Information-driven Inference in Resource-constrained Environments

John Fisher and Jason Williams

Computer Science and Artificial Intelligence Laboratory

Massachusetts Institute of Technology

http://ssg.mit.edu/group/jlwil/publications/Thesis.pdf

Academic Model of a Sensor Network

$$p\left(X|Z\right)$$

We're not always sure of what "P" is (and it keeps changing)

We're not always sure of what "X" is (and it keeps changing)

We think we know what "Z" is (but it keeps changing)

Primary Question

How should we optimize the measurement process in a sensor network for inference problems?

- We can control sensors within resource constraints to obtain different types of information about the underlying phenomenon
- Control choices impact both quality of inference and resource expenditures.

Applications

- state estimation/tracking
- identification
- random field estimation

Outline

- Distributed Inference in Resource Constrained Environments
  - Theoretical Bounds
    - Optimal schemes are infeasible
    - Can we give performance guarantees of approximation schemes as compared to optimal?
  - In-Network Processing with Dynamic Fusion Centers
    - Explicitily trade-off value of information discounted by resource expenditures
    - Incorporate measurement transmission and selection
    - Incorporate coding/transmission of probabilistic models

In-Network Processing: Dynamic Fusion Centers

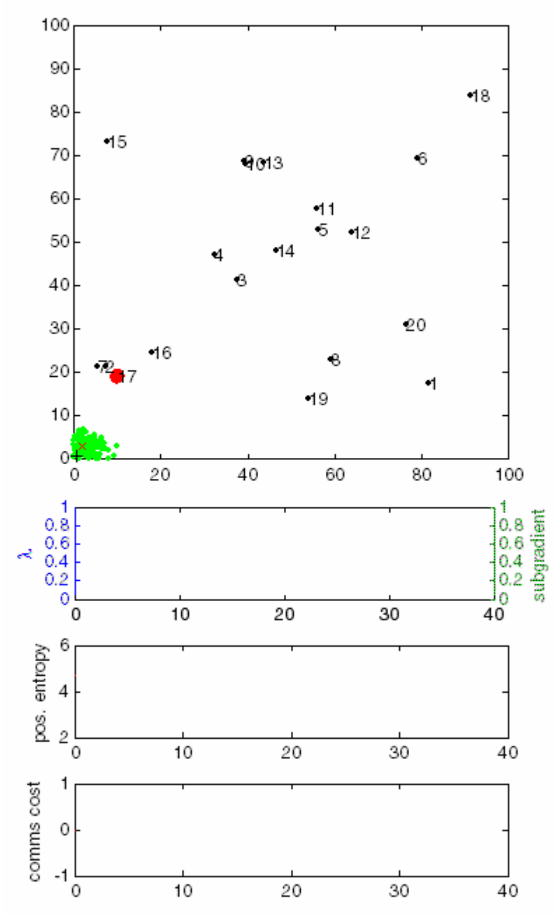

Basic Intuition

- Measurements are not equally useful and incur different resource expenditures.
- Moving the fusion center dynamically is a compromise between centralized and decentralized approaches.
- Information regarding many phenomenon is “local”.

Zhao, Shin, Reich (2002)

- Consider a tracking application in which sensors yield noisy range measurements.
- Utilize the single sensor measurement which minimizes the expected uncertainty at the next time step.
- Perform fusion at the sensor with the highest expected uncertainty reduction.
- Implicitly captures the notion that communications and fusion of all measurements is prohibitive relative to the decrease in uncertainty of the kinematic state.

Maximizing Expected Information Gain

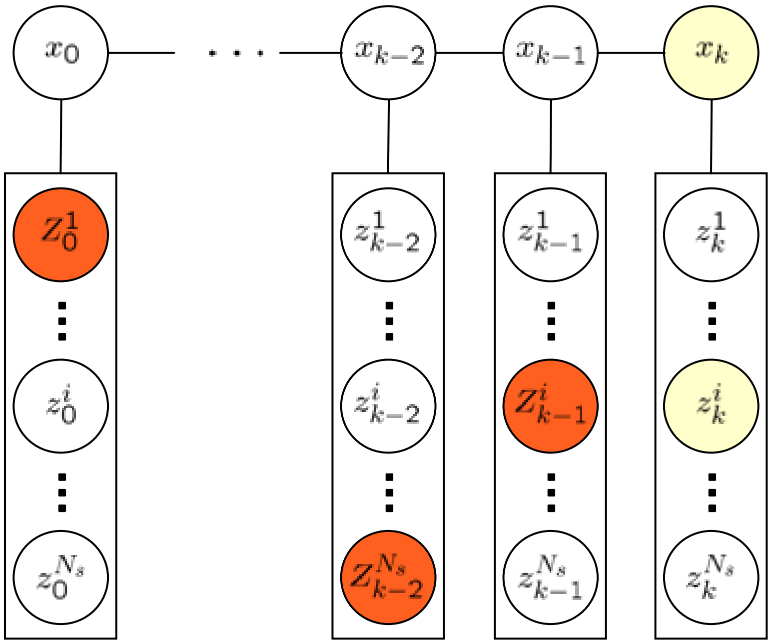

Notation

z_k^i = measurement of sensor i at time k

Z_k^i = measurement **value** of sensor i at time k

$\{z\}_{i,k}$ = selected measurements from time i to k

$\{Z\}_{i,k}$ = selected measurement **values** from time i to k

- Having incorporated previous measurements (or a subset of those available) to compute a posterior

$$p\left(x_k | \{z\}_{0,k-1} = \{Z\}_{0,k-1}\right)$$

choose the sensor whose measurement yields the highest expected information gain.

