

People Tracking with Anonymous and ID-Sensors Using Rao-Blackwellised Particle Filters

Dirk Schulz, Dieter Fox, and Jeffrey Hightower

Dept. of Computer Science & Engineering
University of Washington
Seattle, WA

Abstract

Estimating the location of people using a network of sensors placed throughout an environment is a fundamental challenge in smart environments and ubiquitous computing. Id-sensors such as infrared badges provide explicit object identity information but coarse location information while anonymous sensors such as laser range-finders provide accurate location information only. Tracking using both sensor types simultaneously is an open research challenge. We present a novel approach to tracking multiple objects that combines the accuracy benefits of anonymous sensors and the identification certainty of id-sensors. Rao-Blackwellised particle filters are used to estimate object locations. Each particle represents the association history between Kalman filtered object tracks and observations. After using only anonymous sensors until id estimates are certain enough, id assignments are sampled as well resulting in a fully Rao-Blackwellised particle filter over both object tracks and id assignments. Our approach was implemented and tested successfully using data collected in an indoor environment.

1 Introduction

Accurate and reliable tracking of people using sensors placed throughout an environment is a fundamental problem relevant to several research communities. Knowing the locations of people is of critical importance for research investigating high-level state estimation, plan recognition, and learning of human activity patterns for applications such as work flow enhancement and health monitoring.

Over the years, many location estimation approaches have been introduced using sensors such as cameras, laser range-finders, infrared and ultrasound sensors, and wireless networking infrastructure [Hightower and Borriello, 2001]. A crucial aspect of these sensors is whether they provide explicit information about the identity of a person. *Anonymous sensors* such as radar, reflective ultrasound transducers, and scanning laser range-finders provide accurate location and appearance information, but do not provide explicit identity information. *Id-sensors* like infrared and ultrasound badge systems do provide explicit object identity information, but

with relatively coarse location information [Want *et al.*, 1992; Priyantha *et al.*, 2000]. Various techniques have been proposed for tracking with multiple anonymous sensors or multiple id-sensors, but the problem of integrating anonymous and id sensor information has not been addressed so far. In this paper we present an approach that combines the accuracy benefits of anonymous sensors with the identification certainty of id-sensors.

Our approach uses Rao-Blackwellised particle filters to efficiently estimate the locations and identities of multiple objects. Each particle represents a history of associations between object tracks and observations. For each particle, the individual objects are tracked using Kalman filters. Since the initial id uncertainty makes a sample-based representation of id assignments extremely inefficient, our approach starts by tracking objects using only anonymous sensors and efficiently representing estimates over object id's by keeping track of sufficient statistics. Once the id estimates are certain enough, the approach switches to sampling id assignments as well resulting in a fully Rao-Blackwellised particle filter over both object tracks and id assignments. When applied to anonymous sensors only, our method results in a new Rao-Blackwellised approach to multi-hypothesis tracking, which has gained substantial attention in the target tracking community [Bar-Shalom and Li, 1995].

This paper is organized as follows: Section 2 clarifies the problem. Section 3 then presents our Rao-Blackwellised particle filter approach to tracking multiple objects using only anonymous sensor information and Section 4 extends the approach to incorporate id-sensors. Our implementation and experimental results are presented in Section 5, followed by a discussion.

2 Problem Description

Figure 1 illustrates the problem of tracking multiple people with anonymous and id-sensors. The solid and dotted lines are the trajectories of person A and B, respectively. In the beginning, the identity of the two people is not known. As they walk, the anonymous sensor observes their locations frequently. Since the people are far enough apart, their positions can be tracked reliably using the anonymous sensor. However, until they reach id-sensor areas 3 and 4, both trajectories have the same probability of belonging to either person A or B. Hence there are two different hypotheses for the id's of the two

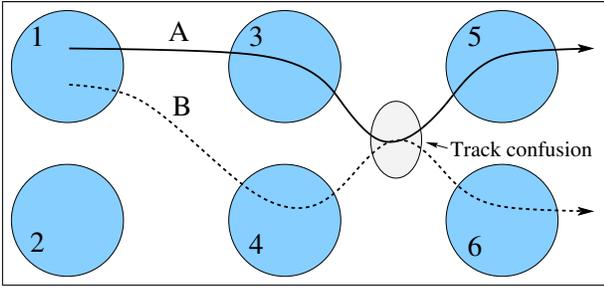


Figure 1: Example scenario: The shaded circles indicate areas covered by id-sensors such as infrared receivers. When a person wearing a badge enters such an area, the corresponding sensor issues a reading indicating the id of the person. Since these sensors provide no information about the person’s location within the area, two people in the same area can not be distinguished. Not shown is an additional anonymous sensor such as a laser range-finder. This sensor provides accurate information at a high rate about the locations of people, but no information about their id’s.

