Information-Based Swarm Management

Chris Kreucher[1,2], Keith Kastella[1], and Alfred O. Hero III[2]
[1]General Dynamics Advanced Information Systems
[2]The University of Michigan, Dept. of EECS

Overview

- Have extended information-based sensor management to swarms
- Basic idea is to use artificial physics ideas to couple swarm behavior to information gain field
- Initial demonstration assumes centralized fusion
- Methods can be extended to fully decentralized control

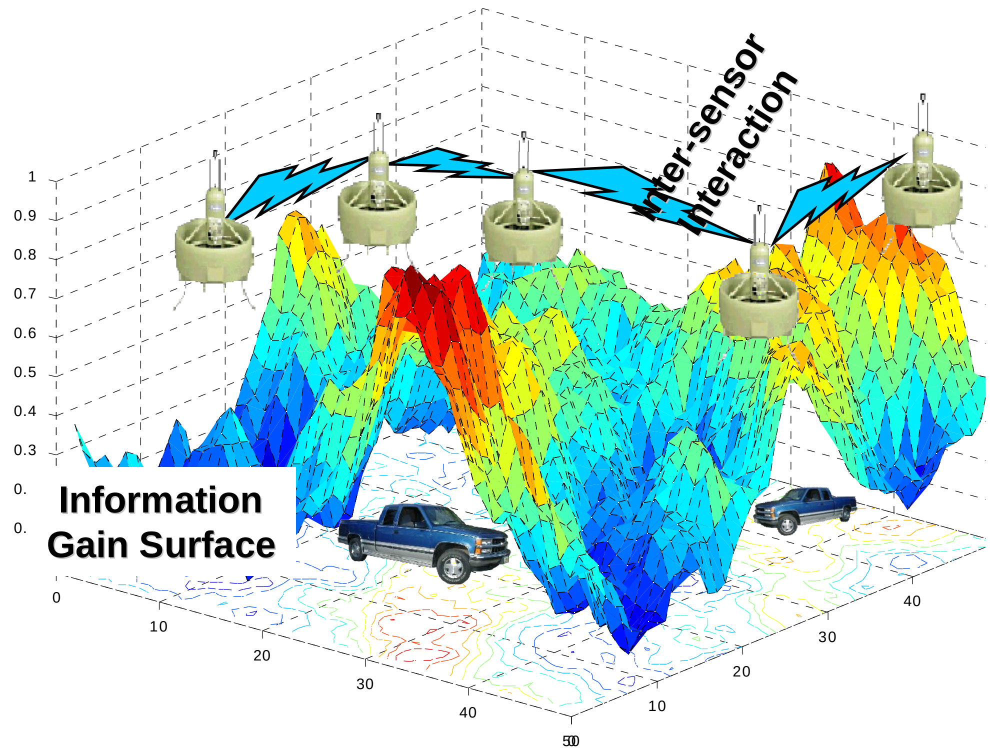

Distributed Swarm Management

Goal: Distributed, Decentralized, Limited Communication swarm management

Model problem:

- Unknown number of mobile ground targets
- Sensors to determine:
  - number of targets
  - state of each (position and velocity)
  - Sensors "hover" at a fixed height, stare directly down
  - Perform measurements with known Pd & Pfa
  - Sensor management problem is to determine the best move for each sensor
- Seek methods to minimize communication with complexity that scales linearly with sensor number

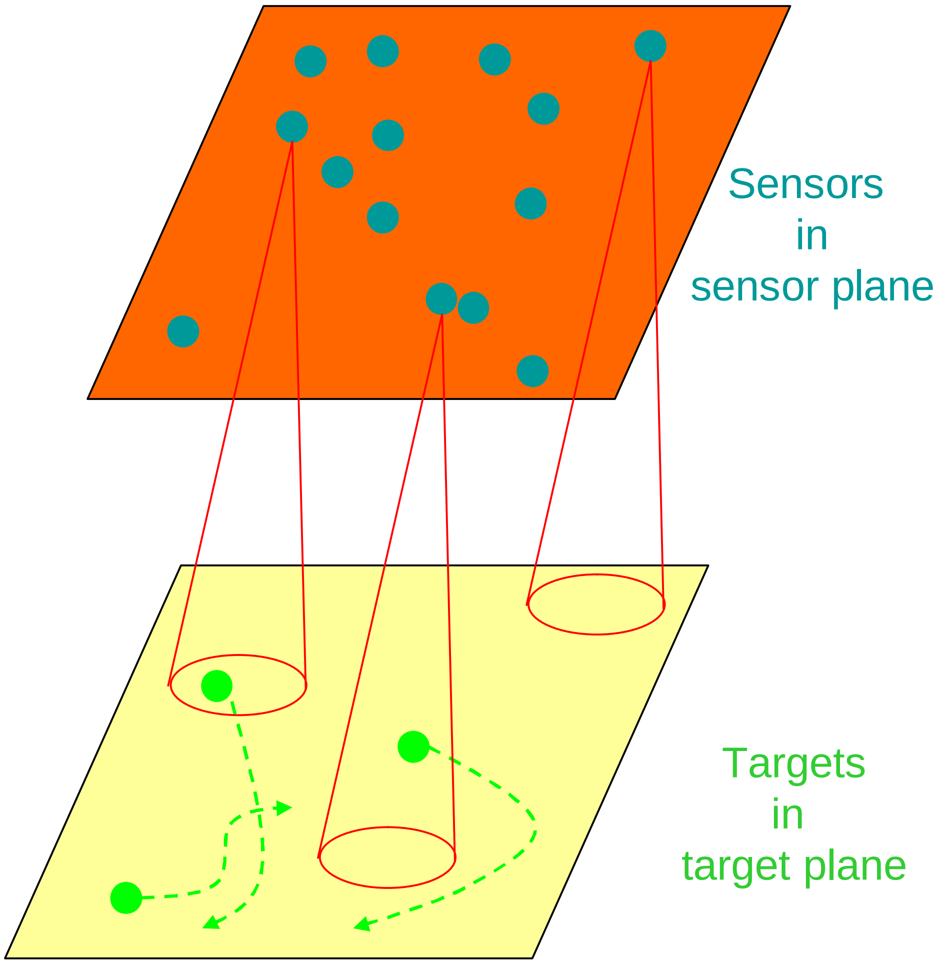

Our Approach

Our approach: Combine artificial physics and information theory to evolve the position of sensors in response to the environment

- Underlying model taken from molecular dynamics
- Sensors obey Brownian dynamics model

- New feature is coupling information field to dynamics
- *Artificial Physics :* A sensor feels an attraction/repulsion from other sensors
- *Information theory* used to compute information gain each possible future sensor location expected to yield

