2014 National Workshop on Research Frontiers in Medical Cyber-Physical Systems

NATIONAL SCIENCE FOUNDATION

February 6th and 7th, 2014
Waterview Conference Center
1919 N. Lynn St., Arlington, Virginia

Final Report on the 2014 CPS Medical Devices Workshop

Scott A. Smolka[1], Susan Trolier-McKinstry[2], and Mani Srivastava[3]

[1] Department of Computer Science, Stony Brook University,
sas@cs.stonybrook.edu
[2] Department of Materials Science and Engineering, Pennsylvania State University,
trolier-mckinstry@matse.psu.edu
[3] Department of Electrical Engineering, University of California at Los Angeles
mbs@ee.ucla.edu

February 6^{th} and 7^{th}, 2014 at the Waterview Conference Center, 1919 N. Lynn St., Arlington, Virgina.

1 Table of Contents

White Paper 2.1: *Multifunction Delegates for Secure Medical Cyber-Physical Systems* - Richard Schantz (Raytheon BBN Technologies)

White Paper 2.2: *Zero-Defect Methods for Surgical Robots* - Yanni Kouskoulas (Johns Hopkins University/APL)

White Paper 2.3: *From Verified Models to Verified Code for Medical Devices* - Miroslav Pajic (University of Pennsylvania)

Breakout Session 2.1: *Medical Device Errors and Misinteractions* - Lead: Insup Lee (University of Pennsylvania)/ Co-Lead: Paul Jones (FDA)

Breakout Session 2.2: *Securing Devices & Data Against Cyber-Physical Vulnerabilities* - Lead: David Kotz (Dartmouth) / Co-Lead: Richard Schantz (Raytheon BBN Technologies) / Co-Lead: Shashank Priya (NSF)

Breakout Session 2.3: *Coping with Alarm Fatigue and Cognitive Overload* - Lead: Julian Goldman (Massachusetts General Hospital / Harvard Medical School) / Co-Lead: J. Perren Cobb (Massachusetts General Hospital) / Co-Lead: Sylvia Spengler (NSF)

Breakout Session 2.4: *Modeling, Simulation, & Verification to Predict Performance & Reliability* - Lead: John Hatcliff (Kansas State) / Co-Lead: Miroslav Pajic (University of Pennsylvania) / Co-Lead: Anindya Banerjee (NSF)

Intervention, Control and Prosthetics

Plenary Talk 3.1: *Moving Beyond Current Approaches to a Broader View of Epilepsy Monitoring* - Steven Schachter (Center for Integration of Medicine & Innovative Technology)

Plenary Talk 3.2: *Implantable Networks of Wireless Nanoelectronic Devices* - Pedro Irazoqui (Center for Implantable Devices, Purdue University)

White Paper 3.1: *Next Generation Soft Wearable Robots to Assist Impaired Patients* - Conor Walsh (Harvard University)

White Paper 3.2: *Smart Architectures for Implantable Wireless Pulse Generators for Gastric Disorder Management: A CPS Approach* - Aydin Farajidavar (New York Institute of Technology)

White Paper 3.3: *Open Platform for High-Acuity Surgical Treatment with Cyber-Physical Systems* - Hawkeye King & Blake Hannaford (University of Washington)

Breakout Session 3.1: *Improving the Man-Machine Interface* - Lead: Conor Walsh (Harvard) / Co-Lead: Art Erdman (University of Minnesota) / Sylvia Spengler (NSF)

Breakout Session 3.2: *Just-in-Time Health Interventions* - Lead: Jim Rehg (Georgia Institute of Technology) / Co-Lead: Wendy Nilsen (NIH)

Breakout Session 3.3: *Medical Robotics* - Lead: Peter Kazanzides (Johns Hopkins University) / Co-Lead: Hawkeye King (University of Washington) / Co-Lead: Tho Nguyen (NSF)

Breakout Session 3.4: *Neural Prosthesis and Interfaces* - Lead: Pedro Irazoqui (Purdue) / Co-Lead: Aydin Farajidavar (New York Institute of Technology) / Co-Lead: Shashank Priya (NSF)

Dinner Keynote Talk

Conclusions and Future Directions

2 Introduction

This report provides a summary of the 2014 CPS Workshop on Medical Frontiers in Research, which took place February 6-7, 2014 at the Waterview Conference Center in Arlington, VA. The workshop brought together multi-disciplinary experts from industry, academia, and government to discuss future CPS research challenges in this space. The participants were experts in the foundational science, engineering, and technology underlying next-generation cyber-physical systems, particularly in the context of the application domain of medicine, including medical devices. Through sessions on CPS research in sensing, diagnosis, and prosthetics, the attendees are defining a compelling research agenda in an area of scientific, technical, and engineering exploration with vast societal impact.

The workshop featured three tracks related to large challenges in cyber-physical systems. Each track had a distinct session comprised of invited plenary talks, short presentations of highlighted white papers, and a number of breakout sessions for group collaboration. These tracks are discussed in detail, each in its own section of the report.

The rest of the document is outlined as follows. Section 4 covers monitoring and diagnosis. Section 5 provides an overview of topics related to verification, modeling, and the underlying principle of "trustworthiness". Section 6 focuses on what can be done in terms of medical intervention and control, including the use of prosthetics. Section 7 covers the dinner keynote talk given at the end of day one of the workshop. Section 8 offers some concluding remarks and future research directions.

3 Workshop Agenda

The workshop was scheduled across two days, and featured a number of invited talks, accepted white papers, and group breakout sessions. The agenda was divided into three tracks, each covering important topics in Medical CPS. Track 2, on modeling, verification and trustworthiness, bridged the two days of the event.

The first session of the workshop ran from 09:25 on February 6^{th} through 12:15 that same day. This morning session featured two invited plenary talks, three presentations on highlighted white papers, and four breakout sessions. All talks and breakout sessions were focused on health and environmental monitoring as well as clinical diagnosis.

The second session rounded out the rest of the first day and springboarded the discussion of the second. This section ran from 13:15 on February 6^{th} through 17:30, and continued for an additional hour on the 7^{th}. Session 2 featured an additional plenary talk, and all discussion centered around trustworthiness as a concept, building models, and verifying properties of models and code.

The final session focused on how we can apply the information gathered with the ideas presented in the previous two sessions in order to correct conditions in patients. Session three ran from 9:45 on the 7^{th} of February through 15:15 that afternoon. Like the first session, the final featured two invited plenary talks.

4 Monitoring and Diagnosis

This session focused on problems related to data collection and usage in medical CPS. Monitoring is a two-fold problem in this field: lots of precise information is needed, but the sheer amount of available data can often "drown out" what is needed in the given moment. We must use CPS solutions both to improve monitoring, and use discovered data in an intelligent manner to help with diagnosis.

4.1 Plenary Talks

PT1.1 Wearable Sensor Systems for Long Term Health and Environmental Monitoring Of the two talks in this session, the first was given by Veena Misra of Nanomachines ERC, ASSIST and NCSU. The talk was entitled *Wearable Sensor Systems for Long Term Health and Environmental Monitoring.*

Veena's work at the ASSIST (Advanced Self-powered Systems of Integrated Sensors and Technologies) center focuses on the use of nanotechnology to build a new class of sensors. These devices push the boundaries of modern wearable technologies by cutting power costs and traditional rigidity while reaching for innovative new solutions in energy harvesting. ASSIST's sensors allow a patient to remain comfortable while providing a large amount of health and environmental sensing data.

The lab has addressed energy concerns via a two-pronged attack plan. First, they aimed to address the actual consumption of the devices in question. The electronics in their sensors are all designed to be "ultra-low power." This ties well into the second prong: energy for the device is harvested from the patient's own body heat and body motion. The devices do not need to be routinely removed for charging- a problem that has plagued the existing generation of wearables.

Finally, Veena and ASSIST have created a pleathora of new, flexible physiological/environmental sensors. These include, but are not limited to, nanowires for gas and overal health sensing, flexible electrodes, and ultrasonic transducers for sensing volatile organic compounds.

PT1.2: How a Wealth of Information Begets a Poverty of Attention: The Special Case of Data Overload in the ICU The second plenary talk on monitoring and diagnosis was entitled *How a Wealth of Information Begets a Poverty of Attention: The Special Case of Data Overload in the ICU* This talk, by J. Perren Cobb, M.D., Director of the Massachusetts General Hospital Critical Care Center, focused on arising problems in the ICU from an over-abundance of data.

Traditionally, it is thought that an increase in available data is always desirable, but in many cases this is not strictly so. Those who have important decisions to make may find themselves victims of so-called "analysis paralysis" - there is too much information available to take decisive action. Unfortunately, this has become the case in the modern Intensive Care Unit. One of the largest challenges in critical care informatics is that often high risk situations are further complicated by the sheer amount of patients, resources, and data all available.

To illustrate this, consider a list of the top 10 technology hazards in an Intensive Care Unit in 2013:

1. Alarm hazards;
2. Medication administration errors using infusion pumps;
3. Unnecessary exposures and radiation burns from diagnostic radiology procedures;
4. Patient/data mismatches in EHRs and other health IT systems;
5. Interoperability failures with medical devices and health IT systems;
6. Air embolism hazards;
7. Inattention to the needs of pediatric patients when using "adult" technologies;
8. Inadequate reprocessing of endoscopic devices and surgical instruments;
9. Caregiver distractions from smartphones and other mobile devices;
10. Surgical fires.

Of this list, only items 8 and 10 are unrelated to data overload. With such an abundance of technological problems, we need to develop strategies to identify and address knowledge and capability gaps.

What is necessary to address these issues is a transformative approach- crossing multiple disciplines, universities, hospitals and organizations. This need has led to the formation of the United States Critical Illness and Injury Trials Group (USCIITG). In 2014 it was comprised of over 200 investigators across 68 ICUs, and focused on four major programs: Prevention of Organ Failure (PROOF); Critical Illness Outcomes Study (CIOS); Early ICU Rehabilitation (PEIR); and Program for Emergency Preparedness (PREP). Overall funding over the previous four years totaled approximately $22 million, and was provided by NIH, CMS, ASPR, and DoD.

The final portion of the talk highlighted two important use cases- public health emergencies and "black boxes" for crashes.

To quote an article in the New England Journal of Medicine entitled *Research as a Part of Public Health Emergency Response*, "Although responses to recent events have typically used the best available science at the time, additional research, done in parallel with and after the response itself, is often essential to address the most pressing knowledge gaps presented by public health emergencies and to ensure that they are addressed by the time another similar disaster strikes. Recent events have also illustrated gaps in planning for, and rapidly executing, scientific research in the context of disaster response."

The last sentence deserves increased emphasis. The "recent event" in the article was the Boston marathon bombing, and it help illuminate the five stages of care: pre-hospital, emergency room, the floor of the operating room, the ICU, and finally rehabilitation. Data collected at each step must help inform those that follow, not be lost or become noise. In hopes to eliminate potential problems, USCIITG proposed the Program for Emergency Preparedness (USCIITG-PREP). PREP is a rapid analysis and data dissemination plan, connecting points-of-care with University, Government and Hospital partners.

The second use case, the "ICU black box," is an attempt to look at health engineering in the year 2032. Which current problems will remain unsolved?

