Toward the Development of a Medical Simulation Training Architecture (MSTA)

Harald Scheirich, MSCS
SimQuest, LLC
Washington, DC
hscheirich@simquest.com

Jeffrey M. Beaubien, Ph.D.
Aptima, Inc.
Woburn, MA
jbeaubien@aptima.com

Rodney Metoyer, Ph.D.
BioMojo, LLC
Holly Springs, NC
rodney@biomojo.com

Gianluca DeNovi, Ph.D.
Massachusetts General Hospital
Boston, MA
DeNovi.Gianluca@mgh.harvard.edu

Timothy Kelliher, BSEE
SimQuest, LLC
Charlton, NY
tkelliher@simquest.com

ABSTRACT

Unlike tactical simulators – which leverage open communication protocols such as High Level Architecture (HLA) or Distributed Interactive Simulation (DIS) to permit large-scale, networked training exercises – most medical simulators have been developed using closed and/or proprietary system architectures. While some open-source simulation platforms do exist – including the Advanced Modular Manikin (AMM), the Open Surgical Simulator (OSS), and the BioGears physiology engine – these platforms are not mainstream. Moreover, each was designed to provide reusable simulation components within a very specific content area. The lack of "across-the-board" simulator interoperability presents numerous challenges for the medical simulation and training communities. For example, it is extremely difficult for a simulated patient's medical condition and physiological parameters to be transferred from one simulator (a realistic patient manikin) to another (a Virtual Reality surgical simulator) during a simulated patient "handoff" between two groups of clinical care providers. Recognizing these limitations, the Program Executive Office for Simulation, Training, Research, and Integration (PEO-STRI) and the Joint Project Manager for Medical Modeling and Simulation (JPM-MMS) have been leading the effort to develop a next-generation medical simulation-based training architecture that will provide an open standard for medical simulation, and are currently developing middleware that will permit integration with HLA, DIS, DDS, xAPI, and related communications protocols. This effort is one part of a much larger DoD effort that seeks to develop a robust Joint Medical System Enterprise (MSE) that spans the continuum from Role 1 (on-site care provided by a Combat Medic) to Role 4 (long-term treatment and rehabilitation that is provided at a CONUS facility) care. The purpose of this paper is to describe the project's vision, goals, approach, current status, future plans, challenges, and lessons learned.

ABOUT THE AUTHORS

Mr. Harald Scheirich, MSCS is a Principal Software Engineer at SimQuest, LLC with more than 25 years of experience designing and developing simulations, games, and other interactive systems. He was the lead developer and architect for SimQuest's Open Surgical Simulation (OSS) platform. His work at SimQuest has been funded by the U.S. Department of Defense, the National Institutes of Health, and the National Institute of Standards and Technology. He holds a Dipl. Ing. (MS equivalent) in Computer Science from the Technical University in Vienna, Austria.

Dr. Jeffrey M. Beaubien is a Distinguished Principal Scientist and Institutional Review Board (IRB) Chair at Aptima, Inc. For the past 20 years, his work has focused on training and assessing leadership, teamwork, and decision-making skills. His research has been sponsored by the U.S. Navy, the U.S. Army, the U.S. Air Force, and the Telemedicine and Advanced Research Technologies Center, among others. Dr. Beaubien holds a Ph.D. in Industrial and Organizational Psychology from George Mason University, a M.A. in Industrial and Organizational Psychology from the University of New Haven, and a B.A. in Psychology from the University of Rhode Island.

Mr. Rodney Metoyer is the Director of Research and Development at BioMojo, LLC. He is an engineer with decades of experience using and developing medical training systems, and was the lead physiology modeler for the BioGears

open-source physiology engine. Prior to joining BioMojo, he served as an Army Combat Medic, an Evacuation Squad Leader, and a Battalion Aid Station Treatment NCO. Mr. Metoyer holds both a B.S. and a M.S. in Aerospace Engineering from North Carolina State University. He is currently finishing his doctoral research in Mechanical Engineering, also at North Carolina State University.

Dr. Gianluca DeNovi is Director of the Medical Device and Simulation Lab at the Massachusetts General Hospital (MGH) in Boston, MA. He received his Master's degree in Electronics Engineering and his Ph.D. in Robotics and Control Systems from University of Bologna (Italy). After spending one year as a visiting researcher at Harvard School of Engineering's Bio-robotics Laboratory, he completed his postdoctoral fellowship in medical simulation at the Department of Imaging of MGH and Harvard Medical School in Boston. In 2011, he became a Visiting Professor at the University Campus Bio-medico of Rome, Instructor at the Harvard Medical School in 2013, Instructor at the Harvard Extension School in 2017, and Director of the Medical Device and Simulation Lab at MGH in 2019. His research interests include computer-based surgery simulation and surgical robotics.

Mr. Timothy Kelliher is the Director of Engineering at SimQuest, LLC. He has over 30 years of experience in systems and software engineering. Mr. Kelliher has worked in military and medical simulation systems for much of his career. He currently serves as Principal Investigator for the MSTA development effort. He is also the Principal Investigator for a project that is developing a prototype of the Joint Evacuation and Transport Simulation (JETS) System for Joint Program Commitee-1. Mr. Kelliher's work has been sponsored by the U.S. Navy, the U.S. Army, the U.S. Air Force, the Defense Advanced Research Projects Agency, and several of the National Institutes of Health. He holds a B.S. in Electrical Engineering from Rensselaer Polytechnic Institute.

