Multi–Occupant Movement Tracking in Smart Home Environments

Masoud Vatanpour Azghandi, Ioanis Nikolaidis(✉), and Eleni Stroulia

Department of Computing Science, University of Alberta, 2-21 Athabasca Hall, Edmonton, Alberta T6G 2E8, Canada
{mvatanpo,nikolaidis,stroulia}@ualberta.ca

Abstract. Recognizing the movement and activities of an individual in an indoor space is a key functionality of smart homes, as a prerequisite to providing services in support of the occupant. Focusing on the particular case of smart homes with multiple occupants, we developed a location-and-movement recognition method using many inexpensive passive infrared (PIR) motion sensors and, a small number of, more costly RFID readers. In our method, PIR sensors, placed throughout the space, recognize movement while RFID readers, placed in key locations, recognize tags worn by individuals as they pass through their coverage area. The RFID readings are used to disambiguate the trajectories constructed based on PIR sensor readings. We evaluate through simulations the effectiveness of our method under different occupancy conditions.

Keywords: Tracking · Activity recognition · Smart homes

1 Introduction

In previous work [4] we used inexpensive wireless passive infrared (PIR) sensors to determine the path of a single occupant in a "smart home" setting. Simple "anonymous" sensors such as PIR sensors are clearly inadequate when more than one individual are to be tracked and their trajectories need to be separated and labeled. In order to track multiple individuals, we are expanding our sensor toolkit to include RFID tag(s) worn by each individual and, correspondingly, RFID readers. The question addressed in this paper is whether a combination of PIR sensor deployment in an indoor space, coupled with the judicious use of RFID readers deployed at certain points in space, is an effective solution to multi-occupant localization.

RFID reader deployment is challenging and a significant contributor to total deployent costs. RFID readers need to be deployed at locations where they can be supplied by a continuous power source, i.e., powering them from batteries is not a viable option. Additionally, the use of relatively large antennas, to produce reliable readings, increases the per-reader cost and results in cumbersome placement. Hence, we are interested to reduce the number of RFID readers and deploy them in locations that are as effective as possible, i.e., where they can add the most in terms of improving the localization accuracy.

A. Geissbühler et al. (Eds.): ICOST 2015, LNCS 9102, pp. 319–324, 2015.
DOI: 10.1007/978-3-319-19312-0_29

Abstractly, the problem at hand is one of sensor fusion for the purposes of tracking individual trajectories, in a mixed environment of anonymous (PIR) and identity (RFID) sensors. The motion sensors are used to determine paths for (possibly groups of) individuals roaming the indoor space, but their paths mix and become ambiguous even if the original locations of each individual was known. The RFID readers help mitigate the ambiguity but are limited because the readers are only present in certain locations and have limited coverage. This leads us to develop a model that can assist the placement of readers using a "skeletal" tree of the paths of motion individuals follow in an indoor space.

Multi–object tracking is a well studied topic in computer vision, e.g., [2], including the fusion with other sensor data [6], but the use of cameras is perceived as privacy–intruding when compared to motion sensors and/or RFID tags and readers. RFID-based sensor fusion (usually with IMU data) solutions have been proposed [1,3], but, contrary to our approach, require individuals to carry cumbersome portable readers. In the following, Section 2 introduces the basic model and metrics used, while Section 3 provides sketches the RFID reader placement heuristic. Section 4 presents some early simulation results and Section 5 concludes with a summary of the main points of the paper.

2 System Model and Performance Metrics

We assume we know the geometry of the indoor space and the geometry of coverage of each sensor type. We also assume the existence (and use) of a heatmap of the *visitation frequency* of each point in the environment, as in [4]. The heatmap is constructed by simulating in advance the paths followed by a potential occupant between areas of interest in the space. This two-dimensional map, includes the location of the walls (W) and obstacles (O). Each of the remaining points is associated with an information utility (I), which is the probability of a person being present at that point. Overall, the heatmap contains N points, (x_1, y_1, l_1), ..., (x_N, y_N, l_N). l_i indicates the group, W, O or I, to which the point belongs. The objective of our method is, given k ($k > 1$) individuals present in and moving around the environment, to reduce the error in inferring the location of each individual in the space at any point in time.

