

Next Generation Neonatal Health Informatics with Artemis

Carolyn MCGREGOR^{a1}, Christina CATLEY^a,
Andrew JAMES^b, James PADBURY^c

^aUniversity of Ontario Institute of Technology, Oshawa, ON, Canada

^bThe Hospital for Sick Children, University of Toronto, Toronto, ON, Canada

^cWomen & Infant's Hospital of Rhode Island, The Warren Alpert Medical School of Brown University, Providence, RI, USA

Abstract. This paper describes the deployment of a platform to enable processing of currently uncharted high frequency, high fidelity, synchronous data from medical devices. Such a platform would support the next generation of informatics solutions for neonatal intensive care. We present Artemis, a platform for real-time enactment of clinical knowledge as it relates to multidimensional data analysis and clinical research. Through specific deployment examples at two different neonatal intensive care units, we demonstrate that Artemis supports: 1) instantiation of clinical rules; 2) multidimensional analysis; 3) distribution of services for critical care via cloud computing; and 4) accomplishing 1 through 3 using current technology without a negative impact on patient care.

Keywords. neonatal intensive care, real-time analysis, clinical rules, multidimensional data, cloud computing

1. Introduction

Neonatal Intensive Care Units (NICUs) deploy state of the art medical devices to monitor and support premature babies; however, neonatologists are unable to process the vast quantities of both manually charted data and data collected from medical monitoring equipment. While there has been a sustained effort to move from paper to electronic charting in critical care, including NICUs, these initiatives have not improved the representation of information that can be derived from that charted data or the translation of that information to knowledge for earlier condition onset warnings.

Recent research is building a strong case for the benefits of real-time data analysis, with clinical events such as late onset neonatal sepsis (LONS) [1, 2] exhibiting early warning signs in physiological data before the clinical impact is sufficient to exhibit current clinical detection indicators. However, that research takes a condition specific, patient specific or physiological data stream type specific approach. The translation of that knowledge is another ‘black box’ clinical decision support system (CDSS) medical device at the bedside. Patients can develop multiple conditions concurrently or over time and each condition can have a set of behaviours with a pattern of occurrence. An infrastructure that can process currently uncharted higher frequency physiological data

¹ Corresponding Author. 2000 Simcoe Street North, Oshawa, Ontario, Canada, L1H 7K4;
E-mail: c.mcgregor@ieee.org

and support the earlier onset detection of multiple conditions has the potential to provide greater knowledge at the bedside than is available today and represents the next generation of informatics solutions for critical care. The provision of this knowledge requires a multidimensional approach as there are multiple conditions and multiple streams of data for which multiple behaviours can exist. In addition, new approaches are needed to enable processing and integration of both real-time synchronous medical device data and asynchronous clinical data to aid in clinical decision-making and improve outcomes for newborn infants.

We present the Artemis framework, a platform for real-time enactment of clinical knowledge as it relates to multidimensional data analysis and clinical research. First implemented at The Hospital for Sick Children, Toronto, in August 2009, Artemis has been running continuously since that time. We discuss, with examples, Artemis' multidimensional approach. Our goal is to provide a comprehensive description of the Artemis platform to date including the introduction of a cloud computing approach to enable distribution and support outsourced service of critical care.

2. Materials and Method

Artemis, shown in Figure 1, provides a flexible platform for the real-time analysis of time series physiological data streams extracted from a range of monitors to detect clinically significant conditions that may adversely affect health outcomes. The Data Acquisition component enables the provision of real-time synchronous medical device data and asynchronous Clinical Information Management System (CIMS) data. This data is then forwarded for analysis within the Online Analysis component which operates in real-time. For this real-time component, Artemis employs IBM's InfoSphere Streams, a novel streaming middleware system that processes data in real-time and then enables data storage within the Data Persistency component. It is capable of processing and then storing the raw data and derived data from multiple infants at the rate they are generated [3]. Stream processing is supported by IBM's Stream Processing Application Declarative Engine (SPADE) language, which is the programming language for IBM's InfoSphere Streams middleware. For the Knowledge Extraction component, Artemis utilizes a newly proposed temporal data mining approach [4]. This component supports the discovery of condition onset behaviours in physiological data streams and associated clinical data. New knowledge, once tested through rigorous clinical research techniques, is transferred for use within the Online Analysis through the Re-deployment component which translates the knowledge to a SPADE representation.

