Tracking Multiple Objects using Sensor Networks and Camera Networks

Songhwai Oh
EECS, UC Berkeley

(sho@eecs.berkeley.edu)

Chess Review
May 11, 2005
Berkeley, CA

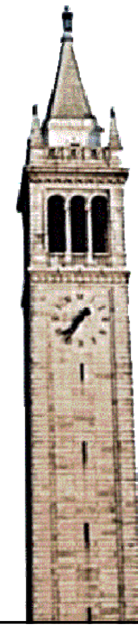

Applications

- Surveillance and security
- Search and rescue
- Disaster and emergency response system
- Pursuit evasion games [Schenato, Oh, Sastry, ICRA'05]
- Inventory management
- Spatio-temporal data collection
- Visitor guidance and other location-based services

Outline

- Multiple-target tracking problem
- Markov chain Monte Carlo data association (MCMCDA) algorithm
- Hierarchical multiple-target tracking algorithm for sensor networks
- Distributed multiple-target tracking using camera networks

Multiple-Target Tracking Problem

- Given
  - Dynamics and measurement models
  - Sensor and clutter models, e.g., detection probability (p_{d}), number of false alarms (Poisson(λ_{f}))
  - Target appearance and disappearance models, e.g., number of new targets (Poisson(λ_{b})), termination probability (p_{z})
  - A set of noisy unlabeled observations $Y = \{y_1, \ldots, y_T\}$, where $y_t = \{y_t^j : 1 \leq j \leq n_t\}$
- Find
  - Number of targets K
  - Initiation and termination times of all targets $\{(t_{\mathrm{i}}^k, t_{\mathrm{f}}^k) : 1 \leq k \leq K\}$
  - Tracks of all targets $\{x_t^k : t_{\mathrm{i}}^k \leq t \leq t_{\mathrm{f}}^k, 1 \leq k \leq K\}$
- Requires solutions to both data association and state estimation problems

Solution Space of Data Association Problem

- Solution space Ω is a collection of partitions of Y such that, for $\omega \in \Omega$,
  - $\omega = \{\tau_0, \tau_1, \ldots, \tau_K\}$, "joint association event"
  - τ_0 is a set of false alarms
  - τ_k is a feasible track ($1 \leq k \leq K$)
- Posterior: ($a_t, f_t \sim$ Poisson)

$$P(\omega|Y) \propto P(Y|\omega) \prod_{t=1}^{T} p_{\mathrm{z}}^{z_t}(1-p_{\mathrm{z}})^{c_t} p_{\mathrm{d}}^{d_t}(1-p_{\mathrm{d}})^{u_t} \lambda_{\mathrm{b}}^{a_t} \lambda_{\mathrm{f}}^{f_t}$$

  - $P(Y|\omega)$ is the likelihood of observations Y given ω, which can be computed based on dynamics and measurement models
  - z_t, c_t, d_t, u_t, a_t are numbers of terminated, continuing, detected, undetected, and new targets, respectively. f_t is the number of false alarms

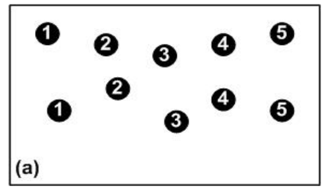

(a) Observations Y

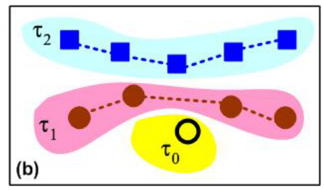

(b) Example of a partition ω of Y

Two Possible Solutions to Data Association Problem

- MAP (maximum a posteriori): find $\omega^* = \arg\max_{\omega \in \Omega} P(\omega|Y)$
  - Surveillance model [Sittler,64]
  - Multiple hypothesis tracker (MHT) [Reid,79; Kurien,90]
- MMSE (minimum mean square error): given a function $X : \Omega \rightarrow \mathbb{R}^n$, estimate $\mathbb{E}(X|Y)$
  - Optimal Bayesian filter: $\mathbb{E}(x_t^k|Y)$
  - Joint probabilistic data association (JPDA) [Bar-Shalom & Fortmann,88]
    * Finds *suboptimal* MMSE estimates
- The complexity of either approach is NP-hard [Poore,95; Collins & Uhlmann,92]

Outline

- Multiple-target tracking problem
- **Markov chain Monte Carlo data association (MCMCDA) algorithm**
  - [Oh, Russell, Sastry, CDC 2004]
- Hierarchical multiple-target tracking algorithm for sensor networks
- Distributed multiple-target tracking using camera networks

