Report on the High-Confidence Medical-Device Software and Systems (HCMDSS) Workshop

February 6, 2006

Insup Lee, University of Pennsylvania, Co-Chair
George Pappas, University of Pennsylvania, Co-Chair

(The document was submitted to NSF as a part of the final report for NSF CNS 0532968.)

EXECUTIVE SUMMARY

Background and Scope

The United States spends about 16% of its gross domestic product on health care—twice the average of most European nations (Health Information Leadership Panel, Final Report, Department of Health and Human Services, March 2005). The rapidly increasing use of software to control medical devices makes the development and production of medical-device software and systems a crucial issue, both for the U.S. economy and for ensuring safe advances in health care delivery. Several federal agencies are interested in identifying the research necessary to improve the design, certification, and operation of medical-device software and systems. The ultimate goal is better and more cost effective medical care.

On June 2 and 3, 2005, the High-Confidence Medical Device Software and Systems (HCMDSS) workshop was held in Philadelphia. An HCMDSS Workshop Planning Meeting (WPM) had been held on November 16 and 17, 2004, in Arlington, Virginia. The WPM was sponsored by the NITRD federal agencies that participate in the HCSS Coordination Group (CG), including FDA, NIST, NSA, and NSF, along with the National Coordination Office for NITRD. Sixty experts participated in the planning meeting, including those in government and industry information technology, software engineers, medical doctors, nurses, and academic researchers (see http://www.cis.upenn.edu/hcmdss-planning/).

The objective of the workshop was to build on the work accomplished at the planning meeting and to identify additional challenges and approaches from other constituencies. More than 90 experts from academia, medical sectors, industry, and government attended the workshop. They represented a complete mix of the relevant stakeholders—including researchers, developers, certifiers, and users—who can help identify emerging systems and assurance needs. This report, a tangible outcome of the workshop, prioritizes recommendations by using a roadmap to determine what, when, and how priorities should be addressed over identified time frames.

Purposes and Format of Workshop

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. The main goal was to develop a roadmap 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 (R&D) 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 HCMDSS workshop included plenary and panel discussions and breakout sessions. The panels and breakout sessions addressed the following six issues essential to HCMDSS:

1. **Distributed Sensing and Control in Networked Medical-Device Systems**. The networking of medical devices for distributed sensing and control can occur at many levels. Although ad hoc growth of network applications for medical devices has occurred,

today's commercial off-the-shelf technologies (COTS) do not produce highly distributed medical-device systems with guarantees of security, privacy, robustness, interoperability, extensibility, mobility, and general patient safety. Research is needed to create medical-device networks with those features and to enable the diffusion of new sensing and control technologies as they become available.

2. **Embedded Real-Time Networked System Infrastructure for Medical-Device Software and Systems (MDSS)**. The next generation of medical systems is envisioned to be a ubiquitous network of networked systems for secure, reliable, privacy-preserving, and cost-effective personalized high-quality health care. It will be a network that improves the quality of life. Although networks of networked medical devices hold many promises and possibilities, they also create challenges.

3. **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, better safety parameters, cost-effective development phases, and ultimately achieving regulatory approval. In the medical-practice domain, modeling and simulation will improve outcomes and quality of care and will provide better utilization of health care costs—with improvements in prevention, intervention, and maximal use of the electronic health record (EHR).

4. **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 methods of software development to ensure reliability and to protect the public health. Exactly how to accomplish that is a major 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 in order to ensure that devices will perform appropriately.

5. **Foundations for Integrating Medical-Device Systems and Models**. 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. Systemic inefficiencies in health care delivery grossly inflate costs and contribute to avoidable medical errors that degrade patient care.

6. **Verification, Validation, and Certification**. Verification and validation (V&V) 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 and more complex and rely much more on embedded software to achieve critical functionality, existing certification processes are being stressed. The results: higher development costs for manufacturers,

longer time to market, and increased chances of device failure—with associated recall or liability costs.

Each working group that participated in the workshop was asked to summarize the state of the art in practice, development, and research in its area, and to identify R&D needs and challenges, along with a roadmap to address the needs and challenges. This report, the full document of the HCMDSS workshop, includes the Executive Summary and six working-group summaries. The presentations of the working groups, keynote speakers, and panelists, along with the submitted position statements of participants, are available on the workshop Web site: http://www.cis.upenn.edu/hmcdss/.

Findings from the workshop are summarized below.

Current State of Affairs and R&D Needs

Advances in computing, networking, sensing, and medical-device technology are enabling the dramatic proliferation of diagnostic and therapeutic devices. Those devices range from advanced imaging machines to minimally invasive surgical techniques, from camera-pills to doctor-on-a-chip, from infusion pumps to implantable heart devices. Although advances in standalone diagnostic and treatment systems have been accelerating steadily, the lack of proper integration and interoperation of those systems produces systemic inefficiencies in health care 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 ensure safe advances in health care delivery. The crucial issue, however, is the cost-effective development and production of reliable and safe medical-device software and systems.

Here are several observations about the state of the art in medical-device software development.

- **Medical-Device Software Development**. Designing bug-free software is difficult, especially in complex devices that may be used in unanticipated contexts. Existing practices have worked as well as they have because industry V&V 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—slows the time to bring devices to market. Engineers often feel overwhelmed by complexity. Because of time-to-deliver pressure and a lack of properly trained software engineers, the development of HCMDSS has, with very 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 ad hoc, however, without standardized integration mechanisms that are commonplace in other domains, such as the highly successful and widely used universal serial bus (USB) 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.

- **Device Interference and Interoperation.** 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 (QoS) and semantic compatibility for interoperation, have not yet been addressed. Also, we need to conduct a systematic study of device interference during integration.

- **Approval and Certification**. FDA device approval centers on a process-driven approach, in which manufacturers obtain approval by showing that they have carried out the process of applying established quality assurance techniques to certain levels of coverage, such as manual code inspections and testing. As a whole, the medical industry does reasonably well in developing and approving standalone devices that have moderate complexity and are based on mature technology. But when considering larger devices with relatively complex functionality, the time and costs associated with V&V tasks such as test generation and execution cause researchers to lose confidence in their ability to bring safe and effective devices to market.

It is important to 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 health care, and although FDA approval processes have mostly kept pace, we now seem to be on the cusp of the types of revolutionary changes in health care 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 health care centers, local hospitals and clinics, the offices of primary-care physicians, home computing, and body-area networks. The health care IT infrastructure will focus on "systems of systems," based on architectures built around middleware, that integrate and blend monitoring and treatment devices. Networks will stream data into medical records that are automatically mined to extract knowledge that drives a host of activities, such as automated treatment and dosing and long-term research into human health and the effectiveness of treatment.

For health care providers, operating rooms and other venues of diagnosis and treatment 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. And as generations of technology-savvy health care consumers enter retirement, they will embrace—and even demand—sophisticated home health care monitoring, treatment, and record systems integrated with national information

databases (such as prescription-drug information systems) and local hospital and primary-care systems.

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

R&D 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. Here are some of the challenges we envision for the next ten years:

- **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 (PnP) device networks. Achieving that while achieving quality-of-service levels that ensure system and patient safety on the one hand and patient security and privacy on the other hand is a great challenge.

- **Critical Infrastructure**. As we head toward an environment where 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 performance of the system must degrade gracefully and safely, and the system must identify, contain, and, if possible, repair faults while providing timely notification to human operators.

- **Design of Embedded Real-Time Systems**. 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 devices' interaction with patients and the environments and contexts 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 uncertain patient models, in both normal and abnormal conditions.

- **Validation and Certification**. Current design practice makes certification and validation an afterthought, 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 based on full systems, not components, 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.

New Research Directions

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

- **Infrastructure for Medical-Device Integration and Interoperation**. The Electronic Health Records initiative needs to be safely and securely integrated with plug-and-play interoperable device networks. We could then fully realize the vision of actively using patient-specific information for optimum health delivery via interoperable medical devices. Interoperability has presented 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 contexts 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 based on models and components. Because of the strong coupling between device and patient, model-based frameworks that explicitly model devices' interaction and limitations with the environment and with the patient would lead to safer, higher-confidence devices and, ultimately, better health care.

- **Component-Based Design Frameworks**. Component-based frameworks have been developed to facilitate the reuse of large and complex systems through the reuse of individually deployable and reusable components. Despite substantial progress in developing component-based frameworks for non-embedded software, such frameworks have rarely been applied to medical-device software and systems. An integrated approach that combines component-based and model-based development is an important research direction that can meet development and maintenance challenges of medical-device software and systems. Component-based development for both design and certification will dramatically affect 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 the environments they live in, it will be important to develop a variety of models and simulators for normal and abnormal patients in a variety of physical and environmental conditions. 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.

- **Adaptive Patient-Specific Algorithms**. Whereas medical devices are typically designed for groups of patients who have similar medical conditions, we could dramatically improve health care by making devices whose operation would adapt to a specific patient's specific medical condition. To achieve that, we need to develop algorithms for medical devices that are certifiably safe for large classes of patients and that can adapt to individual patients or to different environments that patients may be living in.

- **Requirement and Metrics for Certifiable Assurance and Safety**. The development of rigorous requirements for clinical and design purposes, as well as metrics for certifiable assurance, are important research directions. Ideally, it should be possible to extract or convert natural-language clinical requirements to quantified engineering requirements. Such requirements make it possible to develop testing, validation, and analysis techniques with quantifiable guarantees for MDSS.

- **User-Centered Design**. As medical devices permeate cross-sections of society and all educational and technical backgrounds, ergonomics and ease-of-use issues in human-device interfaces should become important factors in design. User-centered design, ergonomics, and ease-of-use issues in human-device interfaces 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.

Research Roadmap

Achieving this grand agenda is not simply a matter of time. It needs planning and support from government agencies. Here we offer a sequence for the most important components of a research roadmap that addresses the research challenges. Finally, this agenda has the potential to create a new scientific community and a new generation of scientists and engineers who integrate computer science, control theory, biomedical engineering, and medicine.

The proposed roadmap, which will dramatically affect medicine and health care, consists of three distinct phases.

Three-Year Roadmap

To develop a coherent body of methods and technologies that can meet the challenges of future MDSS, research needs are not just in information technology but in something that is much more

multidisciplinary and involves significant computer science along with biomedical engineering, device manufacturing, and the medical-care process. In three years, we would like to see—

- An initial set of publicly available open experimental platforms that contain design artifacts, including reference models and usage scenarios of different medical devices, so that researchers can obtain empirical feedback on their ideas about real-world systems
- The development of standards for data, information, and communication to enable plug-and-play medical devices and to support interoperable device networks
- An understanding of the approval process, along with the formalization of clinical and user-centered design requirements, and the development of quantifiable metrics for system assurance and certification
- The formation of an R&D community for medical-device design to raise awareness across different subdisciplines and to foster collaborations between researchers, industry, health care providers, and government agencies

Five-Year Roadmap

For the slightly longer term, we would like to see short-term technologies enter clinical trials. Examples are standards-based compliance specification, verification, and validation technologies and processes for interoperability and QoS. In addition, we should focus on the theoretical and engineering foundations for system engineering aspects of medical infrastructures. In five years, we would like to see—

- The integration of technologies for medical devices that have different QoS requirements into a network-centric system of systems, including management systems for medical information and networked devices in an operating room
- A fundamental understanding of how to carry out the medical-practice-driven design of components and protocols so as to improve the safety of medical devices
- A demonstration of the practicality of model-based frameworks for MDSS development and integration
- The development of certification methods for individual components and networks of devices, and the development of an evidence-based, technology-aware certification process

Ten-Year Roadmap

Further out, we would like to see FDA-approved networked medical devices in wide use, with a medical infrastructure that supports the composition and integration of medical-device components while guaranteeing QoS, security and privacy, validation, and certification. In particular, the research and development of MDSS should culminate in—

- The deployment of a fully integrated hospital intensive-care system using distributed monitoring, distributed control, and real-time wireless networks

- The development of FDA-approved certification methods for medical-device software and systems and networked in-home patient monitoring and assistance
- The availability of high-fidelity organ and patient models for design, testing, and validation and model-based frameworks that support component-based modeling, design, testing, and certification using patient models

DISTRIBUTED SENSING AND CONTROL IN NETWORKED MEDICAL-DEVICE SYSTEMS

Participants: Bruce H, Krogh (chair), Tarek F. Abdelzaher, Timothy Buchman, Rich Craft, Mike Eklund, Sandeep K. S. Gupta, Nagarajan Kandasamy, T. John Koo, Ronald Marchessault, Tom Martin, Douglas Miller, Klara Nahrstedt, Tariq Samad, - Wei Zhao, George Fainekos (observer), Sebastian Fischmeister (observer)

Introduction

The networking of medical devices for distributed sensing and control occurs at many levels, ranging from dedicated networks of devices for individual patients to wireless networks for monitoring residents in long-term-care facilities. These networks may collect data for off-line analysis, generate alarms when critical conditions occur, or close feedback loops for the controlled delivery of drugs. Networks are increasingly being used to distribute stored medical information for remote diagnosis, such as picture archive and communication systems (PACS) that distribute x-rays and other images. In the future, wireless networks will supply distributed data acquisition and control for patients who are free to live relatively normal lives in their communities. Networking is the essential enabling technology for future advances in monitoring, diagnosis, treatment, and effective responses to acute conditions.

