Chapter 17
Automated Workflow Analysis and Tracking Using Radio Frequency Identification Technology

Mithra Vankipuram, Thomas G. Kannampallil, Zhe (Eric) Li, and Kanav Kahol

Introduction

The health care industry faces a number of challenges and arguably one of the most important ones lies in maintaining high levels of patient safety. A much-cited report released by the Institute of Medicine [1] estimates that as many as 98,000 people die each year due to medical errors [1]. The causal determinants of these errors can be traced to a variety of medical, cognitive and social challenges in the clinical workplace. These challenges are exacerbated in critical care environments that are characterized by distributed, interdependent, episodic and non-linear work activities. The dynamic nature of the care process in critical care environment affects the nature and timing of work activities of clinicians, and often increases the possibility

---

Portions of this chapter has appeared in (a) Vankipuram et al., Toward automated workflow analysis and visualization in clinical environments. *Journal of Biomedical Informatics*. 44(3): 432–440, with permissions from Elsevier; (b) Kannampallil et al., Making sense: sensor-based investigation of clinician activities in complex critical care environments, *Journal of Biomedical Informatics*. 44(3), 441–454, with permissions from Elsevier and (c) an article in the Proceedings of the 2009 Annual Symposium American Medical Informatics Association, Vankipuram et al., *Visualization and analysis of activities in critical care environments*.

M. Vankipuram, MS, PhD (✉)
Analytics Lab, Hewlett-Packard (HP) Labs, Palo Alto, CA 94304, USA
e-mail: vankipuram.mithra@gmail.com

T.G. Kannampallil
Center for Cognitive Studies in Medicine and Public Health, New York Academy of Medicine, New York, NY 10029, USA

Z. (Eric) Li, MS
Drchrono Inc., Mountain View, CA 94043, USA

K. Kahol, PhD
Affordable Health Technologies Division, Public Health Foundation of India, New Delhi, India

V.L. Patel et al. (eds.), *Cognitive Informatics in Health and Biomedicine*, Health Informatics, DOI 10.1007/978-1-4471-5490-7_17, © Springer-Verlag London 2014

of errors. Studying the work activities of clinicians in such environments can help in understanding the care delivery process, workflow, and interruptions that affect clinical work.

Exploratory investigations of clinician activities are often performed using observational methods. While these methods provide a descriptive depth that cannot be matched by automated methods, use of participant observation methods [2, 3] in a critical care setting is often challenging, as capturing the work activities of multiple clinicians requires several observers who must be closely synchronized during their data capture sessions. The tools currently used for workflow analysis in clinical environments include methods such as ethnographic observation, shadowing of individual clinicians, surveys and questionnaires [4]. The data collected by these methods can be used to model work activities centered on a particular individual (or a group) and their activities [5]. Such approaches sometimes are inadequate to develop a holistic and complete picture of work activities. For example, observations are gathered from an individual's point of view and may not be adequate to capture multiple activities occurring within a clinical environment. Though it is plausible to capture additional activities by increasing the number of observers, such an approach is highly likely to disrupt clinical activities. Given these constraints in complex critical care environments, there is a need for an unobtrusive alternative that can augment existing methods of data collection and enable piecing together a more complete workflow, understanding the nuances of underlying activities, interactions and dependencies.

Tools that can be used to monitor continuous activities in clinical environments can provide significant insights into the work activities in clinical environments. Radio-frequency Identification (RFID) technology offers a seamless, cheap and effective mechanism for monitoring and tracking events within clinical environments. In this chapter, we describe the potential and use of RFID-based sensors for reliably capturing the activities, mobility and interactions of clinicians. This chapter is based on aggregated results from our previously published work that on the use of RFID technology in critical care settings [6–8].

Background

Complexity and Critical Care Workflow

The study of complex systems draws together emerging approaches from several diverse fields including economics, physics, biology, mathematics and computer science on the common ground of complexity. This interdisciplinary effort seeks to formulate unifying principles of complexity. Several authors have proposed that the healthcare system or elements thereof can be characterized as a complex system [9–13]. For example, Smith and Feied [13] argue that an emergency department is a *paradigmatic complex system*. This argument rests on the unpredictability of both patients' clinical conditions and clinicians' work patterns, the vast decision space and incomplete evidence that complicate clinical decision-making and the inherent unpredictability of the system as a whole.

Several concepts drawn from the complex systems literature are pertinent to the study of a critical care unit as a complex cognitive system. A cogent and readable account of the ways in which concepts from the complexity literature might be applied to social systems has been developed by the Complexity in Social Science [14] project [14]. Complex systems are by their nature non-deterministic and dynamically structured. That is to say, it is not possible to predict the behavior of a complex system by studying the function of its components in isolation, and furthermore the study of the behavior of any such component reveals little about the system as a whole. Likewise, the process of clinical care emerges from a series of dynamic and flexible interactions between patients, health-care providers and outside influences [15]. While this argument applies readily to workflow, it also relates to the cognitive processes that underlie critical care decision making, as the cognitive processes in critical care settings are distributed across the minds of the clinical team and a range of physical media [16]. Given the complex nature of system behavior, it is not possible to predict the knowledge, expertise and information that will be available at the point in time at which clinical decisions are made. Similarly, for transfer of information, it has been observed that within complex social systems the flow of information is determined somewhat serendipitously by the geographical location of team members [17], which is influenced in turn by the complex dynamics of the system as a whole.

These aspects of the critical care workplace present challenges for the human-intensive ethnographic methods that have been previously employed. However, complex systems theory suggests that only limited insight into system behavior can be obtained through the study of component parts. Consequently, there is a role for automated sensors to complement the human-intensive data collection methods that have been employed previously. While not able to capture the depth and richness of representation that are possible through ethnographic methods, these sensors offer certain advantages in that it is possible to collect data concerning a geographically mobile clinical team over an entire shift. This is desirable, as even an exceptionally well-funded research program that may be able to employ multiple well-trained human observers is likely to experience problems integrating a team of observers into a busy clinical environment without obstructing patient care.

RFID Sensor Technology

Recent times have seen a prolific increase in the use of radio frequency identification (RFID) devices in clinical settings. This is driven by early research results that have shown that RFID technology can improve better tracking of patients, more effective and safer drug administration and lower monitoring costs. Potential advantages not withstanding, the widespread adoption has been tempered by the lack of consistent results regarding the viability of real time location systems (RTLS) in clinical settings. RFID tools have been used in a variety of applications including locating healthcare professionals, tracking patient flows, equipment and medication, and improving hospital-wide throughput, bed management, and workflow [18–21].

Sensors typically used for entity activity recognition include passive infrared sensors, radio identification tags and pressure sensors. For example, Fry and Lenert [22] developed a system for location tracking of patients, staff and equipment called MASCAL. RFID sensors were used to track clinicians and equipment during mass casualty events. Sensor tracking data was combined with personnel and clinical information to centralize the management of resources. In a related study, Chen et al. [23] describe the use of RFID sensors to identify patients, and notify clinicians on patient related information that decreased the waiting time for patients in intensive care units.

Sensor technology used for the studies described in this chapter was an active RFID system. The system is composed of *tags* and *base stations* that are used to capture the movement and interactions between the clinicians in critical care settings. Tags are mobile devices that help in the tracking of moving objects. Base stations are stationary devices that provide radio coverage and tracking of the tags. The tags and base stations communicate using a vendor-customized *IP-Lite* radio connection protocol. During data collection sessions, clinicians carried the RFID tags (i.e., the sensors) in the pockets of their coats. Base stations were placed at key locations to capture their movements and the transitions between spaces such as patient rooms. As a clinician carrying a tag comes in close physical proximity with a base station, a ping event is registered with that base station. This is referred to as a *tag-base* ping. The strength of ping event is measured in terms of received signal strength index (RSSI). Additionally, when two clinicians come in close proximity to each other, a *tag-tag* ping is registered. As with the tag-base pings, the relative physical distance between the clinicians is reflected in the signal strength of the tag-tag pings. The tags and base stations send pings at approximately 3-s intervals. In other words, for every 3 s, each tag registered with a corresponding base station in its vicinity.

Figure 17.1 shows the configuration of tags and base stations and how ping events are registered between them. In Fig. 17.1, interactions between three tags and one base station are shown. The tags register pings with each other (tag-tag pings, represented as n1, n2 and n3) and, concurrently register pings to the common base station (m1, m2 and m3 pings). The tag-tag and the tag-base pings are used for the identification of the location of a clinician (or multiple clinicians) and their collaborators at any particular point in time. The tag-tag interactions provide an additional dimension (of co-location of clinicians) through which to interpret the actions of the clinicians.

We use an illustrative example of how some activities in the clinical environments can be captured by appropriate placement of tags and base stations. Consider the scenario representing patient arrival is depicted in Fig. 17.2. First, key members

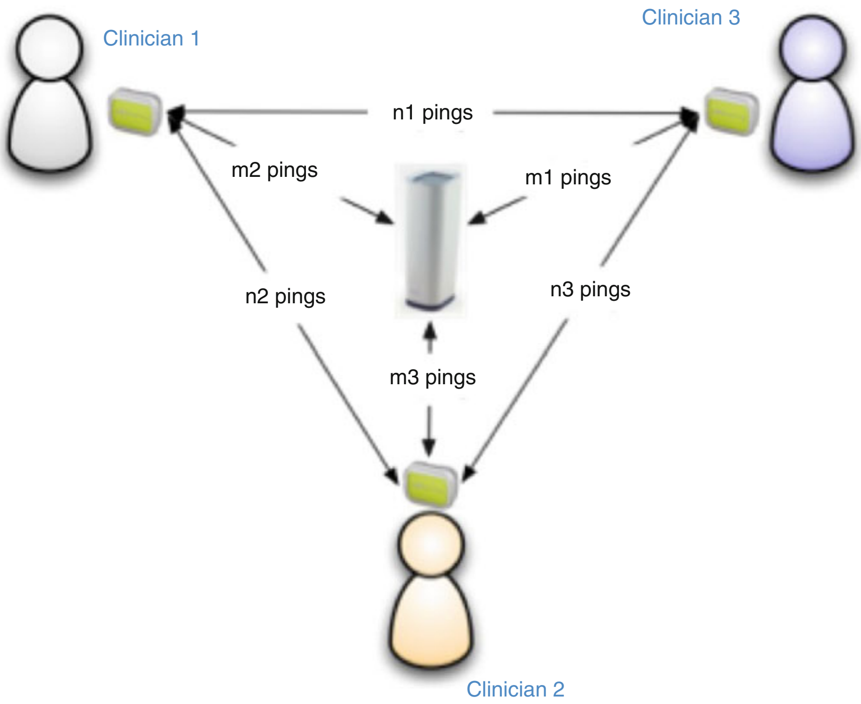

Fig. 17.1 Tag-Tag and Tag-Base configuration. Three tags and one base station is shown in the figure with interactions between them represented by pings (tag-tag and tag-base).

of the patient care team (resident, nurse and so on) gather by the bed of the patient. Following this, examination of the patient takes place. A resident may move to the telephone to consult or the nurse may move to the nurse's station to document details of the encounter. All these activities are linked to entities performing some type of movement in the environment.

