Embedded Networked Sensing: A Technology in Transition

From Smart Dust to Multi-scale, Multi-modal, Multi-user Observing Systems

Mani Srivastava
mbs@ucla.edu
Networked & Embedded Systems Lab
Center for Embedded Networked Sensing
UCLA

Ack: Many slides in entirety or in part are from various CENS colleagues & students

A decade of ENS research: Many First-generation Deployments and System Components

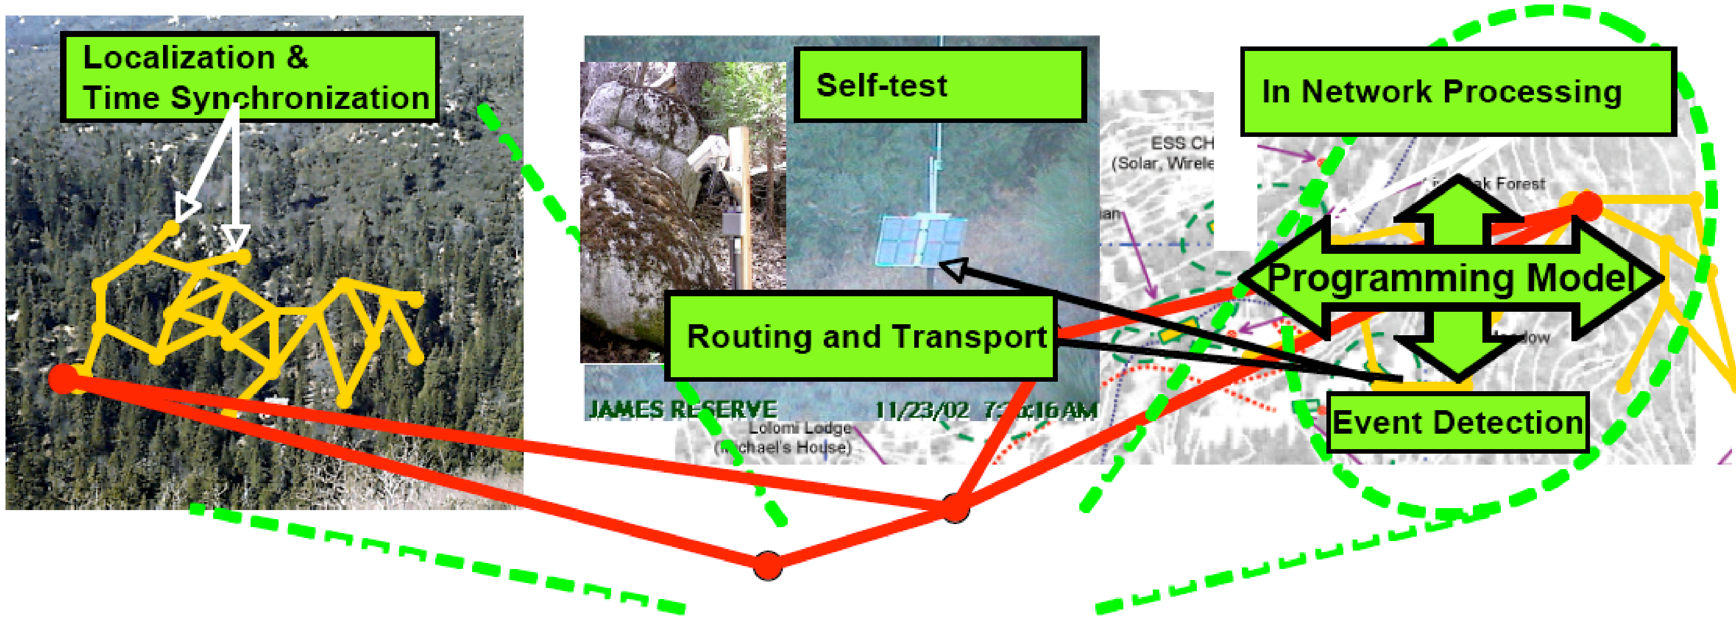

- Reusable, modular, flexible, well-characterized services, tools, and system components
  - Routing, Reliable transport, Mobility, Plug and play
  - Time synchronization, Energy Harvesting, Localization, Self-Test, Calibration
  - In Network Processing: Tasking, Filtering, Triggering, Fault detection, Sample Collection
  - Tools, Programming Abstractions, Application Authoring, Embedded Statistical Tools
  - Development, simulation, testing, management, debugging
- Experience with large (> 100s of sensors), long-term (months-years) deployments
  - James Reserve, Great Duck Island, DARPA NEST, Mexico Seismic Array etc.

Original Drivers of Sensor Networking Research: *Resource Constraints* and *Autonomy*

- Limited battery energy
  - low-power platforms, energy harvesting
- Limited computing, bandwidth, and storage
  - light-weight software frameworks (Tiny*), data centric protocols (diffusion)
- Ad hoc network capacity scaling
  - exploiting correlations among nearby sensors
- Higher cost of communication relative to processing
  - in-network processing to reduce # of bits communicated
- Dominance of Rx over Tx
  - receiver duty cycling with low-power listen
- Self-configuration
  - ad hoc time synchronization and node localization

A Technology in Transition

Early Themes

- Thousands of small “**smart dust**” devices
  - Minimize individual node resource needs
  - Exploit large numbers
- Fully **autonomous** systems
- In-network and collaborative processing for longevity: **optimize communication**

Emergent Themes

- **Heterogeneous** ecology
  - Tiered nodes and networks: optimize system as a whole
  - Inevitable under-sampling (in time or space)
  - Exploit multiple modalities, multiple scales, and mobility
- **Interactive** systems
  - Design for human tier as well... online interaction and tasking
- In-network and collaborative processing for responsiveness, data quality, and data control (privacy): **optimize sensing**
- **Integrity** driven: calibration, self-test, validation

From Single-Purpose Deployments...
... to Multi-Modal Multi-Scale Observing Instruments

FROM...	TO ALSO INCLUDE...
Spatial redundancy	Actuated placement
Planned long-term deployments	Rapid short-lived deployments
Self-configuring autonomy	Interactive exploration
Post-deployment analysis	Live analysis feedback
Pre-calibration	Integrity monitoring
Trusted communications	Safety and robustness
One user and application	Multiple users and applications

New architectural principles

Rich Ecologies beyond Flat & Homogeneous Ad Hoc “Smart Dust” Networks

- Spatially distributed **static nodes**
  - Simultaneous sampling across study volume (dense in time, but possibly sparse in space)
  - Limited energy and sampling rate

- **Collaborative processing arrays**
  - High sampling rate (acoustic, imaging)
  - More computation and communication capacity

- **Actuated nodes** with mobility, articulation
  - Sensor diversity: location, type, duration
  - Allows dense spatial sampling, but possibly sparse in time
  - Adaptive provision of resources (sensors, energy, communication)
  - Enable adaptive, fidelity-driven, 3-D sampling and sample collection

Goal is to have statistical information over entire region with quality similar to high resolution sampling, without applying high resolution to entire region

