Center for Information and
Systems Engineering

Resource Management in Sensor Networks

David Castañón

Venkatesh Saligrama

Boston University

Sensor Networks

- What are sensor networks?
  - Interconnected networks of sensing devices, typically wireless
  - Heterogeneous sensing, communications, processing capabilities per node
  - Spatially distributed to create observability, redundancy
- Important Aspects
  - Power requirements for comms >> local computation
  - Sensing can be controlled
  - Local knowledge, with variation across sensors
- Challenge: Efficient use of available sensing, communications and computation resources to accomplish global sensing tasks
  - Focus on application-level resource management

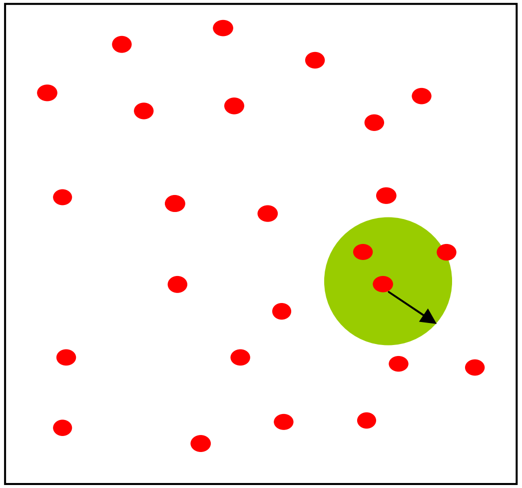

Resource Management

- Resource Management in Sensor Networks
  - Communications: what/when to communicate
    - Tradeoff between cheap local computation, compression vs performance at global center
  - Sensing: where to measure, how
    - Steerable, multi-mode sensors; active vs passive
  - Processing: Where to compute
    - Local vs central processing
  - → Application level resource management
    - Contrast with middleware resource management: power per message, routing, etc given traffic requests
- Issues
  - Heterogeneous sensor networks
    - Different power/processing/sensing capabilities in nodes
  - Diverse excitation
    - Activity not distributed uniformly
  - Distribution of Sensor Management
- Focus of talk
  - Review recent results that address aspects of these issues

Some Ongoing Projects

BU Sargent Camp Test Site
Ecological Monitoring

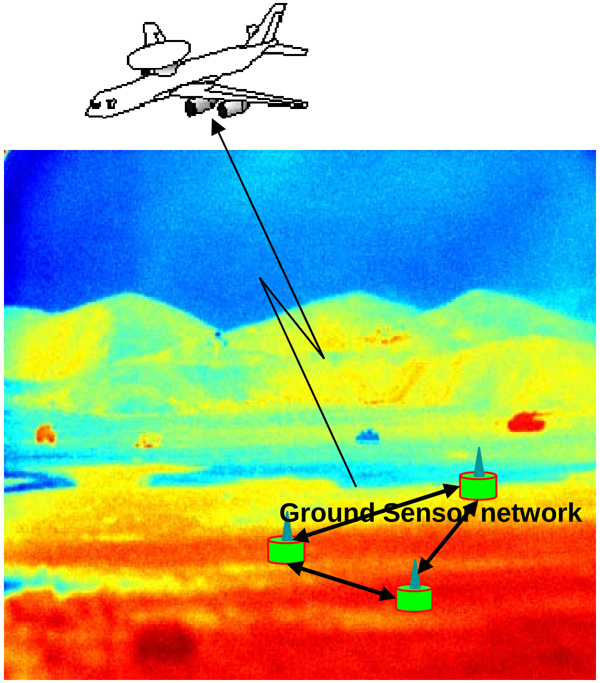

Intelligent Sensing and Relaying

Two Distinct Classes of Problems

Multiple locations view same phenomena
- **Requires information coordination, fusion**