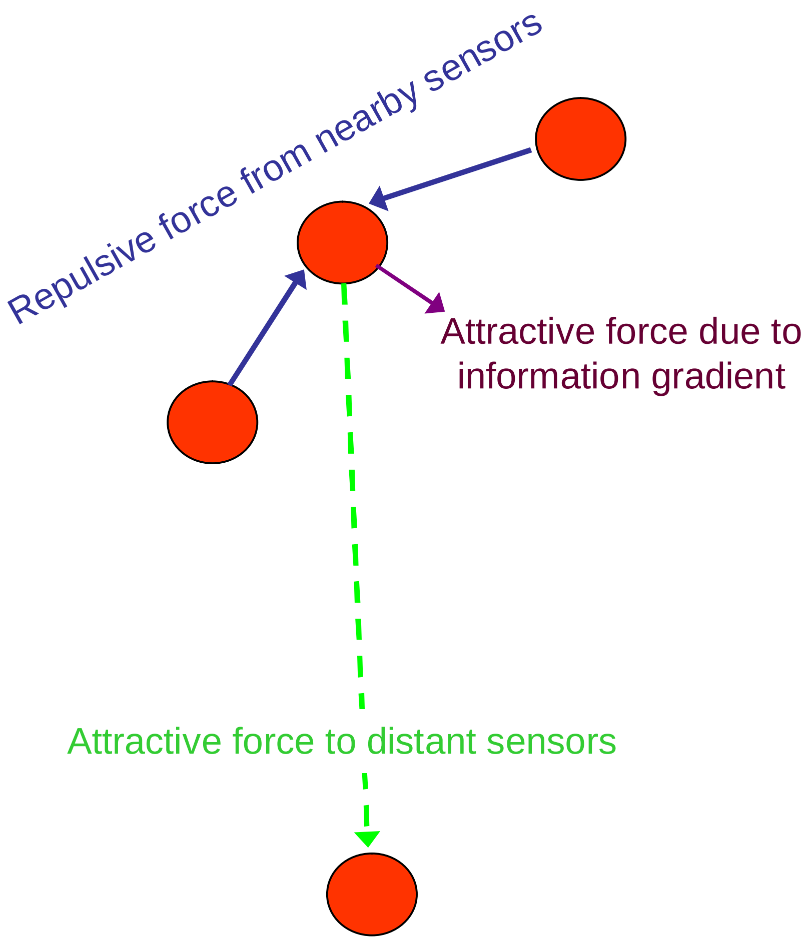

The Ingredients of the Approach

Four ideas:

1. Joint Multitarget Probability Density Estimation
2. Information theory predicts maximally beneficial sensor actions
3. Artificial Physics provides heuristic to drive sensor positions
4. Coupling of information potential and AP potential creates a single force that drives the sensors

The Joint Multitarget Probability Density (JMPD)

- To compute areas of high information gain we must model environment
  - Mathematical model must reflect uncertainty about target number, class, and kinematic state.
- **Central element summarizing system knowledge is *joint multitarget probability density (JMPD)***

$$p(\mathbf{X}^k, T^k|\mathbf{Z}^k) \equiv p(\mathbf{x}_1^k, \mathbf{x}_2^k, ...\mathbf{x}_{T-1}^k, \mathbf{x}_T^k, T^k|\mathbf{Z}^k) \qquad T = 0 \cdots \infty$$

$$= p(\mathbf{x}_1^k, \mathbf{x}_2^k, ...\mathbf{x}_{T-1}^k, \mathbf{x}_T^k|T^k, \mathbf{Z}^k)p(T^k|\mathbf{Z}^k)$$

- **JMPD estimated from noisy measurement sequence**
  - $\mathbf{x}$ is a description of kinematic state (e.g., position and velocity) and ID of a target (e.g., tank)
  - T is the number of targets, which is to be estimated along with the $\mathbf{x}_i$'s
  - $\mathbf{Z}$ is the measurements taken over *k* time steps, e.g.
- The JMPD is hybrid continuous-discrete distribution with normalization

$$\sum_{T=0}^{\infty} \int d\mathbf{x}_1 \cdots d\mathbf{x}_T p(\mathbf{x}_1, \cdots, \mathbf{x}_T, T|\mathbf{Z}) = 1$$

The Joint Multitarget Probability Density (JMPD)

- JMPD temporal evolution described by $p(\mathbf{X}^k, T^k | \mathbf{X}^{k-1}, T^{k-1})$
  - Kinematic model includes how existing targets move, how new targets arrive, and how existing targets leave the region
  - Ancillary information that effects target motion incorporated through this model (e.g., map of hospitability for maneuver)
- Sensor described by $p(\mathbf{z}^k | \mathbf{X}^k, T^k)$
  - Sensor model describes how measurements $\mathbf{z}$ couple to target state $\mathbf{X}$
  - Ancillary information that effects the sensor incorporated through this model (e.g., the visibility map)

- JMPD evolves via Chapman-Kolmogorov-Bayes (CKB) recursion

Temporal Update

$$p(\mathbf{X}^k, T^k | \mathbf{Z}^{k-1}) = \sum_{T^{k-1}=0}^{\infty} \int_{\mathbf{X^{k-1}}} d\mathbf{X}^{k-1} p(\mathbf{X}^k, T^k | \mathbf{X}^{k-1}, T^{k-1}) p(\mathbf{X}^{k-1}, T^{k-1} | \mathbf{Z}^{k-1})$$

Measurement Update

$$p(\mathbf{X}^k, T^k | \mathbf{Z}^k) = \frac{p(\mathbf{z}^k | \mathbf{X}^k, T^k) p(\mathbf{X}^k, T^k | \mathbf{Z}^{k-1})}{p(\mathbf{z}^k | \mathbf{Z}^{k-1})}$$

The Particle Filter Implementation of the JMPD

Particle filtering solves JMPD CKB equations numerically.

a. **Particles represent the density** : The JMPD approximated by a finite set of weighted samples

$$p(\mathbf{X}^k, T^k|\mathbf{Z}^k) \approx \sum_{p=1}^{N_{part}} w_p^k \delta(\mathbf{X}^k - \mathbf{X}_p^k) \qquad \delta(\mathbf{X}^k - \mathbf{X}_p^k) = \begin{cases} 0 & T^k \neq T_p^k \\ \delta_D(\mathbf{X}^k - \mathbf{X}_p^k) & \text{otherwise} \end{cases}$$

b. **The temporal update is done via importance sampling**:

$$(\mathbf{X}_p^k, T_p^k) \sim q(\mathbf{X}^k, T^k|\mathbf{X}_p^{k-1}, T_p^{k-1}, \mathbf{z}^k)$$

c. **The measurement update is accommodated by updating particle weights**

$$w_p^k = w_p^{k-1} \frac{p(\mathbf{z}^k|\mathbf{X}_p^k, T_p^k) p(\mathbf{X}_p^k, T_p^k|\mathbf{X}_p^{k-1}, T_p^{k-1})}{q(\mathbf{X}_p^k, T_p^k|\mathbf{X}_p^{k-1}, T_p^{k-1}, \mathbf{z}^k)}$$

The Information Gain Surface

- The Rényi (α-) Divergence between densities p_1 and p_0

$$D_\alpha(p_1||p_0) = \frac{1}{\alpha - 1} ln \int p_1^\alpha(x) p_0^{1-\alpha}(x) dx$$