The coverage of PIR motion sensors is modelled as a boolean rectangle, assuming that such sensors are placed on the ceiling, facing vertically down, projecting a rectangular base pyramid on the floor. We consider both the 0 degree and the 90 degree rotation rectangular footprints. To address the occlusion caused by, e.g., walls and doors, we use complex polygons to represent sensor footprints (Figure 1(a)). The coverage of RFID readers is a directed boolean sector model, in the sense of [5] (Figure 1(b)). ϕ_0 is called the *orientation angle*, ω is the *angle of view*, r is the *sensing range* resulting in a coverage represented by a circular sector.

In general, there exist locations within the environment, i.e., kitchen counter, one's own bed, etc., that tend to be *destinations* of the occupants' movement. These destinations are potential starting/ending points of paths. Considering

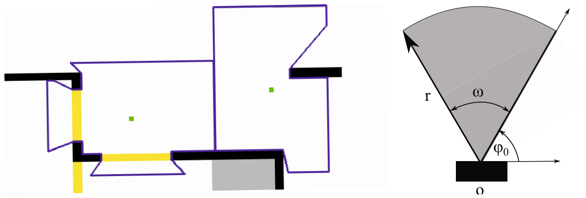

(a) (b)

Fig. 1. Sensor coverage for (a) PIR sensors [4], and, (b) RFID readers

the set $\mathcal{P}$ representing the k individuals ($\mathcal{P} = \{P_1, P_2, \ldots, P_k\}$), the sequence of destinations chosen by each occupant, i.e., their trajectory T_{P_i}, is the sequence of the person's locations-at-time, $l_{P_i,t_n} = (x_{P_i,t_n}, y_{P_i,t_n})$. A *collision* between P_i and P_j means that there exists a timestamp t_{col} when the distance between $(x_{P_i,t_{col}}, y_{P_i,t_{col}})$ and $(x_{P_j,t_{col}}, y_{P_j,t_{col}})$ is less than δ (as a convention we set δ to two times the radius of a circle representing a person's body). The definition of collision is generalized to three or more people. Past a collision point, trajectories reconstructed purely by PIR sensors can be ambiguous, since it is not possible to identify which of the colliding individual(s) continue to which "branch". It is up to the readings from RFID readers to provide the authoritative unambiguous IDs as individuals cross the reader's range; unambigusously recognizing an occupant, at some location, at some point in time, allows us to revisit previous ambiguous trajectories and infer the corresponding individuals. The process is not perfect as certain segments of the trajectories can remain ambiguous. Nevertheless, we are interested to decide on an RFID reader placement that reduces ambiguity of trajectories.

Specifically, we consider two metrics: *ambiguity* and *tracking error*. The ambiguity metric in an indicator of the extent to which the occupants' locations have been incorrectly inferred. The ambiguity for each tracked individual is the fraction of time that the invividual's ID belongs to ambiguous trajectory segments (i.e., segments containing two more more candidate IDs including the ID of the particular individual and we call the sets of two or more candidate IDs the *ambiguity sets*). The ambiguity metric is the average over all individual ambiguities.

The tracking error metric is influenced by the ambiguity metric. It consists of a lower and an upper bound for the localization error. For each individual, there are ambiguous segments that this individual has likely traversed (based on its participation on trajectories that are ambiguous and include the particular ID). Of those, the two paths with most and least Euclidean distance from the person's actual (ground truth) path are considered to describe an upper bound (Tracking Error Upper-bound, TEU) and a lower bound (Tracking Error Lower-bound, TEL) error. The TEU and TEL errors are calculated on a per-individual basis and averaged across time.

