System of Systems Issues for the 2008 U.S. National Healthcare Information Network Remote Patient Monitoring Requirements

Elliot Sloane, Vijay Gehlot, Thomas Way, and Robert Beck
Villanova University
Villanova, Pennsylvania 19085 USA
Email: *{elliot.sloane, vijay.gehlot, thomas.way, robert.beck}@villanova.edu*

ABSTRACT *- This paper describes a number of new System of Systems Engineering (SoSE) issues that must be addressed in order to design and deploy a safe, secure, and private informatics infrastructure for remote patient monitoring systems that are interoperable with the emerging U.S. National Healthcare Information Network (NHIN). Motivations for NHIN's ambitious remote patient monitoring goals – such as reducing the cost of care and improving medical care quality for chronically ill patients – are introduced, and the technological requirements and challenges that arise are described. This paper demonstrates the use of SoSE modeling, simulation, verification and validation techniques similar to the emerging draft of INCOSE's SoSE Engineering Guide to improve the success of such projects. The tools used include event driven modeling and software simulation tools to assist in the design and evaluation of predictive software models that simulate key safety and performance aspects of proposed remote patient monitoring systems.*

Keywords- *healthcare information systems, NHIN, system of systems engineering, remote patient monitoring, telemedicine, population health*

I. INTRODUCTION

President Bush issued two Executive Orders in 2004 ad 2006 that mandated the development and purchase of interoperable medical information systems and medical devices [1], [2]. Fulfillment of these mandates falls to the U.S. Secretary of Health's new Office of the Coordinator for National Health Information Technology (ONCHIT), which is required to rapidly and iteratively design, develop, and implement a National Healthcare Information Network that will ultimately be able to provide nationwide electronic health records for all citizens and their physicians.

Since 2005, each year ONCHIT has managed the NHIN project process with a series of overlapping iterative, one-year analysis-design-prototype cycles. Each year, teams of clinicians, providers, and researchers specify clinical and operational goals and requirements in the form of "Use Cases" that are required [3] Next, several teams of cross-disciplinary team of clinical, management, and technology experts collaborate to identify the best available technical frameworks and standards to accomplish the key data interoperability, clinical capability, informatics goals – including all necessary security and privacy requirements – to allow each Use Case to be fulfilled [4]. It takes approximately a year to define and specify the standards needed to complete each Use Case, and approximately 4-6 Use Cases are developed in parallel by multiple teams every year. All told, each year, many tens of thousands of highly skilled labor hours are poured into the NHIN program, most of which is on a volunteer basis from professionals, societies, hospitals, universities, and vendors from across the country.

The following year all of the specified standards are put into pilot testing across the country to identify problems and necessary refinements while the next year's standards are under development.

2008 marks the beginning of the third round of annual Use Case development. For this year, ONCHIT decided that a major new goal for the third year of NHIN standards development to allow "Remote Monitoring (RM)" of patients outside of the hospital setting, to support the very important emerging field of telemedicine [5]. This Use Case will allow secure, private, and reliable electronic harvesting of data from tens of millions of Americans who are treated for expensive and complex chronic diseases at home or in nursing homes.

The value of RM is multi-fold, including improving timely and proper patient care by making accurate and complete data readily available to physicians and nurses, to detect and ameliorate the harmful and expensive risks of mistaken, inappropriate, or ineffective healthcare quickly, and to extend and enhance nationwide population health vigilance so that common threats like the flu and pneumonia, or new and potentially very aggressive risks such as an avian flu pandemic, can be detected and treated as and effectively as possible.

As designed by ONCHIT, the NHIN is by nature a complex System of Systems (SoS) engineering challenge because contemporary healthcare depends on the expertise of dozens of disparate clinical specialists and computer systems. Each physician is specialized, and must share (e.g., radiologist, cardiologist, or rheumatologist), hospital departments, and care-delivery-providers (e.g., hospital, physician office, or home care), each using specialized computer systems for optimal clinical data and practice management. In addition, telemedicine tools are creating an ever-expanding diversity of points-of-care, creating a growing number of smaller healthcare subsystems that extend to personal, consumer-based health care technologies [6].

All of the tasks that are under way focus on one main goal: to help create a complete, current, and constantly-updated Electronic Health Record (EHR) for every American citizen by 2014 to allow more effective, safe, and economical healthcare. This new remote patient monitoring task will present many new SoSE challenges to consider and overcome.