trajectories. After passing through the coverage of id-sensors 3 and 4, the ambiguity is resolved and both trajectories’ id’s are determined. Then, after the paths cross there is confusion about the continuation of the two tracks. When the people leave the light gray area the anonymous sensor can not determine which observations to associate with which trajectory. This problem is known as the *data association* problem in the multitarget tracking community [Bar-Shalom and Li, 1995]. Were there no id-sensors, it would be impossible to resolve this ambiguity. In our scenario, the ambiguity can be resolved as soon as the people reach the areas covered by id-sensors 5 and 6. To do so, however, it is necessary to maintain the hypotheses for both possible track continuations, A going down and B going up, or B going down and A going up.

The use of a combination of anonymous and id-sensors requires us to solve two types of data association problems. The first problem is the classic multitarget tracking problem of assigning anonymous observations to object tracks. This data association problem has to be solved at each point in time, resulting in $m!^k$ possible associations for tracks of length k involving m people. Fortunately, the probability distributions over these anonymous assignments are typically highly peaked, thereby allowing an efficient, sample-based representation of assignments. The second problem is the one of estimating the id of individual objects/tracks. As with anonymous assignments there are $m!$ possible assignments of id’s to tracks. Fortunately, the number of id assignments does not increase over time since the identity of objects does not change. However, due to the low spatial resolution of id-sensors, the posterior over id assignments is almost uniform during early stages of the estimation process (see Figure 1). For such “flat” distributions, a sample-based representation requires in the order of $m!$ samples, which is certainly not feasible for online tracking.

Figure 2 shows the graphical model for this tracking problem. Here, time is indexed by subscripts and the current time is denoted by k . $x_k := \{x_k^1, x_k^2, \dots, x_k^m\}$ are the current positions of the m people being tracked. Following standard notation in the tracking community, observations are denoted by

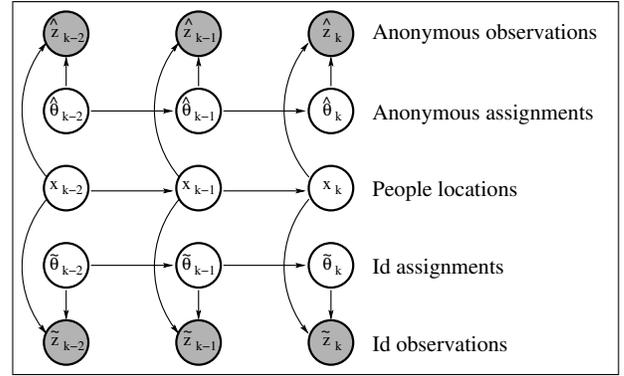


Figure 2: Graphical model for multi-object tracking with anonymous and id-sensors. x_k is the state vector describing the locations of the individual objects at time k . Objects generate anonymous observations \hat{z}_k and id observations \tilde{z}_k . The assignments of individual observations to objects are given by the hidden nodes $\hat{\theta}_k$ and $\tilde{\theta}_k$. Rao-Blackwellised particle filters *sample* assignments $\hat{\theta}_k$ and $\tilde{\theta}_k$ and solve the state updates analytically, using Kalman filters conditioned on the samples.

z_k and the complete sequence of observations up to time k is given by $z_{1:k}$ (we use **hat** and **tilde** to distinguish anonymous from id observations). The associations between object tracks x_k^i and observations are given by assignment matrices $\hat{\theta}_k$ and $\tilde{\theta}_k$. For instance, $\hat{\theta}_k(i, j)$ is one if $\hat{\theta}_k$ assigns anonymous observation \hat{z}_k^j to object x_k^i , and zero otherwise. The observations only depend on the current object positions and assignments.

The goal of tracking is to estimate the posterior over the state x_k based on all sensor information available up to time k . A factored representation of this tracking problem allows us to use Rao-Blackwellised particle filters (RBPF) which sample assignments and track the objects using a bank of m Kalman filters for each sample [Doucet *et al.*, 2000]. However, due to the large number of anonymous assignments and the “flat” distributions over id assignments, a straightforward implementation of RBPFs would require a prohibitive number of samples. In the next section we describe how to efficiently sample anonymous assignments over time. The resulting algorithm, MHT-RBPF (multi-hypothesis tracking RBPF), can track multiple objects using anonymous sensor information. Then, in section Section 4, we will describe how to extend MHT-RBPF to incorporate id-sensors.