The growth of network applications for medical devices using current commercial off-the-shelf technologies (COTS) has been rapid. Development has been ad hoc, however, focusing on particular applications—with little standardization. Current COTS networks do not provide a sound basis for highly distributed medical-device systems that offer guarantees of security, privacy, robustness, interoperability, extensibility, mobility, and general patient safety. Research is needed to create medical-device networks that have such features, in order to improve health care in multiple dimensions and to enable the rapid dissemination of new sensing and control technologies as they become available. The following sections elaborate on these research needs and challenges.

What Can We Do Well?

Currently, data can be collected and sent to a variety of devices, such as—

- Implantable devices, such as pacemakers, defibrillators, nerve sensors, and neurostimulators
- Devices for patient monitoring, such as electrocardiogram, pulse oximetry, transcutaneous oxygen, electroencephalogram, pulmonary testing, and spirometry
- Devices for point-of-care testing and monitoring in assisted-living environments

The use of existing COTS network technologies to collect data from such medical devices is growing. Networks are used to log data and provide real-time alarms to attending health care personnel for patients in intensive-care units. EKG monitors for multiple patients are networked to deliver information to a shared database. In assisted-living facilities, monitoring devices are networked in several ways, including wireless networks that monitor ambulatory patients.

Networking is also facilitating in-home health care. Many people have wireless devices that they can use to request immediate assistance in their homes. In-home monitoring systems regularly and automatically deliver data on blood pressure, heart rate, and glucose levels over phone lines to databases that physicians can access. Home monitoring systems are also available to send alarms for congestive heart failure to emergency response professionals.

In hospitals, raw sensor data is often transmitted over networks, using standards for laboratory information systems, with basic analysis to flag obvious outliers. Networks also play a significant role in emerging systems for therapy and surgery. For example, networks connect sensors, processors, and controllers in radiation therapy systems to perform real-time tumor tracking. Similarly, networks are being used as an alternative to dedicated point-to-point connections of devices in systems for telerobotic surgery.

Why Can't We Declare Victory?

Despite the successes of networking medical devices, they face significant barriers to realizing their full potential. COTS networking technology is typically optimized for bandwidth rather than time delay, making it difficult to optimize the performance of critical real-time feedback loops. For example, one system for the real-time tumor tracking of treatment beam control has to be operated in a degraded mode because its several networked components for signal processing and control provide no timing guarantees.

Also, current systems are fragmented and often incompatible. Although individual devices continue to become more sophisticated, they cannot "talk" to each other. For example, acute-care medicine (including ER, OR, ICU, radiology suites, and anesthesia), which of all medical areas is the best served by networked monitoring, is only now starting to see some data continuity. The lack of standards prevents the integration of information needed to provide the physiological basis for interpreting and acting on the increase in sensing modalities.

Another shortcoming of current networked systems is a lack of significant processing of information. It has been noted that medical systems are now awash in data but are information poor. Monitoring systems deliver data, but human experts need to interpret the data. When interpretation is attempted, it needs to be carefully evaluated. For example, in ICU monitoring systems, false alarms are far too frequent to rely on without extensive human intervention. The ability to glean useful information based on the fusion of data from multiple sources on a network is virtually nonexistent.

The delivery of the correct information will be particularly important in telemedicine, where a network will provide the only direct, real-time information available to the human. Sensor fusion is also needed, to provide correct information to the computer that closes feedback loops. In robotic surgery, for example, the human surgeon must receive haptic feedback, and correct information must be provided to the robotic surgical assistant, which needs precise knowledge of the location and type of tissue in the workspace.

Medical systems must also become more intelligent in the ways they interact with human users. Current devices and systems merely respond to operator inputs in a deterministic way. Beyond

simple limit checking, they have little internal monitoring or analysis of the commands they receive. Given the diversity of training and skills of the individuals who will interact with future medical systems, it is important to consider all aspects of the human factors. As medical systems become more sophisticated and automated, they will need the ability to identify the level of knowledge and skill of the operator—and to respond appropriately. For example, when a system receives inappropriate sequences of commands that could be detrimental to the patient, the system should respond by maintaining a safe operating condition and supplying clarifying information to address the possibility of confusion by the user. Self-monitoring networked devices could adapt to changes in the user over time. As a patient ages, for example, a device might become easier to use by providing more detailed and thorough instructions.

Specific R&D Challenges

For our discussion, we will divide R&D challenges into two categories: component level and system level.

Component Level

The challenges of interoperability and compatibility are perhaps the most important to address immediately so as to facilitate the development of networked medical devices. We need to create standards, based on sound computational principals, with clear semantics. We need to develop middleware, in order to provide standard interfaces between medical devices and networks. That will make it possible to integrate components at the time of use rather than at the time of design.

Advances are needed in embedded-systems technology so that quality-of-service (QoS) guarantees at the component level can guarantee performance at the system level. The relative importance of security, privacy, robustness, interoperability, extensibility, mobility, and general patient safety must be evaluated carefully. Those issues need to be addressed by heterogeneous teams of computer scientists and health care professionals. The undertaking should begin with the establishment of a clear, mutually understood, shared vocabulary.

Devices must operate safely under all possible fault conditions and must provide appropriate information under all conditions. For example, when devices become disconnected from a network, caregivers often conclude that the devices are faulty. Devices and the network infrastructure must be designed so that in such situations, correct default values are delivered to the system. And the individual devices need to have mechanisms that indicate their states when they are not connected to the network.

Model-based design is a significant trend that facilitates the complete development process for embedded systems—from the capturing of requirements to implementation and deployment. The model-based approach builds partly on precise, formal descriptions that define the embedded system in relation to the domain in which it will apply. The medical domain requires a new formal language to support dialog among clinicians, physiologists, and device designers for the development of medical devices that are founded on a common, integrated framework. This formalization of the domain will also lead to new methods of component-level certification for safety-critical software.

A new model-based approach to the design of medical devices will also make it possible to address the many critical constraints in medicine, including form factor, power limitations, and usability. From an information technology perspective, formalizing the requirements and features of networked medical devices will make it possible to leverage emerging technologies much more quickly than we can today. Such technologies include smart sensors, digital video processing, and new methods of power management, such as harvesting energy from the environment.

Systems Level

Connecting components into workable, effective medical systems requires advances in system integration technologies. Distributed sensing and control need standards and architectures that make it possible to create networks of diverse medical devices and medical information systems. The support of real-time critical services must be based on models of end-to-end service that account for the details of sophisticated interconnection technologies.

System-level power management will be a necessary part of the monitoring and control of medical devices for patients beyond the walls of hospitals and health care institutions. Accomplishing the goal will require a full systems engineering approach—one that identifies the parts of the system that must be reengineered to provide the needed QoS for the complex systems of systems that arise in distributed medical monitoring and control.

Research is also needed in medical information and algorithms. Effective systems for model-based monitoring and control must include algorithms for learning and adaptation that adjust parameters appropriately for each patient. Data manipulation schemes must be developed to support the integration of electronic health records (EHR) with information from medical monitoring systems and clinical decision-support systems. Given the rate of innovation in medicine, it is also essential that information management systems are designed to be extensible, so that new types of information can be integrated easily and immediately without requiring any amount of redesign.

System-level research must regard information security and privacy as issues of primary concern. Incorporating support for privacy techniques—such as selective disclosure, auditing the auditors, and encrypted searches—are essential in networked medical systems, because of the safety-critical and personal nature of information. For security and privacy, system-level information authentication schemes need to be developed.

Effective medical monitoring and control can be done at the system level only by drawing on different sources of information to make correct decisions. Besides developing new algorithms to perform sensor fusion for medical applications, research is needed to determine how the diverse sources of information can be used to validate data and even to calibrate distributed equipment. For long-term monitoring, device maintenance can be scheduled by the integrated information system.

Model-based development can help address one of the major human-factor issues: training. Models used for component-level design and system integration could become the basis of

sophisticated simulation-based training systems. The use of simulators for training, which is standard in other areas involving complex technologies—such as aircraft and power systems—could reduce the cost of training and could raise the comfort level and acceptance of new networked technologies in medicine.

Human factors also play an important role in the potential for networked monitoring and control in medicine. Systems must be ergonomic, having human-machine interfaces that encourage acceptance by patients and caregivers and that consider the skill sets of the end users. Local control loops and control over networks need to take into account the full implications of possible human intervention. It is also important that systems be designed so as to enhance the capability of medical professionals to provide high-quality care—and not convey the message that the networked information system is attempting to remove the human from the loop.

New methods are also needed to enhance the safety of medical systems. When abnormal conditions occur, a system should provide timely notification to operators, and performance should degrade gracefully rather than abruptly, especially in life-sustaining medical equipment. This capability requires modeling and implementation technologies that naturally support self-monitoring and prioritized functionality, along with real-time response.

To address the several issues described above, new metrics must be developed to measure the capabilities of systems in multiple dimensions and to evaluate the trade-offs between them. These metrics should deal with the distinctive features of the medical domain, including the diversity of the patient population, the complexities of the conditions being monitored and controlled, and variation in the knowledge and skills of the end-user community. These metrics could support run-time operation strategies as well as design-time evaluation.

Research Strategies and Roadmap

From the discussion above, the following four themes emerge as specifically IT-related research needs:

Standards and Architectures. True interoperability and extensibility for networked medical systems will be possible only when widely accepted standards are embraced. Moreover, an effective infrastructure to support such systems should be built on open-source middleware for distributed systems designed specifically for medical applications.

Formal Models. English-language descriptions of requirements and specifications are too ambiguous for developing networked medical systems. We need a language with formal precision that will eliminate any possibilities for confusion that could arise when clinicians, physiologists, engineers, computer scientists, and others work together to develop safety-critical systems. These formal models can become the basis for the verification and validation of networked medical systems.

Model-Based Design Tools. At both the component and the system levels, effective designs will need model-based tools. Models for the distributed monitoring and control of medical

applications must be developed with representations that support the needs of domain experts as well as methods for analysis and implementation.

Test Beds. We need prototype problems that can serve as test beds for pre-competitive research. Problems from the device level to the system level—including critical features of users (clinicians) and patients—must be defined clearly, and examples must be disseminated widely, so that IT specialists can become fully engaged in developing the fundamental technologies.

To create a clear vision for research, we propose a roadmap, with milestones that stretch well beyond the capabilities of current systems. Those milestones, defined in terms of the capabilities of prototype systems, provide a basis for discussion among the many stakeholders in the medical and IT communities.

Three-Year Milestone: *A deployed networked sensing and control system on a mobile distributed population with more than 1,000 nodes.* In this system, ambulatory patients will be instrumented with continuous EKG tracers and pulse oximeters, and local closed-loop oxygen bottle control will be implemented. Features to be demonstrated include—

- End-to-end QoS design
- Remote processing of data
- Integration with the EHR system
- Multivendor support with autoconfiguration (“plug and play”)
- Smart sensors and smart alarms (few false alarms)

Five-Year Milestone: *Configurable communication, sensing, and control for emergency-room (ER) and geriatric care*. This milestone will have higher bandwidth and more sensing modalities than the three-year milestone. It will introduce additional sensing and control for glucose and insulin. Model-based development tools will provide simulation-based tools for training medical personnel.