USCIITG suggests that the issues of determining which data to collect and monitor, how to predict and model medical problems, and determining automated responses will all remain relevant. On top of that, new problems are likely to develop. However, despite the listed problems, it is likely that "disruptive technology" will help drive solutions in the future.

4.2 Highlighted White Papers

The three white papers in this session were each presented in five minute talks, and each described CPS solutions for existing medical problems.

WP1.1: Need for Technical Solutions Underlying Wireless and Mobile Health Systems The first talk was entitled *Need for Technical Solutions Underlying Wireless and Mobile Health Systems*, and was presented by John Stankovic of the University of Virginia. John focused on the characteristics and underlying research in in-home architecture as well as the connections and possible rifts between the physical and cyber realms.

In-home architecture requires a number of pieces to fall into place. Any such architecture needs to have flexible sensing and actuation to provide utility. It must be heterogeneous, allowing for plug and play applications. The architecture needs to be easy to deploy and maintain, evolve as necessary, and be cloud supported. The research underlying this architecture includes realities for ADLs, anomaly detection, speaker identification, mood detection, reverberation, and dependencies, among others.

While the goal of CPS can be seen as connecting the physical world and cyberspace, there are a number of complications to this process. The physical world is highly complex, with the unpredictable nature of humans, over-simplified assumptions, and new faults leading to blocked or moved sensors. Intervening here requires a comprehensive understanding of all types of human-in-the-loop controls, safety critical systems and properties, differences in system ID versus Machine Learning and the duality of FC and ML.

WP1.2: The Evolution of Software-Programmable Cyber-Physical Digital Microfluidic Laboratories-on-a-Chip The next white paper was presented by Philip Brisk of U.C. Riverside, on *The Evolution of Software-Programmable Cyber-Physical Digital Microfluidic Laboratories-on-a-Chip*. Prior to CPS, digital microfluidic technology consisted of placing a droplet between two electrode-containing plates and using the various control electrodes to split, mix, transport and store the droplet. This technology was somewhat programmable, took up very little space and was quick to run, but there was very little flexibility in control-flow and the computational overhead was immense. Integrating a cyber-physical component to the laboratory has greatly improved the latter three properties with only minor incursions into the former two.

In the cyber-physical scenario (Fig. 1), a PC controller sends data to a microcontroller driver. The driver then exchanges information with the cyber-physical DMFB, until the results are sent from the DMFB back to the PC controller.

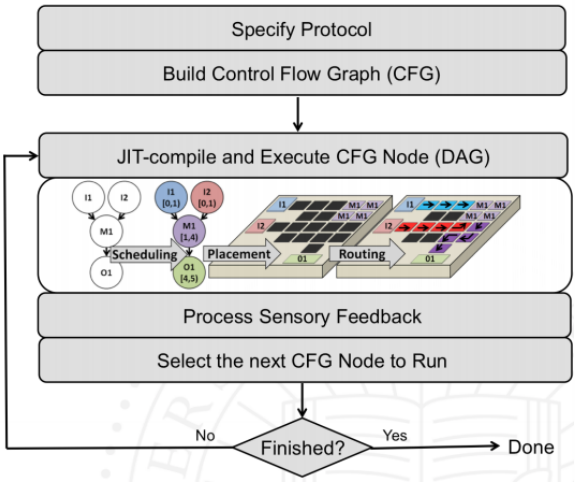

Fig. 1. Cyber-physical Integration for Digital Microfluidic Laboratories.

Cyber-physical integration has so far led to real-time decision-making based on sensory feedback, programming language support, and a fast JIT compilation flow. In the future, it is anticipated that this paradigm will lead to better domain-specific programming languages for automated biology. This will require a syntax deemed acceptable by biologists, express timing of operations and constraints, and must automatically extract and exploit parallelism. Additionally, steps can be taken to automate loading and unloading of samples, including robotic liquid handling and off-chip fluid storage. Finally, Philip expects that future efforts will be in allowing for distributed control over the Internet. This will facilitate multi-site collaboration and ensure timing constraints.

WP1.3: Cyber Physical System Models for Just-in-Time Care Delivery with Mobile Health Sensors Santosh Kumar of the University of Memphis presented the final white paper of this session, a work entitled *Cyber Physical System Models for Just-in-Time Care Delivery with Mobile Health Sensors*. Through the application of computational models, sensor data such as an ECG, respiratory sensor and accelerometers can help determine biomarkers for patient conditions. Those determined in literature so far include markers for stress, smoking, cocaine use, walking and talking. This idea suggests a path towards optimization of care delivery through mHealth sensing. This could potentially mean the realization of P4 and precision medicine tailored for the patient's condition, or help discover predictors of adverse health events. Such sensing could help to find patterns in multivariate biomarker sequences related to such events, along with their specific and sensitive antecedents and precipitants. By identifying predictors in real-time measurement of health, behavior, context, and environment, as well as capturing response to interventions, we could deliver personalized, just-in-time and precision treatments.

What are the CPS foundations for JITAI? Well, we can use mHealth sensors to optimize the timing of JITAI delivery and adapt its content, but in order to systematically analyze their impact on efficacy and safety, we need to establish theoretical foundations for JITAI. These foundations include formal models of health states, behaviors, environment, and interactions among them. This includes predicting any response to potential actions. To top it off, it would require models for human-in-the-loop CPS. In order for this to really be considered a success, JITAI must be eventually shown to work for every individual.

4.3 Breakout Sessions

Each track would conclude with a total of four breakout sessions. For this topic in particular, the sessions were centered upon sensing techniques and data derived from sensing.

BR1.1: Continual Sensing of Health Markers The first breakout session was led by Veena Misra with co-leads Susan Trolier-McKinstry of Pennsylvania state and Tho Nguyen of the NSF. This session closely corresponded with Veena's earlier plenary talk. The focus was upon techniques and challenges for *Continual Sensing of Health Markers*. Any such sensors need long lifetimes, flexibility for comfort, durability to avoid environmental wear and tear, and convenience for the user (no frequent charging stops or other rituals). Wearable sensors offer insights to chronic conditions, health patterns and disease progression over time, and precise data for sudden attacks or trauma.

BR1.2: "Missing" Sensing Modalities Breakout two switched our focus from wearable sensors to *"Missing" Sensing Modalities*. This session, led by Santosh Kumar and co-led by Steven Schacter of the Center for Integration of Medicine & Innovative Technology and Phillip Regalia of the NSF, was an attempt to explore the missing components of modern medical sensing. Santosh's white paper posed the question of how to analyze the efficacy and safety of just-in-time medicine delivery via sensor systems. This breakout bred a discussion focused on modeling the patient and environment with respect to sensors, interactions between sensors, cyber-physical controllers, etc. It also addressed the problem of technology reaching obsolescence before ever being properly tested and analyzed.

BR1.3: Processing and Communicating Sensory Information The third session was led by John Stankovic, who gave the first white paper talk, and co-led by Deepak Ganesan and Thyaga Nadagopal, of the University of Massachusetts and NSF, respectively. The topic was *Processing and Communicating Sensory Information*. Processing can be done at the SP level or at higher levels. The higher-level processing would include machine learning, inference, and decision making under uncertainty.

The research questions addressed were as follows: SP in real environments including data scrubbing; behavioral feedback (long time frames) versus actuation

((multiple dimensions) interventions); Where is data best processed, under complex collections of constraints (energy, real-time, etc.); JIT verification for the "cockpit of the ICU" (processing and communication, alarms, ... dynamic collection of devices); Interoperability/Composition including options for higher level semantics and hybrid systems; making processing context dependent; determining new processing capabilities; handling uncertainty at every level; dependencies and assumptions when creating systems of systems; and challenges pertaining to modeling.

Communication protocols were discussed at both wireless and human levels. Conversations were had about wireless signals co-existing within bands; reliability as a metric, relating to both data delivery and real-time transactions; system scaling; data overload in terms of what data to collect and show; human factors with respect to patients, caregivers, doctors, etc.; and questions of amount of autonomy vs user-in-the-loop properties of this technology. Security, authentication, maintainability, and data provenance were all also left on the table.

BR1.4: Inferences from Complex, Incomplete, and High-Rate Data The final breakout session of this track, entitled *Inferences from Complex, Incomplete, and High-Rate Data*, was led by Greg Pottie of UCLA. His co leads were Philip Brisk, who gave an earlier white paper talk, and Vinay Pai of the NIH. Any student of statistics would know the phrase "garbage in, garbage out" when it comes to data analysis. Unfortunately, when it comes to medical data, we cannot afford to be so nonchalant. Data from a broken or damaged sensor may be our only insights into a traumatic medical event, or incomplete medical history information all that we have for an individual patient. How can we utilize what data we have, in any condition, in order to give necessary treatment while avoiding false-positives that may lead to unnecessary medication or surgery? Even in the case of complete data, how do we appropriately handle medical data with high complexity (nonlinear physiological behavior, numerous medication interactions, etc.)? We must develop strategies to use what information we have to minimize risk to the patient, either through intervention or lack thereof.

5 Modeling, Verification and Trustworthiness

Modeling and verification can be very powerful tools in any field, but they are highly dependent upon assumptions by the developers and users. George Box is oft quoted for the phrase, "all models are wrong, but some are useful." While the universality of his claim can certainly be questioned, the quote itself rings true with regards to unrealistic simplification. The focus of this section is on how do we define trustworthiness in our models and systems, and how we might build models that realistic enough that their verification results are useful.

5.1 Plenary Talks

As mentioned above, we were fortunate to have three plenary talks to lead the second session.

PT2.1: Trustworthy Health and Wellness (THAW) The first of these talks, *Trustworthy Health and Wellness (THAW)*, was given by David Kotz of the Institute for Security, Technology, and Society at Dartmouth College. Healthcare costs 18% GDP; while this is currently the case, it is thought that I.T. can still improve quality while decreasing overall cost. Both mobile and cloud solutions have been deployed for health and wellness problems, but security and privacy is absolutely essential for trust of patients and clinicians to be maintained.

Here is a list of five trends and their implications for security/privacy:

1. *The shifting locus of care.* Instead of traditional offices and hospitals, health care is increasingly delivered through small clinics, elder-care centers, or at the patient's home. This means patients and clinicians must often move among their associated hospital, clinics, and the patient's residence. As such, we have seen an increase in distributed and remote monitoring supported by cloud services.

2. *Accountable care and patient engagement.* We have seen the introduction of Accountable Care Organizations (ACOs). These ACOs need to collect metrics about health of their patient population. On a positive note, ACOs motivate patients to engage and remain healthy.

3. *Continuous patient monitoring outside the clinical setting.* This monitoring has been established for post-discharge of patients returning home, for monitoring and managing chronic conditions and for encouraging healthy behavior/assisting with behavioral health.

4.*Advent of mobile devices and cloud services in health-related applications.* There has been a trend towards smartphones and wearables that measure physiology, activity, and environment. We have also seen the emergence of cloud services to support small and distributed healthcare organizations.

5.*Emerging threats and changing regulatory environment.* The sudden uptick in use of EHRs leads to an increased risk of large-scale privacy breaches. There has been an increasing occurrence of Medical identity theft, and mobile devices and cloud services are increasingly under attack. Additionally, the FDA is planning to regulate some mHealth devices and apps as medical devices, completely altering the path to market.

