A Distributed Deterministic Approximation Algorithm for Data Association

IEEE DCOSS 2011

Songhwai Oh

Department of Electrical Engineering
Seoul National University, Korea

DEPARTMENT OF ELECTRICAL ENGINEERING
SEOUL NATIONAL UNIVERSITY

CPS LAB

What is data association?

Problem domain

- Multiple events, multiple observations
- Source of each measurement is uncertain
  - Observations are correlated
  - A direct estimation of states is not possible

A data association problem is to figure out which observations were generated by which events

A fundamental problem in wireless sensor networks

Examples from Other Areas

Computer Vision

Correspondence Problem

- Structure from motion
- 3D Reconstruction
- Image registration
- Calibration

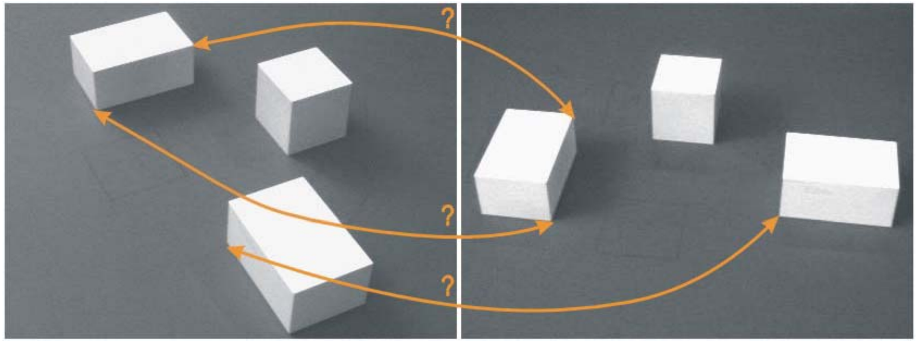

Robotics

Simultaneous Localization and Mapping (SLAM)

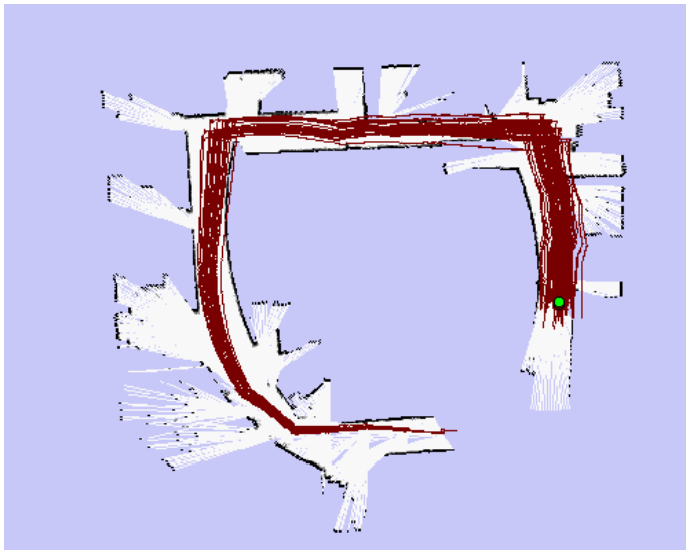

Running Example:
Traffic Modeling by Distributed Camera Networks

CPSLAB (cpslab.snu.ac.kr)

CITRIC
wireless camera mote

Traffic model

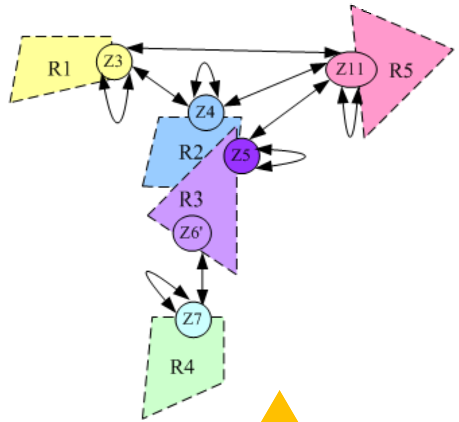

Detect humans

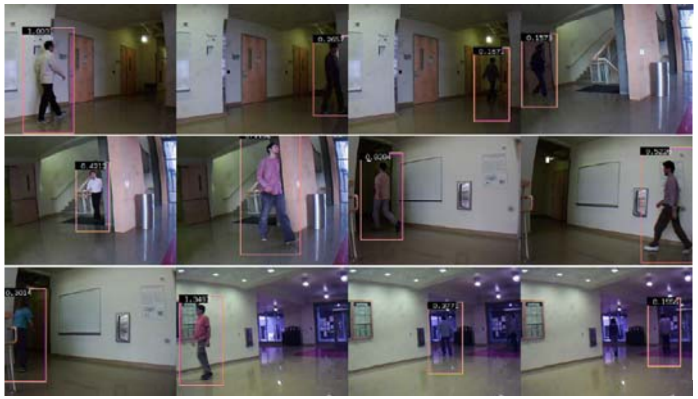

Traffic modeling
- Study the pattern of moving subjects
- Predict the next location of a moving subject
- Estimate the expected travel times
- Determine the relationship between nonoverlapping cameras

Applications:
- Abnormality detection, security
- Surveillance
- Location-based services
- Energy conservation, smart buildings, etc.

