

A Pervasive Computing System for the Operating Room of the Future

Sheetal Agarwal · Anupam Joshi · Tim Finin ·
Yelena Yesha · Tim Ganous

Published online: 27 July 2007
© Springer Science + Business Media, LLC 2007

Abstract We describe a prototype context aware peri-operative information system to capture and interpret data in an operating room of the future. The captured data is used to construct the context of the surgical procedure and detect medically significant events. Such events, and other state information, are used to automatically construct an electronic medical encounter record (EMR). The EMR records and correlates significant medical data and video streams with an inferred higher-level event model of the surgery. Information from sensors such as Radio Frequency Identification (RFID) tags provides basic context information including the presence of medical staff, devices, instruments and medication in the operating room (OR).

Patient monitoring systems and sensors such as pulse oximeters and anesthesia machines provide continuous streams of physiological data. These low level data streams are processed to generate higher-level primitive events, such as a nurse entering the OR. A hierarchical knowledge-based event detection system correlates primitive events, patient data and workflow data to infer high-level events, such as the onset of anesthesia. The resulting EMR provides medical staff with a permanent record of the surgery that can be used for subsequent evaluation and training. The system can also be used to detect potentially significant errors. It seeks to automate some of the tasks done by nursing staff today that detracts from their ability to attend to the patient.

Sheetal Agarwal was a student at UMBC when this work was done.

S. Agarwal (✉)
IBM India Research Lab, New Delhi, India
e-mail: sheetaga@in.ibm.com

A. Joshi · T. Finin · Y. Yesha
Department of Computer Science and Electrical
Engineering, University of Maryland, Baltimore County,
Baltimore, USA

A. Joshi
e-mail: joshi@cs.umbc.edu

T. Finin
e-mail: finin@cs.umbc.edu

Y. Yesha
e-mail: yeyesha@cs.umbc.edu

T. Ganous
University of Maryland Medical School,
Baltimore, USA
e-mail: tganous@verizon.net

Keywords pervasive computing system ·
operating room · electronic medical encounter record ·
medical informatics

1 Introduction

Performing a surgery is an elaborate process and proceeds in progressive stages. The term Perioperative generally refers to the three phases of surgery namely preoperative, intraoperative, postoperative. The pre-operative phase of a surgery involves identifying the patient, determining readiness of the nursing and medical staff, preparing the operating room, and capturing data from incoming medical records regarding vitals, pre-operative medications, tests, and scans etc. In a perioperative setting, hundreds of patients and staff may be flowing through dozens of operating rooms on

a daily basis in a single facility. Some fraction of these patients are unscheduled and identified only on the day of surgery—this fraction can be large in facilities that deal with trauma. The resulting chaos can be overwhelming, even with some form of electronic health record (EHR) system. This is because the orchestration of behavior between information systems is people and paper based. Available automated systems are often dedicated to isolated operations or departments with no automated means to communicate with one other. The limitations of this environment provide great opportunity for process improvement efforts. The work of Sutherland et al. [30] on realtime process improvement in healthcare (RECIPE) states that 30 to 100% improvements in healthcare processes are possible. In the perioperative setting, the information system must keep track of (and archive) what is done to each patient (both diagnostically and therapeutically), how, when, and why it is accomplished, and how the patient responds.

Our goal is to develop a *context-aware perioperative information system* that will automate most of the perioperative support functions such as patient tracking, inventory management, clinical documentation etc with the help of pervasive computing and semantic web technologies. In this paper we present the design of such a system, and a preliminary implementation that tracks supplies, prevents errors, and automatically creates an electronic medical encounter record to document the events occurring during a surgery.

We seek to create a system that can work in two environments. One is the “operating room of the future.” This term is used to describe the concept that is sought to be realized by several surgical groups across the country and involves an IT enabled operating room. A good example is the ORF project at CIMIT (<http://www.cimit.org/orfuture.html>). Another example is the ORF project at the University of Maryland Medical System, which is where our work has been done. The aim here is to assist the surgical team by interpreting events in context. The second is Traumapod [27], a DARPA project whose aim is to develop an automated surgical treatment system that does not require onsite medical personnel on the front lines of battle, and is ready to receive, assess, and stabilize wounded soldiers during the critical hours following injury. The first phase of the program is an effort to develop robotic technology to perform a totally unmanned surgical procedure within a fixed facility. A human surgeon will conduct all the required surgical procedures from a remote location using a system of surgical manipulators. Automated robotic systems provide necessary support to the surgeon to conduct all

phases of the operation. Our system seeks to provide the intelligence and recordkeeping needed in the perioperative environment, that in a normal surgical setting is provided by the scrub and circulating nurses who are assisting the surgeon.

2 Background

Clinical record keeping in high velocity healthcare delivery environments like surgery is a necessary and critical task. It is also a time consuming task that detracts from hands on patient care and contributes to extraordinary labor costs associated with collecting, transcribing and re-keying records throughout the perioperative process. The details of the surgery are documented in patient charts called the *perioperative record*. This record contains information about the patients vital signs at periodic intervals, medicines administered, complications if any, supplies and tools used etc. Errors in medical documentation cost billions of dollars to the health industry every year [19, 33]. Inaccurate records put not only the patient but also the healthcare provider at risk [11, 26]. Similarly, there is a need to track assets, both in terms of equipment and personnel. Perishable assets in particular need to be tracked as they are used, and the usage information needs to be integrated with supply chain management. There are studies that suggest significant costs are incurred in couriering medical supplies—to the tune of billions of dollar in a year. In addition, asset tracking can also help avoid surgical errors by ensuring that supplies such as sponges are properly dispensed and not left in the patient. Finally, there is a need to prevent fairly simple errors, such as the presence of incorrect equipment for the surgical procedure, or an incorrect surgeon or patient in the operating room.

The data recorded during the perioperative process become a part of the patients medical history and is used by physicians to give further treatment to the patient. Data collection in the operating room is complicated due to several reasons. Firstly, multiple providers (e.g., surgeons, anesthesiologists, nurses) record data for a single care event (i.e., the patient’s surgery). Secondly, information collected by one provider is not readily available to another. Thirdly, experienced nurses assess the patients’ condition accurately and provide appropriate treatment, sometimes without documenting these procedures; thus, duplication or differences occur in documentation, data gathering can be cumbersome, and not all details are recorded. Moreover, in situations such as those envisaged in Traumapod, there are no nurses or assistants—a

remotely located surgeon is performing the procedure. Collecting data, maintaining the state of the system, and documenting the process therefore needs to be automated as well.