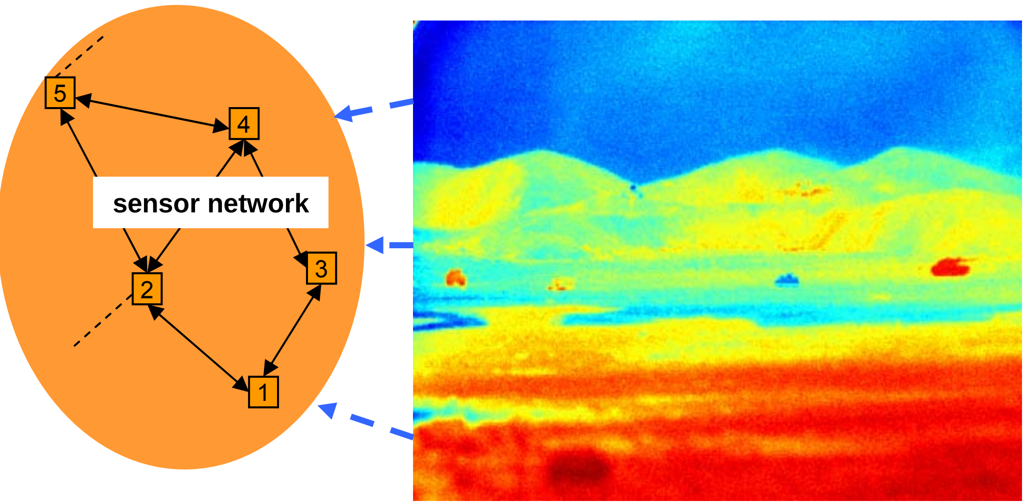

Setup:

- **N sensors, Comm. links**
- **Observing same scene (image)**
- **Discover ``abnormal'' activity**

Examples: Tracking, Detection, Estimation, …

Locations have little overlap in coverage
- **Coordination in processing?**

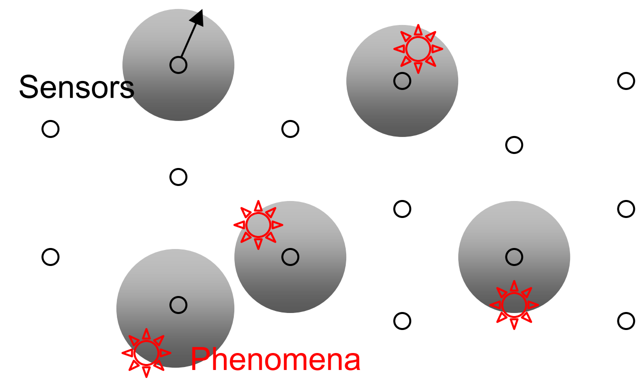

Setup:

- **N sensors, local obs.**
- **multiple, unknown targets**
- **search for targets/locations**