3 MHT-RBPF: Rao-Blackwellised Particle Filters for Multi-hypothesis Tracking

Since the assignments between observations and objects/tracks are not given, we need to estimate the posterior over both object states x_k and assignments $\hat{\theta}_{1:k}$. This joint posterior can be factorized by conditioning the state x_k on the assignments $\hat{\theta}_{1:k}$:

$$p(x_k, \hat{\theta}_{1:k} | \hat{z}_{1:k}) = p(x_k | \hat{\theta}_{1:k}, \hat{z}_{1:k}) p(\hat{\theta}_{1:k} | \hat{z}_{1:k}). \quad (1)$$

The key idea of Rao-Blackwellised particle filters is to compute (1) by *sampling* assignments from $p(\hat{\theta}_{1:k} | \hat{z}_{1:k})$ and then

computing the state x_k conditioned on each sample. More specifically, each sample represents a history of data associations $\hat{\theta}_{1:k}$ and is annotated with a bank of m Kalman filters, one for each tracked person. The Kalman filters are conditioned on the data associations provided by the sample and thus can be updated efficiently using standard Kalman filter update rules for known data association [Bar-Shalom and Li, 1995]. The Kalman filter estimates, or tracks, $t_k = \langle \mu_k^i, \Sigma_k^i \rangle_{i=1, \dots, m}$ are represented by the mean and covariance of the persons' locations.

RBPF generate assignments incrementally by maintaining sample sets containing assignment histories distributed according to the posterior given by the rightmost term in (1). More specifically, at time k a sample set $S_k = \{s_k^{(\iota)}, w_k^{(\iota)} \mid 1 \leq \iota \leq N\}$ contains N weighted samples, where each sample $s_k^{(\iota)} = \langle t_k^{(\iota)}, \hat{\theta}_{1:k}^{(\iota)} \rangle$ consists of a history of assignments and the current position estimates for the m objects. The generic RBPF algorithm generates a set S_k from the previous sample set S_{k-1} and an observation \hat{z}_k by first generating new assignments $\hat{\theta}_k$ distributed according to the posterior $p(\hat{\theta}_{1:k} \mid \hat{z}_k)$. Each such assignment specifies which observations in \hat{z}_k belong to which object track. The final step consists of updating the Kalman filter tracks of each sample using the observations assigned to them by the sample [Doucet *et al.*, 2000].

3.1 Importance Sampling with Lookahead

The efficiency of RBPFs strongly depends on the number of samples needed to represent the posterior $p(\hat{\theta}_{1:k} \mid \hat{z}_{1:k})$. In this section we will devise an efficient algorithm for generating such assignments/samples. Due to the sequential nature of the estimation process, samples must be generated from the assignments of the previous time step. The posterior at time k is given by

$$\begin{aligned} p(\hat{\theta}_{1:k} \mid \hat{z}_{1:k}) &= \frac{p(\hat{z}_k \mid \hat{\theta}_{1:k}, \hat{z}_{1:k-1}) p(\hat{\theta}_{1:k} \mid \hat{z}_{1:k-1})}{p(\hat{z}_k \mid \hat{\theta}_{1:k-1}, \hat{z}_{1:k-1})} \quad (2) \\ &= \frac{p(\hat{z}_k \mid \hat{\theta}_{1:k}, \hat{z}_{1:k-1}) p(\hat{\theta}_k \mid \hat{\theta}_{1:k-1}, \hat{z}_{1:k-1})}{p(\hat{z}_k \mid \hat{\theta}_{1:k-1}, \hat{z}_{1:k-1})} p(\hat{\theta}_{1:k-1} \mid \hat{z}_{1:k-1}) \quad (3) \\ &= \frac{p(\hat{z}_k \mid \hat{\theta}_k, t_{k-1}) p(\hat{\theta}_k \mid t_{k-1})}{p(\hat{z}_k \mid t_{k-1})} p(\hat{\theta}_{1:k-1} \mid \hat{z}_{1:k-1}). \quad (4) \end{aligned}$$

Here, (2) and (3) follow from Bayes rule and the replacement of $p(\hat{\theta}_{1:k} \mid \hat{z}_{1:k-1})$ by $p(\hat{\theta}_k \mid \hat{\theta}_{1:k-1}, \hat{z}_{1:k-1}) p(\hat{\theta}_{1:k-1} \mid \hat{z}_{1:k-1})$, respectively. (4) follows from (3) by the fact that the position tracks t_{k-1} of the objects are sufficient statistics for the previous observations $\hat{z}_{1:k-1}$ and assignments $\hat{\theta}_{1:k-1}$.