The purpose of TH&W, Trustworthy Health and Wellness, is to take advantage of these innovations while addressing the growing issues. TH&W focuses on the advent of mobile and cloud tech, aiming to address issues of authentication, privacy, trustworthy control and accountability. The issues are critical in healthcare domain, but the solutions are applicable to other domains. The project suggests nine *initial* projects on three themes. 1. Usable authentication and privacy tools: clinician friendly authentication; selective reporting of mHealth telemetry; break-the-glass with segmented records; and genomic personal health records. 2.Trustworthy control of medical devices: securing small health networks; and securing remote directives. 3. Trust through accountability: malware detection using power analysis; trust in mHealth data; and audit models and malware detection.

PT2.2: Patient Safety Challenges for Medical Device Innovation The second plenary talk was given by Julian Goldman of Massachusetts General Hospital

and Harvard Medical School. The talk addresses *Patient Safety Challenges for Medical Device Innovation*. When medical device-HIT "system related" adverse events occur, it is often difficult or impossible to find the root cause of the failure. The current medical device-HIT system does not lend itself to complete data logging for later analysis (it does not have black box recording). There is no national reporting pathway for gaps, needs, or failures of heterogeneous medical device (and HIT) systems

Patient-Controlled Analgesia (PCA) *system* safety concerns: Patients can call to request more analgesia, but cannot call for help when over-medicated (Fig. 2). Over-medication can cause respiratory and cardiac arrest. Comprehensive monitoring is *not* typically used due to high false/nuisance alarm rate

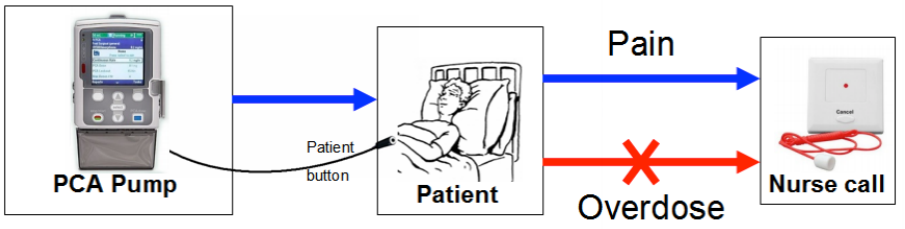

Fig. 2. Patient-Controlled Analgesia (PCA) System Safety Concerns.

According to the Official Journal of the Anesthesia Patient Safety Foundation, with regards to the dangers of postoperative opioids, "A particularly attractive feature may be the ability to automatically terminate or reduce PCA (or PCEA) infusions when monitoring technology suggests the presence of opioid-induced respiratory depression. To facilitate such capabilities, we strongly endorse the efforts to develop international standards for device interoperability and device-device communication." The proposed "smart PCA system" requires the use of physiologic data for earlier detection of respiratory depression and capability to stop medication infusion. But how would you create a smarter PCA system? Sensors: probably SpO_2 and respiratory rate (RR) but there are many potential RR sources. Capnography? ECG (impedance)? Other (acoustic, TBI, thermistry, Vt, etc.)? $ETCO_2$ - can be used for hypercarbia, RR, etc. Which algorithms would be used? How would you used the signals? Threshold? Trends? Would you create a warning alarm or stop the PCA? Do you have enough data to develop the algorithms? What about to personalize sensors and algorithms for different patient populations? Should it support the plug-and-play of devices as needed? Is there other clinical and contextual information needed, like the patient's supplemental O_2 status or possible co-morbidities like sleep apnea? All of these questions would need to be answered.

The talk then went on to discuss the development of a prototype healthcare intranet to improve health outcomes. Such an intranet would be envisioned as an eco-system for interoperability of a medical device and clinical information sys-

tems to support innovation in patient safety and healthcare quality. This would involve four quantum "clinical scenarios." 1. A PCA safety interlock: an example of component-level medical device interoperability to improve safety of medication infusions. This would also encompass multi-parameter "smart alarms." 2. ICU preparedness: an example of the ability to support safer in-hospital patient transfer and dynamic checklists to reduce errors. 3. Tele-health devices within the hospital setting: an example of transferring care from home to hospital and the use of devices for high-acuity care. 4.An FDA regulatory and safety scenario: using sedation for a G.I. procedure as a framework for levels of interoperability and associated levels of hazards and their mitigation.

What could be accomplished in healthcare with open, interoperable medical device app platforms? In terms of innovation, we could see rich contextual data for clinical decision support; means for rapid prototyping of new sensors and algorithms; techniques to facilitate validation for regulatory clearance; and the ability to swap devices as needed to optimize selection. Looking to what is available in other domains, we find crowd-sourcing of "apps". If our device platform is standardized, apps can be developed by the global expert community.

In terms of testbeds, the talk examined how to connect the nonclinical with the clinical. The non-clinical lab medical CPS ICE testbed sees resource development via devices, tools, simulators, and SMEs. Tests occur at the component and system levels, with data and workflow simulation. Assessments are made in terms of compatibility with other devices and networks as well as security. Initial applications on the non-clinical testbed include the MD PnP lab "Virtual Hospital" for Smart America CPS Testbed Challenge as well as device and system verification prior to MGH deployment.

The clinical testbed is used in order to study new sensors and actuators, pharmaceuticals, care pathways, decision support tools and clinical studies that require the testbed. Additionally, it is used to validate apps and identify gaps in existing tools and technologies, and to introduce CS&E to the clinical environment. The initial clinical applications in the MGH pilot include neuro data fusion, smart alarms and closed-loop medication administration.

When connecting the textbeds we see three examples of CPS themes: 1. time synchronized clinical and device data recording; 2. data fusion and patient context (state) including app interactions; 3. command and control in system-of-systems.

Next the OpenICE (Integrated Clinical Environment) platform is introduced to close the loop on safety. OpenICE encompasses medical device interfaces, patient simulation, HIS connectivity, safety assurance, validation and testing, and regulatory pathways. It also includes the ICE Application exchange (ICE AX) for medical apps and device management. The entire system is open source. While it is only in a beta stage of development it currently has useful pieces for clinical and CS research.

Finally, the talk attempts to address a gap in the alignment of healthcare national needs. The recommendation given is to align national patient safety interests with the use of clinical technology. There is a need for a national ap-

proach to evolve the safety and capabilities of healthcare system technologies. There must be centralized reporting, analysis, recommendations, and shared solutions. This includes both regulatory enforcement and market incentives. A Health IT Safety Administration or Board (HITSA) modeled on other national reporting initiatives (NHTSA, ASRS, MedSun, NTSB, ASTERD, PSO, etc.) could oversee adverse event reporting, include FDA regulated and non-regulated (IT) devices, and be multi-stakeholder (regulators, clinical representatives, manufacturers, etc.). Additionally, it recommends a clinical scenario repository. A clinical scenario is defined as a brief description of a clinical situation or event that could be improved through better system solutions. The purpose is to inform technical solutions. A clinical scenario repository would be a web portal to allow clinicians, clinical engineers and other users to enter, revise and annotate clinical scenarios. A place to document and share these scenarios will help to *identify clinical and technical challenges*, address healthcare needs and guide improvements in patient safety and quality of healthcare delivery.

[This section needs lots of figures. The slides were almost exclusively figures.]

PT2.3: Challenges and Opportunities in Controlling the Human Heart The final plenary speaker in this session was Radu Grosu of the Technical University of Vienna and Stony Brook University. His talk was entitled *Challenges and Opportunities in Controlling the Human Heart*. The heart is an electro-mechanical pump and can perform as an error-free (normal operation) system, or an error-prone system (arrhythmia, tachycardia, etc.). But whose problem is it to solve? It is clearly a medical problem. The number one cause of death in the United States in 2006 was heart disease. It lead to 725,192 total deaths, or 30.3% of the total deaths.

The fundamental questions for cardiologists, pharmacologists and patients are two-fold: What is the risk of a patient to develop the disorder? Under what circumstances will such a disorder arise? But when given a disorder-specification and a model of the ventricle the important questions become the following: What is the probability of the model to satisfy the specification? For what parameter-ranges does it satisfy the specification? The question of whose problem it is to solve becomes unclear.

The CMACS project is a multi-institutional, multi-disciplinary team spanning 7 universities, 2 colleges, and 14 departments/schools. Among this group is a large CMACS Atrial-Fibrillation team, focusing on these heart issues.

It is a communication-structure problem: 4 billion nodes interconnected in a very sophisticated way (See Fig. 3)!

It's a cellular problem (Fig. 4). Many of the control structures related to arrhythmia are found at the cellular level. Each cell produces an action potential as ion channels open and close.

But if we are talking about action potentials and ion flows, that clearly indicates that we are dealing with an electrical problem! The rate of change in membrane potential (V) is the sum of physiological currents due to ion flows across the membrane. The physiological model has 67 variables, making it difficult to simulate and formally analyze.

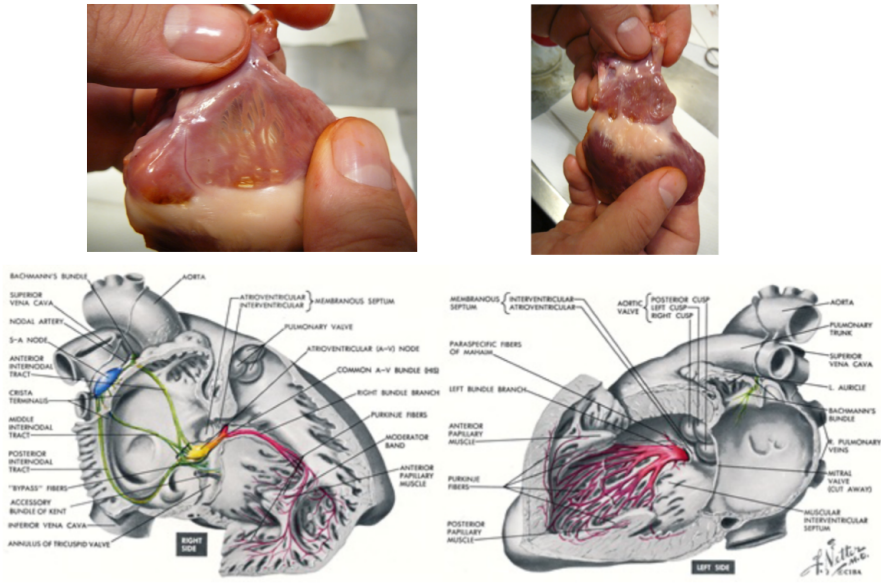

Fig. 3. Topologically Complex Cardiac Structure.

If the complete physiological structure is too complex to model and analyze, then we must deal with abstractions. Therefore heart disease is a cellular-abstraction problem. Though it could also be considered a molecular problem, since the major players are various molecules in solution. So it isn't just the cells being abstracted, but molecules. Clearly, it must be a molecular abstraction problem (Fig. 5).

But in the end, isn't the whole issue that the complexity of the system, including many variables and non-linear dynamics, makes it difficult to simulate or verify? So it must be either a simulation (Fig. 6) or verification problem (depicted in Fig. 7).

It is also important to consider what the results of verification can produce in a non-academic sense. How can these results be put to work in order to help patients? Well, for both atrial and ventricular fibrillation we have found successful defibrillation with LEAP using about 10% of the energy required by the standard 1 shock defibrillation protocol. Furthermore, using high resolution mCT we obtained detailed vessel distribution of the heart and found a scaling law which was used to obtain a theory that explains the mechanism behind LEAP. It's a control problem (Fig. 8).

