Location-Aware Access to Hospital Information and Services

Marcela D. Rodríguez, Jesus Favela, Edgar A. Martínez, and Miguel A. Muñoz

***Abstract*—Hospital workers are highly mobile; they are constantly changing location to perform their daily work, which includes visiting patients, locating resources, such as medical records, or consulting with other specialists. The information required by these specialists is highly dependent on their location. Access to a patient's laboratory results might be more relevant when the physician is near the patient's bed and not elsewhere. We describe a location-aware medical information system that was developed to provide access to resources such as patient's records or the location of a medical specialist, based on the user's location. The system is based on a handheld computer which includes a trained backpropagation neural-network used to estimate the user's location and a client to access information from the hospital information system that is relevant to the user's current location.**

***Index Terms*—Context-aware computing, hospital information systems (HISs), location estimation, mobile collaboration.**

I. Introduction

WORK AT hospital settings requires considerable mobility and coordination due to the complexity of the tasks performed, the intensity of the information exchange, and the fact that information and resources are distributed throughout the premises. A hospital's staff might be distributed in space (i.e., different location within the settings) or time (i.e., working different shifts) and their information needs are highly dependent on their location and other contextual conditions such as their role or time of the day.

Artifacts are used in hospitals to support the staff's coordination or as distributed repositories of information. Whiteboards hung on walls, for instance, help to communicate information regarding patients' conditions and locations, and to infer the location of nurses [2]. Medical records integrate patients' clinical data and constitute a main source of reference for their care. An important trend in medical informatics is the adoption of electronic patient record systems that facilitate access to clinical information and work toward preventing the loss or misplacement of information.

At the same time, physicians are increasingly using handheld computers in their professional practice. It was estimated that 26% of all physicians in the U.S. used a handheld in 2001, a number that was expected to grow to 50% for 2004 or 2005 [12]. In fact, several medical schools in the U.S. are requiring students to have a handheld computer. Such an instance is UCLA's David Geffen's School of Medicine that established this requirement "to enable 'point of contact' access to information resources; and to prepare students for practicing medicine in the 21st century" (www.medstudent.ucla.edu/pdareq/). This trend has generated interest in the development of medical applications for personal digital assistants (PDAs) and evaluating their use [10], [14].

To date, the most popular medical applications on handheld devices are the ones that provide access to reference material, such as drug information databases.

PDAs wirelessly connected to a hospital information system (HIS) can give physicians access to patient medical records from anywhere within the hospital. Even with their limited screen size there are clear advantages from having this increased availability of information.

In this paper, we explore the use of context-aware computing to go one step further. We present a handheld system that provides medical staff with information based on their context of work, mainly, their location. The location-aware HIS discussed in this paper can be used to retrieve medical information relevant to the user's current activity. For instance, a patient's medical record can be made available when the physician is near her bed. The system can also be used to locate peers and devices, tasks performed quite often in hospitals. This work is based on previous efforts we made to support context-aware communication of the hospital's staff [16].

The remainder of the paper is organized as follows. Section II presents the requirements for a location-aware HIS and a scenario used in its design. Section III gives an introduction to the field of context-aware computing. In Section IV, we describe recent efforts aimed at estimating user location within buildings and describe our own approach based on a wireless local area network (WLAN) and a backpropagation neural network. In Section V, we describe how the location-estimation method proposed is integrated into a location-aware HIS. Finally, Section VI discusses our results and in Section VII we present the conclusions.

II. Requirements for Location-Aware Hospital Services

We conducted a field study in a local public hospital. From interviews with the hospital's staff and the observation of work practices, we identified the following information and services that medical workers need to access and depend on the user's location:

Manuscript received August 31, 2003; revised April 13, 2004. This work was supported in part by CONACYT under Grant U-40799 and Grant C03-42447.

M. D. Rodríguez is with the CICESE Research Center, Ensenada, Mexico, and also with the Universidad Autónoma de Baja California (UABC), Mexicali 21280, Mexico (e-mail: marcerod@cicese.mx).

J. Favela, E. A. Martínez, and M. A. Muñoz are with the Centro de Investigación Científica y de Educación Superior (CICESE), Ensenada 22860, Mexico (e-mail: favela@cicese.mx; emartinez@cicese.mx; mike_munoz@hotmail.com).

Digital Object Identifier 10.1109/TITB.2004.837887

1) Documents such as patient's records, laboratory results, and forms to be filled. Considerable coordination efforts are required to locate relevant documents. Clinical records are often misplaced and laboratory results could take hours to be delivered to the person who requested them even to the point of making them useless.
2) The location of patients and colleagues. For instance, a physician might need to locate a specialist with whom to consult a clinical case, or a head nurse might require the help of nurses to attend an emergency.
3) Locating devices. Medical equipment, beds, and other devices need to be moved within the hospital as needed. Identifying the availability and location of these artifacts takes time and effort. For instance, a nurse might be tracking the availability of a bed to transfer a patient to the emergency room, or a doctor might want to display the laboratory results he has just received on the nearest public display.

These needs have shaped our design of the location-aware HIS. We illustrate the desired functionality of the system with a sample scenario:

> While Dr. Diaz is checking the status of a patient (Bed 222), he realizes that he should request a laboratory study. Using his handheld, he makes this request through the patient's electronic clinical record. When the chemist responsible for taking samples for the analysis visits the internal medicine area, his handheld informs him that in Bed 222 a patient requires laboratory analysis. When the chemist stands in front of the patient, his handheld lists the samples he must take and the type of analysis to be performed. Once he performs these analyses, he adds the results to the patient's clinical record. Afterwards, when Dr. Diaz visits the patient on his next round, the laboratory results will be displayed on his handheld. On the basis of these results, he re-evaluates the patient and decides to fill a medical note requesting the nurse in charge to increase the doses of the patient's medication.