- Divergence between JMPD before a measurement has been made and the JMPD afterward

$$D_\alpha \left(p(\cdot|\mathbf{Z}^k) || p(\cdot|\mathbf{Z}^{k-1}) \right) = \frac{1}{\alpha - 1} ln \int_{\mathbf{X}} p(\mathbf{X}^k, T^k|\mathbf{Z}^k)^\alpha p(\mathbf{X}^k, T^k|\mathbf{Z}^{k-1})^{1-\alpha} d\mathbf{X}^k$$

- The information value of being in position **r** is

$$\phi(\mathbf{r}) \equiv \langle D_\alpha \rangle_{\mathbf{r}} = \int d\mathbf{z}^k p(\mathbf{z}|\mathbf{Z}^{k-1}, \mathbf{r}) D_\alpha \left(p(\cdot|\mathbf{Z}^k, \mathbf{r}) || p(\cdot|\mathbf{Z}^{k-1}) \right)$$

- Resulting force

$$F_I(\mathbf{r}) = -\beta \nabla_{\mathbf{r}} \phi(\mathbf{r})$$

Artificial Physics

- Artificial physics is an approximate method used to enforce certain desirable sensor properties
  - aka "physicomimetics", "virtual force methods", "potential field methods"
- This work employs Lennard-Jones (m=6, n=12) force
  - represents a specific choice - - other choices viable
  - Sensors that are close together feel a repulsive force
  - Sensors that are far apart feel an attractive force

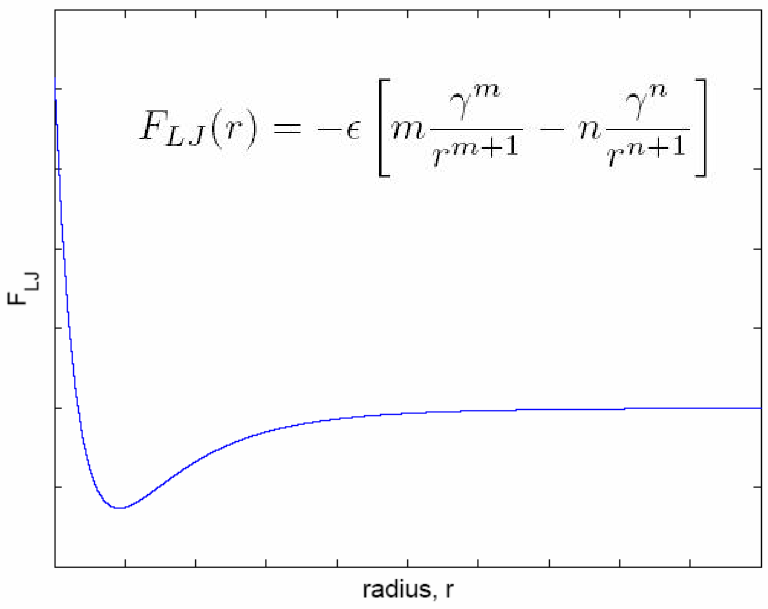

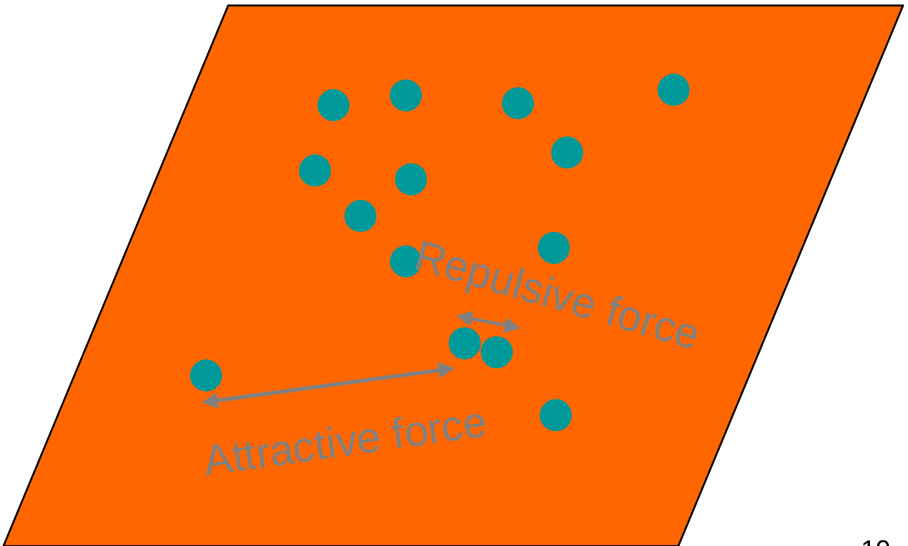

Sensor Dynamics

- Combine information theoretic force and AP derived force on each sensor

$$F_I(\mathbf{r}) = -\beta \nabla_{\mathbf{r}} \phi(\mathbf{r})$$

$$F_{LJ}(r) = -\epsilon \left[m \frac{\gamma^m}{r^{m+1}} - n \frac{\gamma^n}{r^{n+1}} \right]$$

- Combined force on sensor *i* at time *t* is $\mathbf{f}_i(t)$
- Acceleration of a unit mass object obeys the Langevin equation

$$\ddot{\mathbf{r}}_i(t) = -\frac{1}{\tau}\dot{\mathbf{r}}_i(t) + \mathbf{f}_i(t) + d\beta_i(t)$$

- Langevin equation solved numerically
  - Use Brownian dynamics Verlet algorithm is “Industry Standard”

Simulation

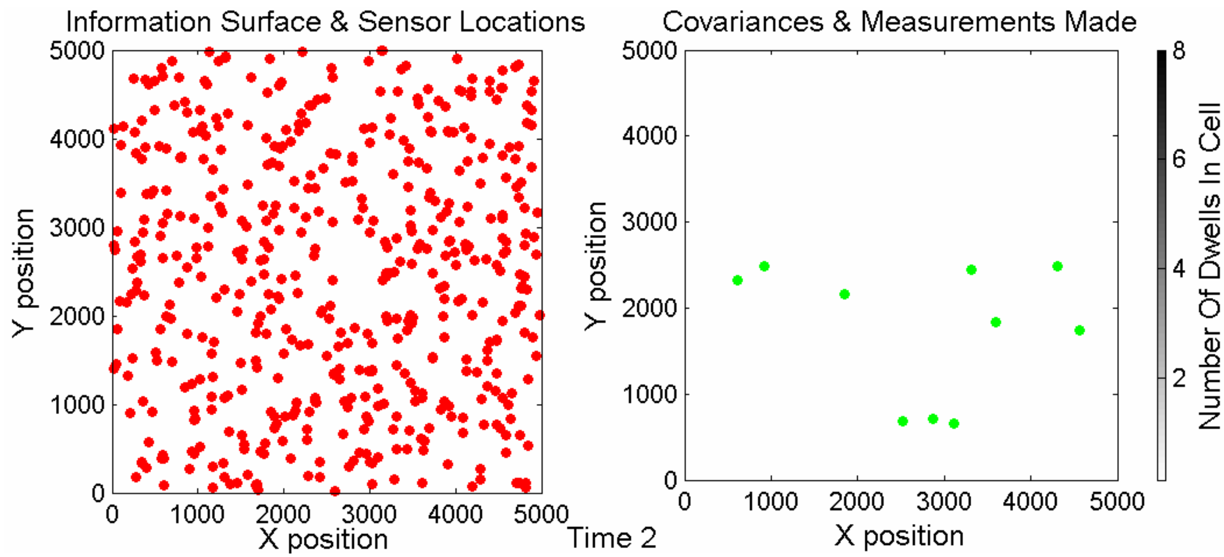

At initialization, the information surface is uniform (lots of uncertainty) and so sensor behavior is dictated by the Lennard-Jones forces : The sensors spread out uniformly through the region