3 RFID Reader Placement

We introduce a *skeleton* of the heatmap (Figure 2(c)), which is the result of a three step process: a) thresholding the heatmap, b) iterative thinning the binary values created by thresholding, and (c) removal of short branches and cycles. The end result is an undirected n-ary tree, $T = (V_t, E_t)$. The vertices of T are also called the branching points and the edges are the actual paths (*i.e.,* they correspond to a sequence of real spatial points, and are not just logical links between vertices). Figure 2 illustrates the conversion of the heatmap to a n-ary tree. Having computed T, all its vertices become potential coverage points for RFID readers as they indicate busy gateways towards different locations, but only some of them will be chosen for reader placement.

We consider the branching points of T as candidate locations for reader placement. The placement heuristic is based on a score function $\mathcal{F}$ that captures (via their product, $\mathcal{F}_v = H_v \times D_v$) two factors. First, the *heat factor* H_v is expressed as the sum of heatmap values of the locations covered by placing a reader at a particular vertex, thus favoring coverage of heavier traffic areas. Second, the *distance factor* D_v expresses the path distance from already placed readers, thus biasing in favor of readers further apart, as placing them near each other results in coverage overlaps producing no noticable advantage.

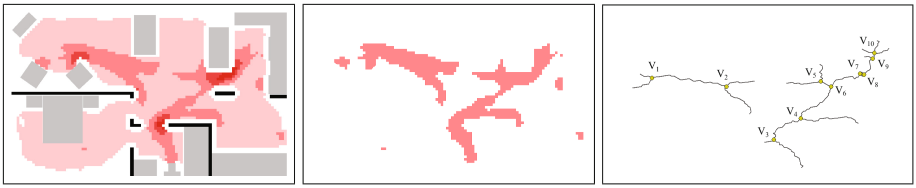

(a) (b) (c)

Fig. 2. The skeletonization process: (a) heatmap, (b) thresholding, and, (c) skeleton

4 Simulation Results

We produced a number of test cases simulating different numbers of occupants moving in the smart home with the floorplan in Figure 2(a) which corresponds to an actual space with an area of $10.60 \times 6.29\,m^2$. In total, we considered 14 different activity patterns, and we simulated the movement of the occupants assuming that each occupant randomly chooses three activities out of the possible 14. On average, each activity takes the occupant 15 seconds of simulated time. In order to emphasize the ability to distinguish trajectories, no pause times were simulated, that is, an individual would be simulated starting a new activity as soon as the previous one had ended. We conducted simulations for $k = 2$, $k = 3$ and $k = 5$ occupants. First, the PIR motion sensor placement by [4] was followed, producing an optimal solution consisting of 11 PIR sensors.

Table 1. Comparing different placement methods for five RFID readers

k=2	Tree-Based	Manual	Random
a_2^* (%)	3.097	3.458	3.864
TEU(m)	0.617	0.621	0.632
TEL(m)	0.613	0.613	0.613

k=3	Tree-Based	Manual	Random
a_2^* (%)	6.075	5.669	7.930
a_3^* (%)	5.073	6.062	6.085
TEU (m)	1.251	1.313	1.389
TEL (m)	0.748	0.755	0.776

k=5	Tree-Based	Manual	Random
a_2^* (%)	0.844	1.746	2.096
a_3^* (%)	3.471	4.389	5.339
a_4^* (%)	7.816	8.987	9.915
a_5^* (%)	11.059	12.190	12.492
TEU (m)	3.768	4.020	4.067
TEL (m)	1.016	1.051	1.056

We compare the placement determined from our heuristic against an "expert" manual placement and a random one. The manual placement aims to cover as much as possible of the entire space but places at least one reader in every room. For the randomized placement, points with a non-zero utility score in the heatmap were randomly chosen. For each RFID reader point, the closest point on a wall was determined and assumed to be the reader's mounting point (the same process was used for mounting points of readers in the tree-based method). For randomized placement we report the average of 10 randomized placements.