Unfortunately, in most cases it is impossible to sample directly from (4). The approach most commonly used in particle filters is to evaluate (4) from right to left in a three stage process [Doucet *et al.*, 2001]: First, draw samples $s_k^{(\iota)}$ from the previous sample set using the importance weights, then draw for each such sample a new sample from the predictive distribution $p(\hat{\theta}_k \mid t_{k-1}^{(\iota)})$, and finally weight these samples proportional to the observation likelihood $p(\hat{z}_k \mid \hat{\theta}_k^{(\iota)}, t_{k-1}^{(\iota)})$. The last step, importance sampling, adjusts for the fact that samples are not drawn from the actual target distribution.

This approach has two main sources of inefficiency. First, the samples of the previous sample set are drawn without considering the most recent observation \hat{z}_k . The second source of inefficiency is that sampling from the predictive distribution in the second sampling step can be very inefficient if the observation likelihood $p(\hat{z}_k \mid \hat{\theta}_k, t_{k-1})$ is highly peaked compared to the predictive distribution $p(\hat{\theta}_k \mid t_{k-1})$ [Pitt and Shephard, 1999]. The second problem is extremely severe in our context since the predictive distribution for assignments is virtually uniform while the posterior is typically concentrated on a small set of assignments [Bar-Shalom and Li, 1995].

MCMC Sampling from the optimal distribution

Let us first discuss how to address the second problem, *i.e.* the problem of drawing samples from the posterior distribution (4), given the previous sample set. Conditioned on a specific assignment $\hat{\theta}_{1:k-1}^{(\iota)}$, the optimal sampling distribution follows from (4) as

$$p(\hat{\theta}_k \mid \hat{\theta}_{1:k-1}^{(\iota)}, \hat{z}_{1:k}) = \frac{p(\hat{z}_k \mid \hat{\theta}_k, t_{k-1}^{(\iota)}) p(\hat{\theta}_k \mid t_{k-1}^{(\iota)})}{p(\hat{z}_k \mid t_{k-1}^{(\iota)})} \quad (5)$$

It is possible to efficiently generate samples proportional to (5) using Markov Chain Monte Carlo (MCMC) techniques [Gilks *et al.*, 1996]. We apply a version of the Metropolis-Hastings algorithm that has been adapted specifically to the data association problem [Dellaert *et al.*, 2003]. In a nutshell, the idea of Metropolis-Hastings is to sample states from an ergodic Markov chain with the posterior as stationary distribution. Such a Markov chain is constructed by choosing a candidate for the next state s' given the current state s according to a proposal distribution $q(s' \mid s)$. This state transition is accepted with probability

$$\alpha(s, s') = \min \left(1, \frac{\pi(s') q(s \mid s')}{\pi(s) q(s' \mid s)} \right), \quad (6)$$

where $\pi(s)$ is the intended stationary distribution. In our case the states are the possible assignments $\hat{\theta}_k$ and $\pi(s)$ is the optimal sampling distribution (5). The efficiency of the Metropolis-Hastings method strongly depends on the choice of the proposal distribution q . We use an efficient approach called smart chain flipping. Smart chain flipping permutes the assignments of a subset of the objects on each transition, where the actual choice already takes the individual assignment likelihoods into account. This approach has been shown to result in improved mixing rates on assignment problems (see [Dellaert *et al.*, 2003] for details).

Assignment lookahead

So far, the samples generated in each Markov chain are distributed *proportional* to the target distribution (5). To generate samples from the desired posterior, we still have to weight all samples in the Markov chain generated from a sample $s_{k-1}^{(\iota)}$ proportional to $p(\hat{z}_k \mid t_{k-1}^{(\iota)})$. Since all samples generated from $s_{k-1}^{(\iota)}$ will get the same importance weight, we can improve the efficiency of the sampler by incorporating this importance factor into the weight $w_{k-1}^{(\iota)}$ of the *previous* sample $s_{k-1}^{(\iota)}$:

$$\hat{w}_{k-1}^{(\iota)} \propto w_{k-1}^{(\iota)} p(\hat{z}_k \mid t_{k-1}^{(\iota)}). \quad (7)$$