Formally we can express this sequence of activities in terms of time as

(i) At time t_1: Patient arrives at the trauma unit and is sent to the trauma bay
(ii) At time t_2: The nurse and a resident check in on the patient
(iii) At time t_3: The resident seeks a phone consult while the nurse heads over to the station to continue with documentation.

In the figure, 'P' refers to the patient; 'N' refers to the nurse and 'R' to the resident on call. The black solid dots denote location of base stations ($B_{1\text{–}6}$). Base stations were placed at various key locations; one at each trauma bay, one near the phone and the other near the computer. For these given sequence of events, the following are the trends we see in the data derived from the tags.

(i) At time t_2: Tags R and N get close to B_1.
(ii) At time t_3: Tag N is very close to B_5 and Tag R is very close to B_6.

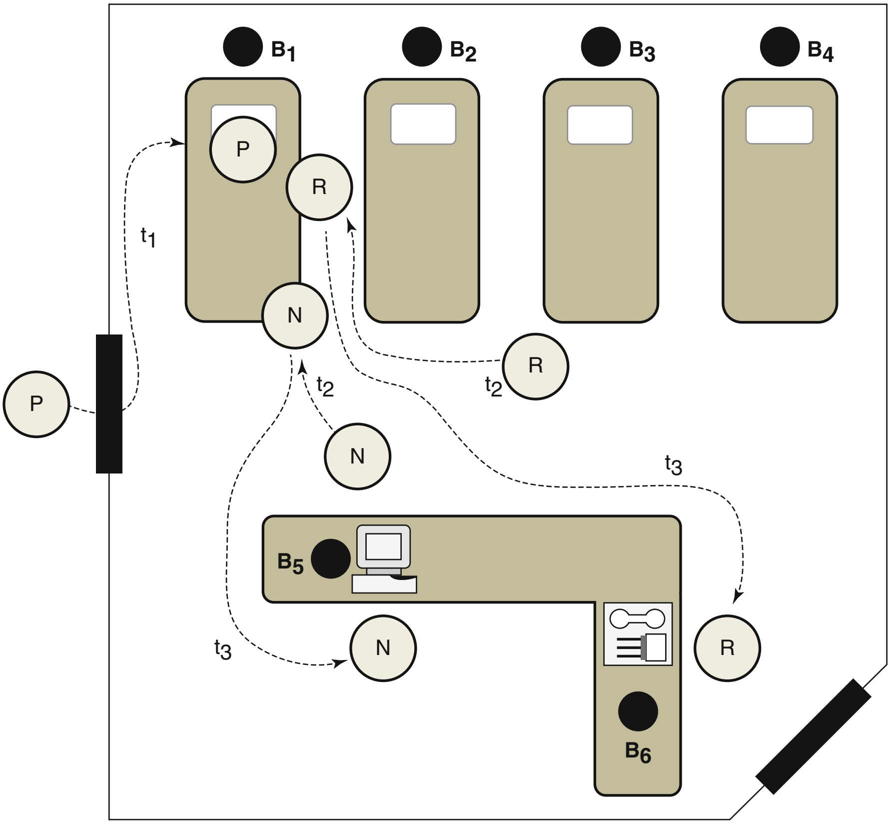

Fig. 17.2 Example scenario: patient arrival at a trauma unit

With the initial setup phase we know that B_1 is trauma bay 1, we can assume that the patient is being managed by the nurse and resident at time t_2 and that the patient arrived at the unit sometime before t_2. Therefore, at time t_3, the system can probabilistically estimate that the nurse was documenting the patient report, and the resident was seeking a phone consult. While the scenario presented is a simplification of the total process, it provides a conceptual view of how we can track activities through tags. In reality, activity models generated can be more complex. The models would be required to handle variations in activities performed while classifying them accurately.

Capturing Clinical Work Activities: Two Evaluation Studies

In the rest of this chapter, we describe two evaluation studies that describe the use of RFID technology. In the first study, sensor data is used to predict a set of clinician activities in a simulated trauma scenario. In the second, sensors are used to characterize the nature of clinician interactions and collaboration in an emergency care setting. Using both these studies we describe the potential scope of using RFID technology in the clinical work environment. We also discuss potential applications of its use in training, monitoring and administration in critical care settings.

Predicting Clinical Workflow from Through Automated Analysis

Many processes produce outputs that may be characterized as observable signals. In the case of RFID tags carried by clinicians, these signals are the discrete received signal strength values captured by the base stations. Hidden Markov modeling is a well-known method for characterizing real-world signals in terms of signal models [23]. The models can provide a theoretical description of the underlying system from which deviations from the norm can be identified.

Activity Modeling Using Hidden Markov Models

Hidden Markov Modeling (HMM) is a probabilistic modeling tool that is usually employed for temporal sequence analysis and has been effectively used in movement analysis, gesture and speech recognition applications. An HMM models a temporal sequence of events (called an observation sequence) in terms of a state machine, in which the current state of the model is probabilistically dependent on the previous states. A well-trained HMM activity model can detect the temporal activities that the HMM has been trained for.

As with any method, HMM based activity recognition has certain advantages and disadvantages. The key disadvantage of HMMs lies in the fact that the amount of data that is required to train an HMM is very large. Another issue with HMMs is that they require positive data to train with, i.e. in order to effectively train an HMM to recognize a class of activities, we require a carefully constructed training set that best describes the activity. However, these disadvantages are outweighed by a trained HMM's capability to handle variations in the final style of execution of an activity. Activities may be performed in a different manner in critical care environments and it is important that the model of activities accounts for these variations. By training the HMM system in a robust manner, it is possible to recognize the motion and some communication activities regardless of the deviations for our application. In addition, HMMs scale well as they can be trained to learn activities incrementally. New activities can be trained for without affecting models of previously learned activities. For these reasons, we chose HMMs for the development of activity models and activity recognition.

Activity recognition using HMMs is a two-step process. It involves (i) *training* HMMs for specific activity models and (ii) *testing* the HMMs for their recognition accuracy with annotated test samples. In order to develop robust activity HMMs, we first require data that describes the activity. This data is obtained from the RFID tags. More specifically, the data utilized is the RSSI value of each tag-base encounter gathered during data collection. We collect this data for the activities of interest in multiple samples. We utilize half of the samples for training the HMMs and retain the rest for testing the developed models. A database of samples for each activity facilitates training the HMMs for each activity, thereby creating a library of HMM activity models for each activity. The training of HMM activity models is achieved using the Baum-Welch algorithm.

Once a library of HMMs is built with one HMM for each activity, the developed models can be tested. The testing of an activity sample proceeds by firstly, estimating the probability that the sample movement belongs to the library. This is achieved using the Forward-Backward procedure for each of the HMM's in the library. The HMM that yields the highest probability for the test sequence is determined to be the type of activity that the movement sequence belongs to. The accuracy of recognition is measured as the ratio of the number of correctly identified test sequences to the total number of test sequences. In this manner, activity models are developed and tested for activity recognition.

Data Collection

We collected two sets of data: (a) Qualitative data from observers, and, (b) Quantitative data gathered from the RID tags.

Both the qualitative data and quantitative data are obtained from standardized sources. While time-stamped quantitative data is retrieved from the RFID tags, observations were generated by observers using an activity tracking software tool. The tool contains a list of commonly occurring activities for the Nurse and Physician. The activities chosen were based on an ontology developed by Zhang and colleagues based on their prior work on analyzing the workflow in emergency departments [24]. Observers may select an appropriate activity from the list provided and add detailed comments a description text box. The observations are then automatically dated and timed and stored in the output observation file. In this way time-stamped data is obtained for both qualitative and quantitative data sources. This makes synchronization of the two data streams possible.

Quantitative data is obtained using *active* RFID tags to gather data. The tags record encounters with other tags (tag-tag encounter) and base stations (tag-base encounter). For each encounter or interaction, the tags record (a) identification number of the tag or base station detected, (b) time and date of encounter, and (c) the received signal strength indication (RSSI) value.