We use this particular scenario because we learned from our site study that laboratory results could take up to 8 h to be delivered once the results are obtained. This is true even though the processing of the samples is almost completely automated and normally takes only a few minutes. We were informed by the hospital staff that often the results are delivered when they are of little or no use.

In this scenario, the system is continuously estimating the location of Dr. Diaz and communicating this information to the HIS, which updates the physician's location on all other users' handhelds. After a few seconds of Dr. Diaz being near Bed 222, the medical record of this patient is displayed on his PDA. Thus, the system adapts to the context, and in particular to the location, of its user. Context-aware computing is a growing area of research that deals with the design of this type of adaptable systems.

III. Context-Aware Computing

Context-aware computing refers to an application's ability to adapt to changing circumstances and respond based on the context of use. A system is context-aware if it uses context to provide relevant information and/or services to the user, where relevancy depends on the user's task [7]. Among the main types of contextual information considered relevant are identity, time, activity, and location which are known as primary context [7], [20]. This information answers the questions of who, when, what, and where, which can be used to identify if a specific information is relevant to establish context. Primary context can be used to derive additional related information which is called secondary context.

Context is difficult to use for several reasons [8]. First, capturing primary context information requires the use of sensors and computing devices. Context must be abstracted to make sense to the application, for instance, the ID of a mobile user must be abstracted into the user's name or role. Finally, context is dynamic, i.e., a mobile tour guide must update its display as the user moves, which require tracking the user's location by gathering information from multiple sensors, and using techniques that estimate the user's location or guess the route that a user will follow, which may introduce uncertainty.

A number of location-based services have been implemented in recent years, such as the Guide systems that provide city visitors with a hand-held context-aware tourist guide [5]. The information presented to visitors is tailored based on their profile and contextual information, such as time and the unit's physical location. For instance, the ordering of the tour recommended by the system can change dynamically when the visitor stays at a location longer than anticipated. "Safe & Sound" is a location-tracking system that allows parents to monitor their children's position [15]. The child has a phone that continuously streams location information to the parent's phone. If the child is outside a secure zone previously defined by the parents, both the parent and the child receive a sound alert and a voice channel between them enables negotiating with the child.

Undoubtedly, location is important to understand the context of mobile users [9]. Location becomes a useful indexing information from which to infer the overall context that a system will use to provide services and information to mobile users. Furthermore, mobile users constantly change their context, most notably their location. This is particularly true in a hospital setting where the staff is constantly moving and the activities they perform are highly dependent on their location. For instance, access to patient's records is most relevant when near the patient's bed. In addition, resources, such as clinical records, laboratory results, or devices are highly distributed and people spend time tracking them.

IV. Estimating User Location in WLANs

A. Location Estimation Methods

Estimating the location of a user has been a subject of considerable attention in context-aware computing in recent years. The first location-based computer applications reported in the literature were based on the Active Badge system which made use of infrared signals emitted by badges and received by infrared sensors located in a building [22]. Advances in Global Positioning Systems allow mobile computers to determine its location with considerable accuracy in outdoor environments. They use triangulation of the signals received from multiple satellites to determine location with an approximate accuracy of 10 m. Other techniques involve the use of ultrasound [13], sensors placed in the floor [17], and the use of multiple cameras and computer vision techniques [6].

Radio frequency identification (RFID) tags have become increasingly popular to track and identify products and people. A major feature of RFID is its read–write capability. Users can seamlessly record and transfer data, such as serial numbers, personal records, account information, etc., from their computer system to an RFID tag (or vice versa). These tags can be attached to a wristband or a smartcard. Several uses of this technology have been identified for the healthcare domain where wristbands with RFID tags are already being commercialized. For instance, RFID systems can help streamline operations and ensure positive patient identification to reduce medical errors by facilitating real-time confirmation of the right patient, right drug and dose. This type of RFID tag is passive since it must be closely coupled with a reader to transmit information. With a read distance generally limited to three feet, a high concentration of tags would be required to continuously track the location of people, as required for the uses we envisioned for this technology. There are also active RFID tags with transmission distances of over 100 feet that use a battery to power the chip's circuitry and broadcast a signal to a reader. This solution, however, requires the installation of a dedicated infrastructure besides the computer network as is the case with most of the solutions discussed above.

Of particular interest are location estimation techniques that make use of an existing WLAN infrastructure, since they have better scalability and less installation and maintenance costs than *ad hoc* solutions, which are requirements for easily surveying a location system using its own infrastructure and components [21]. These methods use the radio frequency (RF) signal strength (SS) between a mobile device and several access points of the WLAN to estimate location. A signal propagation model could be used to estimate this distance, but rather complex models would be required, since the signal is affected by the presence of walls, furniture, and other people and devices. To work around this complexity, empirical methods have been advanced. In these methods, the strength of the RF signal is measured at predefined locations and used to train a pattern recognition model that can then be used to estimate the user's location.