An electronic medical encounter record (EMR), has the potential to reduce documentation errors by minimizing data redundancy and providing accurate details of the ongoing surgery [6, 12]. Formally, an EMR is a medical record or any other information relating to the past, present or future physical and mental health, or condition of a patient, that resides in computers which process this data to deliver more efficient health-related services. The EMR is an essential part of systems like the Traumapod [27] where surgeries are performed by remotely controlled robots and no humans are involved in the process. Only the EMR can provide details of the events occurring during the surgery.

The operating room (OR) has several medical devices that provide information about the patients status. In addition to these devices, we can deploy sensors in the OR that can provide us with better view of the activities occurring in the operating room during a surgery. Unfortunately, most sensing technology detects low level events, and is rather error prone. For instance, consider the use of RFID tags on medicines. Clearly, the sensor will only tell us that the particular medicine is nearby or in the room—this information may or may not be valuable in of itself. What is more valuable is the knowledge that a particular medicine was administered at a given time. Moreover, in an OR with presence of fluids and metal, the accuracy of RFID sensing is itself not very high or robust.

We define a medically significant event as any event that affects or is a part of the surgical procedure. Many systems [20, 22, 28] have been built that monitor physiological parameters of a patient and signal alarming conditions. Healthcare providers use these alarms as cues as it is not possible to maintain a constant vigil over the patients’ health status. The alarms are in the form of an audio alert or a message displayed on the computer screen that can be seen by the healthcare provider.

Most of these signal low-level events such as tachycardia, apnea or any other abnormal pathological state. As we will describe later, such low level events generally do not in of themselves provide any meaningful detail about the patients condition to the surgeon. To provide more meaningful information the alarms or medical events need to be interpreted at a higher level and documented. In addition to physiological data we can make use of data streams from sensors that can be deployed in an operating room to capture additional events such as tools and medicines used and identities

of the members of the clinical staff. In our research we use the radio frequency identification (RFID) system to detect medical supplies, tools and the staff.

3 System architecture

The framework we developed is designed to collect, process and reason over surgical context information to detect medically significant events (Fig. 1). It consists of a three-tier event detection system. Events at the lower levels are processed to infer high-level events. At the lowest level, data is collected from various sensors in the operating room and processed by a stream processor to identify low level events like presence/absence of an RFID tag. These low level events are then processed by the second tier to detect low level medical events such as high blood pressure, apnea etc. Event detection at this level also makes use of patients medical history and knowledge of medicines used to detect events. At the highest level we have the rule based engine that correlates events from the second tier to detect medically significant events such as administration of anesthesia, tension pneumothorax etc.

3.1 Data sources

- *Patient monitors* The operating room has several patient monitoring systems that track the patients physiological parameters. For example, pulse oximeter monitors blood oxygen saturation levels, vitals signs monitors track heart rate, blood pressure etc. To monitor the patient’s condition during the surgery, the surgical team monitors the value and change in physiological parameters. We use

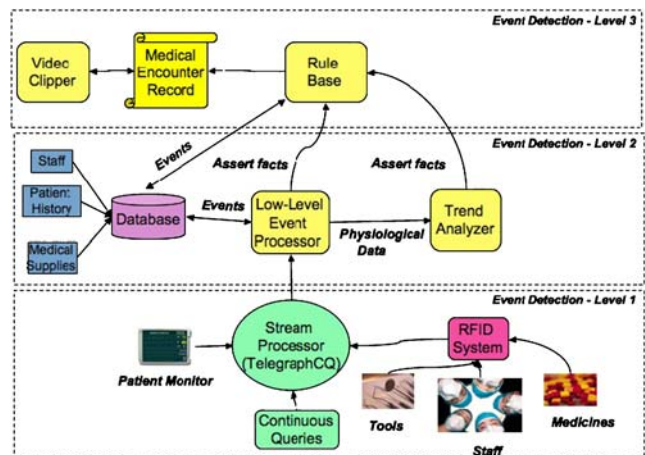


Figure 1 System architecture

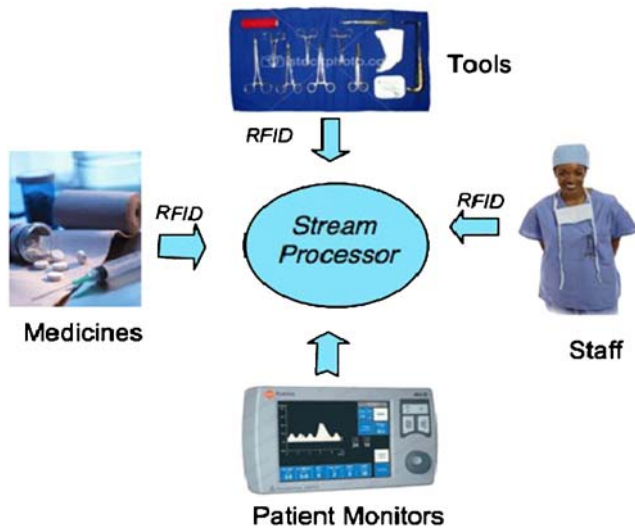


Figure 2 Data sources in operating room

data streams from these patient monitors to determine the state of the patient during the surgery.

- *Asset and personnel tracking systems* The system is designed to incorporate Bluetooth and radio frequency identification to track medical supplies and nursing staff. The architecture can be easily extended to support other tracking technologies. For the prototype developed we used passive RFID tags. Readings from the RFID reader are analyzed to determine the resources used during the surgery and the team performing the surgery. Figure 2 shows the data sources used to acquire contextual information in the OR.
- *Message exchanges* Where the operating room is highly automated, as in the Traumapod system, the nature and sequence of messages exchanged between its constituents can provide a rich source of state data. For example in Traumapod, we can listen to exchanges between the “tool rack” system and the scrub nurse system to figure out which tools have been moved from the rack to the (robotic) nurse. This can be done without any explicit sensing technology that would be needed in a live OR with human surgeons and nurses.

