33rd Annual International Conference of the IEEE EMBS
Boston, Massachusetts USA, August 30 - September 3, 2011

Scalable Patients Tracking Framework for Mass Casualty Incidents

Xunyi Yu, Aura Ganz, *IEEE Fellow*
Electrical and Computer Engineering Department,
University of Massachusetts, Amherst

Abstract—We introduce a system that tracks patients in a Mass Casualty Incident (MCI) using active RFID triage tags and mobile anchor points (DM-tracks) carried by the paramedics. The system does not involve any fixed deployment of the localization devices while maintaining a low cost triage tag. The localization accuracy is comparable to GPS systems without incurring the cost of providing a GPS based device to every patient in the disaster scene.

I. Introduction

Tracking paramedics and patients is a fundamental functionality of various MCI response information systems such as DIORAMA[1], WISSARD[2], CodeBlue[3], and AID-N[4]. Accurate and up-to-date paramedics and patient location is the key for the incident commander to gain situation awareness, allocate resources and coordinate evacuation. In this paper, we focus on tracking patients without the need of fixed infrastructure deployment.

In existing systems, [1,2,3] RF receivers are deployed throughout the incident area as anchors. The density of deployment and the amount of time needed to calibrate the system may vary for different RF localization technologies. However, for rescue efforts in a mass casualty incident, every second wasted on system deployment can mean a difference of life and death. GPS devices can be used to locate paramedics and patients which does not require the deployment of specialized RF localization devices. However, GPS tags are power hungry and costly. In our system each patient is tagged with an active RFID tag denoted as D-tag.

In this paper, we introduce a tracking framework particularly customized for MCI localization by using moving paramedics as mobile anchor points to locate patients. Each paramedic carries a DM-Track device that includes an active RFID reader and a Smartphone (e.g. Android based). The location of the paramedics carrying DM-track devices is determined by the GPS embedded in the Smartphone. The DM-tracks constantly receive active RFID beacons sent by the patients' D-tags. This architecture eliminates the need to deploy numerous localization devices, while keeping the size and cost of patient triage tags down.

Localization using moving anchor points has been investigated in wireless sensor networks. In large wireless sensor networks, it is prohibitively expensive to equip each mobile node with a GPS device. Thus a few nodes equipped with GPS devices are used as mobile anchor points to locate other nodes [5, 6, 7]. However, in wireless sensor networks, it is generally assumed that the nodes to be localized are static once deployed. Therefore, most of the papers developed localization algorithms that are applied repetitively to update node location, with no or limited mobility model for tracking. In contrast, the patients in a mass casualty incident may move to new locations from time to time, and need to be tracked continuously.

The paper is organized as follows. In the next section we introduce the system architecture. The recursive tracking algorithm is presented in Section III. Section IV describes the experiments and Section V concludes the paper.

II. System Architecture

In the triage phase the paramedics determine the triage status of each patient (black, red, yellow and green) and tag them with the appropriate color paper tag along with a D-tag (active RFID tag). The paramedic carries a DM-Track which includes a Smartphone (e.g. Android based platform) as well as active RFID reader.

The active RFID reader collects the D-tags' ID and severity of injury as well as the received signal strength of the RFID beacons. The active RFID reader is interconnected to the Android Smartphone using the Bluetooth interface. The Android Smartphone interconnects to the Internet through the cellular or WiFi network.

The DM-Tracks transmit to the server the following information: 1) the beacon information along with the received beacon signal strength for each D-tag in its range, and 2) the GPS readings obtained from the Android phone.

The recursive tracking algorithm, which is described in the next section, will continuously update the patients location by fusing the D-tags beacon measurements from all available paramedics along with their GPS information.

