

A Platform for Real-time Acquisition and Analysis of Physiological Data in Hospital Emergency Departments

Jason B. Smith-*IEEE Member*, Andrew T. Reisner, Shwetha Edla, Jianbo Liu, Stephanie Liddle, and Jaques Reifman

Abstract— An opportunity exists for automated clinical decision support, in which raw source data from a conventional physiological monitoring system are continuously streamed to an independent analysis platform. Such a system would enable a wider range of functionality than offered by the source monitoring system. Although vendor solutions for this purpose are emerging, we developed our own system in order to control the expense and to permit forensic analysis of the internal core functionality of the system. In this report, we describe a platform that can provide decision support for trauma patients in an Emergency Department (ED). System evaluation spanned 39 days, and included a total of 2200 patient session hrs of real-time monitoring. We highlight the technical issues that we confronted, including protection of the core monitoring network, the real-time communication of electronic medical data, and the reliability of the real-time analysis. Detailing these nuanced technical issues may be valuable to other software developers or for those interested in investing in a vendor solution for similar functionality.

I. INTRODUCTION

Automated alerting and decision support are possible when hospitalized patients receive continuous monitoring, such as physiological alarms embedded within core monitoring systems. A newer form of decision support also exists in which vital-sign data are fed (e.g., using Health Level Seven [HL7] standards) to electronic medical records (EMRs). However, EMR decision-support algorithms usually operate on clinical data that have been reviewed and

verified by clinicians (as opposed to raw source data from the monitoring network).

A third opportunity for implementing decision support exists, in which the raw source data from the monitoring network are continuously streamed to an independent analysis platform, enabling a wider range of functionality than offered by the source monitoring system [1]. A number of emerging vendor solutions are available to support such functionality, including the BedMasterEx (BM) Data Acquisition Software (Excel Medical Electronics Inc., Jupiter, FL) with its StreamingAnalytics platform powered by IBM InfoSphere Streams (IBM, Yorktown Heights, NY); Bernoulli Enterprise Software (Cardiopulmonary Corp., Milford, CT); and DocBox (DocBox Inc., Newton, MA).

Our research team was interested in implementing a set of investigational decision-support algorithms that analyze streaming physiological data in realtime. Our focus is the development of decision support for trauma patients, and this project is named for its intended purpose: Automated Processing of the Physiological Registry for Assessment of Injury Severity in the Emergency Department (APPRAISE-ED). APPRAISE-ED is a follow-up to a similar system previously developed for prehospital air ambulances [2].

Also in prior work, we studied whether the BM system for data acquisition would have any identifiable harmful effects on the hospital's core monitoring network [3]. In that preliminary work, our testing did not reveal any deleterious impact on the core monitoring network, although a telephone poll of customers using the vendor's product revealed that a majority experienced at least one episode of unanticipated failure to archive data due to the difficulties in managing a distributed, network-based data acquisition system. This suggested that the system was safe and effective, but that the support and oversight necessary for the product were often underestimated by novice users.

The next step for our research team was in-hospital real-time data analysis, requiring a platform for acquisition and analysis of physiological signals for automated decision support. We considered using one of the aforementioned vendor solutions, but decided to develop our own solution because these were (in our subjective assessment) relatively expensive, without an established track record of good performance, and also lacked sufficient documentation for us to decisively evaluate their core functionality.

In this report, we describe our system, its evaluation, and the lessons we learned. For those who are considering the development or the purchase of such a system, this report may help them to better understand several key underlying

This work was supported by the U.S. Department of Defense Medical Research and Development Program (Grant No. D10_I_AR_J6_773) and by the Combat Casualty Care Research Area Directorate of the U.S. Army Medical Research and Materiel Command (USAMRMC), Fort Detrick, MD.

J. B. Smith is with the Department of Defense Biotechnology High Performance Computing Software Applications Institute (BHSAl), Telemedicine and Advanced Technology Research Center (TATRC), USAMRMC, Fort Detrick, MD 21702 USA (e-mail: jsmith@bhsai.org).

A. T. Reisner is with the BHSAl, TATRC, USAMRMC, Fort Detrick, MD 21702 USA and with the Department of Emergency Medicine, Massachusetts General Hospital, Boston, MA 02114 USA (corresponding author; phone: 617-726-2241; e-mail: areisner@partners.org).