- Equivalent information-theoretic criterion:

$$\arg\min_j h\left(x_k | z_k^j, \{z\}_{0,k-1} = \{Z\}_{0,k-1}\right)$$

$$\arg\max_j I\left(x_k; z_k^j | \{z\}_{0,k-1} = \{Z\}_{0,k-1}\right)$$

$$\arg\max_j D\left(p\left(x_k | z_k^j, \{z\}_{0,k-1} = \{Z\}_{0,k-1}\right) || p\left(x_k | \{z\}_{0,k-1} = \{Z\}_{0,k-1}\right)\right)$$

$$\arg\min_j D\left(p\left(x_k | z_k^0, \cdots, z_k^{N_s}, \{z\}_{0,k-1} = \{Z\}_{0,k-1}\right) || \right.$$
$$\left. p\left(x_k | z_k^j, \{z\}_{0,k-1} = \{Z\}_{0,k-1}\right)\right)$$

- Such a strategy is **myopic** (searching over a small number of time steps) and ***greedy*** (searching over the single best available measurement at each time step)

Why Use Information-Theoretic Objective Criteria?

Scheffe's Theorem

$$\begin{aligned}\int |p-q| &= 2\sup_A \left|\int_A p - \int_A q\right| \\ &= 2\int_{p>q}(p-q) \\ &= 2\int_{q>p}(q-p)\end{aligned}$$

$$\left|\int_A p - \int_A q\right| \le \frac{1}{2}\int |p-q|$$

Kullback-Csiszar-Kemperman Inequality

$$\int |p-q| \le \min\left\{\sqrt{2D(p||q)}, \sqrt{2D(q||p)}\right\}$$

Bretagnolle-Huber Inequalities

$$\begin{aligned}\int |p-q| &\le 2\sqrt{1-e^{-D(p||q)}} \\ \int |p-q| &\le 2-e^{-D(p||q)} \\ \int \min\{p,q\} &\ge \frac{1}{2}e^{-D(p||q)}\end{aligned}$$

- Closeness in an L_1 sense bounds errors in estimates of event probabilities.
- L_1 is often difficult to optimize, K-L is not in many cases.
- Closeness in K-L bounds closeness in L_1.

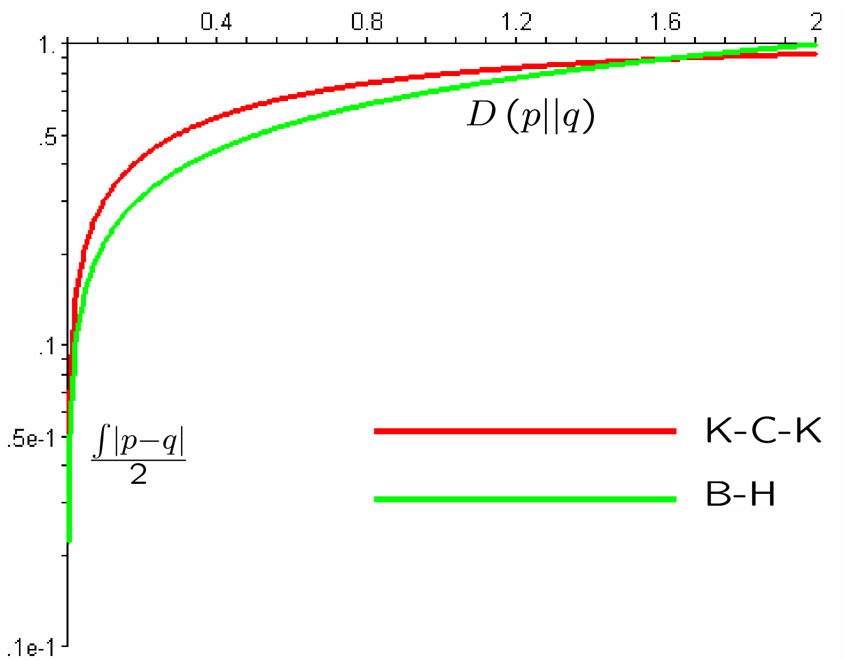

Value of Long Term Planning?

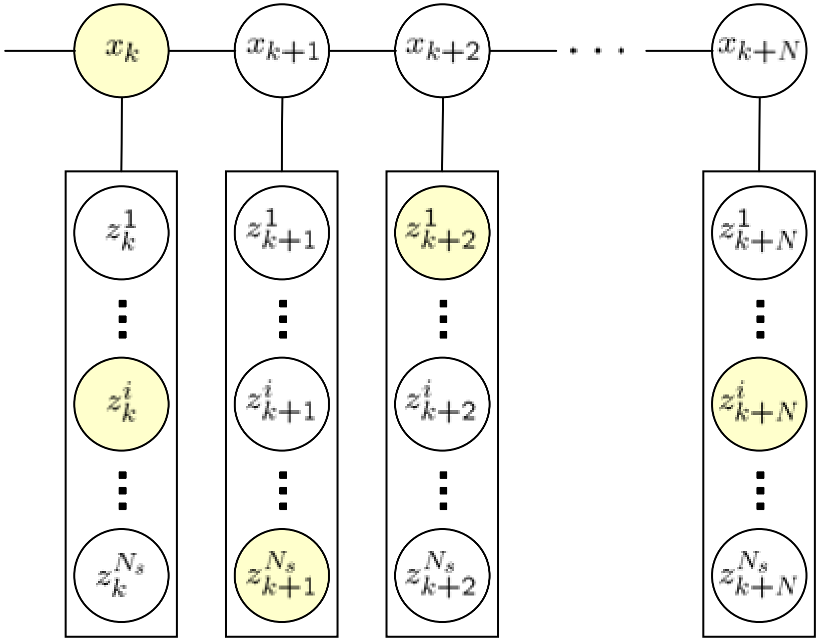

Choosing the optimal set of measurements (sensors) is exponential in the planning horizon.

