

High-Confidence Medical Device Software and Systems

Insup Lee and George J. Pappas, University of Pennsylvania

Rance Cleaveland, University of Maryland

John Hatcliff, Kansas State University

Bruce H. Krogh and Peter Lee, Carnegie Mellon University

Harvey Rubin, University of Pennsylvania

Lui Sha, University of Illinois at Urbana-Champaign

Given the shortage of caregivers and the increase in an aging US population, the future of US healthcare quality does not look promising and definitely is unlikely to be cheaper. Advances in health information systems and healthcare technology offer a tremendous opportunity for improving the quality of care while reducing costs.

The United States spends about 16 percent of its gross domestic product on healthcare, twice the average of most European nations.¹ Given the shortage of caregivers and the increase in an aging US population, the future of US healthcare quality does not look promising and definitely does not look cheaper. Advances in health information systems and healthcare technology offer a tremendous opportunity for improving the quality of our healthcare while reducing healthcare costs.

Advances in computing, networking, sensing, and medical device technology are enabling the dramatic proliferation of diagnostic and therapeutic devices. These devices range from advanced imaging machines to minimally invasive surgical techniques, from camera pills to doctor-on-a-chip, from computerized insulin pumps to implantable heart devices.

Although advances in stand-alone diagnostic and treatment systems have been accelerating steadily, the lack of proper integration and interoperation of those systems produces systemic inefficiencies in healthcare delivery. This inflates costs and contributes to avoidable medical errors that degrade patient care. The use of software that controls medical devices to overcome these problems is inevitable and will help to ensure safe

advances in healthcare delivery. A critical concern, however, is the cost-effective development and production of reliable and safe medical device software and systems.

MEDICAL DEVICE SOFTWARE AND SYSTEMS

The development and production of medical device software and systems is a crucial issue, both for the US economy and for ensuring safe advances in healthcare delivery. As devices become increasingly smaller in physical terms but larger in software terms, the design, testing, and eventual Food and Drug Administration (FDA) device approval is becoming much more expensive for medical device manufacturers both in terms of time and cost. Furthermore, the number of devices that have recently been recalled due to software and hardware problems is increasing at an alarming rate. As medical devices are becoming increasingly networked, ensuring even the same level of health safety seems a challenge.

Several federal and regulatory agencies have identified this growing problem and are interested in establishing a research agenda directed at improving the design, certification, and operation of current and future medical device software and systems. The 2005 High-Confidence Medical Device Software and Systems (HCMDSS) workshop was sponsored by various fed-

eral agencies, including the FDA, the National Institute of Standards and Technology, the National Security Agency, and the National Science Foundation, along with the National Coordination Office for Networking and Information Technology Research and Development. The views expressed in this article are those of the authors, and not necessarily those of the federal sponsors.

The purpose of the HCMDSS workshop was to provide a working forum for leaders and visionaries from industry, research laboratories, academia, and government concerned with medical devices. More than 90 experts from these sectors attended the workshop. They represented a mix of the relevant stakeholders—including researchers, developers, certifiers, and users—who can help identify emerging systems and assurance needs.

This workshop's main goal was to develop a road map for overcoming crucial issues and challenges facing the design, manufacture, certification, and use of medical device software and systems. An additional goal was to identify and form a sustainable research and development community for the advancement of HCMDSS. Of particular interest was the crystallization of technology needs and promising research directions that could revolutionize the way HCMDSS are designed, produced, and validated in the future but that are beyond the range of today's devices because of time-to-market pressures and short-term R&D practices. The presentations of the working groups, keynote speakers, and panelists and the submitted position statements of participants in this workshop are available at www.cis.upenn.edu/hcmdss/.

CRITICAL ISSUES

In the course of this workshop, the participants identified six issues as critical for the future of high-confidence medical devices:

- *Foundations for integrating medical device systems.* Commercial off-the-shelf technologies do not produce highly distributed medical device systems with guarantees of security, privacy, robustness, interoperability, extensibility, mobility, and general patient safety. Advances in computing are instrumental in the development of novel diagnostic and therapeutic equipment and procedures and of widely accessible medical-record systems. Although diagnostic and treatment systems have advanced significantly, they do not work well together. The systemic inefficiencies in healthcare delivery grossly inflate costs and contribute to avoidable medical errors that degrade patient care.