Fulfilling the mission of automated, timely, and accurate transfer of clinical data he EHR goal, though, would not in and of itself achieve all of the clinical needs for patients in remote settings. Indeed, the primary purpose of using remote monitoring for patients is to extend the services of clinicians to reach a larger number of patients, and to reduce overall healthcare costs by limiting use of expensive hospital facilities and resources to optimal, as-needed usage.

The scope of the RM use case is not constrained to any particular illness or setting, though it is not intended to address acute care settings like a hospital or Skilled Nursing Facility (SNF). There are at least three potential main groupings of applications in the non-acute setting, though, if one considers different levels of relative risk. For the purposes of our research, we have suggested the following:

Group A – Wellness monitoring. This group includes situations like remote weight monitoring for patients with eating disorders, Activities of Daily Life (ADL) motion monitoring of home-bound elders, and blood glucose monitoring for chronic but normally stable diabetic patients.

Group B – Vigilance/security monitoring. This group could include variety of potential alarm or alert messages identifying an unsafe, situation and whether that situation is slowly or rapidly emerging. These situations can naturally vary quite widely, and might be reasonably broken into two broad groups: slower-moving, lower risk situations, and faster-moving, higher-risk ones. In the slow-emerging/low-risk category would be problems such as reporting a slightly elevated temperature or respiration rate, starting the last day's supply of IV drug solution signaling a need for resupply, or gradually lowering oxygen concentrator flow levels, indicating a need for filter changes or other inspection or adjustment in the coming week. In the faster-moving, higher-risk category could be events like impending severe drug-overdose or dangerous drug-interaction messages, dead-battery alarms for home-based ventilator-dependent patients, or extremely high fever alarms infants.

Group C – Remote life-support monitoring. This type of setting could include more complex patient care, such as patients with a chronic degenerative illness such as ALS where home ventilator-care may be preferred until it is no longer a safe or viable option, when trying to provide multiple sessions of complex, multi-drug chemo-therapy regimen in a patient's home to avoid expensive and tiring trips to the hospital, or when providing apnea monitoring for a particularly high-risk newborn baby after release from a hospital. Group C situations could even potentially be used to support "hospital-like" home care in extreme cases.

Remote monitoring architectures that are adequate for Group A applications and the slowly-emerging/lower-risk risk situations from Group B might reasonably be handled in much the same way as email is handled telephone, cable, cell phone, or other well-accepted contemporary communication systems. The kinds of messages involved in these situations will not generally be life-critical, and delivery latencies of many minutes or even hours are not likely to cause any appreciable risk. We have lumped these types of situations into a Low Risk Remote Monitoring (LRRM) class to help distinguish them from the next group.

On the other hand, the Group C situations, and many of the rapidly-moving, higher-risk applications found in Group B may have much more stringent requirements. These situations have been lumped into a High Risk Remote Monitoring (HRRM) class for this discussion. For HRRM settings, it may be necessary, for example, to have a separate, dedicated communication line to the home to assure uninterrupted access, and, in the most extreme cases, a backup modality like a dedicated wireless cellular fall-back access facility may be needed.

LRRM situations will typically have fairly modest data demands. Hourly transfers of vital signs data like weight, heart rate, respiration rate, and temperature can be accomplished easily. HRRM situations may be much more complicated, though. For example, if a physician is faced with using telemedicine to attempt diagnosing a serious problem for a ventilator-dependent patient, for a period time he may wish to see continuous breathing, heart, blood oxygen saturation, and other near-real-time clinical data, waveforms, and alarms. This is a much more demanding system requirement, but may be inescapable in order to make decisions like whether or not to transport the patient back to the hospital, and, perhaps to determine if a helicopter-transport is necessary.

The 2008 RM Use Case is intended to facilitate more wide-spread application of a form of telemedicine that allows basic monitoring to be self-managed by patients or their families unless additional intervention is required. If most of the patients are stable and self-sufficient, RM provides both the basic data collection as well as a safety net that "calls for help" when needed. If patient problems happen infrequently, a local hospital might successfully support the remote patients without too much difficulty using existing telephone and/or computer systems.

If, however, the majority of patients are very sick, or seasonal problems like influenza suddenly cause a large, longer-duration surge in remote patient problems, more flexible information management systems may be needed. Such a system might employ several levels of artificial intelligence to help route messages to specialists at other nearby hospitals, at their homes, or even to Europe or India if necessary.

