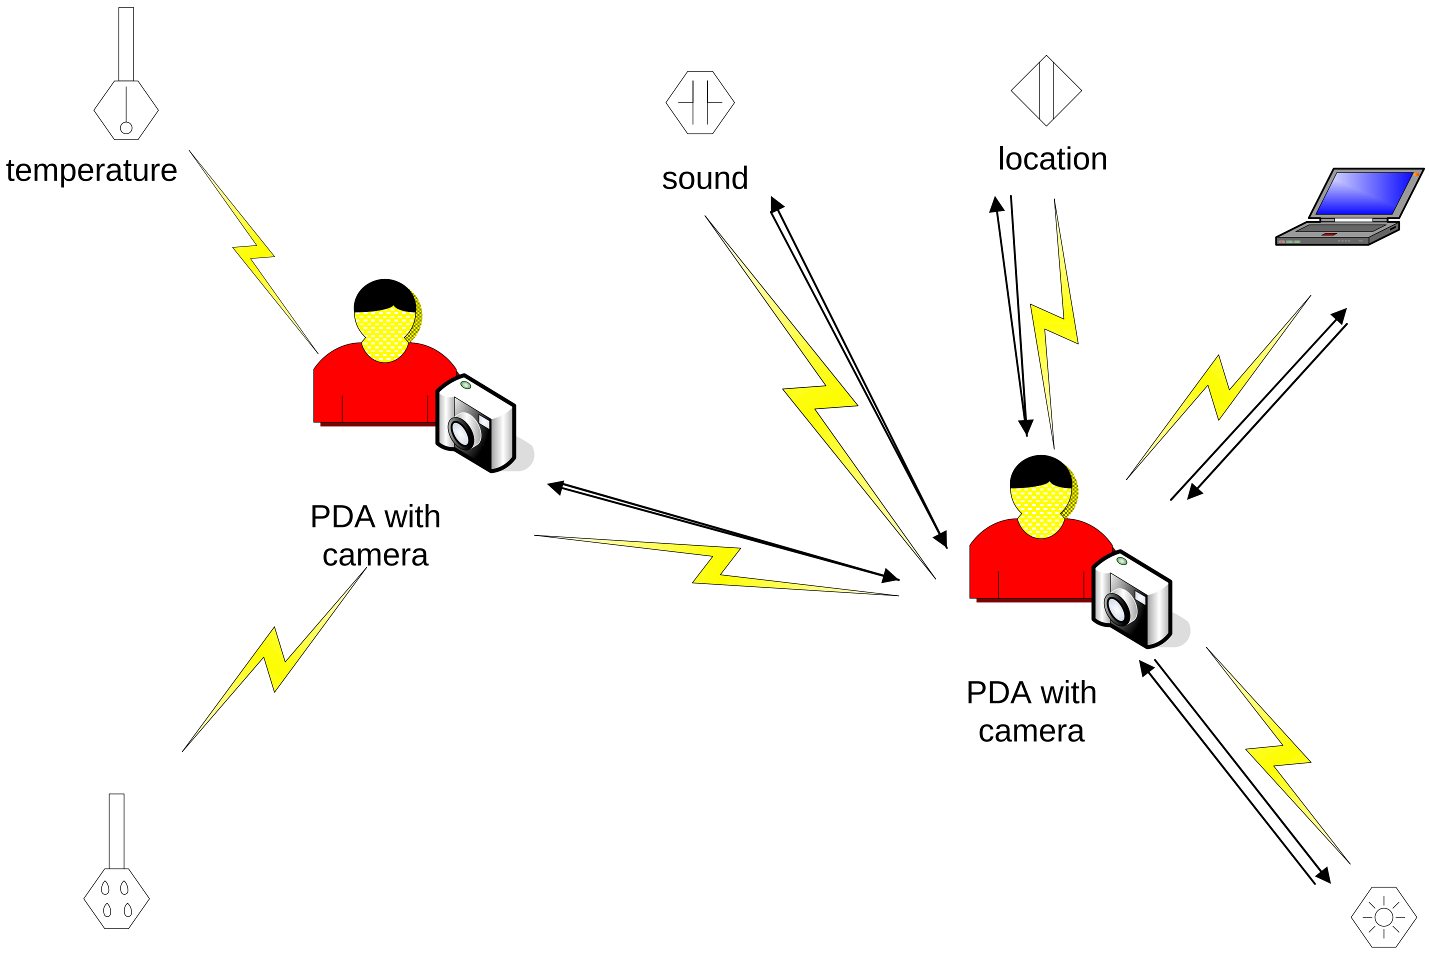

Heterogeneous Networks

- Mobile users (iPAQs & Stargates)
  - User input simulates taking pictures
  - 802.11 WLAN communication
- Environmental sensors (motes & RFIDs)
  - Measure temperature, light, and location
  - RF communication
- Logical sensors (laptops quipped with motes/RFIDs)
  - Communicate with mobile users and environmental sensors
  - 802.11 WLAN communication
  - RF & Bluetooth communication

Metadata Association

```
        procedure Metadata_Association
{
  mark photo-shoot time;
  wait 1 association period after photo;
  determine relevant time interval;
  associate file name and timestamp;
  while ( Pop the smallest offset Data
            within relevant time interval )
        if ( !duplicated (Data) && !filtered (Data) )
            write Data to metadata;
}
```

Context-Aware Image Creation

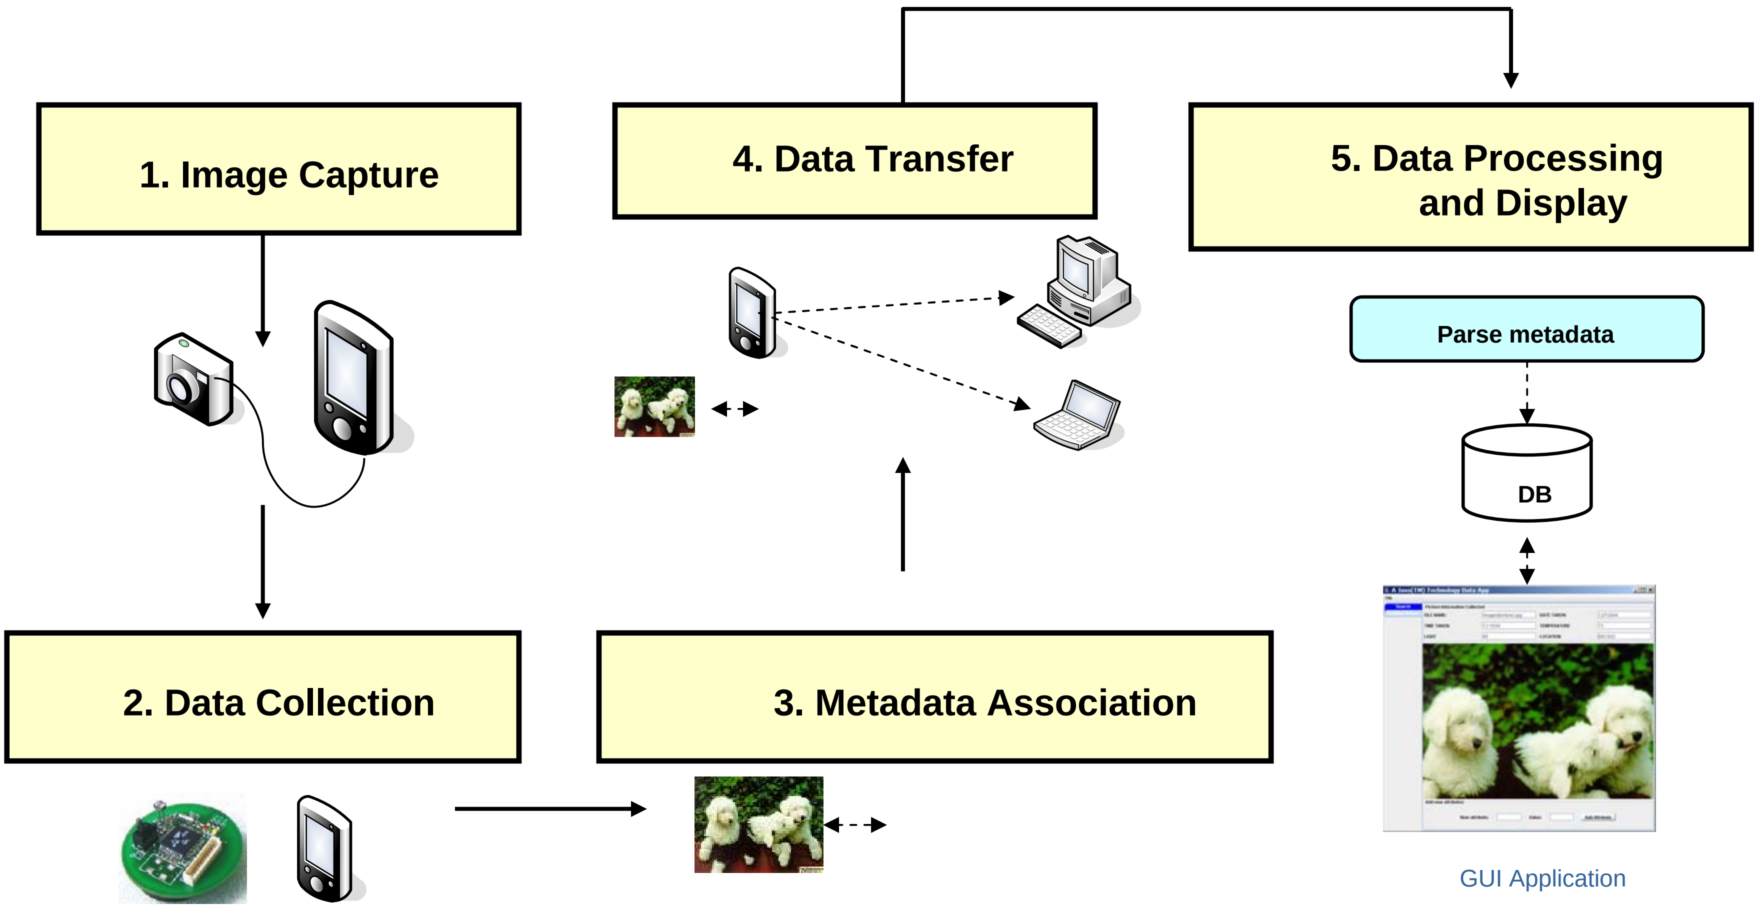

Database and GUI

- Images and associated metadata are transferred to a desk/lap-top PC server
- XML parsed and loaded into the database
- GUI application allows for flexible search and edit

Prototyping and Experimentation

- Testbed
  - 2 mobile nodes (iPAQs)
  - 3 logical nodes
  - 13 environmental sensors
- Users walk around, take pictures, and collect environmental data
  - 1-hour simulation
  - Two users at a time, total of 9 users
- Data collection
  - On-demand
  - Periodic