Although networks of networked medical devices hold many promises and possibilities, they also create challenges.

- *Distributed control and sensing in networked medical device systems.* The networking of medical devices for distributed sensing and control can occur at many levels. These networks can collect data for offline analysis, generate alarms when critical conditions occur, or close feedback loops for the controlled delivery of drugs. Research is needed to create medical device networks with these features and to enable the diffusion of new sensing and control technologies as they become available.
- *Patient modeling and simulation.* Modeling has proved its value in many industries, such as aerospace, automotive, and chemical plants. It has fostered novel product development, increased safety parameters, and ensured cost-effective development phases, ultimately achieving regulatory approval. In the medical-practice domain, modeling and simulation will improve outcomes and quality of patient care, will provide better control of healthcare costs with improvements in prevention and intervention, and will allow maximal use of the electronic health records.
- *Embedded real-time networked system infrastructure.* Researchers envision next-generation medical systems to be a ubiquitous network of networked systems that provides secure, reliable, privacy-preserving, and cost-effective, personalized, high-quality healthcare. Although networks of networked medical devices hold many promises and possibilities, they also create challenges.
- *Medical device software development.* Many medical devices are, essentially, embedded systems. As such, software is often a fundamental—albeit not always obvious—part of a device's functionality. This means that any safety and regulatory requirements for medical devices necessarily call for rigorous software development methods to ensure reliability and to protect the public health. Exactly how to accomplish that is a question, particularly because devices and systems are becoming increasingly complicated and interconnected. We have reached the point where testing as the primary way to gain confidence in a system is impractical or ineffective. Furthermore, requirements and specifications based on medical practice are needed to help ensure that devices will perform appropriately.
- *Validation and certification.* Verification and validation tasks required for the approval of medical devices play a significant role in enabling the FDA to carry out its mandate of approving only “safe and effective” medical devices. Unfortunately, many industry observers believe that we are approaching the limits of current device certification processes. As devices

grow more complex and rely much more on embedded software to achieve critical functionality, existing certification processes are being stressed. This trend results in higher development costs for manufacturers, longer time to market, and increased chances of device failure—with associated recall or liability costs.

CURRENT STATE OF AFFAIRS

Each working group summarized the state of the practice, development, and research in its area, identified R&D needs and challenges, and provided a road map to address their needs and challenges.

Several observations were made about the state of the art in medical device software and system development:

- **Medical device software development.** As a whole, the medical industry does reasonably well in developing and approving stand-alone devices that have moderate complexity and are based on mature technology. In such cases, the domain is well understood and the technology furnishes examples of devices that have been approved. However, designing bug-free software is difficult, especially in complex devices that might be used in unanticipated contexts. Existing practices have worked as well as they have because industry verification and validation personnel and regulators take their jobs seriously.
- **Large-scale, complex devices stress current best practices.** We are still challenged by large-scale, complex devices, such as proton therapy facilities. For these types of devices, the validation procedures and test cases can number in the hundreds of thousands. The burden of validation—in time and costs—extends the time to bring devices to market. Because of time-to-deliver pressures and a shortfall in properly trained software engineers, the development of HCMDSS has—with few exceptions—not kept pace with software assurance techniques practiced in other safety-critical domains such as avionics.
- **Integration of MDSS.** Industry is doing fairly well at integrating products developed by a single manufacturer. Such integrations are largely proceeding in an ad hoc fashion, however, without standardized integration mechanisms that are commonplace in other domains, such as the highly successful and widely used universal serial bus from the personal-computer domain. Because the number of medical devices and systems that are to be networked and integrated is increasing significantly, we must develop standards and regulations for medical device integration.

We now seem to be on the cusp of the types of revolutionary changes in healthcare systems that have transformed other sectors of the nation's infrastructure and economy.