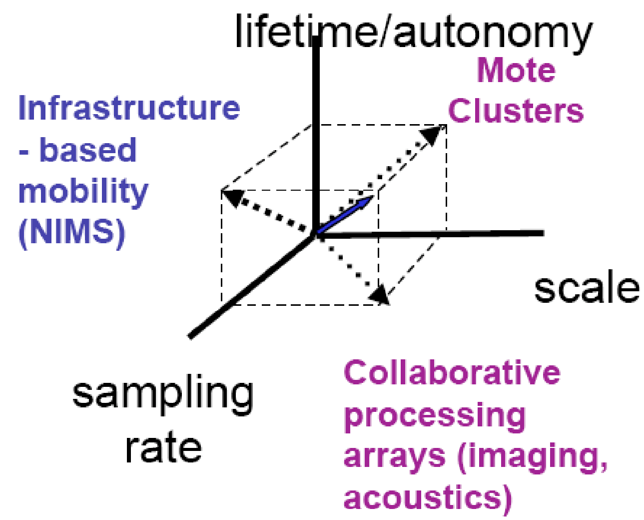

Different information return and trade-offs

Actuation and Mobility as Performance Amplifiers

- While sensor networks are great for dense sensing,
  - The likelihood of under-sampling and communication disconnections is surprisingly high due to obstructions
  - Meeting sampling objectives is often impractical with static nodes
- Mobility, whether controlled or opportunistic, is a critical amplifier of sensing and communication coverage
  - ***Constrained articulation***: magnifies effective range and resolution
  - *Longer-range* ***infrastructure-supported mobility***: enables sensor diversity and adaptive 3D sampling
  - Wide-area ***autonomous mobility***: adds "data mule" capability and increases coverage

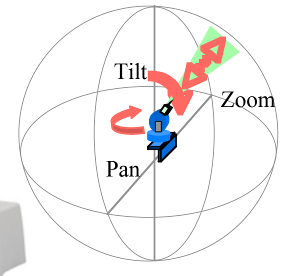

Pan	7.74
Tilt	4.04
Zoom	73
Pan and Zoom	6361
Pan, Tilt and Zoom	226940

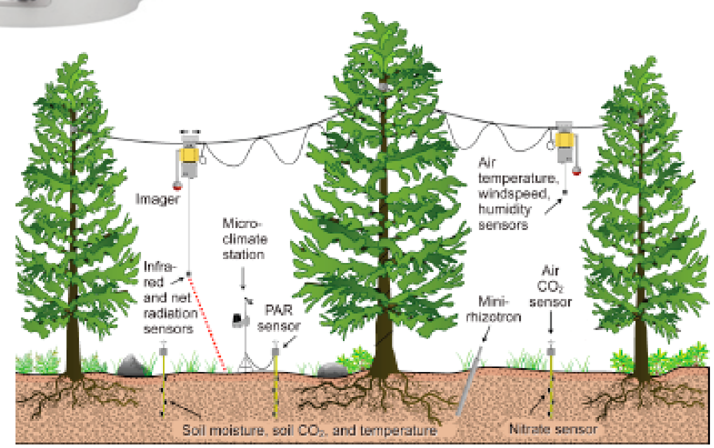

Multi-scale Sensing

- Low-resolution large-field-of-view global sensors guide higher-resolution small-field-of-view local sensors
  - E.g. Image from camera used to guide the actuated NIMS node carrying a high quality PAR sensor yields order of magnitude reduction in area sampled

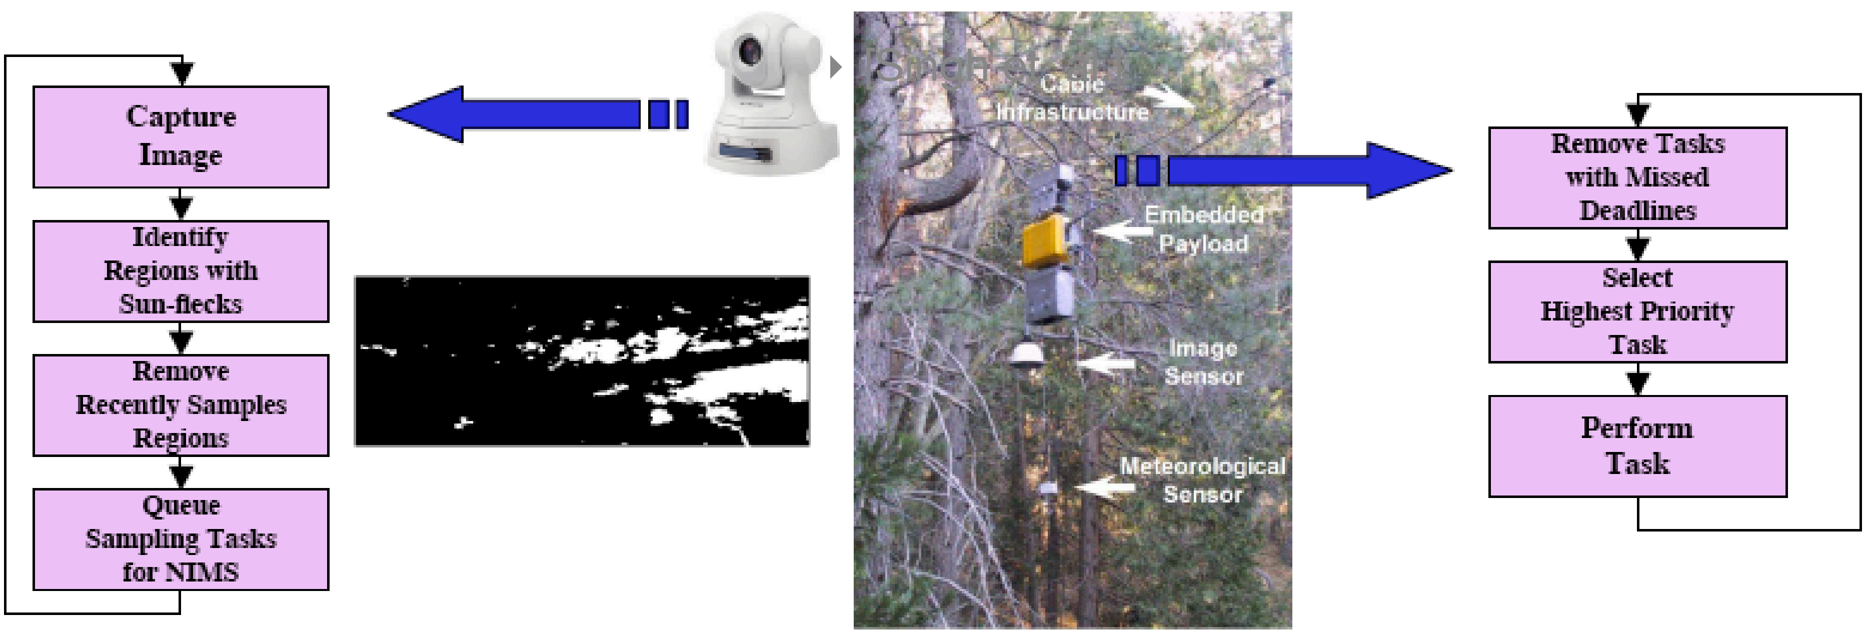

[Singh et. al., 2006]

High Cost of Sensor Data Acquisition

- Early focus on communication
  - “Every bit transmitted brings a sensor node one moment closer to death” (Pottie)
  - Artifact of simple applications requiring low-rate low-complexity sensing
    - thermistors, photodiodes etc.