Experimental Setup

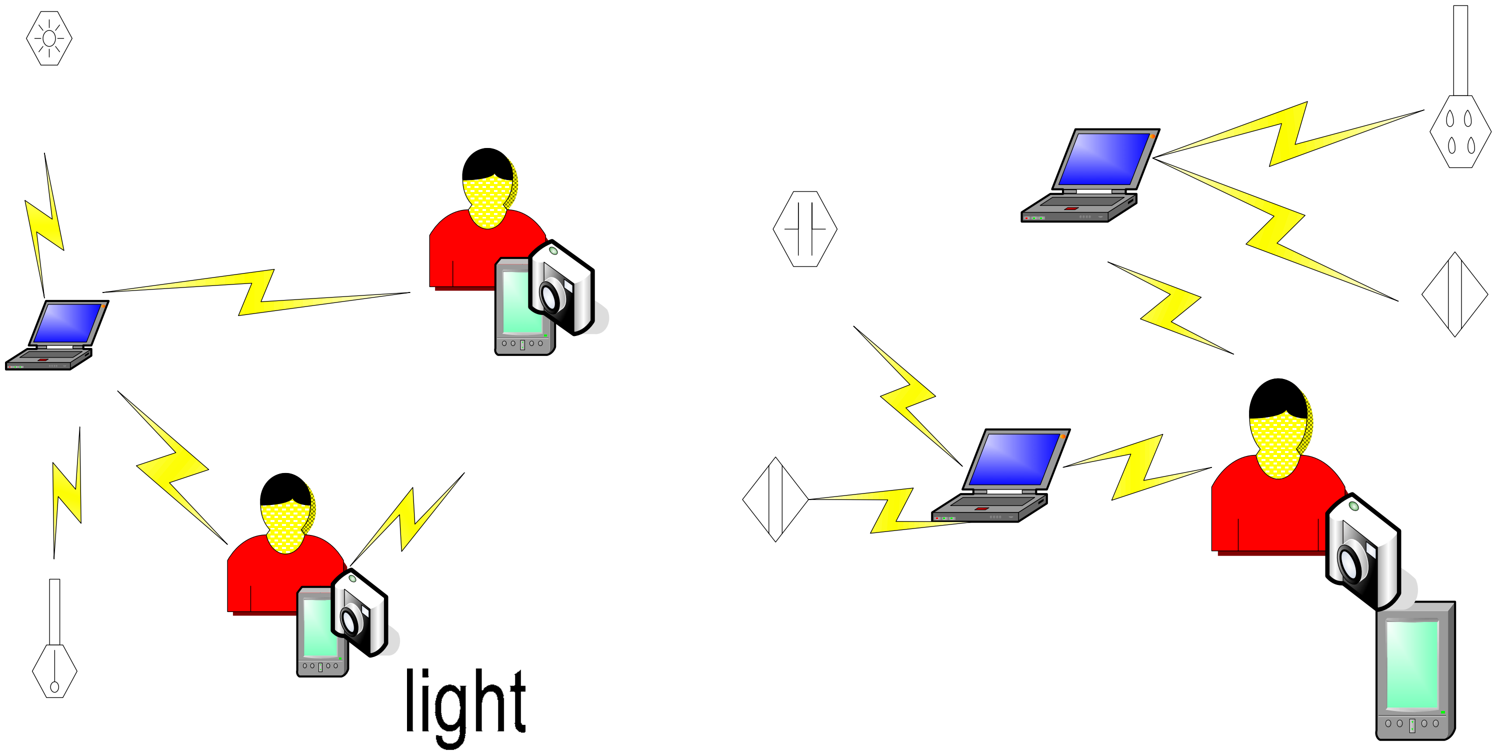

Distributed Location Service

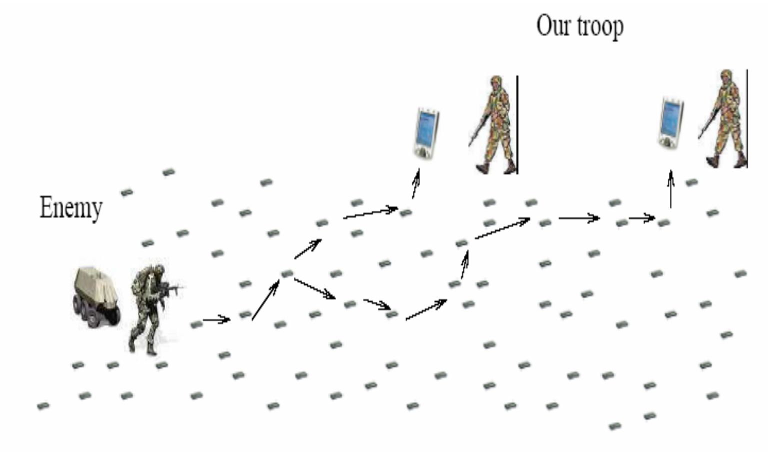

A Typical Scenario:

- Mobile nodes issue queries to the ``static'' sensor network
- Query results are returned to the requester mobiles

When query results are generated:

- Mobile nodes which issued query may have moved away
- Need to route sensed data to a mobile sink!

DLSP: Distributed Location Service Protocol

- What does it do?
  - provides the updated location information of mobile sinks to static sensor nodes

- How?
  - Each mobile independently elects location servers
  - Location info of mobiles is sent to their location servers
  - Other nodes contact the location servers to obtain the location of mobile sinks

Grid Construction

- Level-1 Square
  - Smallest square
- Level-(k+1) Square
  - 4 of level-k square
- No overlap between squares of the same level
- Squares at each level cover the entire coverage area

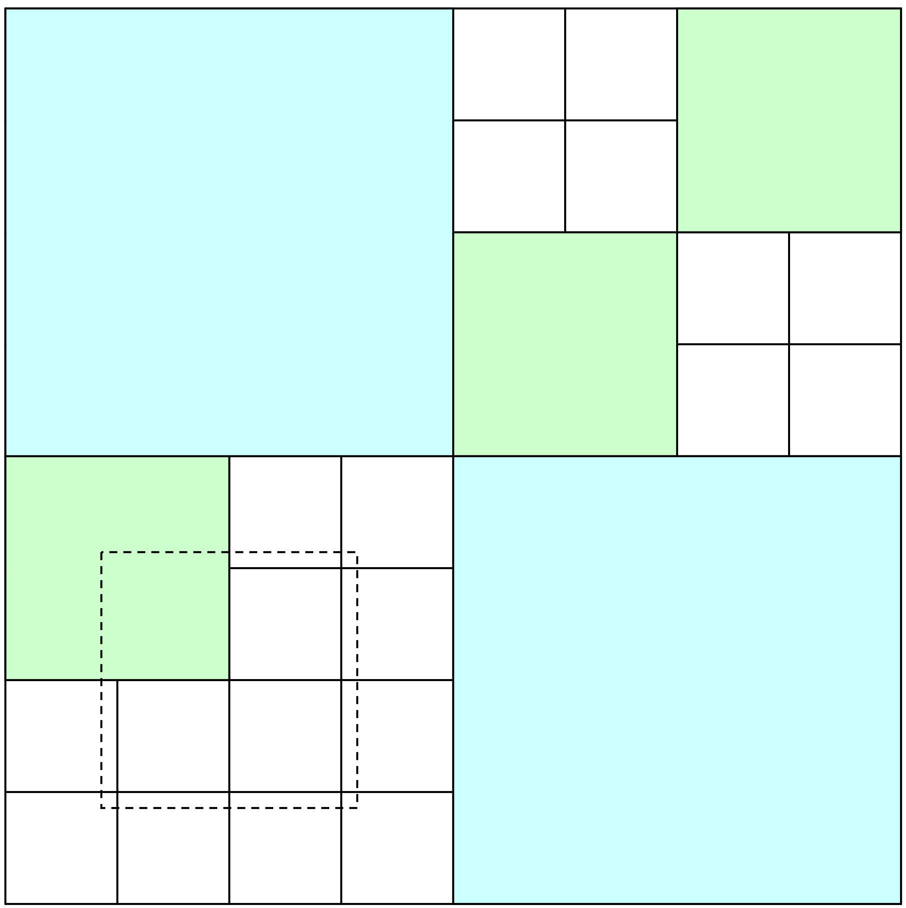

Level-1 square Level-2 square

Level-3 square Not a Level-2 square

Location Server Election

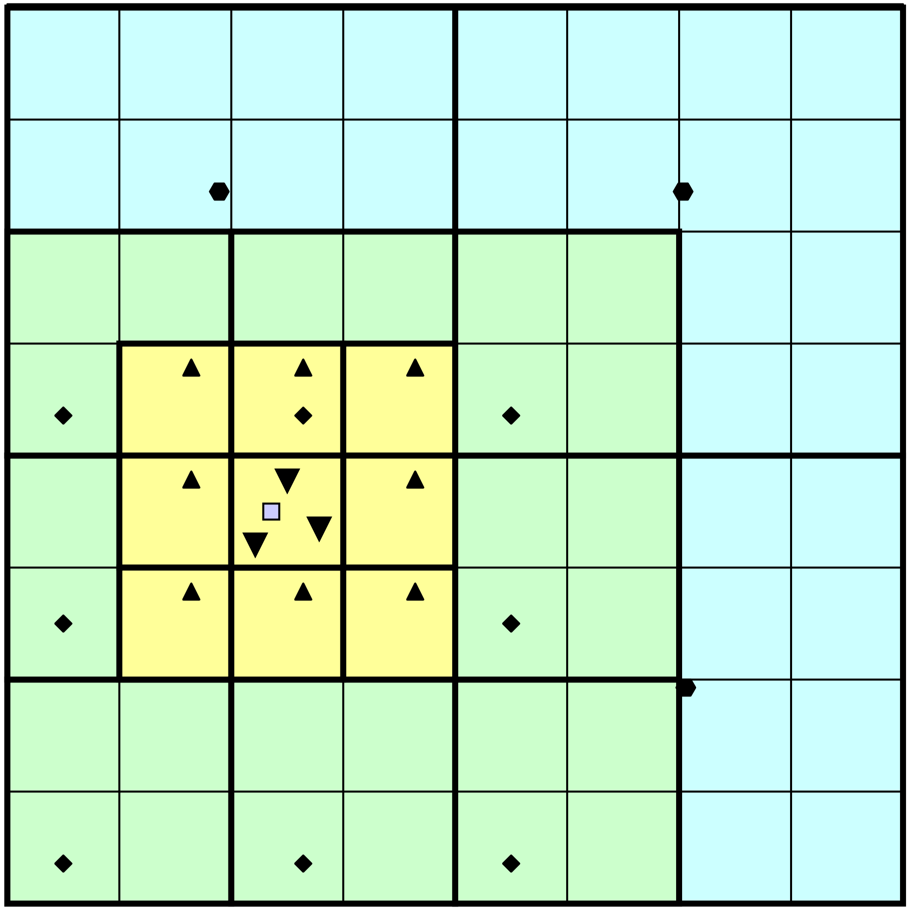

- Level-0 Servers
  - All the nodes within the same level-1 square
- Level-k Servers
  - One from each of neighboring level-k squares
  - Relative location: H(id,k)
- Denser near M and sparser away from M