- **Device interference and interoperability.** Caregivers and clinical engineers report that as devices proliferate and as sophistication and connectivity in hospitals increase, we are becoming lost in a swirl of technology, and we face unanticipated interference between devices. A concerted effort to address interoperability has begun, aiming to develop plug-and-play interoperability standards for the operating room of the future. So far the main concern has been network standards; other essential issues, such as quality of service and semantic compatibility for interoperability, have not yet been addressed. Also, we need to conduct a systematic study of device interference during integration.
- **Certification.** FDA device approval relies on a process-driven approach, in which manufacturers obtain approval by showing that they have applied established quality assurance techniques to certain levels of coverage, using manual code inspections. In particular, collective knowledge and experience within large, well-established companies aid the effort necessary to prepare for the market. But when considering larger devices with relatively complex functionality, the time and costs associated with verification and validation tasks such as test generation and execution cause researchers to lose confidence in their ability to bring safe and effective devices to market.

We must also consider the effectiveness and already high costs of development and certification processes in the context of rapid advances in technology that have fundamentally changed the way many informational, financial, and scientific services are provided.

Although technological advances have contributed to a steady increase in the quality of healthcare, and although FDA approval processes have for the most part been able to keep pace, we now seem to be on the cusp of the types of revolutionary changes in healthcare systems that have transformed other sectors of the nation's infrastructure and economy. Such changes call for a paradigm shift in the development and certification of medical device software and systems.

For example, pervasive networking will enable the integration of national networks, regional healthcare centers, local hospitals and clinics, the offices of primary-care physicians, home computing, and body-area networks. The healthcare IT infrastructure will focus on “systems of systems” that integrate and blend monitoring and treatment devices. Networks will stream data into medical records that are automatically mined to extract the knowledge used to drive a host of activities, such as auto-

mated treatment and dosing and long-term research into human health and the effectiveness of treatment.

For healthcare providers, operating rooms and diagnosis and treatment venues will shift from a collection of fixed monolithic devices to plug-and-play components that enable flexible and rapid reconfiguration of diagnostic, recording, and treatment systems. Advances in minimally invasive medical robotics and real-time high-speed networks will make telemedicine and robotic surgery technologies widely available. As generations of technology-savvy healthcare consumers enter retirement, they will embrace—and even demand—sophisticated home healthcare monitoring, treatment, and record systems integrated with national information databases (such as prescription-drug information systems) and local hospital and primary-care systems.

Although these envisioned innovations hold great promise, they will render current MDSS development and certification processes obsolete. End-user demands inevitably exceed the capability of existing MDSS. Unless researchers develop new certification technologies and unless development and certification processes undergo a paradigm shift, innovation will be stifled because manufacturers and regulators will find the development of HCMDSS too costly—or we will see dramatic increases in security breaches and harmful incidents due to device malfunction.

CHALLENGES

The cross-cutting nature of medical device design—transcending the informational, physical, and medical worlds—along with the possibility of a nationwide networked medical system that actively monitors and regulates the health of our nation’s citizens, raises immense scientific and technological R&D challenges for the IT, medical, and regulatory communities. The challenges envisioned for the next 10 years include the following:

- *System integration.* As we embrace a plug-and-play vision of medical device networks in future digital hospitals and digital homes, we must collectively facilitate the development of medical device systems and coordinate them with the development of standards for the architecture and communication of interoperable plug-and-play device networks. Achieving these goals while establishing quality-of-service levels that ensure system and patient safety on the one hand, and patient security and privacy on the other, is a great challenge.
- *Critical infrastructure.* As we move toward an environment in which all patients are constantly monitored and actively plugged into a nationwide medical

information network, we are creating a new critical infrastructure that will literally monitor the nation’s health. We need new methods to ensure the safety and security of that network, particularly methods involving the active use of information for medical purposes. In the presence of abnormal conditions or attacks, the system’s performance must degrade gracefully and safely, and the system must identify, contain, and, if possible, repair faults while providing timely notification to human operators.

- *Embedded real-time systems design.* Medical devices are embedded not only inside information networks but also inside human patients, whose critical life-functions they monitor and regulate. The design of medical devices is therefore more than an IT issue; it must also include the device’s interaction with the patient and the environment and the context in which they coexist. Thus, we need a fundamental rethinking of medical device design—toward a holistic approach that integrates functional, computational, and communication designs in the presence of highly uncertain patient models in both normal and abnormal conditions.
- *Validation and certification.* Current design practice makes certification and verification an afterthought, taking place 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 certification should be incorporated in early design stages. Furthermore, certification and design frameworks are currently not component-based, resulting in time-consuming and expensive certification of large integrated systems, inefficient certification of incremental or evolutionary designs, and difficulties in maintaining or upgrading legacy systems.