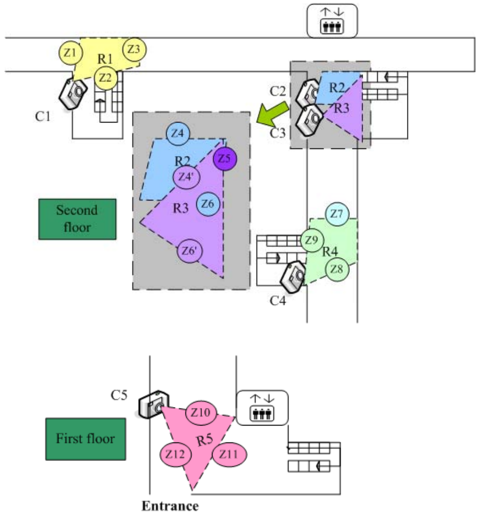

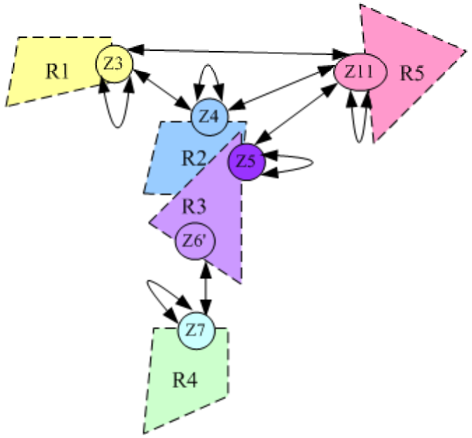

Experiment detail:
- 5 wireless camera motes
- 20+ subjects move freely
- different viewpoints
- 1400+ images for training
- 1100+ images for testing

Human detection
- Histogram of oriented gradient (HOG) based detector

Measurements
- Normalized RGB histograms of the upper torso

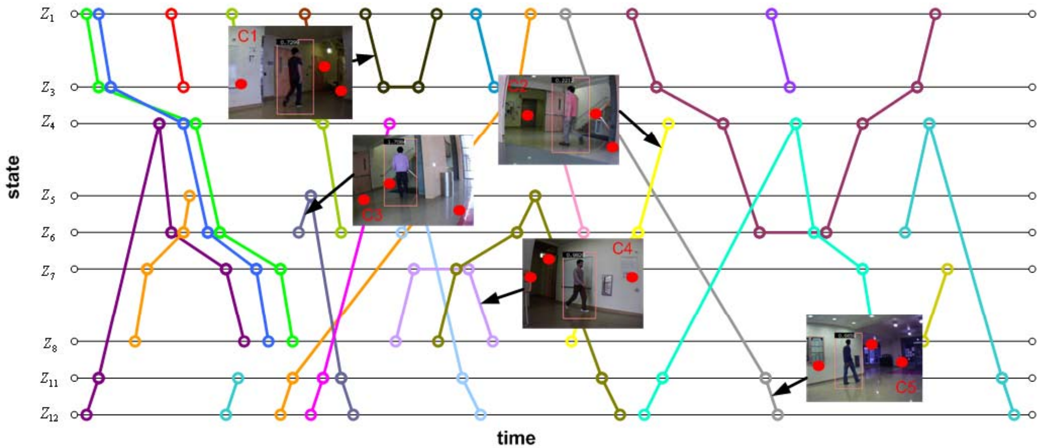

Zaihong Shuai, Songhwai Oh, and Ming-Hsuan Yang, "**Traffic Modeling and Prediction Using Camera Sensor Networks**," in Proc. of the ACM/IEEE International Conference on Distributed Smart Cameras (ICDSC), Atlanta, GA, Sep. 2010.

Outline

- Data association problem and algorithms
- Correlation decay
- Deterministic approximation algorithm for data association
- Distributed data association algorithm for sensor networks
- Simulation results
- Conclusions and future work

Data Association Problem

CPSLAB (cpslab.snu.ac.kr)

A toy problem

- Suppose that we saw two people leaving the scene from cameras
- After some time later, two people are detected by another set of cameras
- Are they the same people that we saw before or new people?