The RADAR location system uses an IEEE 802.11 WLAN and an empirical method based on the nearest neighbor algorithm [1]. Similarly, the Nibble system uses the signal-to-noise ratio (SNR), which is more stable than SS to compute the distance to the access point. Nibble uses a Bayesian network to estimate the probability of the mobile being at one of a set of discrete locations [3].

A recent trend is to estimate the user's location by fusing the readings from different sensor technologies to obtain a more precise estimation of the measured variables; this process is known as sensor fusion. Such an approach is used in [11], which presents an indoor location-measuring system that processes data from infrared and ultrasonic sensors. In spite of their increased accuracy, these procedures require a specialized infrastructure, which may not be suitable for wide deployment scenarios [21].

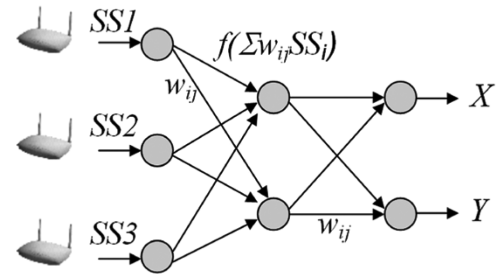

Fig. 1. Simplified architecture of the neural network used to estimate the location of mobile users.

B. Neural Networks for Location Estimation

In our work, we are using backpropagation neural networks trained to map RF signals from a WLAN to two-dimensional coordinates. A backpropagation neural network, or multilayer perceptron, is a supervised nonparametric model that learns from a training set by adjusting the weights that shape the strength of the signals propagated between processing units, or neurons. Once trained, the neural network can be used to classify incoming patterns into labeled classes.

The architecture of a backpropagation neural network is made of two or more layers of processing nodes with each of the nodes of layer i connected to each of the nodes in layer $i+1$. The connections have different strengths (or weights) that represent the influence that a particular node has in a node of a subsequent layer. Each node computes an activation value that is the result of applying a nonlinear function (typically the sigmoid function, when the errors are assumed to be Gaussian) to the sum of the products of the activation weights of the previous layer, times the weight that connect each of those nodes with the unit performing the computation.

A network that is presented with an input pattern in the first layer will then generate a pattern in the output (or last) activation layer. The most important aspect of neural networks, though, is not how they compute these output patterns, but rather how they learn from a set of examples. The backpropagation learning algorithm calculates how the nodes in the internal layers will be penalized (adjusted) when presented with a training pair (a tuple of input and output patterns). The details of the algorithm can be found in [19].

Fig. 1 shows a simplified version of the architecture of the neural network we have used. The SS from each access point is presented to the input layer. We have used a single hidden layer and a two-node output that represents the X and Y coordinates of the location.

C. Experimental Setup and Results

To train the neural network, we measured the SS and SNR from five access points located in a 40 × 20 m building, as illustrated in Fig. 2. We took measurements on 154 different locations within the building. At each location and for each direction (north, south, east, and west), we recorded seven samples of the signals from the access points to the mobile device, the seven samples were averaged to obtain a total of 616 samples.

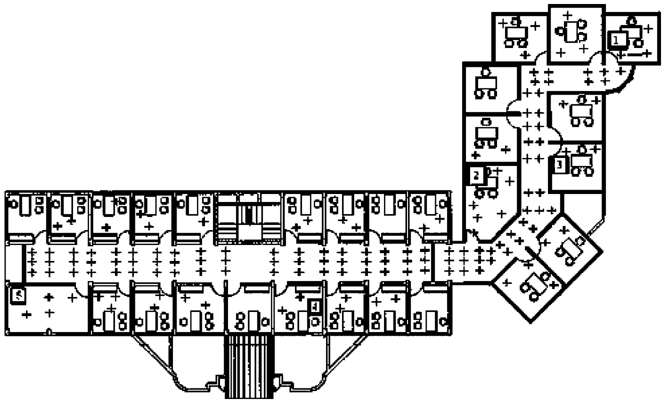

Fig. 2. Map of the building where measurements were taken.

We trained the network with 432 of the 616 patterns and used the remainder 184 to determine the estimation error of the different configurations of neural network that we trained.

We trained several neural networks using different configurations and learning algorithms. In all cases, the output layer had two neurons corresponding to the X and Y coordinates we want to estimate.

In the input layer, we experimented with three different alternative inputs: 1) five neurons with the SS from each of the five access points; 2) five neurons with the SNR from each of the access points; and 3) ten neurons with both the SS and SNR from each of the access points. The best results were obtained with the SNR, which is more stable than SS as has been reported [3].

We used a single hidden layer and performed experiments with 4, 6, 8, and 16 neurons in this layer. The best results were obtained with 16 nodes in the hidden layer. As activation function, we used the sigmoid function on the input and hidden layers, and the identity function in the output layer.

We used backpropagation as the learning method to train the neural network. Although backpropagation is the most popular algorithm for supervised neural networks, it has a poor convergence rate. There are a number of variations of the basic algorithm that are based on other standard optimization techniques, such as conjugate gradient and Newton methods [4]. We compared eight different backpropagation variants, from which, the most efficient was the conjugate gradient with Polak–Ribiére updates. This algorithm performs a lineal search in each iteration to find the learning coefficient in which the average error decreases the most with respect to the previous iteration. Moreover, the conjugate gradient algorithm fulfills several conditions that guarantee that an optimum learning coefficient was found.