Addressing these challenges will dramatically affect the medical device and healthcare industry with significant system integration and development capabilities based on scientific principles and foundations.

RESEARCH DIRECTIONS

Despite the nationwide scale and the heterogeneous nature of the R&D challenges, the following research directions will help us make significant progress toward realizing the outlined challenges.

- *Infrastructure for medical device integration and interoperability.* The Electronic Health Records initiative needs to be safely and securely integrated with

The cross-cutting nature of medical device design raises immense scientific and technological R&D challenges for the IT, medical, and regulatory communities.

plug-and-play interoperable device networks so that we can fully realize the vision of actively using patient-specific information in optimum health delivery via interoperable medical devices. Interoperability presents a major challenge to integrating medical devices from different manufacturers. It will require the development of standards and architectures not only for medical records but also for devices that actively use that information to monitor and regulate patients' medical conditions. Besides unique patient (record) identifiers, which must support the integration of devices from different manufacturers, standards must address data and communication formats as well as the context and environment assumptions in which the information will be interpreted and used.

- **Model-based development.** The multifaceted nature of designing, implementing, and certifying medical devices requires holistic frameworks that are simultaneously model-based and component-based. Because of the strong coupling between device and patient, model-based frameworks that explicitly model a device's interaction with the environment and with the patient would lead to safer, higher-confidence devices and ultimately to better healthcare.
- **Component-based design frameworks.** Despite substantial progress in object-oriented frameworks for nonembedded software, model-based development is a challenging research direction. Component-based development will dramatically affect both the design and certification process: It will enable incremental yet certified compositions of certified components, allowing the safe and rapid reuse of legacy components—models, software, and algorithms. Our goal should be to develop frameworks in which certification is part of the design process rather than an afterthought. Component-based design should also support a variety of standards for communication and security.
- **Patient modeling and simulation.** Medical devices face a unique challenge in model-based design because of the scarcity of patient models and high-fidelity simulators for device design. As future devices adapt to patients, their medical conditions, and their environment, developing a variety of models and simulators for normal and abnormal patients in a variety of physical and environmental conditions will be increasingly important. We must develop models and simulators at various levels of detail, ranging from coarse models for device design to high-fidelity simulators for model validation and virtual validation and testing.

Formalizing requirements will enable precise and transparent translation of natural-language clinical requirements to quantified engineering requirements.

- **Adaptive patient-specific algorithms.** Whereas medical devices are typically designed for groups of patients who have similar medical conditions, we could dramatically improve healthcare by making devices whose operation would adapt to a specific patient's specific medical condition. To achieve that, we need to develop control, optimization, and machine-learning algorithms for medical devices that are certifiably safe for large classes of patients and that can adapt to individual patients or to different environments.
- **Requirements and metrics for certifiable assurance and safety.** Developing rigorous requirements for clinical and design purposes, as well as metrics for certifiable assurance, is an important research direction. In particular, formalizing requirements will enable precise and transparent translation of natural-language clinical requirements to quantified engineering requirements. It will also affect testing by developing frameworks for generating testing scenarios from clinical requirements.
- **User-centered design.** As medical devices permeate cross-sections of society and all educational and technical backgrounds, ergonomics and ease-of-use issues become important design factors in human-device interfaces. These issues should be considered throughout the design process. User and context modeling will result in better interaction between users and devices, minimize unsafe device operation, and result in graceful degradation of performance in the event of user or device failures.

Achieving this grand agenda is not simply a matter of time. It requires synergistic and concerted efforts among healthcare practitioners, medical device developers and manufacturers, and academics from the computer science, control theory, and bioengineering disciplines. Furthermore, it needs planning and support from government agencies.