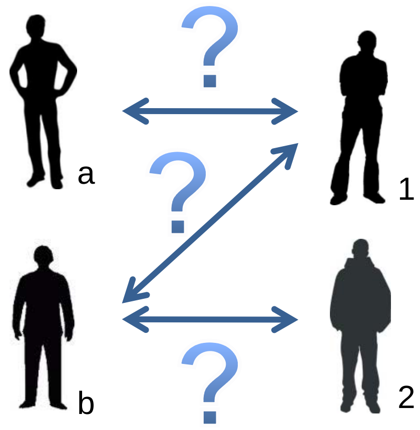

Mutual exclusiveness property

Feasible scenarios

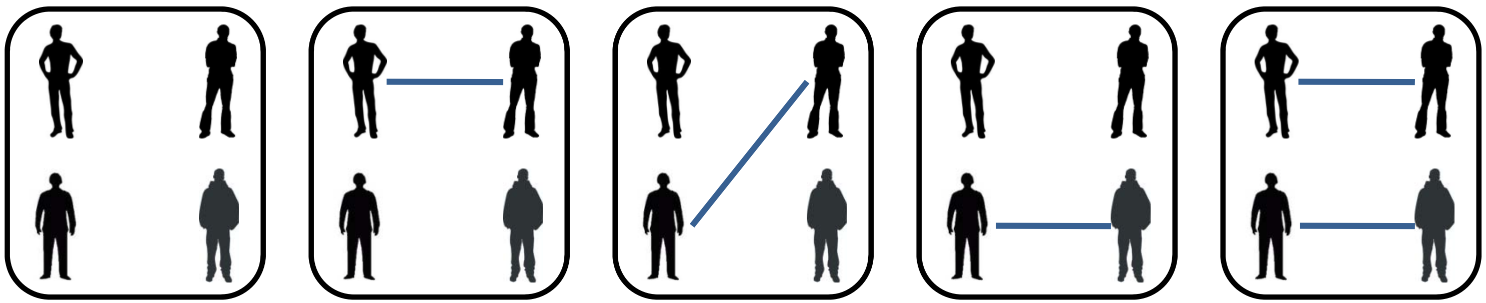

Data Association: Problem Formulation

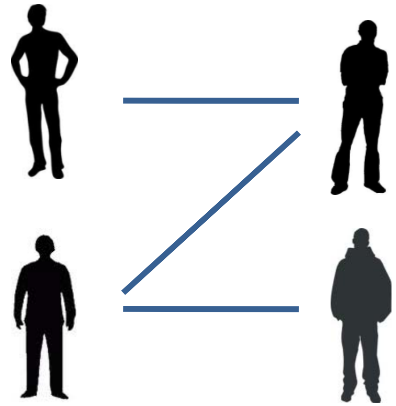

Events Measurements

G = (V,E)

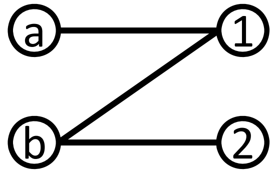

Data association graph

Feasible association events

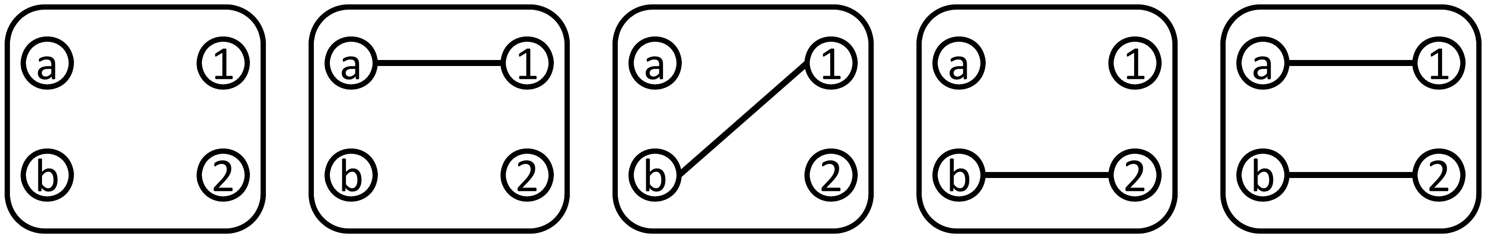

A feasible association event is a matching in G

Association Probability

Solution space, Ω

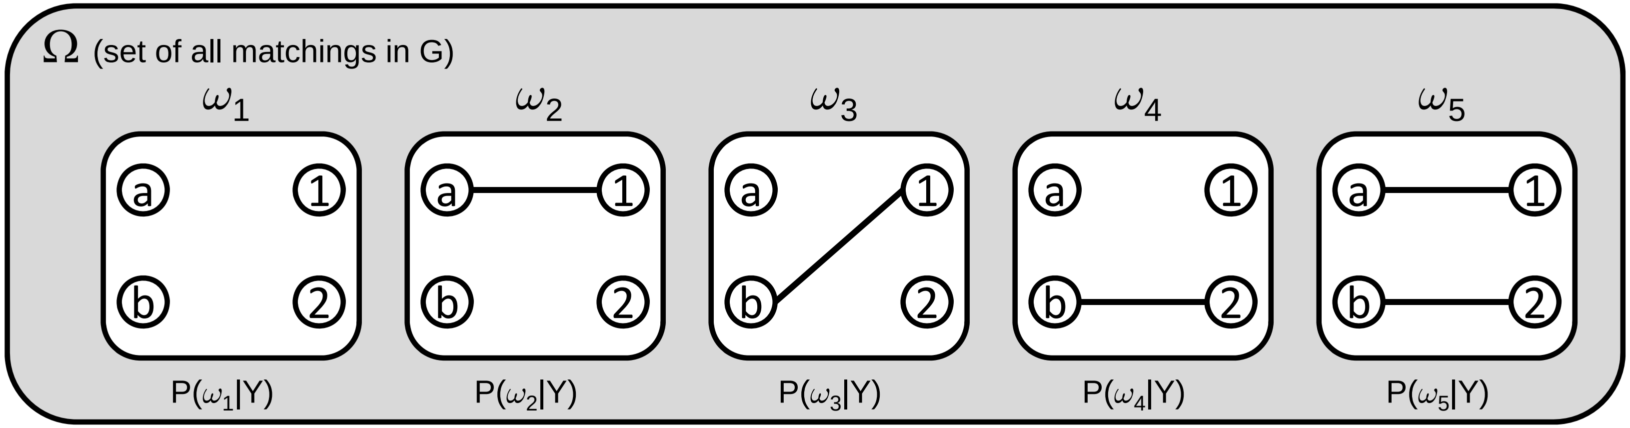

Association probability:

$$\beta_{jk} = \sum_{\omega:(j,k)\in\omega} P(\omega|Y)$$

probability that j-th measurement is associated with k-th event.