It's a CPS problem. We are on the brink of a paradigm shift in the diagnosis and treatment of cardiac disorders. It is up to us to make it happen.

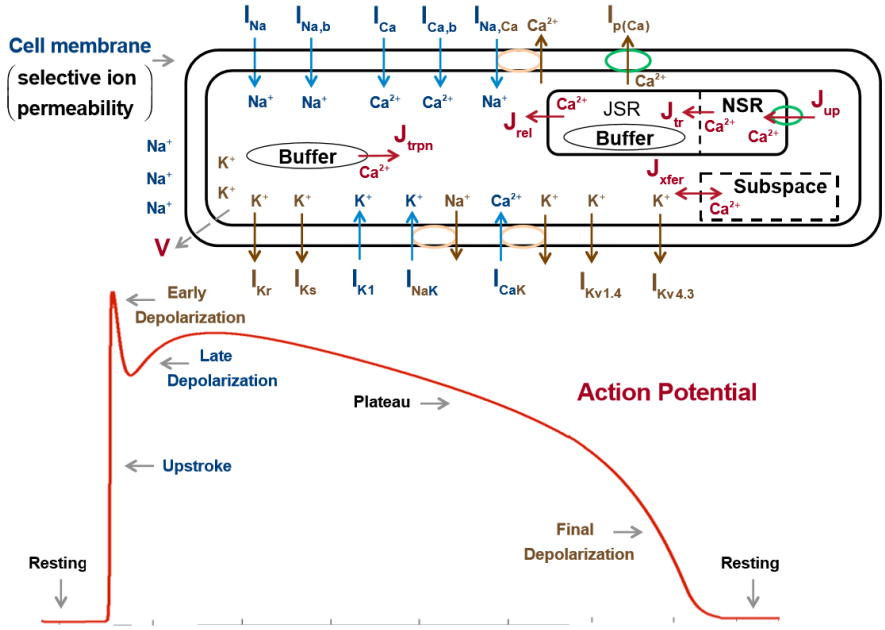

Fig. 4. Ion Channels in Cardiac Myocyte and Corresponding Action Potential.

5.2 Highlighted White Papers

WP2.1: Multifunction Delegates for Secure Medical Cyber-Physical Systems The first white paper of this session was *Multifunction Delegates for Secure Medical Cyber-Physical Systems* by Richard Schantz of Raytheon BBN Technologies. Operating as part of an interconnected eco-system, Multifunction Delegates (MDs) mediate between the embedded devices and information systems for both ingress and egress. Device specific complexity (and security) weaknesses contained below the mediation layer, making the CPS simpler, more interoperable, secure, and privacy preserving. Placed near the embedded devices, MDs mediate all possible interaction between the devices and the larger CPS in either dedicated or multiplexed manner.

The outbound value add of the MD mediation includes: publishing subscribe instead of point to point; giving a uniform time stamp and message format; publishing with 1-1 encryption; publishing with 1-many (attribute and predicate based) encryption; or coordination, sequencing and adaptation. The inbound value add of the MD mediation, on the other hand, focuses on input validation to defend against injection attacks, enforcing appropriate access control and authentication without modifying the end devices, and transformation. Some enablers for transposing and extending are value added infrastructure, real-time cyber-security mediation, interaction privacy, and fully extended privacy in the form of fully homomorphic encryption, and metadata protection.

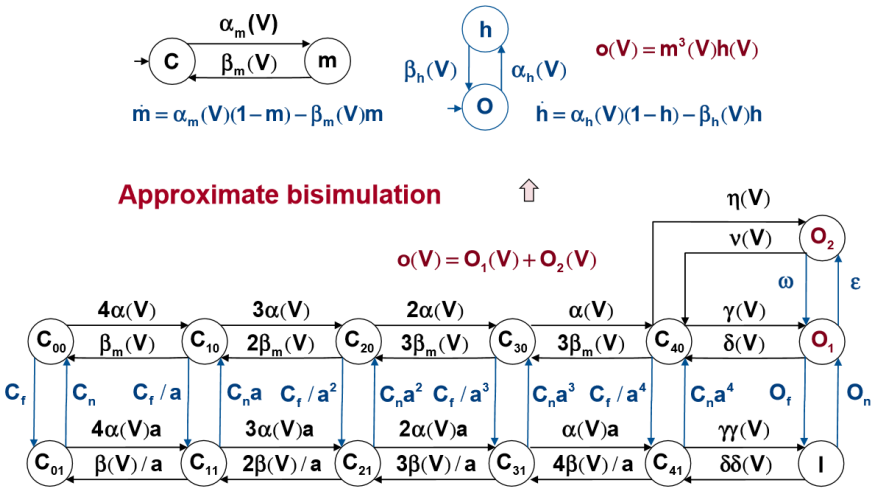

Fig. 5. Markov Chain Displaying Approximate Bisimulation of Sodium Channel Switching Mechanics.

WP2.2: Zero-Defect Methods for Surgical Robots The second white paper, *Zero-Defect Methods for Surgical Robots*, was given by Yanni Kouskoulas of Johns Hopkins University and APL. The primary components when building zero-defect systems are a description of system behavior, a description of system construction, and a logical argument that the two match. In prior work, he applied formal methods to two separate surgical robot software components. The process helped uncover flaws, fix them via redesign, prove that fixes worked, and that there were no more bugs. It was much more powerful than testing approaches. His group managed to guarantee fresh and uncorrupted data transfer by lock-free concurrent data exchange implementation with a surgical Assistant Workstation software library. Additionally, they guaranteed safe enforcement of motion constraints by control algorithm when interacting with robot dynamics in an experimental skull-base surgery robot.

In order to build zero-defect systems, different proofs are needed for different development activities. Developmental guarantees ensure that components implement the correct functionality. Compositional guarantees help ensure existing components interact together to support higher-level objectives. But the following research questions still remain in order to develop a framework that helps us create more comprehensive safety guarantees than are possible today. How do we combine guarantees from different logics at different levels of abstraction into an algorithm for a single component? This would require proof that control algorithm safely restricts movements in the presence of robot dynamics, but also that it provides data to other components through a standard interface. How do we stitch together a web of detailed component guarantees to prove correct emergent-system behavior? In this case we must prove, e.g. “safe and accurate

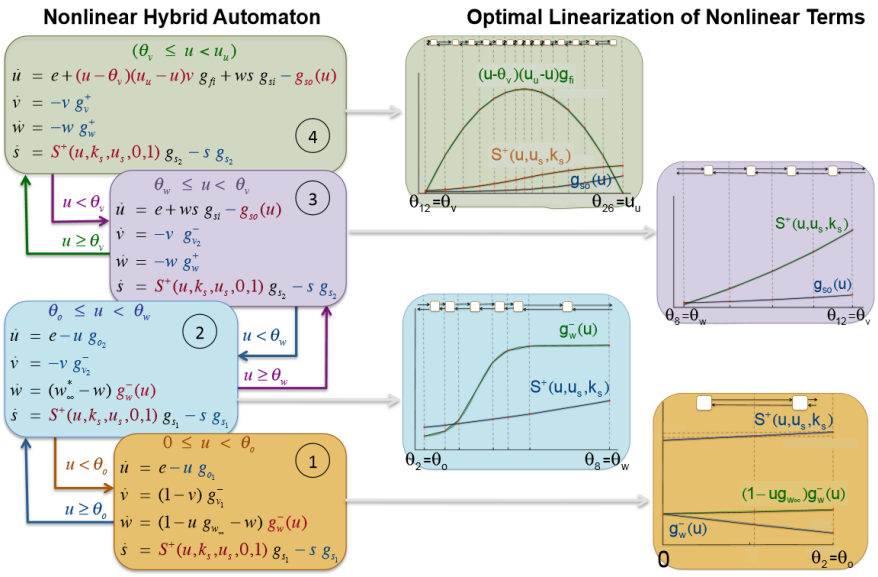

Fig. 6. Picking Optimal Linearizations of Non-linear Modes in Hybrid Model of the Heart.

incisions, according to preoperative plan" and that component-level guarantees are treated as axioms for the proof about emergent behavior but with axioms not necessarily in the same language.

[This next one probably requires the figures to be sensible as well]

WP2.3: From Verified Models to Verified Code for Medical Devices The last white paper of this session was given by Miroslav Pajic of the University of Pennsylvania, entitled *From Verified Models to Verified Code for Medical Devices*. Software failures caused 24%of all medical device recalls in 2011 and more than 1,500,000 software-based medical devices were recalled from 2002-2010. Unfortunately no well-established standards exist for development of software for medical devices. Instead, testing medical device software is currently ad-hoc, open-loop, and very expensive. The approach given by Miroslav is to verify and test the device software in closed-loop with its physical environment. This is model-driven development of pacemakers and PCA infusion pumps.

In the case of implantable medical devices, specifically cardiac pacemakers, closed-loop safety properties are retained through the tool-chain. This leads to the development of verified software from verified models.

To help illustrate the power of this technique, the presentation included a patient controlled analgesia case study. Patient controlled analgesia is a significant source of adverse events. The process combined simulation-based analysis of a detailed system model with model checking of a timed-automata model.

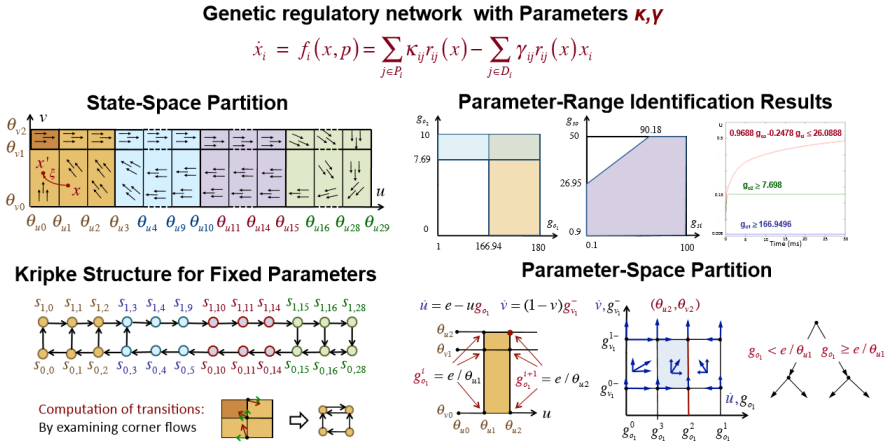

Fig. 7. Identifying Spiral Wave Formation and Breakup as a Verification Problem.

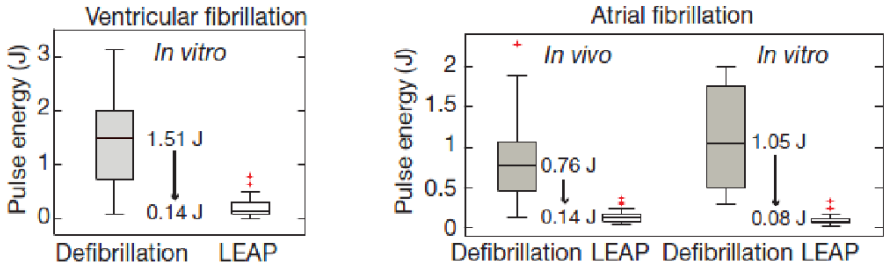

Fig. 8. Low Energy Defibrillation (LEAP) Tested for Canine Hearts.