- Can we bound the difference in performance of approximate (tractable) algorithms as compared to optimal?
- Is there a performance gain if we incorporating planning over a longer time-horizon (non-myopic)

Capturing the Notion of Diminishing Returns

- Assuming that observations are independent conditioned on the state, mutual information is *submodular*

  - If $\mathcal{A} \subseteq \mathcal{B}$ then: $$\begin{aligned} I(X; z^{\mathcal{C}} | z^{\mathcal{B}}) &= H(z^{\mathcal{C}} | z^{\mathcal{B}}) - H(z^{\mathcal{C}} | z^{\mathcal{B}}, X) \\ &= H(z^{\mathcal{C}} | z^{\mathcal{B}}) - H(z^{\mathcal{C}} | X) \\ &\leq H(z^{\mathcal{C}} | z^{\mathcal{A}}) - H(z^{\mathcal{C}} | X) \\ &= I(X; z^{\mathcal{C}} | z^{\mathcal{A}}) \end{aligned}$$

$$\mathcal{B} = \left\{ \underbrace{x_1, x_2, x_3, x_4, x_5, x_6}_{\mathcal{A}}, \ldots, x_N \right\} \qquad \mathcal{C} = \{x_{N+1}, \ldots\}$$

A simple bound…

- Suppose we have *M* stages
  - Each stage involves selection of an observation for a different sensor or a different time step
- Suppose we use the greedy heuristic to select each observation:

$$g_j = \underset{g \in \{1,\ldots,n_j\}}{\arg\max}\ I(X; z_j^g \mid z_1^{g_1}, \ldots, z_{j-1}^{g_{j-1}})$$

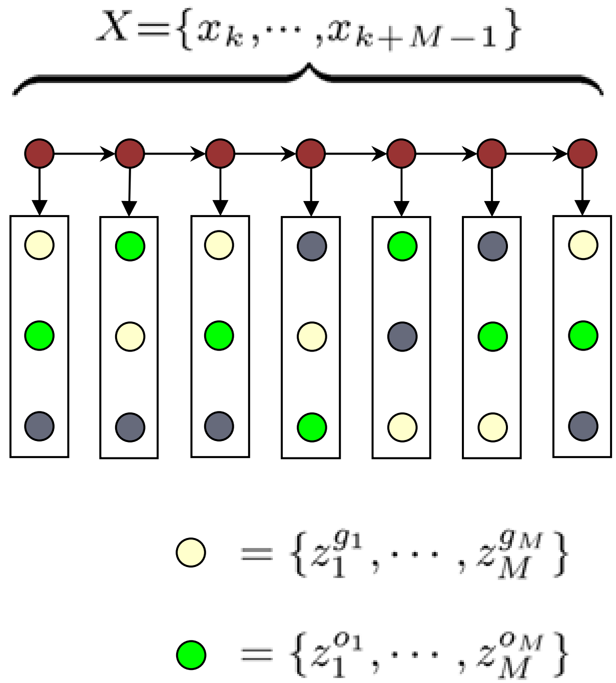

Then…

$$I(X; z_1^{o_1}, \ldots, z_M^{o_M}) \leq 2I(X; z_1^{g_1}, \ldots, z_M^{g_M})$$

A simple bound…

$$
\begin{aligned}
I(X; & z_1^{o_1}, \ldots, z_M^{o_M}) \\
&\leq I(X; z_1^{g_1}, \ldots, z_M^{g_M}, z_1^{o_1}, \ldots, z_M^{o_M}) \\
&= \sum_{j=1}^{M} \left\{ I(X; z_j^{g_j} | z_1^{g_1}, \ldots, z_{j-1}^{g_{j-1}}) + I(X; z_j^{o_j} | z_1^{g_1}, \ldots, z_M^{g_M}, z_1^{o_1}, \ldots, z_{j-1}^{o_{j-1}}) \right\} \\
&\leq \sum_{j=1}^{M} \left\{ I(X; z_j^{g_j} | z_1^{g_1}, \ldots, z_{j-1}^{g_{j-1}}) + I(X; z_j^{o_j} | z_1^{g_1}, \ldots, z_j^{g_{j-1}}) \right\} \\
&\leq 2 \sum_{j=1}^{M} I(X; z_j^{g_j} | z_1^{g_1}, \ldots, z_{j-1}^{g_{j-1}}) \\
&= 2 I(X; z_1^{g_1}, \ldots, z_M^{g_M})
\end{aligned}
$$

Online computable bound

- While the bound is tight, the proof gives rise to an online computable version which may be stronger in particular circumstances
- Let $\bar{g}_i = \underset{\bar{g}\in\{1,\ldots,n_i\}}{\arg\max}\; I(X; z_i^{\bar{g}} | z_1^{g_1}, \ldots, z_M^{g_M})$
- Then:

$$I(X; z_1^{o_1}, \ldots, z_M^{o_M}) \leq I(X; z_1^{g_1}, \ldots, z_M^{g_M}) + \sum_{i=1}^{M} I(X; z_i^{\bar{g}_i} | z_1^{g_1}, \ldots, z_M^{g_M})$$

- This bound is tight in situations where the greedy selection leaves little information behind

Online computable bound – example

- Suppose we model the depth of the ocean as a Gauss-Markov random field thin membrane model (500 $\times$ 100 cells)
- We seek to measure ocean depth using a surface vehicle traveling along a fixed path
  - Available observations depend on current vehicle position