All of these RM solutions are highly dependent on telecommunications from the point of care (POC) to the provider. The research described in this paper focuses on using computer modeling and simulations tools to understand the telecommunication system requirements for creating and maintaining a safe and effective RM environment for Group A, B, or C situations. This paper also describes ways these tools can be used for verification and validation of the safety and performance of a telecommunication system and configuration options prior to purchase and/or following modifications to avoid unnecessary problems and expenses as well as limit potential risks to patients. These approaches are very consistent with the “multi-V-model” procurement approaches described in the draft SoS-Engineering Guide draft that is being developed by the International Council of Systems Engineering (INCOSE) [7].

II. SoSE MODELING AND SIMULATION

The scope and complexity of the proposed RM initiatives suggest a number of interesting SoSE issues that can and should be addressed in the design and deployment of the NHIN’s proposed RM services.

As described in the draft INCOSE SoSE Guide, the complete range of performance and limitations of systems that are themselves made up of many other systems is extremely difficult or impossible to analyze. One major reason for this challenge is the inter-dependencies between one or more systems may create novel, unanticipated problems in a far-distant, loosely-coupled subsystem.

Here is one example: a nursing home may invest in a small, desktop-computer-based smart infusion pump system that can provide drug-overdose vigilance for as many as 6 simultaneous IV pumps. The IV pumps themselves may not contain overdose calculation capabilities, but may depend on alarm messages from the vigilance system to notify the infusion management nursing team who can then decide when, or if, they should manually reduce or stop the infusion regimen. After a couple of years of successful, safe use, the nursing home might decide to expand its IV therapy services to include home-care IV therapies for as many as six homes that are within a short walking range of the facility. The home IV pumps could conceivably be connected to the nursing home’s overdose vigilance system using cellular wireless links in order to avoid telephone line contention. Everything might seem to work well initially, as long as no more than 6 IV pumps do not alarm simultaneously.

A year later, the nursing home may begin using a new, higher-concentration pain management drug that is designed to allow much lower infusion rates, which would also increase the times between IV pump refills. As a matter of convenience and cost management, the nursing home might decide to fill the IV pumps at the beginning of each week, so that pharmacy and nursing staff resources can be optimized. If most of the patient’s infusion rates are approximately the same, however, this combination of decisions could precipitate a new, very dangerous string of unanticipated faults. First, if all of the pumps are mis-programmed one morning, all of them create an alarm state for the vigilance system at the same time. If all 6 of the home-based systems try to send a wireless cellular message to the nursing home, but the nursing home’s cellular receiver may only be set up to handle two simultaneous incoming messages.

In the above scenario, until the combination of changes began to overwhelm the overall system, everything might seem fine. Once problems began to emerge, they could well be intermittent and very hard to pinpoint, depending on the sequence of events. For example, if the nursing home initially did trials with the high-dose drugs in the facility, no problem may be noted. If they extended the trial to a couple of outside homes, they still might not see a problem, especially if the total number of over-dose alarms never exceeded a total of six. They might, however, have a few unexplained “near misses,” when a seventh pump tried to send an alarm and only the patient themselves notice that nobody showed up to help them. Until the cellular phone link was overloaded, though, that fault path may be invisible.

The INCOSE guide describes similar issues in contemporary military procurement processes. It points out that future system changes can never be fully anticipated, and that even the performance of intended near-term system designs may be unpredictable if the system includes multiple newly-developed critical subsystems. The INCOSE guide suggests the best way to avoid serious SoS failures is to build models of each subsystem that can reasonably represent the known or intended characteristics and ranges of operation, and then link those models together to explore a much larger range of potential interactions. These modeling tools can be used to try to identify or anticipate unexpected interactions if they are put through a high-speed range of simulated normal and extreme operating conditions that might conceivably occur over many years of use and/or across many extreme applications.

For this approach to really be valid, each individual model must itself be validated against the real world requirements for each system, and then each model must be verified to ensure it actually performs as expected. As each subsystem is attached to other subsystems, the overall validation and verification process must be repeated, until the entire collection of subsystems that make up the overall system is itself known to be correct.

We have applied the above process of constructing verified and validated modeling and simulation tools for military Service Oriented Architecture applications on behalf of the US Air Force. Those tools have been based on the Colored Petri Nets (CPN) state modeling tool and the Extend discrete event modeling tools, and are described elsewhere in the literature [8], [9].

This present healthcare-focused research is using the same CPN and Extend tools to explore the RM Use Case, and is attempting to use a real-world cellular phone simuation tool, OpNet in order to add the capability to model and simulate the home-care setting as well.