The relationship between the two models manages to preserve the crucial aspect of the timing behavior.

5.3 Breakout Sessions

BR2.1: Medical Device Errors and Misinteractions *Medical Device Errors and Misinteractions* was the first breakout session of the second track, led by Insup Lee of the University of Pennsylvania and co-lead Paul Jones of the FDA.

Model-based Development of Medical Device Systems. Model-based development has emerged as a means of improving software quality of cyber physical systems. The model-based approach allows developers to perform rigorous model verification with respect to safety and functional requirements, and then synthesize code that preserves the verified properties of the model. Modeling and Analysis of User Interfaces and Use Errors. Medical devices such as ventilators, patient monitors and infusion pumps have a user interface that allow users to configure critical parameters, such as patient profile, therapy, and alarm levels.

Failure to configure the device with correct parameters may result in patient harm, e.g., incorrect therapy or incorrect patient information.

It is believed that the modeling and analysis techniques can be used to support early evaluation of user interface design and complement classical human factors analysis techniques. They can be applied to (1) automated verification of user interface design and implementation against interaction design principles, (2) automated verification of user interface design and implementation against mental models, (3) automated verification of user interface design and implementation against tasks carried out by users, and (4) model-based development of the user interface using formal verification tools.

Medical Device Data Loggers. Traditionally, safety critical industry sectors such as transportation rely on flight recorders or data loggers to capture information that permits engineers to reconstruct undesired or adverse behavior and correct the problem so that it doesnt happen again. Since most existing medical devices do not have data loggers the best that can be done at the moment is to look for trends in data reported to regulatory authorities.

To improve this situation, it is essential to be able to obtain medical device adverse event data in a more timely fashion, with much more detailed information, and to perform effective analysis based on the recorded information. A next generation of medical device data logging scheme can greatly help the advance of medical device adverse event analysis. Collected data are useful for (1) forensic analysis, (2) online / real-time detection and prevention of failures, (3) basis for regression tests during development of medical devices.

Assurance Cases for Certification. Certification or regulatory market approval of these complex safety critical systems is dependent on independent reviewer confidence that a device will perform as intended. The use of (safety) assurance case structures (claims, arguments, evidence) can facilitate the presentation of arguments justifying the quality of the device; at different levels of abstraction. In the case where errors or misinteractions have occurred assurance cases may be constructed that argue corrections fixed the problem and didnt introduce any new problems.

BR2.2: Securing Devices & Data Against Cyber-Physical Vulnerabilities Breakout two was led by David Kotz, who gave an earlier plenary talk, with co-leads Richard Shantz of Raytheon BBN Technologies and Shashank Priya of the NSF. The focus of the session was *Securing Devices & Data Against Cyber-Physical Vulnerabilities*. They have provided text as follows.

Introduction

Our group considered medical cyber-physical devices to include a broad range of technologies that measure and/or affect the health and wellness of individuals; these devices may include sensors and actuators, of course, but also include the control mechanisms associated with such devices, and (to a lesser extent) the integrated life-cycle systems that embed sensor and control information. More broadly, medical cyber-physical systems include individual devices, interconnected combinations of devices and back-end systems, and range from strictly clinical systems to remote-monitoring systems to home-care systems to wearable

and implantable devices. Because of their use in support of human health, cyber vulnerabilities in any part of the interconnected systems can adversely affect patient privacy and safety.

Current state

We are presently in a period of sustained, rapid innovation in computing device miniaturization and capability for medical applications, in which new medical devices are released every week. These medical devices, and systems built with such devices, are driven by software and increasingly include some form of network interconnection, whether a local-area network for integrating with other devices or providing off-board services, or a global connection – the whole of the Internet and its open-ended threats. This dependence on software and a network connection, introduces the potential for software- and network-based attacks that can impact data integrity, confidentiality, correctness, and availability of device operation. Simultaneously, the number and sophistication of threats continues to grow, as well as awareness of current deficiencies.

Challenges

The growing demand for interconnected medical cyber-physical systems increases the risks due to cyber-security vulnerabilities. Challenges in this domain often go beyond the typical cyber-security and cyber-defense issues posed by any modern software-based system. We focus on medical-domain aspects that make these challenges particularly acute. First, most medical devices are embedded devices with limited computational resources, and many have limited energy (battery) resources; these limitations make it more difficult to adopt even current, common cyber protections (antivirus and firewall technologies). Second, many medical devices are built on a proprietary software base, making it difficult to leverage existing third-party security software. Third, even those medical devices built on off-the-shelf software environments remain vulnerable to the exploits developed for those systems, because the owner or vendor cannot modify (or patch) those systems for regulatory reasons. Fourth, traditional security principles (such as "default-deny") are often inappropriate for a life-critical medical system where a denial of access can itself lead to patient harm. Thus, these systems require novel solutions that allow trade-offs between security, privacy, availability and safety (for example, providing for accountable, emergency override of security mechanisms). Fifth, many medical cyber-physical systems need to provide long-lived logging of their activity, for auditing individual and aggregate system behaviors, while recognizing the considerable risks to privacy and anonymity that system compromise entails. Sixth, there is a huge variety of legacy devices that need to co-exist with newer devices as part of integrated medical systems. We need to develop efficient and effective ways to add an improved layer of security for these legacy devices in a manner that permits interoperation with evolving subsystems. Seventh, there is an extremely wide range of user sophistication with IT systems, particularly those users who need to manage extra-clinical devices at home. These require secure configuration by non-expert users, and thus the development of easy-to-use interfaces for configuring both device operation and security. Finally, implanted medical devices

pose extra challenges: they are often expected to be long-lived, have extremely limited resources, their host cannot "take them off" to preserve privacy, software upgrade is (at best) inconvenient, and hardware upgrade often requires surgery.

Solution approaches

Ultimately, it will be necessary to develop new standards to raise the bar for the security of medical devices, the privacy of the data they produce, and the security of interconnected systems of devices. Current standards are quite limited and new standards evolve slowly; furthermore, innovation in medical devices (and integrated systems) is occurring rapidly. One complementary viable approach, used effectively in other domains, is the use of "proxy" general-purpose mediators as a tightly controlled virtual overlay to existing devices. These mediator devices interpose on the communication path(s) used by the device, detecting (or deflecting) various forms of attack and abuse, and providing opportunities to monitor and audit device use. This is especially pertinent for legacy devices and systems. Another promising approach is through contextual security, in which the security mechanisms dynamically adapt to contextual parameters such as available resources (low battery or low memory) or current medical situation state (such as an emergency situation). We note there is potential for homomorphic computing (supporting computation directly on encrypted data so it is never exposed to 3rd party services) to provide default privacy for aspects of medical CPS.

Recommendations

The following areas were specifically identified as important and viable areas for NSF investment, aiming for impact over a 5- to 10-year time horizon.

•A general methodology for constructing autonomous, "do no harm", continuously operating, long lifetime, adequately protected systems, inclusive with mechanisms to trade-off across properties to suit the setting or situation.

•Mechanisms suited to stringent protection policies governing the data while at the same time ensuring it is still useful to various stakeholders, in both individual and aggregated forms, involved within a medical eco-system.

•No (or low) maintenance and no (or low) manual configuration for authentication, access control, data sharing, privacy controls, with effective methods of violation detection and management.

•Methods for dynamically varying situational (adaptive & context-specific) notions of security on systems operations and data.

•Developing low-cost, user-friendly methods for verifying and certifying security properties of medical systems, including methods for self-configuration, dynamic adaptation, and leveraging a theory of safe composition.

•Assessing the benefits, risks and potential risk mitigation of building on Commercial Off-The-Shelf (COTS) platforms and infrastructures (such as COTS operating systems), versus robust custom-tailored solutions.

•Assessing trade-offs between adopting a common standards base vs. (or in conjunction with) constant innovation at the device and system interconnect boundaries.

•Developing, certifying and deploying appropriate fidelity models of combined cyber-physical systems operation suitable as an offline and embeddable decision aid, including modeling expected behaviors, misbehavior and compensating actions.

•Developing improved theoretical methods to characterize the potential for information leakage from raw sensor data, especially when combined with other information. Developing improved methods to clearly convey to users/patients what information is collected, where it will be stored, who will have access to it, what they will do with it - and what might be inferred from this shared data. Developing clear interfaces to enable users/patients to set appropriate policies for sharing data.

BR2.3: Coping with Alarm Fatigue and Cognitive Overload The third breakout was entitled /emphCoping with Alarm Fatigue and Cognitive Overload, was led by Julian Golman, who spoke earlier in this session, and co-led by J. Perren Cobb, who gave a talk in the first session, and Syliva Spengler of the NSF. This topic ties in well with Perren's earlier presentation about data overload in the ICU. The premise is this: we live in an age where medical data is increasingly at our fingertips, with sensors monitoring all manner of physiological and environmental information, but the people handling the information have not seen cognitive advancement to keep pace with technology. Doctors and clinicians need to be able to utilize provided data in order to make a proper diagnosis and treatment plan, and often times this requires the knowledge of which data should be ignored.

On the CPS front, we must design sensors and devices with output that is immediately understandable by our users. We need to be weary of making equipment with an abundance of alarms and warnings or risk them going unheeded. It is the classic "car alarm" problem: car alarms were designed to alert the owner and other law-abiding citizens to an unlawful breach of the vehicle. Unfortunately, car alarms had massive false-positives; now when someone hears an alarm going off, the thought is not "someone is attempting to steal a car" but "will someone turn that [expletive] alarm off already?!" With medical device alarms, false positives leading to an alarm being ignored could result in the loss of a patient, or irreparable medical damage.

BR2.4: Modeling, Simulation, & Verification to Predict Performance & Reliability John Hatcliff of Kansas State University led the final breakout of this track, entitled *Modeling, Simulation, & Verification to Predict Performance & Reliability.* John was joined by Miroslav Pajic, who presented a white paper earlier in this session, and Anindya Banerjee of the NSF. The previous talks in this session have demonstrated the power and versatility of modeling, simulation and verification. Models have been built for sensors and controllers, as well as living tissue and organs as they naturally appear. Using these models, we can simulate nominal conditions or disesed behavior, and use formal analysis to prove that properties of the system will always (or never) hold. But what can these fields tell us about the performance or reliability of future devices and techniques?

As it turns out, quite a bit. Models, by their nature, are system abstractions and can very from very simple to extremely complex. They range from basic automata to large sets of partial or ordinary differential equations (PDEs or ODEs). It should come as no surprise that a complex, signal-processing, model can serve as a blue print for device manufacture. Such a translation means that verified properties of the system should hold (with some tolerance for manufacturing errors) for the device. Similarly, simulations done with the model should be our best-estimates of performance in reality. In terms of reliability, verification offers a proof of correctness that cannot be met with even the most rigorous testing in non-trivial cases. Simulation allows a user to get information, often visual, about any manner of complication the system may undergo in practice. The difficult questions tend to revolve around those previously mentioned tolerances, the necessary complexity of the model and time to be spent on simulation/verification.

6 Intervention, Control and Prosthetics

The third track in the workshop revolves around applying the ideas and techniques from the previous tracks and applying them in real, clinical situations. Talks in this sesession focus on robotic and networked medical solutions, preventing attacks and outbreaks through constant monitoring, as well as new, advanced prosthetics.