In order to test the HMM based activity recognition system, we simulated 15 Trauma activities (listed in Table 17.1) in a lab setting, (depicted in Fig. 17.3) with

Table 17.1 Activity list and corresponding clinical descriptions

Activity	Movement	Clinical description
A1	1-to-2	Paged physician/nurse tends to patient on bed 1
A2	2-to-3	Physician/nurse moves to treat patient on bed 2
A3	3-to-4	Physician/nurse leaves trauma through entry/exit 1 after visiting patient on bed 2
A4	4-to-5	Physician/nurse enters trauma through entry/exit 1 and attends to the phone
A5	5-to-6	Physician/nurse after attending to a phone call move to use the computer at the nurse station
A6	6-to-1	Physician/nurse leaves Trauma through entry/exit 2
A7	1-to-4	Physician/nurse enter and leave trauma
A8	4-to-6	Physician/nurse enter trauma through entry/exit 1 and move to use the computer at the nurse station
A9	6-to-2	After using the computer physician/nurse move to treat patient on bed 1
A10	2-to-4	After visiting patient on bed 1, physician/nurse leaves trauma through entry/exit 1
A11	5-to-1	After attending a phone call, physician/nurse leaves trauma through entry/exit 2
A12	1-to-3	Paged physician/nurse attends to patient on bed 2
A13	3-to-5	After visiting patient on bed 2 physician seeks a phone consult
A14	5-to-2	After completing a phone call physician/nurse moves to treat patient on bed 1
A15	3-to-6	After treating patient on bed 2 physician/nurse move to use the computer at the nurse station

ten tags and six base stations. These activities were simulations of clinical activities. In order to simulate potential activities in a lab setting we observed commonly occurring movement tasks in the Trauma unit, an example being “physician moving to phone for a consult” (Activity A13). Figure 17.3 depicts the lab setup for testing.

The setup for the testing involved the creation of a 20 ft by 20 ft grid in a lab setting. Six base stations (depicted by black solid circles) we placed in predefined locations (Base 1 and 4 at Entry/Exit points 2 and 1 respectively; Bases 2 and 3 at Beds 1 and 2; Base 5 at the phone on nurse station; Base 6 at the computer on the nurse station). This is congruous with base station setup in the real-world scenario. We gathered movement data for the 15 sample activities listed in Table 17.1. For each RFID tag-base pair or tag-tag pair an encounter is recorded every 3–4.5 s. This data is captured in a time-modulated manner, i.e., encounter information is communicated by detecting differences in the time of the encounter rather than the frequency. This results in a sparse matrix when considering the entire tag-base station configuration. Figure 17.4 depicts a sample of the matrix generated. The encounter of a tag X with base stations A, B and C (gray filled boxes) are shown in a 60 s long timeline. We use linear interpolation to fill missing data in this sparse matrix. While this methodology provides an RSSI value for all base stations at all instances, it adds some noise to our system that may affect the overall activity recognition accuracy.

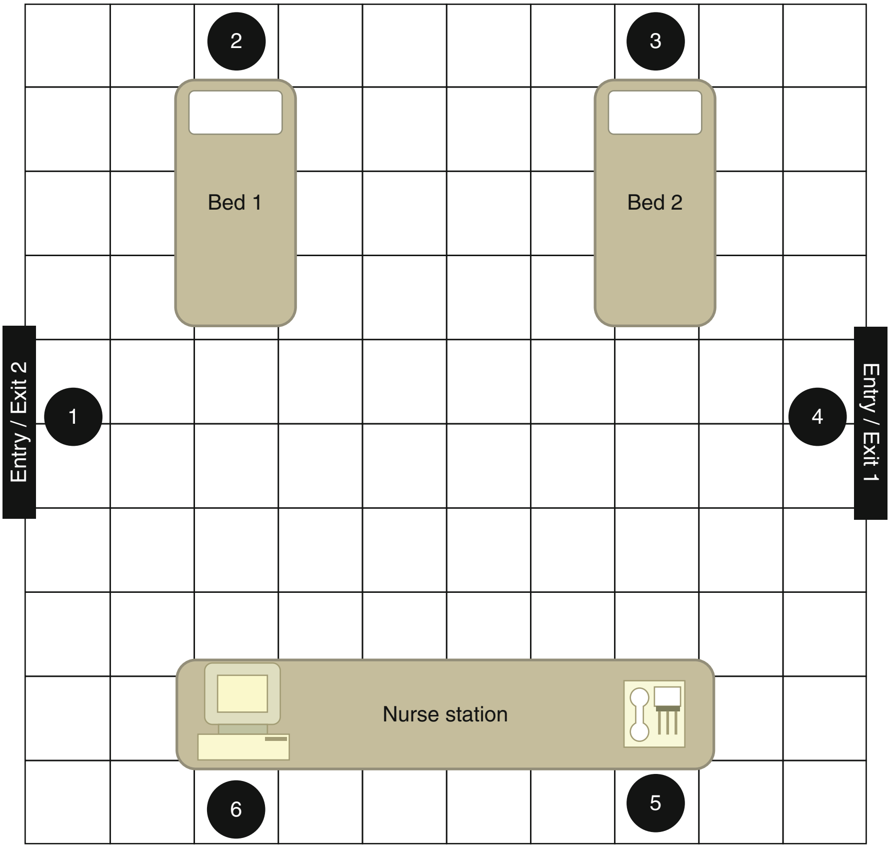

Fig. 17.3 Test setup for simulated clinical activities

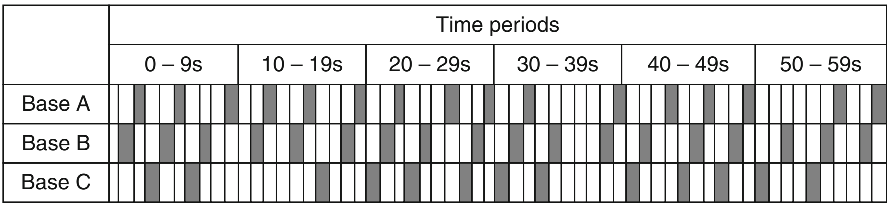

Fig. 17.4 Sparse matrix of tag-base encounters (*gray* fill indicating an encounter record with some tag)

For each of these activities, we gathered ten samples of data. Each sample involved a tagged entity (researcher) following the movement pattern prescribed for the activity. Each sample performed with ten different tags, totaling 100 samples for each activity. This ensured sufficient randomization of activity movements, accounting for inter-tag variability as well. A total of 1,500 samples (15 activities × 10 samples × 10 tags) were gathered for testing. Out of the 100 samples

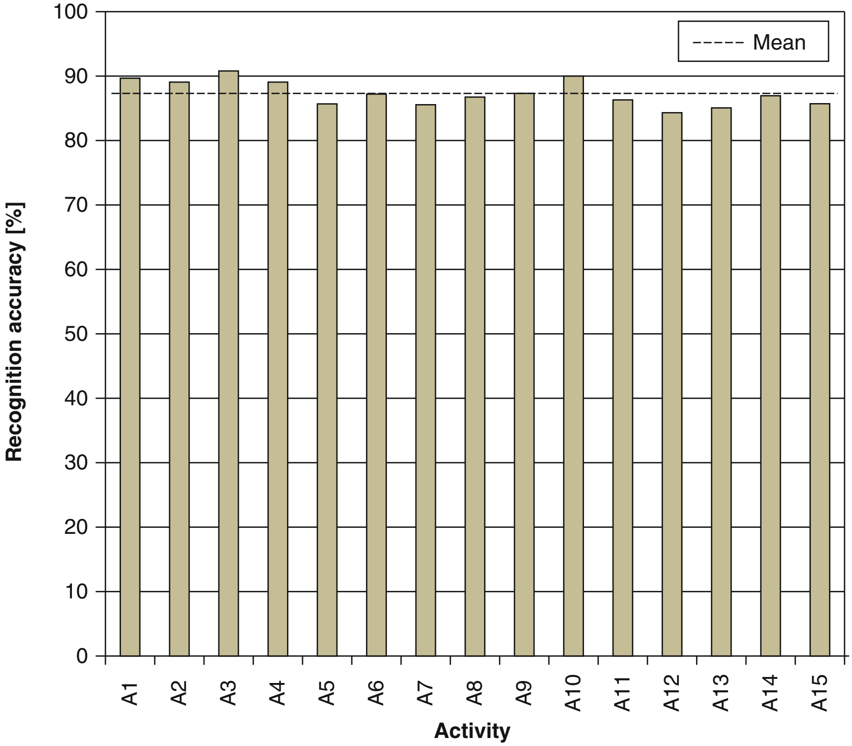

Fig. 17.5 Hidden Markov Model (HMM) based activity recognition results

gathered for each activity, 50 samples were used to train the HMM for activity recognition, and the other 50 were used as a testing set to evaluate the algorithms' accuracy.

Results of HMM-Based Evaluation

Figure 17.5 summarizes the recognition accuracy for the 15 motion patterns (A1–A15). Recognition accuracy is the ratio of the number of activities correctly identified to the total number of activities used for testing. A mean recognition accuracy of 87.5 % was obtained, with a maximum of 90.5 % and minimum of 84.5 %. The analysis of the incorrectly classified test samples revealed that misclassifications were a result of variations in the training set. As discussed previously, HMMs require to be trained on a controlled sample that best represent the activity. Obtaining training data from real-world scenarios are likely to have variations that may compromise the quality of models generated. This is a limitation of utilizing HMM models with RSSI values alone for activity recognition. Additional sensors such as accelerometers could be utilized in conjunction with RFID tags to improve the activity recognition rates.

Summary

RFID sensors were used to record of motion and location of clinical teams, which was used to model activities in critical care environments. A HMM model was developed to identify a set of 15 simulated clinical activities with 87.5 % accuracy. We found that RSSI values, as the only observable signal, were insufficient in identifying activities with the necessary levels of accuracy. With the use of additional sensors such as accelerometers it would be possible to counter the noise levels present in RSSI signals.

Tracking Clinicians During Emergency Care Activities

In many respects, the critical care workplace resembles a paradigmatic complex system: on account of the dynamic and interactive nature of collaborative clinical work, these settings are characterized by non-linear, inter-dependent and emergent activities. Developing a comprehensive understanding of the work activities in critical care settings enables the development of streamlined work practices, better clinician workflow and most importantly, helps in the avoidance of and recovery from potential errors. We used sensor-based technology to capture the movement and interactions of clinicians in the Trauma Center of an Emergency Department (ED). Remarkable consistency was found between sensor data and human observations in terms of clinician locations and interactions. With this validation and greater precision with sensors, ED environment was characterized in terms of (a) the movement patterns of clinicians, (b) interactions with other clinicians and finally, (c) patterns of collaborative organization with team aggregation and dispersion.