S. Edla is with the BHSAl, TATRC, USAMRMC, Fort Detrick, MD 21702 USA (e-mail: sedla@bhsai.org).

J. Liu is with the BHSAl, TATRC, USAMRMC, Fort Detrick, MD 21702 USA (e-mail: jliu2@bhsai.org).

S. Liddle was with the Department of Biomedical Engineering, Massachusetts General Hospital, Boston, MA 02114 USA.

J. Reifman is a Senior Research Scientist and Director of the BHSAl, TATRC, USAMRMC, Fort Detrick, MD 21702 USA (e-mail: jaques.reifman.civ@mail.mil).

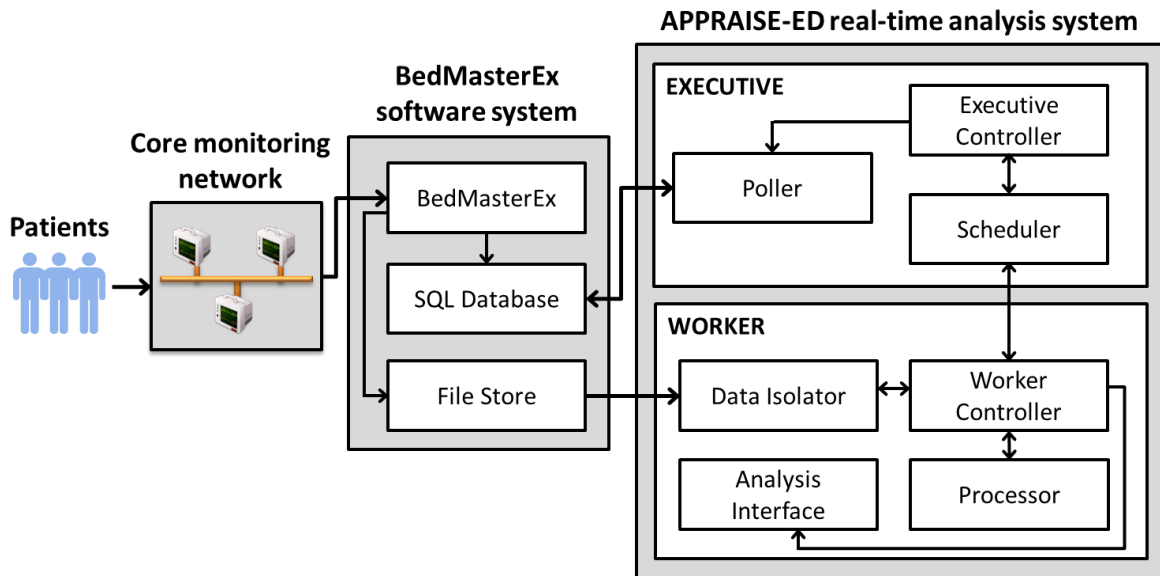


Figure 1. Illustration of the major components that comprise the proposed platform for the real-time acquisition and processing of physiological signals in the emergency department. APPRAISE-ED: Automated Processing of the Physiological Registry for Assessment of Injury Severity in the Emergency Department.

performance issues. We carefully examined two aspects of the system: first, reliability of data communication between the core monitoring network and the novel analysis platform, and second, the reliability of the real-time analysis.

II. METHODS

A. Description

Fig. 1 illustrates the complete system, consisting of three major components: the core monitoring network, which is a system of 16 Solar patient monitors (General Electric [GE], Milwaukee, WI); the proprietary BM software system; and our APPRAISE-ED system.

BM is hosted on a dedicated personal computer (PC) and, as per the manufacturer's specifications, collects physiological data from the GE patient monitors and saves it to binary (STP) files archived on the Windows file system. In parallel, BM archives associated data in an SQL Server (Microsoft Corp., Redmond, WA) database [4] archives associated data, including patient identifying information (i.e., protected health information [PHI]), monitor status, and indices of how data are stored within each STP file. The BM system standardizes waveform and vital-sign data frequencies to 240 and 1 Hz, respectively, and provides a common name for signals, thus offering a uniform means for access and analysis of collected data at the expense of averaging or missing higher frequency sensor data.