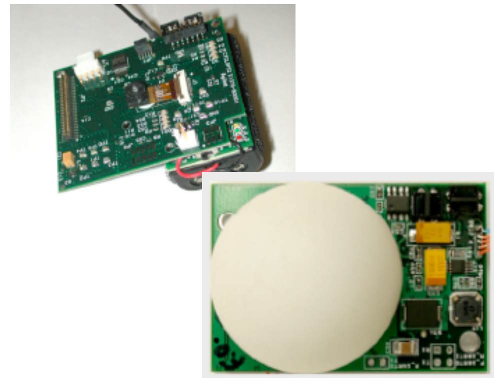

- Emerging applications often require
  - Energy hungry sensing modalities
    - imagers, acoustic arrays, precision sensors etc. (100s of mW)
  - Actuated sensors to cope with 3D spaces and obstructions
    - PZT cameras, robotic nodes (Watts)

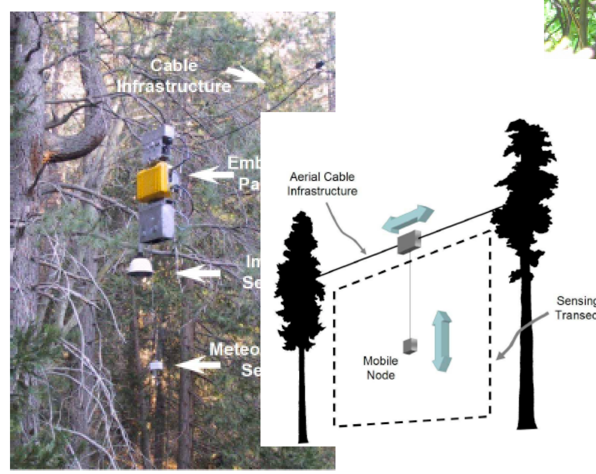

- Architecture implication: **optimize sampling** for required fidelity
  - Adaptive sampling in NIMS, Cyclops
  - Compressive sampling

A barrier challenge: Integrity

How do we monitor the monitors?

Sustaining High Integrity Operation

- Noise and outliers
- Malfunctions
  - faults: calibration, stuck-at
  - bugs: memory corruption, protocol logic
- Malicious adversaries
  - spurious sensor or radio input
  - sensor or comm interference
  - snooping
- Challenges
  - What is the impact on eventual “Quality of Information”?
  - How to detect integrity problems?
  - How to be resilient to them?
  - How to remediate them?