This system will feature—

- Point-to-point communication
- Patient localization
- Service discovery and negotiation within the network
- Secure network reprogramming
- Service virtualization
- Enhanced distributed control of local feedback loops

Seven-Year Milestone: *Integrated portable preoperative and postoperative monitoring for civilian and combat scenarios.* This wide-area distributed system will include advanced capabilities such as—

- Ambulatory ultrasound monitoring
- Blood vessel graft monitoring
- Remote radiologic evaluation

The system will provide the possibility of real-time dosage adjustment based on remote monitoring and diagnosis. The network will operate in an ad hoc mobile environment with a very high bandwidth and guaranteed QoS for time-critical features. Feedback loops will be closed both locally and remotely, using patient-specific models. Treatment planning will be semiautonomous, and information fusion will provide robust smart alarms.

Ten-Year Milestone: The technology roadmap presented here culminates in the following two demonstrations comprising the ten-year milestone.

- A fully integrated hospital acute-care system using distributed monitoring and control
- A demonstration of the impact of long-term deployment of networked in-home patient monitoring and dosage control

These two systems will incorporate all of the advances advocated for distributed sensing and control of networked medical devices, including the application of extensible prototype standards and architectures for future development. The systems will also demonstrate the application of formal models and model-based design tools.

Concluding Remarks

Medical diagnosis and treatment draw on multiple sources of data, ranging from patient records to data supplied by real-time monitors. Current networking technologies have the bandwidth to acquire and deliver the data, but there are many barriers to realizing distributed sensing and control in networked medical-device systems. The barriers include a lack of adequate standards, a lack of effective algorithms to generate important information from the many crucial sources of data, and the unavailability of test data and nonproprietary experimental systems for research and development. The proposed roadmap provides a set of specific aggressive milestones that will drive the basic research and development to realize truly intelligent distributed medical systems in the next ten years.

MEDICAL INFRASTRUCTURE

Participants: Lui Sha (chair), Ashok Agrawala, Chris Gill, Julian Goldman, Jennifer Hou, David R. Jones, Soon Ju Kang, Raj Rajkumar, Majid Sarrafzadeh, Sang Son, Jack Stankovic, Simon Szykman, Russell Taylor, Taieb Znati, Madhukar Anand (recorder)

Introduction

Networks of networked embedded systems create many new possibilities and challenges. The next generation of medical systems is envisioned to be a ubiquitous network of networked systems for secure, reliable, privacy-preserving, and cost-effective personalized high-quality health care. It will be a network that improves the quality of life.

The current system consists of many standalone devices. Even if they are networked, most devices function on their own. For example, a nurse is alerted when a patient's electrocardiogram (EKG) becomes seriously abnormal. Hurrying to the patient's room, however, the nurse might see a contagious-disease warning sign at the entrance and realize that she or he should have put on protective clothing before entering. Memory lapses are common under stress.

In the future, all relevant information—such as the contagious-disease warning and special equipment needs—will be displayed along with the alerts. Thus the EKG machine will no longer be a standalone device and will become an integral part of a network: a network of medical devices and patient management. Patients' medical records and their current state, no matter where they are generated or collected, will be integrated, filtered, and delivered in real time to where they are needed.

During a surgical operation, context information—such as sensitivities to certain drugs—will be automatically routed to relevant devices, such as infusion pumps, to support personalized care and safety management. How the patient's vital signs are reacting to medications and surgical procedures will be correlated to streams of imagery data, selected and displayed in real time and tailored to the needs of medical personnel, such as surgeons, nurses, anesthetists, and anesthesiologists. For some particularly difficult stage in an unusual operation, an expert surgeon could remotely carry out the key steps, using remote displays and a robot-assisted surgical machine, avoiding the need to fly across the country to perform, say, 15 minutes of work. Furthermore, data recording will be integrated with storage management so that surgeons can review operations and key findings for longitudinal studies of the efficacy of drugs and operational procedures.

Networks of networked medical devices hold many promises. Besides enhanced operational capabilities derived from integrated devices and medical-information systems, they allow flexible configuration and deployment, and the collection of more accurate and representative data from natural settings for longitudinal studies to support improved health management. They also raise many challenges.

From the perspective of infrastructure technologies (other than networking technologies for connectivity), much remains to be done.

What Can We Do Well?

Over the past decade, and especially since the development of the World Wide Web, many technologies have been developed to support distributed computing systems. The emergence of middleware enables and simplifies the integration of components developed by multiple technologies: It provides a consistent set of higher-level network-oriented abstractions that are much closer to application requirements than what is done now. That simplifies the development of distributed and embedded systems. Middleware also provides a wide array of common services—such as name services, logging, and security—that have proved necessary to operate effectively in a networked environment. Successful commercial middleware includes J2EE, CORBA, and .NET. They are the foundation on which many Web-based applications are built.

This raises an interesting question: why don't we use these commercially available technologies to network systems?

Why Can't We Declare Victory?

Current medical devices are mostly standalone subsystems that have proprietary designs. Medical workers often must manually enter data and transfer it between machines. Sometimes, they need to set devices manually, so as to emulate the interlock needed between different devices and actions. So if a patient has medication Y, the limits of delivery X in an infusion pump need to be adjusted to Z. Medical workers also need to mentally correlate many paper records and screen displays from various diagnostic and monitoring machines. The process is time consuming, burdensome, and error prone.

The commercial middleware cited above does help, but it is limited to information management for medical records. Many record formats for patient identification depend on particular designs that their manufacturers provide. They cannot be extracted from their native environments and put into a universal patient identification database. Once we develop a common patient ID standard and migration process, we can readily apply commercial distributed-computing technologies.

Beyond record management, the primary two challenges in the development of integrated medical systems are safety and liability. They are intertwined. Many medication devices are safety critical. The FDA approves each device separately, in a specific application context. If we connect them and something goes wrong, who is responsible? Technically, this is a safety question. In the new interactive environment, how is safety specified? Is it specified correctly? Is it implemented correctly and does it perform as specified? Is there a proper training process for users? Do medical personnel use it correctly?

What makes these questions impossible to answer is that current commercial infrastructure software assumes absolutely no liability and has many known and unknown bugs. The development of a certifiably safe infrastructure for networked systems of medical systems is a long-term R&D challenge that involves not only advanced technologies but also a legally sound certification process.

Before we can have a certifiably safe medical network infrastructure, safety-related actions should be limited to computer-aided actions that are closely supervised by medical personnel as an intermediate step in its evolution. Even with this modest goal, we still have a way to go, because the current distributed infrastructure is a best-effort system whose real-time reliability and security cannot be ensured.

In the following sections, we will take a more detailed look at each of the challenges.

Specific R&D Challenges

Designing for Certification

Many medical devices are safety critical and must be certified. At present, the FDA approves medical devices. Certification is desirable but needs R&D to make it possible for medical-device software and systems. Thus, it is important to develop a standards-based infrastructure of certifiable networked medical devices so that we can reduce the costs of development, approval, and the deployment of new technologies and devices.

Certification cannot be an afterthought. We must develop technologies for the specification and the design of verifiable and certifiable medical devices. Certification includes devices' operational environments. In future integrated medical-device and medical-information systems, the application contexts might be quite dynamic. The development of technologies that can formally specify both application contexts and device behaviors is a major challenge for the vision of certifiable plug-and-play medical devices.

It will be vital that researchers work with medical and regulatory agencies to ensure compliance with safety requirements. In addition, they will need to work with the Institute of Electrical and Electronics Engineers (IEEE), the International Organization for Standardization (ISO), and other standard-setting entities to develop voluntary market-driven certification for (1) interoperability and compatibility standards and (2) quality-of-service (QoS) standards—such as real-time fault tolerance, privacy, and security.

Two challenges in this area are—

- How to develop evidence-based safety certification technologies and processes, including technologies for specification and verification and validation
- How to develop open-standard-based compliance specification, verification and validation technologies, and processes for interoperability and QoS properties

Quality of Service

End-to-end QoS is an important concern in the operation of medical devices. This section discusses key QoS attributes.

Managing Safety and Criticality

From operating room to enterprise system, different devices and subnetworks have different levels of clinical criticality. Data streams with different time sensitivities and criticality levels may share many resources of the hardware and software infrastructure. How to maintain safety in an integrated system is a major challenge that consists of many research issues:

- How to develop a safety interlock for the operation of interacting devices
- How to manage the flows of data streams that have different criticality on the same network
- How to mediate and manage the interactions of devices that have different criticality. How to authorize and authenticate. Who can talk to whom?
- How to support the fail-safe operation of individual devices
- How to ensure that erroneous data is contained and does not cascade
- How to specify, design, and verify the safety and efficacy of networked medical systems in the presence of device hardware failures and software errors

Security and Privacy

Medical systems pose new challenges in security and privacy. For example, a patient may need to upgrade software that detects pacemaker arrhythmias using EMF (electric and magnetic fields)–based methods but that is impervious to EMF-based attacks. Emergency personnel may need to access private data on demand anywhere and anytime, especially in wireless environments. We need a better understanding of requirements for security and privacy in medical systems. These are some key aspects of that understanding:

- How to develop a new model of security and privacy tailored to medical needs
- How to develop a modular and flexible architecture that incorporates and evolves technologies for security and privacy
- How to develop cost-effective solutions appropriate to the medical environment

Interoperability

Interoperability has been a major challenge in integrating medical devices from different manufacturers. When a device uses data supplied by other devices, what is involved is not only the format of the data but also the context in which to interpret the data. For example, is the blood pressure measured when the patient wakes up in the morning but is still in bed—or after a workout, or after taking medication? Here are important interoperability challenges:

- How to specify the interface both for the format of medical data and for context information needed to interpret the data

- How to specify, design, and verify the properties and compliance of interoperable hardware and software interfaces that go beyond the data format
- How to develop a standard of unique patient (record) identifiers to support the integration of devices from different manufacturers
- How to represent environmental assumptions implicitly embedded in code and make them machine checkable and user friendly
- How to support the evolution of technologies and maintain interoperability between old and new devices

Real-Time and Scheduling Guarantees

Many medical devices operate in real time with different time constraints and different sensitivities to delays and jitters. In the envisioned network of networked medical devices, many types of real-time and non-real-time data traffic will share the same computing and communication resources. How to ensure the proper scheduling of real-time traffic is an important concern, and here are some of the challenges:

- What should be the policies of resource allocation and scheduling that ensure predictable end-to-end timing constraints and interoperability
- How to provide time-zone abstractions that can support monitoring and control loops that have differing time constraints, such as 5 milliseconds, 10 milliseconds, and 100 milliseconds
- How to support consistent views and actions between distributed and collaborating medical devices within given timing constraints
- When network equipment fails and a system is overloaded, how to ensure that deadlines for critical real-time data streams are met

Medical Information Management

Integrating the operation of medical devices with nationwide medical-information management will give us a better understanding of the effectiveness of medical procedures and provide the right context for treating patients. We need to build in support for integrated medical-information management:

- How to perform the recording, correlation, and analysis of event sequences
- How to provide real-time context awareness for the proper operation and management of medical devices and information
- How to conduct real-time, context-aware alarm processing, filtering, and delivery (at present, false alarm rates are far too high)

- How to integrate with enterprise systems to support record management and long-term studies
- How to manage high-volume data intelligently. Aspects include—
  - User-friendly real-time data collection, filtering, fusion, and delivery
  - Support for storage networks and data mining
  - Capability for TiVO replay during surgery
  - Visualization of massive data sets

Wireless Medical Infrastructure

Wireless networking is an important enabling technology. To provide secure and reliable real-time communication, however, we face many challenges, including—

- How to exploit ultra-wideband (UWB) technologies
- How to improve interoperability and protect against interference
- How to improve security, reliability, and schedulability
- How to support mobility, including programming abstractions that manage mobility
- How to integrate with the wired infrastructure

Research Strategies and Roadmap

In developing a high-assurance medical infrastructure, the priority is to create a system-engineering framework that integrates component technologies with certification technologies. It will be a grave mistake to develop different technologies in isolation. Although it is not overly difficult to develop protocols for safety, security, and real-time reliability and privacy in isolation for particular application contexts, those protocols sometimes interfere with one another unexpectedly when they are used together in different contexts.

