Closing the Loop in ICU Decision Support: Physiologic Event Detection, Alerts, and Documentation

Patrick R. Norris, M.S.[+] and Benoit M. Dawant, Ph.D.[⊕]
Departments of [+]Biomedical Engineering, [⊕]Electrical Engineering and Computer Science
Vanderbilt University, Nashville, Tennessee
http://simon.project.vanderbilt.edu

ABSTRACT

Automated physiologic event detection and alerting is a challenging task in the ICU. Ideally care providers should be alerted only when events are clinically significant and there is opportunity for corrective action. However, the concepts of clinical significance and opportunity are difficult to define in automated systems, and effectiveness of alerting algorithms is difficult to measure. This paper describes recent efforts on the Simon project to capture information from ICU care providers about patient state and therapy in response to alerts, in order to assess the value of event definitions and progressively refine alerting algorithms. Event definitions for intracranial pressure and cerebral perfusion pressure were studied by implementing a reliable system to automatically deliver alerts to clinical users' alphanumeric pagers, and to capture associated documentation about patient state and therapy when the alerts occurred. During a 6-month test period in the trauma ICU at Vanderbilt University Medical Center, 530 alerts were detected in 2280 hours of data spanning 14 patients. Clinical users electronically documented 81% of these alerts as they occurred. Retrospectively classifying documentation based on therapeutic actions taken, or reasons why actions were not taken, provided useful information about ways to potentially improve event definitions and enhance system utility.

INTRODUCTION

Effective medical care processes typically embody a feedback loop, in which care providers continually assess patient condition and take action to improve it. Physiologic data from bedside monitors is one indicator of patient condition, and is a factor in 13-22% of clinical decisions made during ICU rounds[1]. Computerized decision support systems have been developed to monitor physiologic data and alert care providers when events of possible clinical significance occur. However, these systems generally reflect only the information delivery portion of the "loop", in that alerts are delivered but there is no facility for capturing information about related actions or relevance of the alert. Information about patient state and related therapeutic actions at the time of alerts is invaluable if event definitions and are to be progressively improved in a scientific way.

Without such refinement, current systems for physiologic event detection and clinical alerting remain inadequate. One study in 1997 found that only about 23% of physiologic alerts from based on heart rate threshold alarms from physiologic monitors were clinically relevant[2]. While ICU monitoring technology is relatively advanced in terms of technical architecture, information display, variety of sensors, and interfaces to other bedside devices, alert definitions remain primarily restricted to specifying high or low limits of individual monitor parameters, independent of time. Given the substantial variability in patients and clinical environments, providers are faced with a difficult tradeoff in setting these limits: either set a wide range of acceptable values to minimize false alarms, potentially at the expense of timely notification, or set the range narrow to receive earlier notification, at the expense of considerable false-positive alerts. While a variety of event-detection solutions have been proposed including multi-state filters[3], template recognition[4], and fuzzy logic process models[5], these and other advancements are typically not assessed in terms of the relevance of individual alerts generated during actual patient care. Notable exceptions include work by Tate *et al.* to examine effectiveness of alphanumeric pager alerts based on critical lab values[6], and by Shabot and colleagues to study alerts delivered to wireless devices based on physiologic, laboratory and other data[7]. This paper describes recent work on the Simon (Signal Interpretation and Monitoring) project to provide physiologic event detection, alert notification, and documentation capabilities in a working ICU information system, and to study how data entered by clinical users in response to alerts can be used to assess and improve system performance.

METHODS

Since the main purpose was to deliver alerts to care providers over the course of patient care and to capture feedback as alerts occurred, the first step was to implement an architecture that could support progressive development of event detection and alerting mechanisms, while maintaining a level of reliability sufficient for routine clinical use. Existing architectural components[8,9] were deployed, and several new components were added. A schematic of this architecture is show below in Figure 1.