1.	Inputs: $S_{k-1} = \{ \langle s_{k-1}^{(\iota)}, w_{k-1}^{(\iota)} \rangle \mid \iota = 1, \dots, N \}$, observation \hat{z}_k
2.	$S_k := \emptyset$ // <i>Initialize</i>
3.	for $\iota = 1, \dots, N$ do // <i>Generate lookahead assignments using most recent observation</i>
4.	Generate M samples $\hat{\theta}_k^{(\iota, m)}$ proportional to $p(\hat{z}_k \mid \hat{\theta}_k^{(\iota, m)}, t_{k-1}^{(\iota)}) p(\hat{\theta}_k^{(\iota, m)} \mid t_{k-1}^{(\iota)})$ using MCMC
5.	for $\iota = 1, \dots, N$ do // <i>Update importance weights based on average probability of lookahead assignments</i>
6.	$\hat{w}_{k-1}^{(\iota)} \propto \frac{w_{k-1}^{(\iota)}}{M} \sum_{m=1 \dots M} p(\hat{z}_k \mid \hat{\theta}_k^{(\iota, m)}, t_{k-1}^{(\iota)}) p(\hat{\theta}_k^{(\iota, m)} \mid t_{k-1}^{(\iota)})$
7.	for $\iota = 1, \dots, N$ do // <i>Sample $s_{k-1}^{(\iota)}$ using updated weights and draw $s_k^{(\iota)}$ from corresponding set</i>
8.	Sample $s_{k-1}^{(\iota)} = \langle t_{k-1}^{(\iota)}, \hat{\theta}_{1:k-1}^{(\iota)} \rangle$ from S_{k-1} with probability proportional to the updated importance weights $\hat{w}_{k-1}^{(\iota)}$.
9.	Draw an assignment $\hat{\theta}_k^{(\iota)}$ from the corresponding Markov chain generated in step 4.
10.	Update the position estimates $t_k^{(\iota)}$ using Kalman filter updates with \hat{z}_k , $t_k^{(\iota)}$, and $\hat{\theta}_k^{(\iota)}$
11.	$s_k^{(\iota)} := \langle t_k^{(\iota)}, \hat{\theta}_{1:k}^{(\iota)} \rangle$; $S_k := S_k \cup \{ \langle s_k^{(\iota)}, \frac{1}{N} \rangle \}$
12.	return S_k

Table 1: MHT-RBPF algorithm.

That is, the importance weight of the previous sample is updated by the ability of the tracks associated with the sample to *predict* the next observation. Accurate computation of the right term in (7) requires summation over all possible next assignments $\hat{\theta}_k$, as done by [Morales-Menéndez *et al.*, 2002]. Since in our case the number of assignments can be prohibitively large, we estimate (7) using the samples generated in the MCMC step. Unfortunately, computing $p(\hat{z}_k \mid t_{k-1}^{(\iota)})$ is equivalent to computing the normalization factor of a Markov chain, which is not possible in general [Gilks *et al.*, 1996]. In our case, however, we do not need the absolute value of the normalization constant for each chain, but only the value *relative* to the normalizers of the other Markov chains. Since the distributions (5) have similar shapes for all chains (they are highly peaked), we can estimate the relative normalization constants by the average probabilities of the samples in the different Markov chains.

More specifically, let M samples be drawn from each Markov chain. Let $\hat{\theta}_k^{(\iota, m)}$ denote the m -th sample drawn from the Markov chain associated with sample $s_{k-1}^{(\iota)}$. Then the updated weight $\hat{w}_{k-1}^{(\iota)}$ of this sample is given by

$$\hat{w}_{k-1}^{(\iota)} \propto \frac{w_{k-1}^{(\iota)}}{M} \sum_{m=1 \dots M} p(\hat{z}_k \mid \hat{\theta}_k^{(\iota, m)}, t_{k-1}^{(\iota)}) p(\hat{\theta}_k^{(\iota, m)} \mid t_{k-1}^{(\iota)}), \quad (8)$$

where proportionality is such that all weights sum up to one.

MHT-RBPF Algorithm

The algorithm is summarized in Table 1. In step 4, new samples are generated from the previous sample set. The average probability of these samples is used to estimate the lookahead/predictive weights of each sample of the previous set (step 6). This step also involves a normalization so that the weights sum up to one. In step 8, a sample is drawn from the previous sample set S_{k-1} . Then, for each sample drawn from S_{k-1} , we draw an assignment $\hat{\theta}_k^{(\iota)}$ from the posterior re-using the samples generated in the Markov chains in step 4. Step 10 updates the actual position estimates for the individual objects, using the corresponding assignment $\hat{\theta}_k^{(\iota)}$.