III. Recusive Tracking Algorithm

Recursive tracking of a patient using RFID measurements from the DM-track devices carried by the paramedics in the vicinity of the patient has the following three challenges:

Challenge 1) At a particular time, the current estimate of

This project was supported in part by the following grants: Grant numbers DUE-0736888 and DUE- 1003743 from the National Science Foundation and Grant Numbers R21LM008942-01A2 and 3R21LM008942-02S2 (ARRA) from the National Library of Medicine/National Institutes of Health. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Science Foundation,National Library of Medicine or the National Institutes of Health.

the patient location cannot be approximated by a parametric model very well. For example, with one or two paramedics nearby, the distribution of the patient location is of a ring shape distribution and a two modal distribution respectively.

Challenge 2) The measurement model is non-linear even for the free-space attenuation model, and becomes more complicated when the inaccuracy of the paramedic location given by GPS, and likelihood of an obstacle need to be taken into consideration.

Challenge 3) The motion of the patients is too arbitrary to be accurately modeled and poor and intermittent RSSI measurements make it very challenging to maintain the states of the motion model to improve location estimates.

In the next paragraphs we describe how we address each one of the challenges introduced above.

Challenge 1: we sample the possible locations of a patient to form the state space, and recursively update the probability of the patient being at these locations using the Bayesian paradigm. Let $\mathbf{x}_k = [x_{k1}, x_{k2}, ..., x_{kN_k}]$ be N_k sampled locations of a patient at time k. Given new RFID measurements $\mathbf{z}_k = [z_{k1}, z_{k2}, ..., z_{kN_D}]$, which represent the received signal strength of the radio beacons received at N_D DM-tracks at time k, we recursively calculate the probability of a patient at location x_{ki} at time k given all previous $\mathbf{z}_{1:k-1}$ and current RFID measurements $\mathbf{z}_k$ by

$$p(x_{ki} \mid \mathbf{z}_{1:k}) = \frac{p(\mathbf{z}_k \mid x_{ki})\, p(x_{ki} \mid \mathbf{z}_{1:k-1})}{p(\mathbf{z}_k \mid \mathbf{z}_{1:k-1})} \quad , \tag{1}$$

where $p(\mathbf{z}_k \mid \mathbf{z}_{1:k-1})$ is a normalizing factor, $p(\mathbf{z}_k \mid x_{ki})$ is the measurement probability of getting measurement $\mathbf{z}_k$ assuming that the patient is sending radio beacons from x_{ki}. The measurement probability can be determined by a free space RF signal attenuation model.

Challenge 2: the measurement probability can also be calculated using the scene and location specific radio propagation parameters, accounting for signal attenuation due to signal obstruction by buildings, etc. $p(x_{ki} \mid \mathbf{z}_{1:k-1})$ is the prediction probability that the patient is at location x_{ki} given all previous RFID measurements $\mathbf{z}_{1:k-1}$, as

$$p(x_{ki} \mid \mathbf{z}_{1:k-1}) = \sum_{j=1}^{N_{k-1}} p(x_{ki} \mid x_{(k-1)j})\, p(x_{(k-1)j} \mid \mathbf{z}_{1:k-1}), \tag{2}$$

where $p(x_{ki} \mid x_{(k-1)j})$ is the probability a patient moves from the *j*th sampled location at time $k-1$ to the *i*th at time k.

Challenge 3: the motion model of a patient can be determined by a behavior model of patients accounting for crowd movement. For simplicity, we only consider the case where patients' movements are independent of each other.

The motion model has two modes. Normally, we only dilate the previous patient location distribution in all directions to model arbitrary patient movement, which can be done by convoluting the distribution with a Gaussian kernel if location sampling is uniform. The span and variance of the Gaussian kernel is determined by the movement probability and average movement speed of the patients, which can be customized for different triage levels. We use this simple movement model since it is difficult to accurately estimate patient movement direction and speed at a particular moment, given none or limited number of noisy measurements. This movement model is mainly used to reflect the increased uncertainty of patient location distribution over time in case no new measurements are available.