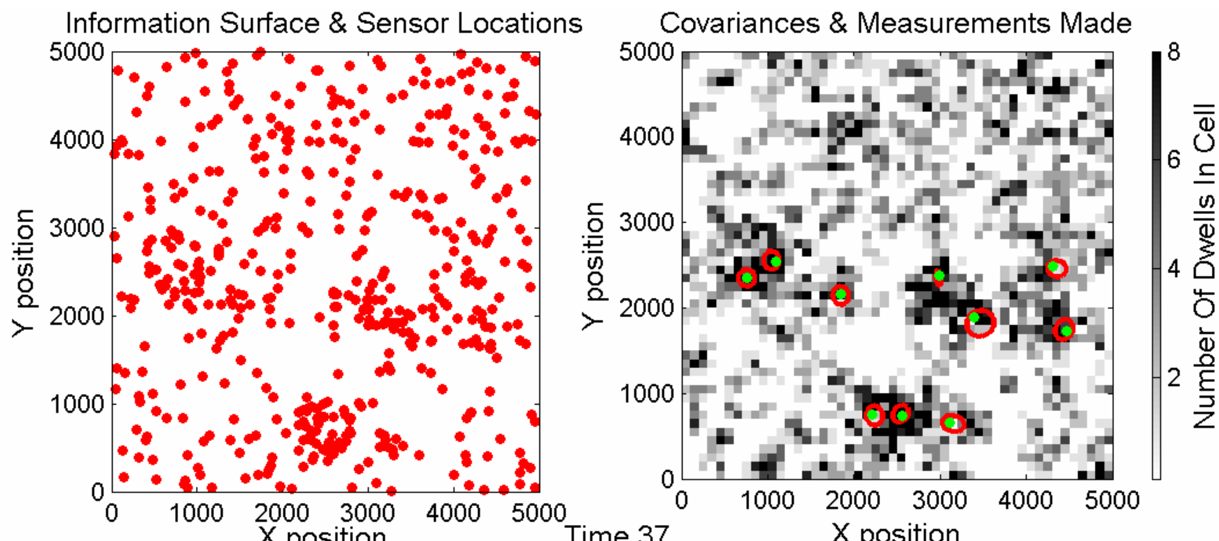

After some time, targets are detected and sensors tend to clump over target locations; however, the Lennard-Jones force ensures sensors still cover the region to address the possibility of new target arrival

Monte Carlo Results

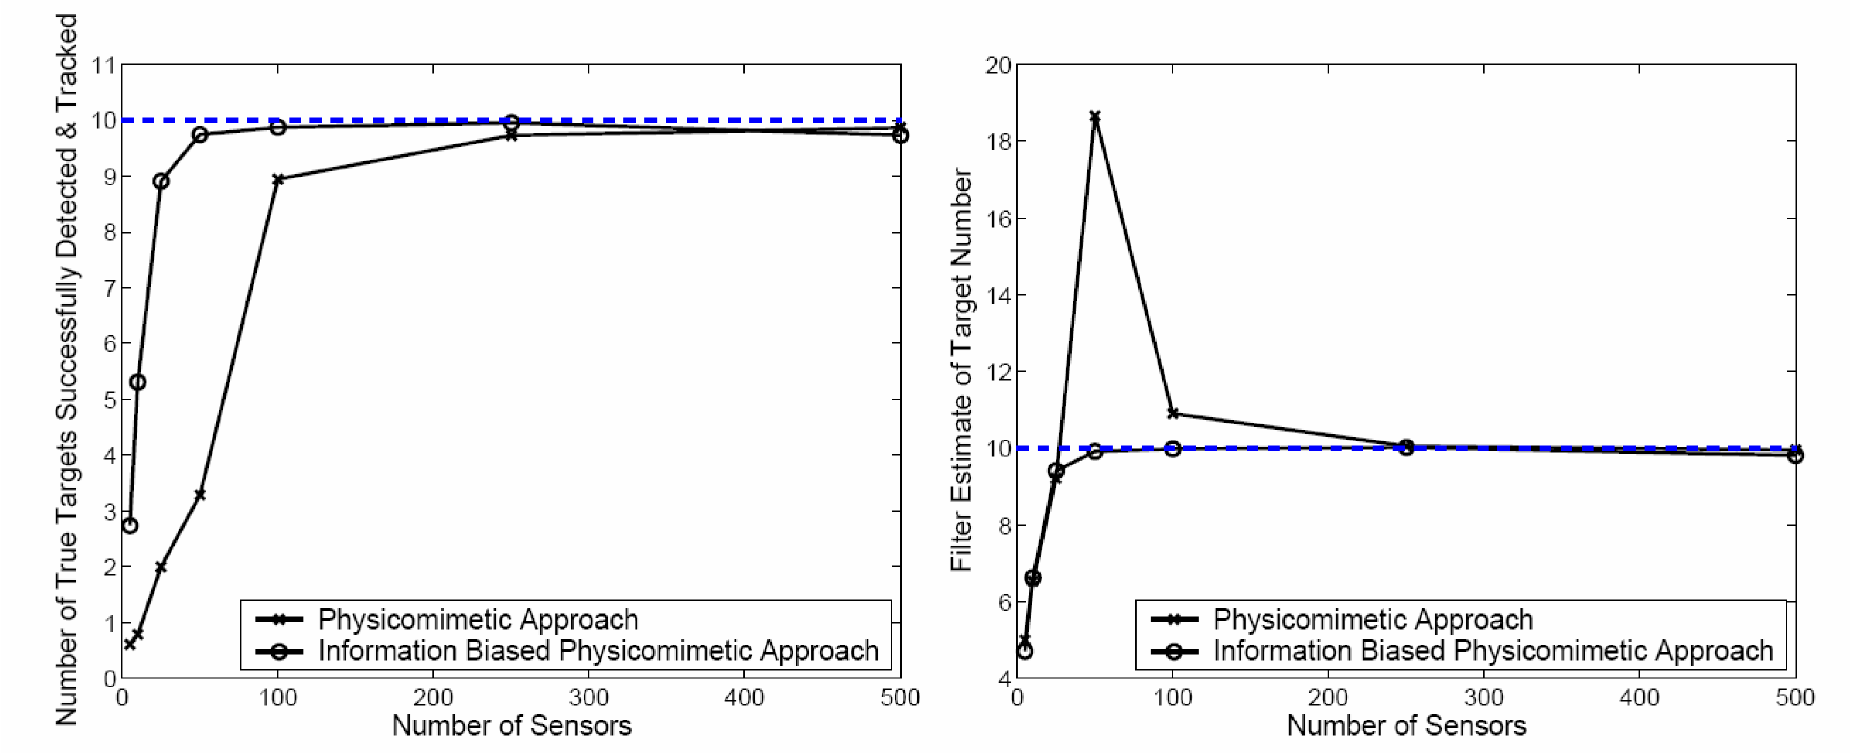

- This plot shows the performance of the Info-based AP method (compared to a purely AP method) at detecting and tracking 10 targets
- Two ways of comparing : The number of true targets successfully detected and the filter estimate of target number
- Coupling to information surface results in factor of 5 to 10 improvement in number of sensors required to meet a performance criteria