3.2 Low level event management

The patient monitoring systems and the RFID reader produces continuous streams of data that need to be processed and analyzed in real-time to detect events. Traditional database systems have been designed to manage finite data sets where client queries are processed immediately against data stored in tables.

In applications that processes continuous data streams, clients require long-running continuous queries that are evaluated as data streams through the application. For example, consider a query that monitors physiological data streams. “Report the heart rate and blood pressure values when heart rate is >100 and blood pressure <70 over a period of 60 s.” Significant work has been done in the area of stream processors. Some of the well know systems are Stanford STREAM data manager, Aurora, Borealis and TelegraphCQ. Each of these systems focus on processing continuous time-varying data.

In our prototype implementation we used, TelegraphCQ [7], developed at University of California, Berkeley to process the physiological and RFID data streams. Data from patient monitors and the RFID reader is pushed to the stream engine continuously. Queries over these data streams are specified over a time window. As new data arrives, the queries are evaluated and results are returned to the client.

3.3 Analyzing RFID data streams

Figure 3 shows the RFID system we used in our prototype. We used the Symbol AR400 900MHz Reader and passive RFID tags. An RFID tag has a unique 96-bit identifier called the electronic product code (EPC). The RFID reader returns the list of EPC codes it detects. We implemented the *byte stream protocol* to interface with the RFID reader. The RFID API we developed provides a layer of abstraction over the low level protocol. The API processes results from the RFID reader.

The RFID module polls the reader periodically to get the list of RFID tags visible. The passive tags use the energy incident from the reader, to return their EPC code. The reader reads the tags at its own internal

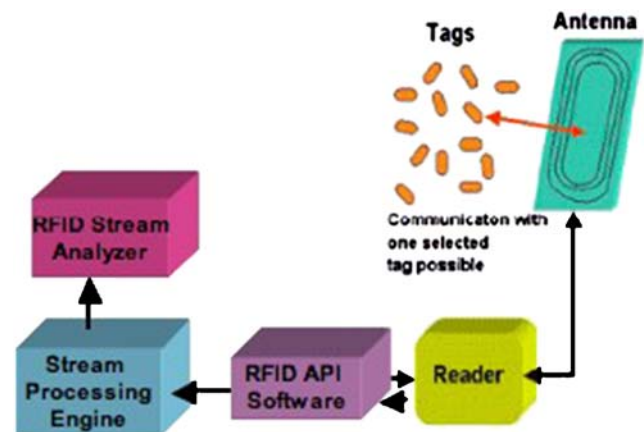


Figure 3 RFID System

frequency. Hence the same tag may be reported more than once in the list of tags detected. When many tags are in close proximity, the signals returned by the tag collide and result in loss of data. Thus a single read from the reader is not sufficient to detect all tags reliably. A continuous query over the RFID data stream, aggregates the number of times a particular tag is reported by the reader. The reader is sampled every 2 s. Experiments show that if a tag is detected at least ten times in a 30 s window i.e. 66.67% of the time, then the tag is visible. However, if a tag goes undetected for 4 consecutive windows i.e. 120 s the tag is determined to have become invisible in the readers field of view.

Tag visibility rules¹

If number of times tag seen ≥ 10 then
Event (Tag Visible)

If tag not seen for ≥ 120 s then
Event (Tag Invisible)

3.4 Analyzing Traumapod messages

As mentioned earlier, the sequence of messages between components of Traumapod can also be used to infer state and create context. Figure 4 shows the Traumapod network. The various systems in the network communicate using the Spread messaging toolkit.² This is an open source messaging system developed at Johns Hopkins University. Each system in the network implements a set of interfaces. Some interfaces are implemented by all subsystems and each subsystem has its own interface defined.

The supply dispensing system (SDS) is a robotic system that contains the medical supplies in form of supply trays. This system is responsible for dispensing the supplies requested by the surgeon. The tool rack system (TRS) is a robotic system that contains the surgical tools. This system dispenses the tools required during the surgery. The scrub nurse system (SNS) moves the supplies and tools between the surgical site and the SDS or TRS. The supervisory control subsystem (SCS) is the master system that controls all the other subsystems. Any request for a tool or a supply is first sent to the SCS. The SCS then executes the task by sending appropriate commands to SDS, SNS and TRS.

¹Expressed as queries over RFID data streams.

²<http://www.spread.org/>

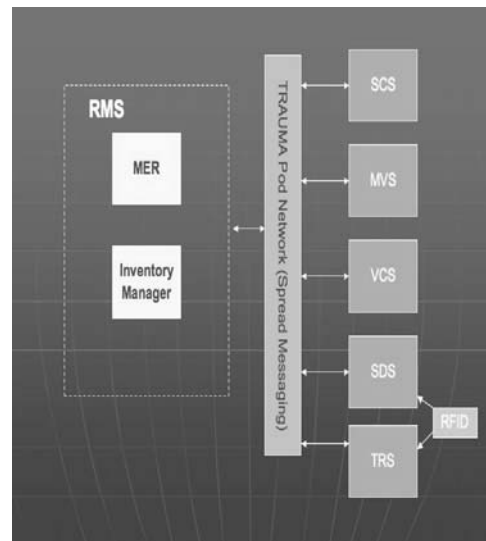


Figure 4 Electronic medical encounter record

The component we add to Traumapod is called the resource management system (RMS). It is a passive system that snoops for messages between SCS, SDS and TRS and based on the messages infers the state of the surgery. For each procedure that is to be performed, the surgeons typically define a surgical template outlining the key steps in the surgery. These steps are broken down into a sequence of messages that would be expected if that step was being performed. For instance, a surgeon will give a voice command to the system request some sutures. This would be directed to the SCS, which would then send messages to the SDS to dispense that suture, and to the SNS to transfer the suture to the surgeon. The messaging template is used as only a hint. The actual surgery may involve additional steps. A rule base is created to map message sequences to a surgical step, and the steps to key events of the surgery. The rules are based on Traumapod events such as supply dispensing, tool requests and tool changes. The current rule base consists of rules to detect events during a “Shunt Procedure,” which is the procedure that the initial phase of the Traumapod project is expected to demonstrate.