To compare the different configurations and variants of the learning algorithm, we used an euclidian error, measured as the distance between the actual and estimated patterns in the test set. The best results were obtained with a 5-16-2 configuration, using the SNR as input, and the conjugate gradient with Polak–Ribiére updates. For this case, we obtained an average error of 2.0947 m. The cumulative distribution function for the best configuration is displayed in Fig. 3. The graph shows the percentage of patterns that fall within a given distance. It can be seen that 80% of the patterns are within 3 m of their target.

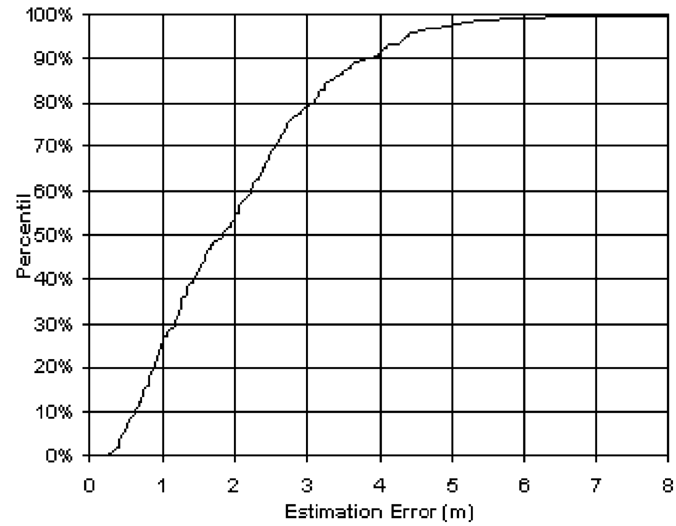

Fig. 3. Cumulative distribution of the error estimated on the location of a mobile device.

D. Implementation of the Location Estimation Module

The training of the location estimation methods was performed using Mathlab and its Neural Network Toolbox.

The measurements for the training and test sets were taken with a laptop computer with an Orinoco Silver 802.11b network card. To collect the samples of the signal intensity from each access point, we used the Orinoco Client Manager software v.2.9, which communicates directly with the network card. The access points used included one Apple Airport model and four Orinoco AP-200. Afterwards, we tested our solution estimating the location of the laptop while it moved using real-time readings of the SNR to the access points.

We implemented the location estimation module in a handheld, our target device. The Windows API was used to read the SNR from a Dell TrueMobile 802.11b card of a Dell Axim handheld with the Pocket PC 2003 operating system. To facilitate the integration of this module within a location-aware hospital system, the neural network was wrapped as a software agent as described in the next section.

V. Supporting Location-Aware Hospital Services

A. System's Implementation

The location-aware hospital system was conceived as an agent-based system that was developed with SALSA, a middleware that provides a set of abstract classes for implementing autonomous agents that act on behalf of users, represent services, or wrap a system's functionality [18]. Agent technology is a useful abstraction for the design of complex systems with distinct and independent components by enabling the aggregation of different functionalities.

A SALSA agent contains several components: a protocol to register the agent with an *Agent Directory*; an instant messaging (IM) client through which users, users' agents, and devices' agents interact by sending extended markup language (XML) messages; and finally, the subsystem that implements the agent's intelligence that includes components for perception,

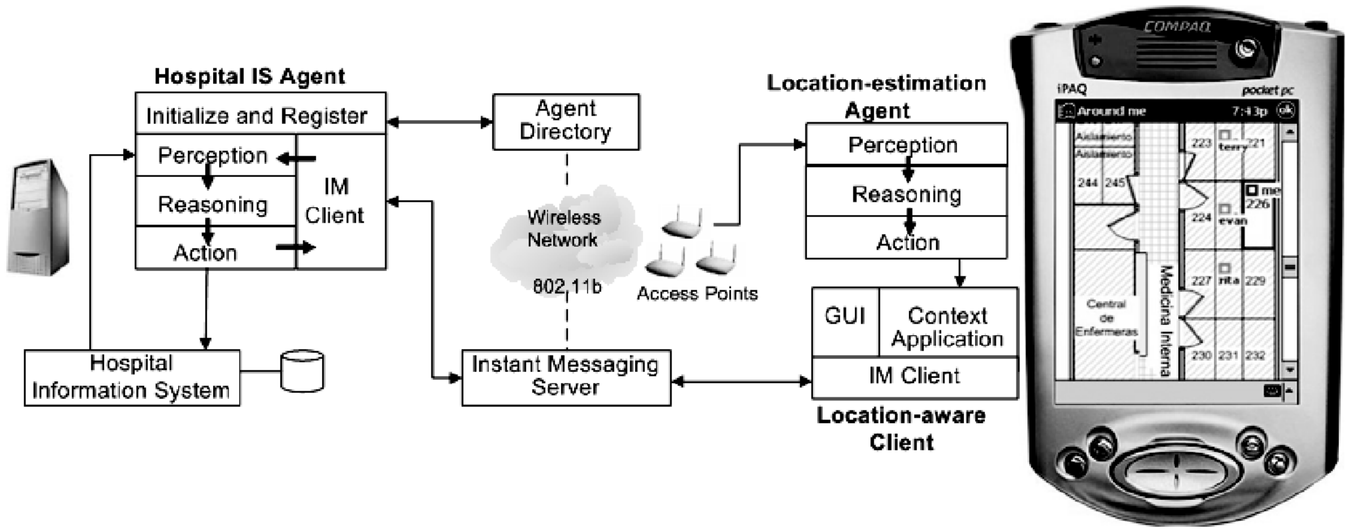

Fig. 4. Architecture of the location-aware hospital system.

reasoning, and action. The perception module gathers knowledge through the IM client from the environment's sensors, other agents, or directly from users or services. The reasoning subsystem governs the agent's actions, including deciding what to perceive next. The action component interacts with the environment by sending messages using a predefined protocol.

