Enabling Context Aware Clinical Applications through Ultra-Wideband Localization

Fabrice Labeau[a], Antoine Buntcha Tchana[b] and Tho Le-Ngoc[a]
[a] *McGill University, Montreal, QC, Canada,* [b] *Faculté Polytechnique de Mons, Mons, Belgium*

Abstract *- This paper presents a preliminary study of the applicability of commercial off-the-shelf localization systems based on Ultra-Wideband (UWB) technology, for the purpose of enabling location-aware applications in next-generation wireless clinical systems. Envisioned applications include proximity-based patient-clinician association and context aware services.*

Keywords: *Ultra-Wideband (UWB); Indoor Localization; RFID*

1. INTRODUCTION

Interest in RFID or positioning systems for clinical applications has been growing recently, mainly as a tool for asset tracking and location [1]. Specialized medical equipment such as infusion pumps have been successfully tracked using such positioning systems, thus reducing the time wasted by clinicians looking for them throughout hospital facilities. Such demonstrated efficiency gains through information technology are now sparking interest in more advanced studies based on location and positioning technology.

Extending the location and positioning technology from only equipment to clinicians and patients opens the door to a new set of possible applications, whose potential for time and cost savings are appealing. In particular, technologies that would automatically report location information about both assets and clinical actors would enable a central server to acquire a key component of the *context* of current clinical interventions and as such would be a basic building block towards an adaptive, context-aware information and communications system. As an example, one could imagine that a *precision* location information system would enable pairing between patients and clinicians by proximity, so as to easily allow access to the patient's electronic health record on the clinician's wireless terminal. Similarly, pairing by proximity of a clinician and his wireless or wired information terminal would allow for faster log on and log off, as well as easy session handover between terminals.

Beyond this sort of online, real-time usage of location information for the development of advanced IT features, it is also important to realize that once deployed, an advanced location and positioning system can be used for offline studies of clinical workflow patterns, and for the non intrusive assessment of bottlenecks in clinical processes. As such, it constitutes an easy way to collect invaluable information about the existing processes within a clinical environment. The precision and extent of data collected will also enable advanced statistical modeling of patient flow, built-in delays, etc. In this context, automatic location and positioning data collection can also be seen as an advanced tool for assessment of clinical workflow efficiencies before and after some intervention.

In this research, we aim to deploy a positioning system to carry out mainly workflow related studies, and also carry out preliminary studies of how real-time online usage of location information could be put in place. The key objectives will be here to deploy, improve and use a precision indoor location system in a clinical environment, so as to allow for advanced studies of workflows, as well as to enable a few targeted basic location-based services for clinicians.

In order to enable advanced location-assisted information technologies into a clinical facility, it is necessary to reach a level of reliability and precision on the location estimate that is commensurate with the sort of applications envisioned. For example, , e.g., a localizing precision of 30 cm to 50 cm is necessary to be able to associate objects to people within a "personal area".

In this work, we consider the use of UWB-based localization in typical indoor environment. Our experimental results indicate that, in non-line-of-sight (NLOS) propagation conditions, UWB-based localization devices suffer from degraded performance [2]. Multipath reflections introduce additional errors, both systematic offsets and increased variance in ranging [3][4]. Our preliminary studies also show that low-cost appropriate post-processing and signal processing algorithms can be applied to enhance such a degraded performance in order to reach the target required by our application.

The remaining of the paper is organized as follows. Section 2 briefly describes the UWB-based localization system under consideration. Section 3 presents the measured results. Section 4 describes the proposed post-processing methodologies and their performance. Conclusions are provided in Section 5.

2. OFF-THE-SHELF UWB SYSTEM

UWB systems are characterized by a very large bandwidth, which allows spreading the signal energy over a large frequency range. It is thus expected that such a system will have a minor impact in terms of interference on other services and, in the case of deployment in a clinical setting, will have a very low probability of creating electromagnetic interference with medical equipment. UWB based on very short pulses also has a very good potential for ranging and localization applications, due to the potentially very time resolution provided by the short pulses.

The available off-the-shelf UWB localization system under consideration computes location based on the TDOA (time difference of arrival) at different receivers of short time-domain pulses emitted by UWB tags, The UWB tags are battery-operated [5], and emit a pulse at regular, programmable intervals. The system is meant to track the position of these tags. Multiple receivers are setup in the space to be monitored; they are daisy-chained through a wired link, and compare the time of arrival of pulses emitted by the tags. Based on the difference between these times of arrival, a multi-lateration algorithm computes the 3-dimensional position of each tag. The calibration of positions is done thanks to the presence of a so-called *reference tag*, whose position must be known to the receivers (typically through some initial manual setup); positions are then computed relative to that reference tag.

