Medical Cyber Physical Systems

Insup Lee Oleg Sokolsky
Computer & Information Science
University of Pennsylvania
{lee,sokolsky}@cis.upenn.edu

***Abstract*— We discuss current trends in the development and use of high-confidence medical cyber-physical systems (MCPS). These trends, including increased reliance on software to deliver new functionality, wider use of network connectivity in MCPS, and demand for continuous patient monitoring, bring new challenges into the process of MCPS development and at the same time create new opportunities for research and development.**

I. Introduction

Medical device industry is undergoing a rapid transformation, embracing the potential of embedded software and network connectivity. Instead of stand-alone devices that can be designed, certified, and used to treat patients independently of each other, we will be faced in the near future with distributed systems that simultaneously control multiple aspects of the patient's physiology. The combination of embedded software controlling the devices, networking capabilities, and complicated physical dynamics that patient bodies exhibit makes modern medical device systems a distinct class of cyber-physical systems, which we refer to as medical CPS (MCPS).

Development of safe and effective MCPS will require new design, verification, and validation techniques, due to increased size and complexity. Model-based technology should play a larger role in the MCPS design. Models should cover devices and communications between them, but also, equally importantly, patients and caregivers.

In addition, MCPS will require a new regulatory procedure to approve their use for treating patients. On the one hand, the traditional process-based regulatory regime used currently by the U.S. Food and Drug Administration (FDA) to approve medical devices is becoming too lengthy and will be prohibitively expensive with the increased MCPS complexity. It threatens to overwhelm FDA resources to process and evaluate submissions in a manner that is both timely and rigorous. At the same time, device manufacturers are often operating under strict time-to-market pressures and soon small innovation-driven companies will not be able to afford the effort of preparing submissions to FDA.

In this paper, we summarize current technological trends in the area and discuss the challenges that arise from these trends. Our view is that the challenges brought by the recent trends present new opportunities for researchers in MCPS and in general embedded and CPS systems.

Research is supported in part by the National Science Foundation grants CNS-0834524 and CNS-0930647.

II. Background and Trends

While software-intensive medical devices such as infusion pumps, ventilators, and patient monitors have been used for a long time, the field of medical devices is currently undergoing a rapid transformation. The changes under way bring new challenges to the development of high-confidence medical devices, but at the same time they open new opportunities for the research community [14]. The main trends that have emerged can be summarized as follows:

a) New software-enabled functionality: Following the general trend in the field of embedded systems, introduction of the new functionality is largely driven by the new possibilities that software-based development of medical device systems is offering. A prime example of the new functionality is seen in the area of robotic surgery, which requires real-time processing of high-resolution images and haptic feedback. Another example is proton therapy treatment. It is one of the most technology-intensive procedures and requires one of the largest-scale medical device systems. Used to deliver precise doses of radiation for cancer patients, the treatment requires precise guiding of the proton beam from a cyclotron to patients, requiring adaptation to even minor shifts in position. Higher precision of the treatment, compared to conventional radiation therapy, allows higher radiation doses to be applied. This, in turn, places more stringent requirements on patient safety. Control of proton beams is subject to very tight timing constraints, with much less tolerance than for most medical devices. To further complicate the problem, the same beam is applied to multiple patient locations and needs to be switched from location to location, opening up the possibility of interference between beam scheduling and beam application. In addition to the proton beam control, a highly critical function of software in a proton treatment system is real-time image processing to determine precise position of the patient and detect any patient movement. In [19], the authors have analyzed the safety of proton therapy machines, however their analysis concentrates on a single system, the emergency shutdown. In general, proper analysis and validation of such large and complex systems remains one of the big challenges facing the medical device industry.

b) Increased connectivity of medical devices: In addition to relying more and more on software, medical devices are increasingly equipped with network interfaces. Interconnected medical devices, effectively, form a distributed medical device system of a larger scale and complexity that has to be properly

designed and validated to ensure effectiveness and patient safety. Today, the networking capabilities of medical devices are primarily used for patient monitoring (through local connection of individual devices to integrated patient monitors or for remote monitoring in a tele-ICU [20] setting) and for interaction with electronic health records to store patient data.