$$\begin{aligned}
\beta_{1a} &= P(\omega_2|Y) + P(\omega_5|Y) \\
\beta_{2b} &= P(\omega_4|Y) + P(\omega_5|Y) \\
\beta_{1b} &= P(\omega_3|Y) \\
\beta_{0a} &= P(\omega_1|Y) + P(\omega_3|Y) + P(\omega_4|Y), \ldots
\end{aligned}$$

State estimation:

$$\begin{aligned}
P(x_k|Y) &= \sum_{\omega\in\Omega} P(x_k|\omega, Y) P(\omega|Y) \\
&= \sum_{j=0}^{N} \beta_{jk} P(x_k|\omega_{jk}, Y)
\end{aligned}$$

But the exact computation of association probabilities is #P-complete [17].

Related Work in Data Association Algorithms

MAP (maximum a posteriori):

- Find $\omega^* = \text{argmax } P(\omega | Y)$
- Multiple hypothesis tracker (MHT) [Reid, 1979; Kurien, 1990]

Bayesian, MMSE (minimum mean square error):

- Given a function $X: \Omega \rightarrow \mathbb{R}^n$, estimate $\mathbb{E}(X | Y)$
- Joint probabilistic data association (JPDA) [Bar-Shalom, Fortmann, 1988]

Complexity of either approach is NP-hard

- [Collins & Uhlmann, 92; Poore, 95], JPDA is #P-complete

Fourier transform based approximation [Huang,Guestrin,Guibas, 2009]

Markov chain Monte Carlo data association (MCMCDA)

- Fully polynomial **randomized approximation** scheme (FPRAS)
- [Oh,Russell,Sastry, 2009]

This paper: First **deterministic approximation** algorithm

- Fully polynomial time approximation scheme (FPTAS)

Correlation Decay

Given a graph $G = (V, E)$. Let $\Omega = \Omega(G)$ be the set of all matchings of G.

Probability distribution on Ω:

$$P_G(\omega) = \frac{\lambda^{|\omega|}}{Z(G)}, \qquad \text{where} \quad Z(G) = \sum_{\omega \in \Omega} \lambda^{|\omega|}$$

partition function

Computation of the partition function Z(G) is #P-complete

Bayati et al. [32] presented a **deterministic** fully polynomial time approximation scheme (FPTAS) for computing Z(G)

Main result [32]: $Z(G)$ can be approximated from recursive computation of

$$\Phi_{\hat{G}}(v, t+1) = \frac{1}{1 + \lambda \sum_{u \in N(v,\hat{G})} \Phi_{\hat{G} \setminus \{v\}}(u, t)}$$

An algorithm $\mathcal{A}$ is a *fully polynomial time approximation scheme (FPTAS)* if, for any $\epsilon > 0$,

$$\exp(-\epsilon) \le \frac{\mathcal{A}(G)}{Z(G)} \le \exp(\epsilon),$$

in time which is polynomial in n, the number of vertices in G, and $1/\epsilon$.

[32] M. Bayati, D. Gamarnik, D. Katz, C. Nair, and P. Tetali, "**Simple deterministic approximation algorithms for counting matchings**," in Proc. of the ACM Symposium on Theory of Computing, San Diego, CA, June 2007.

Deterministic Approximation Algorithm for Data Association

CPSLAB (cpslab.snu.ac.kr)

Probability model for a data association problem

Prior: $P(\omega) \propto (\lambda_f V)^{N-|\omega|} p_d^{|\omega|} (1-p_d)^{K-|\omega|}$

Posterior:
$$\begin{aligned} P(\omega|Y) &= \frac{1}{Z_1} P(\omega) P(Y|\omega) \\ &= \frac{1}{Z_2} \lambda_f^{N-|\omega|} p_d^{|\omega|} (1-p_d)^{K-|\omega|} \prod_{(u,v)\in\omega} L_{uv} \\ &= \frac{1}{Z(G)} \left(\lambda_f^{-1} p_d (1-p_d)^{-1}\right)^{|\omega|} \prod_{(u,v)\in\omega} L_{uv} \end{aligned}$$

Recall:

$$\beta_{jk} = \sum_{\omega:(j,k)\in\omega} P(\omega|Y)$$

$$\begin{aligned} P(x_k|Y) &= \sum_{\omega\in\Omega} P(x_k|\omega, Y) P(\omega|Y) \\ &= \sum_{j=0}^{N} \beta_{jk} P(x_k|\omega_{jk}, Y) \end{aligned}$$

Posterior distribution of data association

$$\begin{aligned} P(\omega|Y) &= \frac{1}{Z(G)} \alpha^{|\omega|} \prod_{(u,v)\in\omega} L_{uv} \\ Z(G) &= \sum_{\omega\in\Omega} \alpha^{|\omega|} \prod_{(u,v)\in\omega} L_{uv} \end{aligned}$$

cf. partition function: $Z(G) = \sum_{\omega\in\Omega} \lambda^{|\omega|}$

Theorem 1:

$$\beta_{uv} = \frac{Z(G \setminus \{u,v\})}{Z(G)} \alpha L_{uv}.$$

Computation of Z(G)