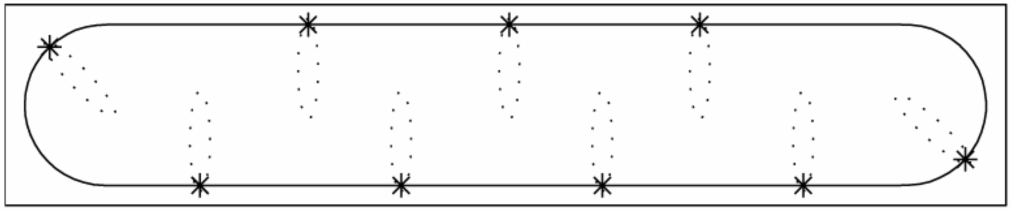

- We need to choose observations at each time

Online computable bound – example

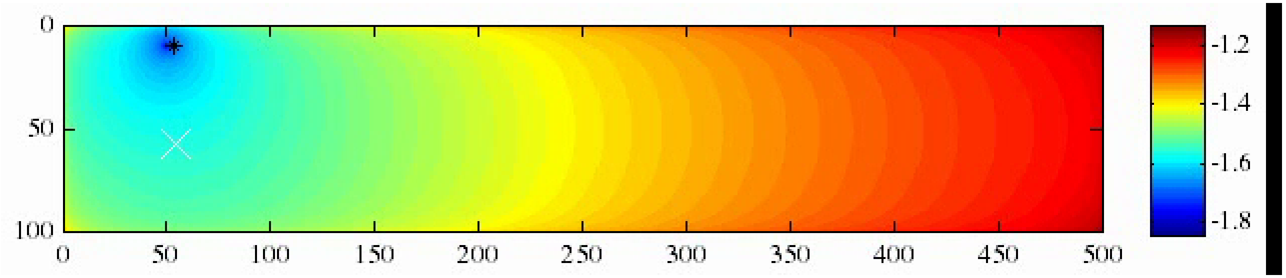

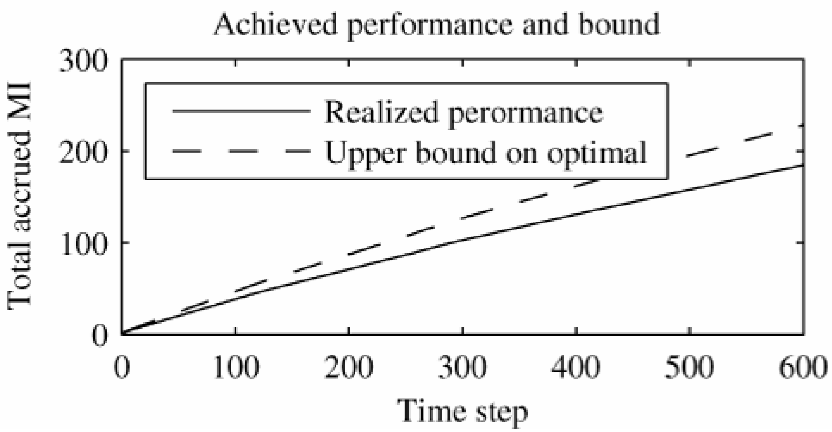

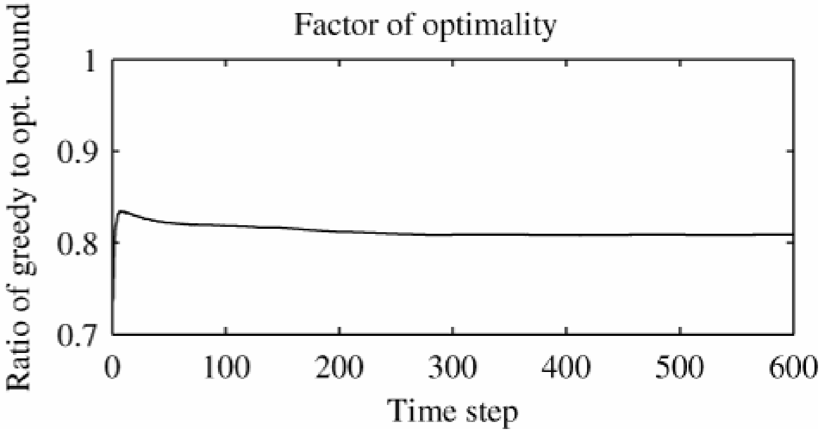

Feasible Information Driven Sensor Planning

- Motivation & assumptions
- Constrained Markov Decision Process formulation
  - Communication constraint
  - Entropy constraint
- Approximations
  - Linearized Gaussian
  - Greedy sensor subset selection
  - n-Scan pruning
- Simulation results

Incorporating Resource Constraints

- Object tracking using a network of sensors
  - Sensors provide localization in close vicinity
  - Sensors can communicate locally at cost $\propto f\,(\text{range})$
- Energy is a limited resource
  - Consumed by communication, sensing and computation
  - Communication is orders of magnitude more expensive than other tasks
- Sensor management algorithm must determine
  - Which sensors to activate at each time
  - Where the probabilistic model should be stored at each time
  - While trading off the competing objectives of estimation performance and communication cost

Greedy Myopic Planning

- Model transmission occurs with little benefit to the inference problem.
- Single measurement used at each time.
- Resources are only implicitly incorporated into the inference problem.

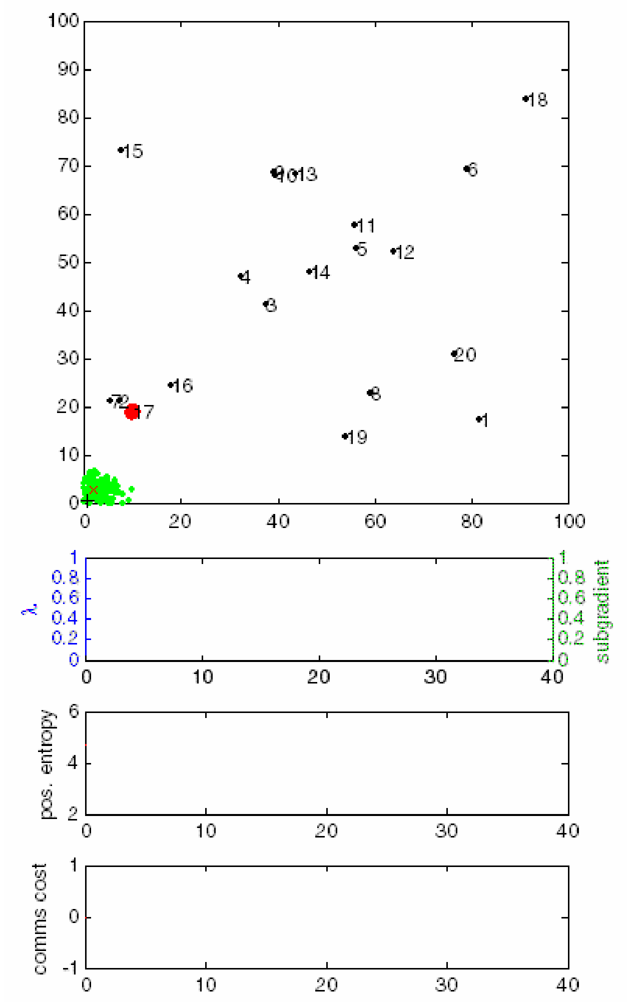

Can we benefit by considering expected information gain discounted by resource expenditures over long horizons?