Information Coupling:
Not too much, not too little, JUST right

Monte Carlo performance study as function of coupling strength shows clear optimal region

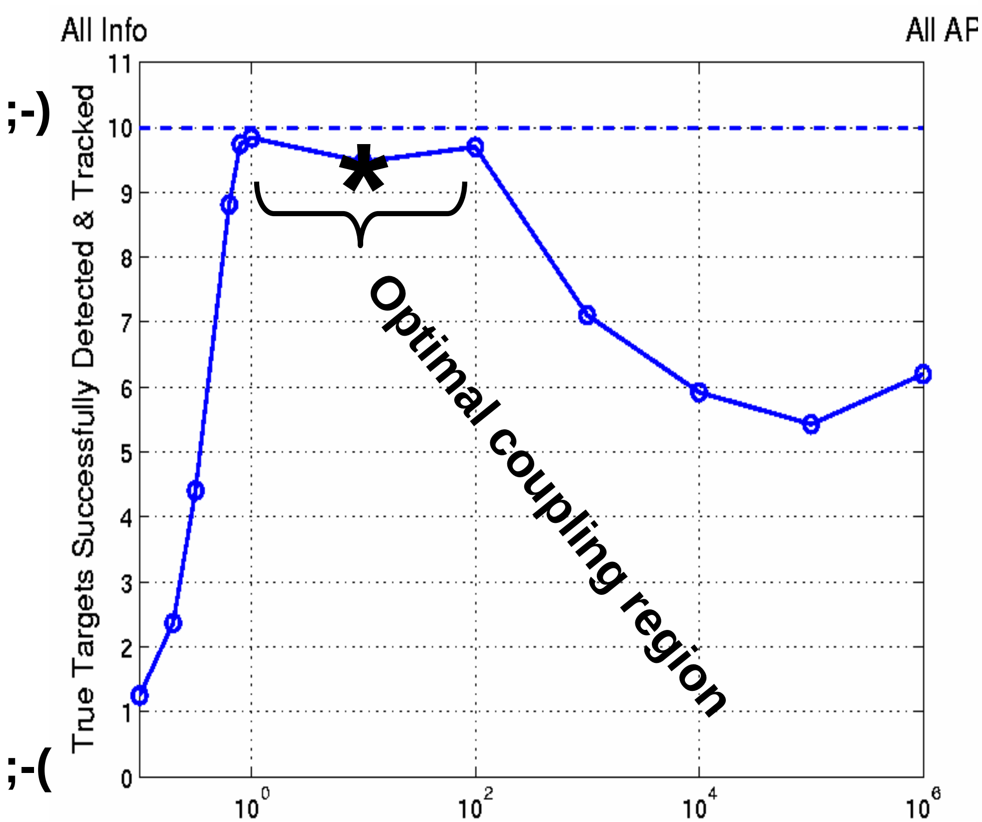

Coupling Strength 1/β

Typical sensor configuration in optimal region *

(400 sensors, single pixel detection)