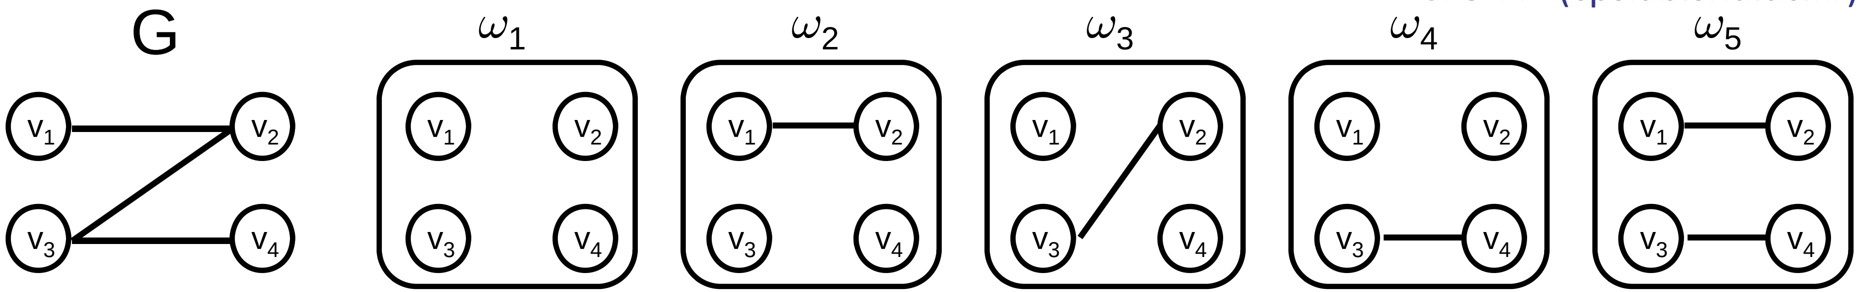

$$
\begin{aligned}
Z(G) &= Z(G_1) = \tilde{P}(\omega_1|Y) + \tilde{P}(\omega_2|Y) + \tilde{P}(\omega_3|Y) + \tilde{P}(\omega_4|Y) + \tilde{P}(\omega_5|Y) \\
Z(G \setminus \{v_1\}) &= Z(G_2) = \tilde{P}(\omega_1|Y) + \tilde{P}(\omega_3|Y) + \tilde{P}(\omega_4|Y) \\
Z(G \setminus \{v_1, v_2\}) &= Z(G_3) = \tilde{P}(\omega_1|Y) + \tilde{P}(\omega_4|Y) \\
Z(G \setminus \{v_1, v_2, v_3\}) &= Z(G_4) = \tilde{P}(\omega_1|Y) \\
Z(G \setminus \{v_1, v_2, v_3, v_4\}) &= Z(G_5) = 1
\end{aligned}
$$

unnormalized posterior

$$P_{G_1}(v_1 \notin \omega) = \frac{\tilde{P}(\omega_1|Y) + \tilde{P}(\omega_3|Y) + \tilde{P}(\omega_4|Y)}{Z(G)} = \frac{Z(G \setminus \{v_1\})}{Z(G)} = \frac{Z(G_2)}{Z(G_1)}$$

$$P_{G_2}(v_2 \notin \omega) = \frac{Z(G_3)}{Z(G_2)}$$

$$P_{G_3}(v_3 \notin \omega) = \frac{Z(G_4)}{Z(G_3)}$$

$$P_{G_4}(v_4 \notin \omega) = \frac{Z(G_5)}{Z(G_4)}$$

Lemma 1:

$$Z(G) = \frac{1}{\prod_{1 \le k \le |V|} P_{G_k}(v_k \notin \omega)}$$

Summary up to now:

- States can be estimated from association probabilities, β_{uv}
- Association probabilities β_{uv} can be computed from Z(G) and Z(G\{u,v})
- $Z(G_k)$ can be computed from $P_{G_k}(u \notin \omega)$

Lemma 2:

$$P_G(u \notin \omega) = \frac{1}{1 + \alpha \sum_{v \in N(u)} P_{G \setminus \{u\}}(v \notin \omega) L_{uv}}$$

Proof:

$$Z(G) = \underbrace{\sum_{\omega : u \notin \omega} \alpha^{|\omega|} \prod_{(u',v') \in \omega} L_{u'v'}}_{\text{matchings without } u} + \underbrace{\sum_{v \in N(u)} \sum_{\omega : (u,v) \in \omega} \alpha^{|\omega|} \prod_{(u',v') \in \omega} L_{u'v'}}_{\text{matchings with } u}$$

$$= Z(G \setminus \{u\}) + \sum_{v \in N(u)} Z(G \setminus \{u, v\}) \alpha L_{uv}$$

$$\frac{Z(G)}{Z(G \setminus \{u\})} = 1 + \sum_{v \in N(u)} \frac{Z(G \setminus \{u, v\})}{Z(G \setminus \{u\})} \alpha L_{uv}$$

Approximating $P_G(u \notin \omega)$

CPSLAB (cpslab.snu.ac.kr)

We approximate $P_G(u \notin \omega)$ by

$$\Phi_{\hat{G}}(u, t+1) = \frac{1}{1 + \alpha \sum_{v \in N(u)} \Phi_{\hat{G}\setminus\{u\}}(v,t) L_{uv}}$$

with $t < n$, where $\hat{G}$ is a subgraph of G and $u \in V(\hat{G})$.