Examples: mine detection, environmental monitoring, …

Image Fusion: What to Communicate

Image at proximal IR Sensor

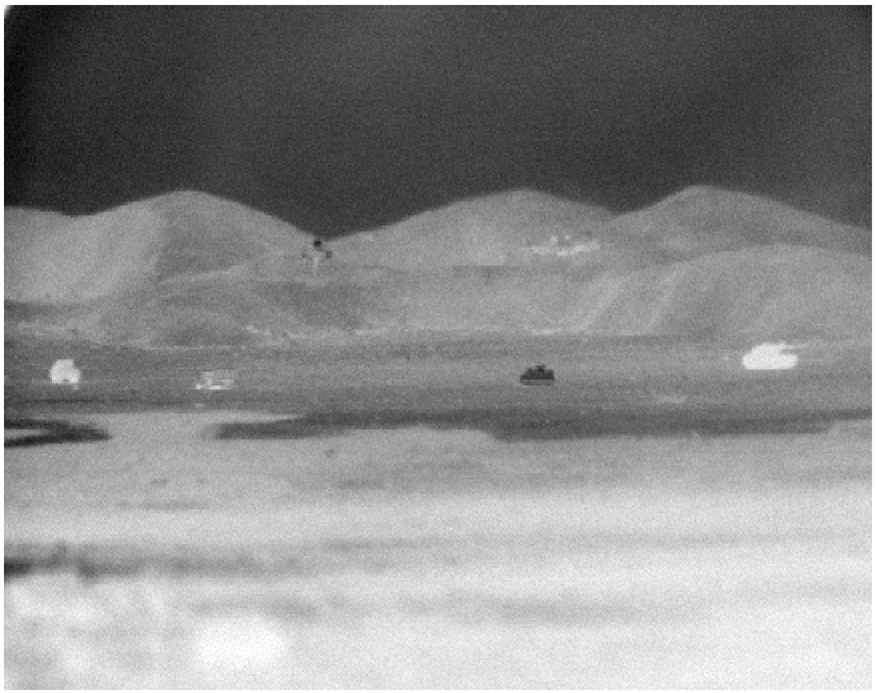

Image at distant IR Sensor

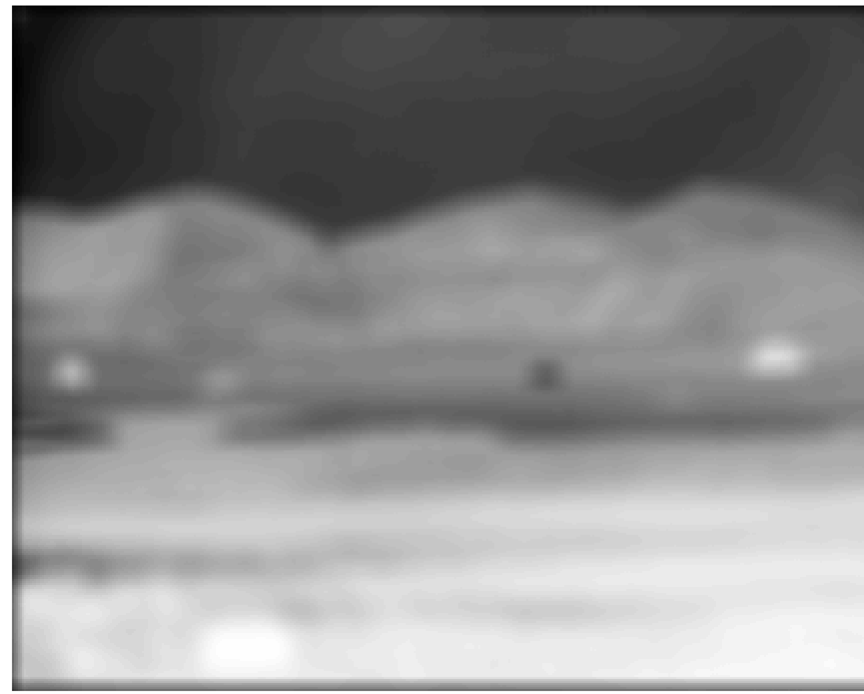

- **Problem:**
  - 2 Sensors, Discover Targets, Min False alarms/miss
- **Issue:**
  - Image, observation contains too much data
  - 1024 X 1024 @ 16bits @ 28Kbps = 10 mins

Transmit only relevant information

Proximal IR Sensor

Distant IR Sensor

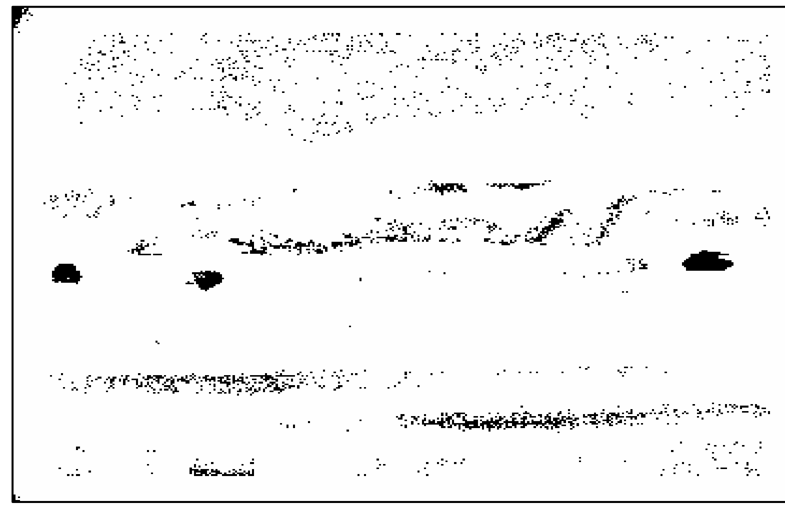

- Most of imagery is of no interest
  - CFAR processing to detect anomalies
- Problem:
  - *what/how much* info to send from nearby sensor to reliably detect at far off sensor

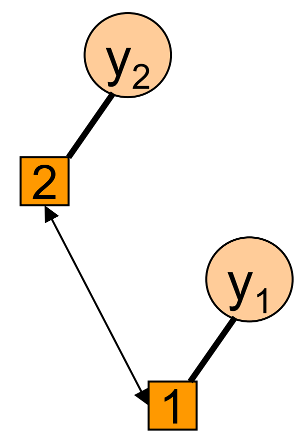

IR Sensor Collaboration: Lessons

- Relevant info transmission can reduce bandwidth
  - Information theory model: Transmitting noisy info through noisy channel with noisy side-information
  - Local processing for CFAR, and coding/compression yields equivalent performance with reduced information rates