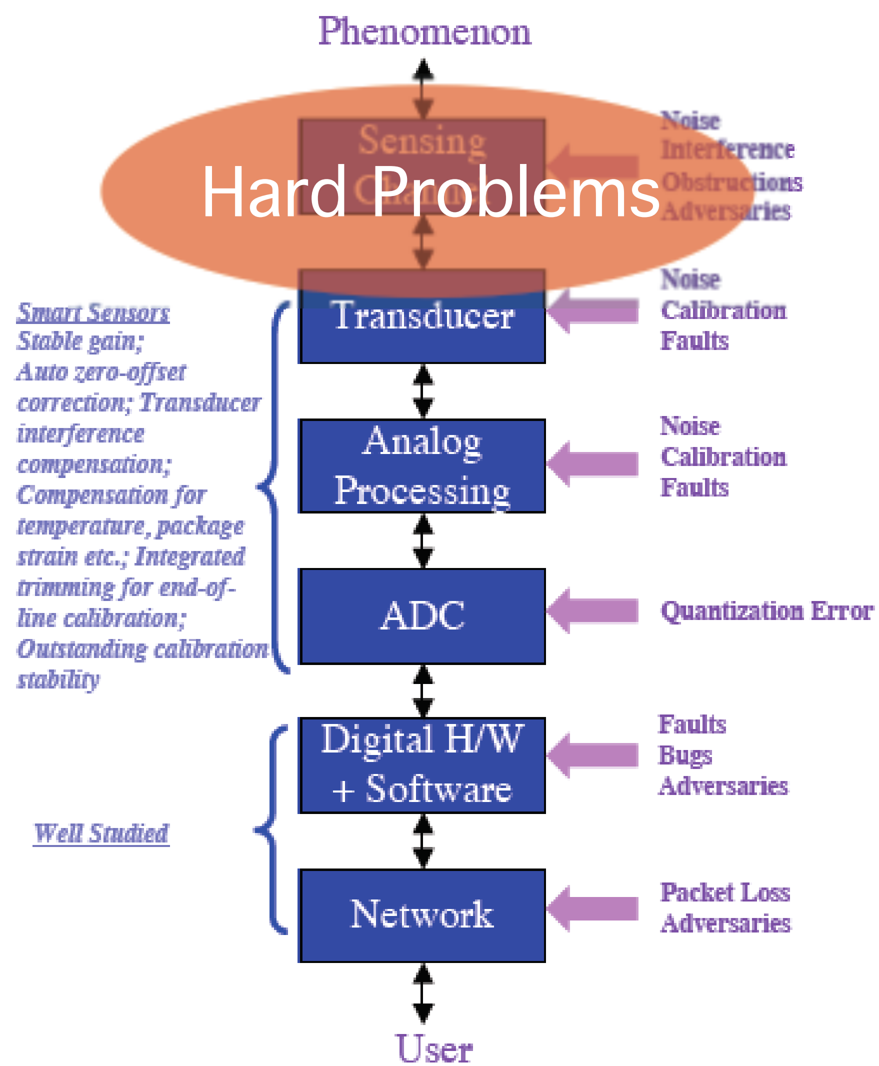

Benign Uncertainties

Adverse Sensing Channels

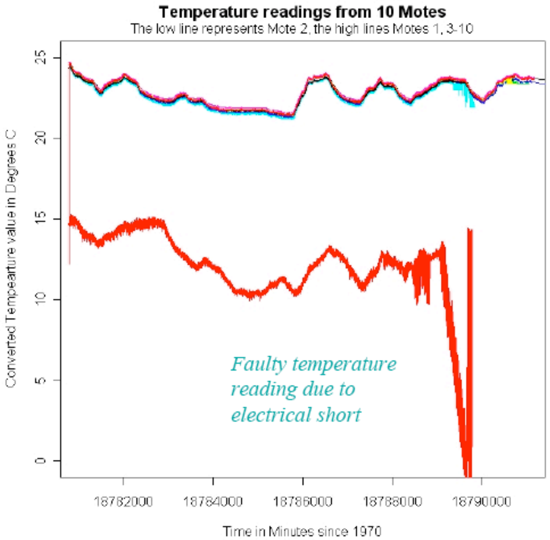

Faults

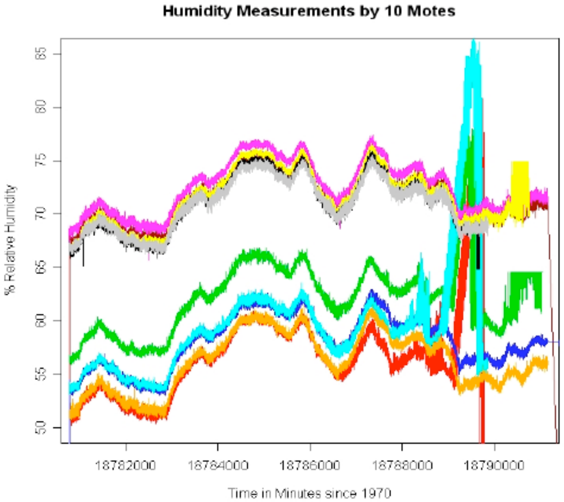

Calibration

Data Faults in Bangladesh Arsenic Study

- Data integrity a show-stopping concern
  - Fault models
  - Detect, diagnose, remedy
- On-line approaches
  - Rule-based
  - Reputation-based

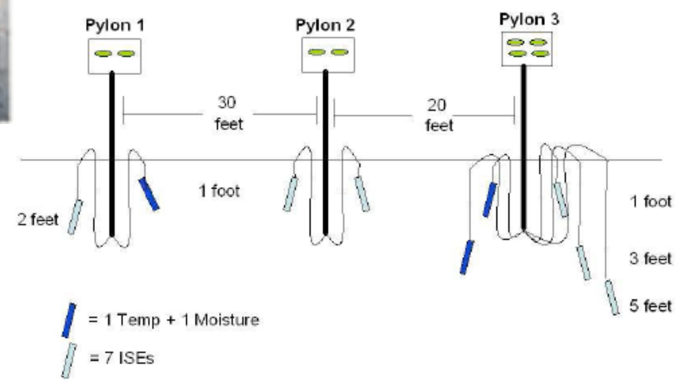

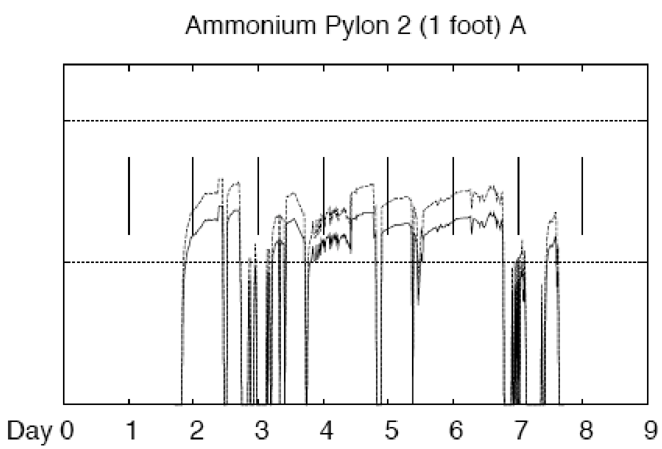

Time-varying Calibration (pre- and post-)

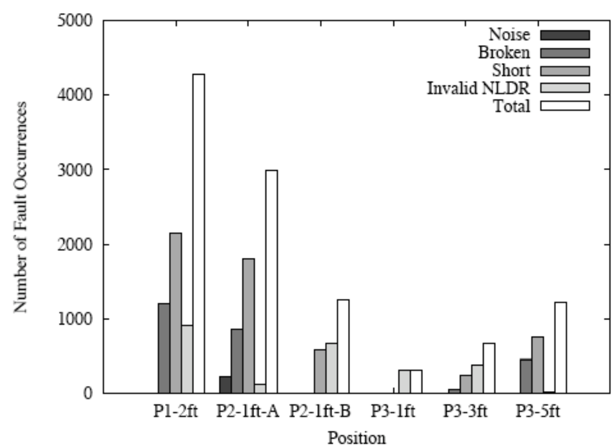

		Bertrand-Krajewski's Reliability Analysis **Reliable**	**Unreliable**
Rule-based	**Not Faulty**	12,138	581
	Faulty	82	8,123