As mentioned above, the location estimation component is implemented as a SALSA agent. The trained neural network to estimate the user's location was wrapped in the reasoning component of the location-estimation agent (LE-a). The perception module reads the SNR from the access points and the action component notifies the estimated location to the location-aware application in the handheld.

Since the SS decays considerably from floor to floor, one estimation agent is trained for each hospital floor. When the strength of the signal from the access points in one floor goes below a certain threshold, a new estimation agent is loaded into the PDA. This will correspond to the agent trained for the floor with the access points that report higher SS.

The components of the LE-a interact to estimate the user's location, as describe next. The agent's perception module includes a `PassiveEntityToPerceive` object, which reads the SNR from the wireless network card, and then notifies it to the `PassivePerception` object. When the SNR value is changed, the `PassivePerception` object sends the new value to the `Reasoning` component, which makes a new estimation of the user's location. This information is communicated by the `Acting` component to the location-aware client. The location-aware hospital application will map the X, Y coordinates to an area identifier (bed number, room, etc.) and will communicate this new location to the rest of the system's agents and users trough its IM client.

Fig. 4 illustrates how the LE-a is integrated in the location-aware hospital system. Besides the LE-a, the handheld computer carried by physicians and nurses includes an application that provides them with information relevant to their location, and allows them to fill forms and communicate with other members of the staff. The interface of the location-aware client is based on the IM paradigm and requires only peripheral attention. Through this interface, users are notified of the availability of other users and their location. This information is displayed in the form of a list (as in traditional IM system) or in a map showing the area surrounding the user (Fig. 4). The location-aware client provides access to the HIS. Information considered relevant to the user's location and role is offered by default, but users can consult additional information either by navigation or querying the system.

The information received in the handheld is obtained from an HIS that manages and stores the patient's clinical records and other data relevant to the hospital, such as what patients are in what beds. An agent (HIS agent) acts as proxy of the HIS, it provides access to, and monitors the state of, the information contained in it. Rules are used to indicate what type of information should be delivered to a user given its current location and role. For instance, considering the scenario explained in Section II, when Dr. Diaz stands in front of the patient, the HIS agent perceives Dr. Diaz' position through its IM client, and analyzes the context information, such as the user's role, and the availability of the patient's laboratory results, and then acts by notifying these results to the doctor's client. This agent runs as a daemon on a computing device with connectivity to an agent directory and the IM server.

Finally, the last component of the architecture is an IM server that acts as an agent broker. Our implementation uses and extends the Jabber open-source IM server (www.jabber.org) and its extensible messaging and presence protocol (currently an Internet Engineering Task Force draft) to report the state of people and agents and to handle the interaction among people, agents, and devices through XML messages.

B. Sample Scenario

We illustrate the use of the location-aware hospital system with the scenario presented in Section II.

Fig. 5 illustrates how the components of the system's architecture interact for this scenario. Dr. Díaz begins his daily routine by visiting each one of his patients. While he moves around