Theorem 2: For every vertex $v \in V$ and every even $t > 0$,

$$\left| \log P_G(v \notin \omega) - \log \Phi_G(v,t) \right| \leq \left(1 - \frac{1}{1 + \frac{L_{\max}}{L_{\min}} \sqrt{1 + \alpha \Delta L_{\max}}} \right)^{t/2} \log(1 + \alpha \Delta L_{\max}).$$

Theorem 3: With $t = O(\log n \cdot 1/\epsilon)$, for any $(u, v) \in E$,

$$\exp(-\epsilon) \leq \frac{\hat{\beta}_{uv}}{\beta_{uv}} \leq \exp(\epsilon),$$

Theorem 4: There is an FPTAS for computing association probabilities.

* See the paper for an exact description of the algorithm.

Data Association in Sensor Networks

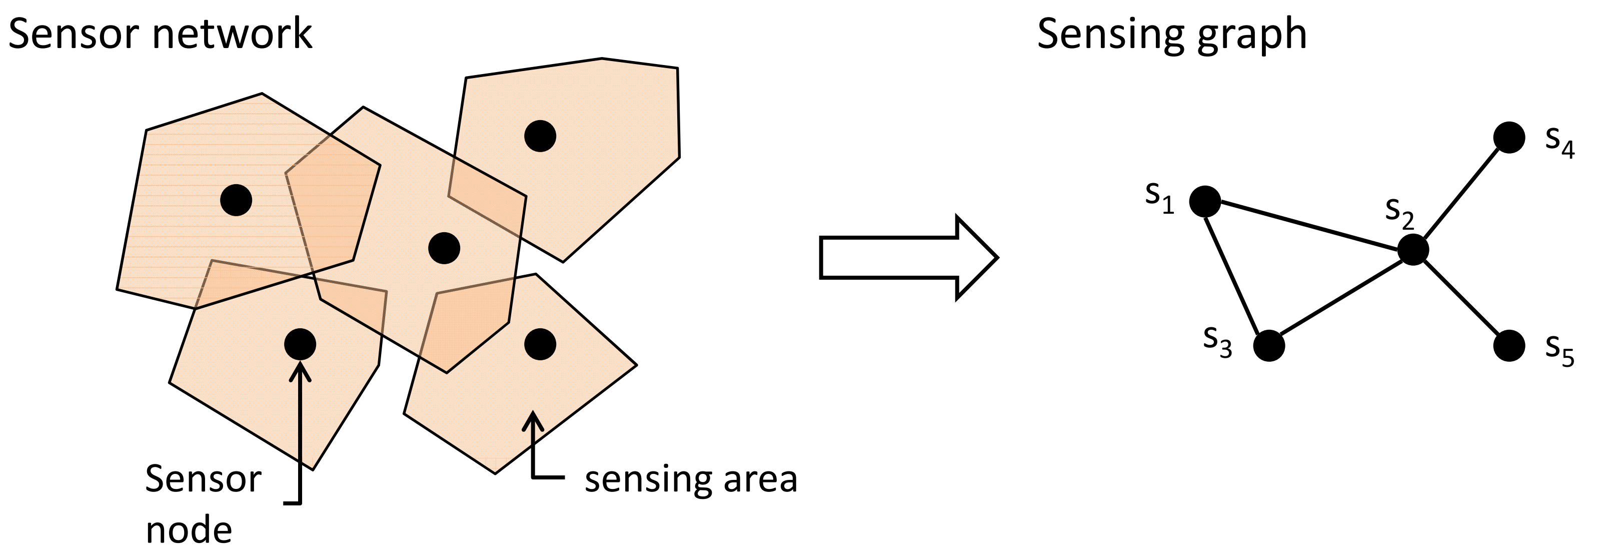

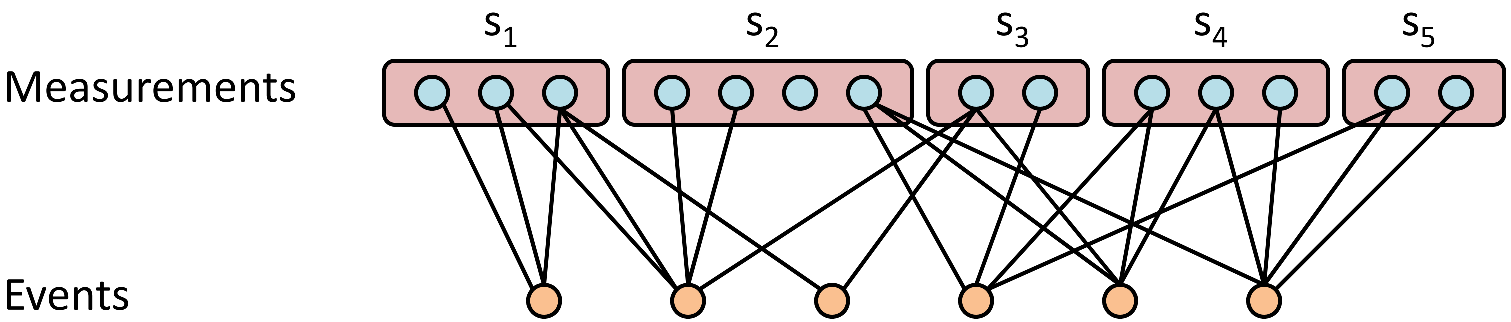

Data association graph

Distributed Data Association Algorithm for Sensor Networks

Sensing graph

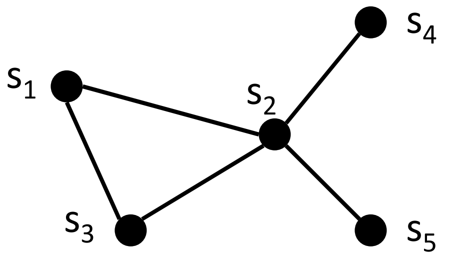

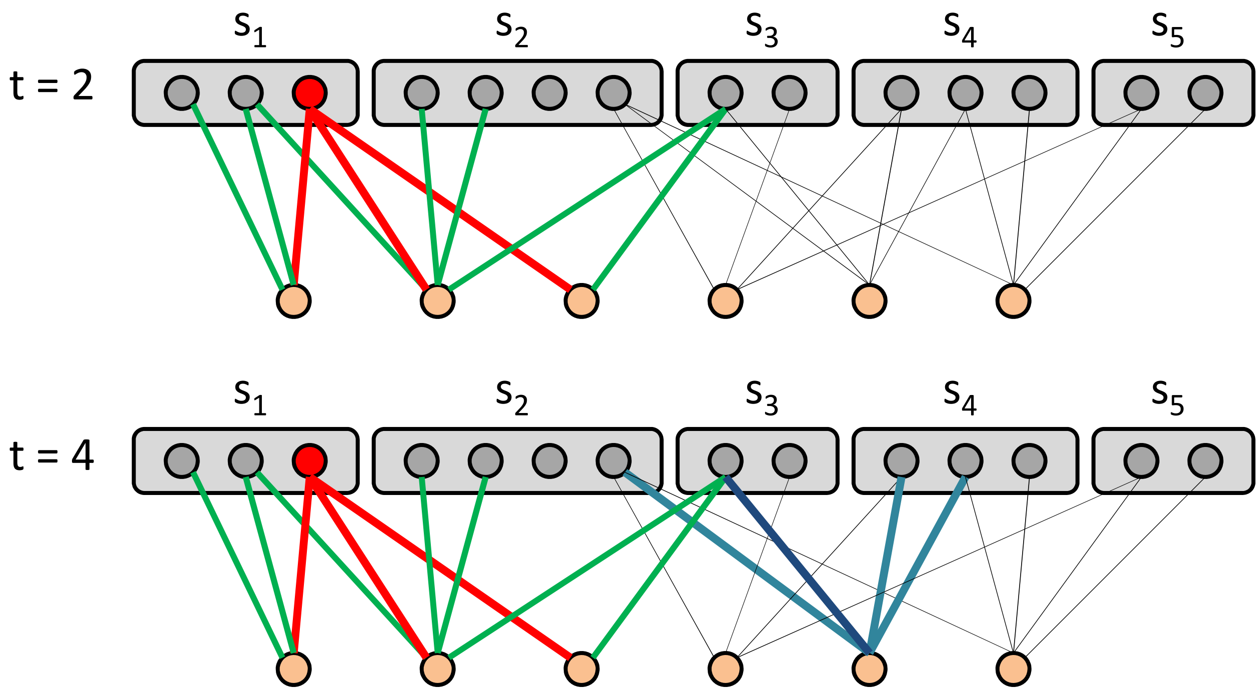

For the approximate computation level of t, we only need measurements from t/2-hop neighbors in the sensing graph.

Naturally decomposes the problem into a set of local problems.

Distributed data association algorithm for sensor networks.

* See the paper for an exact description of the algorithm.

Simulation Results

CPSLAB (cpslab.snu.ac.kr)