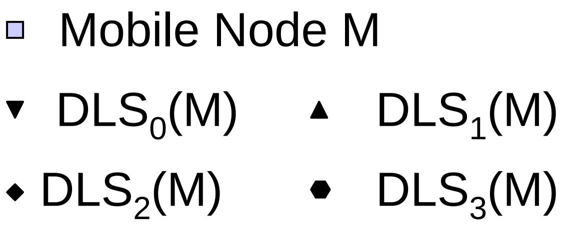

Location Query

- Sink node issues a query if it needs the location of M
- Query is recursively passed to the higher- level (presumed) server

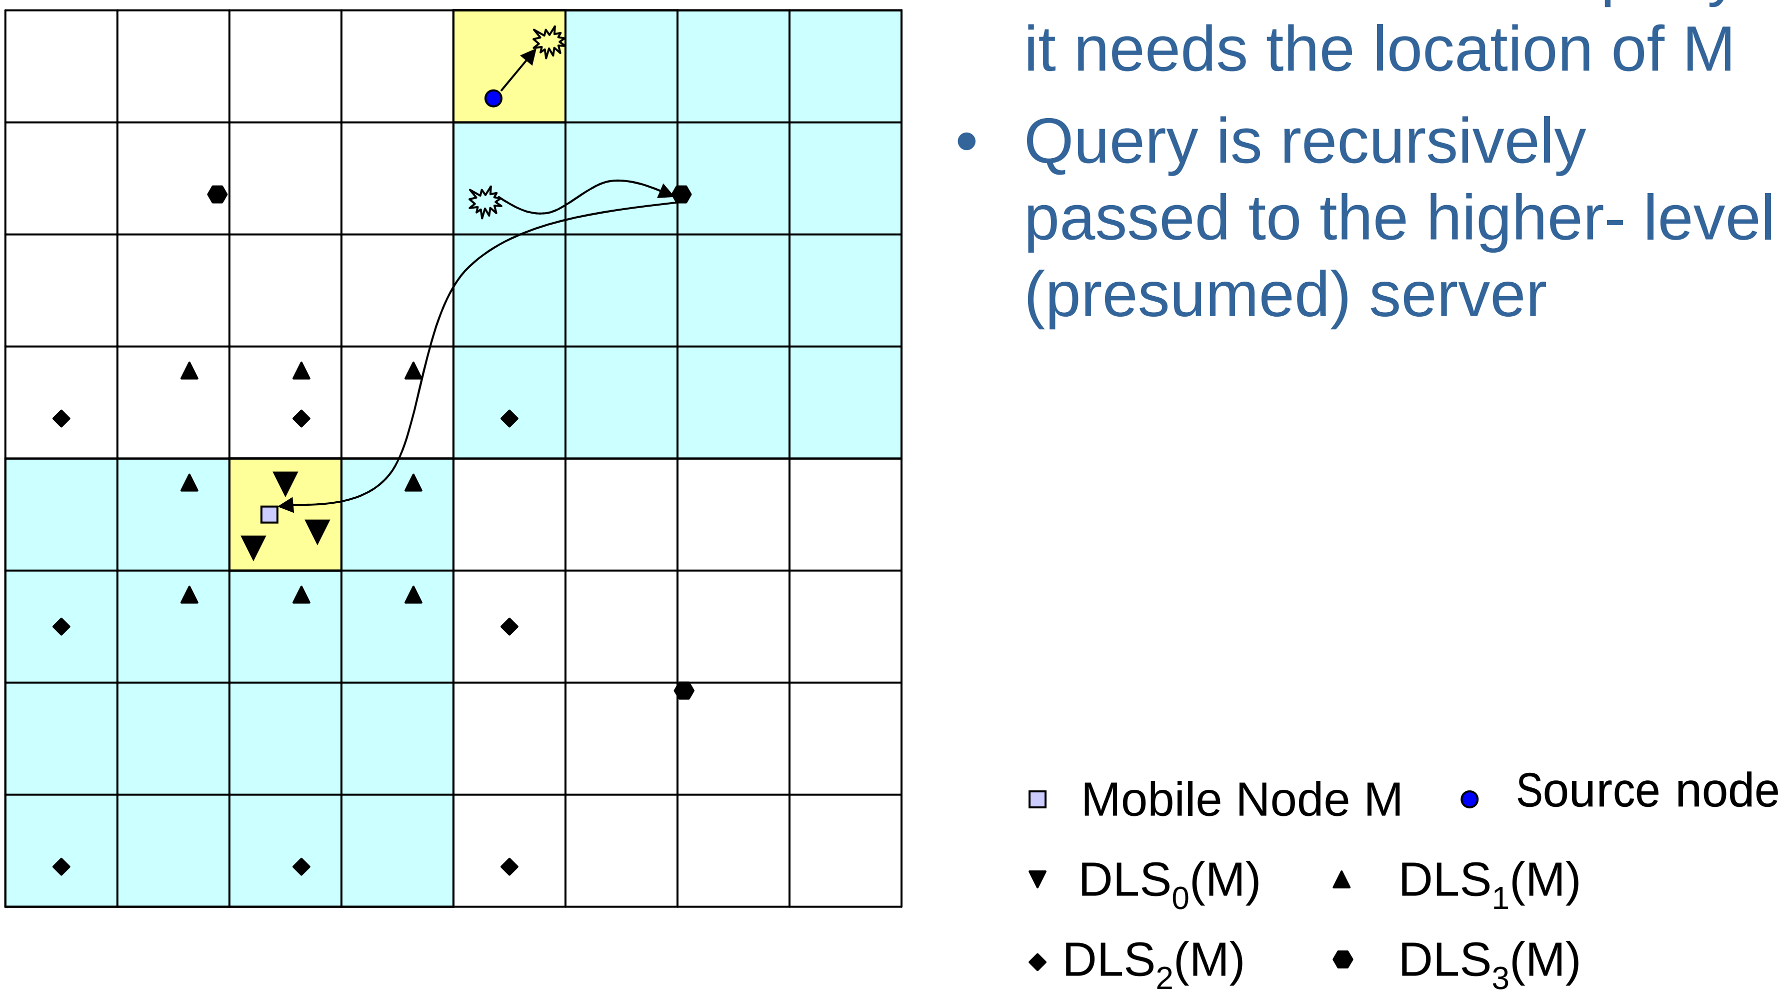

Overhead of DLSP

- Location Query
  - d: distance between src and dst
  - # of msg/query: O(d)
  - delay/query: O(d)
- Location Information Maintenance
  - N: # of sensor nodes, M: # of mobile nodes, L: network size (distance)
  - Mem requirement per sensor node: O(M*log (N)/N)
  - # of msg/mobile node/period: O(L*log (N))

Comparison with Others

- MIT's GLS
  - GLS: Every node is assumed mobile
  - DLSP: Only a small portion of nodes are mobile => more efficient
- Landmark routing
  - DLSP: No need to maintain landmark hierarchy (when nodes move, die, etc.)
- TTDD
  - No overhead for query forwarding, double agent, and local query re-flooding

Security in Networked Embedded Systems

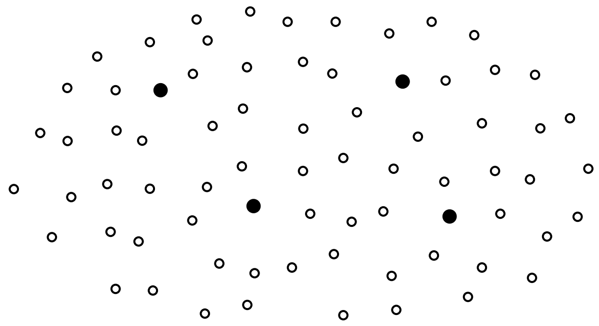

Sensor Network

- Self-organizing, self-healing
- Battery-powered
- Unattended, not rechargeable
- A large number of nodes

Challenges

- Wireless
- Limited Energy
- Large-scale

Our Approach

- Lightweight
  Not sacrificing security level
- Distributed, P2P
- Tailored to Threat/Svc

Threat Model

OUTSIDER

Data Attacks
- Traffic capture/replay
- Spoofing if unencrypted
- Man-in-the-middle (limited)

Radio Attacks
- High-power jamming
- Radio source detection

Physical Attacks
- Reprogram as malicious
- Destroy device
- Extract key materials

INSIDER

Data Attacks
- Traffic injection/flooding
- Unlimited spoofing
- DoS, Man-in-the-middle

Service Disruption on
- Routing (altered/selective)
- Clock synchronization
- Localization

Miscellaneous
- Service/data to adversary
- Malicious service to net

Why LiSP?

THREAT	DEFENSE	PROBLEM	SOLUTION
Attack on Traffic • Eavesdropping • Traffic replay, modification, injection • Service disruption, DoS	**Key Sharing** • Globally • Group-based • Pairwise **Re-Keying** • Periodically • Event-triggered	• Vulnerable to sensor **compromises** • Large **re-keying** overhead • **Transcoding** per hop	**Group-based Key Management** Two-Tier Nets **Distributed Key Management** P2P Nets
Attack on Program The adversary can • capture • reverse-engineer • re-program • clone sensor device(s)	**H/W** Tamper-Resistance **S/W** • Obfuscation • Result Checking • Self-Decryption	Protection of program itself → **Defenseless** once broken	Soft Tamper-Proofing via **Program-Integrity Verification**