The networking capabilities of most medical devices today are limited in functionality and tend to rely on proprietary communication protocols offered by major vendors. There is, however, a growing realization among clinical professionals that open interoperability between different medical devices will lead to improved patient safety and new treatment procedures. Medical Device Plug-and-Play (MD PnP) Interoperability initiative [6] is a relatively recent effort that aims to provide an open standards framework for safe and flexible interconnectivity of medical devices, in order to improve patient safety and health care efficiency. In addition to developing interoperability standards, MD PnP initiative collects and demonstrates clinical scenarios where interoperability leads to improvement over the existing practice.

One example that illustrates how patient safety can be improved by MD PnP is the interaction between an X-ray machine and a ventilator. Consider the scenario taken from [15]. X-ray images are often taken during surgical operations. If the operation is being performed under general anesthesia, the patient is breathing with the help of a ventilator. Because the ventilator cannot "hold its breath" to let the X-ray image be taken without the blur caused by moving lungs, the ventilator has to be paused and later restarted. There have been cases where the ventilator was not restarted, leading to the death of the patient. Interoperation of the two devices can be used in several ways to ensure that patient safety is not compromised, as discussed in [3]. One possibility is to let the X-ray machine pause and restart the ventilator automatically. A safer alternative, although presenting tighter timing constraints, is to let the ventilator transmit its internal state to the X-ray machine. There typically is enough time to take an X-ray image at the end of the breathing cycle, when the patient has finished exhaling until the start of the next inhalation. This approach requires the X-ray machine to know precisely the instance when the air flow rate becomes close enough to zero and the time when the next inhalation starts. Then, it can make the decision to take a picture if enough time – taking transmission delays into account – is available.

c) Physiologically closed-loop systems: Traditionally, most clinical scenarios have a caregiver – and often more than one – controlling the process. For example, an anesthesiologist monitors sedation of a patient during an operation and decides when an action to adjust the flow of sedative needs to be taken. There is a concern in the medical community that such reliance on "human in the loop" may compromise patient safety. Caregivers, who are often overworked and operate under severe time pressure, may miss a critical warning sign. Nurses typically care for multiple patients at a time and can be distracted at a wrong moment. Using an automatic controller to provide continuous monitoring of the patient state and handling of routine situations would be a big relief to the caregiver and can improve patient care and safety. Although the computer will probably never replace the caregiver completely, it can significantly reduce the workload, calling the caregiver's attention only when something out of the ordinary happens.

Scenarios based on physiological closed-loop control have been used in the medical device industry for some time. However, their application has been mostly limited to implantable devices that cover relatively well understood body organs, such as the heart in the case of pacemakers and defibrillators. Implementing closed-loop scenarios in distributed medical device systems is a relatively new idea that has not made its way to the mainstream practice.

A clinical scenario that can easily benefit from the closed-loop approach is patient-controlled analgesia (PCA). PCA infusion pumps are commonly used to deliver opioids for pain management, for instance after surgery. Patients have very different reactions to the medications and require very different dosages and delivery schedules. PCA pumps give the patient a button to press to request a dose when they decide they want it rather than using a schedule fixed by a caregiver. Some patients may decide they prefer a higher level of pain to the nausea the drugs may cause and can press the button less often, while patients who need a higher dose can press it more often. A major problem with opioid medications in general is that an excessive dose can cause respiratory failure. A properly programmed PCA system should not allow an overdose because it is programmed with limits on how many doses it will deliver, regardless of how often the button is pushed. However, this safety mechanism is not sufficient to protect all patients. Some patients still receive overdoses if the pump is misprogrammed, if the pump programmer overestimates the maximum dose a patient can receive, if the wrong concentration of drug is loaded into the pump, or if someone other than the patient presses the button (PCA-by-proxy), among other causes. PCA infusion pumps are currently involved in a large number of adverse events, and existing safeguards such as drug libraries and programmable limits are not adequate to address all the scenarios seen in clinical practice [18].

In [4], we studied a PCA system that contains a supervisor to monitor patient data for the early signs of respiratory failure. The supervisor can stop the infusion and sound an alarm if the patient experiences an adverse event. We use a pulse oximeter device that receives physiological signals from a clip on the patient's finger and processes them to calculate heart rate and SpO_2 outputs[1]. Figure 1 shows the devices and essential data flow in this control loop. The pulse oximeter receives physiological signals from the patient and processes them to produce heart rate and SpO_2 outputs. The supervisor gets these outputs and makes a control decision, possibly sending a stop signal to the PCA pump. The PCA pump delivers a drug to

[1] SpO_2 is the measure of blood oxygenation, an important indicator of the heart activity.