Markov Chain Monte Carlo (MCMC)

- A general method to generate samples from a complex distribution
- For some complex problems, MCMC is the only known general algorithm that finds a good approximate solution in polynomial time [Jerrum, Sinclair, 1996]
- Applications:
  - Complex probability distribution integration problems
  - Counting problems (#P-complete problems)
  - Combinatorial optimization problems
- Data association problem has a very complex probability distribution

MCMC Data Association (MCMCDA)

- Start with some initial state ω_1 2 Ω

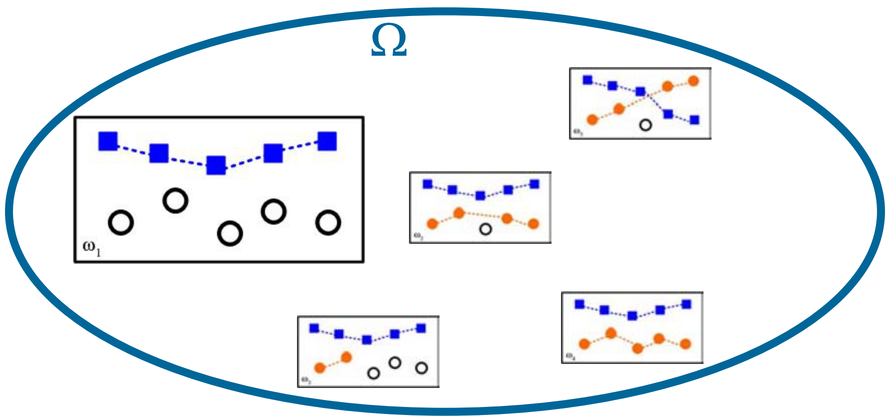

MCMC Data Association (MCMCDA)

- Propose a new state ω' » $q(\omega_n,\omega')$
- q: Ω £ 2^{Ω} ! [0,1], proposal distribution $q(\omega_n,\omega')$ = probability of proposing ω' when the chain is in ω_n

- $q(\omega_n,\omega')$ is determined by 8 moves:

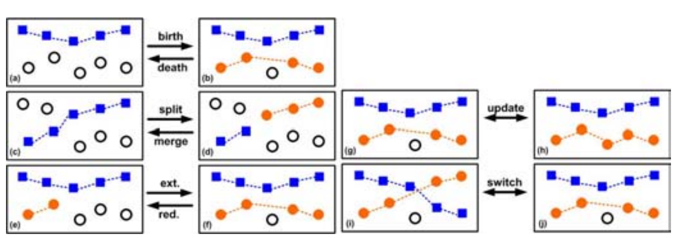

MCMC Data Association (MCMCDA)

- Accept the proposal with probability

$$A = \min\left(1, \frac{\pi(\omega')q(\omega', \omega_n)}{\pi(\omega_n)q(\omega_n, \omega')}\right)$$

$\pi(\omega) = P(\omega|Y)$, Y = observations

- If accepted,

$\omega_{n+1} = \omega'$

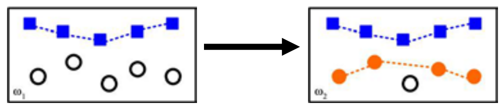

- If not accepted,

$\omega_{n+1} = \omega_n$

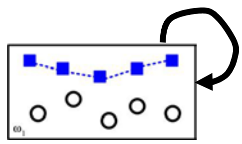

MCMC Data Association (MCMCDA)

- Repeat it for N steps

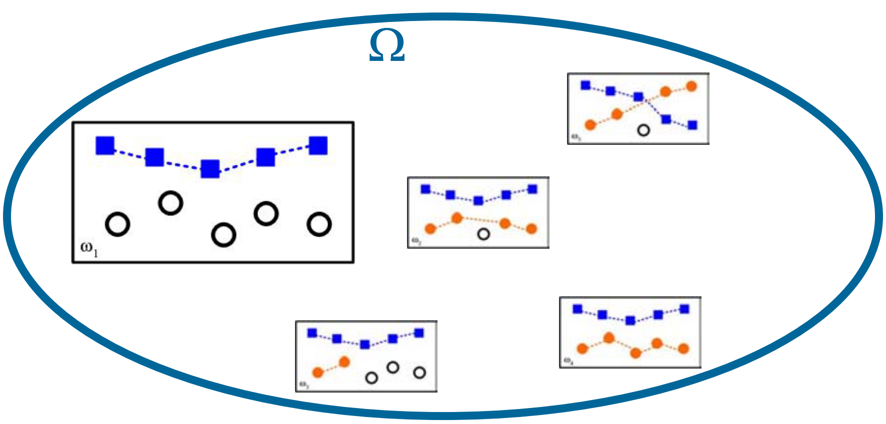

MCMC Data Association (MCMCDA)

Given $X : \Omega \to \mathbb{R}^d$, compute

$$\hat{X} = \frac{1}{N} \sum_{n=1}^{N} X(\omega_n),$$

where ω_n is the state at the n-th iteration.