Bits	0	0.1%~1sec	0.4%~ 4sec	Image Coding with Side Info 7min	Full 10 min
False alarms	7	4	0	0	0
miss	1	1	1	1	1

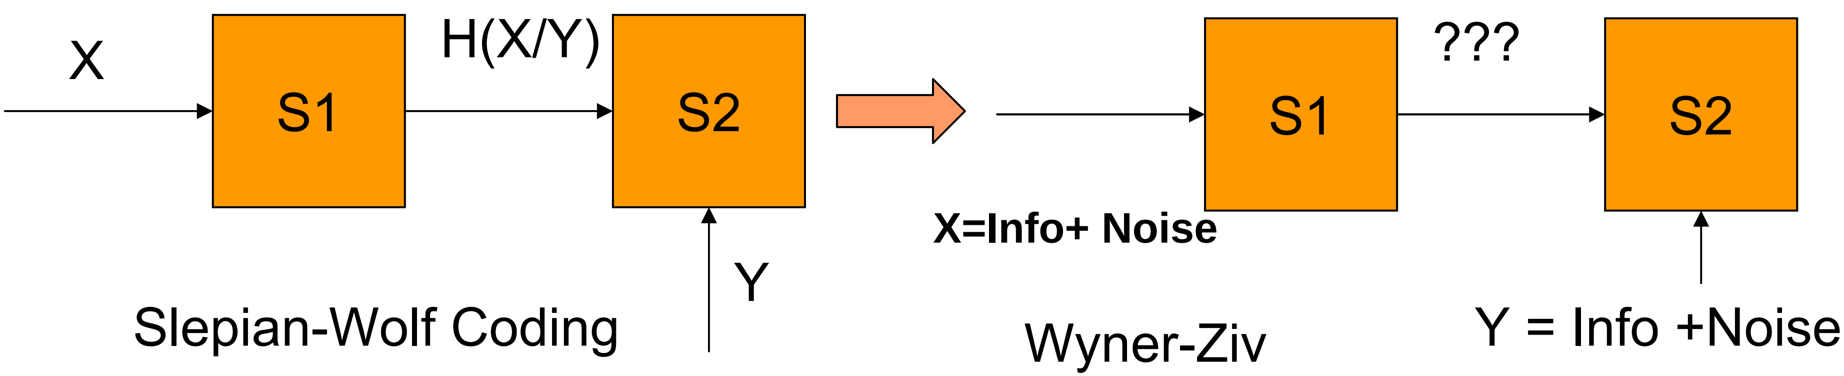

Extension to Networked Settings

- Strategy:
  - decisions: $m_{j\rightarrow k}(t)$
    - transmit relevant info that is new to Neighbor nodes
    - belief propagation variation
- Advantages:
  - Local transmission
  - Local models
- Issues:
  - No fusion center
  - Loopy graph: info feedback, convergence?
- Objective
  - Seek consensus, Centralized performance, min. energy/node

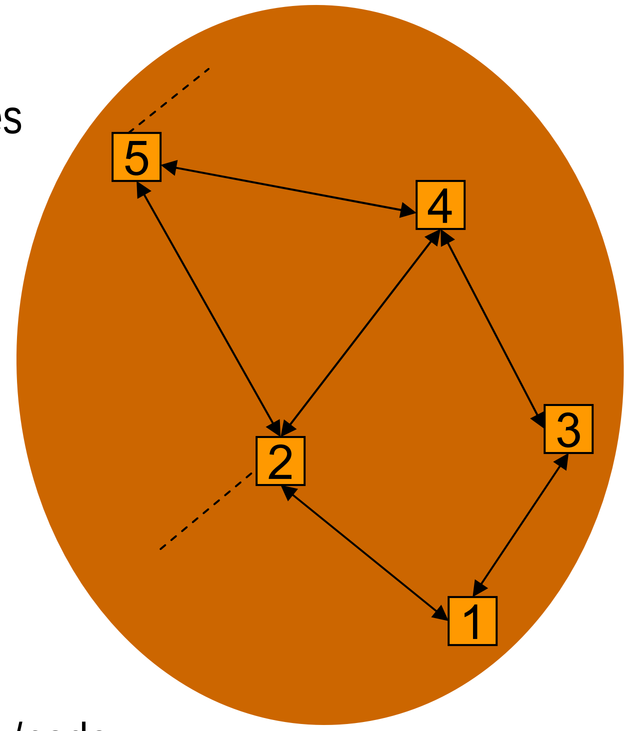

Network Result

- Theorem: For any arbitrary graph topology there exists a distributed localized algorithm that:
  - Achieves consensus at each sensor.
  - Consensus = MAP estimate of phenomena
    e.g. Most likely classification
- Can stop communications based on info theory concepts: value of information
- Example:
  - network of 400 sensors
  - max. of 10 Targets
  - noisy information
  - Convergence to common estimate
- Energy Scaling
  - (Hierarchical → log(N) joules/node)