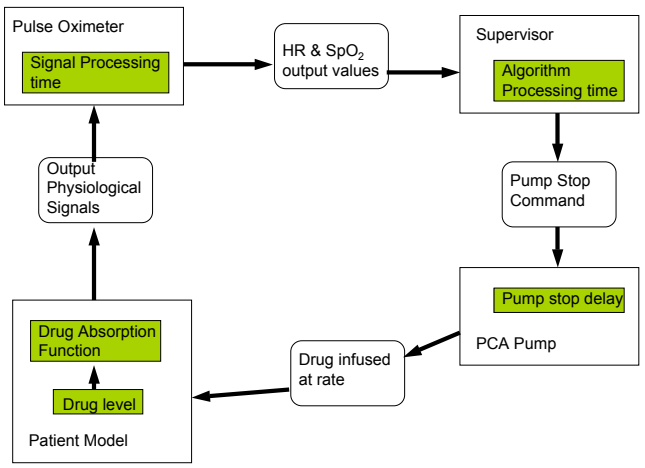

Fig. 1. PCA System Control Loop

the patient at its programmed rate unless it is stopped by the supervisor. The patient model gets the drug rate as an input and calculates the level of drug in the patient's body. This in turn influences the physiological output signals through a drug absorption function. Figure 1 also identifies the sources of delays in the control loop that the supervisor needs to account for. The supervisor also needs to be tolerant to faults that interfere with the control loop, in particular communication failures between the devices.

d) Continuous Monitoring and Care: Due to a high cost associated with in-hospital care, there has been increasing interest in alternatives such as home care, assisted living, telemedicine, and sport-activity monitoring. Mobile monitoring and home monitoring of vital signs and physical activities allow health to be assessed remotely at all times. Also, there is a growing popularity of sophisticated technologies such as body sensor networks to measure training effectiveness and athletic performance based on physiological data such as heart rate, breathing rate, blood-sugar level, stress level, and skin temperature. However, most of the current systems operate in store-and-forward mode, with no real-time diagnostic capability. Physiologically closed-loop technology will allow diagnostic evaluation of vital signs in real-time and make constant care possible.

III. Challenges and Opportunities

As can be seen from the trends described in Section II, the cross-cutting nature of Medical CPS (MCPS) transcends the informational, physical, and medical worlds, and raises significant scientific and technical challenges for the IT, medical, regulatory communities. Here are some of the challenges that are envisioned for the next ten years. These challenges in turn provide opportunities for R&D communities.

e) Executable clinical workflows: The trend towards increased interconnectivity and interoperability of medical devices opens the way for the dynamic construction and deployment of MCPS to implement custom clinical scenarios that best suit the needs of a given patient. Dynamism in MCPS deployment, in turn, poses a new challenge for ensuring patient safety in these custom scenarios. While safety analysis of dynamically created scenarios is an open problem, one can envision a possible path to the solution based on rigorous modeling of clinical scenarios and their subsequent analysis. A language for describing clinical scenarios should specify

- devices necessary for the implementation of the scenario;
- requirements for data flows between the devices and the patient;
- caregiver roles required for the scenario;
- operational procedures for each caregiver role, including actions necessary to compose the devices, program them (e.g., set infusion rates), and carry the scenario out; and possibly
- decision logic for the closed-loop control between devices.

Clinical scenario workflows should be given precise operational semantics. Analysis of such precise descriptions of a scenario will allow to make sure that instructions for caregivers are unambiguous and cover all possible situations; ensure that devices can interact with each other as desired; explore the effects of faults and user errors. Furthermore, a model of the scenario can be compiled into run-time components that will provide decision support for caregivers, detect device incompatibilities, and help recover from faults.

f) Model-based Development: With executable clinical workflow specifications, MCPS present a unique opportunity in the area of model-based development. We can introduce modeling beyond individual devices or even device systems, to the level of clinical scenarios that would serve as top-level system requirements. As discussed above, there is a trend towards dynamic composition and deployment of a medical device system for a given clinical scenario. Analysis of scenario models will allow us to assess patient safety in a scenario before a device system for the scenario is built, and generate requirements for the devices that can be safely used in a scenario implementation and interconnections between them. Such requirements can then be checked during deployment, ensuring safety of the implementation. The challenge, however, is precisely specifying the interface between static and dynamic safety checks.