In the experiments detailed below, we have deployed up to 5 receivers in a dedicated laboratory space, with a varying number of tags in different configurations. The aim of these experiments was to determine how accurate the locations reported by the system could be; how reproducible the measurements were; how much post-processing was needed to make the system useful; and how much the room configuration and tag visibility to receivers influenced the performance of the system.

3. RAW MEASUREMENTS

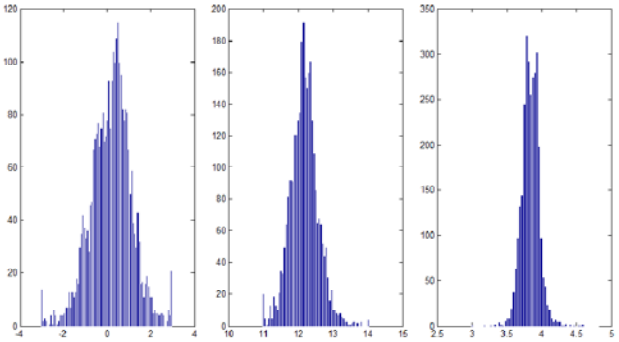

(a) x direction **(b) y direction** **(c) z direction**

Figure 1- Histograms of errors (meters) in UWB localization (static tags), in 3 main directions (x,y,z).

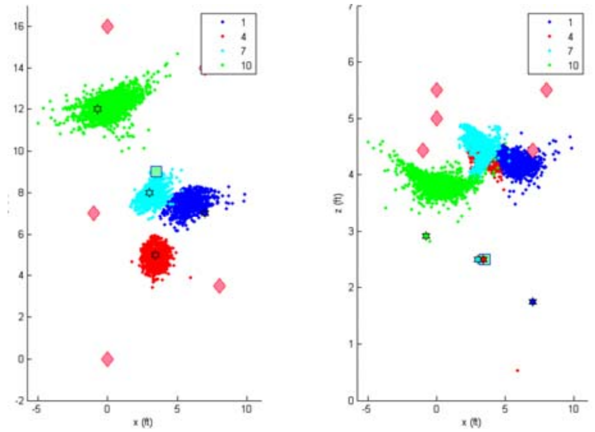

(a) $p_x - p_y$ plane **(b) $p_x - p_z$ plane**

Figure 2: Initial location test for the PAL system, without any filtering. (*Diamond*: receiver position; *Square*: *reference* tag position; *Star*: *actual* tag location; *Dot*: *measured* location over a period of 1h23min).

Initial measurements were taken in different configurations. Figures 1 and 2 show illustrative results. Figure 1 shows the distribution of positioning measurements for a given tag, in the 3 directions of an orthogonal system of axes. Figure 2 illustrates the readings obtained from the system for four characteristic tags (tags labeled 1, 4,7 and 10). In this figure, each dot represents a position computed from the system; positions are only represented for pulses that have been detected by at least 4 receivers. It is clear from the test results that (*i*) even when the tags are not in motion, the measurements are very noisy; (*ii*) the location in the vertical direction (figure 2(b)) is much less accurate than in the horizontal plane, and seems to present some bias. These experiments were carried out with tags emitting a signal every second. The variation of location over time shows that, as is, the information from the system is far from suitable for applications where sub-meter precisions are necessary. For instance, if a clinician and a terminal are both equipped with a location tag, the fluctuation on the reported location of these two tags would not allow for a stable association of clinician and terminal, and would result in a series of associations/dissociations that would hinder the clinical workflow being executed. One additional comment is in order: the quality of the location results are very strongly affected by the visibility of tags by receivers. When receivers lose a direct line-of-sight to the tags, it affects dramatically the precision of their location reading.

Based on these raw measurements, it seems clear that some form of smoothing or filtering is needed. In the next section, we will explore the Kalman filter as a possible post-processing tool.

4. KALMAN FILTERING

Kalman filtering has been successfully used in many application areas to smooth or predict data. It provides a smoothed or predicted version of the state of a system that is optimal in the mean-squared sense, under some assumptions.

The application of Kalman filtering to the case at hand follows from a modeling of the localization problem in terms of linear stochastic equations, as follows.