We also use a second mode for the actual movement of patients. There may be an extended period of time when the patient move a significant distance to a new location, yet there were no paramedics near a patient to collect new measurement, and update the patient location distribution accordingly. As the dilating parameter for the normal case is chosen conservatively to account for those static patients, the patient location distribution will not dilate fast enough, and the patient's new location will have a significantly lower prediction probability $p(x_{ki} \mid \mathbf{z}_{1:k-1})$ than his previous location in the next Bayesian update. Thus it will take a very long time before the algorithm can converge to the patient's new location. To detect such movements, we calculate the ratio between the measurement probability at current point location estimate and the highest measurement probability of all locations.

$$R = \frac{p(\mathbf{z}_k \mid \bar{x})}{\max_i (p(\mathbf{z}_k \mid x_{ki}))} \tag{3}$$

where $\bar{x}$ is a point estimate of patient location, which can be the average or mode of the patient location distribution $p(x_{(k-1)i} \mid \mathbf{z}_{1:k-1})$ in the previous step. When R drops below a threshold T_R it indicates that the previous location distribution may be out-dated. We add a uniform component to the distribution and renormalize it.

$$p(x_{(k-1)i} \mid \mathbf{z}_{1:k-1}) = p(x_{(k-1)i} \mid \mathbf{z}_{1:k-1}) + C \tag{4}$$

Note that occasional outlier measurements can also cause a drop in R. Lower threshold T_R and higher C enable the algorithm to converge faster in case of unobserved patient movement, while reducing the effective number of measurements that can be used to improve localization accuracy for static patients.

We can use the variance of the patient location distribution as a factor to evaluate the accuracy and timeliness of the location estimate. If there were no paramedics nearby a patient to collect up-to-date RFID beacons, we would not be able to tell if the patient have moved to a new location or not. The motion model of the patient dilates the patient location distribution, increasing its variance. The patient location distribution variance decreases after paramedics move close to the patient, and new measurements are taken. Due to log-distance attenuation of the RF signal, the closer a paramedic is located to a patient, the more concentrated the measurement probability. Thus multiple diverse measurements taken close to a patient enable us to locate the patient within a small area, with low variance of the patient location distribution.

For an update at time k, the complexity of the algorithm is $O(N_k N_t)$, where N_k is the number of locations we sampled, and N_t is the number of targets. We can choose to only sample locations nearby a patient, to limit the number of states N_k, and make it independent of the size of the incident area. N_k can be limited to hundreds with 10 feet sampling granularity. Since the number of states is low, simple uniform sampling can be used, without resorting to more complex importance sampling and particle filtering. The constant factor in the O() expression depends on specific model we choose to model measurement likelihood, target motion, resample the state space, and specific implementations. A discrete sampling method significantly reduces the complexity of these steps as opposed to a parametric method. Our algorithm scales very well to large number of targets and large incident sites. The algorithm can run in real-time for a scenario with hundreds of targets on a typical PC with an update interval of one second or larger.

IV. EXPERIMENT

We carried out field experiments and simulations to evaluate the static and dynamic performance of the recursive tracking algorithm. In the static evaluation, we assume that patients stay still after the beginning of the test, and the patient location distribution is initialized to uniform distribution. In the dynamic evaluation, the patient will be moving and then stop in continuous cycles. The system needs to detect the movement of the patients, and apply corresponding movement models such that the location estimate can converge to the new location of the patient as soon as possible.

Field experiment: The field experiment was carried out in a 100ftX100ft area, with 20 patients who have been triaged with D-tags. Six patients have moved to a new location after they were initially triaged. In this experiment, the paramedic will first evacuate 14 static patients to the staging area located at the at the origin (0,0). After the 14 static patients are evacuated, we will evaluate how accurate are the location estimates for the 6 patients that moved. In Figure 1, we plotted the location of the static patients, and the actual and estimated location of the patients that moved. All 6 patients are localized within 20 feet of their actual location. Note that localization accuracy of patients can be even better than localization accuracy of GPS in some cases, as GPS localization error averages out when we use multiple RFID measurements taken at different paramedic locations.