4 Detecting medically significant events

The highest level of event detection in the system consists of a rule-based system, JESS [13] with FuzzyJ library. This system define the rules to detect the medically significant events. In this section we will describe the rule elicitation process and the techniques we used to detect events.

Physiological parameters reflect a patients' health status. Interpretation of physiological data to infer the patients condition is a challenging problem. Some of the earliest data analysis systems used simple limits on physiological parameters for basic interpretation. The more advanced systems considered dependencies between parameters to provide more meaningful interpretations [20, 22, 28]. The problem with such systems is the high rate of false positive or false negative events. The poor performance was due to the fact that physiological data was interpreted independently of the clinical conditions in which the data was acquired. Physiological parameters not only depend on the physiologic processes but also on factors such as patients current condition, medical history, medicines administered and the sequence of occurrence of other events. Most systems consider only a subset of these factors. Since our system can directly detect or infer many of these factors, we can do a better job of interpreting the parameters.

4.1 Rule elicitation process

The knowledge base for our system was developed by gathering information directly from an anesthesiologist (Dr. Geoffrey Kuzmich) through interviews, and reviewing medical literature [4, 15, 16, 25, 29, 31, 32] which describes methods to analyze and interpret physiological data.

The anesthesiologist is the primary person who is responsible for treating and regulating changes in critical life functions—breathing, heart rate, blood pressure—as they are affected by the surgery being performed. These medical specialists are the doctors who will immediately diagnose and treat any medical problems that might arise during the surgery.³ We chose to develop a rule base for trauma care events since (1) The events that occur in the trauma care scenario are very common in surgical scenarios (2) The human patient simulator we used to evaluate our system is currently programmed for trauma care surgeries, and (3) the Traumapod system focusses on Trauma cases, as the name implies.

The anesthesiologist helped us shortlist some of the trauma care scenarios that we would be using to evaluate our system.

4.1.1 Trauma care scenarios

Each scenario outlines the following details

- *Pre-op diagnosis* The condition of the patient before he is brought into surgery. This includes details of the external injuries and an initial assessment of the condition. For example, a patient that has suffered excessive blood before being operated on, is said to have hypovolemia.
- *Event list* The scenario describes the various states, such as hypovolemia, tension pneumothorax etc., that a patient goes through and the treatment given.

For each of the events in the scenarios we took note of the physiological parameters that need to be monitored. Each physiological parameter has a range of values and given a data value, there is no set threshold that will deterministically classify the value as normal or abnormal. Also the interpretation of a parameter varies with the clinical context. In order to help us understand how the physiological parameters change for each event, we ran some simulations on the human patient simulator for the selected scenarios. We examined the changes in parameter values in the data sets generated from the simulations. Based on the interviews we established parameters that need to be monitored for each event, and for each parameter defined the following:

- Range of values. For example: heart rate 0–200
- Rate of change of the parameters. For example: Heart rate can be “Increasing,” “Decreasing,” etc.
- Range for the rate of change of parameters. For example, heart rate change by 20–30 is an abrupt change

These observations were not absolute. There is no fixed combination of parameters and their values that can detect an event. To capture this uncertainty we used fuzzy rules. The fuzzy rule functions used were chosen based on the above findings. For example, the range for heart rate is 0 to 200. When the heart rate is between 80–90 the heart rate is normal. However heart rate of 79 or 91 does not mean that the heart rate is definitely low or definitely high. Hence we define a Trapezoid-FuzzySet [23] to partition the range of values for the heart rate. To defuzzify the values we use the maximum defuzzification function. In this method the mean of the x values, with maximum membership values over the entire set of FuzzyValues, is calculated.

In the interviews we also discussed how the values and the change in the parameters are affected by the patients medical history. To incorporate the diverse

³<http://www.asahq.org/patientEducation/know.htm>

patient population in our knowledge base, we set the range of normal values for a parameter based on the patient medical history. This data is predetermined from the patients profile. It is not a part of our system to analyze the medical history to establish these values. When the system initializes, it loads the patient information into the system which includes values for the range of the parameters. For all the data sets we used, the anesthesiologist helped us establish the value of the parameters.

4.2 Techniques to detect events

Given the data value of a physiological parameter and its rate of change, the techniques we used to correlate low level events are:

- *Pre-op diagnosis* Typically, before the patient is brought into the surgery, the patient's condition is evaluated. The evaluation includes taking note of the vital signs, any medical care provided, a physical examination and any other notable medical condition. The pre-op diagnosis is used to initialize the event history. The actions taken during the surgery also depend on the pre-op diagnosis, so we used this information in detecting events. For example, if the patient was bleeding excessively prior to the surgery, detecting fluid infusions during the surgery is more accurate. Currently we use pre-op diagnosis to only provide clues about the patients' condition before the surgery starts.
- *Multi-variable analysis* Monitoring a physiological parameter in isolation does not give much information about the state of the patient. Coleman et al. [9] state that "each physiologic state variable is intimately related directly and indirectly to many others by relationships." Thus interpretation of a variable varies with its relationship with other variables. For example low and decreasing blood pressure, does not give signify too much detail about the patients condition. However, low and decreasing blood pressure with high and increasing heart rate implies potential loss of fluids. Monitoring a physiological parameter along with its relationship with other parameters helps determine more meaningful events.
- *Event history* Event history is a set of low-level and high-level events already detected. A high level event is often a composition of low-level events and potentially other high level events. The composition can be a conjunction or disjunction of events

[21]. Often, events need to be considered in the context of the event history and both a high and low level event.

Example: event conjunction⁴

```
If (TensionPneumothorax)
    and (Systolic BP "low" and
        "increasing")
    and (Heart Rate "high"
        and "decreasing") then
    Event (Decompression)
```

- *Effect of medicines* The medicines administered during the surgery, may or may not have a significant effect on the patient's physiology. The time the medicine was administered, effects expected, duration of effect and time to affect are some of the factors that need to be taken into account to detect their effect in the physiological parameters. Determining the event of medicine administration is extremely difficult without the context of the medicine used. In our system, the medical supplies are tagged using RFID tags. Thus we can acquire the contextual information and use this information to detect medicine administration events. Detection of the medicine does not imply that the medicine was actually administered.