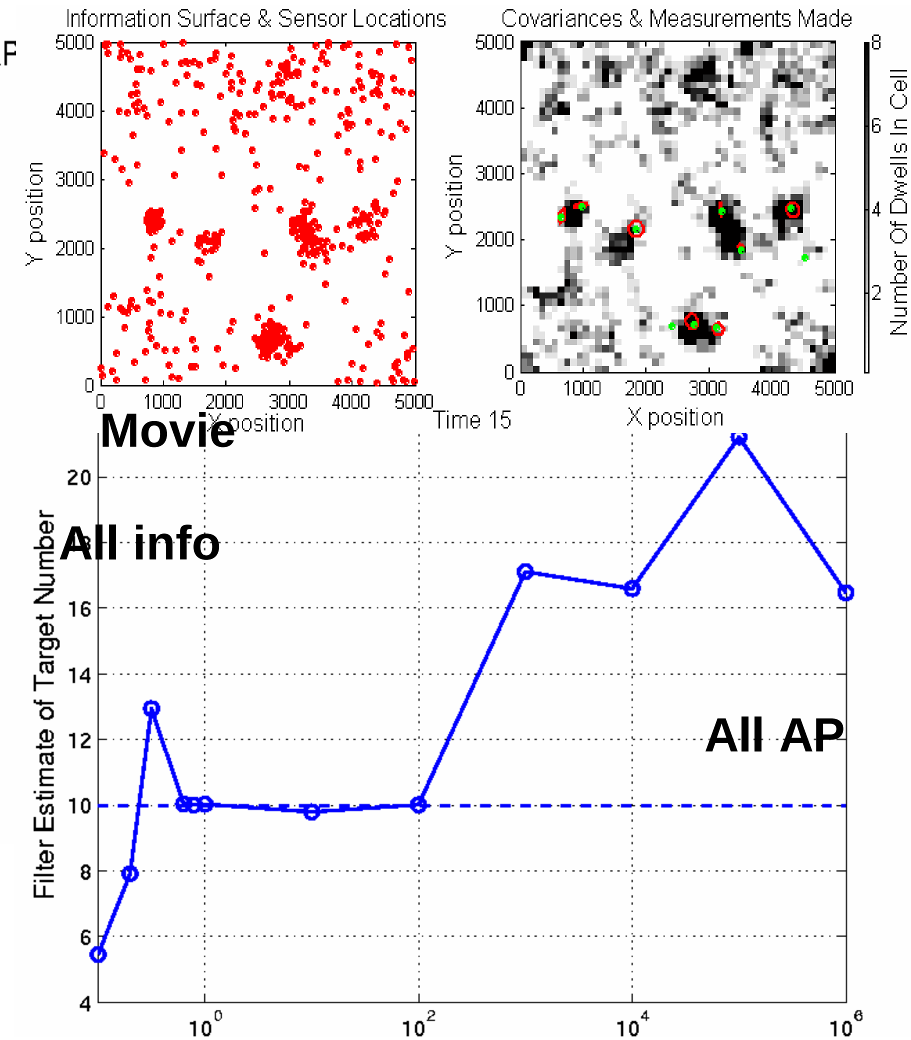

On the Choice of Mixing Parameter β

β=.01

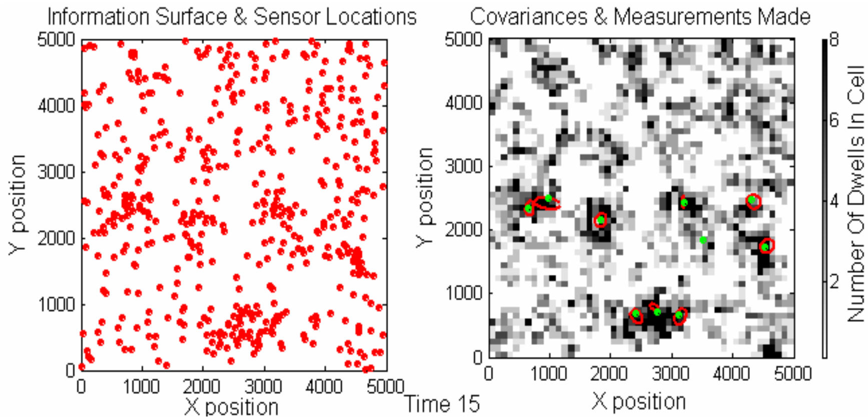

β=.02

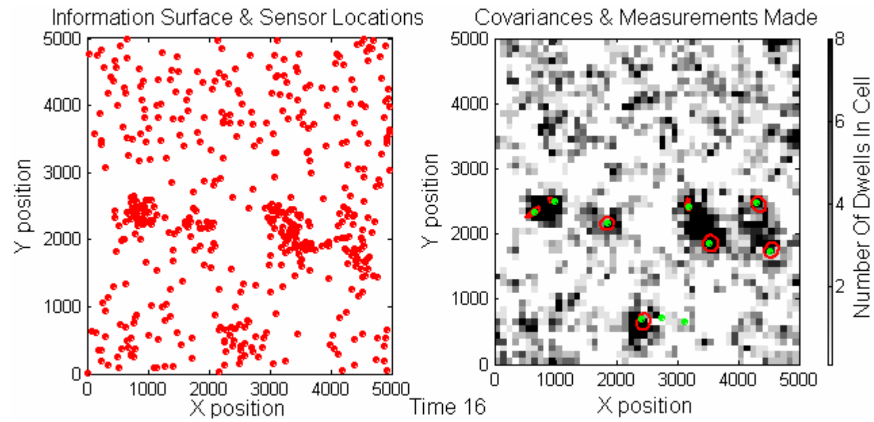

β=.04

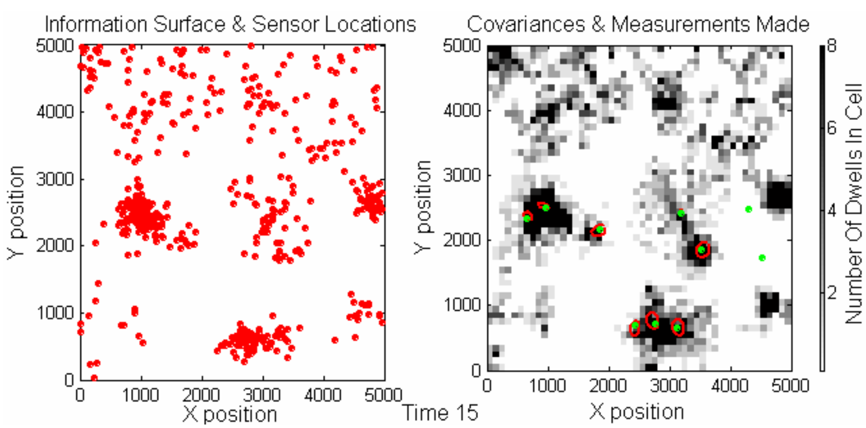

β=.08

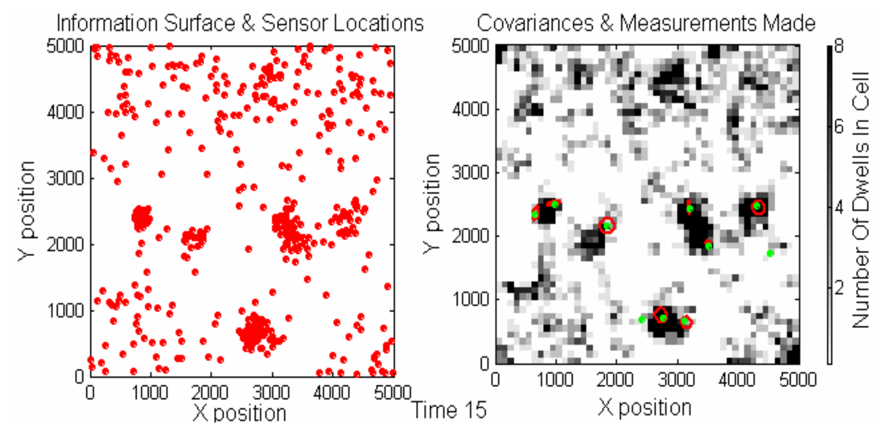

Implementation in a Distributed Decentralized Environment

- Develop architecture with each sensor responsible for making its own decisions
  - Each sensor estimates surveillance region state
    - Has high-fidelity local representation
  - Each sensor uses estimate to drive its actions
    - Base on information theory and artificial physics
  - Each sensor decides how, when and with whom to communicate
  - When communicating, a sensor will send selected *measurements* to selected other sensors
    - Reduces loopy communication problem