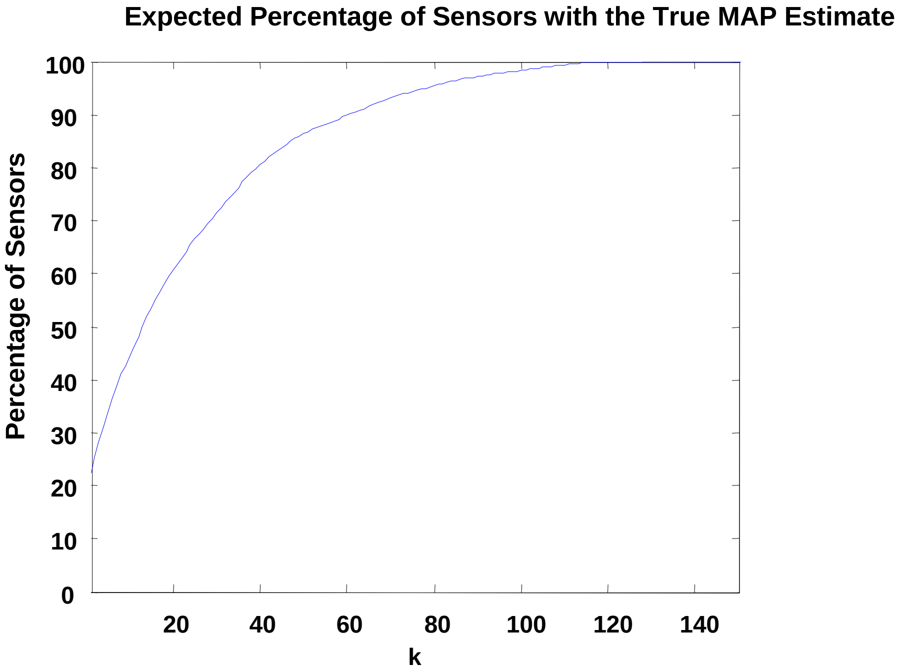

Estimation with Unreliable Comms

- N sensors observing common dynamical system
- Comms links have reliability p_i (.8, not .999)
  - Messages dropped, not retransmitted
- Protocol: transmit measurements
  - Result: Thresholds relating p_i to boundedness of expected steady state estimation error (covariance) (Poolla, Sastry, Goldsmith, …)
- Alternative: Local processing, transmit statistics
  - Based on distributed estimation theory (Speyer, Chong, Willsky, Castañon, Bar-Shalom …)
  - Balance of local and global processing
  - Result: Much better estimation performance, can tolerate much lower p_i before losing boundedness of expected error
  - Result 2: Can drop data rate 50%, still achieve significant improvements in estimation performance, and tolerate lower reliability links

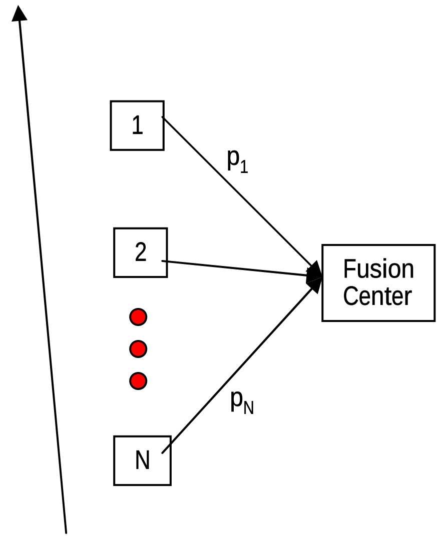

What to Measure

- Local problem: Multi-modality, multi-mode sensors
  - Passive modes with small energy consumption
  - Multiple active modes with greater energy and bandwidth consumption
  - Constraints on total energy/bandwidth consumption
  - Multiple areas to monitor for detection and classification
- Related problem: Sensor Management
  - Limited theory available for multi-sensor, multi-modality sensing
  - Typical solutions based on incremental, greedy approaches
  - No performance prediction