Three-Year Roadmap

To create a coherent body of technologies that are certifiably safe, we must first gain a deep understanding of the context of medical applications. It is difficult to develop effective technologies without knowing the constraints. Three years from now, we would like to see—

- An initial set of publicly available model application problems that allow researchers to understand user requirements and to test their ideas
- An initial suite of coherent QoS protocols designed and rigorously specified and verified for correctness and compatibility

- An initial set of experimental prototypes demonstrating the feasibility of new QoS and wireless technologies

Five-Year Roadmap

We would like to see short-term technologies enter clinical trials. Examples are standards-based compliance specification, verification, and validation technologies and processes for interoperability and QoS. In addition, we should focus on the theoretical and engineering foundations for system engineering aspects of medical infrastructures. In five years, we would like to see—

- Network timing abstraction methods that support monitoring and control loops with multiple timing constraints, ranging from 1 millisecond to seconds to minutes.
- Integration technologies for medical devices that have different QoS requirements into a network-centric system of systems, including management systems for medical information and networked devices in an operating room
- Fundamental understand on how to design protocols that make certification easier and how to develop an evidence-based, technology-aware certification process

Ten-Year Roadmap

In 10 years, we would like to see FDA-approved networked medical devices in wide use, with medical infrastructures supporting the composition and integration of medical-device components while guaranteeing QoS, security and privacy, validation and certification.

- How to model and reason about the interactions between protocols for safety, interoperability, real time, reliability, security, and privacy
- How to design protocols that are compatible rather than interfere with one another
- An evidence-based and technology-aware certification process defined for approval of medical device software and systems by regulating agencies

PATIENT MODELING AND SIMULATION

Participants: Harvey Rubin (chair), Ruzena Bajcsy, Scott L. Bartow, Amit Bose, M. Cenk Cavusoglu, Robert C. Kircher, Douglas Rosendale, Charles Taylor, Russ Taylor, David Arney (recorder)

Introduction

Convincing successes in other fields—such as aerospace, chemical plants, and the automotive industry—confirm the value of modeling. Among the benefits we see are novel product developments, increased safety parameters, cost effectiveness in development phases, and ultimately a higher rate of regulatory approval.

Patient models exist and evolve over five levels of the spatial scale. At each scale, the models involve heterogeneous structures and physical processes, and each model evolves over time. On the atomic and molecular level—at the angstrom scale—the models are manifest in biochemical and genetic systems. At the next level, the cellular level, models are actively being generated and investigated. These two levels are integrated into models of organ structure and function. Organ models are then integrated into models of whole-body function. At the highest level, societal interactions are modeled, where the models must accommodate the uniqueness of each patient and also must permit the aggregation of populations.

In the medical-practice domain, we anticipate that high-level modeling will result in improved health care, with better outcomes and a higher quality of care. We also anticipate that modeling will provide better use of health care dollars, with improvements in prevention, intervention, and maximal value and utilization of the electronic health record (EHR).

In product development, an argument can be made that patient modeling is a highly efficient tool in developing devices. The reasons are multifold. Among them: human studies are expensive; and device manufacturers need models for sophisticated protocols and the effective execution of procedures, planning, monitoring, and control. Currently, however, the barriers to developing specific models are high.

In training and professional certification, patient modeling provides outstanding opportunities for patient education and guidance in clinical decision-making.

In research, patient modeling offers a rich investigative platform on which to develop a wide variety of tools and techniques.

What Can We Do Well?

Functions and operations that patient modeling communities do well are found at the ends of the modeling spectrum. For example, although many parameters in enzymatic reactions have yet to be determined, the mathematical modeling of enzyme reaction kinetics is well established. Even reactions that follow stochastic rules are tractable with high-performance computing.

Another advanced example of patient modeling is found in medical imaging. When used properly, imaging has clearly been shown to be clinically cost effective and to have positive measures of outcome. Examples include seizure focus ablation, arrhythmia focus ablation, image-guided biopsies, and radiation therapy mapping. Imaging is also an important component at the procedural level, providing training environments and programs that are otherwise unavailable, costly, or dangerous. Some systems are already in use, indicating the potential for commercial success.

Remarkable advances in patient modeling have been made over the past five years at the subcellular level. For example, transcriptional analyses of a wide variety of diseases have been enabled by the rapid accumulation of microarray data coupled with the explosion of genome sequencing data. More complete mapping of the metabolome in various normal and disease states is accumulating as well. These data have been complemented by new techniques in data mining, analysis, and integration. Finally, modeling of the epidemiology of disease is a mature discipline, with well-developed tools and algorithms.

Convincing preliminary data shows that physiology-based modeling is effective. The critical-care domain and the intraoperative domain are two areas in which physiologic modeling is widely used. These areas are driving the development of even more sophisticated modeling software and devices. It is anticipated that the home-care and institutional-care markets will be the next to rely heavily on sophisticated patient modeling software and devices.

Patient models are the focus of a number of national and international funding agencies. Large government-funded programs have begun to support the development of tools necessary for patient modeling, and some already exist. For example, the Physiome project and the DARPA BioComp program support development at the biochemical, genetic, and cellular levels. The ITK open-source NIH-funded image processing toolkit (http://www.itk.org/) focuses on tools for modeling at the organ and whole-body levels.

Other government programs in this realm are NASA's "digital astronaut" project, which is in the planning stage, and DARPA's Virtual Soldier project (http://www.darpa.mil/dso/thrust/biosci/virtualsoldier.htm).

Why Can't We Declare Victory?

Patient models involve heterogeneous structures—atoms, molecules, cells, organs individuals, societies—**and physical processes that evolve in time and space**, posing difficult computational problems. We will call this the *multiscale/multistructure* problem.

Patient models must be accessible to a wide variety of communities. Fully integrated models must be available to large and heterogeneous communities, including practitioners, investigators, device developers, and regulators. Those communities are not tightly integrated; they do not generally participate in the same meetings; they do not share journals; and they are spread across

industry, government, and even different departments in academic centers. We will call this the *communication* problem.

Mechanisms need to be developed to share data, models, tools, and results. The interoperability of models and the maintenance of privacy are two of the most challenging problems facing the field. In addition, both commercial and academic institutional barriers limit the sharing of data and tools. In the academic domain, the reward systems of appointments and promotions continue to rely heavily on independent contributions to knowledge creation and communication (research and teaching), although a number of national initiatives, such as the NIH Roadmap, address this issue. Commercial success clearly depends on resolving the issues of sharing, interoperability, and privacy. We will call this the *interoperability* problem.

Ultimately any patient model—whether molecular or societal—**is an abstraction of the real situation**. Therefore it is imperative to understand the theoretical possibilities and limitations of each domain-specific abstraction. This is the *abstraction* problem.

Improved computational techniques for assessing clinically relevant variability in measurements are needed. This is the *variability* problem.

Experimental validation of models using ex-vivo and biomimetic materials and systems, animal models, and clinical data is needed. This is the *validation* problem.

Specific R&D Challenges

Besides the six challenges mentioned above, we face least two infrastructure challenges:

- Patient models need ongoing research and support.
- Issues of policy—privacy, security, and legal and regulatory issues—must be investigated and recommendations developed.

Research Strategies and Roadmap

We suggest that the six challenges be initially attacked in two groups of three and that the following strategies and roadmap be developed:

Multiscale/multistructure, data variation, and abstraction problems. These challenges have been recognized as important and to some extent are funded by a number of government research agencies, including NSF, NIH, and DARPA. We recommend that participants in the relevant programs at each agency meet as the focus of a working group in a national symposium to review progress and define future directions with high-confidence medical devices and systems.

Communication, interoperability, and validation problems. We recommend that these challenges be addressed by developing a series of demonstration cases. To build the foundation for an open-source environment that (1) addresses issues of ontology, (2) includes links to available models, data, and device sources, and (3) develops protocols for validation, we

recommend the creation of a "Knowledge Portal," consisting of a human anatomical atlas and a protocol manual, over the next two to five years.

Extensive data for the atlas exists, and new data may be readily available. The Veterans Affairs system already employs computerized medical systems, so the VA may be an excellent source of new data. The atlas should combine information from multiple patients and generate a coordinate system to "place" each patient. It should be searchable and should be able to generate statistical analysis. Ideally, the atlas would then be used to help predict outcomes based on individual characteristics and statistical outcomes. The atlas could help device companies to project the range of scales and sizes necessary for clinically useful devices.

The protocol manual would consist of detailed written descriptions of specific interventions, along with metrics to evaluate each intervention.

Three-Year Roadmap

Develop common ontologies:

- Descriptions of blood vessel branching for predicting cardiovascular surgery outcomes
- Descriptions of activities of daily living for safe performance in the home by the elderly

Five-Year Roadmap

- Develop statistical and analytical tools to analyze "on the fly" randomized trials
- Develop risk analysis tools that link procedures to outcomes
- Develop statistical methods for characterizing variability, abnormality, and anatomical variance
- Build multidisciplinary academic and industry teams for the production of high-confidence medical devices that will develop work plans; prioritize specific models; carry out preclinical and clinical trials to validate models and publish results of those studies; support the FDA approval process for the model and for medical-device validation; and maintain the model and support device manufacturers that use the model for FDA submissions

Although the above seem like complicated tasks, there is an existing example. The SRI/Stanford consortium consists of seven medical-device manufacturers to develop a model of a femoral artery stent. The consortium does data acquisition and modeling, and publishes the work, and the work product can be used for certification. Companies buy in and get prepublication data.

We recommend that the FDA, NSF, NIH, and NIST encourage public-private partnerships among academia, industry, and government.

Funding, privacy, security, and legal and regulatory issues are high-level, cross-cutting issues for the entire HCMDSS community. They should be addressed by an executive committee with representatives of academia, federal funding and regulatory agencies, and industry.

Concluding Remarks

Successful patient modeling will require the solution to a massive problem of information fusion and analysis. But the payoff to solving the problem is enormous. It will lead to better patient care at every level: better clinical results, better disease prevention, and, in principle, better use of costly medical interventions and potentially scarce resources. Fortunately, interest in the problem is intense across a wide range of biomedical, commercial, patient advocate, and regulatory communities. The challenge now is to integrate and channel that interest so that each community can contribute to the solution, recognizing that there is new knowledge to be generated, communicated, and implemented—knowledge that in the end could have a profound impact on the fabric of life represented in better personal and public health.

MEDICAL-DEVICE SOFTWARE DEVELOPMENT

Participants: Peter Lee (chair), Steve van Albert, Rajeev Alur, Douglas Barton, Ralph DePalma, Matthew Dwyer, Steven Getz, LeRoy Jones, Peter Kuzmak, Rami Melhem, Jens Palsberg, John Regehr, Victoria Rich, Michael Robkin, Gregg Rothermel, Vish Sankaran, Bow-Yaw Wang, Arvind Easwaran (recorder)

Introduction

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 involve the verification and validation of software-based systems. Exactly how to accomplish that is a major question, particularly because devices and systems are becoming increasingly complicated and interconnected. We may already have reached the point where testing as the primary means to gain confidence in a system is impractical or ineffective.

Further, the lack of uniform standards in the engineering of medical-device software leads to many deficiencies, such as the lack of precise system specifications. Both lacks inhibit even the best testing approaches and make it difficult for domain experts (cardiologists, neurosurgeons, and so forth) to ascertain whether a device will perform appropriately—even if it has been thoroughly tested. The lack of standards goes beyond software engineering practices and development technologies. Although the physician who performs an operation must be licensed, the developers who create the software used in the devices have no licensing requirements. Thus it is especially important that designs and software both be subjected to rigorous analysis. Such analysis is fundamental to effective and efficient testing, the analysis of systems-of-systems, and the determination and mitigation of risk.

In today's medical-device industry, embedded software and applications are often designed and developed by medical specialists—not by experienced software engineers. Although medical specialists are sometimes effective and experienced programmers, they often lack the software design expertise that has been developed in software engineering over the past 40 years. Typically, shortcomings in software analysis and design lead to rigid software and nonstandard platforms that are not resilient to change. The lack of resilience ultimately reduces the operational life of the underlying applications.