6.1 Plenary Talks

PT3.1: Moving Beyond Current Approaches to a Broader View of Epilepsy Monitoring The first plenary talk of the final session, given by Steven Schater of the Center for Integration of Medicine & Innovative technology, was entitled *Moving Beyond Current Approaches to a Broader View of Epilepsy Monitoring*. This talk is broken into three sections: the first describes in-home monitoring with video-EEG; the second examines needs unmet by current approaches and emerging technologies to surpass them; the final mentions opportunities and challenges for future development work.

Video-EEG offers a new solution to in-home monitoring of patients. By recording events of interest in an environment where patient has the episodic behaviors, clinical suspicion of epilepsy and seizure types can be more easily confirmed. The monitoring also allows for the documentation and quantification of seizures of which patients may be unaware. A modern home-based video-EEG system can record up to 96 hours of EEG and 48 hours of HD video on one memory card. It features a low-light camera setting and uses Bluetooth to synchronize the EEG and video footage. Finally, the device can connect to a PC via USB and upload 24 hours of EEG in under 5 minutes.

While the process of epilepsy monitoring has come a long way, there are still many unmet needs in this domain. Monitors cannot currently provide warning of an impending debilitating seizure. An ideal monitor would serve as input for

closed-loop system in order to lessen risk of injuries. Injuries linked to epileptic episodes include fractures, intracranial hematomas, burns, and even death via accidents, aspiration, drowning or SUDEP. According to patients, the unpredictable nature of seizures is the worst aspect of their epilepsy. Better tools would allow doctors to monitor the long-term impact of treatments on seizure control and co-morbidities. Patient diaries in clinical trials of anti-seizure therapies could be replaced with more accurate physiologically-derived data. Such a monitor would improve adherence to medication and lifestyle issues, such as sleep. It could assess patient safety during and after seizures, even summoning help or calling for an emergency evaluation when needed. Better monitoring tools could identify patients at the greatest risk for developing epilepsy, given a history of traumatic brain injury, stroke, or Alzheimer's disease. Such a technology could correlate specific symptoms with EEG or AED levels in real-time. Given these unmet needs, monitoring is evolving from EEG- and video-based devices for short-term use in establishing a diagnosis to long-term diagnostic and treatment systems incorporating a range of technologies to manage epilepsy.

We have recently seen emerging developments in both hardware and software domains. Hardware has evolved along two tracks- EEG based and non-EEG based. New developments in EEG hardware include wearable electrodes, many of which have been designed for use in gaming or for brain-computer interfaces. We have also seen the advent of subdermal electrodes. These devices include intracranial systems for seizure detection (Responsive Neurostimulation System (Neuropace)) and prediction (Seizure Advisory System (NeuroVista)).

Non-EEG based hardware development has seen an uptick in accelerometers in either wristwatches or werable sensor networks. One popular tool is the SmartWatch. It consists of a GPS module and a proprietary accelerometer/gyroscopic sensor to detect the excessive and repeated motions of tonicclonic seizures. Buttons on the device allow users to cancel a false alert, should one occur. Other sensor developments include the electrodemal activity sensor, Epiband. This sensor measures skin resistance (marker of sympathetic tone), motion, and skin temperature. It communicates wirelessly to the patient's smartphone and has been tested as a method to detect GTCS and autonomicly correlate to postictal EEG suppression after CPS and GTCS. This is vital because postical EEF supression in this case may be a risk factor for SUDEP. Other sensors include near-infrared, mattress movement monitors, ECG, and video-based analysis of movement detection.

But there are still many opportunities and challenges for further development and clinical adoption. Steven has been networking with IT and engineers from numerous disciplines. His group has assembled development teams featuring clinical expertise, technology and patient input. Opportunities for early financial support have included the Epilepsy Therapy Project of the Epilepsy Foundation (see Shark Tank competition). There is still a need for strategies to demonstrate proof of principle, and favorable cost:benefit to financial backers. There is also the critical matter of patient and physician acceptance/adoption. Above all, any solution must allow integration across technologies to customize and individu-

alize systems for specific patients. No single sensor (EEG, accelerometer, etc.) will be sufficient for all clinical situations, so interoperability is desirable.

PT3.2: Implantable Networks of Wireless Nanoelectronic Devices The final plenary talk of the session, and of the workshop as a whole, is *Implantable Networks of Wireless Nanoelectronic Devices*, given by Pedro Irazoqui of the Center for Implantable Devices at Purdue University. It is expected that the medical marketplace will reach 19.5% of GDP by 2017, according to statistics by the US Department of Health and Human Services. The goal of enabling multi-pathology treatment through physically distributed networks of nanoscale wireless sensors and actuators has caused a paradigm shift in medical treatment. Microchips are designed and assembled at Purdue for all manner of testing. Those highlighted here include glaucoma IOP monitoring, seizure detection, prosthetic arm control (TMR), and heart failure warning. Such systems, however, require a network of implantable wireless devices, CMOS compatible nano-biosensors, packaging, and energy harvesting.

Enter NEEDS: Nano Engineered Electronic Device Simulation- led by Purdue partnered with MIT, UC Berkley and Stanford. The vision of this program is one of a new era of electronics enabled by the new capabilities of emerging nano devices. Their mission is to connect nano-material and device research to new applications with physics-based circuit simulation models. Areas of study include miniaturization, magnetically driven insertion, and energy harvesting. The idea of NEEDS was to combine the implantable devices, nanotechnology and clinical partnerships at Purdue and fill key gaps in faculty and facilities in order to achieve preeminence. NEEDS has been pushing the envelope in development of networked wireless nano-implants.

6.2 Highlighted White Papers

WP3.1: Next Generation Soft Wearable Robots to Assist Impaired Patients To start off our white papers for this track, we have *Next Generation Soft Wearable Robots to Assist Impaired Patients* by Conor Walsh of Harvard University. While recently there have been leaps and bounds in the development of exoskeletal robotic systems to increase physical output for medical, military, and commercial purposes, most of these devices feature rigid frames. Such exosuits are often heavy, bulky, difficult to align and ajust, and feature other technical problems. The bandwidth costs of these units are high, the torques imposed by the system can be difficult on the body, and the units tend to alter the normal kinematics of walking. While some of these are non-issues in the military and commercial realms, they are not ideal for medical devices.

Conor and his colleagues at Harvard have been working on soft, wearable exosuits for use by the physically impaired. The suits are made of structured textiles that spread load across their surface. They are easily adjustable, only apply tensile forces, and only apply forces at or below normal human levels. The suit features an integrated sensor network to detect kinematics, and has several modes of control available to the patient- assistive force generation, force control

and force-based position control. On top of all of this, the devices are thin and light-weight and have little to no negative impact on walking kinematics.

WP3.2: Smart Architectures for Implantable Wireless Pulse Generators for Gastric Disorder Management: A CPS Approach The second highlighted white paper of this session if *Smart Architectures for Implantable Wireless Pulse Generators for Gastric Disorder Management: A CPS Approach* by Aydin Farajidavar of the New York Institute of Technology. This work focuses on the virtual stomach model (VSM, Fig. 9), a whole-organ electrophysiological model with an ICC Network. Sensor inputs are fed to both a pacemaker (AFSM) and the stomach model (FSM). The pace maker forwards control signals to the stomach where they interact with sensors and signal processors. Fig. 10 shows that if the signal processing out of the stomach and the desired signals out of the model are incongruent, this is considered an error and a signal is propagated to the pacemaker.

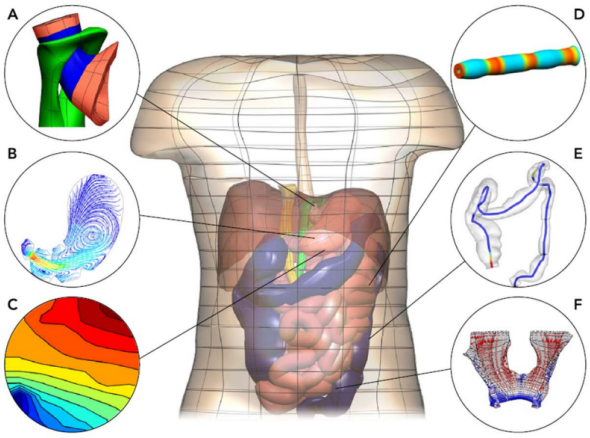

Fig. 9. Virtual Stomach model.

WP3.3: Open Platform for High-Acuity Surgical Treatment with Cyber-Physical Systems The final white paper of the workshop was given by Hawkeye King and Blake Hannaford of the University of Washington. The paper is entitled *Open Platform for High-Acuity Surgical Treatment with Cyber-Physical Systems.* Telerobotics describes a computer mediated interaction with a remote location. In this field, it is the role of the "cyber" portion of CPS to improve safety and task performance, allow remote telepresence and automation, and to compensate for any time delay. A major opportunity for new research comes in the development of the Raven II open telesurgery research platform. This effort sees NSF CRIs at 6 US locations, 4 individual contracts in the US, Canada and France and an applied dexterity lab in South Korea.

There are several research needs we have yet to meet in CPS with regards to surgical systems. New architectures for combined human/machine operation

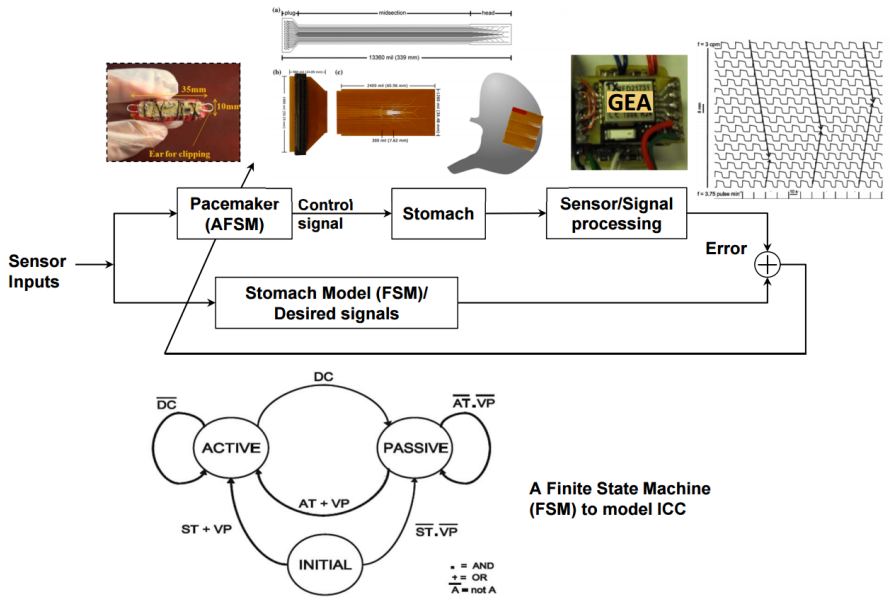

Fig. 10. Modeling Connections Between the Stomach Model, Sensors, and Physiological Controllers.

are still desired. There is a lack of detailed modeling of soft-tissue interaction. Human/robot interaction paradigms for traded human/'cyber' control are not yet effective and safe. There is a need for engineering and control design for simultaneous safety and task goals. i.e. tissue aware controllers. Finally, there must be Raven community software support, as the effort is meaningless in practical terms unless the results see use.