The APPRAISE-ED software runs on a dedicated server. The software consists of an executive module responsible for managing overall APPRAISE-ED system functionality. Also, for each monitor/bay under analysis, it creates an instance of a dedicated worker module.

The executive module includes three key sub-modules. The executive controller sub-module directs the tasks to be performed by the other sub-modules. The poller sub-module determines which monitor/bays are in active use by querying the BM SQL Server database. In addition, as new

physiological data are continuously accumulated through time, the poller sub-module determines the timing information about when each BM STP file was updated, as well as the location of the new data within the updated STP file. The poller's queries are performed every 5 s (query intervals are configurable), and this information is passed on to the scheduler sub-module.

The scheduler sub-module creates instances of the worker modules for each monitor/bay in active use. The scheduler also coordinates the transfer of newly acquired physiological data to each worker module, depending on its estimate of the time required to perform each analysis.

It is the individual worker modules which, as directed by the scheduler, are responsible for actually extracting physiological data from BM and then performing analyses. The worker controller sub-module in each worker directs the operations to be performed by each of the sub-modules. Each worker's data isolator sub-module locates and extracts the newest physiological data from the STP file. We rely on a software tool, Stp2Xml, available from the BM vendor, to transform the STP file into an intermediate data format. The data isolator sub-module then reads this intermediate XML file and deletes PHI from the extracted data. Finally, the data isolator stores the physiological data in a new portable binary file format (HDF: Hierarchical data format [5]).

Each worker's processor sub-module performs additional data handling and monitoring, such as tracking update times and verifying that new physiological data were obtained. The processor also checks the physiological data for indications that the patient was off-monitor and/or that a new patient may have been swapped into that monitor/bay. We found that, in a busy ED, there were often cases where patients were being removed from the monitor/bay without being electronically discharged, or patients were swapped without any update of the patient identity within the GE monitoring network. Accordingly, when no physiological data are detected for 5 min (configurable) the active session is

terminated. If data re-appear after that, a new session handled by a new instance of the worker module is started.

Each worker's analysis interface sub-module communicates with a MATLAB process (The Mathworks Inc., Natick, MA) to perform an analysis of isolated data via the MATLAB application programming interface (API). Algorithms are used to determine the reliability of waveform (e.g., electrocardiogram) and vital-sign data (e.g., heart rate) [6-8]. A primary focus of our research is early identification of patients with hemorrhage, and we have investigated a methodology involving multivariate classification [9] with the sequential probability ratio test [10] (the latter is an established technique for identifying abnormal patterns in a series of repeated measurements). The analysis interface has the ability to simultaneously run multiple instances of the analytic algorithms on the data from a given session at the same time, so that results from different algorithms can be compared. We constructed an analysis viewer to allow for real-time viewing of analysis results from a remote, networked location. At the end of each session, the HDF files and the analysis results are saved on the storage server for post-hoc review.

In prior work, we examined the impact of BM on the function of the core GE monitoring network and its constituent monitors [3]. For the current project, it was a priority that any newer functionality should not pose any additional risk to the same core monitoring network. Accordingly, the core GE monitoring network remains an isolated network without direct connection to the internet or the APPRAISE-ED software, except indirectly through the BM host PC, which serves as a bridge between the two networks. The APPRAISE-ED software resides on a separate server which communicates (read-only) with the BM host PC through select open ports. Queries from the APPRAISE-ED software to the SQL Server database (on the BM host PC) are designed to be low impact, i.e., minimum number of queries. Other aspects of query design were to limit the size of returned data corresponding to a given query, and to minimize the number of connections to the database. The core GE monitoring network thus remains a closed network, with its data travelling first to the BM host PC and then, through highly restricted ports and read-only access, to the APPRAISE-ED server and storage systems. In order to minimize any exposure to malware, neither the BM host PC nor the APPRAISE-ED analysis server is used for any other purpose than the aforementioned data processing.

B. Testing

Overall, testing was similar to that employed by Khitrov et al. [2], although the current system adds the complexity of up to 16 bays/monitors to analyze at a time, rather than a single transport monitor. We tested the APPRAISE-ED software component by component by using a unit-testing approach that exercised each module and sub-module. Next, we tested integrated system function during real-time operation by implementing a "simulated" ED consisting of three networked GE monitors, a virtual BM installation bridging the GE network and the general laboratory network, and a virtual server installation hosting the APPRAISE-ED

software. Patients were simulated using Netech MiniSim 1000 (Netech Corp, Farmingdale, NY) patient simulators.