If the chain is aperiodic and irreducible, then by the ergodic theorem [Roberts, 1996], as $N \to \infty$,

$$\hat{X} \to \mathbb{E}_\pi(X|Y).$$

But how fast does it converge?

Polynomial-Time Approximation to Joint Probabilistic Data Association*

Assuming K is fixed, the state X_t^k of target k at time t can be computed as (ω_{jk} is the event j-th observations in Y_t is from k-th target, $Y_{1:t} = \{Y_1, \ldots, Y_t\}$):

$$\begin{aligned} \mathbb{E}(X_t^k|Y_{1:t}) &= \sum_{\omega\in\Omega} \mathbb{E}(X_t^k|\omega, Y_{1:t}) P(\omega|Y_{1:t}) \\ &= \sum_{j=1}^{N_t} \underbrace{\mathbb{E}(X_t^k|\omega_{jk}, Y_{1:t})}_{\text{single target est.}} \underbrace{P(\omega_{jk}|Y_{1:t})}_{\beta_{jk}} \end{aligned}$$

Theorem: Let N_t be the number of measurements, $0 < \epsilon_1, \epsilon_2 \leq 1$, and $0 < \eta < .5$. Then, with at most

$$N = O(\epsilon_1^{-2}\epsilon_2^{-1} \log \eta^{-1} N_t(N_t \log N_t + \log(\epsilon_1^{-1}\epsilon_2^{-1})))$$

samples, MCMCDA finds estimates $\hat{\beta}_{jk}$ for β_{jk} with probability at least $1-\eta$, such that, for $\beta_{jk} \geq \epsilon_2$, $\hat{\beta}_{jk}$ estimates β_{jk} within ratio $1+\epsilon_1$, i.e., $(1-\epsilon_1)\beta_{jk} \leq \hat{\beta}_{jk} \leq (1+\epsilon_1)\beta_{jk}$, and, for $\beta_{jk} < \epsilon_2$, $|\hat{\beta}_{jk} - \beta_{jk}| \leq (1+\epsilon_1)\epsilon_2$.

In other words, MCMCDA finds a good approximation of β_{jk} in polynomial time.

* [Oh,Sastry, ACC 2005]

MCMCDA Highlights

- Optimal Bayesian filter in the limit
- Provides approximate solutions to both MAP and MMSE
- Avoids the enumeration of all feasible events

- Single-scan MCMCDA approximates JPDA in polynomial time with guaranteed error bounds [Oh,Sastry, ACC 2005]
- Outperforms Multiple Hypothesis Tracking algorithm [Oh, Russell,Sastry, CDC 2004]

- Statistically sound approach to initiate and terminate tracks
  - Can track an unknown number of targets
  - Suitable for an autonomous surveillance system

- **Easily distributed and suitable for sensor networks**

Outline

- Multiple-target tracking problem
- Markov chain Monte Carlo data association (MCMCDA) algorithm
- **Hierarchical multiple-target tracking algorithm for sensor networks**
  - **[Oh, Schenato, Sastry, ICRA 2005]**
- Distributed multiple-target tracking using camera networks

Challenges in Sensor Networks

- Limited capabilities of a sensor node
  - Limited supply of power
  - Short communication range
  - High transmission failure rates
  - High communication delay rates
  - Limited amount of memory and computational power
- Inaccuracy of sensors
  - Short sensing range
  - Low detection probabilities
  - High false detection probabilities
- Inaccuracy of sensor network localization

Hierarchical MTT Algorithm

Requirements

- Autonomous
  - Ability to initiate and terminate an unknown number of tracks
  - Required for distributed tracking
- Low computation and memory
- Robust against
  - transmission failures
  - communication delays
  - localization error
- Scalable
- Low communication load

Hierarchical MTT (HMTT)

← **MCMCDA** (Autonomous; Low computation and memory; Robust against)

← **Hierarchy** (Scalable)

← **Local data fusion** (Low communication load)