Balzano & Kohler, 2006

Faults at James Reserve

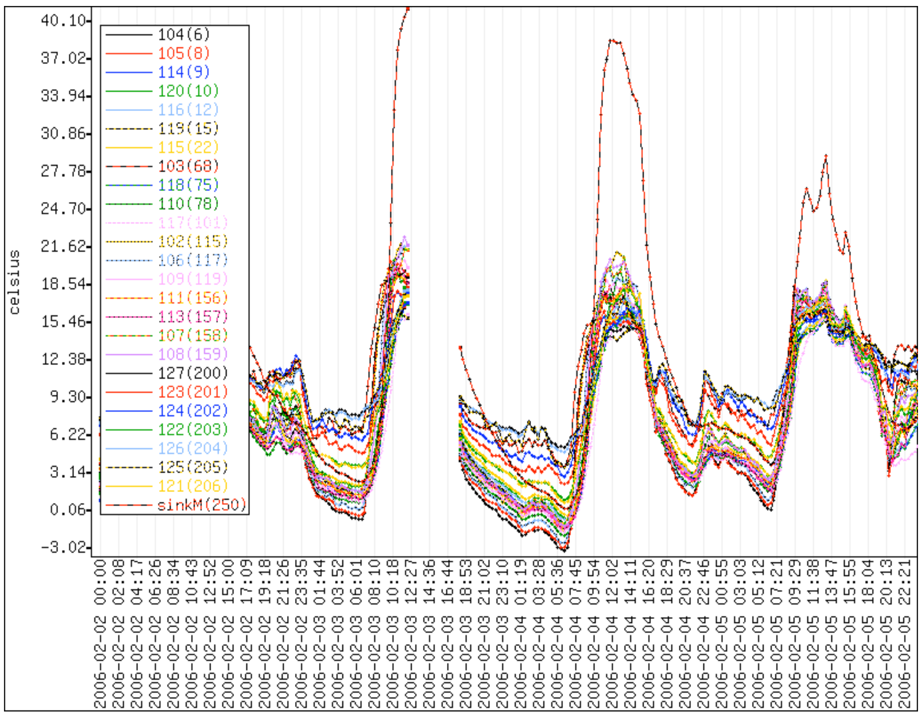

Unreliable Sensor Network

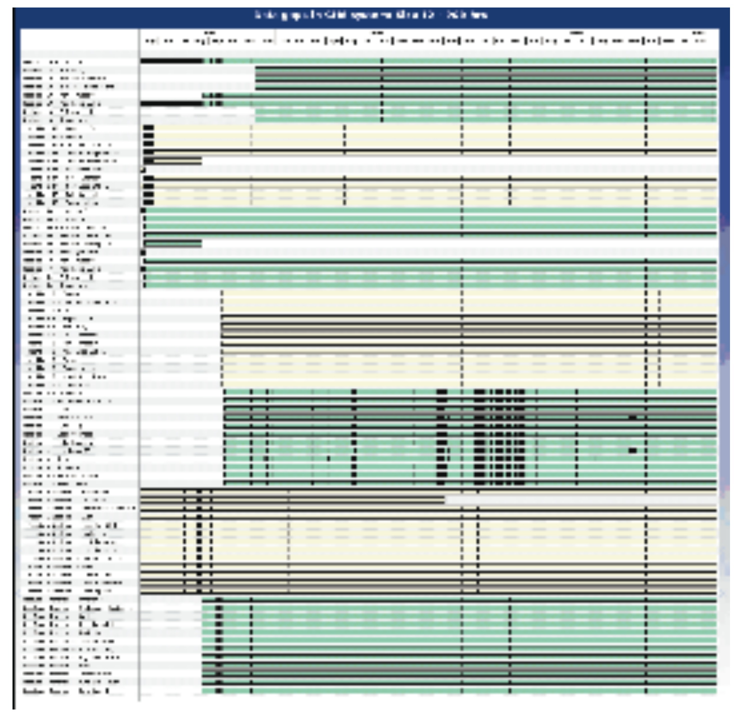

Even Data Loggers Not Reliable!

Balzano & Kohler, 2006

Faults on a Volcano in Ecuador [WLJ+06]

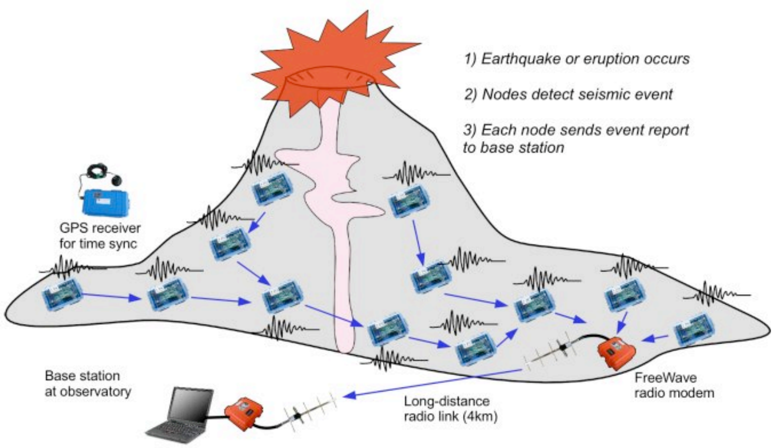

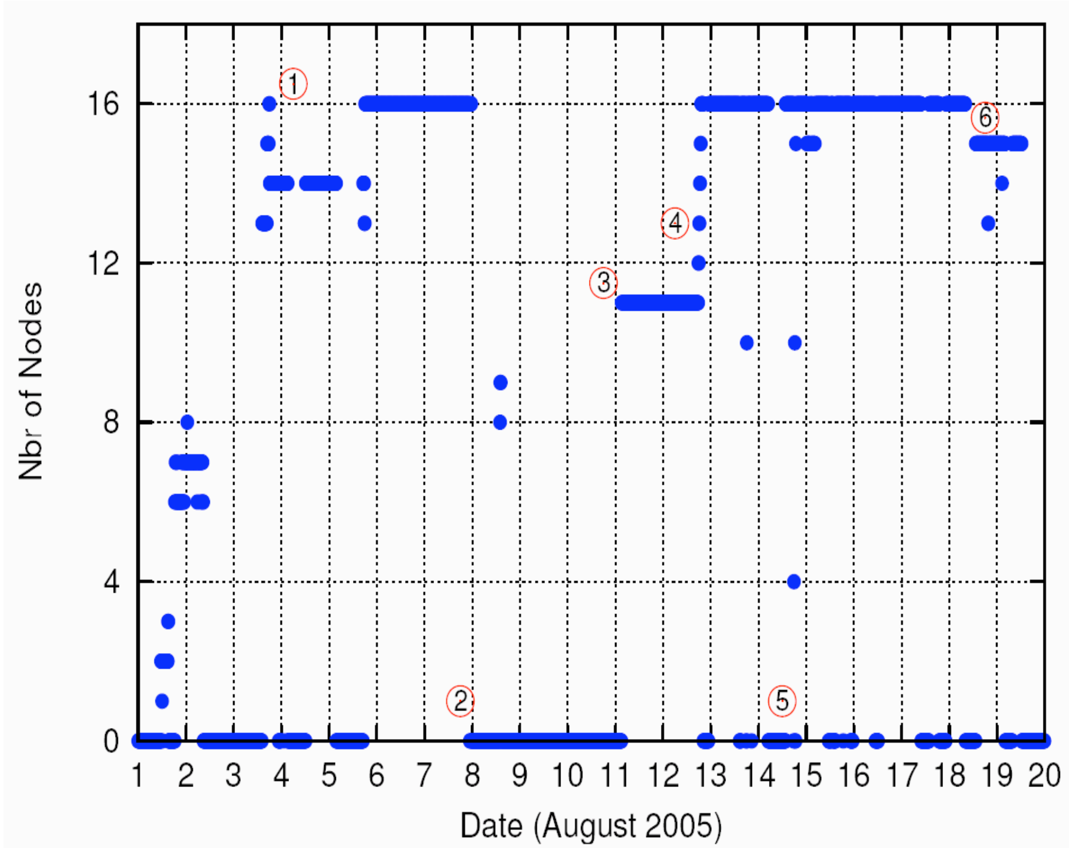

Balzano & Kohler, 2006

Robustness in Embedded Sensing Systems

- People pay for robustness in other systems
  - Higher quality hardware
  - Technicians to monitor the data
  - Wired infrastructure
- In sensor networks when we pay, we pay for scale
- The burden on software and algorithms has increased
- Robustness in sensor networks requires research and engineering

Balzano & Kohler, 2006

What is needed for integrity monitoring?

- Models
  - Modeling of faults, drifts, offsets etc.
    - model system anomalies so that they can be identified
  - Phenomenon Modeling
    - model the physical phenomenon being observed so as to obtain prior information about expected measurements
- Algorithms
  - **Detection**: to identify occurrence of integrity problems
    - source scoring and signature analysis; reference samples and sensors; actuated auditors; self-awareness sensors
  - **Resilience**: to tolerate occurrence of integrity problems
    - scoring of sensor data; robust estimation, aggregation and fusion; multi-scale and multi-modal algorithms; reputation-based mechanisms; data cleansing
  - **Diagnosis & remediation**: to identify the cause and fix integrity problems
    - reorient, reposition, or re-calibrate sensors; replace or add sensors; reconfigure software
    - how often or when to do this?
- Systems
  - Node platform hardware support
    - What hardware features will make detection, resilience, and remediation easier?
  - System software support
    - What system software mechanisms and protocols will make it easier to create resilient system, and easier to recover?