LiSP Architecture

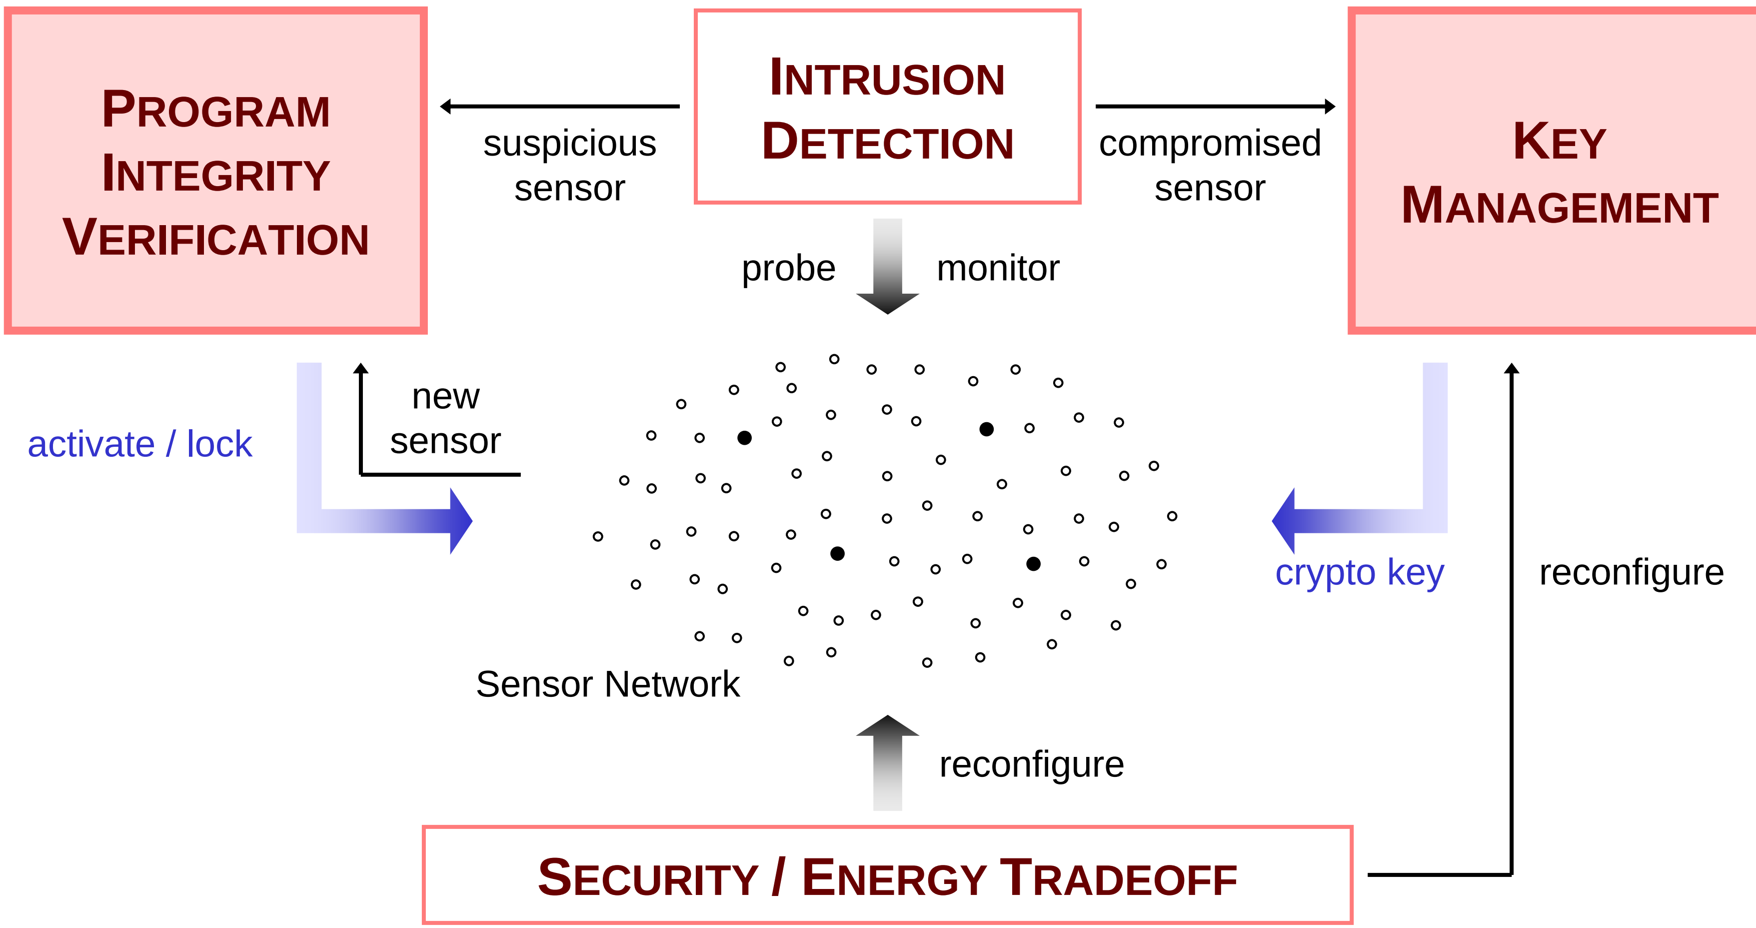

Sensor Networks Research at UCB

Miniaturization – Pister (SmartDust)

- The Goal: Autonomous millimeter-scale r
  - Sensing
  - Computation
  - Communicatio
  - Power
  - Motors

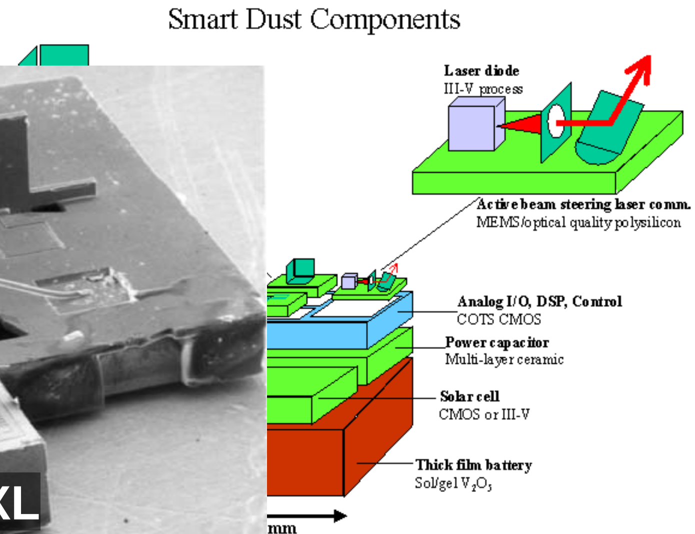

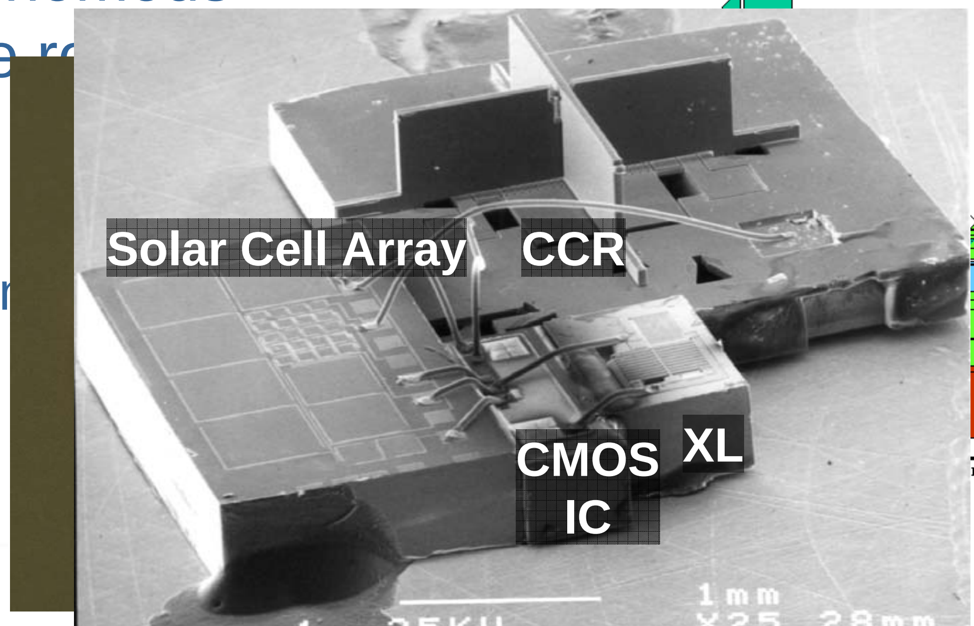

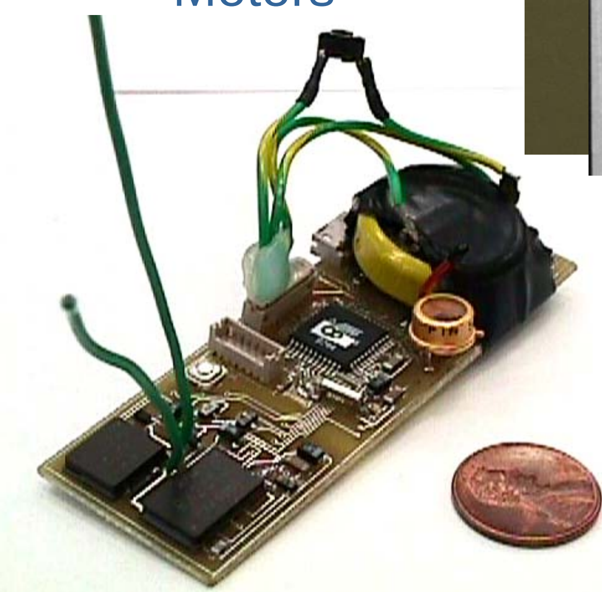

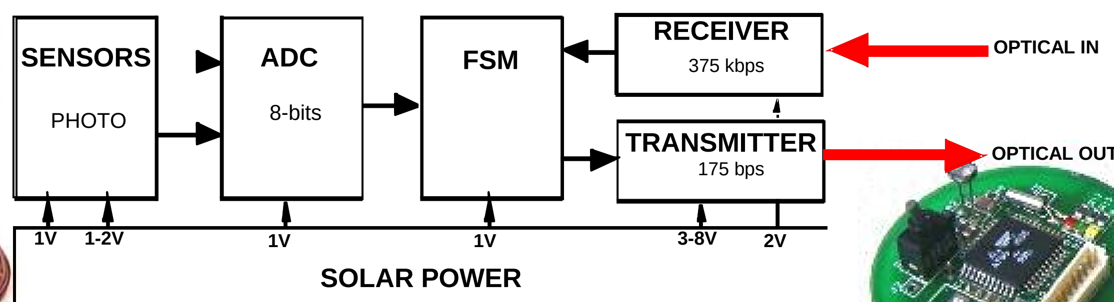

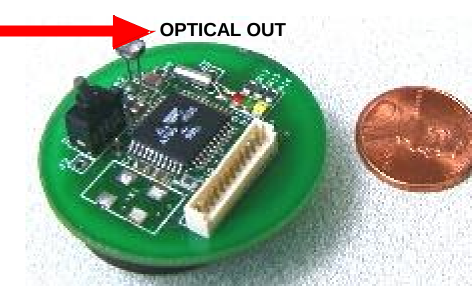

Low Power RF – Rabaey (PicoRadio)