Algorithm Overview

- Assume a few supernodes, e.g., Intel's Stargate, iMote2
  - Longer communication range
- Regular sensors are grouped by supernodes

1. Sensors detect an object and fuse local data
2. Fused data are transmitted to the nearest supernode
3. Each supernode estimates tracks by running the online MCMCDA
4. Supernodes exchange tracks with each other
5. Track-level data association by MCMCDA to resolve duplicate tracks

Simulation

- To measure the performance of the algorithm against localization error, transmission failure and communication delay
- Setup
  - 100x100 grid sensor field
  - Single supernode at the center
  - Separation between neighboring nodes = 1 unit length
  - Signal strength sensor model (noisy range data only)
  - Data fusion: weighted sum of sensor positions of neighboring sensors with overlapping sensing regions (weighted by signal strengths)

Robustness against Localization Error

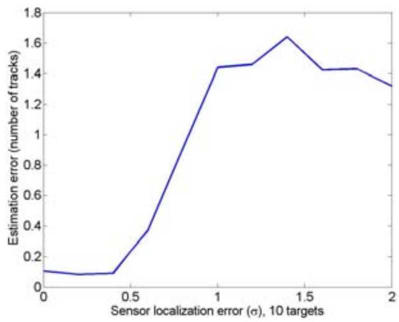

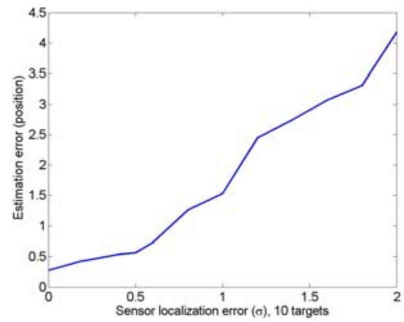

- Gaussian localization error
- **No performance loss up to the average localization error of .7 times the separation between sensors**

Robustness against Transmission Failures

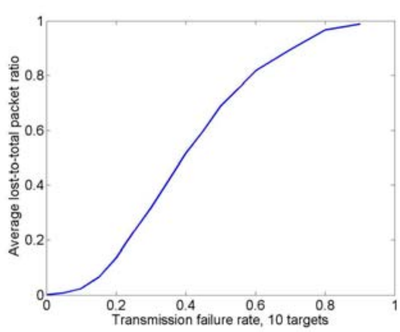

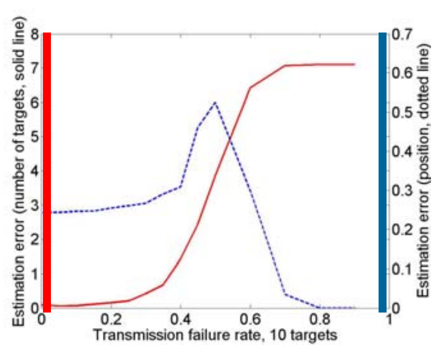

- An independent transmission failure model is assumed
- **Tolerates up to 50% lost-to-total packet ratio**

Robustness against Communication Delays

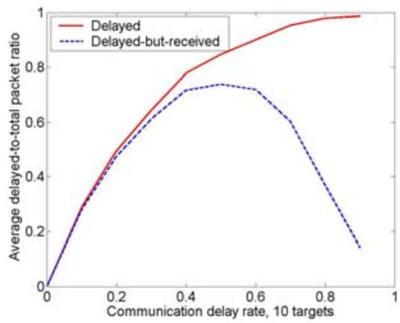

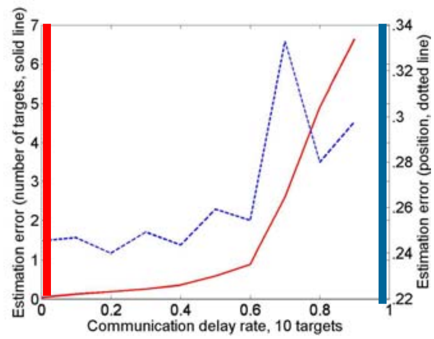

- An independent communication delay model is assumed
- **Tolerates up to 90% delayed-to-total packet ratio**