g) Physiological close-loop control: The use of automatic control in clinical scenarios raises the stakes for the application of control theory in medical applications. Medical device systems for patients with complicated conditions may involve application of several treatments simultaneously, which affect several body systems in complicated and often insufficiently understood ways. These treatments also can interfere with each other. Effects of each treatment can differ widely from patient to patient. Critical variables are often not directly observable, adding to the uncertainty. Control-theoretic methods designed to operate under high parametric uncertainty, such as supervisory adaptive control [17], may be helpful in this context.

h) Patient Modeling and Simulation: A closely related challenge is that of patient modeling. Patient models are needed for the design of closed-loop control, as well as for the safety analysis of scenarios. For example, the closed-loop PCA scenario requires us to model drug absorption by the patient body as well as the relationship between the drug dose and concentration, on the one hand, and patient vital signs, such as heart and respiratory rates, on the other hand. Pharmacokinetic models of drug absorption are known from the literature (e.g., [16]), and there is statistical data on the effect of the drug on vital signs. However, comprehensive models are too complex to be used in the design and analysis, thus development of new abstraction techniques is paramount for addressing this challenge. At the same time, high-fidelity models and efficient simulators for them are needed for testing and validation of MCPS.

i) Adaptive Patient-Specific Algorithms and Smart Alarms: Medical devices are typically designed for groups of patients with similar medical conditions. However, the staggering range of patient responses to the same treatment and variation of vital signs for the same condition make this approach very generic and inefficient. For example, most medical devices are capable of triggering alarms when a potentially dangerous situation is detected. However, since alarm triggering conditions are aimed at an "average" patient, we are faced with a proliferation of false alarms that try to divert caregiver's attention to insignificant issues. The result is the well-known *alarm fatigue* that caregivers commonly experience, which makes them stop paying attention to device alarms and potentially missing important cases.

Network connectivity in medical devices and increasing availability of electronic health records (EHR) makes it possible to develop adaptive algorithms that will be attuned to the unique parameters of a given patient. For example, well-trained athletes can have heart rates that would be considered abnormal in most patients. Having the patient's exercise history from the EHR will let the system adjust alarm thresholds, reducing false alarms.

Another opportunity offered by interconnected devices is to correlate alarms raised by different devices in the system and use multivariate trends to provide "smart alarms," further reducing false alarms. To give a trivial example, a sudden drop in SpO_2 readings may mean that a patient is experiencing a heart failure. But if blood pressure readings remain normal, the more likely cause cause of the problem is a disconnected wire – which is a problem that needs to be fixed but has a much lesser degree of emergency than a heart failure.

j) User-Centered Design: Caregiver errors in using medical devices are a major source of adverse events [9], [23]. Undoubtedly, some of these errors are due to stress and overload that caregivers experience daily. However, a large number of these errors can be attributed to poor user interface design. If a device is hard to operate, has a counterintuitive interface, or responds to user inputs in an unexpected manner, user errors are much easier to occur. Design and validation of medical devices needs to take into account user expectations. To use model-based design for interactive medical devices, we have to incorporate models of caregiver behavior. Such user modeling is a notoriously challenging problem. However, incorporating information about likelihood of certain actions into caregiver models opens the way to quantitative reasoning about device safety.

k) Infrastructure for Medical-Device Integration and Interoperation: Currently, distributed MCPS are built by a single manufacturer using proprietary communication protocol. While this approach may make regulatory approval easier, it limits the benefits of inter-device communication and stifles creativity of medical professionals [11]. Open interconnectivity standards for MCPS, such as the ICE [5] standard proposed via the MD PnP initiative, lay the groundwork for medical device interoperability. Yet, for these standards to be effective, development and deployment platforms should be developed. A recent development in this area is the open-source MCDF toolset [13]. MCDF has been shown to be useful in the development of closed-loop MCDF. However, it is based on the relatively heavy-weight Java Messaging System. More nimble deployment platforms, designed to be amenable to verification, may ease regulatory approval of MCPS.

l) Compositionality: Since interoperable network-enabled medical devices will increasingly be composed into MCPS in a dynamic fashion, compositional reasoning is the only rigorous way to ensure safety of such systems. Techniques like temporal induction [21], which allows sound circular compositional reasoning between mutually dependent interacting devices in a scenario, may be useful in this case. A particularly challenging problem is predicting the possibilities of unexpected interactions between devices in the system. In particular, devices providing different treatments to the same patient may incur radio interference because of close proximity to each other. More importantly, treatments themselves can interfere with each other by affecting physiological responses [8]. Discovering these interferences, of course, is the subject of biomedical research. However, MCPS designers should be aware of these interferences and ensure that the system providing a treatment is made aware of potentially interfering treatments through sufficient context information.