Study Setting

The study was conducted in a certified Level 1 Trauma Center in the Emergency Department of a large teaching hospital located in the United States. The hospital provides 24/7 emergency and trauma care to approximately 52,000 patients a year. The ED is separated into distinct units caring for pediatric patients, general medicine patients and those requiring trauma care. The physical set-up of the trauma side of the ED includes eight trauma patient beds and five urgent care beds. In times of high patient volume, additional chairs and beds are placed in the open spaces as needed. The care team for trauma ED typically includes one attending physician, two resident physicians and two trauma nurses, an urgent care nurse, a charge nurse, one technician, and a respiratory therapist shared by the entire ED. The trauma center is also supported by a dedicated trauma team, consulting physicians and the staff from other units of the ED (including off-service providers) as needed.

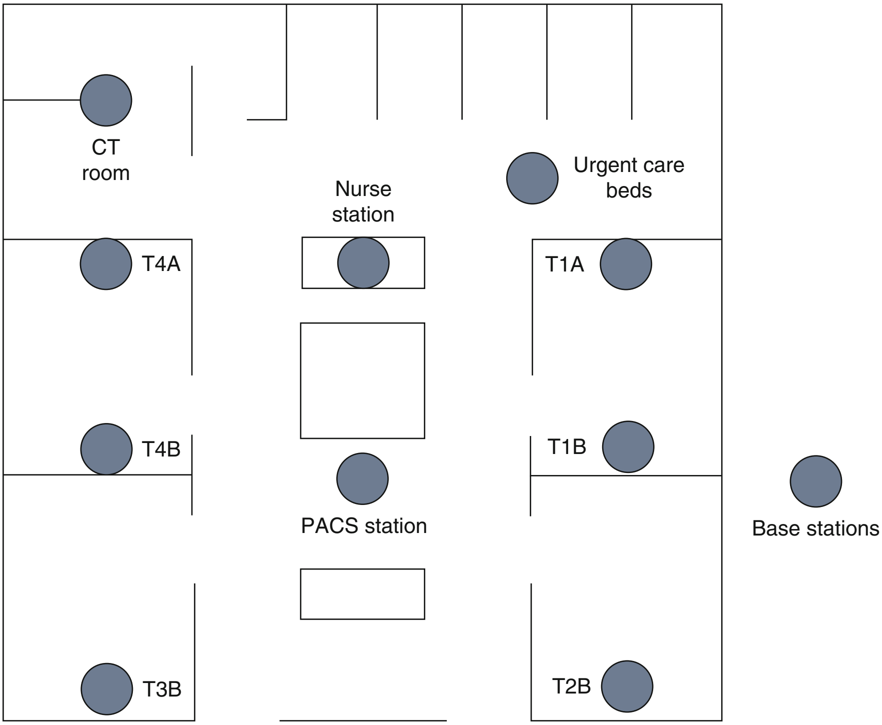

Fig. 17.6 Spatial orientation of the base stations in the ED. Each circle represents a base station at that location. The locations were CT room (*CT*), Nurse station (*NS*), Image browsing station (PACS), Trauma bed 1A (*T1A*), Trauma bed 1B (*T1B*), Trauma bed 2B (*T2B*), Trauma bed 3B (*T3B*), Trauma bed 4A (*T4A*), Urgent care beds

Participants

Observation and tagging occurred on four separate shifts over a 2-month period at the trauma center. During each observation session, the attending physician, two resident physicians, and two trauma room nurses, were solicited for participation. Informed consent was obtained from all participants before the start of each data collection session. Participants were instructed to go about their usual activities.

Sensor Setup

A total of ten (10) base stations were placed across the trauma rooms, physician station, nurse station, CT room and urgent care rooms. The tags were distributed among the attending physician (1), residents (2) and nurses (2). The sensor data included the tag-tag and tag-base pings along with their corresponding signal strength and time-stamp. Sensor data on the tags and base stations was then formatted and uploaded to a MySQL database server. The spatial orientation of the base stations is shown in Fig. 17.6.

One of the critical factors in effectively using the sensor technology is the calibration of the sensors to filter "good" signals from noise. Prior research has used a variety of mechanisms to filter the sensor signals. In general, *threshold signal strength* is often established as a baseline measure. In our experiments, we used a RSSI signal strength value of −70 dB (decibel) as our cut off signal strength. This value was based on the manufacturer's specification and our calibration tests verified this threshold.

Shadowing

In order to validate and complement the information provided by the sensor data, human observers shadowed the "tagged" clinicians. The purpose of shadowing the clinicians was twofold: first, to confirm the accuracy of the location estimations made by the tags and second, to get additional information on the activities of clinicians. The attending physician was shadowed for two sessions, while in the other sessions, a resident and nurse was followed.

To assist observers with their shadowing tasks, we used the UObserve suite of data logging tool [25]. UObserve is a mobile platform that provides researchers with the ability to conduct field observations using standard templates to ease data collection, and importantly the capacity to precisely record the time of recorded events. The UObserve tool is based on the work domain ontology of the ED environment. The use of UObserve allowed for precision and ease in capturing events (e.g. time, place, participants, activities) and synchronization with the tagging data. For this study, observers were provided with a version of UObserve, which had a list of ED-specific locations (based on the base-station locations) and collaborating clinicians at that location. At every instance when the tagged subject changed location, the observer noted the location on the UObserve tool. Additionally, other clinicians who came in direct contact with the shadowed-clinician were also noted.

For each location selection, a time-stamp was automatically added by the system. This time-stamp was synchronized with the time-stamps on the sensors. The data from UObserve was uploaded from the mobile device to an encrypted server. A companion application was developed to export the data in customizable data formats. A sample screen shot from one template in the UObserve interface is shown in Fig. 17.7.

Data Collection

Four data observations of the core trauma care team (1 attending, 2 residents, 2 trauma nurses) occurred over a 2-month period. One clinician was shadowed per session by an observer. Prior to collecting the data in the ED, the tags and

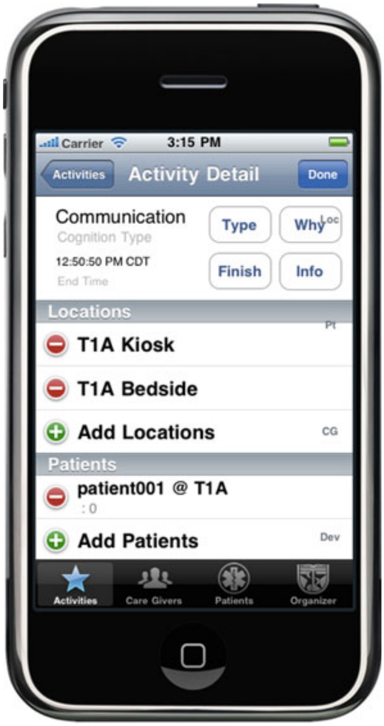

Fig. 17.7 UObserve iPhone interface with the location details and activity template is shown

base stations were extensively tested in the laboratory and in the ED (in pilot experiments) to ascertain their accuracy and effectiveness. During each of these sessions, both sensor and shadowing data were captured. On average, each of the data collection sessions lasted about 3 h (mean = 3.2 h, s.d. = 0.14 h) and was conducted from the start of attending shifts during both afternoon and night periods. While all five team members wore RFID tags, only selected team members were shadowed. Clinicians varied across the sessions.

Data Analysis

In this section, a detailed explanation of the various measures that were used to analyze the sensor and observation data is provided. Particular attention is given to the manner in which the data from the sensors are extracted, processed and analyzed. We specifically investigate two characteristics of clinician activities: *movement of clinicians* and their *interactions* with each other. Based on these two specific characteristics, we investigate the following: time spent at a location, time spent with other clinicians, transition between various locations and collaborative work activities.

Time Spent at a Location and in Proximity to Other Clinicians

Collaborative work is often done within the specific context of location and people. By ascertaining the location of a clinician and subsequently the time spent at that location, it is possible to make preliminary judgments on the work activities of the clinicians.

The location of a clinician is determined based on the tag-base pings and the shadowing data. For determining the time spent by the clinicians at a location, we use the tag-base ping events that were retrieved from the base stations. The time spent by a clinician in proximity to a base station is determined by aggregating the tag-base pings at each identified base station with the highest threshold signal strength value at that particular time. Like time spent in a location, time spent in proximity to others is measured by pings over the threshold response level. Unlike time at a location (tag-base pings), time spent in proximity to other clinicians is computed as an aggregate of the tag-tag pings. If there were multiple tag-tag pings at a particular time, then all possible pairs of tag-tag pings were aggregated for this computation.

Transitions Between Locations

One of the ways to investigate the workflow of clinicians is to trace the movement patterns of the clinicians. As explained earlier, work activities are often context (and location) dependent. In other words, locations can be used as a general proxy for certain types of activities. For example, the presence of an attending or resident at a trauma bedside can be considered as a "patient care" activity. Similarly, a physician at a physician workstation can be construed as the physician performing a documentation task. On account of the hands-on nature of clinical work in this setting, transitions between locations provide a preliminary account of the workflow in a collaborative setting. For example, the movement of the attending physician across various locations within the ED over the period of a shift can be used to gauge their work pattern. If the attending physician was at their workstation for most of a shift, then we can make predictions about the low degree of activity during that shift. In contrast, if there is significant amount of movement by the attending physician across various trauma rooms, then we can make predictions about the high degree of activity during a shift. While these examples are extreme scenarios, it is important to note that transitions between different locations can be used as a basis for determining the nature of activities in the ED. In short, the transition between locations provides a trace-based illustration of the workflow.