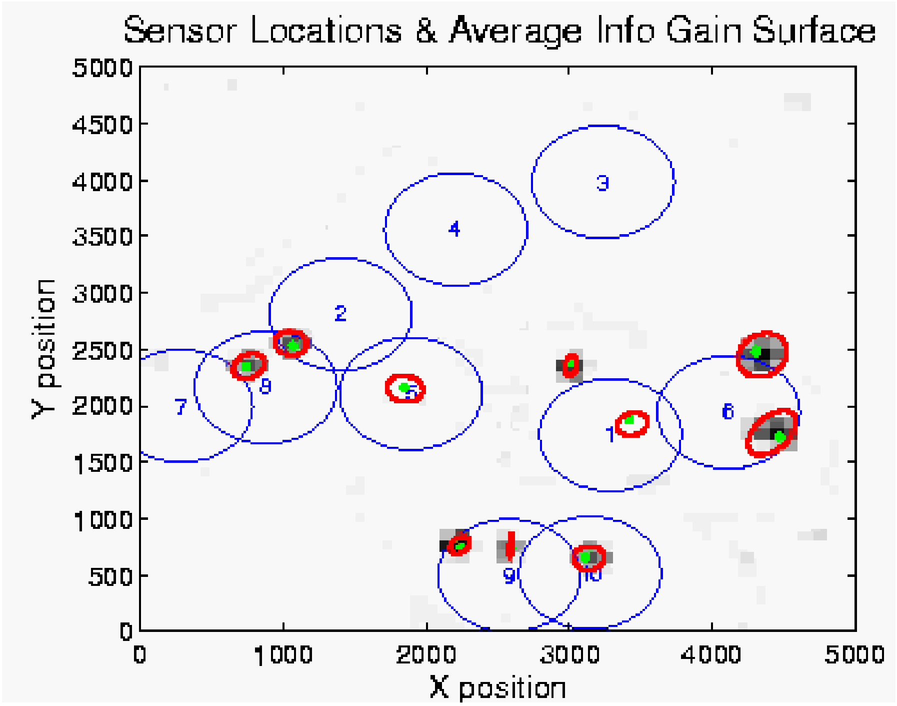

The Sensor

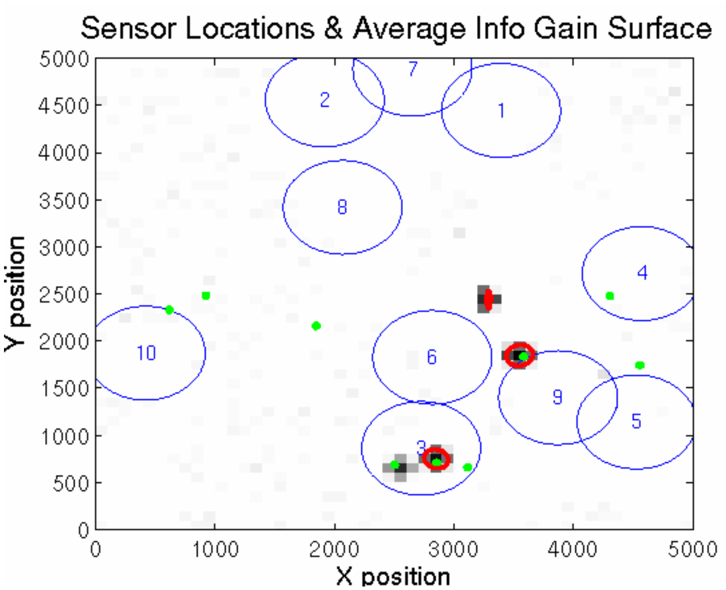

SNR

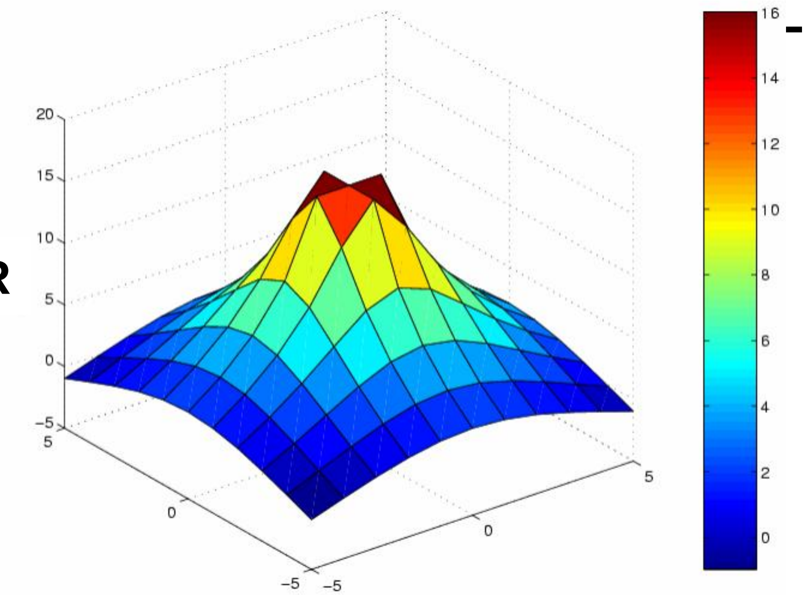

P_d

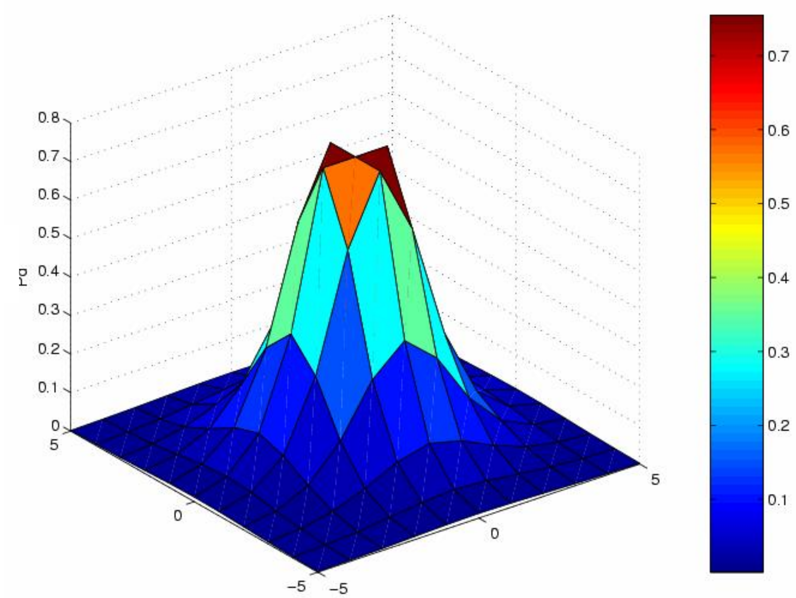

Model Problem : A Day in the Life of a sensor

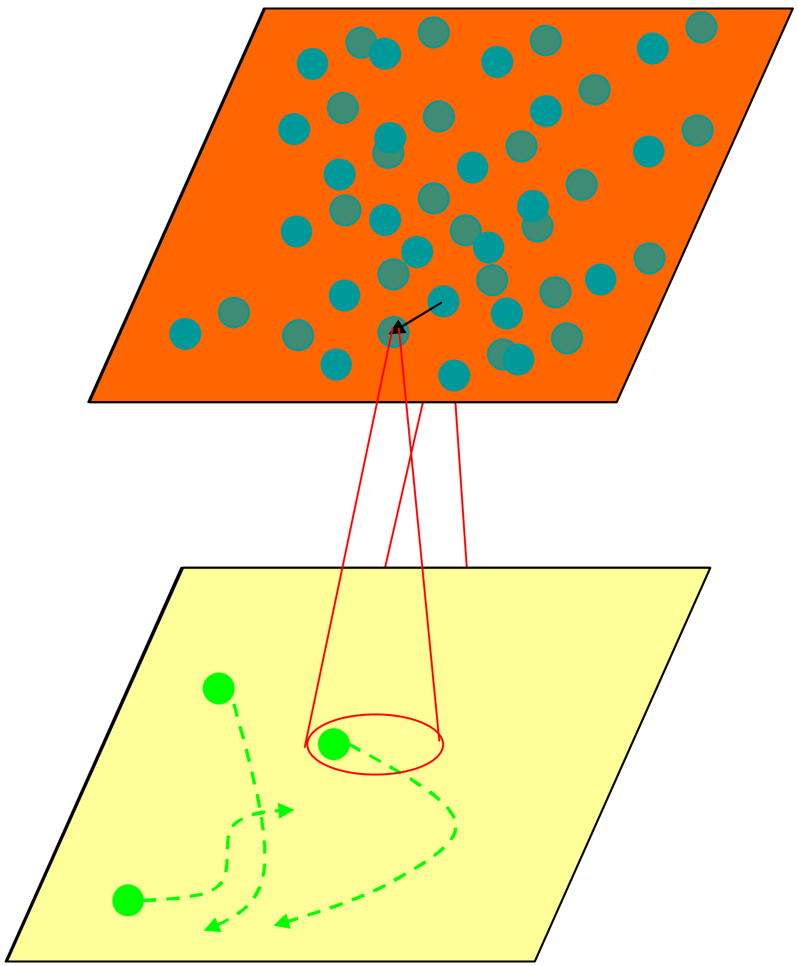

[illegible]