First, a time-varying vector x_k representing the state of the system is defined. Typically, this state will contain the quantities to be estimated, i.e. in our case at a minimum the three-dimensional spatial coordinates of the UWB tag to be located. In the case where tags are moving, it might also be convenient to include velocities in this state vector.

The evolution of the state vector over time (represented here by the subscript) is described by a linear state equation:

The Kalman filter recursively gives an estimate of the current state x_k based on observations up to time k through the following well-known equations:

$$x_{k+1} = A_k x_k + w_k, \tag{1}$$

where A_k is the state transition matrix, that might be time-varying depending on the modeling assumptions, and w_k is the so-called process noise, which is assumed to be zero-mean, white with covariance matrix Q. On the other hand, the observed quantities are assumed to be a noisy version of some linear combination of the state:

$$y_k = B_k x_k + \xi_k, \tag{2}$$

where ξ_k is a measurement noise assumed to be zero mean, white, with covariance matrix R.

$$\mathrm{x}_\mathrm{k}^- = \mathrm{A}_\mathrm{k}\mathrm{x}_{\mathrm{k}-1}^+ \tag{3}$$

$$\mathrm{P}_\mathrm{k}^- = \mathrm{A}_\mathrm{k}\mathrm{P}_{\mathrm{k}-1}^+\mathrm{A}_\mathrm{k}^\mathrm{T} + \mathrm{Q} \tag{4}$$

$$K_k = P_k^- B_k^T (B_k P_k^- B_k^T + R)^{-1} \tag{5}$$

$$x_k^+ = x_k^- + K_k(y_k - B_k x_k^-) \tag{6}$$

$$P_k^+ = (I - K_k B_k) P_k^- \tag{7}$$

Equations (3)-(7) form the basis of the Kalman update. K_k in equation (5) is the Kalman gain.

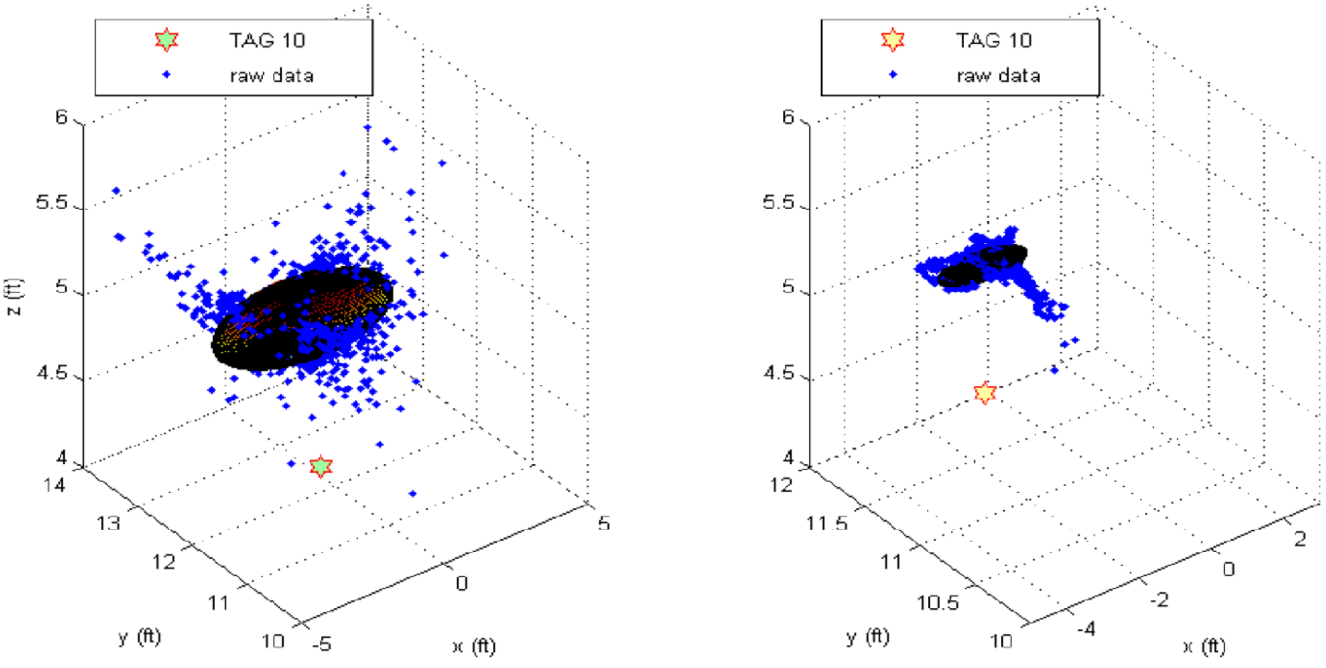

(a) raw measurement

(b) Kalman filtering

Figure 3: 3-D representation of the localization of a given tag (tag 10), without (a) and with (b) Kalman filtering. (Note: *Confidence ellipsoids* represent ±3σ intervals aligned with eigen-directions under assumption of Gaussian location data).