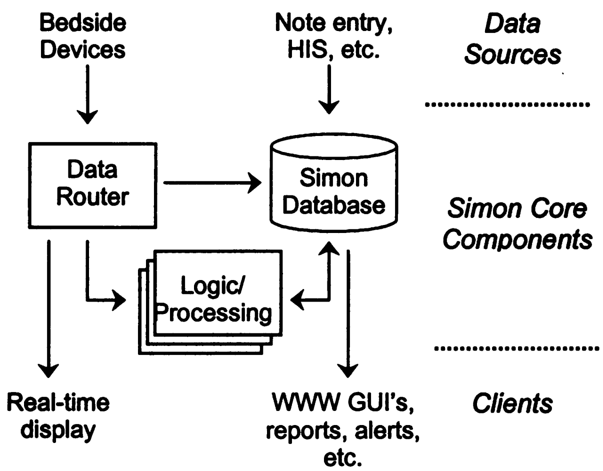

Figure 1: Simon Architecture

For performance reasons, the architecture separates high bandwidth data streams generated by bedside devices from other data channels such as user notes and hospital information system (HIS) data. The central point of access for all bedside device data is a *data router*, which relies on a publish-subscribe mechanism implemented over TCP/IP sockets to rapidly relay data from bedside medical device interfaces ("publishers") to processing modules and other "subscribers" requiring immediate, continuous access to bedside device data sampled every few seconds. All other data access is accomplished through open database connectivity (ODBC) connections to a relational database on a separate machine, implemented using Microsoft SQL Server. Components developed as part of this work included:

1. A database archive module to store all physiologic data to the database in "batches";
2. Processing/logic modules to compress this data for long term storage, associate data with patients based on hospital census information, monitor and manage system components, notify the researcher if system problems occurred, and deliver alerts to care providers via email and/or alphanumeric pager;
3. A java-based note application for entering free text notes in response to alerts, modified from existing code developed at VUMC;
4. Clients for WWW display of current data, alphanumeric pager alerting, and recently generated alerts.

During the test period, the implementation of this architecture supported data collection and processing from up to six different medical devices on each of four trauma ICU beds, and a variety of other sources and clients. As of July 2001, six additional beds had been added.

After implementing the basic architecture, an attending trauma surgeon with over 20 years of critical care experience was asked to define a set of physiologic events for testing that he thought might have clinical significance. He choose to generally express events in the form of threshold conditions over time, for example *intracranial pressure > 25 mmHg for 15 minutes*. Some events did not have a duration requirement, such as *cardiac index < 2.5 l/min/m*^{2}. An event detector was developed by another member of the project team[9], and tested off-line on actual data. During initial testing, several factors were noted that influenced design of event detection algorithms as well as the choice of which events to study: 1) Some parameters were not always available, due to device configuration requirements that unit staff were trained to do; 2) Noise and artifacts would require more advanced processing than simply monitoring current values in the data streams; and 3) False positive alerts would likely be present even with better processing, due to external variables that could not be sensed by the system, including lab data and patient status (i.e. organ donor, DNR).

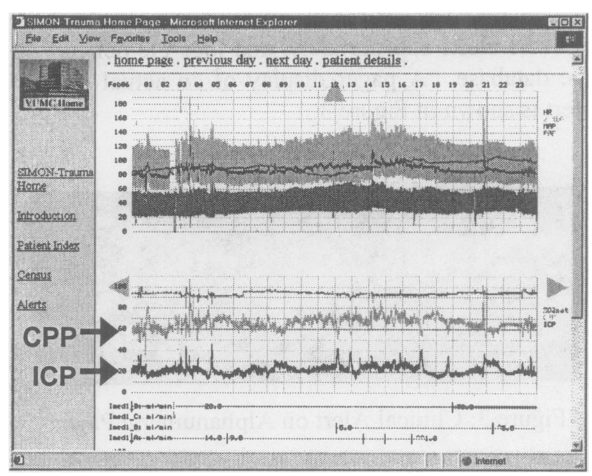

Figure 2: Portion of Simon WWW GUI (annotated)

As a result, initial efforts focused on detecting events in intracranial pressure (ICP) and cerebral perfusion pressure (CPP), two parameters the system could reliably acquire at all times. The corresponding event definitions were:

ICP > 25 mmHg for 15 minutes
CPP < 60 mmHg for 15 minutes

Typical ICP and CPP signals are noted above in Figure 2, showing a full day of data for a single patient. Each arrow points to the respective threshold value. In addition, a multi-state detection algorithm was defined to improve performance in the presence of noise, short data dropouts, and other artifacts.