4 Tracking with Anonymous and ID-Sensors

In principle, MHT-RBPF's can be readily extended to include id-sensors. Instead of only sampling anonymous assignments, it is possible to sample both anonymous and id assignments. Such a straightforward extension, however, results in an infeasible increase in the number of samples needed during early stages of the estimation process. This has two reasons. First, each hypothesis (sample) of MHT-RBPF has $m!$ possible assignments of id's to object tracks. Second, due to the low spatial resolution of id-sensors, the posterior over id assignments is initially very uncertain and sample-based representations of such flat distributions are inherently inefficient. To overcome these difficulties we instead use a two-stage estimation process. During the initial stage, only anonymous sensors are used for object tracking while the id-sensors are simply used to estimate the identity of the different objects. Once these estimates are certain enough, the process moves into the full Rao-Blackwellisation phase, during which both anonymous and id assignments are sampled. The two phases are discussed below.

Identity estimation phase

During this phase only anonymous sensors are used to track the objects. The id-sensors are used to estimate the identity of the different objects. More specifically, for each hypothesis of the MHT-RBPF, there are $m!$ possible assignments $\tilde{\theta}_k$ of identities to tracks. In order to avoid estimating distributions over this potentially too large number of assignments, we only keep track of sufficient statistics that allow us to recover distributions over assignments. Such sufficient statistics are given by the m^2 individual assignments $\tilde{\theta}_k(i, j)$ of id's j to tracks t_k^i . The probabilities of these individual assignments can be updated recursively using the most recent id observation:

$$p(\tilde{\theta}_k(i, j) \mid \hat{\theta}_{1:k}, z_{1:k}) \propto p(\tilde{z}_k^j \mid t_k^i) p(\tilde{\theta}_{k-1}(i, j) \mid \hat{\theta}_{1:k-1}, z_{1:k-1}) \quad (9)$$

Here, \tilde{z}_k^j is an id observation corresponding to person j . To determine the assignment probabilities, we have to normalize these values by considering all possible assignments:

$$p(\tilde{\theta}_k = \tilde{\theta}^j) = \frac{\prod_{(i,j) \in \tilde{\theta}^j} \hat{\theta}'(i, j) p(\tilde{\theta}_k(i, j))}{\sum_{\tilde{\theta}''} \prod_{(i,j) \in \tilde{\theta}''} \tilde{\theta}''(i, j) p(\tilde{\theta}_k(i, j))} \quad (10)$$

Again, $\tilde{\theta}'(i, j)$ is one if $\tilde{\theta}'$ assigns id j to object track i , and zero otherwise. The computational complex computation of (10) can be avoided by *sampling* id assignments using Metropolis-Hastings, based on the individual values $\tilde{\theta}_k(i, j)$. This approach works identical to the method used to sample anonymous assignments in the MHT-RBPF algorithm.

To summarize, in the identity estimation phase, each sample $s^{(i)}$ consists of m Kalman filters and an $m \times m$ matrix $\tilde{\theta}_k$ representing the sufficient statistics $\tilde{\theta}_k(i, j)$ of the id assignments. Whenever needed, the posterior over id assignments can be computed using MCMC sampling. The id estimation stage is ended as soon as the posterior over id assignments is sufficiently peaked to allow an accurate representation with a reasonable number of samples. Currently, we estimate this condition by determining the average number of hypotheses in the Markov chains run for id assignments.

Full Rao-Blackwellisation phase

Once the id's are estimated accurately enough, we begin estimating the joint posterior of both anonymous and id assignments. This posterior is given by:

$$p(\hat{\theta}_{1:k}, \tilde{\theta}_{1:k} | \hat{z}_{1:k}, \tilde{z}_{1:k}) \propto p(\hat{z}_k | \hat{\theta}_k, t_{k-1}) p(\tilde{z}_k | \tilde{\theta}_k, t_{k-1}) \cdot p(\hat{\theta}_k | t_{k-1}) p(\tilde{\theta}_k | \tilde{\theta}_{k-1}, t_{k-1}) p(\hat{\theta}_{1:k-1}, \tilde{\theta}_{1:k-1} | \hat{z}_{1:k}, \tilde{z}_{1:k}). \quad (11)$$

Note that id assignments are sampled only once during the complete estimation process, since after an id assignment is sampled, the identities of the attached object tracks are fixed. From then on, id observations serve two purposes: First they provide information of object positions, and, second, they provide information for weighting anonymous hypotheses, thereby improving the object estimates considerably.