The results presented in Table 1 are averages over 100 runs for each reader placement. To properly appreciate the results, we note that even in the ground truth there is a small probability that individual trajectories collide (same location at the same time). The ground truth for the case $k = 2$ shows 2.986% collision between the two occupants. For $k = 3$, two people trajectories are colliding 10.547% of the time, and three colliding 0.301% of the time. The percentages for $k = 5$ are 14.627%, 1.209%, 0.137% and 0.001%, for collisions involving 2, 3, 4 and 5 trajectories, respectively.

In each sub-table of Table 1, two rows present the TEU and TEL. An additional $k-1$ rows show the average ambiguity (a_2^* to a_k^*) for an individual's trajectory, as a function of the cardinality (from 2 to k) of the *ambiguity set* (defined earlier). In all cases the tree-based method surpasses the other two methods: it exhibits smaller tracking error and smaller ambiguity. However, while the merits of the tree-based method become increasingly apparent as k increases, the absolute tracking error deteriorates quickly from approximately 60cm for two occupants, to between 75cm and 120cm for three occupants, to 3.5m in the case of five occupants (making it virtually unusable). Note that as k increases the value of a_k^* is inflated, because disambiguating ("teasing apart") the trajectories of each individual in a large group is less likely to accomplish, than from a small group. Additionally, for large populations of occupants, given a budget limitation to five RFID readers only, it is increasingly likely that, after approaching (and read by) an RFID reader, an individual's trajectory will collide with some other trajectory before it gets a chance to be read again by an RFID reader, hence ambiguous trajectories abound in large occupant population scenarios.

5 Conclusion

We introduced and evaluated through simulations an indoor multiple-person tracking system based on a combination of PIR and RFID technologies. The ambiguity of PIR-based tracking for mutiple occupants is mitigated by the use of information from RFID readers. Due to the relatively costly and challenging deployment of multiple RFID readers, we devised a RFID reader placement heuristic aiming to produce good tracking results. A reading from an RFID tag worn by an occupant provides unambiguous location information, subsequently used to disambiguate segments of occupants' trajectories that were, up to that point, unknown to which occupant they corresponded. The proposed heuristic favorably compares against random as well as naive manual placements that attempt to cover the whole space.

Acknowledgments. This work has been partially funded by IBM, the Natural Sciences and Engineering Research Council of Canada (NSERC), Alberta Innovates - Technology Futures (AITF), and Alberta Health Services (AHS).

References

1. House, S., Connell, S., Milligan, I., Austin, D., Hayes, T.L., Chiang, P.: Indoor localization using pedestrian dead reckoning updated with rfid-based fiducials. In: 2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC, pp. 7598–7601. IEEE (2011)
2. Milan, A., Schindler, K., Roth, S.: Detection-and trajectory-level exclusion in multiple object tracking. In: 2013 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3682–3689. IEEE (2013)
3. Ruiz, A.R.J., Granja, F.S., Prieto Honorato, J.C., Rosas, J.I.G.: Accurate pedestrian indoor navigation by tightly coupling foot-mounted imu and rfid measurements. IEEE Transactions on Instrumentation and Measurement **61**(1), 178–189 (2012)
4. Vlasenko, I., Nikolaidis, I., Stroulia, E.: The smart-condo: Optimizing sensor placement for indoor localization. IEEE Transactions on Systems, Man, and Cybernetics: Systems **45**(3), 436–453 (2015)
5. Wang, B.: Coverage problems in sensor networks: A survey. ACM Computing Surveys (CSUR) **43**(4), 32 (2011)
6. Yu, C.R., Wu, C.L., Lu, C.H., Fu, L.C.: Human localization via multi-cameras and floor sensors in smart home. In: IEEE International Conference on Systems, Man and Cybernetics, 2006. SMC 2006, vol. 5, pp. 3822–3827. IEEE (2006)