Simulation: We use simulations to further evaluate the localization accuracy of the recursive tracking algorithm under various movement patterns of the paramedics. Note that our patient tracking algorithm tracks each patient independently. Thus, in the following simulation, we only use one patient to evaluate the localization accuracy, which also makes it easier to define paramedic's movement patterns.

<u>Simulation 1: static patient.</u>

We use this simulation to evaluate the location accuracy of the static patient under different paramedics' movement patterns. The paramedics' movements are in straight lines of fixed length of 50 feet around a patient. The angles of the movement lines are random. The measurements taken from these lines are used to locate the patients. The GPS offset and RFID measurements used in the simulation are synthesized from field measurements. To obtain these field measurements, we place on the patient located at the origin a triage tag. Then, a tester wearing the mobile DM-device walks in a grid pattern around this patient. In this way, we can collect the GPS location estimate offset at different locations around the patient, and the measured RFID signal strength at these locations. In the simulation, we synthesize different line movement patterns, and use GPS location estimate and RFID signal strength from the field measurements. This hybrid approach can help us capture the temporal and spatial correlation of GPS localization error and RFID noise. We assume that the beacons are sent at 2 seconds intervals. In Figure 2, we plot the average localization accuracy in terms of the shortest distance from the lines to the patient, and the number of lines where RF measurements are taken.

We observe that the localization accuracy improves as the number of lines (either obtained by a single paramedic in multiple passes or by multiple paramedics) increases and as the distance between the paramedic path and the patient decreases.

The algorithm obtains an accuracy of less than 20 feet in a number of cases: if we have a few readings, we need to obtain the readings close to the patient (less than 15 feet away). If we have 4 or more lines of readings the paramedic(s) can be as far as 50 feet away from the patient at all times.

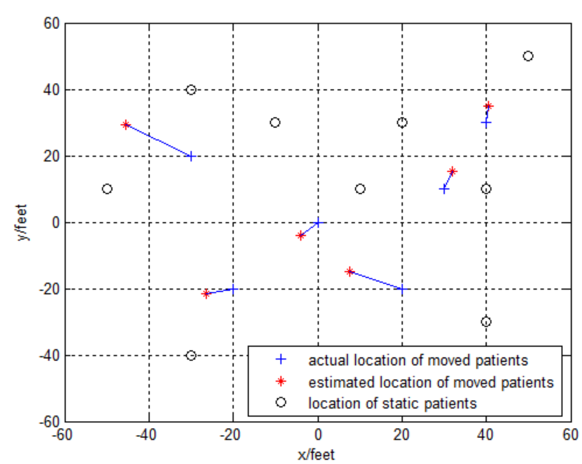

Figure 1. Location estimates of patients that moved

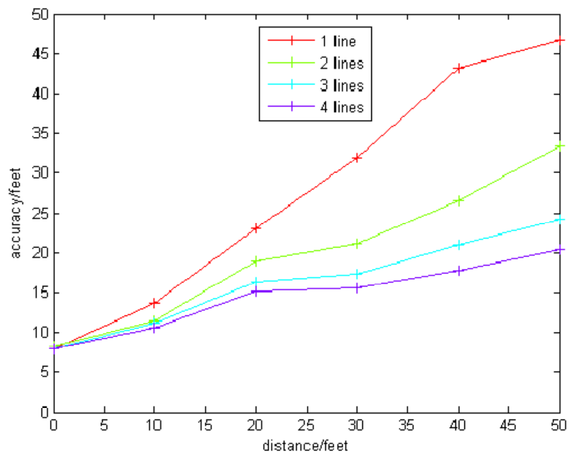

Figure 2. Localization accuracy as a function of average distance between paramedic paths and the patient