If doing self-localization, the sensor uses timing information from received messages to update its estimate of position

This in turn is used to create an “information force” that drives it to move towards the area of highest information gain

Performance Examples

- 15 sensors monitor the area for 250 time steps
- Each sensor is completely autonomous
  - SM decisions made based on received positions of other sensors plus information gain computed from sensor's own estimate of surveillance region
  - Communication done to nearest neighbors (r=500m) with double-hop and pedigree. Measurements are sent based on sender's estimate of target density.

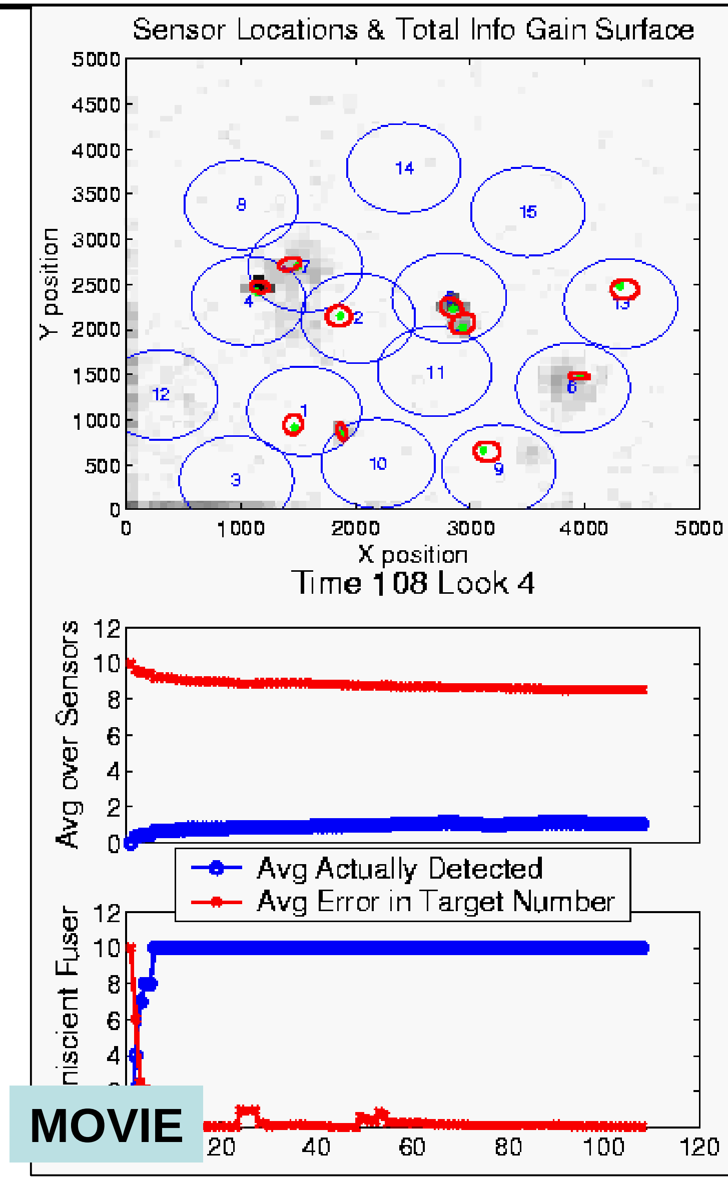

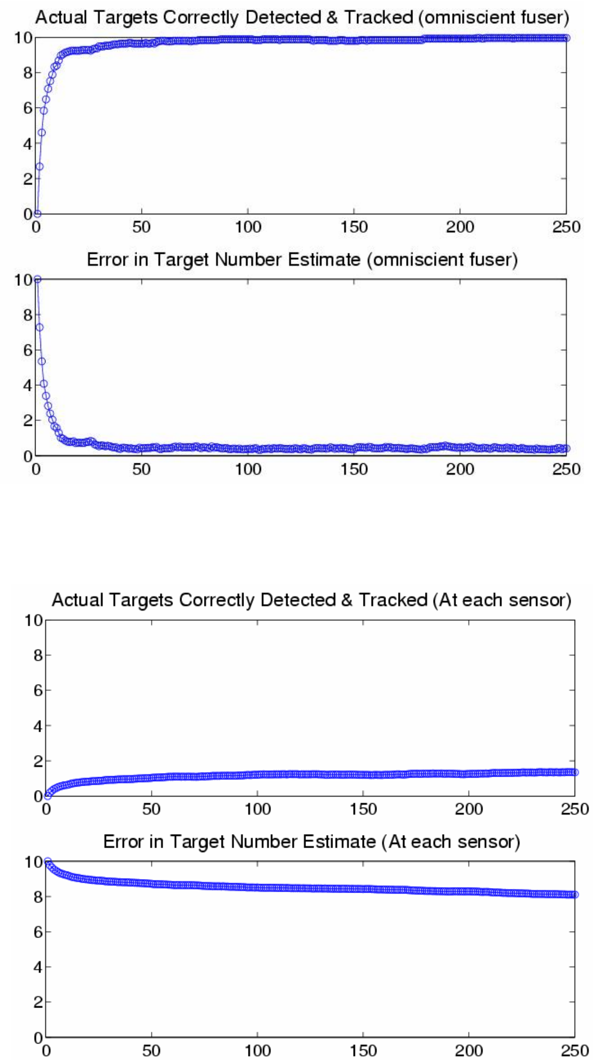

Monte Carlo Analysis of Performance

Conclusions and Summary

- This talk has given a strategy for controlling the positions of a large number of simple sensors
- Method combines information theoretic optimization and artificial physics
  - Computes value of future positions based on the impact on knowledge about the targets under surveillance
  - Enforces a regularity to target spacing
- Main benefit is tractability
  - Combinatorial problem is avoided
- Performance guarantees?

A few references

Several papers available at Al Hero's website:

- Multiplatform Information-Based Sensor Management (SPIE Aerosense 05)
- Multi-target tracking using the Joint Multitarget Probability Density (IEEE AES 05)
- Sensor Management Using and Active Sensing Approach (Signal Processing 05)