- CMOS
  - Cheap, Integrated
- mW -> sub mW
- Simple
- Advantage in Numbers

BWRC

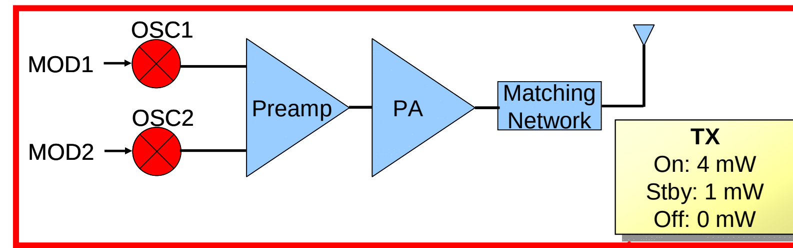

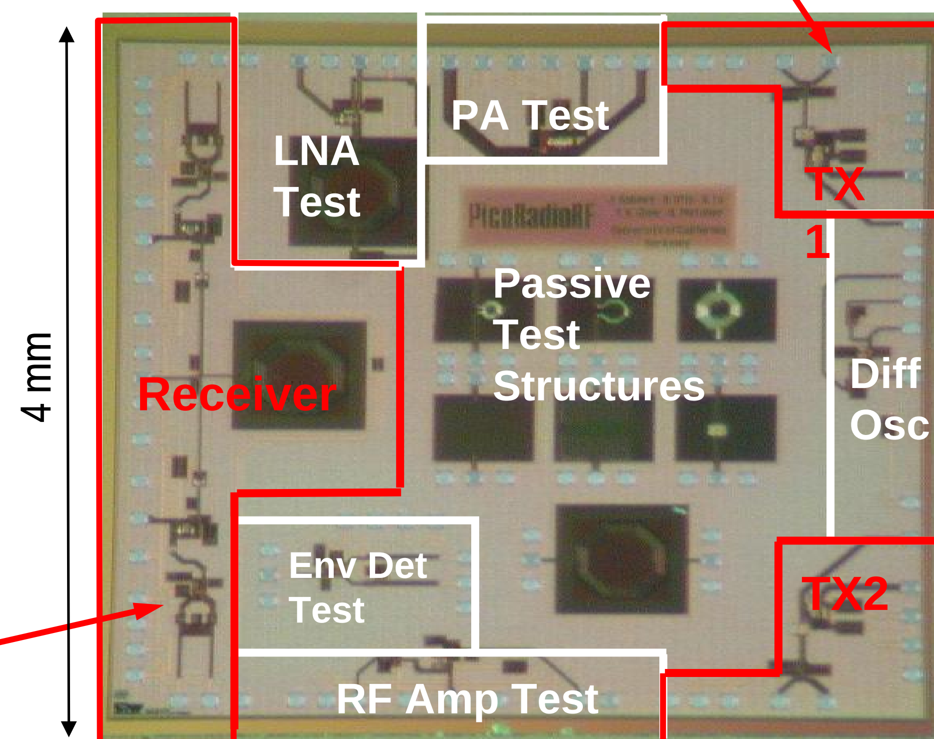

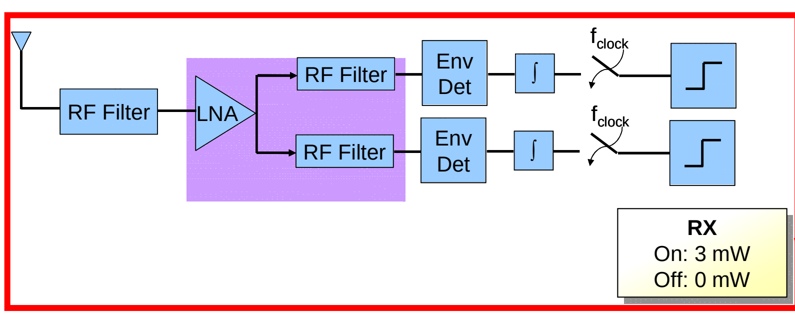

System/Networking/Programming – Culler

Services

Networking

TinyOS www.tinyos.net

WeC 99
“Smart Rock”

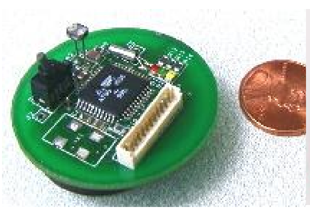

Small microcontroller

- 8 kb code,
- 512 B data

Simple, low-power radio

- 10 kb

EEPROM (32 KB)

Simple sensors

Rene 11/00

Designed for experimentation

-sensor boards

-power boards

DARPA SENSIT, Expeditions

Dot 9/01

Demonstrate scale

- Intel

Crossbow

Mica 1/02

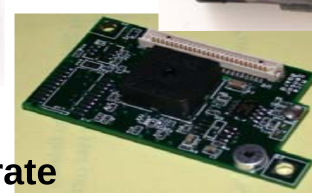

NEST open exp. platform

128 KB code, 4 KB data

50 KB radio

512 KB Flash

comm accelerators

- DARPA NEST

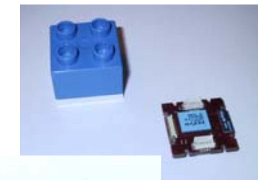

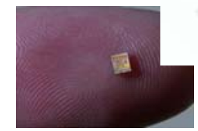

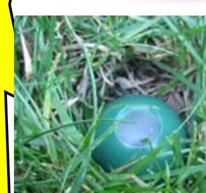

Structural Monitoring – Glaser, Fenves

- Dense Instrumentation of Full Structure
  - Cost is all in the wires
- Leads to in situ monitoring
- Self-inspection and Diagnosis

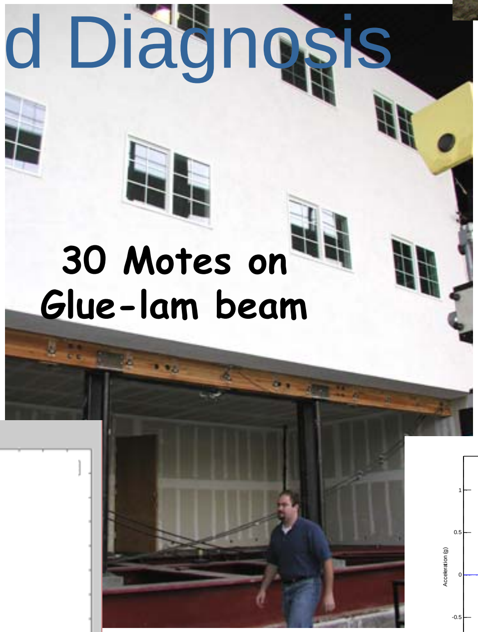

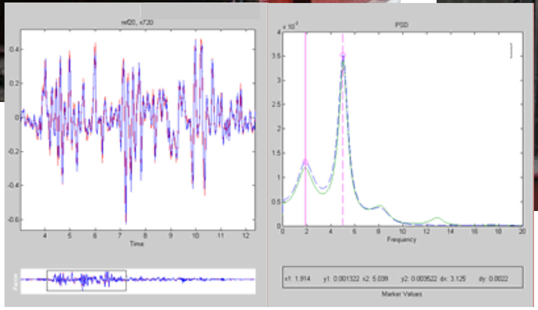

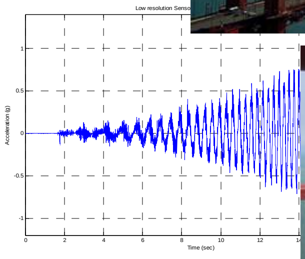

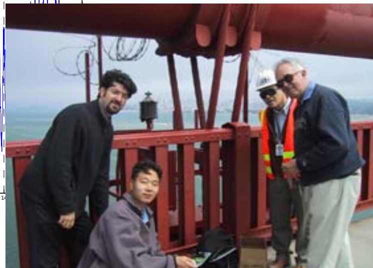

Protection – Sastry, Culler, Brewer, Wagner

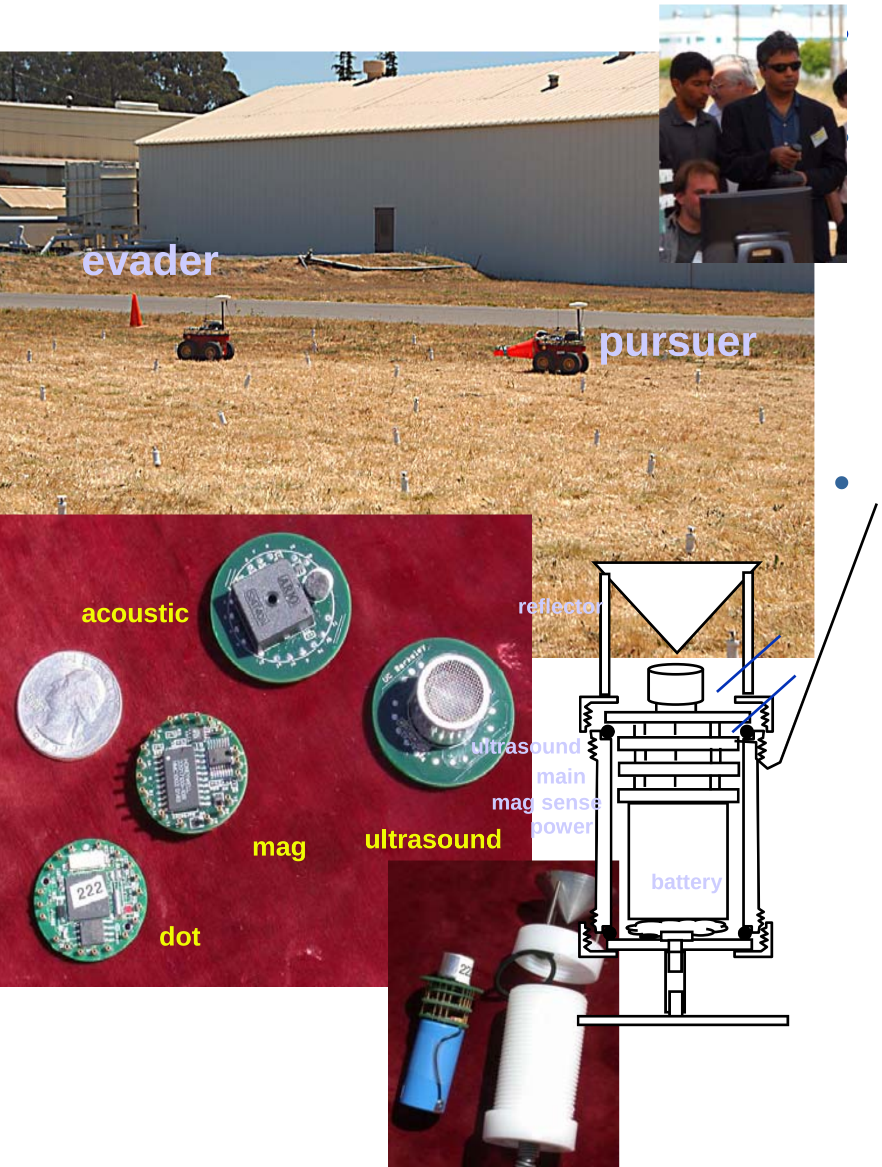

- Detect vehicle entering sensitive area, track using magnetics, pursue and capture by UGV.
- Components
  - 10x10 array of robust wireless, self-localizing sensors over 400 m^2 area
  - Low cost, robust ‘mote’ device
  - Evader: human controlled Rover
  - Pursuer: autonomous rover with mote, embedded PC, GPS
- Operation
  - Nodes inter-range (Ultrasonic) and self localize from few anchors, correct for earth mag, go into low-power ‘sentry’ state
  - Detect entry and track evader
    - Local mag signal processing determines event and announces to neighbors
    - Neighborhood aggregates and estimates position
    - Network routes estimate from leader to tracker (multihop)
  - Pursuer enters and navigates to intercede
    - Motes detect and estimate multiple events
    - Route to mobile Pursuer node
    - Disambiguates events to form map
    - Closed inner-loop navigation control
    - Closed information-driven pursuit control

Sensor Net Databases – Hellerstein, Franklin

- Relational databases: rich queries described by declarative queries over tables of data
  - select, join, count, sum, ...
  - user dictates what should be computed
  - query optimizer determines how
  - assumes data presented in complete, tabular form
- database operations over streams of data
  - incremental query processing
- process the query in the sensor net
  - query processing == content-based routing?
  - energy savings, bandwidth, reliability