6.3 Breakout Sessions

BR3.1: Improving the Man-Machine Interface Conor Walsh, who spoke earlier this session, led the breakout *Improving the Man-Machine Interface* alongside Art Erdman of the University of Minnesota and Sylvia Spengler of the NSF. Conor's earlier presentation on exoskeletal robotics focused on minimizing the invasiveness of the device by engineering it to be as light and thin as possible. They have provided the following report.

Introduction

Our group considered the term human-machine interface quite broadly in the context of medical cyber-physical systems. In particular we asked the questions "Who is the human and who is doing the interfacing?" For example, it could be a clinician (surgeon, nurse, etc.) interacting with some touch screen display, a patient monitoring their own health through a website or sensors detecting the intent of a wearer of an exoskeleton system. These are all clearly very different types of systems but the group agreed that all CPS have some type of

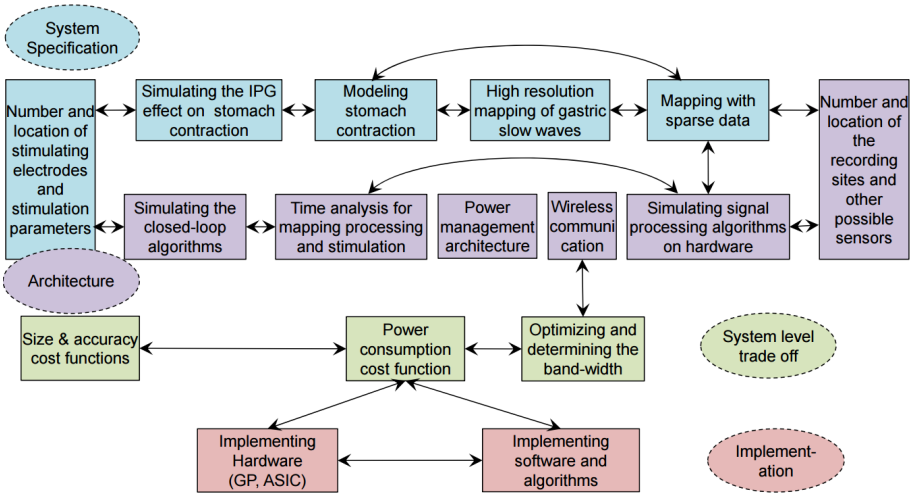

Fig. 11. Connecting Specification, and Implementation, with Consideration of System-level Trade-offs.

human-machine interface and that it should be clearly identified as such in the development and evaluation of these systems. For all of these cases, we determined that the interface was a very critical part of the CPS in order for the appropriate delivery of care and so that the CPS and the user (patient, physician etc.) can operate synergistically and effectively. It should be both intuitive to the user and also as transparent as possible so it does not add a large physical or cognitive burden.

Current state

We asked the question "What are the different types of human-machine interface?" At a high level, these could be data user interfaces (1D, 2D, 3D) for displaying and transmitting information (e.g physiological, neural, etc.) or physical machine interfaces (e.g. vibrotactile for haptics or physical bracing/attachments for wearable devices). Some interfaces can be bi-directional (in that information can flow both ways (e.g. sensing and stimulation with neural implant), some may have logging, reporting or decision support capabilities, some may have a direct connection to the human and some may have an intermediate device. The group also brought up the fact that the type of data flow was important, in particular whether data was displayed in a discrete or continuous fashion.

Challenges

A key item raised by the group is that there are a diverse set of expertise required for developing human-machine interfaces. In particular there are social and behavioral scientists (e.g. psychologists, sociologists, anthropologists, human factors researchers, industrial designers and cognitive scientists), engineers and computer scientists, biological scientists/engineers (e.g. physiologists, neuroscientists, biomechanists and experts in motor control) in addition to input from

end users (doctors, patients, etc.). For all cases of development, it is important to have a very clear understanding of the environment where the interface will be used and need to undertake testing in as close to real world as possible (e.g. a simulated environment). The group recognized that often a user interface is the part of a system that is developed last but should ideally be thought about from the start of the development process. Also, once developed, the group agreed that we need better quantitative metrics for evaluating the performance and functionality of human-machine interfaces. Also, the human-machine interface component of the CPS needs to work in cases where patients may not have full cognitive or physical capabilities (e.g. when a person is having an adverse event).

Considerations for human-machine interfaces

The group made a list (not complete) of many of the issues that need to be considered when designing a human-machine interface. These were: be secure and safe; accommodate different patient abilities; operate in the presence of uncertainly or variability; level of transparency between the user and the CPS; degree of portability required; level of intuitiveness (for ease of use); operate over multiple time-scales; be reliable and robust; level of reconfigurability; level of appropriate reporting; connectivity to larger data bases; level of invasiveness.

Open Research Questions

The following questions were specifically identified as important to address and viable areas for NSF investment, aiming for impact over a 5- 10 year horizon.

•How do we determine how frequently we need to monitor and display information with human-machine interfaces?

•Is there a way to provide universal signals to convey information in human-machine interfaces for CPS?

•Are there best practices for what sensors/sensor fusion algorithms to use?

•Can rewards be used to increase patients motivations for interacting correctly with human-machine interfaces?

•What are the correct techniques for handling massive amounts of data and correlating it with the individuals physiology (each person is unique)?

•What are new methods to combine discrete and continuous mathematics (e.g. stochastic approaches) to improve human-machine interface functionality?

•How to deal with uncertainly in human-machine interfaces for CPS?

•What are experimental approaches to characterize the performance of human-machine interfaces for human-in-the-loop CPS?

•What are appropriate methods to use to deal with multiple time scales and data abstraction levels?

•What are new approaches to modeling in order to simulate the interaction between the human and devices/systems and how can these be validated with experimental data?

BR3.2: Just-in-Time Health Interventions The second breakout session was led by Jim Rehg of the Georgia Institute of Technology and co-led by Wendy Nilsen of NIH. The session was entitled *Just-in-Time Health Interventions*. For many medical conditions even a small delay in treatment can lead to complications or even death. One of the most well-known instances of this problem is the case of

defibrillation: the earlier it is used, the more likely the patient will survive. In the case of patients with chronic conditions, such as epilepsy or heart complications, worn or implanted medical devices can potentially apply treatment as soon as the precursors of an episode are sensed.

Jut-in-time interventions, along with customized treatments, are considered to be an absolute end-goal for wearables and implanted devices. Constant monitoring of the patient's physiology and environment, along with heuristics tuned to the patient's medical history can act as a force-multiplier in effectiveness of treatment. Unfortunately, there are a number of complications before JIT treatments can become a universal reality. First and foremost, what happens when the device has an electronic or mechanical failure? Is the device secure from attack by a hostile actor? What if the device detects a problem, and applies intervention as such, when no problem exists? What happens when the device fails to act? These are all real issues that impact the proposition of JIT interventions at the moment. As a CPS community, we must be able to address these problems with reliable, redundant devices that always act when the should and never act when they should not.

BR3.3: Medical Robotics The breakout section on medical robotics, led by Peter Kazanzides and co-led by Hawkeye Kinf and Tho Nguyen, prepared an in-depth report as follows.

Intro

Medical robotics represents the ultimate merger of cyber- and physical- systems. Sensors and data create a health-care model for the patient, algorithms and human-computer interaction create a therapeutic course of treatment, and mechatronic arms and effectors physically treat the patient. Ultimately the machine affects a person's life on a profound level.

The most successful and widely recognized among medical robotics are telesurgical robots exemplified by the daVinci robot by Intuitive Surgical. Surgical robots are leading a revolution in surgical care. At the same time, however, medical robotics includes other applications like rehabilitation, prosthetics and exoskeltons, interventional radiology, robotic catheters, colonoscopy, and home health care.

What do we know and do well today?

The most mature CPS medical robotic systems are local master-slave teleoperation robots where the surgeon and patient are in the same operating room. FDA approved systems from companies like Intuitive Surgical (Sunnyvale CA) and Hansen Medical (Mountain View, CA) use this paradigm. Other systems use a cooperative control paradigm, where the surgeon and robot share control of the instrument, as in the FDA approved system from Mako Surgical (Ft. Lauderdale, FL.)

Autonomous and human-machine collaborative surgical robotics has been demonstrated in static surgical environments, e.g., orthopedics and neurosurgery applications. These have the advantage of operating in well-known environments, where the anatomy does not deform significantly.

Medical imaging techniques are highly mature and provide excellent sensing and environment mapping that can be used with robotics

Non-medical robotics have many proven technologies that could be highly beneficial in the medical realm. For example, precise repetitive motion, miniaturization, and neural interfaces are all available robotics techniques that could be brought to bear.

What are the research challenges and approaches?

Modeling is important for planning and executing safe effective therapy. Soft-tissue modeling for surgery. Human behavior and therapy modeling for rehabilitation. Modeling whole procedures for stateful robot action and human-machine ergonomics.

Sensing and adaptation: real-time environment sensing is limited by needs for sterilizability, miniaturization. More sensing of the robots actions and environment will lead to more effective procedures. Miniaturized, sterilizable force, blood oxygenation, 3D imaging, and robot-guided ultrasound sensors would improve MIS robotic surgery.

Novel actuation will enable, for example, miniaturized surgical manipulators, intrinsically safe home health-care systems, or in-body robotics.

We need better validation platforms to develop and compare systems. One example is realistic phantoms that are standard and sharable, such as a realistic CPS beating heart. Shared ground-truth/reference data sets provide the raw material for new data-driven research. Access to shared, powerful simulation platforms (e.g. computation-intensive cloud based physical models) could help validate new methods in simulation providing new training platforms and shortening the hardware development cycle.

Common benchmarks for medical robotics, both surgeon and robot performance, is key to validating new technologies.

Human-robot interaction technologies with rich haptic, visual, verbal, and neural interfaces will improve our control of existing teleoperation systems. Non-surgical patients will benefit from richer feedback, e.g., haptic and neural interface prosthetics.

Human-machine collaborative control of teleoperation systems will combine human capabilities (intelligent, dexterous, safe) and robotic ones (precise, abundant). Augmented reality (haptic, visual) telepresence will improve human perception of and action in the treatment space.

Presently there are no robot interaction paradigms that dictate interaction between the care team, the robot, and the patient. How can information flow to all the caregivers in the robotic OR when the room itself is acting autonomously? How will a home healthcare robot inform the patient about their cooperative rehabilitation activities?

Educational Needs

First and foremost cross-disciplinary training and collaboration among clinicians and engineers is critical. This focuses engineering problems on positive therapeutic outcomes, and enlightens medical providers with new possibilities in delivering care.

Open platforms for training and experimentation will quickly bring neophytes up to the state of the art. This includes open simulators, open-access hardware such as Raven and the daVinci Research Kit, and possibly remote-access platforms- open high-performance centralized simulation sandboxes.

More courses in medical device development will train clinicians and engineers about cutting edge practices.

Hardware Needs

In CPS Surgical Robotics the primary hardware needs are for miniaturized, biocompatible, sterilizable sensing and actuation. Three elements together would provide a new level of immersive robotic telesurgery: rich force and tactile sensing at the end-effector, highly dexterous multi-dof end-effectors, and advanced haptic human-machine interfaces.