In the current version of the system we record all the medicines detected by the RFID system. For those medicines whose effect is observed in the physiological parameters, an event is signaled to indicate that the medicine was actually used. For example, when a paralytic medicine is administered as a part of the anesthesia process the respiratory rate of the patient is expected to drop to very low values. This is normal when the medicine is administered. However, without the context the event is signalled as a false alarm. On the other hand, when a pain killer is administered, it does not produce any significant changes in the patients physiology. Thus the medicine is only signalled as detected and not administered. In some cases, it has been suggested that a dedicated short range RFID type sensor exist and that the medicine to be administered be flashed before it to indicate administration. Where practicable, such a system can be of significant help.

⁴Expressed as JESS rule in the knowledge base.

Example:⁵

```

If (Respiratory Rate "abrupt decrease"
and Systolic BP "abrupt decrease"
and Heart Rate "stable")
  If (Time Anesthetic detected
    < 40 sec)
    then Event (Start Anesthesia)
  
```

5 Results

We have created a prototype of the system described above and evaluated it in a surgical training setting. We first describe the EMR Interface that is actually available to the surgeons and trainees to interact with our system, and then show results about how accurate our event detection system is.

5.1 System interface: electronic medical record

The EMR can be displayed on a computer screen in the operating room. It provides a summary of the patient profile, the pre-op diagnosis, laboratory reports, and radiology/imaging reports. Before surgery can commence, obvious errors such as right patient, right surgeon, right equipment and the presence of consent form (assumed to be RFID tagged) are checked. Any noted errors are announced. If no errors are detected, permission to proceed is provided.

During surgery, the vital signs of the patient are streamed to the screen during the surgery. The event list gets populated as events are detected. Typically, only clinically significant events are populated on the screen. A part of the screen is used to show the medicines and the surgical staff as detected by the RFID system. As members of the surgical team enter and leave the OR, the screen is updated to show only those present in the room. For each event we save the vital signs of the patient at that instant of time. In many surgical situations videotaping is done. Our system can accept and display the feed. In addition, post surgery, video clips of defined length for each of the clinically significant events are created and the corresponding video URL is stored in the medical encounter record. When reviewed at any time after the surgery, the surgeon or a trainee can interact with the record to see the vital signs of the patient at the major points in the surgery and the associated video data. Instead of viewing the entire video footage of the surgery, the

Date	Start Time
5/1/06	17:52:34

Patient Name	Procedure
Alexander	Multiple Trauma Surgery

Pre-op Diagnosis		
Vital Signs		
Parameter	Value	
RR	12.0	
HR	75.0	
SBP	100.0	
DBP	90.0	
O2SAT	98.0	

Labs		
Date	Report	Result
2006-03-28	Urinalysis	Normal

Events		
Event Name	Annotation	Time
Surgeon Ent...	Dr Park ente...	17:53:20
Scrub Nurse...	Sae Eliot ent...	17:53:55
Hypovolemia	Hypovolemia	17:53:56
Anesthesiolo...	Dr Evans ent...	17:53:56
Tension_Pne...	Tension_Pne...	17:54:36
Decompress...	Decompress...	17:55:26
Scrub Nurse...	Sae Eliot left...	17:55:30

Intra-op Vital Signs		
Parameter	Value	
Heart Rate	93.3	
SBP	100.9	
DBP	52.2	
O2SAT	98.0	
Respiratory Rate	0.0	
PaO2	76.0	
MAP	66.0	

Medicines		
Name	Type	Time
Fentanyl	Pain Killer	17:55:45

Staff		
Role	Name	Enter Time
Surgeon	Dr Park	17:53:20
Anesthesiolog...	Dr Evans	17:53:55

Figure 5 Electronic medical encounter record

surgeon can browse through the key parts of the video by selecting an event from the event list. Figure 5 shows a snapshot of the EMR.

In addition, separate interfaces are available that allow access to the state information relating to supplies that the system maintains. In other words we can see what supplies and tools were used during the procedure, and what numbers remain. We hope that this information can be subsequently integrated into the supply chain management system at the hospital.

5.2 Event detection accuracy

In this section we describe the test environment we used to evaluate our system. We used physiological data sets from the human patient simulator (HPS)⁶ called Stan, manufactured by the METI. It is a complex system that emulates the human body response to medical treatment. The simulator is used to train medical students.

Stan, shown in Fig. 6, can be loaded with various patient profiles. For example, the doctor could create an asthmatic patient with chronic heart disease who is taking a handful of certain drugs and is currently experiencing anaphylactic shock, a severe allergic reaction. The medical students in turn have to figure out how to treat the patient. If medication is required, the drugs are “administered” by scanning the barcode/RFID tag at a designated reader. The simulator produces in Stan the physiological response that the drug would have produced in a patient with that medical condition. The

⁵Expressed as JESS rule in the knowledge base.

⁶<http://www.meti.com/>

Figure 6 Human patient simulator



HPS remains in each of the states in a given scenario for a fixed period of time after which it transitions to the next state. The changes in the physiological parameters of the simulator are logged constantly and the parameters vary according to the current state of the HPS.

This system was made available to us by the Air Force Simulation Center at University of Maryland Medical School. Since Traumapod [27] focuses on trauma care on the battlefield, we chose to use trauma related scenarios to evaluate our system. Several variants of the trauma scenario were tested, we describe a key scenario here.

5.2.1 Scenario: blunt trauma multiple injuries

This scenario consists of a patient who has been wounded in a battlefield. In this scenario the patient goes through the following states during the course of trauma care:

- Hypovolemia (excess blood loss)
- Tension pneumothorax
- Decompression
- Fluid Infusions

These states constitute the set of medically significant events we wish to detect from the low level sensor data streams. This scenario was simulated on different patient profiles. Each patient has different medical history and pre-op diagnosis. The scenarios were simulated with five different profiles. The scenarios were chosen to be realistic in that the underlying vitals did not directly translate to a specific state. For instance, conditions to detect hypovolemia are increasing heart rate and a decreasing blood pressure. However, these conditions occur during Tension pneumothorax also. As such, state detection needs to take into account current vitals, their history, pre-op diagnosis, past states detected etc.