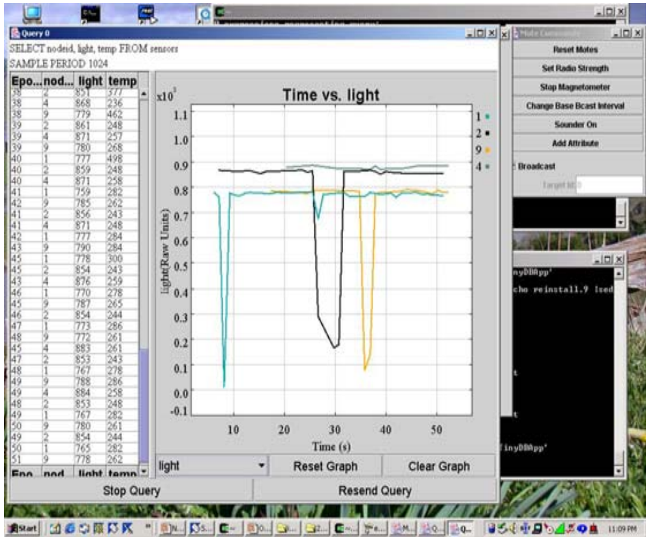

SELECT AVG(light)
GROUP BY roomNo

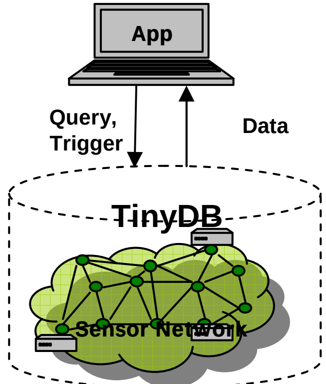

Security - Wagner

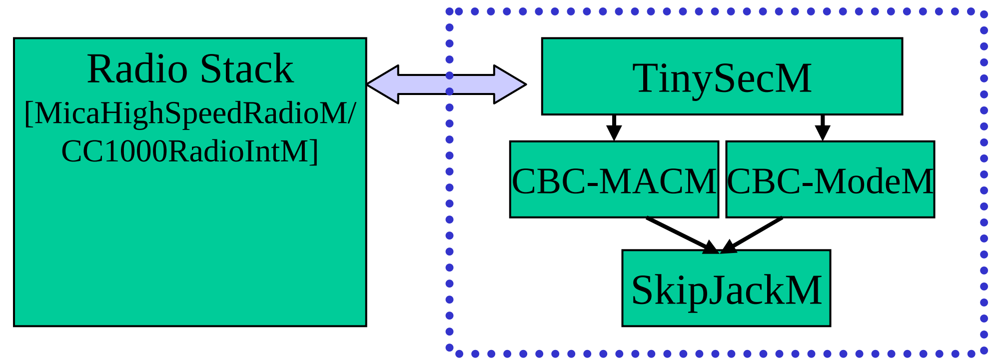

cellphones

1980 analog cellphones: AMPS

analog cloning, scanners
- fraud pervasive & costly

digital: TDMA, GSM

1990

TDMA eavesdropping [Bar]

more TDMA flaws [WSK]
GSM cloneable [BGW]
GSM eavesdropping [BSW,BGW]

2000

Future: 3rd gen.: 3GPP, …

wireless networks

1999 802.11, WEP

2000

2001 WEP broken [BGW]
WEP badly broken [FMS]
- attacks pervasive

2002

2003 WPA

Future: 802.11i

Let's get it right the first time!

sensor networks

Berkeley motes

2002 TinyOS 1.0

2003 TinyOS 1.1, TinySec

VigilNet
University of Virginia

Energy Efficient Surveillance Syst

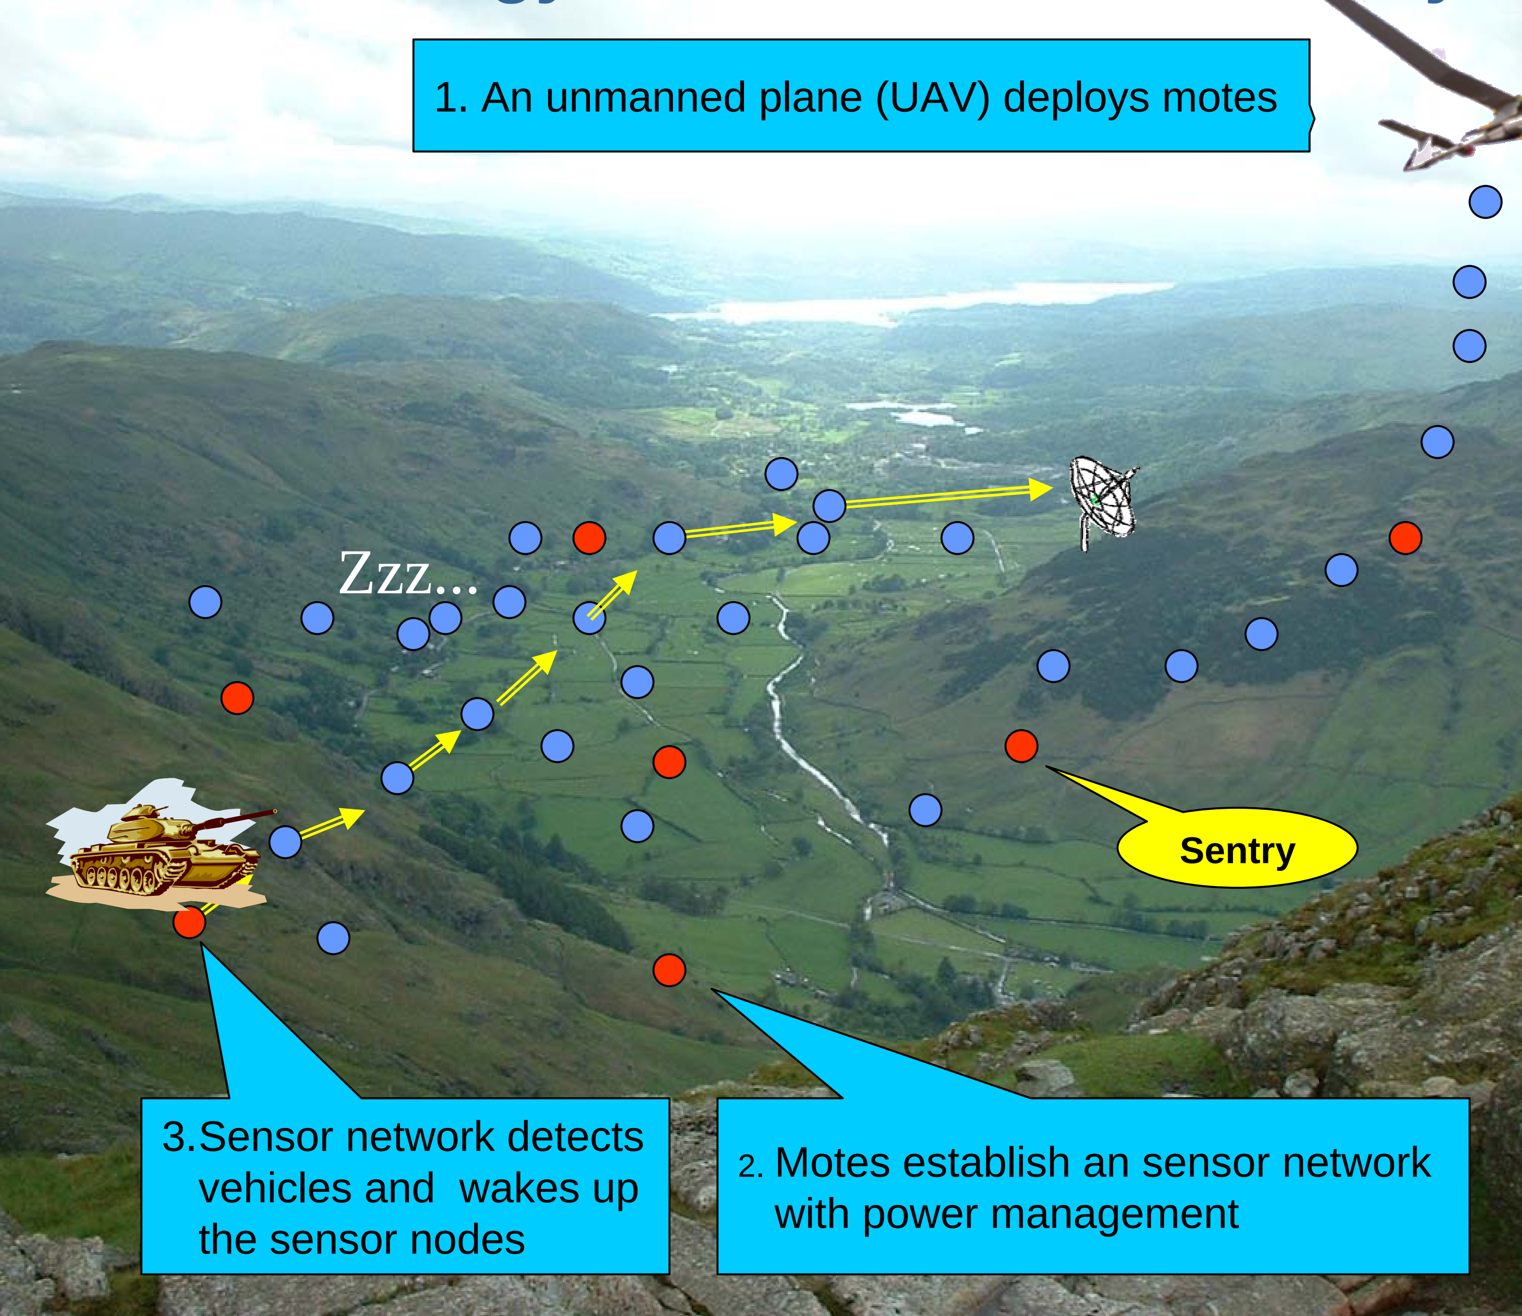

- Diffusion Routing
- Neighbor Discovery
- Time Synchronization
- Parameterization
- Sentry Selection
- Coordinate Grid
- Data Aggregation
- Data Streaming
- Group Management
- Leader Election
- Localization
- Network Monitor
- Tripwire Service
- Reconfiguration
- Reliable MAC
- Leader Migration
- Scheduling
- State Synchronization
- ……

Goals

- Develop an operational self-organizing sensor network of size 1000
- Cover an area of 1000m x 100m
- Stealthy
- Lifetime 3-6 months
- Timely detection, track and classification
  - Large or small vehicle
  - Person, person with weapon
- Wakeup other devices when necessary
  - Extend the lifetime of those devices as well
- Exhibit self-healing capabilities