In order to develop the transitions between locations in the ED, we identified ten locations in the ED where the base-stations captured significant signal strength. These locations were: CT Room (CT), Nurse Station (NS), Image Browsing Station (PACS), Trauma Bed 1A (T1A), Trauma Bed 1B (T1B), Trauma Bed 2B (T2B), Trauma Bed 3B (T3B), Trauma Bed 4A (T4A), Urgent Care Beds. Based on the tag-base pings at these locations, we first developed a transition probability matrix of location transitions for each clinician.

A location-based transition probability matrix represents the transitions between a set of selected locations. Each cell in the matrix represents the total count of the transitions between the two locations. For example, if the cell value between the CT room and the Nurse's Station for the attending physician was 25, it means that the physician moved from the CT room and the Nurse's Station a total of 25 times during the shift. The transition probability matrix is also often referred to as an *antecedent-consequent* matrix, since it provides the counts of the number of transitions between the antecedent and consequent events. We developed a 10×10 matrix for the location-based transitions (for the ten locations described earlier in this section) for each of the clinicians, per session.

In order to develop the transition probability matrix, we first filtered the tag-base pings that were above the threshold value. Using a sliding window with an interval of 15 s, we temporally collected the locations of all clinicians within this time-window. The location with the highest RSSI strength per clinician was then separated out. This process was applied to the entire data set till all locations of all clinicians were obtained over their entire shifts. The temporal sequence of locations was then converted into a matrix of location-based transitions for further analysis.

Collaboration: Aggregation and Dispersion

Highly complex environments are often characterized by collaborative interactions to maintain the continuity of work activities. The collaborative interactions can be characterized in terms of three key concepts: the *size* of the collaborating team of clinicians, the *length* of their collaboration and the *location* at which the interactions of the team occurs. The knowledge of these three concepts is useful in developing a "blueprint" of the collaborative activities within the ED. We use the tag-tag pings between clinicians to estimate the collaborative interactions between them. Using physical proximity as an indicator for interaction, we identify the following: first, the pair-wise interactions between all the clinicians and the locations at which these interactions take place were identified (based on tag-base pings). Then, the location and size of the largest group of clinicians are detected using matrix-based algorithm.

While, we use the term "interaction" in a general sense, meaning physical proximity between clinicians, it can be argued that close proximity at a particular location in an emergency care setting (e.g., at a specific trauma bed) would indicate that the clinicians are together for a common purpose or goal (e.g., care for a patient at a location). Thus, even though the clinicians may not be verbally communicating with each other, a common goal of being at the same location can be considered as a measure of a shared collaborative activity. We use this concept to measure the degree of team aggregation and dispersion in the ED.

As explained earlier, we first identify the pair-wise interactions between all pairs of clinicians. For the sensor data, we focus primarily on the pair-wise interactions of the attending physician, as they are central to controlling the workflow in the ED. For ascertaining the pair-wise interactions, the sensor data was first "chunked" into intervals of 30 s, after testing with intervals ranging from 30 to 180 s. To be considered as a "valid tag-tag ping" at a particular location several conditions were first evaluated. We describe these conditions with an example. Consider two tags, tag1 and tag2 at a location B1 (base station location). A valid tag-tag ping between these two tags would involve the following interactions: tag1-tag2 ping, tag2-tag1 ping, tag1-B1 ping and tag2-B1 ping. Additionally, all these pings have to occur within the selected 30-s interval.

After obtaining the pair-wise interactions (and their locations), we evaluated the formation (aggregation) and dissipation (dispersion) of larger clinician groups. The identification of large groups was progressively more complex than the pair-wise comparisons. Since groups (size > 2) take longer time to form (and disperse), we considered time intervals of 100 s for this analysis. The time period of 100 s was arrived after testing with various "time-chunks", discussions with ED attending physicians and our own observation data. Based on our observation data and discussion with ED clinicians, we evaluated the average group formation (for groups of different sizes) time across each shift. A 100-s interval was found to be an appropriate time-span for capturing the formation (and dispersion) of groups of sizes varying from two to four. The groups were ascertained in the following manner: first, the presence of a group within the considered time interval was determined. Second, it was verified whether the interactions were occurring within the same location. We explain the aggregation algorithm with an example.

For every 100-s interval that we considered, we developed a two-dimensional matrix similar to the one shown in Fig. 17.4. There are two types of information that is encoded in the matrix: the tag-tag interactions (represented as a binary operator between tags T1–T5 in the left half of the matrix) and the tag-base interaction (represented as a binary operator between base stations B1–B10). From the example matrix (see Fig. 17.4), we generate all possible tag-tag interactions. In this case, the only tag-tag interactions are with tag 1 (T1) with (T2 and T3). The interactions of all other tags (T2, T3 and T4) are with only with T1. Thus, the direct interactions in this period of time are {T1, T2, T3}. Next, we investigate the reverse tag pings (i.e., from T2 to T1, T3 to T1, etc.). For this, we evaluate the column values for T1: {T2, T3, T4}. The intersection set between direct and reverse set of tag-tag pings gives us the set of tags that were interacting in this time period. In our case, we get the set of tags as {T1, T2, T3}. This means that the clinicians carrying the tags T1, T2 and T3 were in close physical proximity to each other.

The last step in the algorithm is to establish the location where the clinicians were together. For this, we use the identified set of tags and compare it with the common set of locations at which these tags were present. In other words, we explore the columns for the base stations (B1–B10) that have non-zero values in the cells for the set of identified interacting tags. In the case of the example provided, the only location where the base station has non-zero value is for the column

	Tags					Base stations									
	T1	T2	T3	T4	T5	B1	B2	B3	B4	B5	B6	B7	B8	B9	B10
T1	0	1	1	0	0	1	1	1	1	0	0	0	0	1	1
T2	1	0	0	0	0	0	1	0	0	1	0	0	1	1	0
T3	1	0	0	0	0	0	1	1	0	1	0	1	0	0	1
T4	1	0	0	0	0	0	0	0	1	0	0	0	0	0	0
T5	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0

Fig. 17.8 Mechanism for identifying group aggregation and dispersion. The matrix has two components: the tag-tag interactions (between tags T1–T5) and between tags and base stations (B1–B10). The presence of a tag-tag or tag-base interaction between tags/base stations is denoted by 1 in the corresponding cell

pertaining to B2 (see Fig. 17.8). Consequently, a group will only be considered as such if all members ping one another, as well as the location base station during the same 100-s time period. Thus, we can identify the largest group during this time period as {T1, T2, T3} at location {B1}. The highest signal threshold values were taken into consideration if there were multiple possible locations for the identified group. There was less than 5 % incidence of multiple locations for a group across all sessions. We computed the size of the largest group for every 100-s interval for the all the four sessions.

Results

In this section, we report on the results from the sensor and observation data. First, we validate the correlation between sensor and observed data. Based on this validation (i.e., the plausibility of using tags as a data collection mechanism), we investigated the relative entropy of the ED system. Then we report on the workflow of the ED clinicians based on their location transitions and interactions with other clinicians. Finally, we describe the formation and dispersion of teams as a measure of collaboration in the ED.

Validating Sensor and Shadowing Data

In order to evaluate the degree of association between the sensor and shadowing data, we computed the correlation between these data sets for both mobility and interactions among clinicians. A high correlation between the sensor and observed data validates the accuracy of the sensor data in capturing the location and interactions among the clinicians. We computed the Pearson moment-correlation between the location determined by the sensor data and location determined by human observers. We obtained a statistically significant correlation between the observed and sensor-based location data ($p < 0.01$, $R = 0.96$) (See Fig. 17.9a).

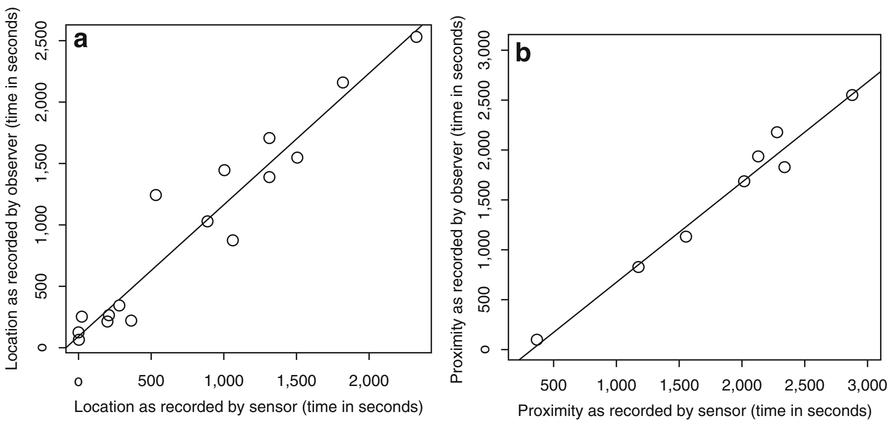

Fig. 17.9 (**a**) Correlation between location as determined by the sensor and location as determined by the shadowing observer over time and (**b**) correlation between interactions between physicians as determined by the sensor and as determined by the shadowing observer over time (across all four data collection sessions)

Similarly, we also computed the correlation of proximity between the clinicians as determined by the sensors and shadowing observer. Based on Pearson product-moment correlation, we found significant correlation between co-location of the physicians as determined by the sensors and by the observers (*R* = *0.98, p <0.001*) (See Fig. 17.9b). In other words, physicians (attending and the two residents) were more likely to be co-located than the nurses. The inherent lack of co-location of nurses can be attributed to the significant percentage of nurse activities are often performed in isolation from other physicians (e.g., documentation, care coordination). Hendrich et al. [26] reported similar results where they found that that nurses spend significant amount of their time at nurse stations performing documentation and care coordination activities. The mobility and interaction correlations were computed from data across all sessions.

The significant correlation between the sensor and observed data provides an initial validation for the accuracy of the sensor data in capturing the location and interactions of clinicians in the ED. A comprehensive knowledge about the location and interactions is instrumental in real-time monitoring of emergency environments. Such monitoring can provide useful insights into the activities around specific events such as arrival of a patient with severe acuity or a mass emergency event (e.g., a train accident) and for the study of errors. These concepts are further explored in the discussion section.