5 Experiments

To validate our approach we captured trace logs of six people simultaneously walking around the cubical areas of the office environment outlined in Figure 3. Each person was wearing a small id-sensor track-pack consisting of two infrared badges and an ultrasound badge. Infrared and ultrasound receivers were installed throughout the ceiling. The entire scene was continually observed by two wall-mounted laser range-finders scanning at chest height just above the cubical partitions. The duration of the log was 10 minutes, during which the individual people moved between 230 meter and 410 meter. In this challenging data log, the paths of people frequently crossed each other and there were situations in which up to 4 people were occluded by others. To validate our tracking algorithm quantitatively, we carried out a series of experiments based on this data log and on additional simulation runs (see www.cs.washington.edu/robotics/people-tracking/ for visualizations).

Tracking ability

The path of the six people as estimated by our system is shown in Figure 3. This result was obtained with an MHT-RBPF using 1000 samples and Markov chains of length 100. With this setting, the algorithm was able to reliably track the six people if no lookahead was used. After determining these parameters, we carried out trial runs for different variants of our algorithm. A run was considered successful if at the end the

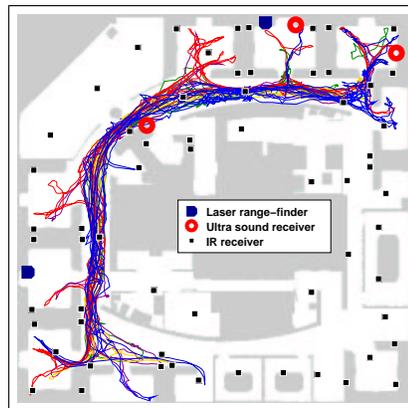


Figure 3: Outline of the Intel Research Lab Seattle. The environment is equipped with ceiling mounted ultrasound and infrared receivers. Cubical partitions are half-height (about 1.3 meters high). The two laser range-finders scan at chest height just above the partitions. Also shown are the paths of the six people as estimated by our system.

sample set contained at least one hypothesis with correct locations and identities. Each method was tested on 10 trial runs using the real data set. Without identity estimation phase, the algorithm was never able to successfully complete the data set. Next, we tested the approach when not switching to the full Rao-Blackwellisation phase, *i.e.* the algorithm remained in the identity estimation phase. Without lookahead, the update times of the approach were prohibitively large (more than 10 times real time). The lookahead resulted in significant speedup and the results are shown in the first row of the table below. Finally, we tested our two-phase approach with and without lookahead. The algorithms were able to successfully track the complete data in most cases, as can be seen in the first column. The other three columns give average time, standard deviation, and maximum time per update in seconds. The results demonstrate that both the lookahead and the two-stage process improve the performance of the tracking algorithm.

Method	Succ.	Avg.	Std.	Max.
Id estimation w. look.	10	0.062	0.125	2.9
Two-phase RBPF	10	0.154	0.170	1.5
Two-phase RBPF w. look.	9	0.036	0.060	1.2

In another set of experiments, we compared our MCMC based assignment generation to a deterministic sampling scheme, as used in traditional MHT algorithms. Here, assignments are enumerated in decreasing order wrt. their likelihood [Cox and Hingorani, 1996]. Note that generating assignments by decreasing likelihood results in estimates that are strongly biased towards more likely assignments. We performed extensive tests using real data and data simulating 20 people and found that our approach is slightly more efficient while achieving the same robustness. Since such results depend on implementational details, we did not consider them significant enough. However, the fact that smart chain flipping works at least as good as ranked assignments in practice is very encouraging, since the MCMC approach results in much less biased estimates.

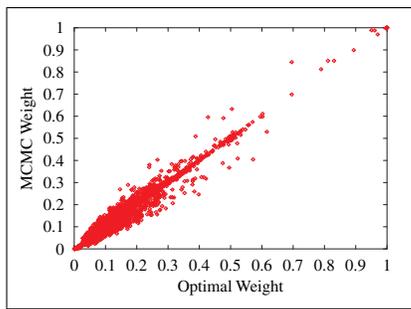


Figure 4: Optimal RBPF sample weights plotted against the sample weights computed using relative MCMC normalizers.

MCMC weight estimation

In Section 3.1, we introduced an approach to using the most recent observation in order to weight samples *before* drawing them for the next time step. So far it is not clear whether this estimation based on the probabilities of samples drawn from the Markov chains is a valid approximation. To test the quality of this approximation, we computed the optimal predictive weights by enumerating all possible assignments of the next time step. Figure 4 shows the MCMC sample weights computed in (8) plotted against the optimal weights. The similarity to $y = x$ suggests that our method of estimating Markov chain normalizers results in accurate weight estimates.