We determined the sensitivity and specificity of detection for each of the events by our system on data from 30 different simulations on the HPS system. One

key criterion of success is low false positives (and negatives). We found that for most events, we had very high true positives (in the 90+ percent range), and very low false negatives. The exception was events such as Hypovolemia and Fluid Infusion, which depend on conditions for which we do not have all vitals available, and where vitals alone are often not determinative of the event. However, even here the use of event history and pre-op diagnosis significantly increased the true positives over just using the available vitals. For instance, knowing that a patient has undergone hypovolemia makes it more likely that subsequent change in vitals might signal fluid infusion. Similarly, the events “paralytic administered” and “reducing anesthetic” show 100% specificity and sensitivity as they are triggered only when the appropriate medicine is detected as being administered by the RFID reader. Without RFID, it is difficult to detect these events by just monitoring the physiological data.

The other concern is the latency between the occurrence of an event and its detection by the monitoring algorithm plays an important role in the performance of the system, as does the order of detection. We detected all events in proper order. The latency varied from a few seconds all the way upto 56 s, and the variance was generally high. This was especially true for events where the response of the vitals “develops” over time and is not instantaneous. It is also reflective of the variance in individual patient physiologies. For instance, in a normal patient, hypovolemia will cause a rapid drop in blood pressure. In a hypotensive patient however, the drop is more gradual.

6 Discussion

Our rule base currently has 27 rules. Adding and retracting facts from the knowledge base is an expensive operation. In our system, size of the working memory changes slowly, once the initial set of facts has been added. As the parameters change their values, the existing facts are updated to reflect the change. We avoid

adding and retracting facts frequently as this affects the performance of the system.

The system evaluation is extremely positive, with many events having 100% true positive and almost zero false positives. However, admittedly the data set was small and was obtained from a Human Patient Simulator. The results may vary with real patient data. They will also likely vary as more scenarios are evaluated. Also as the knowledge-base grows, addition of new rules to detect more events may increase the number of false positives.

Not every rule in the knowledge base will be appropriate for each surgical scenario. We can take a template based approach where the rule base is initialized with only those rules that are relevant to the current clinical scenario. Determining the set of relevant rules is again a challenging problem. We will address the problem of scalability of the knowledge base in our future work.

Currently we use simple queries over the data streams to detect low-level events. We maintain a state variable model and use various techniques to correlate these low-level events to infer more meaningful events. Some of the event correlation can be done by using appropriate queries on data streams. The current version of TelegraphCQ does not provide support for sub-queries and access to historical data. With support for sub-queries and access to archived data in the subsequent version we can move some of the event correlation rules to the stream processing level.

We use RFID to detect staff and medicines in the operating room. The use of RFID in healthcare presents a number of critical issues unique to healthcare in addition to the basic limitations of the technology.

- *Electromagnetic interference* The healthcare environment is already full of safety critical devices that are sensitive to radiation at various frequencies.
- *Tagging medical supplies* We conducted a feasibility study of using RFID to tag medical supplies. The current state of art is not sophisticated enough to allow tagging of all medical supplies. The smallest passive tags available are 1" × 1". With tags of this size it is difficult to tag items like surgical tools, medical supplies like cotton balls, sponges, gauze etc. Tags that are of the size of a grain of rice are also available. But these tags are designed to embed under the skin of cattle or humans. These are not suitable to tag medical supplies.
- *Environment hazards to tags* The healthcare industry presents a unique challenge to the physical integrity of RFID tags because of its pervasive

infection control measures. Supplies like sponges, gauze become wet with fluids. Tags attached to clothes may be damaged when they are washed. The RFID tags were originally designed to tag objects for supply chain management and are not capable of withstanding harsh medical environments.

Due to the above limitations, in the current version of the system we make use of heuristics to determine tag visibility. The simulations to evaluate the system were conducted in the laboratory. The latency of events depending on RFID input is likely to change in an actual operating room setting due to the factors mentioned above. In spite of the above the limitations, the use of RFID in healthcare is expected to rise rapidly. According to a report published by IDTechEx [14] “the market for RFID tags and systems in healthcare will rise from \$90 million in 2006 to \$2.1 billion in 2016. Primarily, this will be because of item level tagging of drugs and real time locating systems (RTLS) for staff, patients and assets to improve efficiency, safety and availability and to reduce losses.” The technology is expected to evolve to address the requirements of RFID in hospital environments.

Another important aspect to consider is that the EMR documents only those events that are detected by the system. A complete perioperative record has several details such as physical observations of the patients body, devices implanted, exact amounts of fluids infused etc. These are the kind of details that cannot be deduced from the data sources we currently use.

7 Ongoing work

In this section we describe some of the ongoing extensions we are implementing and enhancements that could improve the system.

7.1 Domain-based medical ontology

A knowledge-based system represents relationships between objects, entities and concepts that exist in a domain of interest. Ontology is a specification of such concepts. The relationship between the objects is specified in a vocabulary that is used by the knowledge systems to represent knowledge [10]. Within health informatics, ontology is a formal description of a health-related domain.

The use of ontologies in medicine is mainly focused on the representation and (re-)organization of medical

terminologies. Physicians developed their own specialized languages and lexicons to help them store and communicate general medical knowledge and patient-related information efficiently. Such terminologies, optimized for human processing, are characterized by a significant amount of implicit knowledge. Medical information systems, on the other hand, need to be able to communicate complex and detailed medical concepts (possibly expressed in different languages) unambiguously.

In the perioperative environment, use of a standardized language decreases patients' risk for injury by eliminating inconsistency of language or meaning. This is a difficult task and requires a detailed analysis of the structure and the concepts of medical terminologies. But it can be achieved by constructing medical domain ontologies for representing medical terminology systems.

The benefits of using a medical ontology are:

- Ontologies can help build more powerful and more interoperable information systems in healthcare.
- Ontologies can support the need of the healthcare process to transmit, re-use and share patient data.

Constructing the medical encounter record using a domain-based ontology will make the record usable by other health-informatics systems for further processing. Several groups, such as GALEN [24], CIMIT [8], SNOWMED-CT (<http://snomed.org/>), have developed medical ontologies to represent medical concepts. Most groups focus on a domain within medicine and have their ontology represent concepts relevant to the domain. The unified medical language system (UMLS) [5] is a meta-thesaurus created by the National Library of Medicine (NLM) that integrates the ontologies developed by various groups.