Rule-based On-line Tools for Data Integrity: *Sympathy & Confidence*

	Sympathy	**Data Integrity**	**Confidence**
Detect	Track end-to-end data *quantity*	Track end-to-end data *quality*	Both
Diagnose	*Data Flow Rules* identify locations data could be lost	*Hardware Rules* identify locations data could be corrupted	Both
Refine & Adapt	-----	-----	Action-Refinement Probes + Database
Remediate	User Actions	User Actions	Both

Nithya Ramanathan

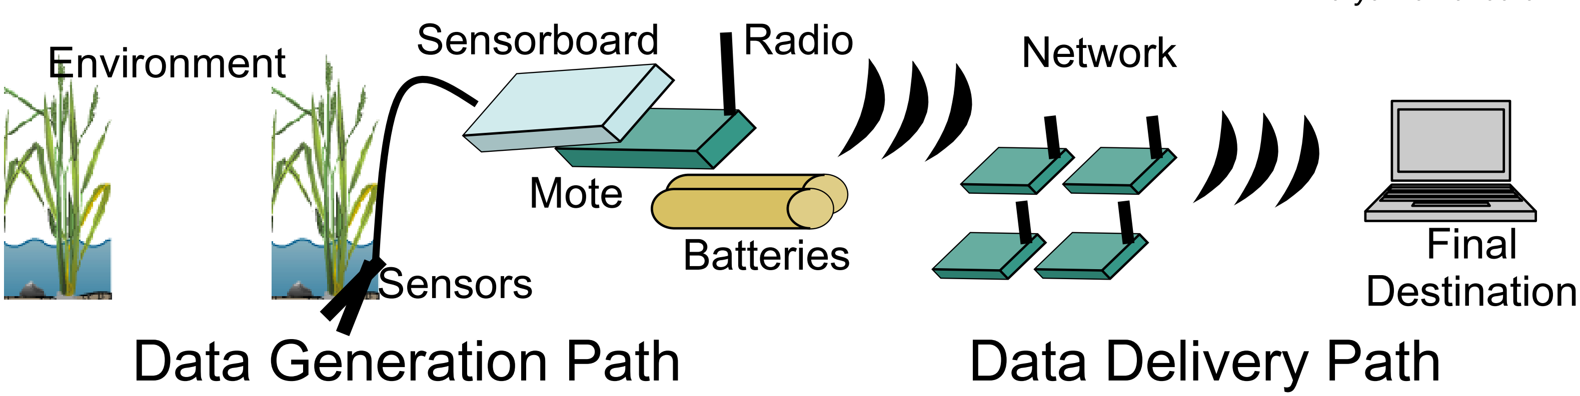

Ramanathan, 2006

Fault Detection and Diagnosis

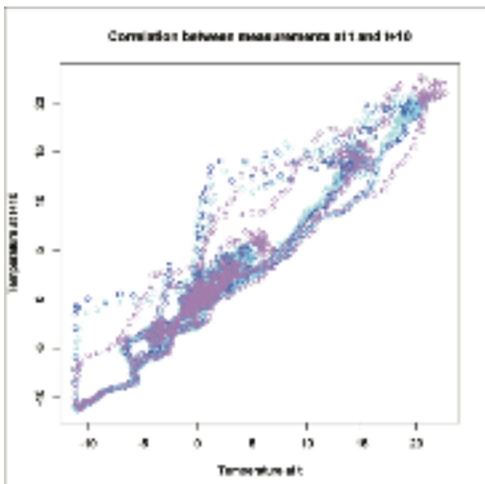

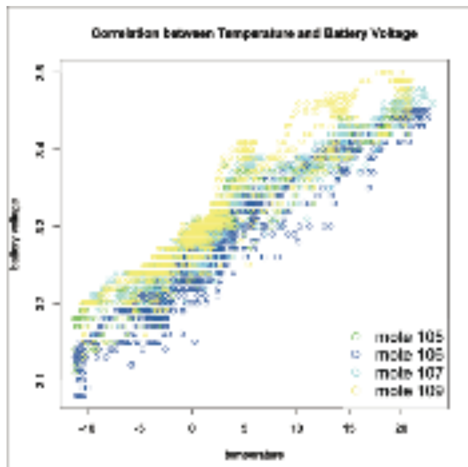

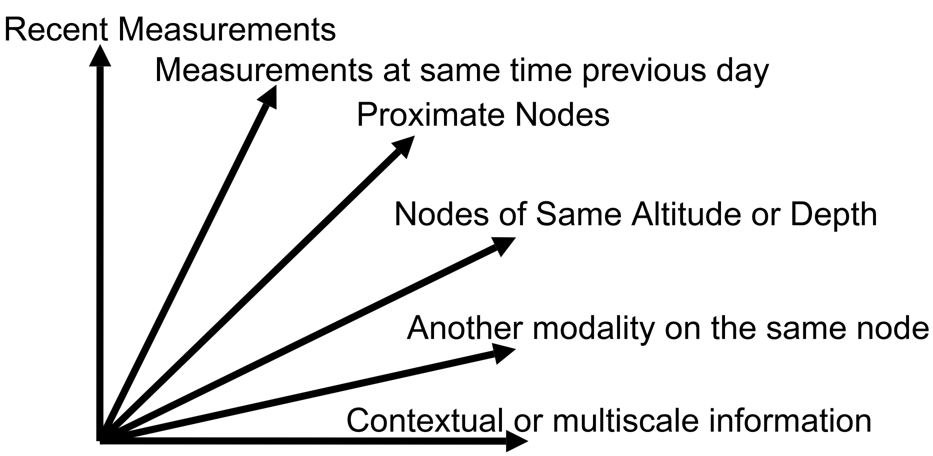

- Recognize a potential fault when sensor data
  - deviates from what is plausible
  - matches something implausible
- Establishment of a reference
  - external high-quality sensor
  - injecting controlled stimulus
  - model of the phenomenon
- Challenge: model-based approaches
  - physics-based constraints and statistical correlations among different variables
  - declare a fault when sensor reading violates the constraints or correlations
  - variable space is high dimensional, and signal processing techniques may provide the needed efficiency

Balzano & Kohler, 2006

Model-based Calibration of Soil Moisture Measurements [Balzano & Margulis]

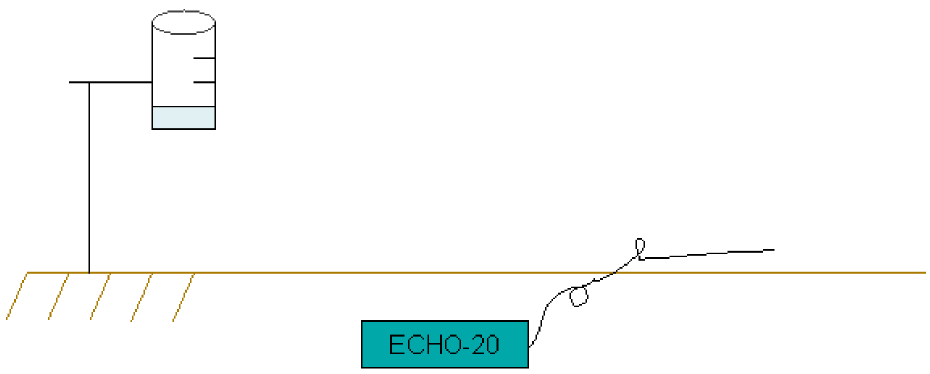

- Sensor measurement model
  - Gain and bias: $M_{est} = \alpha * V_{meas} + \beta + \text{noise}$