Time Spent at Locations and in Proximity with Other Clinicians

Based on the tag-base pings, we computed the time spent by the clinicians in various ED locations. As described earlier, the time spent was computed based on

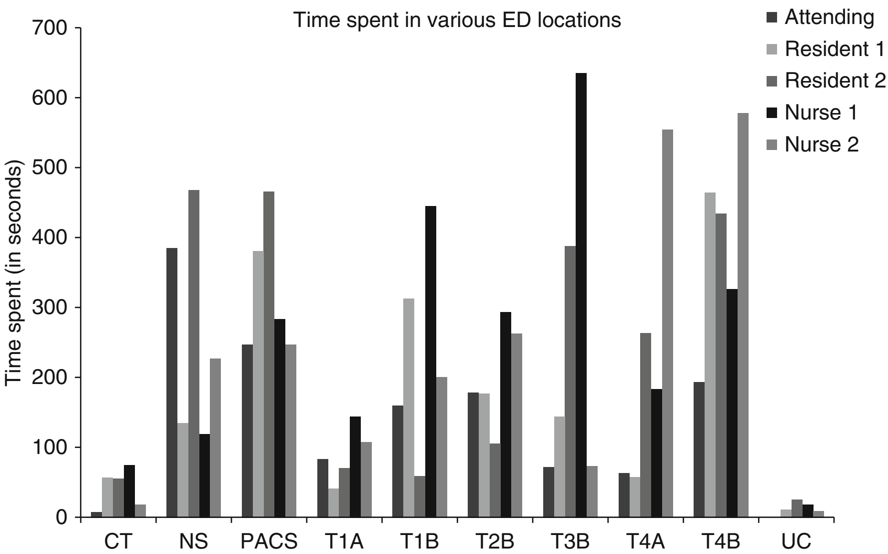

Fig. 17.10 Time spent by all clinicians at various ED locations across all sessions. The locations are *CT* CT room, *NS* Nurse station, *PACS* Image browsing station (and physician station), *T1A* Trauma bed 1A, *T1B* Trauma bed 1B, *T2B* Trauma bed 2B, *T3B* Trauma bed 3B, *T4A* Trauma bed 4A, *UC* Urgent care beds

the aggregation of tag-base pings at each location over time. Figure 17.10 shows the time spent by the clinicians at the various locations in the ED. The x-axis shows the different ED locations (same as those marked up in Fig. 17.2) and y-axis is the time spent at each location in seconds. From Fig. 17.10, we found that: clinicians spent most of their time in the trauma rooms (at the various trauma beds 1A, 1B, 2B, 3B, 4A and 4B); the residents and nurses spent significantly more time in the trauma rooms (i.e., beside the patients) than the attending physician. This is primarily a function of the care process in large teaching hospitals where residents (along with the support of nurses) manage the care process under the supervision of the attending physician.

In a similar manner, we also computed the time spent by the attending physician with other clinicians based on the tag-tag pings. We found that the attending physician spent considerably more time with other physicians (residents) compared to time spent with nurses ($p < 0.01$). This was expected considering as the study was conducted at a teaching hospital.

Transition Between Locations

In order to investigate the clinician workflow we traced the transitions between various locations by the clinicians. The transitions were determined based on the transition probability matrices. Figure 17.11 shows the counts of transitions between

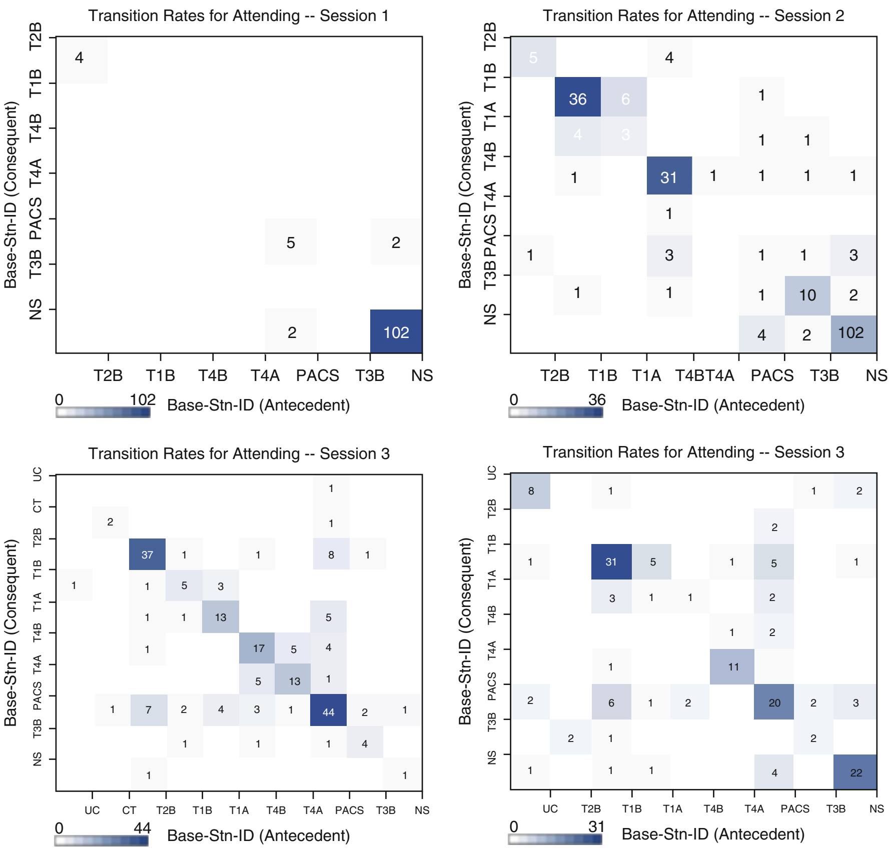

Fig. 17.11 Location transition matrix for the attending physician from the four sessions. The x-axis shows the originating location and the y-axis shows the terminating location during the transition. The counts in the diagonal matrix shows the instances where the attending physician did not move in consecutive time intervals

various locations by the attending physician in the four sessions. The x-axis represents the originating location and y-axis represents the terminating location for each transition. The diagonal of the matrix represents instances where the attending physician was in the same location for consecutive time intervals. Significant differences in the transition patterns can be gleaned from the analysis of the four graphs. In session 1, the attending physician was fairly sedentary at the nurse station (NS). This was probably due to a relatively slow shift.[1] In sessions 2–4, we can see that the attending physician moved across the various trauma rooms and had a "foot print" across all the locations in the ED. It can also be observed that a significant

[1] In fact, our observation data shows that during this session, the attending physician spent a considerable portion of this slow shift teaching the residents at the Nurse's station.

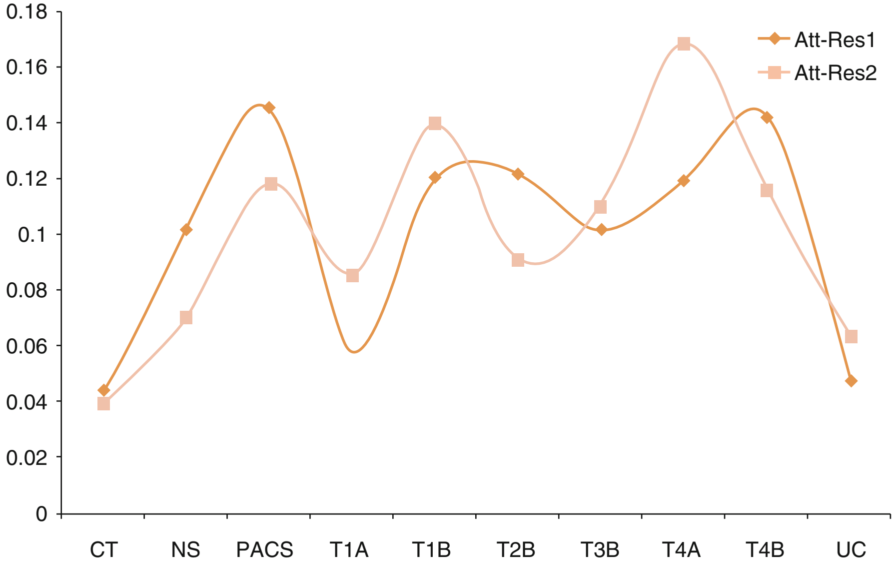

Fig. 17.12 Pair-wise co-location probability between the attending and the residents

time was spent in the trauma rooms (darker squares in the cells representing the trauma rooms).

We also developed similar location matrices for other clinicians. In the case of residents, we found that the transition pattern of one resident was complementary to the other. In other words, we found that, one resident was invariably present at a set of trauma rooms (and absent from the rest of the trauma rooms), while the second resident was present at the remaining trauma rooms. This is consistent with the demands of their shared workload and division of patient care duties. This is further investigated in the next section on collaborative patterns. We found no consistent patterns in the location transitions among the nurses.

Collaboration: Aggregation and Dispersion

We computed all pair-wise co-occurrences between the attending physician and other clinicians. As expected, we found consistent co-location of the attending and the residents in the trauma rooms. This was further confirmatory evidence for the likely complementary role that each of the residents took for the patient care activities. In other words, we found that one resident had a prominent "role" with respect to the treatment of a specific patient. This can be seen in terms of the pair-wise co-location probability (see Fig. 17.12) where one resident is more likely to be present along with the attending physician in a trauma room. The high co-location probability of one resident was highly correlated with a low co-location probability of the other resident being in the same trauma room. We did not find any consistent patterns with respect to the co-location between nurses and the attending physician.

While the interaction between pairs of clinicians is interesting, complex settings are characterized by a significant amount of collaborative activity. Consequently, we were interested in the behavior of the team as a unit, in addition to that of individual clinicians or clinician pairs. We investigated the formation and dispersion of larger groups (>3) in the ED. Based on our algorithm described earlier we computed the team size dispersion over the data collection sessions. On average we found that there was a high percentage of two and three-clinician groups across all sessions. We found several interesting patterns with respect to the aggregation and dispersion of teams across the ED.