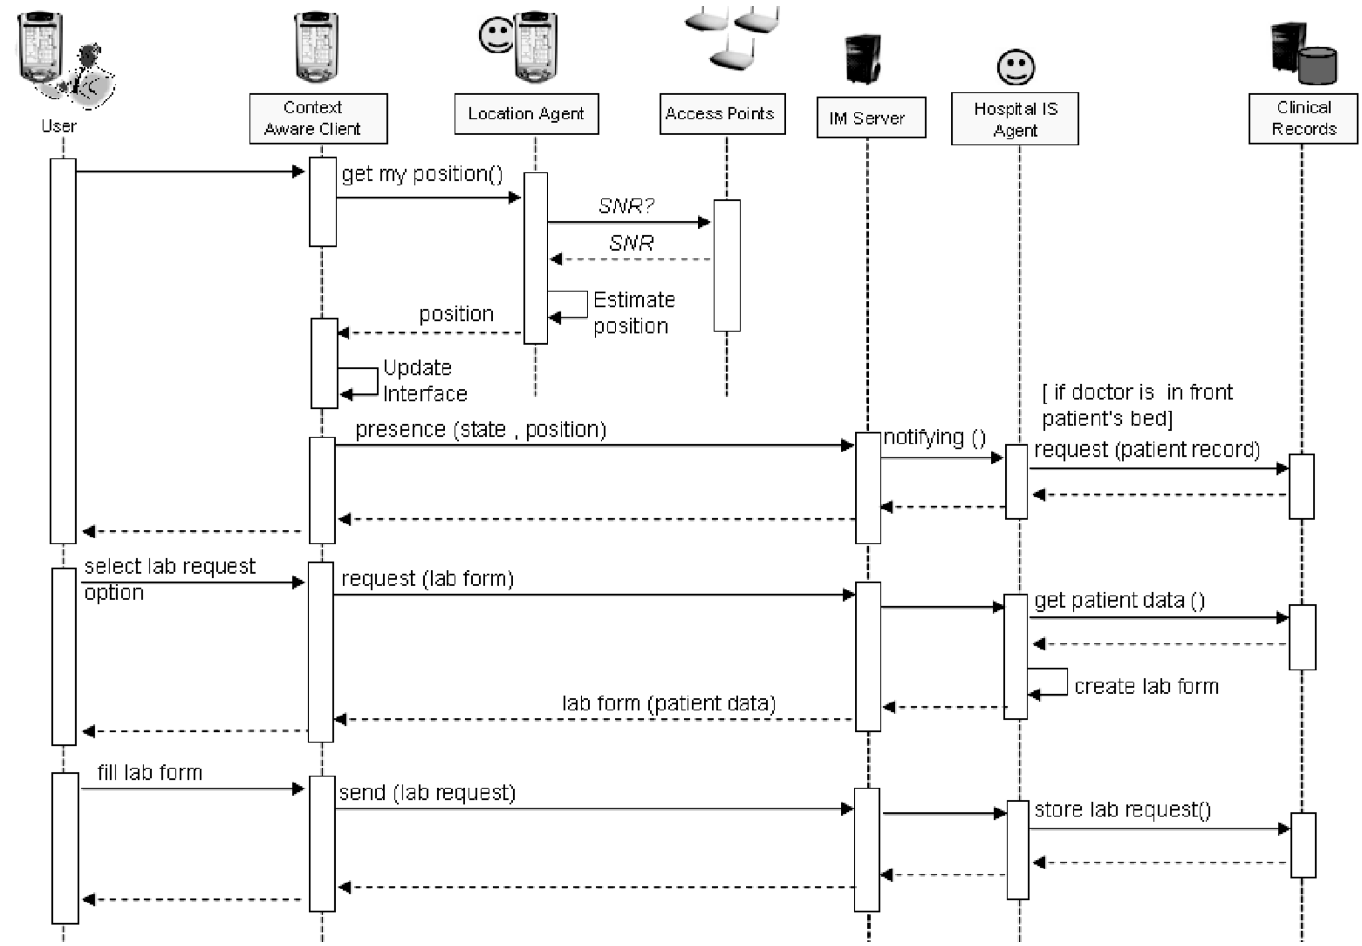

Fig. 5. Sequence diagram for the request of laboratory analysis.

the patient's rooms, the context-aware client in his PDA communicates with the LE-a to constantly update his position. When his location changes, the LE-a sends, through the IM server, the doctor's position to all users and agents who have him registered in their rosters, such as the HIS agent. Then the IS agent verifies if its contextual conditions match the new context, i.e., the user's role and location, in order to send him a message through which he can directly retrieve the patient's clinical record.

After consulting both the record and the patient, Dr. Diaz decides to request a laboratory analysis by using the form illustrated in Fig. 6(a), which is customized for him based on his identity (Dr. Diaz), role (physician), and current location (Bed 222). Once he selects the "Laboratory Study Request" option, the HIS creates the laboratory form, which includes some of the patient's data such as his name and bed number, as shown in Fig. 6(b). Dr. Diaz fills the laboratory form and sends it to the HIS to be added to the patient's clinical record.

The HIS will inform the chemist of the analysis to be performed and the medical samples to be taken. Once the chemist performs the analyses, he adds the results to the patient's clinical record.

The next time Dr. Diaz is near the patient, the HIS will send him a message indicating that the laboratory results are available and offering to display them on his handheld, as shown in Fig. 7(a). After analyzing these results, the doctor decides to fill

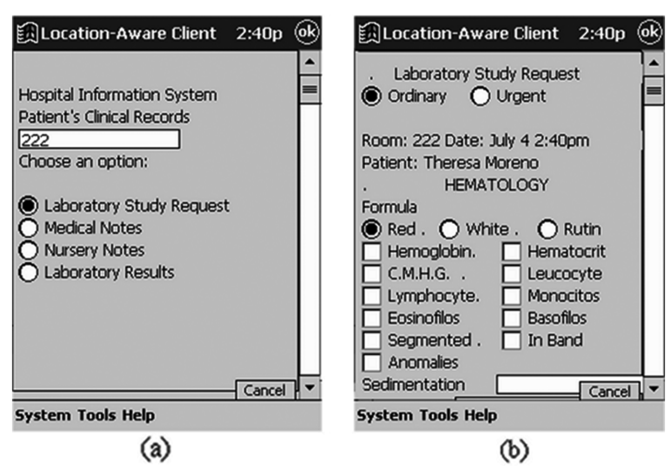

Fig. 6. Requesting a laboratory analysis for the patient.

a medical note requesting an increase in the doses of the patient's medication [see Fig. 7(b)].

VI. Results

Based on the room size and bed concentration of the hospital we have used as our site study, we had set as target a maximum error of 2 m in order to deliver location-aware information, such as patient's records, to the hospital's staff. For the purpose of

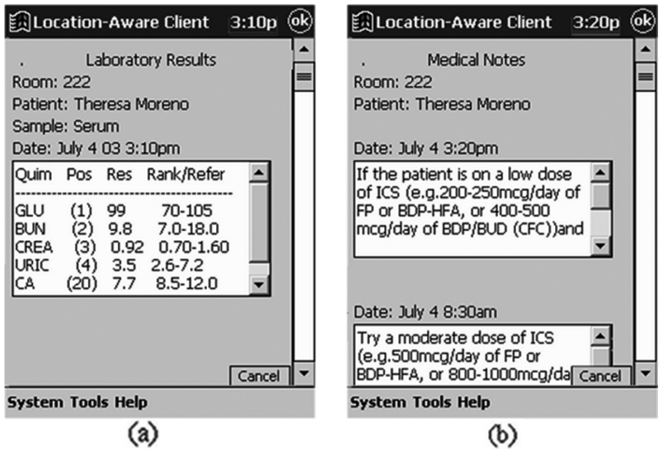

Fig. 7. (a) Physician receives the laboratory results and (b) as a result of his diagnosis fills a medical form.

locating people within the hospital, a 4-m maximum error is considered adequate, since within this range a person can be reached visually or by calling her name.