To initiate this research agenda, medical device manufacturers must provide open experimental platforms to the academic community, and the regulating agencies must better educate everyone about the current medical device approval process. Understanding the approval process along with formalization of clinical and user-centered design requirements will be critical for subsequently developing quantifiable metrics for system assurance and certification. Standards for data, information, and communication will enable plug-

and-play, interoperable device networks operating on robust real-time infrastructures.

It is becoming increasingly evident that model-based frameworks that support component-based modeling, design, testing, and certification can have a dramatic impact. The one aspect of this agenda that might require a longer horizon, however, is the development of high-fidelity organ and patient models for design, testing, and validation.

This agenda has the potential to create a new scientific community and a new generation of scientists and engineers that integrate computer science, engineering, and medicine. ■

Acknowledgments

We thank all the participants in the HCMDSS workshop, held 2 and 3 June 2005 in Philadelphia. In particular, we thank Wei Zhao and Helen Gill from NSF; Paul Jones from the FDA; David Hislop from ARO; Brad Martin from NSA; LeRoy E. Jones from ONCHIT; Shankar Sastry from UC Berkeley; and Simon Szykman, Sally Howe, and Frankie King from NCO/NITRD. The workshop was supported in part by NSF CNS 0532968 and the School of Engineering & Applied Science, University of Pennsylvania.

Reference

1. Dept. Health and Human Services, "Health Information Technology Leadership Panel: Final Report," Mar. 2005; www.hhs.gov/healthit/HITFinalReport.pdf.

Insup Lee is the Cecilia Fitler Moore Professor in the Department of Computer and Information Science at the University of Pennsylvania. His research interests include embedded and real-time systems, medical device systems, and model-based development. Lee received a PhD in computer science from the University of Wisconsin, Madison. He is a Fellow of the IEEE and a member of the ACM. Contact him at lee@cis.upenn.edu.

George J. Pappas is an associate professor in the Department of Electrical and Systems Engineering at the University of Pennsylvania. His research interests include control systems, hybrid systems, and embedded systems. Pappas received a PhD in electrical engineering and computer sciences from the University of California, Berkeley. Contact him at pappasg@ee.upenn.edu.

Rance Cleaveland is a professor of computer science and the Executive and Scientific Director of the Fraunhofer USA Center for Experimental Software Engineering at the University of Maryland at College Park. His research interests include formal methods for software verification, mathematical modeling of software, and embedded systems. Cleaveland received a PhD in computer science from Cornell University. He is a member of the IEEE and the ACM. Contact him at rance@cs.umd.edu.

John Hatcliff is a professor in the Department of Computing and Information Sciences at Kansas State University, where he also leads the Specification, Analysis, and Transformation of Software (Santos) Laboratory. His research interests include model-driven development, program analysis and verification, and software engineering for safety- and mission-critical systems. Hatcliff received a PhD in computer science from Kansas State University. Contact him at hatcliff@cis.ksu.edu.

Bruce H. Krogh is a professor of electrical and computer engineering at Carnegie Mellon University. His research interests include design and verification of embedded control systems and hybrid dynamic systems. Krogh received a PhD in electrical engineering from the University of Illinois. He is a Fellow of the IEEE. Contact him at krogh@ece.cmu.edu.

Peter Lee is a professor of computer science at Carnegie Mellon University. His research interests focus on the design, analysis, and implementation of programming languages with applications in networking, security, and proof-carrying code. Lee received a PhD in computer and communication sciences from the University of Michigan, Ann Arbor. He is a fellow of the ACM and a member of the Board of Directors of the Computing Research Association. Contact him at peter.lee@cs.cmu.edu.

Harvey Rubin is a professor of medicine, microbiology, and computer Science at the University of Pennsylvania. His research interests focus on molecular mechanisms of bacterial pathogenesis and computational biology. He received a PhD in molecular biology from the University of Pennsylvania and an MD from Columbia University. Contact him at rubinh@mail.med.upenn.edu.

Lui Sha is the Donald B. Gillies Professor in the Department of Computer Science at the University of Illinois at Urbana-Champaign. His research interests focus on dependable real-time systems. Sha received a PhD in computer science from Carnegie Mellon University. He is a Fellow of the IEEE and the ACM. Contact him at lrs@cs.uiuc.edu.