Notation

Time index: k

Position of sensor s: $\boldsymbol{l}^s$

Leader node: l_k

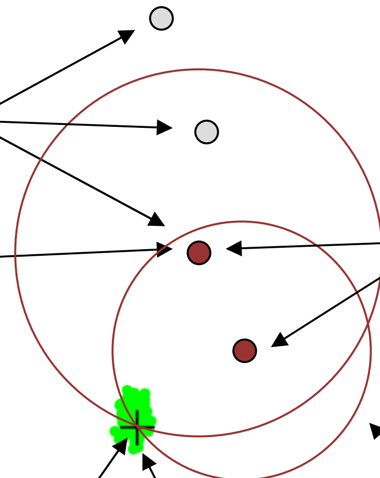

Activated sensors: $\mathcal{S}_k \subseteq \mathcal{S}$

Object state / location: $\boldsymbol{x}_k$ / $\mathbf{L}\boldsymbol{x}_k$

Probabilistic model: $\mathbb{X}_k = p(\boldsymbol{x}_k|\boldsymbol{z}_{0:k-1})$

Sensor s measurement: $\boldsymbol{z}_k^s$

History of incorporated measurements: $\boldsymbol{z}_{0:k-1}$

Example – observation/communications

Observation model:

$$z_k^s = \frac{a}{||\mathbf{L}\boldsymbol{x}_k - \boldsymbol{l}^s||_2^2 + b} + \boldsymbol{v}_k^s$$

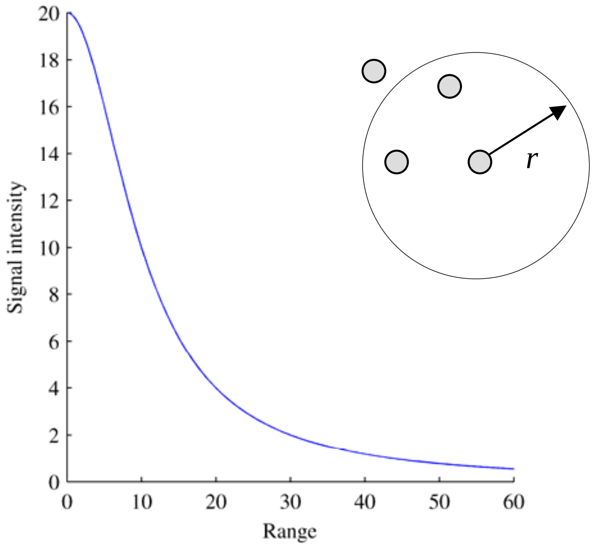

Communications cost:

$$\tilde{C}_{ij} = ||\boldsymbol{l}^i - \boldsymbol{l}^j||_2^2$$

$$C_{ij} = \sum_{k=1}^{n} \tilde{C}_{i_{k-1} i_k}$$

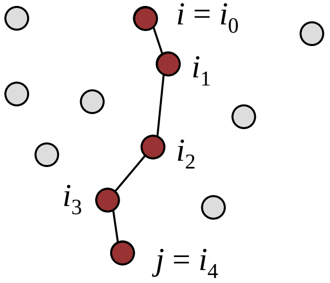

Formulation

- We formulate as a constrained Markov Decision Process
  - Minimize uncertainty subject to a constraint on communication cost
  - Minimize communication cost subject to a constraint on estimation performance
- State is PDF ($\mathbb{X}_k = p(\boldsymbol{x}_k | \boldsymbol{z}_{0:k-1})$) and previous leader node (l_{k-1})
  - Dynamic programming equation for an N-step rolling horizon:

$$J_i(\mathbb{X}_i, l_{i-1}) = \min_{u_i} \left\{ g(\mathbb{X}_i, l_{i-1}, u_i) + \mathop{\mathrm{E}}_{\mathbb{X}_{i+1}|\mathbb{X}_i, u_i} J_{i+1}(\mathbb{X}_{i+1}, l_i) \right\}, \; i \in \{k : k+N-1\}$$

$$J_{k+N}(\mathbb{X}_{k+N}, l_{k+N-1}) = 0$$

  - such that $\mathrm{E}\{G(\mathbb{X}_k, l_{k-1}, u_{k:k+N-1}) | \mathbb{X}_k, l_{k-1}\} \leq 0$

Communication-constrained formulation

- Cost-per-stage is such that the system minimizes joint expected conditional entropy of object state over planning horizon:

$$g(\mathbb{X}_k, l_{k-1}, u_k) = -\sum_{j=1}^{|\mathcal{S}_k|} I(\boldsymbol{x}_k; \boldsymbol{z}_k^{\mathcal{S}_k^j} | \boldsymbol{z}_{0:k-1}, \boldsymbol{z}_k^{\mathcal{S}_k^{1:j-1}})$$

- Constraint applies to expected communication cost over planning horizon:

$$G(\mathbb{X}_k, l_{k-1}, u_{k:k+N-1}) = \sum_{i=k}^{k+N-1} \left[C_{l_{i-1}l_i} + \sum_{j \in \mathcal{S}_i} r C_{l_i j} \right] - C_{\max}$$

Communication-constrained formulation

- Integrating the communication costs into the per-stage cost:

$$\bar{g}(\mathbb{X}_k, l_{k-1}, u_k, \lambda) = -\sum_{j=1}^{|\mathcal{S}_k|} I(\boldsymbol{x}_k; \boldsymbol{z}_k^{\mathcal{S}_k^j} | \boldsymbol{z}_{0:k-1}, \boldsymbol{z}_k^{\mathcal{S}_k^{1:j-1}}) + \lambda \left[C_{l_{k-1} l_k} + \sum_{j \in \mathcal{S}_k} r C_{l_k j} \right]$$