The results we have obtained indicate a probability that we would satisfy our criteria 92% of the time for locating people, but only 50% of the time for the delivery of patient related information. This would indicate that the location-estimation method we have used is suitable for the first task, but not for the second. However, these results are the errors from independent estimations; they do not take into account the fact that people move smoothly seldom exceeding a speed of 1.5 m/s when walking in closed public spaces. When people move at speeds higher than 1.5 m/s, they are normally in a hurry and a precise estimation of their location may be of little importance for a location-aware hospital system. In fact, for our purposes, hospital staff will spend at least several seconds in the area of interest (a patient's bed, a laboratory, or an office) and, thus, a better estimation of its location can be obtained.

We have modified our algorithm to include a simple criterion that takes into account this fact. Of this, we make four readings in one second of the SNR from each access point and estimate the location for each of them independently. The resulting estimation is computed from the average of the four estimations. However, if one of the estimations is more than 1.5 m away from the average, it is eliminated and the location is calculated using the other three estimations. Using this criterion, with the same data we were able to reduce the average estimation error from 2.09 to 1.87 m. For this case, 96% of the estimations fall within 4 m and 58% are within 2 m. This is however, just a simple strategy for estimating continuous movement. Hidden Markov Models or algorithms for time series prediction, including neural networks, could be used to improve location estimations. Furthermore, continuous location estimation can also be used to determine the direction of the user being tracked, which can be used to infer her current task or availability.

We have compared our results with the k-Nearest Neighbor algorithm and obtained equivalent precision. Using the continuous estimation criterion mentioned above, for instance, we obtain an average error of 1.9 m versus the 1.87 we have obtained with the use of neural networks. Neural networks, however, offer the advantage of being less memory intensive, an important feature given the limitations of PDAs. With the configuration and data we have used, for instance, the nearest neighbor approach will require the storage of 27 times more data than neural networks. This is important not only in terms of storage, but also for the transfer of this data through the wireless network when a user moves to a new floor.

Another important issue raised by context-aware computing is that of privacy. Our approach deals with this issue in two fronts, from the technological and the social point of view. From the technical side, location is estimated at the handheld device, thus, this information is shared with other users and the HIS only if, and with whom, the user whishes to do so. Additionally, at the server side, the area can be configured in a way that prevents from sharing the location of users in certain rooms. For instance, it could be decided that the location of a subject will not be shared if he is in the lounge or the bathroom.

From the social point of view, people who work in hospitals are expected to have a certain degree of availability to their peers, and in fact locating them through various ways, including the use of speakers or sending text messages to their cellular telephones, is common practice. The location-aware system we are proposing does not seem to provide a threat to their privacy.

Security is also an important concern when dealing with hospital information. The system authenticates PDA users requiring them to login when they start the location-aware hospital system and will only provide access to information relevant to their work. This would not prevent someone from stealing a PDA and accessing patient records if he knows the user's login and password. However, there are commercial PDAs that incorporate fingerprint scanners to identify its user; the use of these devices would prevent intruders from accessing the hospital server.

VII. Conclusion

Hospitals are complex work environments where information and people are distributed, thus requiring considerable coordination and communication among the professionals that work in such settings. Electronic medical records are an important step toward providing adequate access to clinical information. However, the most precious resource is the attention of the medical personnel. With adequate support to estimate the context of work, context-aware systems can deliver information that is relevant to the user's location, identity, and/or role. In particular, location is an important factor to establish the information that is relevant to a given user and, thus, reduce the burden of locating people or data, or even worse, making decisions without it.

We have presented a handheld-based HIS that can be used to deliver information based on the location of medical staff. A location estimation method, based on a backpropagation network, is used by the system to locate the mobile device and the user that carries it. Our results show that estimation errors are adequate to locate people but higher than required to deliver relevant patient information. Additional work is required to improve the estimation by tracking the user over time rather than relying only on individual samples.

REFERENCES

[1] P. Bahl and V. N. Padmanabhan, "RADAR: an in-building RF-based location and tracking system," in *IEEE INFOCOM 2000*, Tel-Aviv, Israel, 2000.

[2] C. Bossen, "The parameters of common information spaces: the heterogeneity of cooperative work at a hospital ward," in *Proc. ACM Conf. Computer-Supported Cooperative Work*, New York, 2002, pp. 176–185.

[3] P. Castro, P. Chiu, T. Kremenek, and R. Muntz, "A probabilistic room location service for wireless networked environments," in *Proc. 3th Int. Conf. Ubiquitous Computing (UbiComp 2001, LNCS 2201)*, Atlanta, GA, 2001, pp. 18–34.

[4] C. Charalambous, "Conjugate gradient algorithm for efficient training of artificial neural networks," in *Proc. Inst. Elect. Eng., Part G*, vol. 139, 1992, pp. 301–310.

[5] K. Cheverst, N. Davies, K. Mitchell, and A. Friday, "Experiences of developing and deploying a context-aware tourist guide: The GUIDE project," in *Proc. Mobile Computing and Networking*, 2000, pp. 20–31.