While the event detector was being tested, a note-entry application was implemented that allowed users to enter free-text notes in response to alerts. It was important to be able to link this application with existing information sources, and to be able to easily deploy it on three bedside laptops dedicated to the project. An existing template-based note entry application, written in Java at VUMC, was adapted to run via a web browser, and to directly interface with the Simon database via Java Database Connectivity (JDBC). Four nurses initially tested this application over a period of six weeks, and performance issues were addressed by adding additional memory to the three bedside laptops dedicated to the project.

Finally, mechanisms were needed to notify care providers of alerts and to tie event, alert, and note data together. Database tables were defined to store events, alerts, and notes, as well as configuration data such as who should receive alerts for a particular bed. When an event is detected and added to the database, a notification engine looks up any number of email or alphanumeric pager recipients for the particular bed, delivers alerts appropriately, and logs the delivery in the database. A pager display of an actual clinical alert is shown in Figure 3.

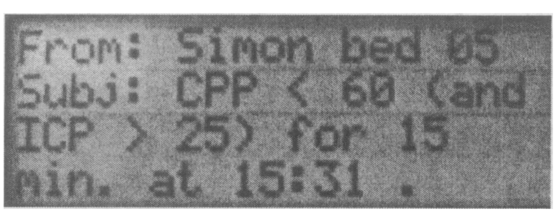

Figure 3: Clinical Alert on Alphanumeric Pager

A WWW page (Figure 4) lists all alerts generated over the past 12 hours and whether notes have been entered for each alert. Clicking on the alert hyperlink brings up the note entry application with patient demographics, timestamps, and alert information automatically entered.

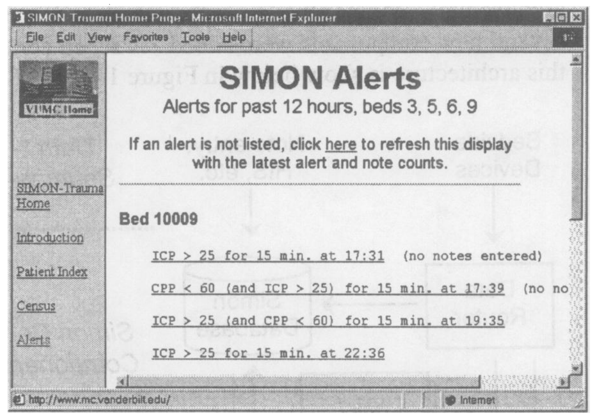

Figure 4: Alert status GUI

After all components were developed and tested, a group of clinical users was identified who would enter documentation in response to each alert. The VUMC Division of Trauma staffs a clinical nurse supervisor (CNS) position with a licensed nurse practitioner at all times, an ideal test group for several reasons. First, they are generally aware of the state of the most critically ill patients in the trauma unit. Second, their duties permit them to review alerts for any patient, assess the situation, and enter documentation in a timely manner. Most important, they generally have significant clinical experience and interest in the project, and were willing to act as "filters" between the system and the other members of the care team.

In July 2000 users attended a brief presentation and demonstration of the system, and the CNS on duty began receiving alerts via alphanumeric pager. After only a few days, they suggested an important refinement to the alert definitions. While users wanted to receive ICP alerts at all times, they only felt CPP alerts were significant in the presence of increased intracranial pressure, so the system was modified accordingly. Data from August 8, 2000 through Jan. 18, 2001 were reviewed. Data corresponding to patients that arrived in a Simon bed before the start of this interval, or left after the end, were not considered. For each patient with ICP or CPP alerts, the total monitored time was computed as the duration of ICP monitoring less data gaps greater than five minutes. Since an important aspect of the work was in assessing the feedback side of the process, documentation for each alert was subjectively classified into one of five areas according to the type of clinical action, or reason why no action was taken, in relation to the alert.

RESULTS

Over the study period, 530 ICP and/or CPP alerts were detected in 14 different patients, corresponding to approximately 2280 total hours of ICP data. Four additional patients had ICP/CPP monitoring for a significant time (> 30 minutes), but no alerts were detected. Figure 5 shows the incidence of ICP and CPP alerts by patients who had at least one alert.