VigilNet Architecture V1.3

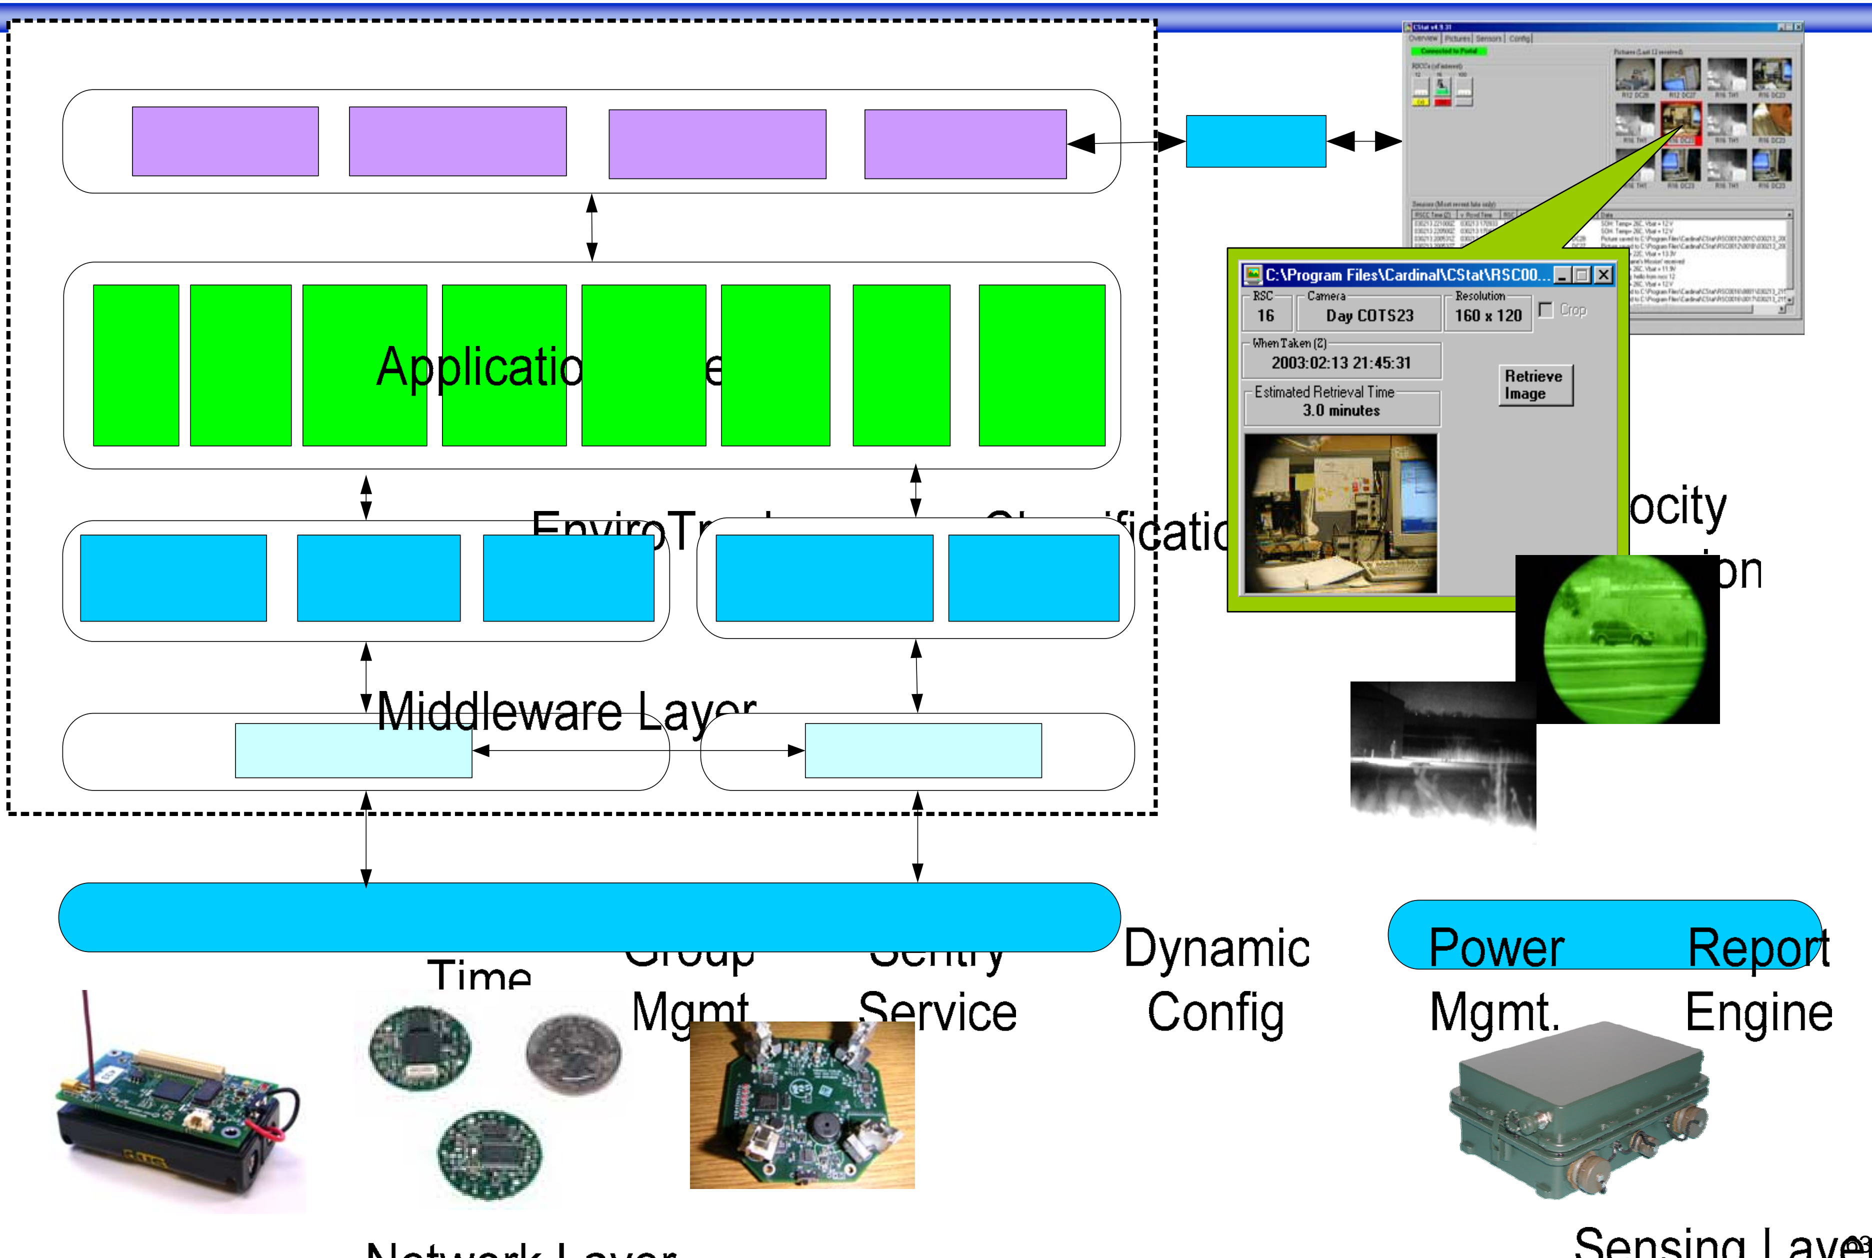

Tripwire-based Surveillance

- Partition sensor network into multiple sections.
- Turn off all the nodes in dormant sections.
- Apply sentry-based power management in tripwire sections
- Periodically, sections rotate to balance energy.