[6] T. Darrell *et al.*, "Integrated person tracking using stereo, color, and pattern detection," in *Proc. Conf. Computer Vision and Pattern Recognition*, Los Alamitos, CA, 1998, pp. 601–608.

[7] F. D. Davis and V. Venkatesh, "Measuring user acceptance of emerging information technologies: An assessment of possible method biases," in *Proc. 28th Hawaii Int. Conf. System Sciences*, 1995, pp. 729–736.

[8] A. K. Dey, "Understanding and using context," in *Personal and Ubiquitous Computing*. London, U.K.: Springer-Verlag, 2001, vol. 5, pp. 4–7.

[9] A. Dix, T. Rodden, N. Davies, J. Trevor, A. Friday, and K. Palfreyman, "Exploiting space and location as a design framework for interactive mobile systems," *ACM Trans. Compter-Human Interaction*, vol. 7, no. 3, pp. 321–385, Sept. 2000.

[10] S. Fischer, T. Stewart, S. Mehta, R. Wax, and S. Lapinsky, "Handheld computing in medicine," *J. Amer. Med. Inform. Assoc.*, vol. 10, no. 2, pp. 139–149, Mar./Apr. 2003.

[11] D. Fox, J. Hightower, L. Liao, D. Schulz, and G. Boriello, "Bayesian filtering for location estimation," *IEEE Pervasive Computing*, vol. 2, pp. 24–33, July/Sept. 2003.

[12] "Harris Interactive. Physicians' use of handheld personal computing devices increases from 15% in 1999 to 26% in 2001," *Harris Interactive Health Care News*, vol. 1, no. 25, Aug. 15, 2001.

[13] M. Hazas and A. Ward, "A novel broadband ultrasonic location system," in *Proc. 4th Int. Conf. Ubiquitous Computing (UbiComp 2002)*, Goteborg, Sweden, 2002, pp. 264–280.

[14] A. Lapinsky, J. Weshler, S. Mehta, M. Varkul, D. Hallett, and T. Stewart, "Handheld computers in critical care," *Critical Care*, vol. 5, no. 4, pp. 227–231, 2001.

[15] N. Marmasse and C. Shmandt, "Safe & Sound—A wireless leash," in *Proc. CHI 2003 Extended Abstracts on Human Factors in Computing Systems*, pp. 726–727.

[16] M. Muñoz, M. Rodriguez, J. Favela, V. M. Gonzalez, and A. I. Martinez-Garcia, "Context-aware mobile communication in hospitals," *IEEE Computer*, vol. 36, pp. 60–67, Sept. 2003.

[17] R. J. Orr and G. D. Abowd, "The smart floor: A mechanism for natural user identification and tracking," in *Proc. 2000 Conf. Human Factors in Computing Systems (CHI 2000)*, New York, 2000.

[18] M. Rodríguez and J. Favela, "Autonomous agents to support interoperability and physical integration in pervasive environments," in *Proc. Atlantic Web Intelligence Conf. (AWIC 2003)*, 2003, pp. 278–287.

[19] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning internal representations by error propagation," in *Parallel Distributed Processing: Explorations in the Microstructures of Cognition, Vol. 1: Foundations*, D. E. Rumelhart and J. L. McCelland, Eds. Cambridge, MA: MIT Press, 1986.

[20] B. N. Schilit and M. M. Theimer, "Disseminating active map information to mobile hosts," *IEEE Network*, vol. 8, no. 5, pp. 22–32, Sept./Oct. 1994.

[21] J. Scott and M. Hazas, "User-friendly surveying techniques for location-aware systems," in *Proc. Ubicomp 2003*, Berlin Heidelberg, 2003, pp. 44–53.

[22] R. Want, A. Hoopper, V. Falcao, and J. Gibbons, "The active badge location system," *ACM Trans. Inform. Syst.*, vol. 10, no. 1, pp. 91–102, 1992.

Marcela D. Rodríguez received the M.Sc. degree in computer science from CICESE, Ensenada, Mexico, and the B.Sc. in computer engineering from the Autonomous University of Baja California (UABC), Mexicali, Mexico. She is working toward the Ph.D. degree in computer science at CICESE.

She is a Professor in Computer Engineering at the UABC. Her research interests include ubiquitous computing, autonomous agents, and CSCW.

She is a Student Member of the Association for Computing Machinery.

Jesus Favela received the B.S. degree from the National Autonomous University of Mexico (UNAM) and the M.Sc. and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, MA.

He is a Professor of Computer Science at CICESE, Ensenada, Mexico, where he leads the Collaborative Systems Laboratory and heads the Department of Computer Science. His research interests include CSCW, ubiquitous computing, and information retrieval.

He is a Member of the IEEE Computer Society, Association for Computing Machinery, and the current President of the Sociedad Mexicana de Ciencia de la Computacion (SMCC).

Edgar A. Martínez received the B.Sc. degree in Computer Engineering from the Tecnológica de la Mixteca University, Oaxaca, México, in 2004.

During the summers of 2002 and 2003, he was a Research Assistant in the Department of Computer Science, CICESE. His research interests include wireless communications, neural networks, and context-aware computing.

Miguel A. Muñoz received the B.Sc. degree in computer engineering from the Instituto Tecnológico de Zacatecas (ITZ). He received the M.Sc. degree in computer science at CICESE's Research Center, Ensenada, Mexico.

His research interests are in context-aware computing, ubiquitous computing, and CSCW.