- Per-stage cost now contains both information gain and communication cost

Information-constrained formulation

- Cost-per-stage is such that the system minimizes the energy consumed over the planning horizon:

$$g(\mathbb{X}_k, l_{k-1}, u_k) = C_{l_{k-1}l_k} + \sum_{j \in \mathcal{S}_k} rC_{l_k j}$$

- Constraint ensures that the joint entropy over the planning horizon is less than given threshold:

$$G(\mathbb{X}_k, l_{k-1}, u_{k:k+N-1}) = -\left\{ \sum_{i=k}^{k+N-1} \sum_{j=1}^{|\mathcal{S}_i|} I(x_i; z_i^{\mathcal{S}_i^j} | z_{0:i-1}, z_i^{\mathcal{S}_i^{1:j-1}}) - I_{\min} \right\}$$

Evaluating the DP

- DP has infinite state space, hence it cannot be evaluated exactly
- Conceptually, it could be evaluated through simulation

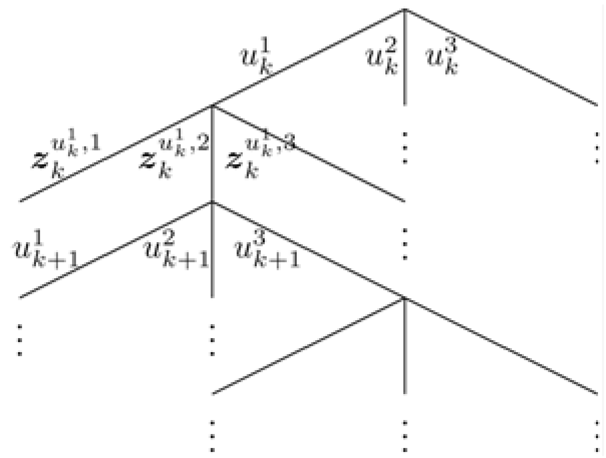

- Complexity is $O([N_s 2^{N_s}]^N {N_p}^N)$

Branching due to measurement values

- The first source of branching is that due to different values of measurement
- If we approximate the measurement model as linear Gaussian locally around a nominal trajectory, then the future costs are dependent only on the control choices, not on the measurement values
- Hence this source of branching can be eliminated entirely

Greedy sensor subset selection

- For large sensor networks complexity is high even for a single look-ahead step due to consideration of sensor subsets
- We decompose each decision stage into a generalized stopping problem, where at each substage we can
  - Add unselected sensor to the current selection
  - Terminate with the current selection
- The per-stage cost can be conveniently decomposed into a per-substage cost which directly trades off the cost of obtaining each measurement against the information it returns:

$$\bar{g}(\mathbb{X}_i, l_i, \mathcal{S}_{i,i'}, s_{i,i'}, \lambda) = \lambda r C_{l_i s_{i,i'}} - I(\boldsymbol{x}_i; \boldsymbol{z}_i^{s_{i,i'}} | \boldsymbol{z}_{0:i-1}, \boldsymbol{z}_i^{\mathcal{S}_{i,i'}})$$

Greedy sensor subset selection

- Outer DP recursion

$$\bar{J}_i(\mathbb{X}_i, l_{i-1}, \lambda) = \min_{l_i} \left\{ \lambda C_{l_{i-1} l_i} + \bar{J}_{i,0}(\mathbb{X}_i, l_i, \{\emptyset\}, \lambda) \right\}$$

- Inner DP sub-problem

$$\bar{J}_{i,i'}(\mathbb{X}_i, l_i, \mathcal{S}_{i,i'}, \lambda) = \min \left\{ \underset{\mathbb{X}_{i+1}|\mathbb{X}_i, \mathcal{S}_{i,i'}}{\mathrm{E}} \bar{J}_{i+1}(\mathbb{X}_{i+1}, l_i, \lambda), \right.$$
$$\left. \min_{s_{i,i'} \in \mathcal{S} \setminus \mathcal{S}_{i,i'}} \left\{ \bar{g}(\mathbb{X}_i, l_i, \mathcal{S}_{i,i'}, s_{i,i'}, \lambda) + \bar{J}_{i,i'+1}(\mathbb{X}_i, l_i, \mathcal{S}_{i,i'} \cup s_{i,i'}, \lambda) \right\} \right\}$$

n-Scan approximation

- The greedy subset selection is embedded within a n-scan pruning method (similar to the MHT) which addresses the growth due to different choices of leader node sequence