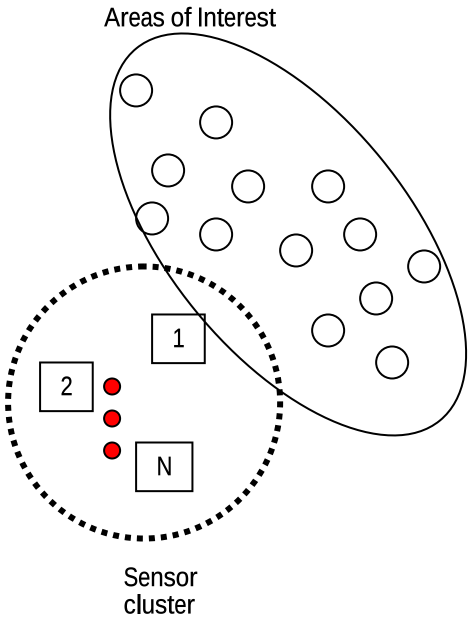

New Results

- Approach: Solve as dynamic decision problem with uncertainty and feedback
  - Allows sequencing of multi-sensor actions plus local processing
  - Underlying information state: information available on potential contents, class of area of interest
  - Control evolution of information state through sensing actions and processing of measurements
- Problem
  - Large dimension of potential information states: exponential in number of areas!
  - Need scalable approach
- Result: New formulation that allows for hierarchical decomposition
  - Linear scalability with number of areas of interest
  - Caveats: idealized sensor models with narrow fields of view
  - Provides bound on best achievable performance!!!
  - Provides algorithms that perform close to bound, better than alternative sensor management ideas

Hierarchical Approach

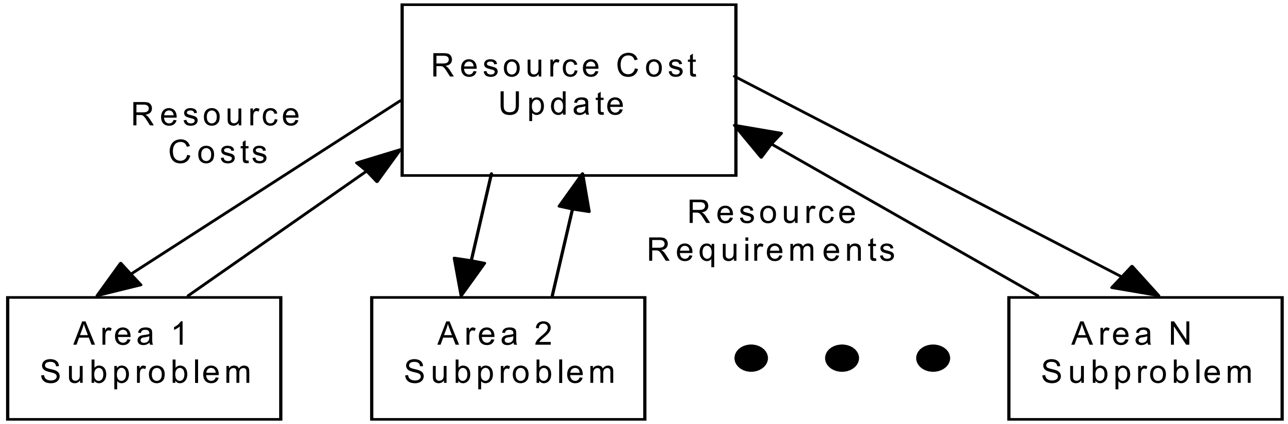

- Resource costs: dual variables for consuming sensor resources for different sensors
- Uses combination of nonlinear programming, integer programming and stochastic dynamic programming
- Centralized resource cost updates

Results

- 100 Areas of Interest
  - Detect and classify objects in areas
  - Two sensor modes
  - Limited resources