Moreover, ontologies and Semantic Web languages such as OWL can be used to capture domain knowledge and rules explicitly and in a machine interpretable way. This will allow for significantly greater interoperability at a semantic level.

7.2 Fine grained tracking with RFID

Supply counting is an important procedure during a surgery. It is the responsibility of the surgical team to ensure that no supply is left within the patients body at the end of the surgery. RFID can be used to perform supply counts provided all supplies can be tagged. Since RFID tags cannot be localized, as an alternate solution we can use low frequency readers to detect tags in a

particular zone of the operating room. The ability to divide the operating room in zones will allow us to track the supplies in the operating room and ensure no supply is left within the patient's body. Lui et al. [18] describe a system *Ferret* that determines the approximate position of a passive RFID tagged objects. The algorithm employed depend on the mobility of the RFID reader to detect the tags from different vantage points. Repeated intersection of the region of influence of the reader at each point refines the location of the object. In the hospital environment we could instead use a network of RFID readers that share the tag readings and determine the approximate location of an object.

Tracking supplies at this granularity can also be useful in inferring events that are not detectable through physiological data streams. For example, if the surgeon is holding a vascular clamp and the surgery involves placing a shunt, we can estimate the time that the clamp was used to clamp the blood vessels. With the current system, such events are not detectable.

8 Related work

We developed a context-aware system that monitors and analyzes the data streams from various medical equipments and create an electronic medical encounter record, according to the inferences made by analyzing the data streams, in a perioperative environment. The surgical team can see the record being populated in real-time which ensures that everyone is aware of the progress being made and of the patients health status at all times. The system was designed to detect events during trauma care and general anesthesia scenarios.

Automated analyses of a patients' physiological data to detect alarming conditions has been a subject of research for over a decade. Several patient monitoring systems have been developed that alert the healthcare provider to alarming conditions. InCare [20], is one of the earliest automated systems to detect events in post-cardiac operated patients. InCare had a rule-based system that used multi-variable and trend based analysis of physiological data to detect events. Similarly, Schecke et al. [28] designed a knowledge-based decision support system for patient monitoring in cardio anesthesia. The medications used and progress of the surgery was fed into the system manually by one of the members of the surgical staff.

Hewlett Packard Labs has recently developed a framework that allows development of scalable soft-

ware systems to monitor and analyze continuous streams of data [3]. A prototype system BioStream was implemented to show its use in remote patient monitoring. BioStream is built on top of stream data processing architecture for real time processing of physiological signals. They use a database-oriented approach to analyzes data streams. The streams are subjected to “operators” that belong to a part of a patient plan. The current prototype is capable of identifying simple pathological conditions by monitoring ECG signals.

Bardram et al. [1, 2] developed a context-aware infrastructure to build context-aware applications for a hospital environment. The infrastructure includes sensors to detect presence of the nurse in the room, a smart pill container, a smart hospital bed to identify the patient. Radio frequency identification (RFID) is used to detect the people and the medications being used. A context-aware electronic patient record was designed to present an user interface that adapts based on the current context. However, the work is in the preliminary stages and the focus is more on human computer interfaces.

Levine et al. [17] have proposed the development of a computer automated perioperative situational awareness system that captures and records data from various medical devices and provides an integrated display to allow the operating team to visualize the data. The focus of this work so far has been on data capture and facilitating data visualization to provide context sensitive information and improve real time access to data. Our system focuses on making use of the data captured to detect medically significant events and creating an EMR.

Though the individual components of our system such as the algorithms to analyze physiological data, stream processing of data have been studied in previous systems, to the best of our knowledge no system has yet been developed to create an EMR in the perioperative environment.

9 Conclusion

We presented a prototype of a context-aware system that analyzes data streams in an operating room to detect medically significant events and document them in an electronic encounter record. The system uses technologies like radio frequency identification to acquire contextual information such as resources used and the staff present in the OR. We explored the use of medical

history and effect of medicines on physiology and we conclude that such techniques help us detect complex and more meaningful medical events. The system can be extended to detect events over larger number of scenarios.

Acknowledgements This work would not have been possible without significant help and support from colleagues at the ORF project at UMMS. In particular, we would like to thank Ivan George, as well as Drs. Park, Kavick and Nagy. UMBC students Palani Kodeswaran and Patricia Ordonez helped with elements of the code and GUI. We would like to thank Col Grissom at the Air Force Simulation Center, University of Maryland Medical School, for the use of the HPS METI system to evaluate our system. We would also like to thank Dr. Geoffrey Kuzmich for the invaluable clinical advice on developing the medical rule base without which this work would not be possible. Finally, thanks are also due to Dr. Franklin and his students at University of California Berkeley, for their support and expertise in using TelegraphCQ.

This work supported in part by a subcontract from UM Medical School in context of the Traumapod Project.

References

1. Bardram JE (2003) Hospitals of the future ubiquitous computing support for medical work in hospitals. In: Proceedings of UbiHealth 2003 the 2nd international workshop on ubiquitous computing for pervasive healthcare applications
2. Bardram JE (2004) Applications of context-aware computing in hospital work examples and design principles. In: Proceedings of the 2004 ACM symposium on applied computing
3. Bar-Or A, Goddeau D, Healey J, Kontothanassis L, Logan B, Nelson A, Thong J (2004) BioStream: a system architecture for real-time processing of physiological signals. Technical report, Hewlett-Packard Labs
4. Bloom M (1994) Techniques to identify clinical contexts during automated data analysis. *J Clin Monit Comput* 10(1): 17–22
5. Bodenreider O (2004) The unified medical language system (UMLS): integrating biomedical terminology. *Nucleic Acids Res* 32:267–270
6. Briggs B (2003) Electronic records. protection on the road to patient safety. *Health Data Manag* 11(5):34–8, 40
7. Chandrasekaran S (2003) TelegraphCQ: continuous dataflow processing for an uncertain world. *CIDR 2003*
8. CIMIT (2006) For integration of medicine and innovative technology. <http://www.cimit.org/>, accessed April 2006
9. Coleman WP, Seigel JH, Giovanni I, DeGaetano A, Goardarzi S, Tacchino RM (1990) Probability and the patient state space. *J Clin Monit Comput* 7(4):201–215
10. Corazzon R (2006) Ontology, a resource guide for philosophers. <http://www.formalontology.it>, accessed March 2006
11. Edwards M, Moczygemba J (2004) Reducing medical errors through better documentation. *Health Care Manag* 23(4):329–333
12. Electronic Medical Records (2006) http://www.expertsystem.com/medical_errors.htm, accessed April 2006