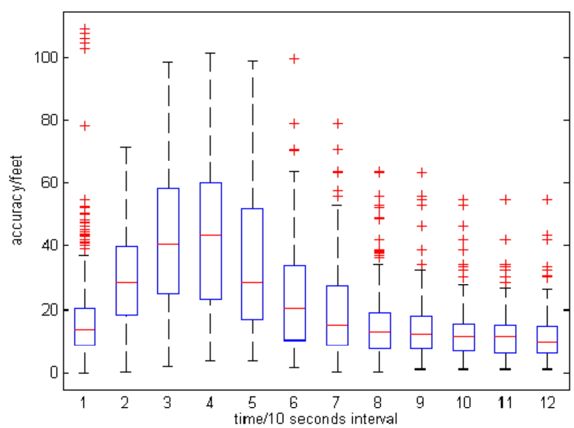

Figure 3: Box plot of localization accuracy

Simulation 2: moving patients.

We use the following simulation to evaluate the dynamic performance of the algorithm and its convergence speed. A patient in a 100ftX100ft area moves in the first 30 seconds of a cycle, and then stops for 90 seconds. There are two paramedics in the area moving randomly. Their DM-tracks collect RFID readings from the patients' D-tags used to estimate the location of the patient. The simulation runs for 100 cycles continuously. In this simulation, we assume that the GPS location estimate is uniformly distributed in a 40 by 40 feet box centered on the actual paramedic location. For simplicity, we model the temporal and spatial correlation of GPS localization error and RFID noise. To reflect the correlation in the simulation, we under-sample RFID measurements from 2 seconds interval to 6 seconds interval. Figure 3 shows the box plot of the localization errors for the 10 seconds interval of a cycle. As the patient starts moving, the system will experience delays in updating his new location due the availability of paramedics nearby, and localization errors peaks in the 20-30 second interval. Then, as the patient stops moving, more measurements are collected, and the system refines patient location estimation gradually. Within 30 seconds after the patient stopped moving, the 75 percentile localization error drops below 20 feet. As time passes, the localization accuracy improves, and the number of location estimates with errors above 30 feet drops constantly. These large errors are primarily due to the lack of enough measurements taken close to the patient.

In summary, the experiment and simulation results demonstrate that in some scenarios mobile readers can replace densely deployed fixed readers with equivalent performance. Moreover, the localization accuracy depends on the availability and closeness of paramedics to nearby patients.

V. Conclusion

In this paper we introduce a patient localization system using a mobile infrastructure (i.e., the DM-tracks) carried by the paramedics. The proposed system takes advantage of the fact that the paramedics at a disaster site move close to the patients that need to be triaged and evacuated. This proximity will result in relatively high localization accuracy (comparable to GPS localization) given the fact that there is no fixed infrastructure. No fixed infrastructure requirement translates in reduced cost as well as significantly less deployment time. This technology is the first of its kind to enable outdoor localization of patients tagged with only an active RFID tag and requires no fixed infrastructure.

References

[1] A. Ganz, et al., "Dynamic Information Collection and Resource Tracking Architecture", *IEEE Engineering in Medicine and Biology Society*, September 2010, pp. 386-389.

[2] WIISARD: Wireless internet information system for medical response in disasters. https://wiisard.org

[3] CodeBlue: Wireless sensor networks for medical care. http://www.eecs.harvard.edu/~mdw/proj/codeblue/

[4] AID-N: Advanced health and disaster aid network. Available: https://secwww.jhuapl.edu/aidn/

[5] C-H. Ho Ou, "A Localization Scheme for Wireless Sensor Networks Using Mobile Anchors With Directional Antennas", *Sensors Journal, IEEE*, 2011, vol. 11, pp.1607-1616

[6] K.F. Ssu, et al. "Localization with mobile anchor points in wireless sensor networks", *Vehicular Technology, IEEE Transactions on*, 2005, vol. 54, pp.1187-1197

[7] H. Chen, "Mobile element assisted cooperative localization for wireless sensor networks with obstacles", *Wireless Communications, IEEE Transactions on*, 2010, vol. 9, pp.953-963