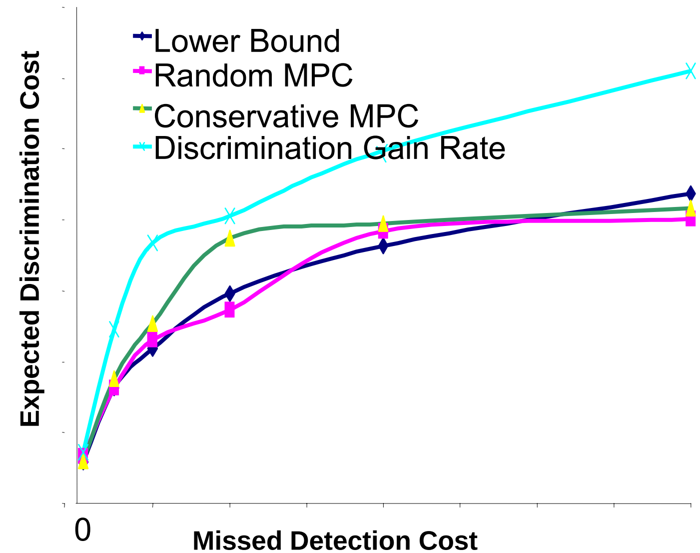

- Mixed strategies
  - Strategies use H (cheap) measurements first to prune candidates for I (expensive) measurements

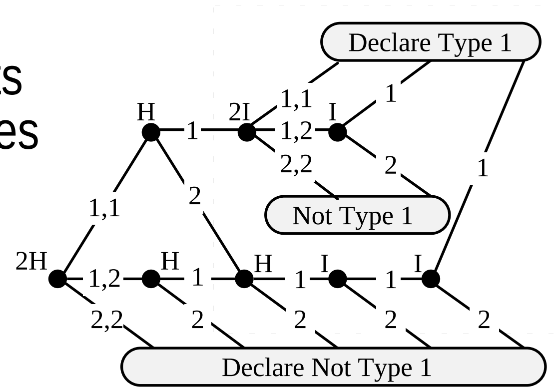

High Resource Decision Tree, $u_{high}(t)$

$r\left(u_{high}(t)\right) = 0.4$ sec/object

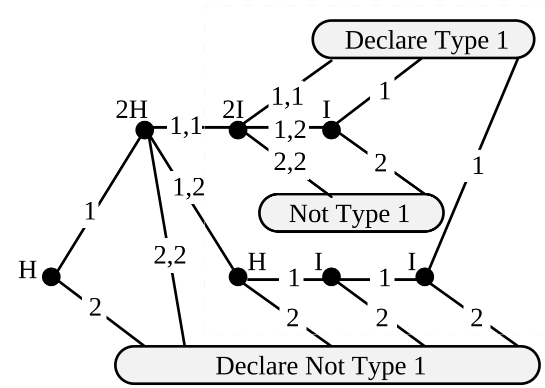

Low Resource Decision Tree, $u_{low}(t)$

$r\left(u_{low}(t)\right) = 0.35$ sec/object

Where to Compute

- Network of passive sensors tracking targets
  - Which sensors should report measurements?
  - To what processing center?
  - Williams, Fisher, Willsky,
- Objective: Maximize information gained
  - Subject to comms constraints
- Problem
  - Locate where fusion will take place, and which sensors will report
  - Switching fusion sensor involves handover of track information → comms cost
- Result: New algorithms using approximate dynamic programming
  - Selects most informative sensors, migrates fusion center in anticipation of new information

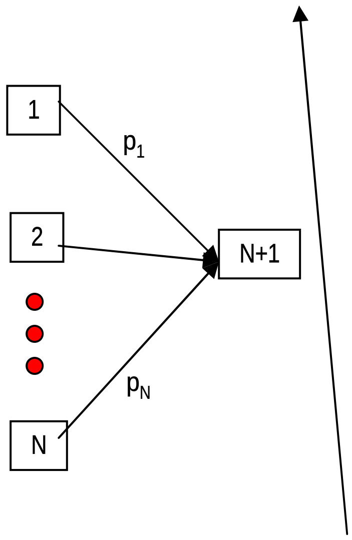

Conclusion

- New ideas available for resource management in sensor networks
  - Exploitation of local computation and distribution of computation for communications reduction, robust performance
  - New algorithms for distributed fusion with guaranteed performance properties
  - Approaches for selecting which sensors, modes to use, where to communicate under resource limitations
- Common theme: exploit local processing
- Main limitation: Efficient resource management requires global information about sensor statistics
  - Can't determine local processing without knowing about global topology, capabilities
  - May require side information: partial feedback of fusion results
  - Also: focus on academic models so far → Need extension to realistic sensor models, suites for applications