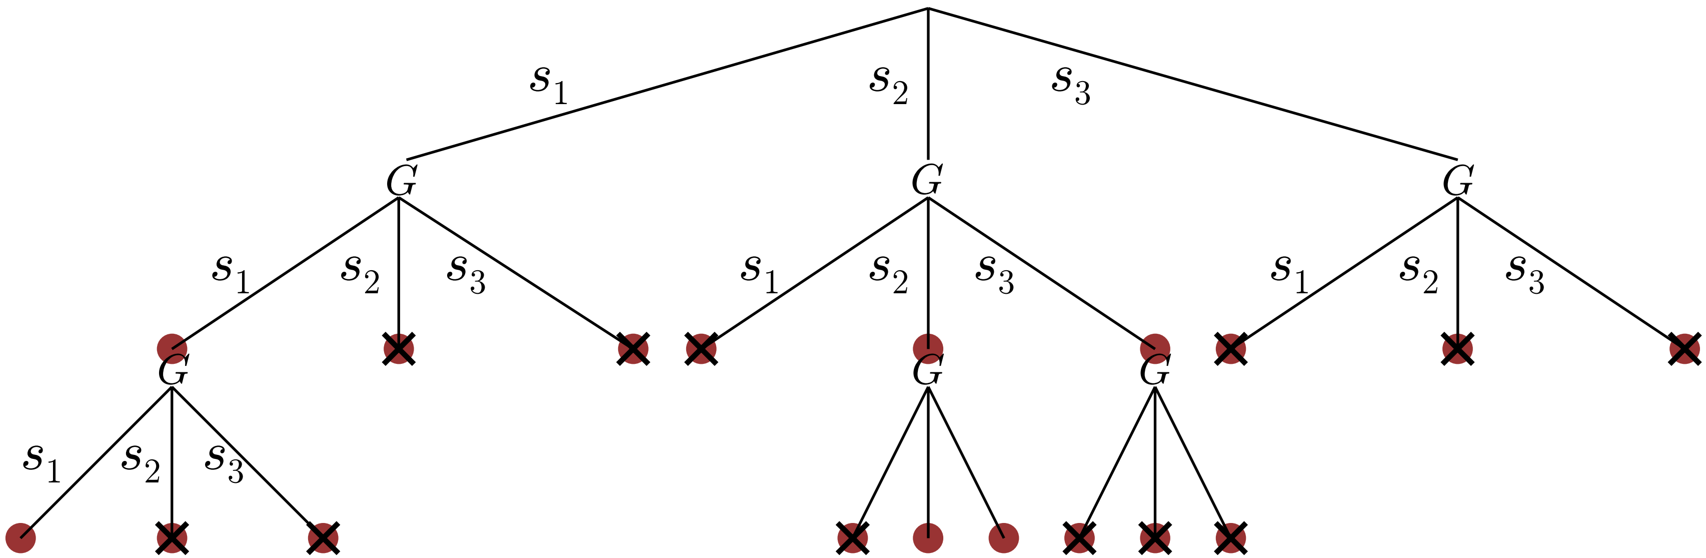

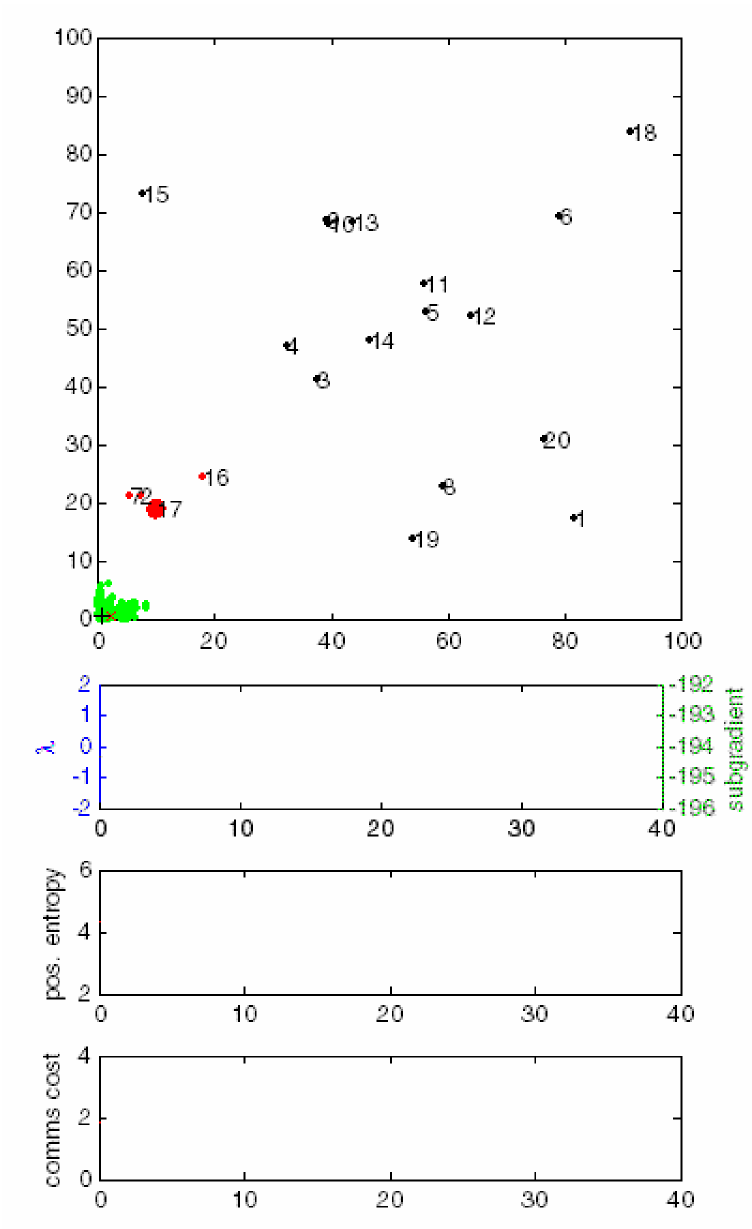
18
15
6
9
10
13
11
5
12
4
14
3
20
16
7
2
17
8
1
19
λ
subgradient
pos. entropy
comms cost