As an illustration, consider the following "mixed criticality" scenario. Measurement of mean arterial pressure (MAP) depends on the relative position of the patient and sensor. Thus, when the patient's bed (which is a Class I medical device, lowest criticality in the FDA classification) is raised, the MAP reading changes. If the MAP sensor is part of a system providing continuous monitoring that follows trends in patient vital signs, the sudden change may trigger false alarms or unwarranted actions. This problem may be addressed by supplying additional context information to the monitoring system. For example, an additional event can be added to signify that bed height has changed. The system, then, can correlate that event with the change in the MAP reading and suppress the alarm. The design challenge is to identify and provide all sources of interactions as explicit inputs to the

system.

m) Security and Privacy: While networking capabilities let medical devices acquire functionality that was never possible previously, they also open the door to a host of new potential problems. Security and privacy concerns are some of those new problems [1]. An attacker who penetrates an MCPS network has the potential to harm and even kill patients by reprogramming devices [7]. The extreme approach, taken by most device manufacturers today, is to limit the functionality that can be invoked through the network interface. In most cases, the device can send out data, such as sensor readings or event logs, but not accept commands from the network. Although such an approach improves security of the system, it severely limits the ability to deploy closed-loop scenarios. Finding the right balance between flexibility and security is an important challenge for MCPS. Extensive solution have been developed to address security and privacy concerns for electronic health records. These solutions are currently being extended to MCPS [22], but the problem is far from being solved.

n) Verification, Validation and Certification: Current design practice places certification and verification at the end of the design cycle, when it is frequently too late to change design choices. As medical devices become more complex and more interconnected, it is becoming increasingly evident that verification and certification should be incorporated in early design stages. This can be done in two ways: on the one hand, the "design for verification" approach [2] can help verification techniques scale better and make generation of verification evidence easier. On the other hand, model-based generative techniques allow to perform verification early in the design and then extend the guarantees provided by verification to the implementation through code generation.

Throughout the domain of embedded and CPS systems, a new regulatory approach to certification has been advocated [12], based on collecting and reviewing evidence that the system achieves its goals. Model-driven techniques can help with the transition to evidence-based certification, from the current process-based approach. Using compositional modeling techniques and assume-guarantee reasoning may enable *incremental certification*, which would allow us to re-certify MCPS after component upgrades without reconsidering the whole assurance case from scratch.

IV. Conclusions

The domain of MCPS offers a unique set of challenges, distinct from any other CPS domain [10]. The area is about to undergo a substantial transformation, both in terms of doctors' and caregivers' expectations of what MCPS can do for them, and in terms of how these systems are developed and approved. The challenges facing MCPS are formidable, yet they present vast opportunities for research with immediate practical impact.

In this paper, we summarized the challenges and outlined the most promising, in our opinion, research directions. Modeling and model-driven engineering, which increasingly take hold in many other domains, will need to take the leading role in MCPS development as well.

References

[1] M. J. Ackerman, L. P. Burgess, R. Filart, I. Lee, and R. K. Poropatich. Developing next generation telehealth tools and technologies: Patients, systems, and data perspectives. *Telemedicine and e-Health*, 16(1):93–95, Jan/Feb 2010.

[2] K. Alexander and P. Clarkson. Good design practice for medical devices and equipment, Part II: design for verification. *Journal of Medical Engineering & Technology*, 24(2):53–62, 2000.

[3] D. Arney, J. M. Goldman, S. F. Whitehead, and I. Lee. Synchronizing an x-ray and anesthesia machine ventilator: A medical device interoperability case study. In *BIODEVICES 2009*, pages 52 – 60, January 2009.

[4] D. Arney, M. Pajic, J. Goldman, I. Lee, R. Mangharam, and O. Sokolsky. Toward patient safety in closed-loop medical device systems. In *Proceedings of the 1^{st} International Conference on Cyber-Physical Systems*, Apr. 2010.

[5] ASTM International. *STAM F2761-2009. Medical Devices and Medical Systems — Essential Safety Requirements for Equipment Comprising the Patient-Centric Integrated Clinical Environment (ICE), Part 1: General Requirements and Conceptual Model*, 2009.