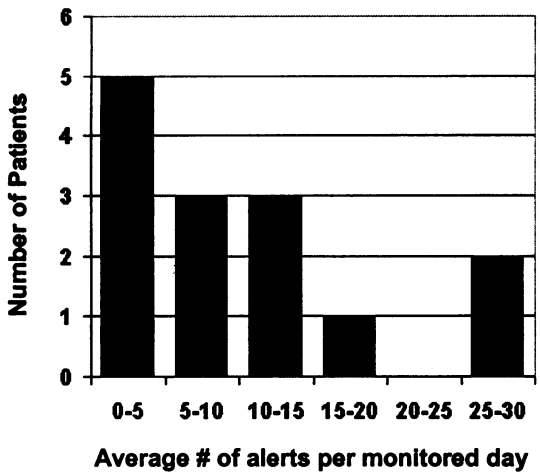

Figure 5: Frequency of ICP and CPP alerts in four trauma critical care beds, 8/00 – 1/01

Figure 6 shows the breakdown of documentation types entered in response to alerts during the study period.

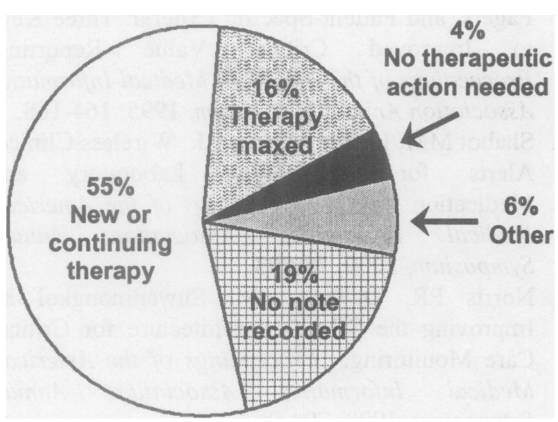

Figure 6: Types of documentation entered in response to 530 ICP and CPP alerts

While users were not given specific instructions about what type of information to include in notes, almost all documentation referenced new or continuing therapy related to the alert condition, or gave reasons why therapeutic actions were not taken. Such reasons included: 1) Medical therapy maximized – no additional medical treatment options were available due to patient lab values, or 2) No therapeutic action was needed because the condition spontaneously resolved, or a determination was made to discontinue treatment for ICP/CPP, as in the case of patients with do not resuscitate (DNR) orders. Six percent of notes could not be classified into these groups, and included a wide variety of information such as the patient being under care of another service, to technical feedback and suggestions. A few in this category were reports of disbelief by the bedside nurse that the alert condition occurred, when asked by the CNS who received the alphanumeric page. No notes were recorded for 19% of alerts, although almost all of these were for a single patient over a four-day period, indicating a short-lived technical or user issue that went undetected.

Technically, the system generally performed well during the test period. There was one known extended outage, of five days duration starting on 11/23/00. Outages were otherwise limited to a few hours per month on average, usually due to network and/or power glitches. Subjective user feedback was fairly positive, although users expressed dismay in a few cases where alerts continued after DNR orders were in place, as the system had no way of sensing this situation. Finally, successive notes included substantial repetition, discussed in more detail below.

DISCUSSION

Overall, the system was effective at providing clinical alerts to users, as well as capturing data about how alerts were related to clinical therapy or lack thereof. Users were remarkably good at ensuring a note was entered for each alert. If the one period mentioned above is omitted (where notes for one patient were not entered over a four-day period) more than 95% of all alerts were documented. Factors that contributed to this high response rate likely included a simple, easy-to-use interface for documentation, as well as substantial enthusiasm for the project from the trauma division director. From a technical standpoint, alerts were reliably detected for ICP and CPP because of fully automatic operation with no special connection or configuration requirements. Reliably detecting events in other signals, such as those from portable bedside devices that must be physically connected to the system each time or specifically programmed to send data, might be more difficult and require substantial user training.