First, the incidence of larger clinician groups (4 and above) was very low. On average, there were less than 15 such group occurrences. These clinician groups always included the physician, both residents and one of the nurses. The low occurrence of the larger groups was probably due to a combination of factors: first, such large groups would entail the majority of the care team. From observations, we know that these large groups typically come together during a major trauma and quickly disperse to care for the other patients in the ED center. During occasion of lower patient volume, large groups might congregate in central locations with team members entering and exiting freely. These circumstances of high demand and low volume are relatively infrequent. Second, our algorithm that determined the presence of teams was extremely stringent in terms of the requirements that ascertained the presence of a group (multiple tag-tag and tag-base pings within a short interval). While, this may ignore extremely slow forming groups, we believe that the ED is an extremely fast-paced environment where the formation and dispersion of groups are in response to rapidly emerging situations.

Second, larger clinician groups (size greater than or equal to four) always congregated in one of the trauma rooms. This is highly likely in ED settings where the arrival of a patient with high acuity levels triggers significant activity around that patient. While, we cannot directly verify the acuity of the patient at the times where the larger groups congregated, in our future work we plan to retrospectively investigate the arrival acuity levels of patients for the sessions in which we collected sensor data. Third, team size of three almost always (90 % of the cases) involved at least one resident and a nurse. The third participant in such three-person groups was either the resident or the attending physician. About 60 % of such three-person groups were formed in the trauma rooms, while the rest were primarily split between the nurse station (NS) and physician station (PACS). Two clinician pairs were very common and we found significant variability among these pairs. But, about 50 % of the two-clinician groups identified consisted of the physician and one of the residents. This is typical considering the dual role of the attending physician in patient care and medical education.

An example of how the overall size of the largest ED team changes over a data collection session is shown in Fig. 17.13. The x-axis shows the time while the y-axis represents the size of the largest group at that point in time. As can be seen from the figure, the size of the group varies between 2 and 3 and for a short time a group of size 4 congregates together.

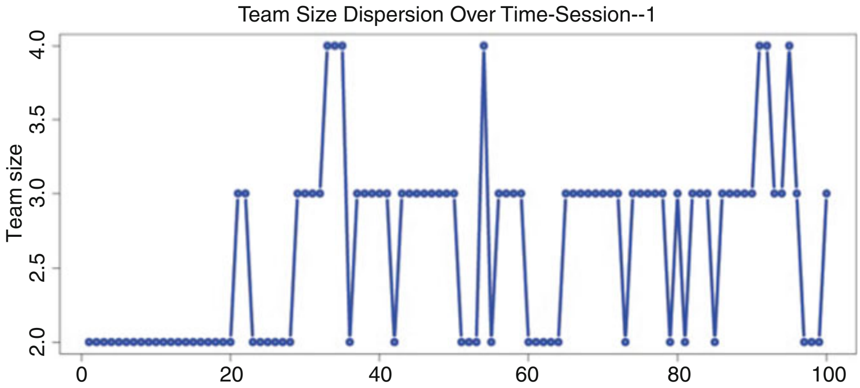

Fig. 17.13 Team size dispersion across a data collection session (session 2). The x-axis shows the time distribution (i.e., the time over a data collection session). The y-axis shows the size of the team

Discussion

We used RFID sensors to simulate and predict workflow and to capture the work activities of clinicians in critical care settings. The results from two studies reported in this chapter show the *appropriateness of using sensors to study work activities in complex critical care environment.* While, we used limited data collection sessions, our results provide significant support for more extensive use of sensors for studying complex activities. Though human observers are definitely required to collect highly nuanced information about the activities in complex environments, sensors are a reasonably reliable complementary data collection mechanism. Combining sensor data with other readily available clinical information (such as patient arrival information, condition, acuity, etc.) can help in developing flexible mechanisms for monitoring and managing the resources of complex environments. We further describe potential applications and uses of sensor technology including its role in visualization and training, management of resources and tracking of errors in critical care environments.

Visualization of Workflows

Visualizing workflow in 3D enables researchers and clinicians alike to easily grasp the activities that make up the workflow. In addition to enabling researchers review workflow in a novel way, the configurable virtual reality (VR) visualizations can also be employed for educational purposes. For example, a resident would be able to go experience a trauma from the perspective of the attending or nurse. This kind of configurability would enable the cross-training of clinical teams.

Fig. 17.14 Virtual trauma unit for workflow visualization

The visualizations can also be used to educate clinicians by illustrating cases of optimal workflow in relation to error-prone workflow.

In the domain of healthcare, virtual reality has been used to develop simulations for training of cognitive and psychomotor surgical skills and clinical decision making skills [27–29]. However, there is a lack of VR-based solutions for visualization of workflows and error scenarios even though such systems may have a major role to play in error prevention and mitigation. We can employ online VR environments such as Second Life® (http://secondlife.com/) and Active Worlds® (http://www.activeworlds.com/) for such visualizations. In this stage of the work, we have developed a standalone system that could be employed for such visualizations employing an open source gaming engine called Irrlicht (www.irrlicht.net).

A sample virtual trauma unit (see Fig. 17.14) was developed to mimic the trauma unit at Banner Good Samaritan Medical Center, which is the site of development for the project. The virtual trauma room consists of four trauma pods or beds. The nurses' station faces the trauma pods. A computer and phone are key components that are included in the design of the nurses' station. Two exit doors are present in either side of the trauma room. These details are synchronous with the test and real world set up. The current simulation contains three basic characters – the patient, resident and the nurse. The number and type of models to be utilized depend on the entities studied in the real-world. Models of the characters are built using modeling software (Maya and 3dMax; http://usa.autodesk.com/). Once the models are developed they can be controlled in the simulated world programmatically.

In order to obtain VR simulations of the workflow, the system generates a list of activities making up the workflow. These activities are then manually fed into the visualization engine to create the simulations. Currently, this stage of visualization process is completed offline. VR simulations created in this manner present a simulated view of real-world events. This is valuable to clinicians and researchers in highlighting the main events in the workflow within the context of the clinical environment.

Recent research [7] has reported on the potential of online 3-D virtual environments for medical education and learning. Online virtual environments provides an informal environment in which the learners can understand the norms, practices and challenges of working in a complex environment and integrate such information through repetition and group interactions.

Real-Time Monitoring of Activities and Resources in the ED

Sensor technology has been significantly useful in the remote and real-time monitoring of activities in various environments such as nursing activities, elderly care and telemedicine. Monitoring and management of resources in a highly dynamic and complex setting requires significant amount of data with respect to the activities and happenings within that setting. Data from the sensors (both mobility and interaction) provide information regarding the clinician (in terms of their location and co-location with other clinicians) with great precision and detail. Additionally, this information is time-sequenced. As a result, a real-time feed from the sensor data can be used to develop a trace of events in the ED. For example, the rapid formation and dispersion of large teams at different trauma beds may indicate the possible arrival of several patients with high acuity. Hospital administrators can use the data from the sensors to ascertain the "status of the ED". This information is critical in deploying additional resources, both in terms of personnel and equipment, to the ED. Additionally sensor data can have potential applications when changes are introduced in a critical care environment. For example, the introduction of new health information technology (HIT) creates significant changes in work activities.

Framework for Studying Errors

The study of errors in emergency care settings has received significant attention in recent times. While sensor technology has been minimally used in the investigation of origin and propagation of errors in the ED, it is a viable mechanism for this purpose. From our sensor data, we developed normative and predictive models of clinician activities in the ED. These activities can be retrospectively used to investigate the temporal events and activities that surround reported error incidents.

What is missing from most prior studies on the tracing of errors in critical care environments is the detailed information regarding clinician activities around the time at which the error occurred. The continuous monitoring using sensors provides a large database of clinician location, movement and interaction events. Using the methods described earlier (e.g., transition patterns, group formation and interactions), it is possible to re-create the distribution of attention and resources in the ED around the time at which the error was reported. Such a “replay” of events can help in tracing potential activities that could have been avoided and may have contributed to the error. We will use an example to describe this.

Consider that an attending physician self-reports an error regarding the delayed administration of a drug to a patient in trauma-bed 4 at 530 ET on June 1, 2010. The error report also includes the arrival condition of the patient, history and other patient-relevant information. There are two sets of information that can be used to develop a trace of the events that happened prior and after the error occurred. The sensor data can be used to identify the patterns of interactions, movement and collaboration among the clinicians around the time at which the error happened (say, from 5 to 6 PM on June 1, 2010). The clinical information on the patient along with observation (audio or field notes) can be used as complementary evidence to develop a much richer perspective of the activities surrounding the reported error event. Thus, a detailed sequence of events can be used to track the possible contributory activities that possibly led to the error event. This framework, which combines sensor data and clinical data, for studying errors is shown in Fig. 17.15.

This framework for investigating the origin and propagation of errors has several advantages. First, the data collected from using the sensors can be retroactively combined with the clinical data. Self-reported errors in an emergency setting are usually very low. As such, it is important to be able to trace the events that happened around the time the error incident was reported. Sensors provide a viable mechanism by which data can be collected for extended periods of time and then be retrospectively used for evaluation and analysis. Second, sensors can be used as a passive data collection mechanism with minimal interference with the clinician’s work activities. Third, the relatively long battery life of most sensors makes it feasible for running long data collection sessions (e.g., 20–30 days) without any breaks in data capture. Such an arrangement with human observers is extremely costly and labor-intensive. Our future research work involves the use of the framework to investigate the activities of clinicians in the ED around self-reported errors.