First, this paper tests whether the Artemis platform can enable the instantiation of clinical rules. Second it demonstrates how this platform can support multidimensional analysis. Third, we propose that this platform can be provided not only through an in-house installation but also through cloud computing providing a service of critical care [5]. This is particularly of interest for remote hospitals whose infrastructure for information technology technical support is much more limited than larger urban centre healthcare organizations. In this way, raw physiological streams and related clinical data can be transmitted securely over the Internet, with de-identified patient identifiers, for processing at the cloud-computing site. Finally, we show that the current technology is capable of supporting the platform without a negative impact on patient care. The Hospitals' Research Ethics Boards approved this research.

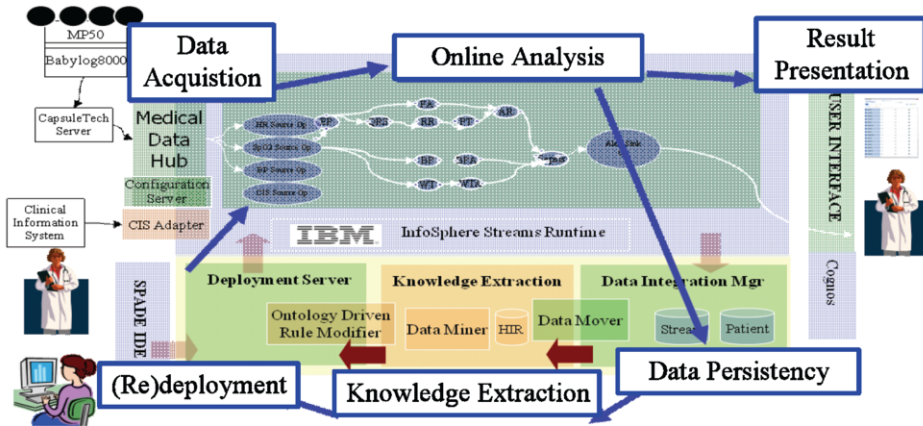


Figure 1. Artemis Framework (modified from [3])

3. Results

Our first implementation of Artemis is located at The Hospital for Sick Children (SickKids), Toronto. Real-time synchronous data is being acquired from the Philips Intellivue MP70 Neonatal monitors. Asynchronous data is being acquired from CIMS. Clinical protocols require that electrocardiogram derived HR (ECG-HR); transcutaneous oxygen saturation (SpO₂); respiration rate (RR); and impedance respiratory rate (IRW) data streams are constantly collected. When available, we also receive the systolic, diastolic and mean blood pressure. We have deployed SPADE code within the Online Analysis supporting our research into early detection of LONS. Data Persistence occurs at SickKids and data is replicated daily to the University of Ontario Institute of Technology (UOIT) where Knowledge Extraction research supports our clinical research into new earlier onset detection of LONS. Artemis has collected data on 174 patients, representing 4.1 patient years of data; all raw and derived data has been stored for retrospective research. Currently supporting eight concurrent patients and collecting approximately 1250 readings a second, Artemis at SickKids is deployed on three laptops: 1) for data acquisition; 2) for online analysis; and 3) for stream persistence. An incremental backup of the data is made to a persistence storage mirror at UOIT and used by the knowledge extraction component.

In April 2010 a second Artemis instance began collecting data from the Women and Infants Hospital in Rhode Island (WIHRI), USA. The WIHRI has successfully used a cloud-based deployment, where spot readings taken each minute are collected from the bedside SpaceLabs devices and fed in raw form to the Data Acquisition component, implemented in Mirth, of the Artemis platform running at the UOIT through a secure internet tunnel. In this setting all components of the platform are housed in the Health Informatics Research Laboratory at UOIT. Knowledge Extraction research supports our clinical research into new earlier onset detection of LONS on slower frequency physiological data. To date WIHRI has enrolled 203 patients, representing 10.6 patient years of data.

We have implemented a third installation of Artemis that contains only the Data Persistence, Knowledge Extraction and Redeployment components. Using the Knowledge Extraction component, we are performing retrospective data mining on a

dataset of nearly two years of 30 second spot reading data, obtained from 1151 patients, to further inform our refinement of a clinical rule for earlier detection of LONS which can ultimately be deployed by the Re-deployment component. We have successfully instantiated clinical rules though their implementation in SPADE for deployment by the Online Analysis component for LONS [3], apnoea [6], and hypoglycemia [7].

The three different implementations demonstrate that the platform can support multiple dimensions, shown in Figure 2, including: multiple locations, care providers, patients, conditions, data streams, and data stream behaviors. By care providers we mean that the platform can provide different temporal data summaries to different providers. For example, with apnoea, the neonatal nurse responds to alarms for extended respiratory pauses and falling SpO₂ and HR levels indicative of potential apnoea events. Our goal is not to generate further alarms for discrete events, but rather to create integrated temporal summaries of events from multiple data streams. For instance, a single mild apnoea event may not be clinically relevant; however, clusters of such events could be indicative of LONS and this information should be available to the neonatologist.