If one wants to only localize tags that are not in motion, then the state to be estimated is only a 3-dimensional vector of positions in space, and the corresponding state equation is very simple ($A_k$becomes an identity matrix).

In practice, for moving tags, it is better to augment the state with velocities, so that the corresponding state vector x_k is made up of six components, namely the positions p_i along the 3 Cartesian coordinates and the velocities v_i along these coordinates:

$$\mathrm{x_k} = \left(\mathrm{p_{x,k}} \;\; v_{x,k} \;\; \mathrm{p_{y,k}} \;\; v_{y,k} \;\; \mathrm{p_{z,k}} \;\; \mathrm{v_{z,k}}\right)^{\mathrm{T}}.$$

In this case, the state transition matrix A_k becomes:

$$A_k = \begin{pmatrix} 1 & \Delta t & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & \Delta t & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & \Delta t \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \qquad (8)$$

where Δt represents the time interval between measurements. This matrix can only faithfully represent a linear movement at constant speed, but we will see in later sections how the filter can be made robust to changes in speed and direction. It is clear that more complex state transition matrices or state representations can be used to more faithfully represent non uniform motion.

Figure 3 shows an example of results obtained with Kalman filtering, in a scenario similar to the one in Figure 2 (no motion of the tags). Measurements for one tag are shown, along with ellipsoids depicting the extent of the variance of measurements before and after Kalman filtering. The ellipsoids are obtained by calculating the 3 × 3 covariance matrix of the measurements, and taking an eigenvalue decomposition of this matrix. Eigenvectors of this matrix represent eigendirections of the distribution (assuming it is Gaussian), and its eigenvalues represent variances along these eigendirections. The ellipsoids are drawn along these eigendirections, with main axes lengths of three times the standard deviation along these axes. This figure clearly shows the decrease in variance of the measurements for the fixed tag. On the other hand, it is clear also that the bias in the vertical direction is not dealt with properly.

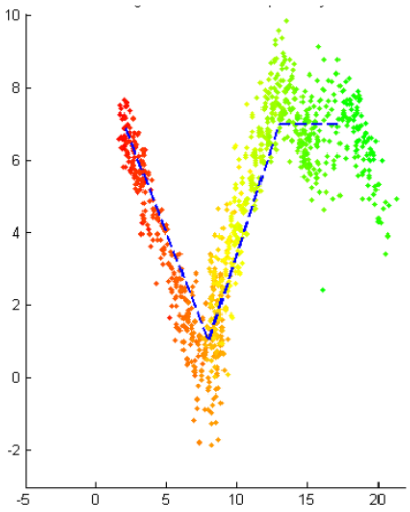

Figure 4: Raw measurements with a mobile tag (p_y*vs.* p_x, in feet; *dashed line*: actual trajectory; *red*: first measurement, *green:* last measurement).