- Autoregressive model of surface moisture at a point near the surface
  - moisture drainage, model error, precipitation

$$y_k = \delta y_{k-1} + \sqrt{dt}\sigma\rho q_k + precip_{k-1}$$

- Ensemble Kalman Filter used to track the pdf of the state of the dynamical system using Monte Carlo

Time Series Forecasting for Joint Fault Detection and Efficient Data Collection [Tulone]

- Goal: answer queries at a sink together while detecting faulty sensor data

- Approach
  - each sensor node learns a local AR model for its measured time series of samples
  - model parameters sent to sink
  - sensor uses AR model to
    - detect outliers and potential malfunctions
    - decide when to update the model and send new parameters to sink
  - sink uses models to
    - cluster sensors making similar measurements and select a cluster leader
    - answer queries using cluster leader's model and periodic readings received from it
    - verify fault report against global view, and diagnose cause and geographic scope

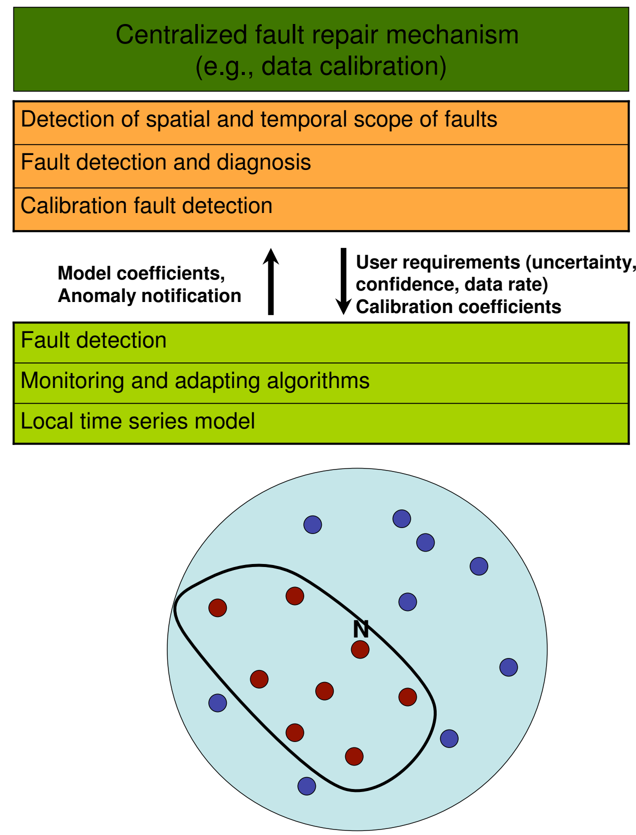

Not just Data Integrity, but <u>Control</u>

Towards Participatory Sensing Applications

From Science Problems to Human Concerns

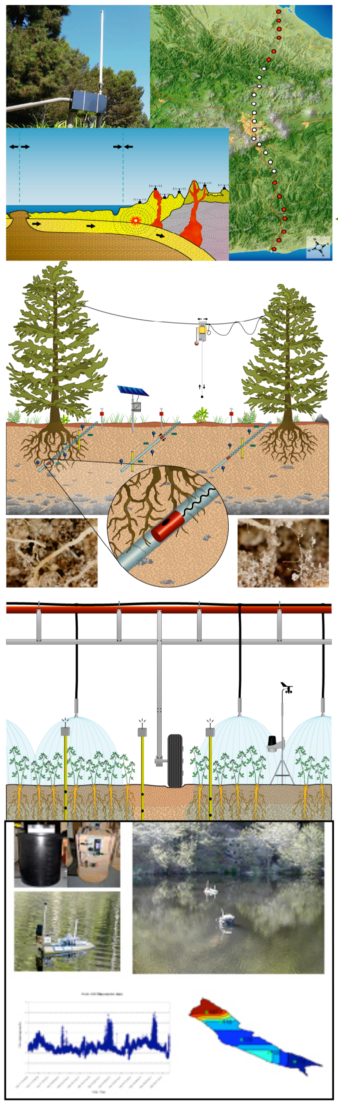

- **ENS is revealing the previously unobservable in science applications**
  - Multi-scale data and models to achieve context, and in network processing and mobility to achieve scalability (communication, energy, latency)

- **Automatically geocoded and uploaded participatory sensing data promises to make visible human concerns that were previously unobservable…or unacceptable**
  - Data collection & documentation vital in public health, urban planning, natural resource management, culture etc.
  - Urban sensing applications will leverage the billions of cell phone acoustic, image, bluetooth-connected location-aware sensors
  - Searchable sensor feeds and blogs with geotime tags to achieve context, and in network processing for privacy and control

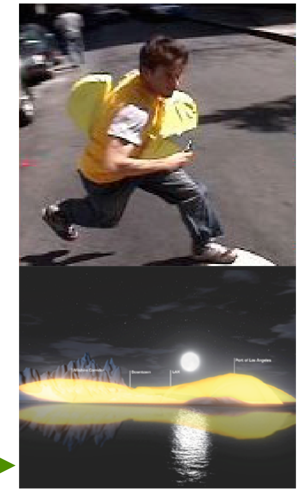

NOKIA Lifeblog

Range of Participatory Sensing Applications: *Urban, Social, Personal...*

Towards an internet of public, private, personal observatories

'Citizen-initiated' sensing, publishing, sharing, analyzing

- Every-day user in the act of gathering, analyzing, and sharing local knowledge
- Participatory urban planning (traffic, sound, road conditions etc.)
- Slogging (sensorbase.org+ESPML)
- Place-aware social networking
- Distributed documentary – journalism
- Community-built histories, the new 'local library'

'Directed' sensing campaigns

- Professionally authored observation campaigns for community data gathering
- Self-administered health diagnostics
- Eco-PDA
- Public/community health

NSF FIND and Nokia SensorPlanet projects @ CENS: PARTISAN, CAMPAIGN

Data Privacy and Quality

- Data contributors like to exercise control over resolution and sharing rules
  - Reveal location only in terms of ZIP code
  - Reveal time only in terms of hour or day
  - Reveal only as part of a large enough group
  - In some cases, the object of sensing may have a greater stake
- Data and its context more valuable if verified
  - Subscribers want to know when and where the measurements were taken, and whether they can be corroborated
  - Verification at a resolution the contributor is comfortable with
- Privacy and quality in terms of controlling the resolution of revealed information under different sharing rules, increased assurance via attestation of sensor values and context, and ability to verify authenticity and validity via audit trail
  - Must be automated system components (in-network, back-end)
  - Must be cryptographically secure against malicious threats