Real-time functionality, including BM data extraction, algorithm processing, and result archiving, were compared to offline analysis of the same raw data (sourced from the BM archive). We confirmed that the software met all the aforementioned functional specifications (see *Description* above).

We conducted an exploration of several potential failure scenarios. We attempted to operate the GE Solar monitors in unusual fashions (e.g., turning monitors off in the middle of a session; swapping patients without formally discharging the initial patient within the GE network, etc.). Also, we simulated network failure scenarios during an ongoing session, such as interrupting APPRAISE-ED access to the SQL Server by blocking the SQL Server port, and blocking access to the STP file location by unmapping the shared drive on the PC running BM.

After successful laboratory testing, the system was tested in clinical use, in the Massachusetts General Hospital's ED, where we compared the system's resultant HDF files and real-time analysis results to offline analysis of data sourced directly from the BM archive. We also reviewed Windows Performance Monitor to assess the function of the BM host PC during clinical use.

III. RESULTS

Here, we summarize the main results and present several notable findings. During laboratory testing, we confirmed that the software met all the aforementioned functional specifications. The system was able to begin and end analysis sessions when new patients were placed on or removed from the monitor. The system was able to process simultaneous patient sessions as intended. During simulated network communication interruptions, the system appropriately logged the events and recovered from those interruptions as designed.

We performed clinical testing on weekdays, mostly during morning hrs, over a span of 39 days. Over a total of 230 hrs (2200 patient session hrs) of real-time ED operation, we did not observe network communication errors between the components of the overall system. Real-time analysis was executed as per our design, on up to 16 monitors/bays simultaneously, without any operational errors.

Based on the data within the HDF files (both vital-sign numeric as well as waveform data), we ascertained that the data passed from BM in real time matched the source BM record, with one exception. We uncovered one trivial but consistent discrepancy in the data values for the pulse oximetry (SpO₂) waveform that were passed to the real-time system. Specifically, the first three samples of the waveform (representing the first 0.0013 ms) at the beginning of each 60-s segment received by the real-time system did not match the source BM data. Based on our internal analysis and in discussion with the BM vendor [11], we learned that the Stp2Xml data export tool applied a moving-window average after extracting SpO₂ waveform data excerpts from the BM STP file. This moving-window average distorted the data at the beginning of the excerpt where it lacked preceding SpO₂

waveform for proper averaging. We addressed this issue by using the Stp2Xml tool to transform the previous SpO2 waveform excerpt along with the current segment so all data within the BM extraction window would be available.

When we compared real-time analysis in the ED (the results of which were available within the HDF data archive) versus offline retrospective analysis of the same source data (as archived by the BM system), there was convincing agreement. The mean standard error between these methods was 0.00. There were no episodes of unusual operation of the BM host PC, according to the logs of the Windows Performance Monitor.

IV. DISCUSSION

We have successfully developed, validated, and deployed the APPRAISE-ED system for prospectively testing real-time decision-support algorithms during clinical operations. The goal of this report is to highlight the technical issues that we confronted. Details of these issues may be valuable to other software developers or to those interested in procuring a vendor solution for similar functionality, which is often a six-figure investment. These issues may not be readily apparent to clinicians, administrators, and researchers interested in acquiring this functionality.

First, we felt it was important to carefully consider the integrity of the core GE monitoring network. In prior work, we assessed whether the BM archiving system could alter the functionality of the core GE monitoring network [3]. By design, the newly added functionality did not interact with the core monitoring network, but used the BM host PC as the indirect communication bridge through which real-time physiological data were obtained. Interactions between the APPRAISE-ED server and the BM host PC were kept to a minimum (i.e., read-only access, query frequency minimized, query date returns minimized, and working within elevated levels of BM host PC security).

Second, we felt it was important to consider the reliability of the communication between the software components. In our system design, we added functionality to log interruptions and gracefully recover from such interruptions automatically. We also identified at least one condition in which the data passed from the BM system in real time were not exactly the same as the data actually archived by BM for retrospective analysis. Although the differences were trivial, this issue—the integrity of data communicated for real-time analysis—is not trivial. It will be essential to consider this issue for any and all interoperable systems if such healthcare decision-support functionality becomes normative in the future.