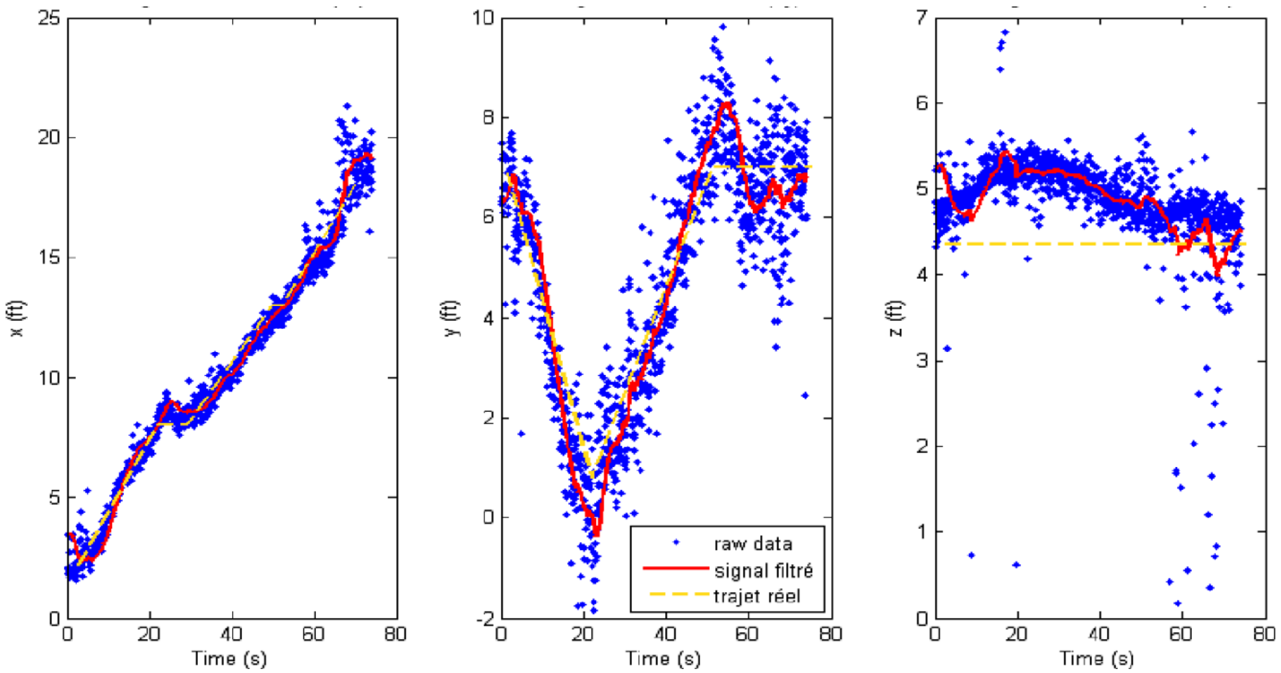

Figure 5: Kalman filtering in the case of an irregular motion. (*Blue dots*: actual measurement; *red continuous line*: Kalman filter output; *yellow dashed line*: actual trajectory).

Figure 4 shows the raw measurements taken for the pursuit of a mobile tag, with non-constant speed and changes in direction. Note that this is not taken into account by the model of the state transition matrix in equation (8), which assumes constant velocity. Figure 5 represents the corresponding results with Kalman filtering. The results are reassuring for this type of tracking application, where clearly the Kalman filter can provide an adequate amount of smoothing.

For this dynamic context, the Kalman filter model can be made robust to changes with respect to the model, by adapting the matrices R and Q that model the amount of noise in the process (equation (1)) and in the observation (equation (2)). In general, setting a relatively large Q and a relatively small R implies that the measurements are more trusted than the model. Inversely, with a large R and a small Q, the model is more trusted than the measurements. The adjustment of these two quantities along time is shown to be the key for proper tracking (e.g., in phases of motion that are happening at constant velocity, a small Q can be used to provide extra measurement noise smoothing).

Devising appropriate adaptation algorithms for R and Q is the topic of current research.

5. CONCLUSION

This paper has explored the usage of an off-the-shelf UWB-based indoor localization system to provide location information for future context-aware clinical applications. Through experiments, we have shown that indoor NLOS propagation in presence of multipath fading degrades the ranging performance, and Kalman-based smoothing with additional automatic adaptation algorithms can be used as a candidate low-cost post-processing scheme to enhance such a degraded performance in order to achieve the required target.

6. REFERENCES

[1] T. Ostbye, D. Lobach, D. Cheeseborough, “Evaluation of an infrared/radiofrequency equipment-tracking system in a tertiary care hospital”, *J Med Syst*, Vol. 27, pp. 367-380, 2003.

[2] B. Denis, J. Keignart and N. Daniele, “Impact of NLOS propagation upon ranging precision in UWB systems”, *IEEE Conf. Ultra Wideband Systems and Technologies*, pp. 379-383, 2003.

[3] S. Venkatesh and R.M. Buehrer, “Power Control in UWB Position-Location Networks”, *IEEE Int. Conf. Comm.*, vol.9, pp.3953-3959, 2006.

[4] M. Tuchler, V. Schwarz, and A. Huber, “Location accuracy of an UWB localization system in a multi-path environment”, *Int. Conf. UWB*, pp. 414-419, 2005.

[5] Multispectral Solutions Incorporated, *PAL650 User’s Guide*, Revision number 2.01, June 2004

ACKNOWLEDGMENTS

The authors would like to acknowledge Nortel’s support.

CORRESPONDENCE ADDRESS:

Fabrice Labeau is with the Department of Electrical and Computer Engineering at McGill University, Montreal, Canada.

Address: 3480 University St. #753, Montréal, QC H3A2A7 Canada.