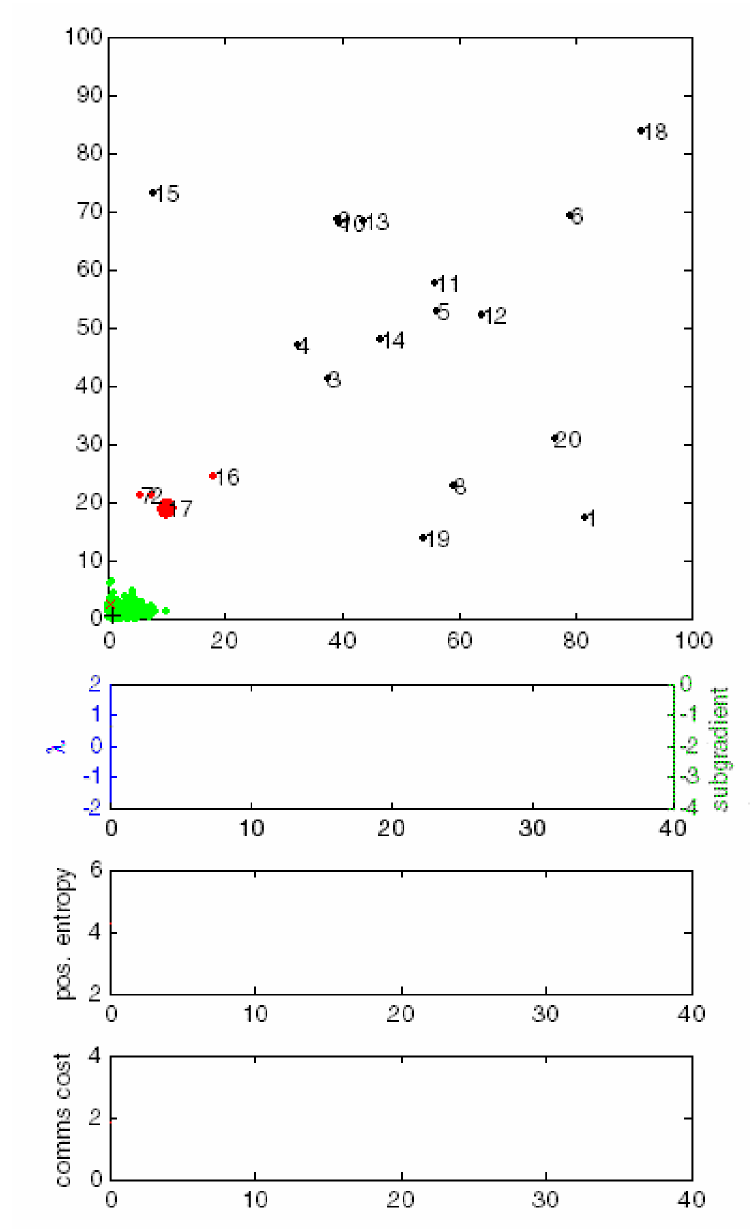
18
15
6
9
10
13
11
5
12
4
14
3
20
16
7
2
17
8
1
19
λ
subgradient
pos. entropy
comms cost

Simulation results

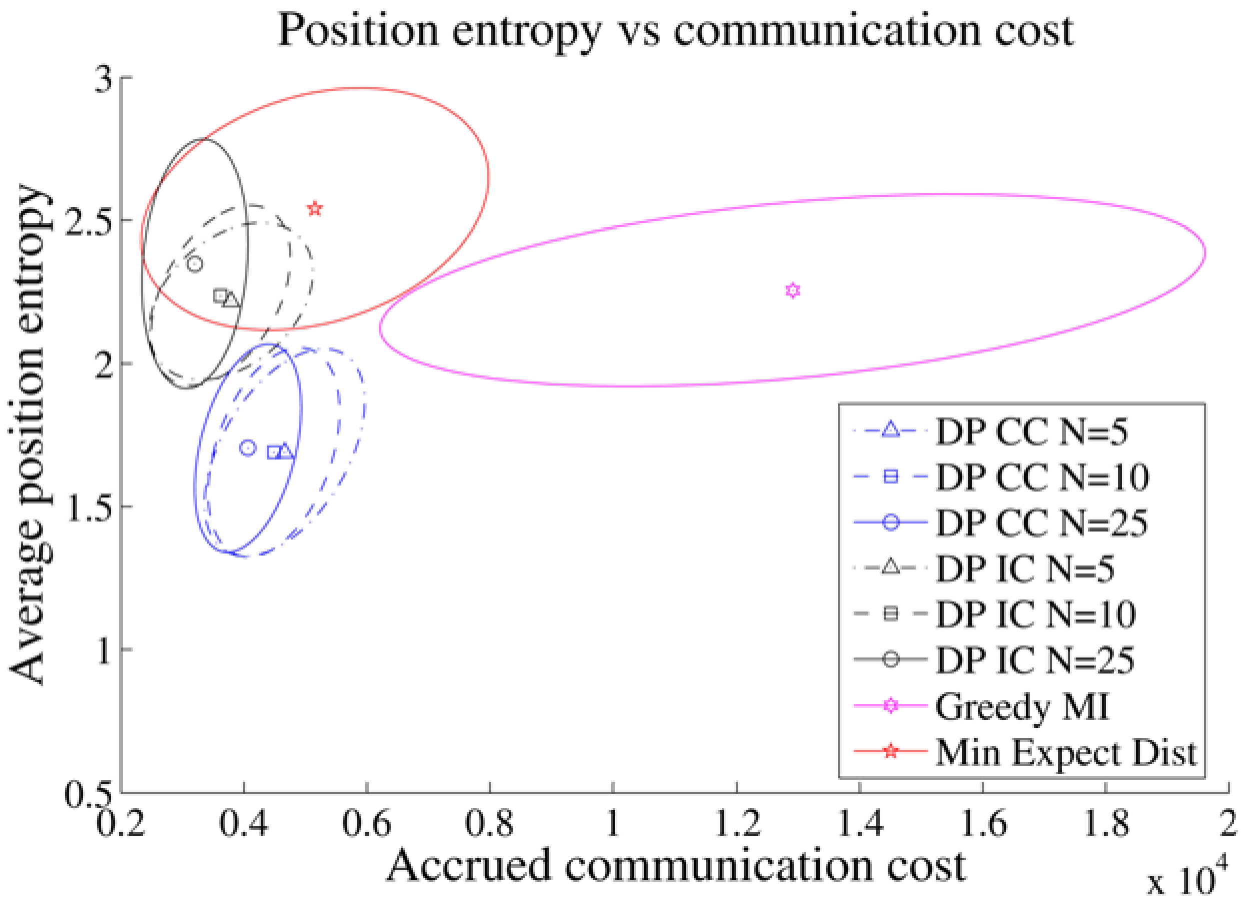

Significant Cost Reduction

- Decomposed cost structure into a form in which greedy approximations are able to capture the trade-off
  - Complexity $O([N_s 2^{N_s}]^N N_p{}^N)$ reduced to $O(N N_s{}^3)$
  - $N_s = 20$ / $N_p = 50$ / $N = 10$: 1.6×10^{90} → 8×10^{5}
  - Strong non-myopic planning for horizon lengths > 20 steps possible