To build the models, we recognized that any RM system will have to be designed to handle single- and multiple-component and system fault modes without endangering patients. This is true regardless of whether the patients are in the LRRM or HRRM situation, although the HRRM situation may require much more robust characteristics.

Also, RM depends on one or more telecommunication systems to transfer clinical and alarm data from the home, and they typically have to use whatever telecommunications services exist in the home. In rural settings like remote Alaskan towns, for example, the only choices might be using the phone lines, which are only suitable for low-speed modems. If higher-volume data transfer capability is needed, that might be accomplished by multi-hour data transfers through the night, but that would at best only support certain LRRM situations.

Many potential telecommunication network configurations could be used, and, across the country it is fairly likely that a huge number of network variations will be used. The simplest configuration might seem to be point-to-point communication using the Plain Old Telephone System, or POTS. This system is used for simple telephone calls or fax transmissions, but it is relatively inefficient because it ties up national, regional, and local wiring for very infrequent use. Today, more and more Voice over Internet Protocol (VoIP) and cellular voice and data services are using digitization techniques to simulate those services over the Internet or over wireless cellular phone networks. These approaches allow a large number of voice, fax, and computer data users to share the wired or wireless systems simultaneous. Unfortunately, VoIP and digital cellular networks may give rise to a new class of problems for RM applications. When using POTS, data or voice transmissions do not need to compete with other users and applications. Once VoIP or cellular wireless networks are deployed, though, data or voice communication can become delayed, choppy, distorted, or lost. Although this is annoying for voice users, and confusing or delaying for computer applications, VoIP could potentially create dangerous problems in HRRM applications.

These types of problems can be exacerbated if/when multiple LRRM and HRRM applications need to share the same VoIP, as might be the case in the large elder-living communities that are being developed. Such communities not only have lots of active seniors living in condos or apartments on the grounds, but also offer assisted living and nursing care areas as residents' health decreases and higher levels of assistance or nursing care are needed.

Our present research focuses on this type of setting, in which multiple LRRM and HRRM situations can reasonably be expected to be deployed simultaneously for the residents.

III. PROCESS MODEL SOLUTIONS

Solving these challenges will require controlled simulation, modeling, and deployment management so that the various interdependencies can be foreseen, detected, and accommodated as the system expands. Software tools like Colored Petri Nets (CPN), or Event-driven Process Models using tools such as OPNET, can help explore data loading in advance of actual deployment, especially for slow-moving illnesses like certain diabetes or chronic obstructive lung disease (COPD) patients. Furthermore, careful deployment to disaggregated points of care, such as nursing homes can allow system-level verification and validation of software tools while limiting or controlling confounding variables.

The CPN and Extend modeling techniques have already been used model and simulate a hospital-wide, wireless IEEE 802.11x patient monitoring alarm system of systems that incorporates realistic parameters and simulations to identify important issues in topological design, risk identification, and configuration management. In that research, life-critical alarm delays exceeding 2 minutes or more were observed, which, in real life could have threatened a patient's life.

Because remote patient monitoring systems are being adopted and deployed very rapidly, and they will have to use whatever wired or wireless communication modality that is available, including cellular wireless networks, they will be subject to the same life-threatening limitations that have been discussed above.

The CPN, Extend, and OpNet modeling and simulation tools that we are testing can help ensure that these new system applications can be thoroughly designed and tested before deployment, and the modeling and simulation approach reported in this paper can be used to identify potentially critical performance problems, reduce cost, and ameliorate patient risk in the design and deployment of a complex, wireless patient monitoring system of systems.

IV. CONCLUSIONS

The NHIN's remote patient monitoring project is ambitious, and its robust, correct and cost-effective implementation is intended to improve healthcare costs, effectiveness, and safety. Delivering safe and effective RM will require careful system of systems thinking and planning in order to succeed, because it is being done in an area where little prior research, experience, or expertise exists. Combining a careful consideration of issues and predictive analysis using state-of-the-art modeling and simulation techniques can ameliorate risk and provide an effective means to design a workable solution.

Future related research is feasible in other aspects of the RM initiative. One of the first issues is that the monitoring devices themselves are quite heterogeneous. The devices are manufactured by hundreds of vendors, and because they are provided to patients and nursing homes by thousands of private medical equipment rental companies across the country, the vendor, model number, configuration, and maintenance and calibration history varies widely. In addition, the use and programming of each device may be under the control of multiple parties, including the patient, family, visiting or staff nurses, or even non-clinical aides. Many new consumer-product companies claim to be poised to introduce new, low-cost “health monitoring” devices for home use, which could, in theory, make “wellness” monitoring available in virtually every home in the country [3].