[6] J. Goldman, R. Schrenker, J. Jackson, and S. Whitehead. Plug-and-play in the operating room of the future. *Biomedical Instrumentation and Technology*, 39(3):194–199, 2005.

[7] D. Halperin, T. Heydt-Benjamin, K. Fu, T. Kohno, and W. Maisel. Security and privacy for implantable medical devices. *Pervasive Computing*, 7(1):30–39, January-March 2008.

[8] M. B. Happ. Treatment interference in acutely and critically ill adults. *American Journal of Critical Care*, 7(3):224–235, May 1998.

[9] R. W. Hicks, V. Sikirica, W. Nelson, J. R. Schein, and D. D. Cousins. Medication errors involving patient-controlled analgesia. *American Journal of Health-System Pharmacy*, 65(5):429–440, March 2008.

[10] High Confidence Software and Systems Coordinating Group. High-confidence medical devices: Cyber-physical systems for 21st century health care. A Research and Development Needs Report, NCO/NITRD, February 2009.

[11] J. Hotchkiss, J. Robbins, and M. Robkin. MD-Adapt: A proposed architecture for open-source medical device interoperability. In *Joint Workshop on High Confidence Medical Devices, Software, and Systems and Medical Device Plug-and-Play Interoperability (HCMDSS-MDPnP)*, pages 167–170, July 2007.

[12] D. Jackson, M. Thomas, and L. I. Millett, editors. *Software for Dependable Systems: Sufficient Evidence?* National Academies Press, May 2007. Committee on Certifiably Dependable Software Systems, National Research Council.

[13] A. King, S. Procter, D. Andresen, J. Hatcliff, S. Warren, W. Spees, R. P. Jetley, P. L. Jones, and S. Weininger. An open test bed for medical device integration and coordination. In *ICSE Companion*, pages 141–151, May 2009.

[14] I. Lee, G. J. Pappas, R. Cleaveland, J. Hatcliff, B. H. Krogh, P. Lee, H. Rubin, and L. Sha. High-confidence medical device software and systems. *Computer*, 39(4):33–38, April 2006.

[15] A. S. Lofsky. Turn Your Alarms On. *APSF Newsletter*, 19(4):43, 2004.

[16] J. X. Mazoit, K. Butscher, and K. Samii. Morphine in postoperative patients: Pharmacokinetics and pharmacodynamics of metabolites. *Anesthesia and Analgesia*, 105(1):70–78, 2007.

[17] A. S. Morse. Supervisory control of families of linear set-point controllers – Part 2: Robustness. *IEEE Transactions on Automatic Control*, pages 1500–1515, Nov. 1997.

[18] T. K. Nuckols, A. G. Bower, S. M. Paddock, L. H. Hilborne, P. Wallace, J. M. Rothschild, A. Griffin, R. J. Fairbanks, B. Carlson, R. J. Panzer, and R. H. Brook. Programmable infusion pumps in ICUs: An analysis of corresponding adverse drug events. *Journal of General Internal Medicine*, 23(Supplement 1):41–45, January 2008.

[19] A. Rae, P. Ramanan, D. Jackson, and J. Flanz. Critical feature analysis of a radiotherapy machine. In *International Conference of Computer Safety, Reliability and Security (SAFECOMP)*, Sept. 2003.

[20] A. Sapirstein, N. Lone, A. Latif, J. Fackler, and P. J. Pronovost. Tele ICU: paradox or panacea? *Best Practice & Research Clinical Anaesthesiology*, 23(1):115–126, Mar. 2009.

[21] M. Sheeran, S. Singh, and G. Stålmarck. Checking safety properties using induction and a sat-solver. In *FMCAD '00: Proceedings of the Third International Conference on Formal Methods in Computer-Aided Design*, volume 1954 of *LNCS*, pages 108–125, 2000.

[22] K. Venkatasubramanian and S. K. S. Gupta. Security for pervasive healthcare. In Y. Xiao, editor, *Security in Distributed, Grid, Mobile, and Pervasive Computing*, chapter 15, pages 443–464. CRC Press, Apr. 2007.

[23] K. J. Vicente, K. Kada-Bekhaled, G. Hillel, A. Cassano, and B. A. Orser. Programming errors contribute to death from patient-controlled analgesia: case report and estimate of probability. *Canadian Journal of Anesthesiology*, 50(4):328–32, 2003.