Participatory Sensing System Architecture

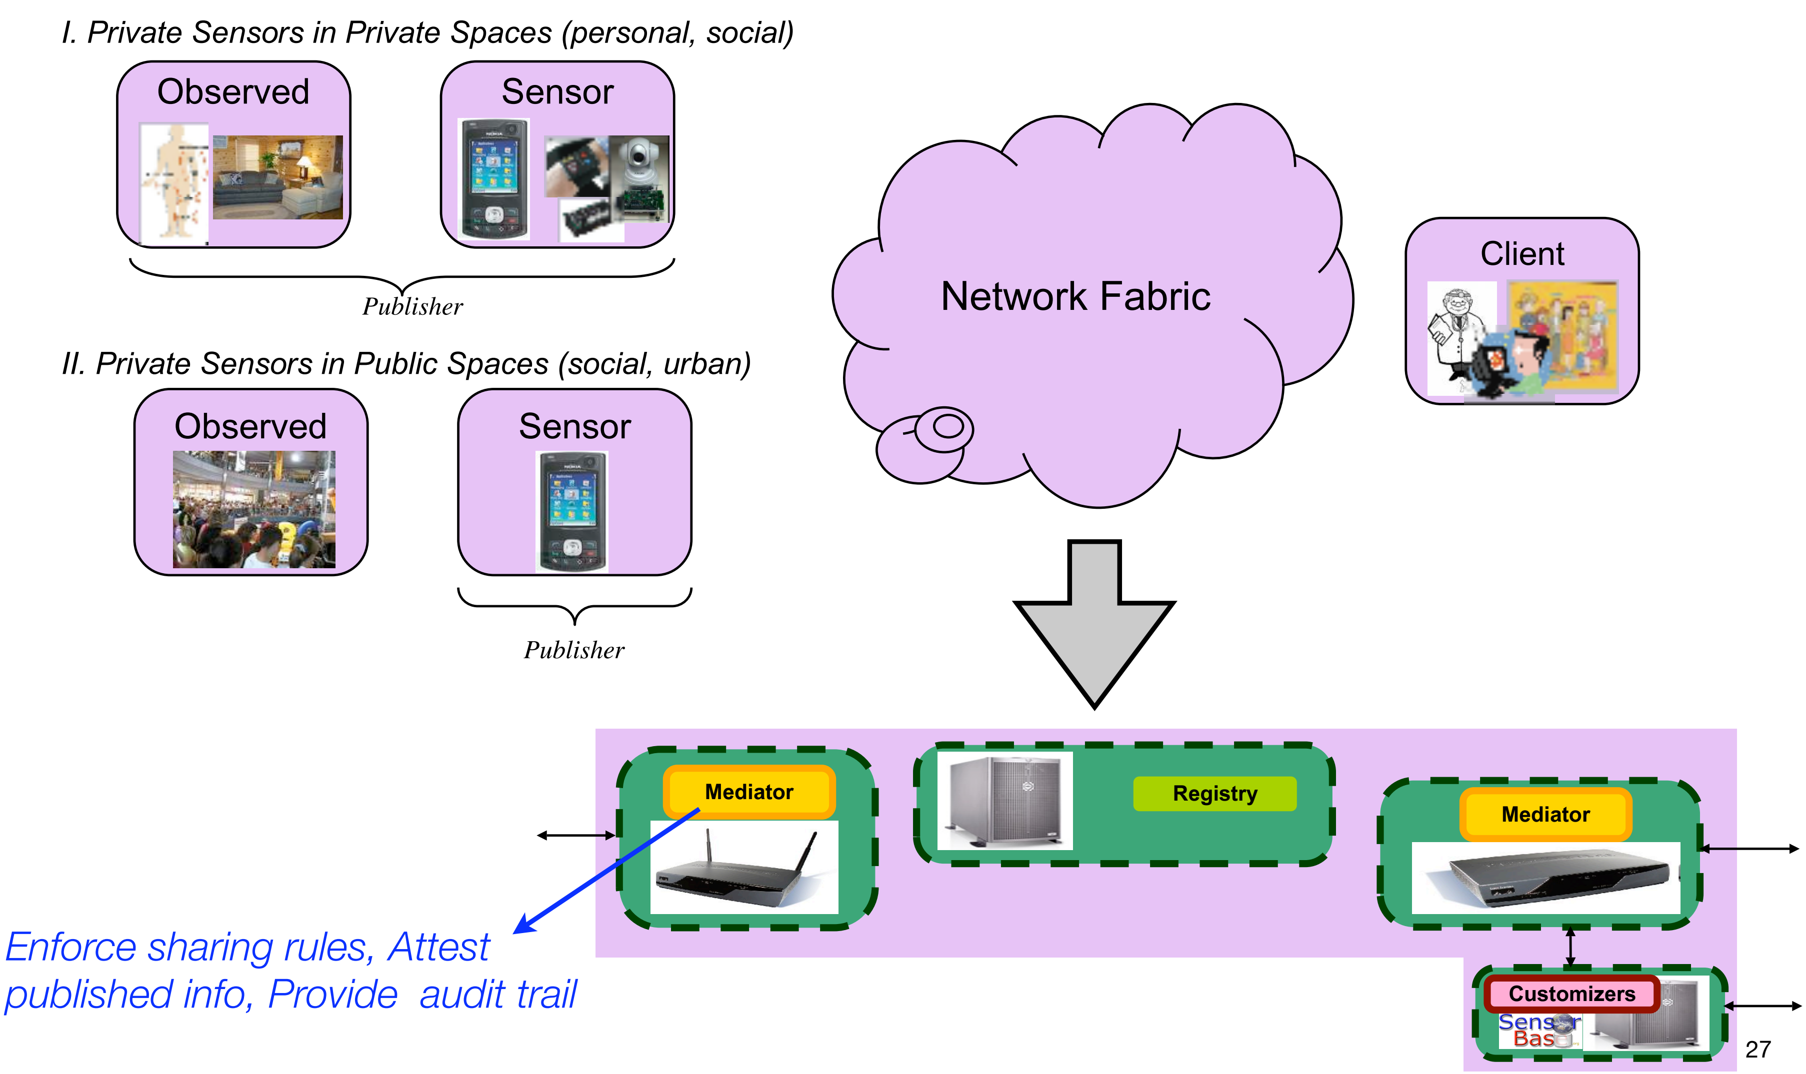

Enforce sharing rules, Attest published info, Provide audit trail

Conclusions

- A decades worth of sensor network research has yielded lots of interesting science and military applications

- In a transition now:
  - Smart dust ⇒ Mobile, Multi-scale, Multi-modal
  - Deployments ⇒ Multi-user Observing systems
  - Science, military ⇒ Urban, social, personal, enterprise
  - Resources, communications, autonomy ⇒ Integrity, sensing, participatory and interactive

- Profound impact of cell phone: a wireless sensor with two-legged mobility
  - Nokia's SensorPlanet effort

- Implications on and integration with Internet
  - Publishing and sharing sensor data: “slogging”
  - Architectural support for verification, privacy, selective sharing
  - Application authoring

- Algorithmic challenges lie in integrity and privacy