The heterogeneity of the devices used for RM is also going to raise novel interoperability challenges. For example, the encoding of clinical and alarm data varies very widely from brand to brand because industry-wide standards are only now being developed [IHE]. Because data from multiple different clinical devices made by multiple different vendors may need to be integrated for an individual patient, system testing tools too examine probable and uncommon data set combinations could help explore and identify potential problems. NIST, in fact, has helped develop some software modeling and simulation tools for this purpose, though at present they only simulate a modest subset of potential real-world data streams.

As another example of future research interest, in a “perfect world,” it would be ideal if the remote patient monitoring data could simply be directly integrated into the patient’s EHR. In reality, however, the NHIN systems will need to be sophisticated enough to segregate, retain, and provide differential access to related-but-different patient monitoring data that spans hospital, nursing home, and at-home care. If the disparate sources of data cannot be segregated, life-threatening and inefficient misinterpretation may occur if the data is erroneously pooled for decision making. As an example, if it is determined that a patient’s home-owned blood pressure monitoring device is severely out of calibration, that data may need to be transformed or even deleted from certain reports and records. This issue will emerge more clearly later in 2008, and in the pilot testing of the RM standards that will occur in 2009. It is feasible that novel modeling and simulation tools can be developed to support this need.

There are also “emergent behavior” challenges and risks that can be quite substantial. In SoSE literature, emergent behavior generally refers to individual or community behaviors that may unexpectedly emerge following adoption of technology. In this RM telemedicine application, for example, it is not unlikely that some patients or home-care vendors will defer maintenance, repair, or supply restocking to the bare minimums, which, in some ways would seem consistent with many contemporary “lean engineering” approaches. If, however, these reductions are based on seat-of-the-pants estimates, they may fail to properly anticipate seasonal fluctuations and risks such as severe weather, labor and supply shortages, holiday communication system overloads, or common traffic gridlock periods. These emergent behaviors could, in the worst case, push some patients into the HRRM category due to external factors that could not be readily identified or remedied in the initial design.

Other reasonably likely emergent behaviors that have been seen in other healthcare settings include over-reliance on monitoring and alarm systems by patients or families to excuse longer and longer care-giver trips away from the patients that exceed the safety margins built into the system, over- or under-dosage of drugs by patients to try to save money or increase the perceived therapeutic (or side) effects, and even systemic “cheating” or lying, such as consuming candy or liquor up, but not quite beyond, the system alarm settings.

REFERENCES

[1] G. W. Bush. *Incentives for the Use of Health Information Technology and Establishing the Position of the National Health Information Technology Coordinator.* Executive Order 13335 of the President of the United States, 31 April, 2004.

[2] G. W. Bush. *Promoting Quality and Efficient Health Care in Federal Government Administered or Sponsored Health Care Programs* Executive Order of the President of the United States, 22 August, 2006.

[3] American Healthcare Information Communities (AHIC) Workgroups web site at U.S. Department of Health & Human Services. www.hhs.gov/healthit/ahic/workgroups.html.

[4] Healthcare Information Technology Standards Panel (HITSP), American National Standards Institute (ANSI), available at: www.ansi.org/hitsp.

[5] NHIN Use Cases at the Health Information Technology web site, U.S. Department of Health & Human Services. Accessed at: www.HHS.gov/healthit/usecases/.

[6] Continua Health Alliance, digital healthcare interoperability, available at: www.ContinuaAlliance.org.

[7] INCOSE System of Systems Engineering Guide draft, available at www.acq.osd.mil/sse/ssa/docs/SoS-SE-Guide-v1_0-WORKING-DRAFT.pdf

[8] E. B. Sloane, T. Way, V. Gehlot and R. Beck. *Conceptual SoS Model And Simulation Systems For A Next Generation National Healthcare Information Network (Nhin-2)*. Proceedings Of The 1st Annual IEEE Systems Conference, Honolulu, HI, April 9-12, 2007.

[9] V. Gehlot, E. B. Sloane. "*Ensuring Patient Safety by using Colored Petri Net Simulation in the Design of Heterogeneous, Multi-Vendor, Integrated, Life-Critical Wireless (802.x) Patient Care Device Networks*. Proceedings of the 27th Annual International Conference of the Engineering in Medicine and Biology Society, 2005. IEEE- EMBS 2005. Shanghai, China, 1-4 September 2005, pp: 162-165.