Third, we felt it was important to carefully consider the validity of the real-time paradigm. Our paradigm involved “quasi” real-time processing, where there were brief but non-zero delays in the frequency of checking for new real-time data, then additional brief delays in processing those data. (We felt that delays of 2 min or less were acceptable when seeking to identify a physiological condition that is unlikely to progress substantially in that time frame.) Overall, we validated that this integrated system was able to perform as intended, within a 2-min analysis latency. Obviously, if such

real-time analyses become normative in healthcare, it is important to consider the “worst case scenario” in terms of latency for any decision support upon which tomorrow’s clinicians grow to depend. An important corollary is that the latency and overall performance are a function of the computational complexity of the algorithms; a system that performs suitably with one set of algorithms may not perform well with a different set of algorithms.

Lastly, although largely out of the scope of the current report, it is important to consider the “meta-data” of the system. For instance, is there any risk of clock error, or of associating data with the wrong patient? An expanded discussion of these issues can be found in [12].

DISCLAIMER

The opinions and assertions contained herein are the private views of the authors and are not to be construed as official or as reflecting the views of the U.S. Army or of the U.S. Department of Defense. This paper has been approved for public release with unlimited distribution.

REFERENCES

- [1] *Medical Device Plug-and-Play (MD PnP) Interoperability Program* [Online]. Available: http://www.mdppnp.org/uploads/Oct12_MD_PnP_White_Paper.pdf [accessed March 7 2014].
- [2] M. Y. Khitrov, M. Rutishauser, K. Montgomery, A. T. Reisner, and J. Reifman, "A platform for testing and comparing of real-time decision-support algorithms in mobile environments," *Conf Proc IEEE Eng Med Biol Soc*, vol. 2009, pp. 3417-20, 2009.
- [3] S. Liddle, L. Grover, R. Zhang, M. Khitrov, J. C. Brown, J. P. Cobb, J. Goldman, J. Chou, D. Yagoda, B. Westover, and A. T. Reisner, "Safety evaluation of a Medical Device Data System," *Conf Proc IEEE Eng Med Biol Soc*, vol. 2012, pp. 5899-902, 2012.
- [4] *BedMasterEx Operators Manual*, Excel Medical Electronics Inc., Jupiter, FL, 2010.
- [5] *The HDF Group. Hierarchical data format version 5, 2000-2014.* [Online]. Available: <http://www.hdfgroup.org/HDF5> [accessed March 7 2014].
- [6] L. Chen, T. McKenna, A. Reisner, and J. Reifman, "Algorithms to qualify respiratory data collected during the transport of trauma patients," *Physiol Meas*, vol. 27, pp. 797-816, Sep 2006.
- [7] A. T. Reisner, L. Chen, T. M. McKenna, and J. Reifman, "Automatically-computed prehospital severity scores are equivalent to scores based on medic documentation," *J Trauma*, vol. 65, pp. 915-23, Oct 2008.
- [8] C. Yu, Z. Liu, T. McKenna, A. T. Reisner, and J. Reifman, "A method for automatic identification of reliable heart rates calculated from ECG and PPG waveforms," *J Am Med Inform Assoc*, vol. 13, pp. 309-20, May-Jun 2006.
- [9] L. Chen, T. M. McKenna, A. T. Reisner, A. Gribok, and J. Reifman, "Decision tool for the early diagnosis of trauma patient hypovolemia," *J. Biomed. Inform.*, vol. 41, pp. 469-478, 2008.
- [10] L. Chen, A. T. Reisner, X. Chen, A. Gribok, and J. Reifman, "Are standard diagnostic test characteristics sufficient for the assessment of continual patient monitoring?," *Med Decis Making*, vol. 33, pp. 225-234, Feb 2013.
- [11] J. B. Smith, Department of Defense Biotechnology High Performance Computing Software Applications Institute, Fort Detrick, MD, private communication, July 16 2013.
- [12] A. T. Reisner, M. Y. Khitrov, L. Chen, A. Blood, K. Wilkins, W. Doyle, S. Wilcox, T. Denison, and J. Reifman, "Development and validation of a portable platform for deploying decision-support algorithms in prehospital settings," *Appl Clin Inform*, vol. 4, pp. 392-402, 2013.