Figure 2. Multidimensional approach

Our current implementations at SickKids and WIHRI have no impact on care at the bedside. We are collecting and comparing when behaviours in the physiological streams are noted for comparison with current clinical observation and treatment practices. Due to the volume of data collected there were initial concerns expressed by the hospital's Information Services group about network traffic; however we have found that Artemis consumes less than 0.5% of network bandwidth. We have demonstrated that such a platform is capable of keeping up with the data collected at the speed at which it is received. The two other Artemis environments are currently running at UOIT spread across four servers.

4. Discussion

The Artemis platform uses currently available technology to support next generation health informatics, through online analysis and knowledge extraction of currently uncharted higher frequency data. In addition to the three implementations presented here, new implementations of the Artemis platform are in the planning stages for another NICU in Canada, as well as two NICUs in China and one in Australia.

Artemis provides clinical decision support in a flexible and transparent manner. Flexibility results from the ability to receive any asynchronous physiological data,

support the generation of multiple clinical rule representations as autonomous or interrelated SPADE graphs for Online Analysis, and perform multiple clinical research studies within Knowledge Extraction for clinical event analysis. The use of SPADE to represent the clinical rule enables transparency of the representation of the knowledge processing. This is in direct contrast to many CDSSs based on complex mathematical processing, such as artificial neural networks, which from the clinicians' viewpoint operate as black boxes. While a growing number of studies indicate that properly designed and effectively used CDSSs have the ability to improve quality of patient care [8], black box approaches raise concerns about the possible negative effects of CDSSs, including: potential de-skilling effects if system users do not understand how results were generated; a lack of flexibility and overly prescriptive outcomes; promoting over-reliance on software; and difficulty in evaluating outcomes [9]. Artemis is not a black box solution; rather it provides a means to instantiate clinical knowledge into the information processing pathway. From a clinical policy perspective, a number of international regulatory bodies are mandating that CDSSs require regulatory approval [10]. Canada has recently introduced new regulations classifying patient management software as a medical device that must be regulated [11]. The impact this will have on the clinician's ability to perform updates to CDSSs based on new evidence-based medicine is not yet clear.

Acknowledgements: This research is funded by the Canada Research Chairs program, Canadian Foundation for Innovation, an NSERC Discovery Grant, and an IBM First of a Kind award.

References

- [1] Flower AA, Moorman JR, Lake DE, Delos JB. Periodic heart rate decelerations in premature infants, *Experimental Biology and Medicine* 235 (2010), 531-8.
- [2] Griffin MP, Lake DE, O'Shea TM, Moorman JR. Heart rate characteristics and clinical signs in neonatal sepsis, *Pediatric Research* 61 (2007), 222-227.
- [3] Blount M, Ebling M, Eklund J, James AG, McGregor C, Percival N, et al. Real-time analysis for intensive care: development and deployment of the Artemis analytic system, *IEEE Eng Med Biol Mag* 29 (2010), 110-8.
- [4] McGregor C. System, method and computer program for multidimensional temporal data mining. Patent # 089705-0009; Canada, Gatineau Quebec (2010).
- [5] McGregor C, Eklund JM. Next generation remote critical care through service-oriented architectures: challenges and opportunities, *Service Oriented Computing & Applications* 4 (2010) 33-43.
- [6] Catley C, Smith K, McGregor C, James A, Eklund JM. A Framework to model and translate clinical rules to support complex real-time analysis of physiological and clinical data, *Proc. 1st ACM International Health Informatics Symposium* (2010), 307-315.
- [7] Kamaleswaran R. *A SOA method for the integration of heterogeneous data models for decision support*, Master's thesis (in progress), University of Ontario Institute of Technology (2011).
- [8] Wright A, Sittig DF, Ash JS, Sharma S, Pang JE, Middleton B. Clinical decision support capabilities of commercially-available clinical information systems, *JAMIA* 16 (2009), 637-44.
- [9] Open Clinical, DSS Success Factors (2005). Accessed January 2011 from: <http://www.openclinical.org/dssSuccessFactors.html>
- [10] Berner ES. Legal and regulatory issues related to the use of clinical software. In Greenes RA, ed. *Clinical Decision Support, The Road Ahead*. Elsevier Inc., 2007.
- [11] Health Canada, Software Regulated as a Class I or Class II Medical Device (2010). Accessed January 2011 from: http://www.hc-sc.gc.ca/dhp-mps/md-im/activit/announce-annonce/md_notice_software_im_avis_logiciels-eng.php