Along with the software-quality issue, research evidence suggests that the "frequency and consequences of medical device use errors far exceed those arising from device failures" (http://www.fda.gov/cdrh/humanfactors/important.html). This in turn suggests that techniques for *user-centered design* have not yet made a significant impact on medical-device design. Surveys reveal that medical-device manufacturers give variable attention to usability issues in device design. Also, modern medical practice continues to see dramatic increases in the already enormous amount of knowledge that clinicians must absorb and use. Medical devices often demand detailed knowledge of increasingly sophisticated and often fragile science and engineering. The trend reduces the probability that clinicians will understand the inner workings and essential mental models that underlie the technologies used in complex medical devices.

What Can We Do Well?

Recent years have ushered in many improvements in the technology and processes of software development. Even relatively simple tools and platforms, such as modern IDEs (integrated development environments), promote certain best practices (such as source revision controls) and are suitable for medical devices and systems.

Increasingly, manufacturers worldwide have adopted disciplined processes—such as those offered by the Capability Maturity Model Integration (CMMI) project—to help improve, assess, and sustain the quality of their products, along with their engineering, management, and methods of quality assurance. Those advances can have a great impact on the quality of medical devices, although they are not sufficient. Further research advances are necessary.

Medical Devices vs. Avionics Systems

A natural question is the extent to which the standards and practices used in avionics systems apply to software for medical devices. Although the two domains share the need for safety-critical software, their standards and practices have large differences. Perhaps the most striking is the almost complete lack of regard, in the medical-device software domain, for the specification of requirements. Those in the medical-device community often say, in essence, “We don’t need requirements in developing medical devices.”

Although such statements may seem startling, the business models and incentives for medical devices lead to the development of highly proprietary technologies. This necessarily decreases the interaction among development teams and diminishes the perceived value, at least in the short term, of specifying requirements. It also presents a significant barrier to academic researchers’ participation in medical-device technology—in sharp contrast to their participation in avionics.

Other technical metrics, such as availability and mean time to failure, also seem to be largely absent or appear only in diminished form in the development of medical-device software.

Culturally, the sense is that medicine is a “people-intensive” activity that necessarily has highly complex and individually tailored workflows. Although software is a critical enabler of those workflows, it is often viewed as a minor contributor to failures in the system, failures that can happen in many ways. So although information technology is a key enabler to lowering the cost and improving the quality of patient care, the software quality itself is not perceived to be, and in fact may not be, the single most critical issue.

It is unclear whether differences between avionics and medical-device software will continue to be significant. The development of each medical device as a separate “stovepipe” seems untenable for the future. Thus we anticipate that the development of medical devices will use practices much more familiar to the safety-critical software community, particularly as medical-device systems are used more and more in closed-loop situations—in which medical-device systems adapt to changes in patients’ conditions without caregiver intervention.

Why Can't We Declare Victory?

The current state of the practice in medical-device software has several key problem areas. These are the most critical:

Poor Quality of Software and Software Architectures

The development of complex safety-critical software is the subject of much research in the computer science and software engineering communities. Even so, numerous IT-based medical systems have clearly identified faults that can be addressed without any advances in research. In some cases, easy technical solutions exist, but other barriers (usually nontechnical) prevent their application. Often those barriers involve the lack of standards, of interoperability, and of metrics in the field. Issues with legacy systems also hamper progress.

One system, for example, has a problem in which a Web form for one patient was automatically—and incorrectly—filled in with information for a different patient, most likely because of an interaction with "cookies" in the Web interface. In other cases, specific networking failures arise.

In still other cases, the faults lie in poor technical standards for interfaces and architectures. For example, the bar-coding system of the Bar Code Medication Administration (BCMA) lacks a uniform standard across blood suppliers. Furthermore, the bar-code tags themselves are physically unreliable. Those problems, coupled with faulty or incompatible bar-code readers, have led directly to patient deaths.

As yet another example, the identification system of the Digital Imaging and Communications in Medicine standard Health Level 7 (DICOM/HL7) is neither standardized in the industry nor properly implemented in deployed devices. This has led directly to the loss of medical records.

Some of those problems have immediate technical solutions that repair the faults and defects. But because they are legacy systems, applying the technical solutions poses many practical difficulties, even if the device manufacturers are made aware of the solutions. Furthermore, the market may or may not offer incentives to apply the solutions.

And so in each case, we have situations where specific algorithms and technical solutions may be known to the research community. How to apply them in the real world is a major question.

Feature Creep

Industry perceives that the market demands more and more features. This "feature creep" not only increases risk and complicates use but also leads to software interactions that are hard to specify and even harder to analyze. The result is that no one, not even the developers, understands how the systems behave.

Even in the research community, the value of testing and techniques for validation in the face of new features is poorly understood. Furthermore, developers lack metrics for reliability and

usability. This makes it impossible to assess the costs and benefits of various validation approaches.

Usability Problems

Workflows and operations have become overly complicated because of the proliferation of devices and technologies. That proliferation has provided clear benefits to the quality and cost of patient care. But it can introduce significant and often unnecessary complications to hospital procedures. And it can lead to situations that are confusing to caregivers (such as task overload) and unacceptably perilous to patients.

Practitioners describe being on a "gadget treadmill," with constantly changing workflows and increasing distractions in the operating room. Some practitioners report difficulty knowing whether the operating room has been fully "switched on." In one OR, a nurse had the responsibility of making sure that the room was ready to go and that the operation could proceed, but the nurse was so focused on entering data that the main procedure was almost an afterthought! Such intensive distraction necessarily puts some patients in peril.

Beyond having to understand how to operate numerous devices, we see reported issues with version skew—situations where a device has two, three, or even more versions, whose interactions cause a system or network failure. Finally, when something fails, medical devices often have no audit capability, or "black box," that permits a post-mortem analysis of the failures.

Lack of Knowledge Sharing

"Knowledge saves lives" is a clear refrain in the medical community. Data can be transformed into information, and information into knowledge. IT is thus a key enabler, because even simple advances such as faster data entry can have a measurable impact on the number of lives saved. The sharing of information is hampered by the lack of interoperability standards and technologies and by confusion over privacy regulations, specifically the Health Insurance Portability and Accountability Act (HIPAA).

Aggravating the problem is the medical-device industry's proprietary nature. Each medical device or system is developed largely on its own, with little regard to the sharing of critical information across systems. Manufacturers seem to have clear commercial incentives *against* standardization, which presents yet another barrier to information sharing.

Lack of a Systems Engineering Perspective

Integrating technology into the clinical environment—which includes practitioners, workflows, and specific devices—often lacks a holistic, systems perspective. Many medical devices are designed, developed, and marketed largely as individual systems or gadgets. Device integration, interoperability, and safety features are not considered during development, acquisition, or deployment.

The situation is not unique to medical devices. In both research and development, the tendency is to look at systems in isolation, ignoring how well they will work in a larger setting. Workflow

issues, including people and clinical procedures, could be reengineered to make safer and more optimum use of some technologies, but that is rarely considered.

The lack of a systems perspective results not only in safety problems but also in lost opportunities for exploiting advances in information technology. One issue that is in dire need of progress and for which there are exemplars in other industries is version skew. The automotive industry, for example, uses analysis and modeling tools that help manage version skew, but similar tools have not been widely adopted by the medical-device industry.

Specific R&D Challenges

Medical and assistive devices must be dependable. Dependable devices work as intended, are highly available even when not well maintained, and do no harm when they fail or are misused. Devices should also be customizable and easy to use, upgrade, and maintain.

Although those criteria apply to many domains of software application, we find that some research challenges and needs are either specific to medical devices or deserve additional emphasis compared with other domains. We summarize the challenges below.

Formal and Model-Based Analysis, Design, and Implementation

Methods of software engineering that are increasingly applied in other application domains go a long way toward improving the quality of software-centric systems. The methods range from relatively informal (though still disciplined) process models to rigorously formal methods based on mathematics and logic. One unifying theme is formal modeling.

At the relatively informal end of the spectrum, modeling languages such as the Unified Modeling Language (UML) and development methods such as Model-Driven Architecture (MDA) confer an important level of discipline and organization on the software design process. That discipline facilitates the early discovery of design problems and helps ensure that independently developed components work well together.

Although model-based approaches are normally applied to standard application software, experience shows that they also benefit embedded systems. But two features of medical systems—and of embedded systems in general—stand in the way of realizing the full potential of model-driven development. Correctness requirements for medical systems often include detailed operational requirements. Thus, commonly used modeling languages such as UML, which concentrate on a system's structural properties, fall short of the design needs. To make matters even more complex, medical devices work in complicated, dynamically changing environments. Without adequate modeling of the environment, it is impossible to make the model-based design of medical systems reliable. Also, the safety-critical nature of medical embedded devices requires a higher degree of validation than most software engineering processes provide.

Therefore, some amount of formal verification and validation needs to be at the core of modeling technology for medical software. We need to develop methods of modeling that naturally support operational specifications and enable rigorous verification, and we need to incorporate those

methods into existing design processes. The quest for such modeling techniques and practical formal methods will probably be the biggest challenge to the design of software-based medical systems in the future.

Medical Practice–Driven Design and Validation

Like a sound foundation for integration, better tools and environments as described above can help us *build systems right*. Equally important is *to build the right system,* as pointed out by Wears and Berg in their March 2005 *Journal AMA* article "Computer Technology Still Waiting for Godot." The authors suggest that the root cause of errors such as the medication errors introduced by physician order-entry systems is that "the pattern of [IT] use is not tailored to the workers and their environment."

Although research has been performed on iatrogenic injuries (caused by a doctor or other caregiver), researchers rarely investigate medical errors involving poor device design. Besides patient injury, poor design causes other maladies, such as reductions in treatment efficiency and effectiveness, excessive training or maintenance costs, stress and confusion for users, and other attending complications. Any study based on patient harm, however, will vastly underestimate the systemic consequences and costs of poor integration and interface design. Specific usability challenges for overcoming those issues include the following:

- **Metrics for usability impact**. Studies have shown that, for example, the average software program has 40 design flaws that result in lost productivity. But little work has been done on the impact of medical-device design. We need studies and data collection methods to better understand the impacts and costs of poor and suboptimal device design.

- **Capture and dissemination of usability guidelines**. A wealth of literature exists on design principles for medical devices. But the knowledge is often relatively useless because it lacks information that attaches principles and guidelines to usability contexts. In addition, many guidelines are rarely updated. Given the fast pace of technological change, we need dynamic methods that capture and disseminate emerging knowledge.

- **User-centered design methods**. Adhering to even the best usability guidelines and principles is not sufficient to guarantee appropriate device design. Diverse usability settings dictate that designers must work closely with users to match device design with specific practices and conditions of use. User-centered design methods have been created for design stages including requirements and inception, summative evaluation, end-user support, and post-release instrumentation and maintenance. What is missing is a means to unify those methods so as to better understand usability phenomena. In addition, existing user-centered techniques are not designed to address issues stemming from complex interacting systems (systems of systems) comprising users, devices, and use environments.

- **Diverse contexts of use**. User interface design is difficult partly because interfaces are influenced by many situational variables. They include environmental context (such as offices visits, surgery, and emergency room), user populations (doctors, nurse practitioners, nurses, anesthesiologists), devices (internal, monitoring, analysis), interface

type (device controls, screens, handhelds), and interaction type (selection sequences, parameterization, programming), to name a few. Design methods must therefore have the flexibility to accommodate the highly diverse and situation-dependent nature of interface design.

- **End-user programming and customization**. Interface diversity is a direct reflection of the diversity of individuals and their medical conditions. Medical devices are designed with varying levels of end-user customization, from setting up parameters to tailoring environments involving forms of end-user programming that individualize therapies carried out by medical devices. As devices become more flexible, the potential for error increases. Medical-device interfaces must be designed to minimize and avoid life-threatening errors and parameter-setting combinations by, for example, detecting anomalies or outliers in parameters and making sure that programming directives do not break parameter invariants.