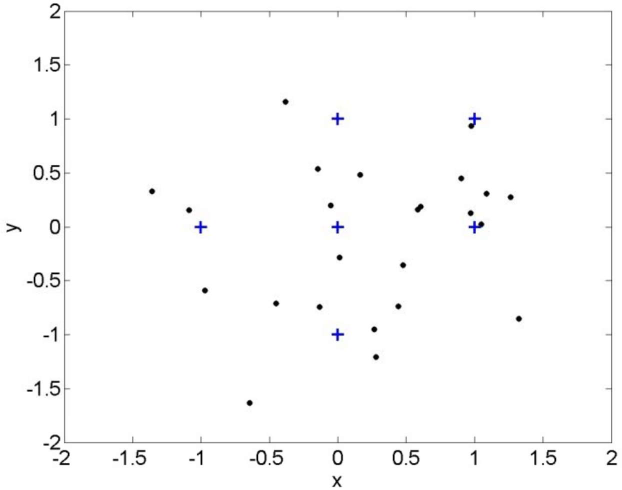

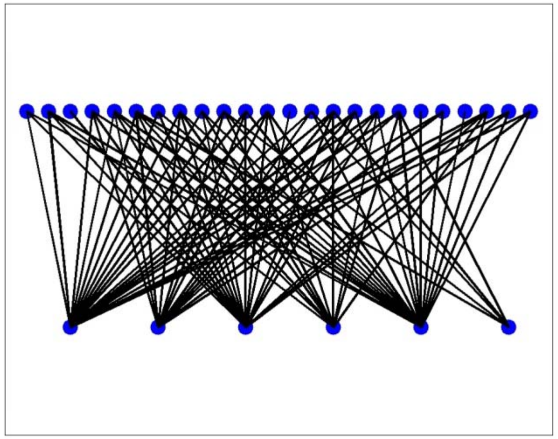

- 6 events
- 24 measurements

- 8,517,625 possible association events
- Exact computation of association probabilities take over 11 hours

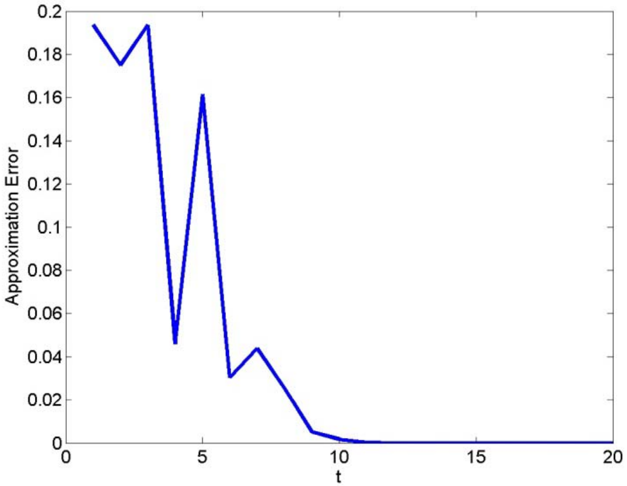

- After t = 9, the approximation error is less than 0.005 and its running time is about 26% of the exact algorithm.
- If the desired error bound is at most 0.05, then t = 6 is good enough and the corresponding running time is 0.7% of the exact algorithm.

Sensor Networks

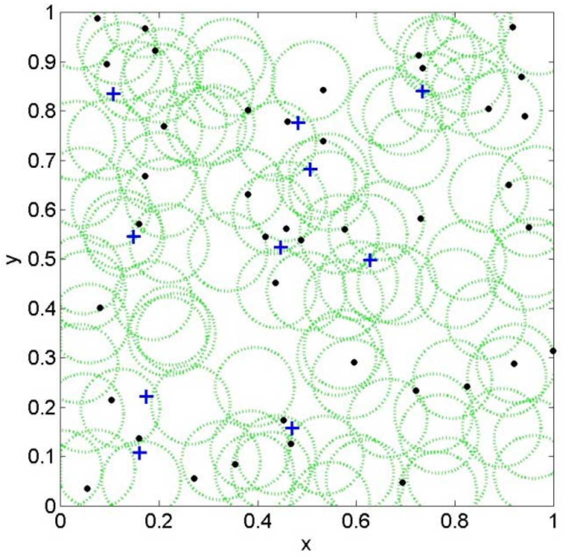

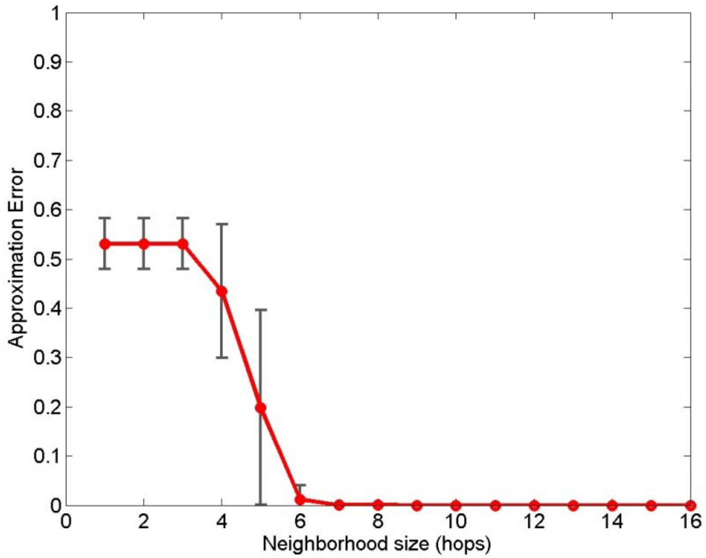

10 random scenarios of:
- 100 sensor nodes
- 10 events
- 40 measurements

- Approximation error is close to zero after 6 hops.
- A good approximation is possible using only local information from 6- hop neighbors.

An analysis of correlation decay for a complex problem can reveal how an efficient distributed approximation algorithm can be constructed for the problem.

Conclusions and Future Work

First deterministic approximation algorithm for data association

- Based on the idea of correlation decay
- Scalable: naturally decomposes a problem into a set of local problems

The technique can be applied to other problems to derive distributed approximation algorithms for NP-hard problems

Future work

- Computation between levels
- Multiple detections of the same event by different sensors
- Experiments using distributed wireless camera networks
- Parallel processing on GPUs

Thank you

http://cpslab.snu.ac.kr

Seoul National University