Challenges and Lessons Learned

In summary, there are several potential research and applied opportunities for the use of sensor technology in complex critical care environments. In spite of the significant challenges for designing, calibrating, collecting and analyzing sensor data, we believe that sensor technology has exciting prospects for developing

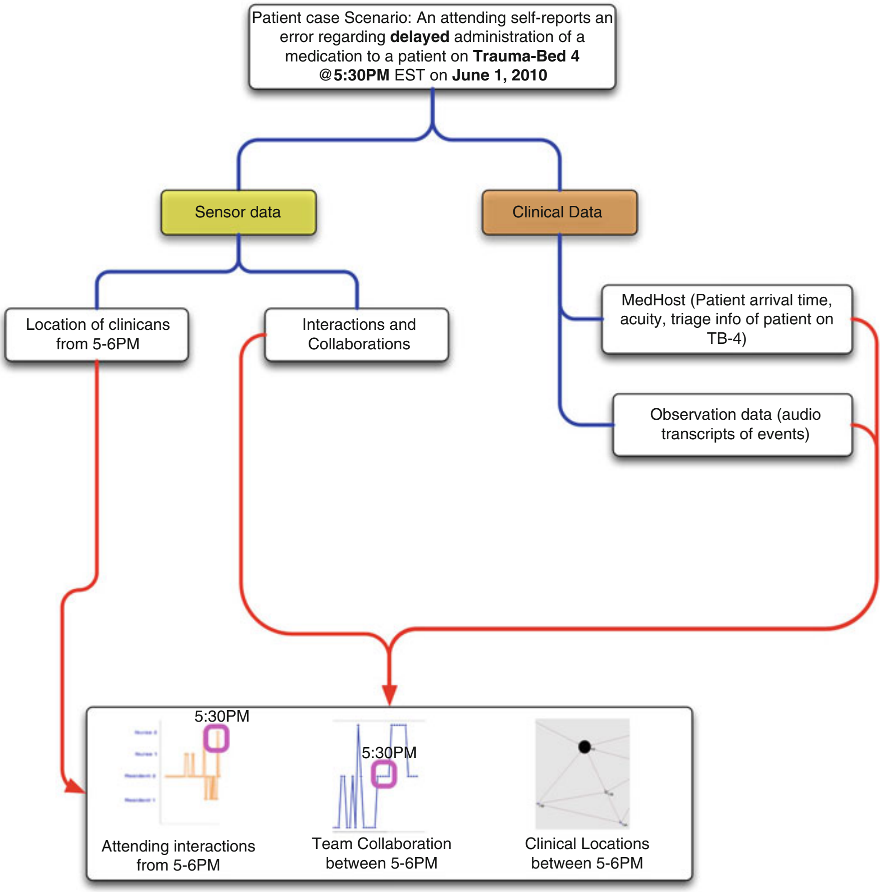

Fig. 17.15 Framework for studying errors

insights of the work of complex critical care environments, which would otherwise be impossible due to significant time and cost burden of using human observers. The calibrating and setting up of the sensors often requires extensive pilot testing to ascertain the exact positioning of the base stations to get maximum coverage. We also had to ensure that our technology did not cause adverse effects on medical equipment and devices. Per our manufacturer's description, our sensor technology operates in the same frequency range as the WiFi (Wireless), which is ubiquitous in hospital settings. While, we did not extensively test for adverse effects of sensors, we believe that our technology does not cause adverse effects on medical devices as argued by van der Togt et al. [30]. Some clinicians were concerned about their privacy issues due to the use of sensors during their shifts. We collected no physician or patient-identifying information and all IRB-regulated protocols were followed for assuring data protection and privacy. For example, all data was saved on an encrypted drive and all identifying information (e.g., time) was removed prior to data analysis. Another significant challenge that we faced was the cost involved in managing the sensor technology. Due to the significant amount of data generated from the sensors, we developed algorithms for compressing and storing the data. This volume of data also required us to develop computationally efficient algorithms for analysis.

Discussion Questions

1. What are the challenges of tracking clinical workflow in critical care settings? What are some of the potential solutions for collecting high-fidelity data in such settings?
2. One of the major challenges with capturing micro-level data (e.g., using sensors) is the significant volume of data. What are some of the approaches to streamline data collection using sensors?
3. How can we minimize the "noise" in sensor data? What are some of the algorithmic approaches for doing so?
4. There are several activities that take place in a hospital setting that may be of interest from the patient safety point of view. Hand washing is one example. Can you provide other activities related to patient safety that would be interesting to track and quantify?

References

1. Institute of Medicine (IOM), U.S. Committee on Quality of Health Care in America. In: Kohn LT, Corrigan JM, Donaldson MS, editors. To err is human: building a safer health system. Washington, D.C: The National Academies Press; 2000.
2. Abraham J, Kannampallil T, Patel VL. Bridging gaps in handoffs: a continuity of care approach. J Biomed Inform. 2012;45(2):240–54.
3. Laxmisan A, Hakimzada F, Sayan OR, Green RA, Zhang J, Patel VL. The multitasking clinician: decision-making and cognitive demand during and after team handoffs in emergency care. Int J Med Inform. 2007;76:801–11.
4. Horsky J, Gutnik L, Patel VL, editors. Technology for emergency care: cognitive and workflow considerations. AMIA Annu Symp Proc. 2006; 334–8.
5. Malhotra S, Jordan D, Shortliffe E, Patel VL. Workflow modeling in critical care: piecing together your own puzzle. J Biomed Inform. 2007;40:81–92.
6. Kannampallil TG, Li Z, Zhang M, Cohen T, Robinson DJ, Franklin A, et al. Making sense: sensor-based investigation of clinician activities in complex critical care environments. J Biomed Inform. 2011;44(3):441–54. Epub 2011/02/25. eng.
7. Vankipuram M, Kahol K, Cohen T, Patel VL. Visualization and analysis of activities in critical care environments. AMIA Annu Symp Proc. 2009;2009:662–6.
8. Vankipuram M, Kahol K, Cohen T, Patel VL. Toward automated workflow analysis and visualization in clinical environments. J Biomed Inform. 2011;44(3):432–40.

9. Bar-Yam Y. Improving the effectiveness of health care and public health: a multiscale complex systems analysis. Am J Public Health. 2006;96:459–66.
10. Patel VL, Cohen T. New perspectives on error in critical care. Curr Opin Crit Care. 2008;14(4): 456–9.
11. Plsek PE, Greenhalgh T. Complexity science: the challenge of complexity in health care. BMJ. 2001;323(7313):625–8.
12. Plsek PE, Wilson T. Complexity, leadership, and management in healthcare organisations. BMJ. 2001;323(7315):746–9.
13. Smith M, Feied C. The emergency department as a complex system: New England Complex Systems Institute; 2006. Cited 30 July 2010. Available from: http://www.necsi.edu/projects/yaneer/emergencydeptcx.pdf.
14. COSI. COSI http://www.irit.fr/COSI/. Cited 20 June 2010. Available from: http://www.irit.fr/COSI/.
15. Berg M. Patient care information systems and health care work: a sociotechnical approach. Int J Med Inform. 1999;55:87–101.
16. Cohen T, Blatter B, Almeida C, Shortliffe E, Patel VL. A cognitive blueprint of collaboration in context: distributed cognition in the psychiatric emergency department. Artif Intell Med. 2006;37:73–83.
17. Celler BG, Earnshaw W, Ilsar ED, Betbeder-Matibet L, Harris ME, Clark R, et al. Remote monitoring of health status of the elderly at home. A multidisciplinary project on aging at the University of New South Wales. Int J Biomed Comput. 1995;40:147–55.
18. Alwan M, Dalal S, Mack D, Kell SW, Turner B, Leachtenauer J, et al. Impact of monitoring technology in assisted living: outcome pilot. IEEE Trans Inf Technol Biomed. 2006;10(1): 192–8.
19. Østbye T, Lobach DF, Cheesborough D, Lee AMM, Krause KM, Hasselblad V, et al. Evaluation of an infrared/radiofrequency equipment-tracking system in a tertiary care hospital. J Med Syst. 2003;27(4):367–80.
20. Schrooyen F, Baert I, Truijen S, Pieters L, Denis T, Williame K, et al. Real time location system over WiFi in a healthcare environment. J Inf Technol Healthcare. 2006;4(6):401–16.
21. Stanford V. Using pervasive computing to deliver elder care. IEEE Pervasive Comput. 2002;1(1):10–13.
22. Fry EA, Lenert LA. MASCAL: RFID tracking of patients, staff and equipment to enhance hospital response to mass casualty events. AMIA Annu Symp Proc. 2005:261–5.
23. Chen C, Liu C, Li Y, Chao C, Liu C, Chen C, et al. Pervasive observation medicine: the application of RFID to improve patient safety in observation unit of hospital emergency department. Stud Health Technol Inform. 2005;116:311–5. IOS Press.
24. Brixey JJ, Robinson DJ, Turley JP, Zhang J. The roles of MDs and RNs as initiators and recipients of interruptions in workflow. Int J Med Inform. 2010;79:e109–15.
25. Li Z, Robinson DJ, Zhang J. UObserve: a mobile app for the study of emergency department workflow. Ann Emerg Med. 2010;56(3):S121. American College of Emergency Physicians Research Forum.
26. Hendrich A, Chow M, Skierczynski BA, Lu Z. A 36-hospital time and motion study: How do medical surgical nurses spend their time? Perm J. 2008;12(3):25–34.
27. Seymour NE, Gallagher AG, Roman SA, O'Brien MK, Bansal VK, Andersen DK, et al. Virtual reality training improves operating room performance – results of a randomized. Double-blinded study. Ann Surg. 2002;236(4):458–64.
28. Kahol K, Vankipuram M, Smith ML. Cognitive simulators for medical education and training. J Biomed Inform. 2009;42(4):593–604.
29. Ahlberg G, Heikkinen T, Iselius L, Leijonmarck CE, Rutqvist J, Arvidsson D. Does training in a virtual reality simulator improve surgical performance? Surg Endosc. 2002;16(1):126–9.
30. van der Togt R, van Lieshout E, Hensbroek R, Beinat E, Binnekade JM, Bakker PJM. Electromagnetic interference from radio frequency identification inducing potentially hazardous incidents in critical care medical equipment. JAMA. 2008;299(24):2884–990.