Clearly, we must adopt the user-centered approach, which calls for considering users (including caregivers, patients, and service personals) throughout the acquisition, design, implementation, and evaluation of requirements. Research is needed for user-centered design and quality assurance, including the creation of a library of user scenarios, user models, and their environments in the context of medical devices and systems based on real data. We also need standards similar to the International Standards Organization's ISO 13407 (User-Centered Design Process) to guide user-centered derivation of requirements and design and evaluation against the requirements.

Achieving User-Centered Designs for Medical Devices

Many techniques for modeling and analyzing user-interfaces are well known and have been investigated in various domains. But medical-device software is unique in many respects and brings concerns about safety and dependability that have not been adequately researched. We need empirical investigations so that we can better understand those specialized needs.

For example, it is important to ensure that at every step in a medical intervention, the supporting medical software reflects a valid state. At design time, standard medical procedures should be modeled in such a way that the use of the software is understood and documented. This may require the development of workflows and use cases for medical procedures—and some efforts have been started in that direction. Later, preoperative models should be used in real time to confirm the steps necessary during the medical procedure. Monitoring approaches of this sort should also incorporate support for error-handling and fault tolerance. With solutions to these challenges, the effectiveness of medical processes for common procedures can be evaluated and generally enhanced.

We also need state-based control standards, mechanisms, and diagnostic tools for medical software. Those items will help connect the actions in a medical operation with the corresponding actions of the supporting software. In addition, we need general validation approaches and real-time diagnostic tools to support preoperative and postoperative procedures.

As a next step, general validation models should be investigated to validate medical processes during medical procedures.

Because context is vitally important to the development of high-confidence user interfaces, collecting and analyzing field data are crucial. Research is needed on technologies that capture not only data from laboratory studies, surveys, and interviews but also data from the field on operational software. Software instrumentation and profiling techniques, for example, could capture how users supply incorrect information and interact erroneously with devices.

Component-Based Design and System Integration

Contemporary software development emphasizes components that have clearly specified application programming interfaces (APIs). A static API for a software component such as a Java library class consists of all the (public) methods, along with the types of input parameters and returned values that the component supports. This promotes a clear separation between the specification of the component and its implementation. Such static APIs can be enforced using type systems of programming languages. But although type systems are indispensable, they offer only a partial solution to designing bug-free software, because they do not capture constraints on resources, real-time guarantees, and other quality-of-service aspects. Consequently, they offer little assistance in “system” integration. This is an important issue, not only for being able to derive system-level performance and correctness guarantees, but also for being able to assemble components cost-effectively.

The interface for a device that interacts with a patient must incorporate information about timing delays and continuous parameters such as threshold levels. Capturing the notion of quality-of-service abstractly, and having mechanisms that can enforce the adherence to interfaces and that can check compatibility between interfaces, is already an emerging and challenging trend in research on embedded systems. The additional concerns with medical practice–driven design and formal analysis increase the need for advanced research. On the other hand, by identifying key component types relevant to medical devices, we open new opportunities to apply what is already known. Providing an impetus to the research community to understand and address the issue is a critical need.

Open-Research Test Beds

Today we have open-research platforms that provide highly effective support for the widespread dissemination of new technologies and even the development of classified applications. The platforms also provide test beds for research collaborations involving both researchers and practitioners. One spectacular example is the Berkeley Motes system with the TinyOS operating system.

The medical-device community could benefit from the existence of such open-research platforms. They would enable academic researchers to become engaged in directly relevant problems while preserving the need for proprietary development by the industry. (TinyOS facilitates academic input even on government-classified technology, which is an example of what is possible.)

Taking steps to enable the creation of such test beds should be considered immediately.

Research Strategies and Roadmap

To achieve assurance guarantees for medical devices, we need a paradigm shift—one that includes a formal approach to design, implementation, and analysis. Model-based and formal methods have been successful in targeted applications such as microprocessor designs. Recently, formal verification tools have fully analyzed and verified a number of low-level (device-driver level) and embedded avionics applications. We believe that the same success is feasible in the domain of medical devices. Such an approach would comprise the following steps:

1. Define requirements of the system in a mathematically precise notation.
2. Design a high-level model of the control algorithm for the medical device.
3. Design a high-level model of the environment in which the device will operate.
4. Subject the device model, together with the environment model, to powerful analysis techniques, such as simulation, optimization, and verification.
5. Either generate code automatically from the device model or verify that hand-written code is consistent with the model.

Three-Year Roadmap

In view of these R&D challenges, it seems clear that research needs are not just in information technology but in something that is much more multidisciplinary and involves a systems approach. It includes significant computer science along with biomedical engineering, device manufacturing, and the care process.

- Develop open, experimental platforms and standardized research data sets to foster collaboration between academic, industry, and industry
- Extract models from clinicians and their workflows, analyze change effects, measure system performance, and so on, to support a systems engineering approach
- Academia, industry, and the government should work together to address the technically straightforward, well-known deficiencies in today’s medical IT systems. Specific examples are universally unique identifiers for medical-information objects and standardized bar-code technologies in blood bank applications and administration.

Five-Year Roadmap

In five years, it should be possible to perform experimental evaluations of the effectiveness of integrated systems:

- Develop medical practice-driven design and validation, including metrics for usability evaluation and user-centered design methods

- Device component-based designs and systems-integration-based formal models of medical devices, control algorithms, and workflows of medical practice
- Create open-research test beds to disseminate new technologies and develop a set of applications

Ten Years

In ten years, we expect to have model-based techniques for producing complex integrated systems cost-effectively. The integrated systems will be subjected to quantifiable assurance measures that can be the basis for certification:

- Develop model-based validation and certification of systems of medical device systems
- Develop medical-device software and systems that can prevent misuses by caregivers, adapt to changing patient conditions, and can be safely used in closed-loop situations

FOUNDATIONS FOR INTEGRATING MEDICAL-DEVICE SYSTEMS AND MODELS

Participants: W. Rance Cleaveland II (chair), M. Brian Blake, Andrew Casertano, Sherman Eagles, Scott Henninger, David W. Hislop, Zachary Ives, Tomasz Petelenz, Jane W. S. Liu, Tom Martin, Gregory Sharp

Introduction

Medical technology is in the midst of a profound technological revolution. On the one hand, advances in computing have led to the development of novel diagnostic and therapeutic equipment and procedures, ranging from advanced imaging machines to minimally invasive surgical techniques and implantable heart devices. On the other hand, low-cost computing equipment and improved software have enabled the automation of many administrative aspects of health care, including hospital management and insurance-claim processing. The result has been operational efficiencies that open the door for the delivery of lower-cost, higher-quality care.

At the same time, the national dialog on health care continues to focus on the inexorable rise in costs and on concerns about the quality of care. The press lauds advances in standalone diagnostic and treatment systems—and rightly so. But systemic inefficiencies in health care delivery grossly inflate costs and contribute to avoidable medical errors that degrade patient care.

Those inefficiencies are primarily caused by a lack of mechanisms to create integrated medical systems that coordinate the collection of information and the delivery of medical services. The health of all Americans would benefit greatly from research into technologies that support the seamless and secure sharing of health information; enable the integration of hospital-based monitoring, diagnostic, and therapeutic devices; and provide capabilities for remote monitoring and the treatment of chronic disease.

The technical recommendations in this chapter derive from a core observation: The systems of medical devices needed to implement the data-sharing and distributed-care delivery described above are impossible to build in a cost-effective, timely manner. This stems largely from an ad hoc approach to designing, developing, and deploying such systems. There is a pressing need for well-understood models, theories, and tools for reasoning about medical-device interfaces, the composition and integration of systems from component devices, and systems of systems that support the predictive analysis of end-to-end system properties that are important to medical applications. Among those properties are—

- Efficacy
- Safety
- Security
- Privacy

- Traceability
- Confidence

What Can We Do Well?

This section characterizes the state of the art for medical-device and medical-system integration and modeling in terms of technologies used in three types of integration scenarios:

1. **Medical-information systems**. Such systems manage patient and caregiver information. The integration of such systems is intended to reduce the overhead and risks associated with information management.
2. **Point-of-care-based monitoring, diagnostic, and therapeutic device and systems**. These systems are typically overseen by medical experts. They are involved in the treatment of specific patient disorders. Integration issues include individual device design (Will a device perform according to requirements? Is a device easily and appropriately integratable with others) and system design (Will one device interfere with another? Can system behavior be easily inferred from the properties of an individual device?).
3. **Extramural monitoring, diagnostic, and therapeutic devices and systems**. These systems may be home based or implanted and thus are beyond the immediate control of medical personnel. Integration issues include the ones mentioned above for hospital-based systems. Other issues involve the delivery of information from these systems to caregivers in a manner that is timely and secure, and that respects privacy.

Communication Standards

Efforts to develop integration frameworks for medical devices and systems have focused on defining communication protocols and data-format standards for medical data. Some standards have focused on specific areas in medicine, a prominent example being the DICOM (Digital Imaging and Communications in Medicine) standard for medical imaging. DICOM defines a format for the digital storage of images that is independent of imaging technology. It is in widespread use.

Other work has focused on issues related to the exchange of medical data among medical-information systems. For example, the mission of the Health Level 7 (HL7) standards organization is *"To provide standards for the exchange, management and integration of data that support clinical patient care and the management, delivery and evaluation of healthcare services. Specifically, to create flexible, cost effective approaches, standards, guidelines, methodologies, and related services for interoperability between healthcare information systems."* Standards developed by HL7 include the Version 2.5 Message Standard, which defines message formats and protocol standards governing patient control (admission, discharge), order entry (dietary, pharmaceutical), financial management, and the like. According to the HL7 Web site (http://www.hl7.org), Version 2.5 is the most widely implemented communication standard

in the world. As part of the Version 3 Message Standard, HL7 has developed the Reference Information Model (RIM) for clinical and administrative health care data.

Other standards bodies have focused on point-of-care communication protocols. The IEEE 1073/ISO 11073 specification, for example, defines rules for connecting bedside devices so that that they can exchange data. The standard includes components for a variety of medical devices, such as respirators, defibrillators, and electrocardiographs. The IEEE 1073 Web site (http://www.ieee1073.org) refers to 1073 as the basis for plug-and-play medical-device systems.

At present, no widespread standards exist for so-called extramural medical devices. The main reason is the relatively recent emergence of a wide variety of such devices that improvements in power consumption have enabled. Improved radio and infrared technology makes the communication of data from devices to external medical equipment feasible, but data standards have yet to be developed. Similarly, we know of no standards governing communication between point-of-care systems and medical-information systems. Although communication and data standards exist, the bulk of clinical and administrative information remains encoded either on paper or in proprietary computer systems.

Standards and Regulations for Medical-System Development

The development of medical devices and systems is a safety-critical undertaking. Faulty devices can endanger patients' well-being and even lives. Most efforts aimed at improving the development of such devices have taken the form of standards and regulations governing best development practices.

The IEC 60601 family of standards is a case in point. The standards specify requirements for electrical medical devices (defined as those that have at most one connection to a power main and that transmit energy to a patient). They contain general information about electrical shock, radiation, and fire hazards as well as device-specific guidelines for electrical safety. The FDA recognizes IEC 60601 as a consensus standard and advises development organizations that adhering to the standard provides reasonable assurances of electrical safety.

The ISO 13485 standard for quality-management systems defines rules for assessing how robust a medical-device developer's quality-assurance processes are. The standard is an adaptation of ISO 9001 (a general quality-management system) for the medical industry and reflects the industry's considerable regulatory concerns. The revised Quality System Regulation (21 CFR 820, Oct. 7, 1996) (QSR) is also based on the ISO 9001 and ISO 13485 quality-system standards.

Other standards deal with risk management for medical devices. Emblematic of them is ISO 14971, which defines the components of a medical-device risk-management process, including policy development, training, and techniques such as Failure Mode and Effect Analysis (FMEA) and Fault-Tree Analysis (FTA), which are standard techniques for determining root causes of failures in systems.

Other governmental bodies have established regulatory frameworks to ensure medical-device quality. Of particular note are the Medical Device Directive (MDD, European directive

93/42/EEC), and guidelines for active implantable medical devices (90/385/EEC). These standards and regulations deal with individual medical devices, not systems of devices.