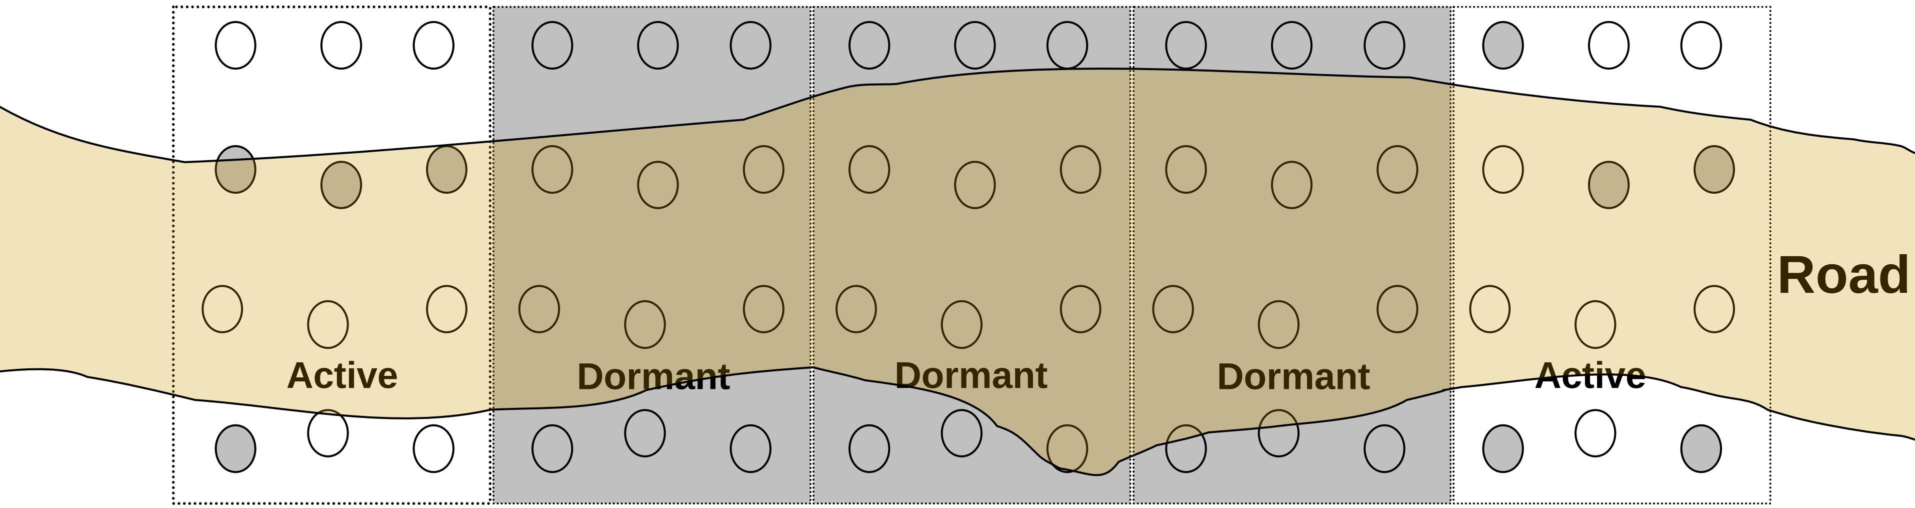

System Test with 203 Nodes

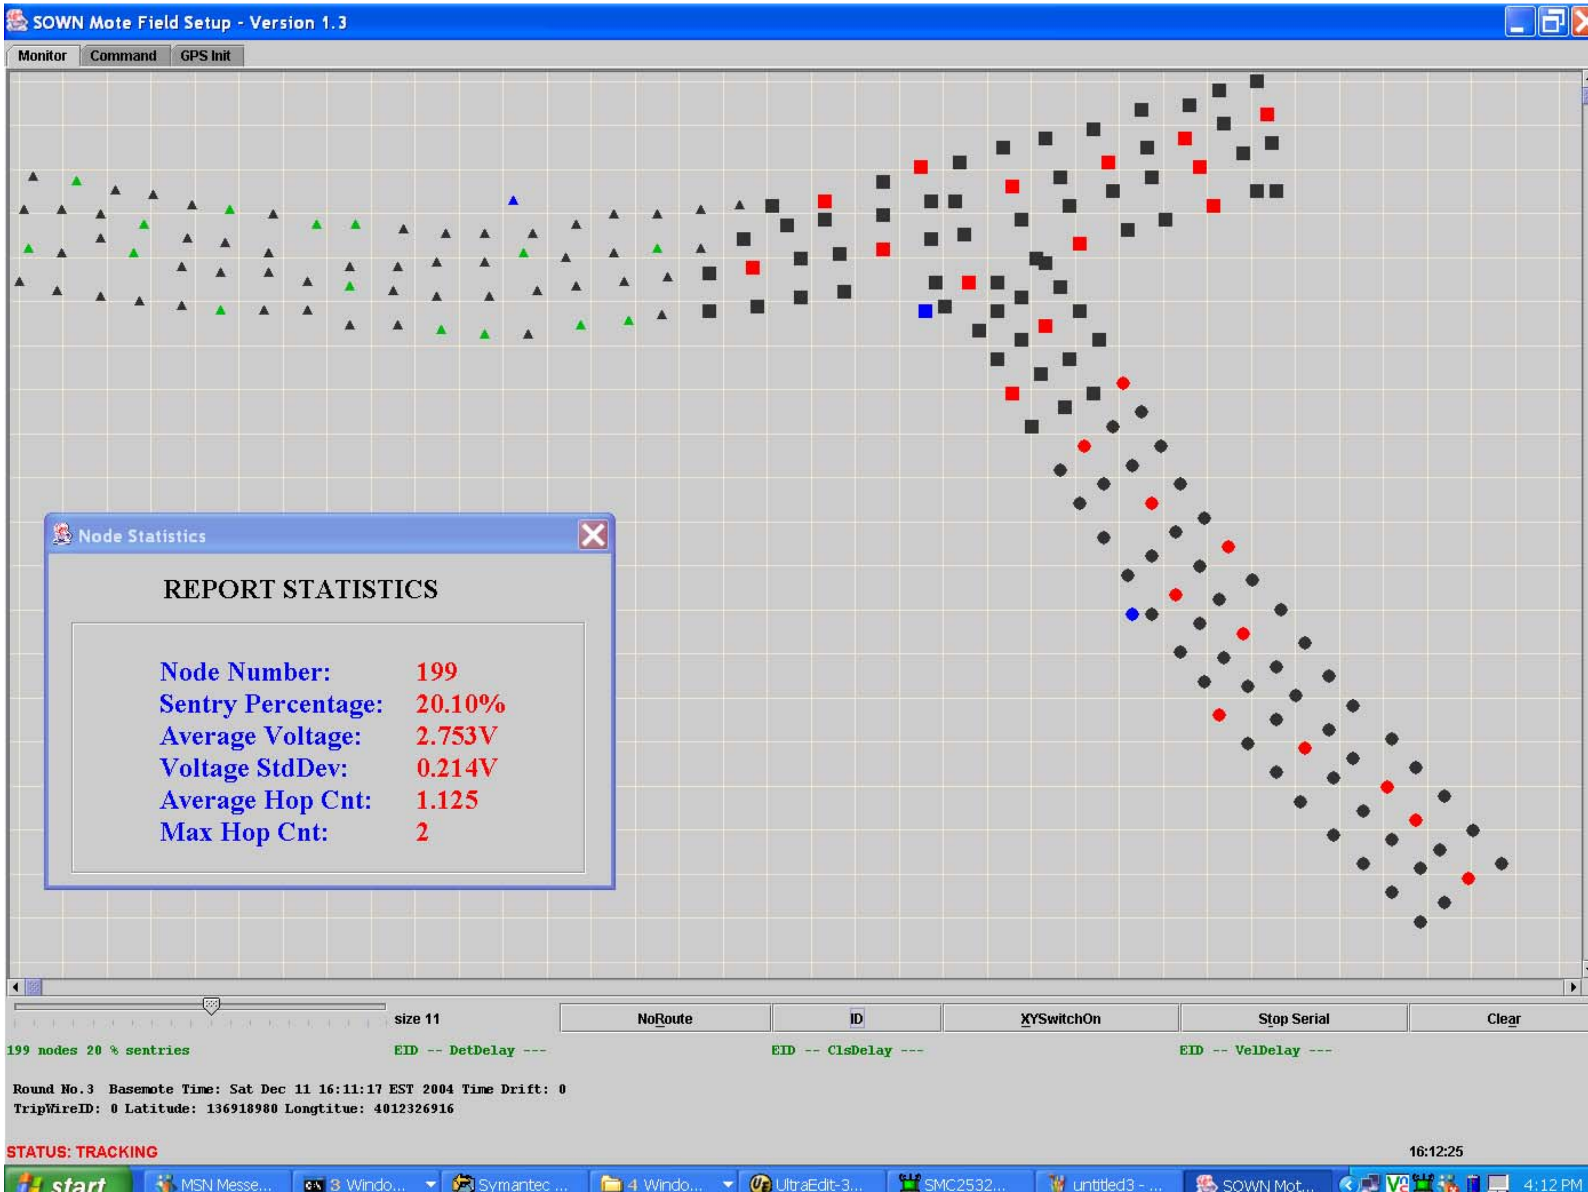

3-Tier Classification

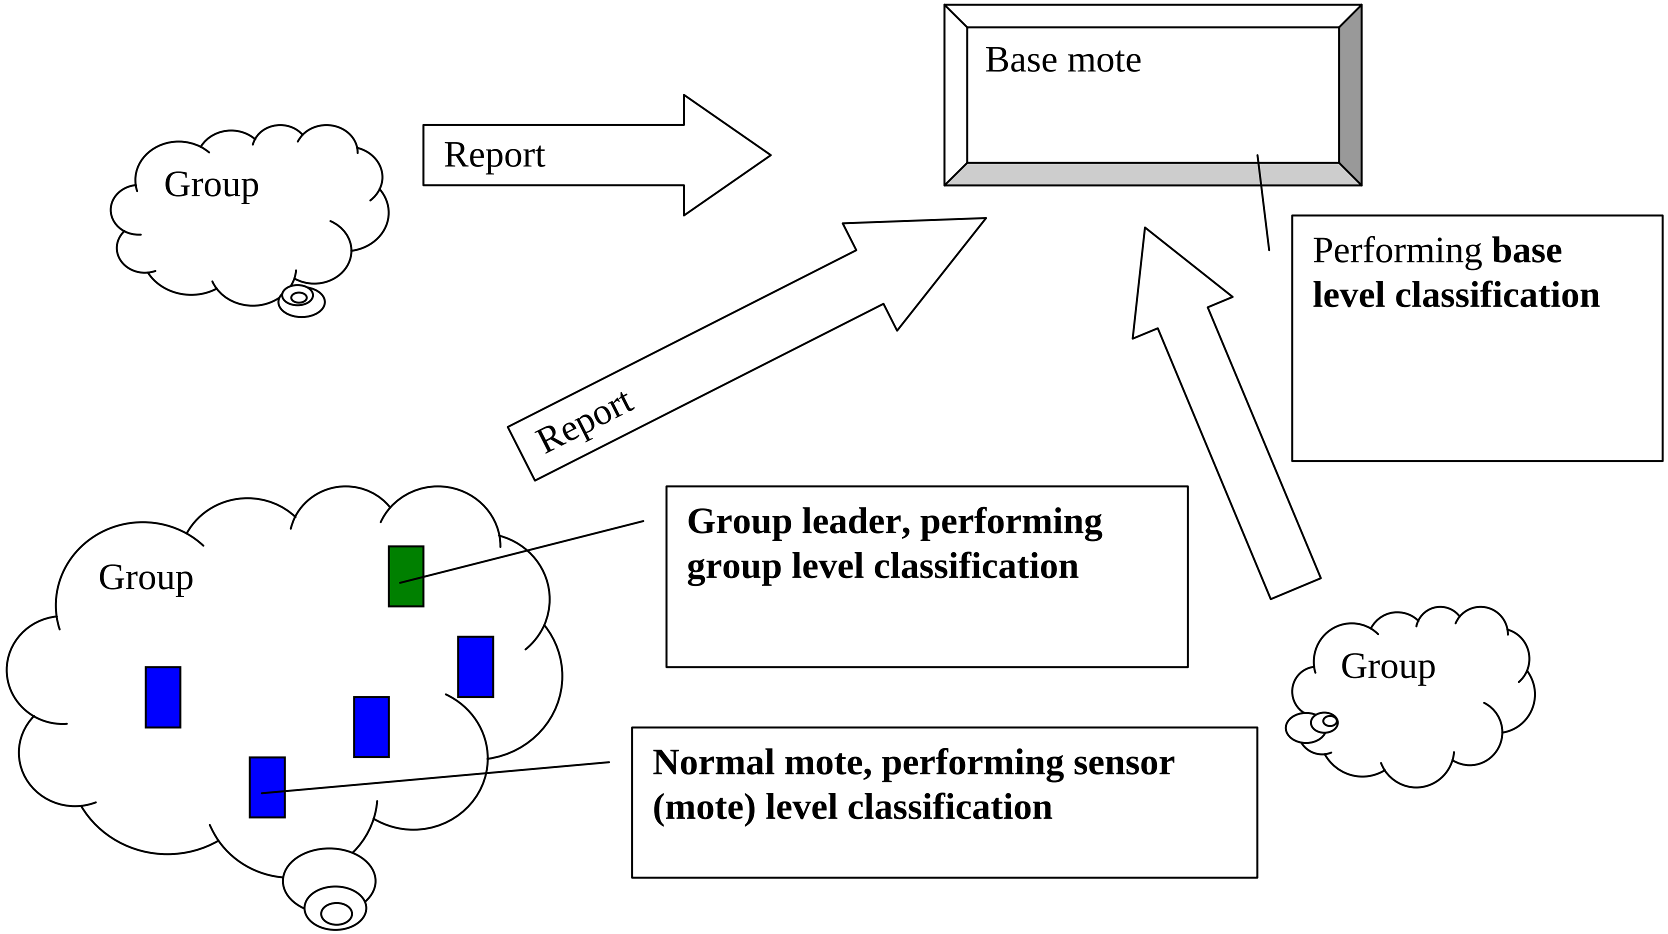

Concluding Remarks

- Sensor networks provide an inexpensive vehicle for exploring various (old and new) research issues
- Commercial applications with RFIDs as leader
- Current and future directions: query processing using geostatistics, sensor network security; tradeoffs among perf, security, reliability and resource consumption; extreme scaling and other DoD/commerical apps.