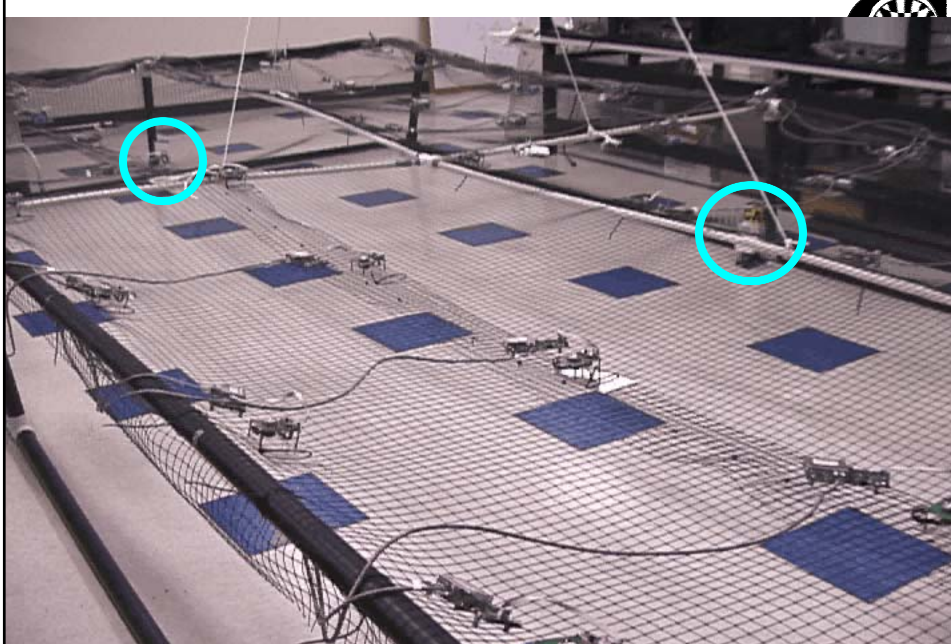

* Joint work with Phoebus Chen

Experiment Results

Scenario

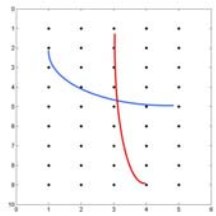

Tracks from HMTT

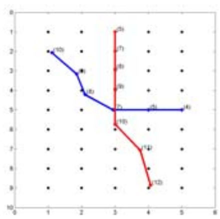

Execution of HMTT

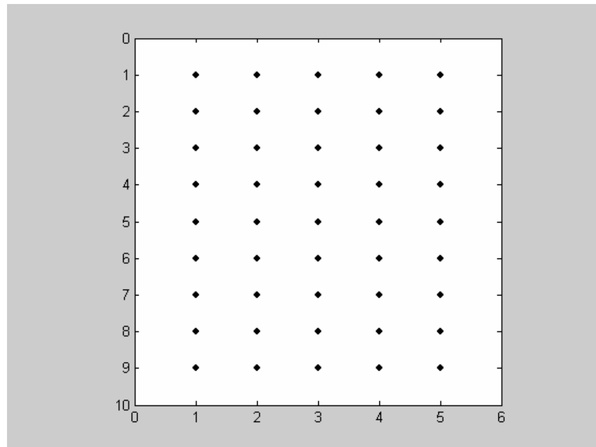

detection

potential tracks

estimated tracks

Outline

- Multiple-target tracking problem
- Markov chain Monte Carlo data association (MCMCDA) algorithm
- Hierarchical multiple-target tracking algorithm for sensor networks
- **Distributed multiple-target tracking using camera networks**
  - **In progress**

Camera Networks

- Concept: sensor network of cameras
- Benefits
  - Less installation cost (no wires)
  - Rich set of measurements (color, shape, position, etc.)
  - Reliable verification of an event
- Applications
  - Surveillance and security
  - Situational awareness for support of decision making
  - Emergency and disaster response

Cory Hall Camera Network Testbed

- DVR on 1st floor operations room
- 4 omnicams/12 perspective cameras, focus on 2 busier hallways

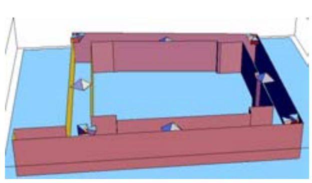

Video

Conclusions

- Multiple-target tracking is a challenging problem
- It gets more challenging when data is collected via unreliable networks such as sensor networks
- MCMCDA is an optimal Bayesian filter in the limit and provides superior performance
- Hierarchical multiple-target tracking algorithm
  - No performance loss up to the average localization error of .7 times the separation between sensors
  - Tolerates up to 50% lost-to-total packet ratio
  - Tolerates up to 90% delayed-to-total packet ratio
- Camera networks
  - Presents a new set of challenges
  - Hierarchical multiple-target tracking algorithm is applied to track multiple targets in a distributed manner