Design, Development, and Validation Technology for Software

As in other industries, the development of software for medical devices and systems remains largely a systems-based, standards-driven, document-intensive process. Standards such as those mentioned in the previous section are used to organize development organizations and to manage risk. The testing of device software is done primarily at the system level, using test benches containing hardware and software components. For radiological devices, “phantoms”—bags of fluid each about the size of humans, with sensors for measuring radiation exposure and the like—are used as mockups of human bodies to assess safety and performance.

Some organizations are beginning to use model-based design and development. Examples are use-case modeling of requirements for information systems and computer-assisted surgery, and state machines to model infusion pumps and implanted cardiac devices. Such technologies are rarely mentioned in standards documents, which points to their immaturity in the medical arena.

Why Can’t We Declare Victory?

The state of the art in medical-device and medical-system integration consists of two key components: low-level data-transfer protocols and standards, and standards and regulations specifying best practices for device development. Most medical information remains spread throughout hospitals and caregivers’ offices and in insurance-company files. Patient bedsides consist of an array of devices connected, if at all, by cables that are easily disconnected and cause hazards. Operating rooms contain a hodgepodge of equipment and displays that can overwhelm medical personnel with data. Implanted devices suffer from recalls and operate in an information vacuum. After insurance paperwork, the first form a patient completes when visiting a doctor is still a medical history, which the patient may or may not recall. The transfer of data from the point of care to patient records still relies on human intervention.

Compare this situation with the following scenario. Everyone has an electronic health record (EHR) containing a complete account of his or her health history and current insurance information. That information can be queried by caregivers—no more haphazard history taking—and medical devices themselves, which could provide additional safety interlock features based on a patient’s history. Bedside devices would communicate wirelessly—no cables—and could be configured to forward collected data automatically to a patient’s medical record. Access to patient records would no longer be a bottleneck to consultations among caregivers. An operating room would feature a single integrated display for all equipment in the room, which would communicate with the display wirelessly—no cables to trip over—and would record information in “black boxes” for subsequent review of poor outcomes. Implanted devices and wearable monitoring equipment could record information in a patient’s EHR, making it available for caregivers’ review.

That vision represents a radical departure and entails significant investment in the research areas listed below. The benefits to human health, however, would offer a many-fold return on the investment.

Specific Research Challenges

Realizing the vision described above will require substantial research into technologies for medical-device and medical-system integration. The technologies are described below.

Electronic Health Records (EHR)

One backbone of the preceding vision is an accurate and complete electronic health record. Significant technical challenges confront the development of such a record, because of the extremely personal nature of individual health information and the many institutions that would need access to some, but not all, of the information. To be widely adopted, EHRs must be secure (to prevent unauthorized tampering) and provide privacy guarantees, so that control may be exercised over who has access to different components. In addition, EHRs must contain audit-trail mechanisms to record who has seen what. At the same time, those mechanisms must be minimally intrusive and must support the incorporation of legacy health data, such as paper records.

Plug-and-Play Network Devices

Another enabling technology for the aforementioned vision is the development of plug-and-play networking technology for medical devices. Plug-and-play capability is needed to ease the setup of integrated point-of-care and extramural arrays of medical devices that communicate with a patient's electronic health record.

Devising the technology would require addressing concerns about privacy, security, safety, regulations, and technology. In hospital settings, for example, networks would form and reform frequently, as patients are admitted and discharged. Technology for the rapid formation of ad hoc networks needs developing. At the same time, authentication mechanisms would be needed to ensure that a device is on the correct network—and not, say, incorrectly attached to the one next door. Communication among devices would need to be made more secure than current wireless technology supports, and the problem of incorporating legacy devices must be addressed. Regulatory approaches ensuring the safety of open networks would need to be developed. Finally, for portable and implanted devices, plug-and-play technology would need to minimize power consumption, implying the need for low-power-consumption communication protocols.

Ergonomic and Ease-of-Use Issues in Human-Device Interfaces

The integration of devices in a care-giving setting requires careful attention to considerations of human factors—especially in extramural settings, where nonprofessionals would interact with devices. Issues of mode confusion and data fusion (combining data from various sources into a coherent form) will also need to be addressed, as point-of-care facilities become more complex and begin to resemble cockpits in airplanes.

Monitoring and Post-Intervention Analysis

In the integrated medical-device and medical-system setting considered here, faults and errors will occur: Devices will fail; caregivers will make mistakes; networks will crash. To cope with such problems, technologies must be developed for error handling, fault tolerance, fault diagnosis, and fault isolation in ad hoc wireless networks. “Black-box” standards—similar to those in aircraft—should be developed so that when systems do fail, lessons can be learned and responsibility assigned.

Component-Based Design and Validation

The technologies outlined in the previous sections will find use only as long as integrated systems of medical devices can be cost-effectively safe, robust, and secure. Current design and validation mechanisms for medical system generally do not support open systems that share data while guaranteeing end-to-end privacy and security. Technologies for lifting reasoning about such properties from the single-device, single-component level to integrated component-based systems must be developed, most likely via notions of safety and security interfaces that would specify the safety and security properties that a component guarantees. Design patterns and service-oriented computing paradigms tailored for medical applications, with their specialized mix of security and privacy concerns, must also be built, and standardized modeling paradigms for open systems of communicating medical devices should be devised, so that “virtual” device networks can be built and simulated for validation purposes. Such simulation models would also be useful in training personnel.

To further support such virtual validation activities, accurate control-theoretic patient models must be developed so that medical-device model behavior can be simulated and assessed before trials on humans. Models of caregiver behavior should be developed for devices so that close-loop modeling of device, patient, and caregiver can be undertaken. Such simulations would yield useful information about likely caregiver errors in using devices, and that information could be used to improve user interfaces to devices and lower the likelihood of device misuse due to caregiver error. To guide subsequent system testing and risk analysis, we need to develop methods to quantify residual risk from virtual validation activities—how accurate are the data derived from simulations, and how can the analysis be used to focus efforts in subsequent clinical-trial testing? Strategies for in-vivo test planning based on virtual testing would open the door for additional efficiencies. Issues of data integrity and integration must also be studied. Efficient techniques for converting models into executing systems must be developed.

Open Experimental Platforms

The final component of an effective research program would be an open experimental platform for use by researchers investigating technologies to integrate medical systems and devices. Such a platform would contain design artifacts, including reference models and scenarios about the use of different medical devices, so that researchers could obtain empirical feedback on their ideas about real systems. An open experimental platform would be vital to the success of the research program, because concerns about intellectual property would otherwise preclude the sharing of device and system information.

Research Strategies and Roadmap

This section offers roadmaps for the research program outlined in the previous section.

Three-Year Roadmap

The goal in the first three years will be to begin developing conceptual frameworks for two core issues: the security and privacy of medical data.

- Reference models for privacy and security should be developed, with a view toward defining the stakeholders (such as patients, caregivers, insurance providers, and government) and the levels of access and control they have on patient information.
- Precise mechanisms for specifying the safety and security properties of medical software components must be developed, and use cases and reference models for standard medical devices such as infusion pumps, radiological devices, and cardiac devices must be elucidated. That software work will entail adapting software-design notations to medical purposes.
- Control-theoretic models of human biological processes should be developed, and work on caregiver modeling initiated.

Five-Year Roadmap

The next phase of the research program will involve the development of tools and methods for technical support of the conceptual framework mentioned above. (The conceptual framework will continue to evolve.) Specifically—

- Technologies will be developed for electronic health records that enforce the privacy and security policies for medical data.
- Techniques for establishing end-to-end safety and security properties of systems based on component properties will be defined, and prototype automated tool support implemented.
- Protocols for ad hoc secure networks of medical devices will also be defined, to ensure adherence to the security and privacy policies of medical data by collected data.
- Virtual validation environments for medical devices, based on models of devices, human biology, and caregiver behavior, will be developed, and pilot studies will be conducted on the medical-device information contained in the developing open experimental platform.
- Human-factors studies for medical-device interfaces will be initiated.
- Service-oriented paradigms and design patterns for medical-device integration will be defined, and prototype tool support for the design of medical-device control systems begun.

Ten-Year Roadmap

The final phase of the research effort will focus on further developing the technologies and major case studies:

- Implementations of electronic health records will be assessed empirically for utility and safety.
- Improved control-theoretic models of human biological and physiological processes will enable a more thorough analysis of systems of medical devices, allowing the efficient development of multifunctional systems that enable the coordinated treatment of disorders with multiple causes (such as diabetes-induced heart disease).
- Clinical trials will allow the quantification of residual risk in virtual, simulation-based validation approaches to medical-device systems, and as a result the virtual validation approaches will be refined.
- Service-oriented design paradigms for medical systems will be codified, and design and validation tool support for them will be developed.
- Secure, ad hoc networking protocols for medical applications will be implemented in both hardware and software, laying the basis for a commercial market for medical-device middleware.
- Human-factors studies will feed into this work so that systems of medical devices can be assembled easily, and with user interfaces providing maximum support for caregivers and patients responsible for their use.

VERIFICATION, VALIDATION, AND CERTIFICATION

Participants: John Hatcliff (chair), Peter Coronado, Steve Dimmer, Jacob Flanz, Elsa L. Gunter, Mark P. Jones, Paul L. Jones, Brad Martin, Rosemary C. Polomano, Stacy Prowell, John Rushby, Rich Schrenker, Oleg Sokolsky, Aaron Evans (recorder)

Introduction

Verification and validation (V&V) 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. Many industry observers believe that we are approaching the limits of reliable device V&V processes. As devices grow more and more complex and rely much more on embedded software to achieve critical functionality, existing methods are being challenged. The results: higher development costs for manufacturers, longer time to market, and increased chances of device failure—with associated recall or liability costs.

Today, V&V activities account for as much as half the cost of bringing devices to market. Moreover, it is important to consider the effectiveness and already high costs of verification, validation, and certification 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 health care, and although V&V 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 the domain of health care systems that have transformed other sectors of nation's infrastructure and economy.

For example, pervasive networking will enable the integration of national networks, regional health care centers, local hospitals and clinics, the offices of primary-care physicians, home computing, and body-area networks. And as generations of technology-savvy health care consumers enter retirement, they will embrace—and even demand—sophisticated home health care 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 V&V processes obsolete. Unless new certification technologies are developed and unless V&V processes undergo a paradigm shift, innovation will be stifled, because manufacturers and regulators will find the V&V of systems too costly—or we will see dramatic increases in security breaches and harmful incidents due to device malfunction.

What Can We Do Well?

Designing bug-free software is difficult—especially in complex devices that may be used in unanticipated contexts. Existing practices have worked as well as they have because industry V&V personnel and regulators take their jobs seriously.

We know how to develop and regulate **standalone and embedded medical devices** that have moderate complexity and are based on mature technology. In such cases, the domain is generally

well understood, and the technology provides a level of confidence because the evolution of devices is incremental.

Wallace and Kuhn, in “Lessons from 342 Medical Device Failures,” clearly indicate that many failures could be prevented just by systematically employing known quality assurance techniques.

Industry seems to perform the following V&V activities reasonably well:

- Gathering requirements
- Coding
- Testing
- Performing hazard analysis

In particular, collective knowledge and experience within large, well-established companies aid the effort necessary to prepare for validation activities.

The adoption of state-of-the-art techniques for software development is not uniform across the industry. Some companies are beginning to incorporate advanced design techniques into their development processes. Guidant, for example, uses rich modeling tools and product-line architectures to develop its pacemaker software. Clean-room software engineering and modeling have also had some success.

Automated tools for V&V activities—ranging from homegrown configuration management systems to static analysis tools—can have wide-ranging benefits. They can raise overall quality by reducing the risk of human error in applying quality assurance techniques and by reducing the time and effort needed to apply the techniques. A reduction of time and effort can free up resources that can be applied to obtain greater and denser coverage of code and to encourage innovation and technological advances by speeding up time to market.

Why Can’t We Declare Victory?

Despite the pluses listed in the previous section, existing technology faces many challenges:

- **Lack of Tools that Automate Certification Tasks**. Generally, we know how to perform activities such as writing requirements, prototyping, and testing, but we need to perform them more accurately and in a more automated way. We also need to achieve greater cross-leveraging of techniques and artifacts—that is, we need to create architectures and processes that allow the information produced by one tool to supplement and increase the effectiveness of other tools. Currently, almost all tools are standalone tools and are not context aware.