Roboticized miniature sensors like ultrasound, Optical Coherence Tomography and 3D imaging will help surgeons see and feel the surface and internal structures the operate on.

Non-invasive sensing and actuation such as HIFU, OCT, Gamma Knife could make a tremendous impact on medical care. Robotic control of these systems will be absolutely necessary for targeting treatment sites.

System Needs: Interoperability

The field of robotics is unique in its system-level integration of many varied components. Incorporating heterogenous sensors, actuators, and devices with distributed computation and human-interaction requires a wide variety of technologies to work together.

Future medical robotics research should use and contribute open source software to minimize redundant development, and leverage existing work. Robot Operating System (ROS) is a popular and useful opne-source software framework. Surgical Assistant Workstation (SAW) is also open-source purposefully designed for surgical robotics.

Non proprietary data exchange standards should be more widely adopted to ensure compatibility of developed robots. Several ad-hoc research standards for robotic systems integration are likely candidates, including OpenIGTLink, a data interface designed for image guided therapy. Interoperable Teleoperation Protocol (ITP), another possible standard, was specifically designed for master-slave teleoperation and telesurgery robots.

Furthermore, devices entering the clinical setting should make efforts towards a standardized user interface to simplify training and reduce the cognitive load required to operate many machines.

Safety

CPS Medical Robotics presents a unique safety challenge. A robot has a direct effect on a patient's physical state, and compromise of a robot can do immediate harm.

At the same time, physical properties of the system give a convenient method for validation of robot behavior. Pure teleoperation systems' output should not exceed the performance envelope of the human operator, providing validation of intended movement commands. Similarly, a rehabilitation system should not

apply more force to a patient than the patient can withstand. As a result, physical models of system behavior will provide safety checks (in addition to other potential benefits mentioned above).

Methods for validating CPS medical robotics might include formal methods such as a model checking and deductive verification techniques, or through extensive testing.

Safety vulnerability in the network and networked control loop must be addressed. Unique, CPS robotics specific attack on the network and network based control loop such as delaying, dropping or spoofing feedback messages need to be addressed. Addressing these attacks will rely on security mechanisms proposed for networked control, while other mechanisms will be based on the knowledge about physical constraints and dynamics of the system. For example, knowing the allowed operating region of the remote manipulator will allow for quick detection of spoofed feedback messages.

Roadmap

Medical robotics development in the future will have the following themes: increased access with decreased invasiveness, continued merging of human and machine control, richer feedback and dexterity.

5 Years

Surgery: medical imaging augments surgeon's view of the patient. Novel, miniaturized robotic systems provide fully-insertable manipulation and visualization platforms. Miniaturized articulated arms provide dexterous single-port surgery capabilities.

Rehabilitation: first robots in the home help patients comply with rehab, monitor progress to help with future diagnosis, receive updates from physician.

Prosthesis: prosthetic devices will move towards closed loop control. Not only will better control algorithms be implemented (both EMG and neural - EEG, BCI, spinal), but sensory feedback via haptics and targeted muscle reinnervation will also provide an important component in control.

10 Years

Surgery: immersive telepresence surgery with 360 degree visualization of anatomy. Augmented reality overlays merge medical imaging with vision to automatically identify anatomical structures. Surgical workflow heuristics and pre-operative plans guide surgeons through procedures and minimize mistakes. Partial automation helps with precise dissection, suturing, manipulation. Tissue interaction modeling estimates and minimizes unnecessary tissue damage.

Prostheses: Progress towards osterointegrated prosthetic implants.

General medical: Robotic devices take over many manual tasks, like blood pressure, blood draws, immunizations.

Home healthcare: In home telepresence robots provide doctor housecall consultation, assistance with healthy behaviors and lifestyle choices, first-responder capabilities, and act as rehabilitation partner.

Security: new network security features for robotics use operator or task behavior signatures to identify intended/erroneous action.

Imaging: swallowable, steerable cameras.

20 Years

Surgery: high degree of automation in a mostly robotic OR. Imaging, planning, and actuation are all highly integrated and robot assisted.

Prostheses: cyborgs, direct interface of neurons to microcontrollers. Growing neurons on chips embedded in some biocompatible material/prosthetic limb.

BR3.4: Neural Prosthesis and Interfaces The final breakout, *Neural Prosthesis and Interfaces*, was led by Pedro Irazoqui, who gave the second plenary talk of this session. His co-leads were Aydin Farjidavar, who presented an earlier white paper, and Shashank Priya of NSF. While the CPS field has help propel great advances in prosthetics, neural implants and other neural prosthetics are still in their infancy. The human brain is the most complex organ in the body, with a large degree of its exact functionality still a mystery. It is also perhaps the most dangerous organ to interfere with, as it is the epicenter of control signals sent throughout the body.

Nonetheless, neural prosthetics and neural interfaces for other, existing, prosthetics are full of incredible promise. With advances in nano technology, we have seen the advent of tiny, energy-harvesting, implants offering swaths of new options in sensing and control. As our understanding of neural control signals evolve, we are able to hijack, redirect, duplicate, or generate our own pseudo-natural signals to seamlessly control prosthetic devices. Neuraly implanted sensors have granted new insights into epilepsy and its warning signs. The addition of control devices offers the possibility of just-in-time treatment to prevent or respond to an episode. While we might lack the understanding necessary to completely replace portions of the human brain with machines, as our knowledge grows, so do the possibilities.

7 Dinner Keynote Talk

At the end of the first day of the workshop, we were privileged to welcome Richard Satava, a M.D. in the Department of Surgery of the University of Washington. Richard's keynote talk was entitled *Advanced Technologies Beyond the Horizon:Considerations for Strategic Planning*. He began the talk with a quote from the late, great Yogi Berra, "The Future is not what it used to be." This set the frame for the concept of "disruptive vision"- what we expect of the future is often vastly different than what occurs due to an unforeseen development that alters the course of history. Richard brought up another famous quote by Henry Ford, "If I has asked people what they wanted, they would have said a faster horse."

When it comes to the future of medical care, we often have expectations of developments yet to come. At the same time, no matter how vast our expectations, they will never meet with the reality of the challenges set before us. The rest of the talk addresses the current state-of-the-art in medical care, what seems to be on the horizon, and what challenges seem to lie before us. This is all, of course, keeping in mind that disruptive technology can and likely will alter development trajectories in the future.

In terms of surgery, there are a number of innovations that are beginning to come together to form an operating room without human operators. Throughout this report, we have talked about advances in telesurgery via robotics as well as advances in simulation and modeling. Combining these technologies with advances in minimally invasive and open surgical techniques, pre-operative planning through detailed 3D models, and inter-operative navigation using the same models as overlays on top of camera feeds, we can have precise surgeries without a surgeon anywhere in sight. The first instance of such an OR was at SRI International in Menlo Park, CA in January of 2007.

There is an emerging field of micro robotics for surgical use. These machines are locomotive sensors with either numerous legs, submarine-style rotor-propulsion, or even vibration-propelled motion.

Other surgical advances include the femtosecond ($1 * 10^{-15}$ sec) laser for precise ablation, consoles for cellular surgery, and atomic force microscopy. In fact, some surgical advances lead us to question what qualifies as "surgery." Cyberknife, for example, allows for solid organs to be ablated in a non-invasive procedure. This poses the question: is this strategy actually surgery or radiology?

We can also ask "when does care begin?" We know that the earlier care is applied to the patient, the less risk of complications. Imagine a humanoid robot, complete with on-board medical devices, cutting through a wall to rescue and treat the victim of a disaster. The Rescue Bot- DARPA Robotics Challenge was the first step towards making this a reality. The challenge has led to great advancements in robotic agility, including the famous Boston Dynamics robots.

Looking at current trends, we have a rough idea of what might happen in the future. Trends in surgery have seen us going from open surgery to endoluminal, to minimally invasive, and now multi-invasive surgical strategies. In the future, we expect more of these procedures to become completely non-invasive. We might also see a progression from traditional surgery to flexible endoscopy, laparoscopic techniques, robotic surgeries, to directed energy. The "energy weapons" of science fiction, lasers and plasma, will be directed inwards instead of outwards; directed energy will allow for faster, more precise surgery.

But we know that disruptive technologies can throw everything off of the rails. What might those technologies be? There has been some signs of progress towards developing a brain machine interface. Such a technology would change the idea of prosthetics and control. With such a device, one could control a foreign object as if it were a part of their own body. But do we even need complex control of medical devices to come from the user? Not necessarily. Many motor functions can be achieved by computer systems in "intelligent prostheses." Or, avoiding the artificial altogether, we can use tissue engineering to grow human tissue on a mouse or other surrogate. We're also beginning to see advances in cryogenics, stereo-lithography, and genetic engineering of food and drugs.

But all of these advancements, both here and around the corner, bring up new ethical dilemmas. Technology is advancing at a rate at which society cannot keep up. What are the moral implications of cloning, or designer genes for humans? What about artificial life extension or organ replacement? At what point does a

human cease to be human? Could advances in life prolongation lead to a more diseased population, or could we see the end to disease overall? Finally, who will be allowed to make use of this technology? These ethical questions will need to be answered as we progress scientifically.

8 Conclusions and Future Directions

The 2014 CPS Medical Devices workshop brought together leading experts in biology, medicine, computer science, and engineering. Three sessions across two days saw discussion on topics of monitoring and diagnosis, modeling, verification and trustworthiness, and intervention, control and prosthetics. Cyber-physical systems offer unique challenges computationally, logistically, mechanically, and in human-computer interaction. The solutions presented in the talks and break-out sessions are still early steps in a rapidly developing field.

The primary outcome of this workshop was that it provided a multidisciplinary and multiorganizational forum in which new research directions in the CPS area were discussed and explored. In terms of the engineering of high-confidence medical devices, the following findings and recommendations emerged from workshop discussions.

There seems to be three main barriers to engineering medical devices with lower error or miss-interaction rates. The first barrier is that there are no good "plant" models to engineer medical devices against; i.e., there is no equivalent to physics for human physiology function or human computer interfaces.

The second barrier is due to the lack of appropriate standards of practice and adequate (trusted) tool support for developing safety critical CPS.

The third barrier has to do with how engineers can learn from mistakes. If a patient dies or is injured, it is not clear if it is due to a device error (malfunction), misinteraction, or because they were already too sick.

Future Directions in Monitoring and Diagnosis :
•New, better domain specific languages for automated biology.
•Automation in preparation and use of wet lab tools.
•Distributed control systems over the Internet with timing constraints and QoS guarantees.
•Individualization- custom monitoring and diagnosis for each individual patient.

Future Directions in Modeling, Verification and Trustworthiness :
•Further innovation in mobile- and cloud-computing systems. Must be ensured private, verified correct.
•Nationally shared comprehensive logging and analysis.
•Efforts to open and crowd-source existing platforms.
•Diagnosis and treatment of cardiac disorders using controllers based on formally verified models.
•Proving correctness of combinations of verified components, especially with regards to emergent behaviors and properties of the combined system.

Future Directions in Intervention, Control and Prosthetics :
•Integration of existing epilepsy monitoring technologies to improve interoperability and customize solutions for patients.
•Architectures for human/machine combined operations in telesurgery and other semi-autonomous fields.
•Paradigms for switching control between humans and robots.
•Detailed modeling of soft-tissue interaction.
•Safety-aware controller engineering and design.