Also, success in detecting alerts and eliciting documentation from users does not necessarily imply clinical usefulness. Subjective review of note content indicated that users probably considered many of the alerts redundant, due to substantial numbers of notes reading "ditto", "see previous entry", or similar phrases. This effect is in part due to the close coupling of ICP and CPP. CPP is the difference between mean arterial pressure and ICP, so an increase in ICP is usually accompanied by a decrease in CPP. In many cases this would trigger two alerts within a few minutes of each other, with similar user documentation. While it is difficult to assess utility given the fact that supervisors entered all data and were not specifically asked to rate usefulness in any controlled way, the presence of duplicate notes and the short time interval between the corresponding alerts implies some unnecessary redundancy in alert delivery.

However, note type and content may be used to prioritize system improvements to reduce this redundancy, making the system more suitable for routine use by bedside nurses. Note content was very important during initial testing, in terms of deciding to only deliver CPP events in the presence of an increased ICP event. During the study period, the types of alerts entered suggest additional enhancements. Since 16% of notes indicated that therapy was maximized based on clinical lab values, only notifying care providers when therapy could be resumed based on lab data, or by not sending alerts in cases where no therapy could be provided might improve usefulness. In addition, a facility to tell the system not to send alerts for patients with DNR orders would be helpful. Notes referencing corresponding therapy were most prevalent, suggesting that a mechanism to incorporate information about drug administration and other therapies might be most beneficial. In this case, the system would not generate alerts if appropriate therapy was being administered. Defining "appropriate therapy" in terms of computational algorithms may be challenging, requiring higher-level knowledge than the fairly simple event definitions described here. However, by closing the loop and evaluating not only alerts generated to care providers but also related therapeutic actions, such modifications can be progressively implemented and evaluated to improve performance.

Acknowledgements

Support for portions this work was provided by the VUMC Division of Trauma, and National Library of Medicine Training Grant LM00053-01A1.

The authors would like to thank John A. Morris, Jr., MD, for providing event definitions, and the nursing staff of the Vanderbilt University Trauma Unit for their participation in system evaluation.

References

1. Bradshaw KE, Gardner RM, Clemmer TP, Orme JF, Thomas F, West BJ. Physician Decision-making—Evaluation of Data Used in a Computerized ICU. *International Journal of Clinical Monitoring and Computing,* 1984; 1(2): 81-91.
2. Tsien, CL, Fackler JC. An Annotated Data Collection System to Support Intelligent Analysis of Intensive Care Unit Data. *In Advances in Intelligent Data Analysis,* X. Liu, P. Cohen, and M. Berthold (eds.); 111-121. Springer-Verlag, 1997.
3. Sittig DF, Cheung KH. A Parallel Implementation of a Multi-State Kalman Filtering Algorithm to Detect ECG Arrhythmias. *International Journal of Clinical Monitoring and Computing*, 1992; 9(1):13-22
4. Haimowitz IJ, Le PP, Kohane IS. Clinical Monitoring Using Regression-Based Trend Templates. *Artificial Intelligence in Medicine,* 1995; 7(6):473-96.
5. Becker K, Thull B, Kasmacher-Leidinger H, Stemmer J, Rau G, Kalff G, Zimmerman H-J. Design and Validation of an Intelligent Patient Monitoring and Alarm System Based on a Fuzzy Logic Process Model. *Artificial Intelligence in Medicine,* 1997; 11(33-53).
6. Tate KE, Gardner RM, Scherting K. Nurses, Pagers, and Patient-Specific Criteria: Three Keys to Improved Critical Value Reporting. *Proceedings of the American Medical Informatics Association Anuual Symposium,* 1995; 164-168.
7. Shabot MM, LoBue M, Chen J. Wireless Clinical Alerts for Physiologic, Laboratory and Medication Data. *Proceedings of the American Medical Informatics Association Anuual Symposium,* 2000; 789-93.
8. Norris PR, Dawant BM, Suwanmongkol K. Improving the SIMON Architecture for Critical Care Monitoring. *Proceedings of the American Medical Informatics Association Annual Symposium,* 1998; 794-98.
9. Suwanmongkol K. Simon: A Distributed Real-time System for Critical Care Patient Monitoring. *Masters Thesis in Electrical and Computer Engineering, Vanderbilt University.* 2001.