- **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—slows the time to bring devices to market. Engineers often feel overwhelmed by complexity.

- **Limited and Ad Hoc Systems Integration**. We are able to integrate products developed by a single manufacturer. But in a clinical environment, people are incorporating devices from different manufacturers, and technical personnel must be trained to recognize potential incompatibilities.

- **Inability to Capture and Model Clinical Environments and Processes**. The correct operation of many devices relies on assumptions about the clinical environment or process in which the devices will be used. We are experiencing device failures due to unexpected interferences between devices and environments. Clear specifications of environmental assumptions and processes are needed. Formal models are preferred because they would enable tool integration.

- **Lack of Component-Based Approaches to Certification**. Current V&V procedures focus on approving complete devices or systems—not components of systems. The emphasis is on end-to-end testing against device requirements. No regulatory mechanism allows reusable, individually deployable components or infrastructure to be certified in a way that reduces overall certification costs when the components are used in larger systems.

- **Inadequate Coverage of Security Issues in Certification**. In today's V&V of devices, security issues are almost never considered. Yet the increasing blending of devices, medical records, telemedicine, and home, local, regional, and national networks will make security a central concern in the certification of future systems.

- **Widely Varying Knowledge and Training Across Vendors**. Although some device vendors use state-of-the-art development practices and quality assurance techniques, many are uninformed about best practices for techniques such as test case generation, testing to different coverage metrics, and capturing requirements. Indeed, studies have concluded that the nature of many reported device faults indicates that known practices may have been misused or may not have been used at all. Moreover, clinical engineers receive little formal training—they gain knowledge through experience.

- **Device Interference and Poor Integration**. Industry is doing fairly well at integrating products developed by a single manufacturer (such as Varian's Linac, used in concert with imaging devices and treatment planning systems). As the benefits of system integration are realized, however, and as integration mechanisms become commonplace in other domains, such as the highly successful and widely used universal serial bus (USB) from the personal-computer domain, the number of attempts to connect and integrate devices is increasing significantly. Such integrations are largely proceeding ad hoc, with little or no documentation and no systematic training.

Specific R&D Challenges

Based on the assessment presented above, we now lay out specific R&D challenges associated with certification.

Formulating Requirements

The development process begins by formulating requirements. The requirements formulation phase is the most important and one of the most-challenging-to-execute phases of development for medical devices.

We need technologies that support more automated and rigorous methods of eliciting and capturing requirements as described by experts in their respective domains. We also need technologies that can leverage requirements, including technologies that—

- Directly perform semantic querying or simulation of use cases
- Derive test cases more effectively and automatically from requirements
- Generate visualizations and natural-language descriptions from requirements
- Generate artifacts and reports needed for certification directly from requirements

Not only must we capture requirements that are appropriate; we must also automate capabilities for assessing and evaluating requirements.

Modeling and Formalizing Clinical Environments and Processes

One trend in the practice of medicine is a proliferation of sophisticated devices in complex environments where exceptions are the norm. We must develop a more rigorous approach to capturing, visualizing, and reasoning about clinical environments and processes. In particular, research should address the following questions:

- How do we incorporate variability in clinical environments?
- Can we mine rigorous descriptions of processes, workflows, and so forth to guide the development of device requirements or detect problematic device interaction in a clinical environment?
- Can we use those models to clarify goals and focus the tasks to achieve more rigorous clinical validation?

Benefiting from Automated Tool Support

The FDA has relied on device manufacturers to validate the tools they use to develop their devices. Generally, manufacturers accomplish this by validating that a finished device performs as intended and by making sure that tool patches are up to date. But the scope of the validation process is limited. As devices become more complex and interconnected, more attention will need to be paid to tools' comprehensiveness and trustworthiness. We have opportunities to

borrow from other industry sectors. For example, the FDA might recognize tools from the aviation domain that have been DO-178B certified or that conform to European MISRA coding and compiler standards.

Moreover, the lack of integration or synergistic interplay between tools and the encompassing development process means that some potential benefits of automated tool support are unrealized. Tools could be enhanced with standardized reporting formats that the FDA recognizes and that could support automated querying and auditing by FDA regulators—greatly increasing regulators' efficiency and accuracy.

Reorienting Procedures Toward Component-Based Certification

The development of component-based software—software that has individually deployable and reusable components—can reduce software development costs. Some vendors are effectively developing large-scale software systems by using *product-line* methods, in which the cost of developing *families* of similar systems is reduced by identifying commonalities across all systems in a family and then developing common components. Current development trends, along with the need to support emerging plug-and-play devices, provide a clear mandate to develop component-based approaches to certification.

Developing Safety-Critical Middleware

Closely tied to component-based certification and the use of product lines is the need for high-assurance and safety-critical middleware. Middleware is system software that resides between applications and the underlying operating systems, network protocol stacks, and hardware. It is often described as the "glue code" or "plumbing" that hooks multiple applications together and routes data and information transparently between different back-end sources of data.

The development of sophisticated plug-and-play devices, devices integrated through networks, and systems of systems as in operating rooms of the future will all require reusable infrastructure code that deals with the complexities of distributed systems. Middleware technology such as Common Object Request Broker Architecture (CORBA) provides a development solution. Unfortunately, most middleware implementations use complex object-oriented design patterns that are difficult to validate. We need safety-critical middleware implementations, which vendors are reluctant to pursue because of an uncertain market. Changes in regulatory guidelines might encourage the production of safety-critical middleware.

Certifying in the Presence of Change

Obtaining certification for upgrades or changes to already "certified" products is costly and error prone. Many failures occur because of the inadequate recertification of modified devices. The state of the art for medical device recertification does not include the use of tools from programming language and software engineering—such as program slicers, dependence analysis, and impact analysis—that could be applied to determine exactly which sections of a software implementation are affected by modifications to code and to requirements.

Researchers should pursue more sophisticated techniques for dependence analysis, which could help determine the impact of changes. Potential examples are techniques that manage certification artifacts and automatically detect where rework in the presence of change is necessary. FDA recognition of tools that perform impact analysis and traceability could reduce the effort needed for retesting and recertification.

Monitoring Deployed Devices

We are able to report and track device "incidents," but reporting, tracking, publicizing, and interpreting the root causes of device incidents across manufacturers is problematic. We need to incorporate some notion of a "flight recorder" black-box reporting mechanism to gather data about events. Some devices, such as the Varian Linac, already incorporate "flight recorder" technology. Research is needed into runtime monitoring approaches that can be used in real-time embedded systems to gather event and fault data from medical devices. Statistical and sampling-based approaches to mining data collected from devices are also important.

Moving from Process-Driven Approval to Evidence-Based Certification

FDA device approval centers on a *process-driven* approach, in which manufacturers obtain approval by showing that they have carried out established quality assurance techniques such as code coverage, manual code inspections, and test cases. The approach often fails to account for innovations in development and verification techniques and does not encourage them. Moreover, as systems become more complex, we expect that best-effort processes will increasingly fail to catch subtle errors.

Research is needed into the alternate paradigm of an *evidence-based* approach, championed by John Rushby in his HCMDSS position paper, "Goal-Based Certification for Medical Devices." Instead of seeking to provide evidence that a process has been planned for and followed, the evidence-based approach seeks to generate independently verifiable evidence that a system satisfies its requirements. An example of a certification process based on an evidence-based or goal-based approach is the UK Defense Standard 00-56.

Formalizing Environment Models and Assumptions About Context

Device manufacturers often work under the assumption that devices will be used in certain fixed, step-by-step processes and will have no interaction with other devices. But when devices are used in environments that differ from those envisioned by the manufacturers, the devices often behave in unanticipated ways. An example is patients who have characteristics that deviate significantly from the norm (such as severely obese patients).

The state of the art relies on resolving undocumented and unanticipated interactions on the fly and in an ad hoc manner. Specifically, constraints on interactions are not captured in the validation process. As the number of devices in a clinical environment increases, and as the sophistication of devices and their connectivity methods (such as wireless Internet and Bluetooth radio) increases, clinical technicians will become overwhelmed with the task of assembling a safe and effective environment. Research is needed on formalizing models of clinic

environments, clinical processes, and assumptions made about the contexts in which devices will be deployed.

Dealing Effectively with the Certification of Security

Security-related research is needed in the following areas:

- Specification formalisms for specifying security properties at the level of system requirements, designs, and implementations
- Tools and techniques, such as those based on static analysis and lightweight theorem-proving, that can help verify that implementations satisfy specifications for security
- Regulatory guidelines, architecture recommendations, and techniques for security certification (Common Criteria and software architectures, for instance), to achieve standards such as the Multiple Independent Levels of Security (MILS) architecture

Research Strategies and Roadmap

It is likely that individual research teams can make progress in addressing the R&D challenges described above. Yet overall research progress and an exchange of ideas about research approaches, domain knowledge, and the effectiveness of competing techniques will be significantly hampered unless we can create significant infrastructure to support research on high-confidence medical devices.

Specifically, we need an infrastructure that facilitates interaction between device vendors and academic researchers. Researchers will be better able to address certification challenges facing V&V technologies if they have a detailed understanding of the certification process. Currently, researchers understand little about the certification process, and it makes little sense for them to engage in research about tools and techniques to aid certification or to propose changes in certification techniques without such an understanding.

These are two ideas to help researchers obtain a deeper understanding of the certification process:

- FDA and industry personnel hold training workshops on the review and approval process using example artifacts.
- FDA or industry performs a mock “red team” review process for products and artifacts developed by researchers. For example, the FDA or industry personnel could review medical device software built by academic researchers to demonstrate their V&V approaches. The review would use the same processes and standards that are applied to actual device submissions to the FDA.

The most important advances will come from the establishment of a research infrastructure, such as one or more open experimental platforms (OEP) for medical devices. An OEP is a publicly

available test bed. In this instance, the OEP would provide realistic and complete examples of the following artifacts for a specific medical device:

- Requirements (including hazard analysis)
- Implementation or partial implementation
- Results of fault analysis, test cases, and other V&V activities
- Examples of material that would be submitted to the FDA for device approval

The OEP should also include a list of challenge problems, written by FDA and industry personnel for academic researchers, that indicate specific areas on which researchers should focus.

Strong consideration should be given to having national funding agencies contract with device vendors to—

- Develop OEP artifacts and challenge problems as described above
- Field questions from academic researchers
- Help evaluate tools and techniques that academic researchers have produced for realistic products and in realistic development settings

Five-Year Roadmap

The five-year roadmap involves the further development and maturation of tools that can provide automation for V&V and certification activities, including an initial exploration of V&V, the development of V&V techniques, and a demonstration of their usefulness. It would include the following steps:

- Develop an OEP for medical devices
- Apply existing process modeling languages to model clinical environments and processes
- Develop suites of sophisticated requirements for capturing, simulating, and querying requirements
- Incorporate more tools into the certification effort (adding value)
- Certify development tools (analysis and traceability tools) to reduce the burden of certification and recertification
- Enhance conventional formal method tools (static analysis and model-checking) to produce a variety of artifacts, including test cases and natural-language description of traceability steps

- Demonstrate the pervasive use of model-based development techniques with automated reasoning for—
    1. Component conformance to interfaces
    2. Component capability based on checking interface capability
    3. End-to-end reasoning of system behavior
    4. Managing, integrating, and automatically generating certification artifacts

Ten-Year Roadmap

The ten-year roadmap focuses on refining tools for use by manufacturers and regulatory agencies, including component-based certification, reuse of certification assets, and reuse of "precertified" high-assurance middleware and other infrastructure:

- Reorient the certification process toward component-based certification
- Develop certified components as commodities
- Arrive at the pervasive use of secure, QoS-aware, fault-tolerant, certified middleware
- Achieve integrated, end-to-end, model-based development frameworks dealing with composition, evolution, and change
- Effectively demonstrate the metrics and other items necessary to change industry and regulatory practices
- Establish a wide body of interoperability standards
- Build more sophisticated tools for compositional analysis
- Design tools for compositional hazard analysis