6 Conclusions and Future Work

We have presented a solution to the problem of tracking multiple people using a combination of anonymous and id sensors. The approach inherits the advantages of both sensor types, thereby being able to accurately track people and estimate their identity. Our technique uses Rao-Blackwellised particle filters to make the estimation problem tractable. We introduced several improvements to the vanilla particle filter. First, the estimation process is separated into two stages, a first stage of identity estimation and a second stage of full Rao-Blackwellisation. A second improvement is in using the most recent observation *before* sampling from the previous sample set. In contrast to [Morales-Menéndez *et al.*, 2002], the state space of our problem can become too large to allow an accurate estimate of the predictive quality of samples. Therefore, we estimate this quality using Markov chains generating samples distributed according to the posterior. We demonstrate the robustness of our approach in a challenging experiment involving six people walking through a confined office environment. We also show that our MCMC prediction is an accurate approximation of the optimal weighting function.

The approach introduced in this paper is just the first step towards a reliable and efficient tracking system. Currently, the transition between the identity estimation stage and the full Rao-Blackwellisation stage is based on a simple heuristic, namely the average number of different assignments in the Markov chains. We intend to replace this measure by a more fundamental approach such as the overall complexity of the distribution [Fox, 2002]. Another source of potential improvement lies in the handling of hypotheses. In our current system, the number of hypotheses grows extremely fast whenever several people are close to each other. We intend to overcome this

problem by clustering people into groups for which we do not attempt to estimate the individual id's, as already introduced in a similar context by [Rosencrantz *et al.*, 2003]. Finally, the recovery from tracking failures is another important issue for future research. Especially in the full Rao-Blackwellisation phase, the current approach can not recover from losing the correct id hypothesis, since id hypotheses do not change over time.

Acknowledgments

This work has partly been supported by the National Science Foundation under grant number IIS-0093406, and by DARPA's SDR Programme (contract number NBCHC020073). We are also very grateful for the help of various people at the Intel Research Lab Seattle for collecting the data needed for our experimental evaluation.

References

- [Bar-Shalom and Li, 1995] Y. Bar-Shalom and X.-R. Li. *Multitarget-Multisensor Tracking: Principles and Techniques*. Yaakov Bar-Shalom, 1995.
- [Cox and Hingorani, 1996] I. J. Cox and S. L. Hingorani. An efficient implementation of reids multiple hypothesis tracking algorithm and its evaluation for the purpose of visual tracking. *IEEE Transactions on PAMI*, 18(2), 1996.
- [Dellaert *et al.*, 2003] F. Dellaert, S.M. Seitz, C. Thorpe, and S. Thrun. EM, MCMC, and chain flipping for structure from motion with unknown correspondence. *Machine Learning*, 50(1-2), 2003.
- [Doucet *et al.*, 2000] A. Doucet, J.F.G. de Freitas, K. Murphy, and S. Russell. Rao-Blackwellised particle filtering for dynamic bayesian networks. In *Proc. of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2000.
- [Doucet *et al.*, 2001] A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte Carlo in Practice*. Springer-Verlag, New York, 2001.
- [Fox, 2002] D. Fox. KLD-sampling: Adaptive particle filters. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14 (NIPS)*, Cambridge, MA, 2002. MIT Press.
- [Gilks *et al.*, 1996] W.R. Gilks, S. Richardson, and D.J. Spiegelhalter, editors. *Markov Chain Monte Carlo in Practice*. Chapman and Hall/CRC, 1996.
- [Hightower and Borriello, 2001] J. Hightower and G. Borriello. Location systems for ubiquitous computing. *Computer*, 34(8), 2001. IEEE Computer Society Press.
- [Morales-Menéndez *et al.*, 2002] R. Morales-Menéndez, N. de Freitas, and D. Poole. Real-time monitoring of complex industrial processes with particle filters. In *Advances in Neural Information Processing Systems 15*, 2002.
- [Pitt and Shephard, 1999] M. K. Pitt and N. Shephard. Filtering via simulation: auxiliary particle filters. *Journal of the American Statistical Association*, 94(446), 1999.
- [Priyantha *et al.*, 2000] N.B. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In *Proceedings of MOBICOM 2000*, Boston, MA, 2000. ACM Press.
- [Rosencrantz *et al.*, 2003] M. Rosencrantz, G. Gordon, and S. Thrun. Locating moving entities in dynamic indoor environments with teams of mobile robots. In *Second International Joint Conference on Autonomous Agents and Multiagent Systems*, 2003.
- [Want *et al.*, 1992] R. Want, A. Hopper, V. Falcao, and J. Gibbons. The active badge location system. *ACM Transactions on Information Systems*, 10(1), 1992.