13. Friedman-Hill E (2003) *Jess in action: Java rule-based systems*. Manning Publications Co., Greenwich, CT, USA. <http://herzberg.ca.sandia.gov/jess/>
14. Harrop P, Das R (2007) *RFID for Healthcare and Pharmaceuticals 2007–2017*. Technical report, IDTechEx. <http://www.idtechex.com/products/en/view.asp?productcategoryid=101>
15. Hunter J, McIntosh N (1999) Knowledge-based event detection in complex time series data. In: *AIMDM '99: Proceedings of the joint European conference on artificial intelligence in medicine and medical decision making*. Springer, London, UK, pp 271–280
16. Krol M, Reich DL (1999) The algorithm for detecting critical conditions during anesthesia. In: *CBMS '99: Proceedings of the 12th IEEE symposium on computer-based medical systems*. IEEE Computer Society, Washington, DC, USA, p 208
17. Levine W, Meyer M, Brzezinski P, Robbins J, Sandberg W (2005) Computer automated total perioperative situational awareness and safety systems. *Comput Assist Radiol Surg* 1281:856–861
18. Liu X, Corner MD, Shenoy P (2006) Ferret: RFID locationing for pervasive multimedia. In: *Proceedings of Ubicomp 2006*, pp 422–440
19. Agency for HealthCare Research and Quality (2000) *Medical errors: the scope of the problem*. Fact sheet, Publication No. AHRQ 00-P037. Rockville, MD. <http://www.ahrq.gov/qual/errback.htm>
20. Mkivirta A, Sukuvaara T, Koski E, Kari A (1993) A knowledge-based alarm system for monitoring cardiac operated patients: technical construction and evaluation. *J Clin Monit Comput* 10(2):117–126
21. Nagargadde A, Srihar V, Ramamritham K (2007) Representation and processing of information related to real world events. *Knowl-based Syst* 20(1):1–16
22. Navabi M, Mylrea K, Watt R (1989) Detection of false alarms using an integrated anesthesia monitor. In: *IEEE engineering in medicine and biology society 11th annual conference*
23. Orchard B (2006) Fuzzy sets. http://www.iit.nrc.ca/IR_public/fuzzy/fuzzyJDocs/FuzzySet.html, accessed April 2006
24. Rector A, Rogers J (1996) The GALEN ontology. In: *Medical Informatics Europe (MIE 96)*
25. Roberts J, Stalow C, Hedges J (2005) *Clinical procedures in emergency medicine*. Elsevier
26. Roediger J, Salmon P (2006) Making changes in charts: do's and don'ts. <http://www.muschealth.com/professionals/ppd/chartshanges.htm>, accessed November 2006
27. Satava R (2005) Telesurgery, robotics, and the future of telemedicine. *Eur Surg* 37(5):304–307
28. Schecke T, Langen T, Popp H, Rau G, KŁsmacher H, Kalf G (1993) Knowledge-based decision support for patient monitoring in cardioanesthesia. *J Clin Monit Comput* 10(1): 1–11
29. Sukuvaara T, Koski E, MŁkivirta A, Kari A (1996) Clinicians' opinions on alarm limits and urgency of therapeutic responses. *J Clin Monit Comput* 12(2):85–88
30. Sutherland JV (2004) *RECIPE for real time process improvement in healthcare*. In: *13th Annual PHYSICIAN-COMPUTER CONNECTION symposium*
31. Tsien CL (2000) Event discovery in medical time series data. In: *American Medical Informatics Association (AMIA) symposium*
32. Tsien CL (2000) *Trend finder: automated detection of alarmable trends*. Technical report, Massachusetts Institute of Technology
33. Thompson TG, Brailer DJ (2004) *The decade of health information technology: delivering consumer-centric and information-rich health care—framework for strategic action*. Technical report, US Department of Health and Human Services



Sheetal Agarwal is a technical staff member at IBM India Research Laboratory. Her research interests include mobile and pervasive computing environments and information and communication technologies for developing regions. Sheetal earned an MS in computer science from the University of Maryland, Baltimore County.



Anupam Joshi is professor in the Department of Computer Science and Electrical Engineering at the University of Maryland, Baltimore County. His technical interests include mobile and networked computing, intelligent systems, data mining, and the Semantic Web. Joshi received a Ph.D in computer science from Purdue University. He is a member of the ACM, the IEEE, and the IEEE Computer Society.



Tim Finin is a Professor of Computer Science and Electrical Engineering at the University of Maryland, Baltimore County. He has over 30 years of experience in the applications of artificial intelligence to problems in information systems and intelligent interfaces and is currently working on social media, the semantic web, intelligent agents and mobile computing. He holds degrees from MIT and the University of Illinois and has also held positions at Unisys, the University of Pennsylvania, and the MIT AI Laboratory.



Tim Ganous is the Executive Director of the Institute for Defense Research Participation, a not-for-profit organization based in Alexandria, VA, that establishes programs for communities with unique economic development needs and aligning them with research and engineering goals of the U.S. Defense research establishment. Prior to IDRП, he was managing partner for Strategic Knowledge Solutions, LLC specializing in decision support and project management related to the introduction of new technologies into healthcare organizations. Mr. Ganous has been involved in the healthcare research and healthcare administration for twenty years. He received a BS degree in Political Science from the University of Delaware and a MPA in Health Planning from the State University of Pennsylvania.



Yelana Yesha is a Professor of Computer Science and Electrical Engineering at the University of Maryland, Baltimore County. She has 20 years of experience in research and development of advanced database and information systems. She received the B.Sc. degree in Computer Science from York University, Toronto, Canada in 1984, and the M.Sc. and Ph.D degrees in Computer and Information Science from The Ohio State University in 1986 and 1989, respectively.