Toward the Development of a Medical Simulation Training Architecture (MSTA)

Harald Scheirich, MSCS
SimQuest, LLC
Washington, DC
hscheirich@simquest.com

Jeffrey M. Beaubien, Ph.D.
Aptima, Inc.
Woburn, MA
jbeaubien@aptima.com

Rodney Metoyer, Ph.D.
BioMojo, LLC
Morrisville, NC
rodney@biomojo.com

Gianluca DeNovi, Ph.D.
Massachusetts General Hospital
Boston, MA
DeNovi.Gianluca@mgh.harvard.edu

Timothy Kelliher, BSEE
SimQuest, LLC
Charlton, NY
tkelliher@simquest.com

BACKGROUND

Unlike tactical simulators, which leverage open communication protocols such as High Level Architecture (HLA) or Distributed Interactive Simulation (DIS) to permit large-scale, networked training exercises, most medical simulators have been developed using closed and/or proprietary system architectures. While there are a few notable exceptions – such as the Advanced Modular Manikin (Advanced Modular Manikin, 2015), the Open Surgical Simulator (A. G. Gallagher & O'Sullivan, 2011; Kelliher, Lenoir, Novotny, & Scheirich, 2014), and the BioGears physiology engine (Bray et al., 2019) – these simulation platforms are not mainstream. Moreover, each was designed to provide reusable simulation components within a very specific content area. The lack of "across-the-board" medical simulator interoperability presents numerous challenges for the medical simulation and training communities. For example, it is extremely difficult for instructors to: 1) network multiple simulators together when training patient handoffs from one clinical care team to another; 2) to develop novel simulator-based metrics of learner proficiency (that were not hard-coded by the simulator manufacturer), and; 3) to network the various simulators with external tools, such as a centralized Learner Record Store (LRS), to track learner proficiency over time.

Consider the following realistic training scenario: Using a realistic patient manikin, the emergency care team at a Battalion Aid Station (BAS) practices treating a Soldier with a gunshot wound (GSW) to the shoulder. They have successfully stopped the bleeding, administered blood products, and stabilized the patient's vital signs. The simulated patient is then handed off to an aeromedical evacuation crew (again using the manikin) for transfer to a Combat Support Hospital (CSH), where a team of trauma surgeons will repair the damaged tissues (using a Virtual Reality based simulator). Given the current state-of-the-art in medical simulation, the instructors need to do a lot of "hand waving" to make this scenario play out. For example, at some point during the training event, one simulator (the manikin) must be physically replaced with another (the VR simulator). Because the two simulators are not networked, the VR simulator's initial physiological parameters will need to be manually set by the instructor, and may differ somewhat from those of the manikin. Moreover, because the two simulators are likely using fundamentally different physiology models, the surgical team may be presented with misleading or contradictory fidelity cues – visual representations, vital signs, and other physiological parameters – that they would not normally expect to see given what they observed on the manikin while preparing the patient for surgery. This can potentially result in negative transfer (to the extent that the surgical team incorrectly acts on the misleading or contradictory fidelity cues), confusion (which degrades the quality of the overall learning experience), and frustration (which decreases the learners' satisfaction with training). With this in mind, the long-term goals of the this project are to support the next generation of medical simulation and training, thereby proving our wounded Warfighters with the best possible medical and surgical care.

PROJECT VISION AND GOALS

Recognizing the inherent limitations of the current state-of-the-art in medical simulation, the Program Executive Office for Simulation, Training, Research, and Integration (PEO-STRI) and the Joint Project Manager for Medical Modeling and Simulation (JPM-MMS) have been leading the effort to develop a robust Joint Medical Simulation

Enterprise (MSE) that will span the continuum of care from Role 1 (on-site care that is provided by a Combat Medic or Navy Corpsman) to Role 4 (long-term treatment and rehabilitation that is provided at a CONUS-based rehabilitation facility). The Medical System Training Architecture (MSTA) is a key part of the larger MSE. When complete, the architecture will provide an open standard for next-generation medical simulation and training that leverages custom-developed middleware to integrate the various system components, such as simulators, performance measurement engines, visualization dashboards, and Learner Record Stores (LRS). While the initial software reference implementation is being developed using the Data Distribution Service (DDS) communications protocol, the system is fundamentally technology-agnostic. As a result, it will work with other communications protocols including HLA, DIS, and xAPI.

Before going further, we briefly enumerate the system components that will eventually be subsumed within the larger ecosystem (Scheirich, 2017). They include: 1) a web-based portal that provides learners with point-of-demand (POD) instruction and a means to assess their medical knowledge; 2) multiple simulation platforms for practicing critical medical and surgical skills; 3) custom-developed middleware, based on the DDS standard, for conveying patient physiology parameters and other data in a standardized format; 4) a networking architecture and protocol that allows the various simulators to interact with one another; 5) an instructor control panel which allows instructors to select from a library of training scenarios, and to dynamically modify scenarios on-the-fly, as needed; 6) electronic performance assessment tools for computing unobtrusive simulator-based metrics of learner proficiency; 7) services for interacting with the performance measures and for recording video examples of the learners' performance during the AAR; 8) an interactive dashboard for conducting the post-training AAR; and 9) a centralized LRS for tracking and trending learner proficiency over time. The research team has developed a detailed domain model, which contains 21 different "views" of the proposed system. One of those views is depicted in Figure 1.

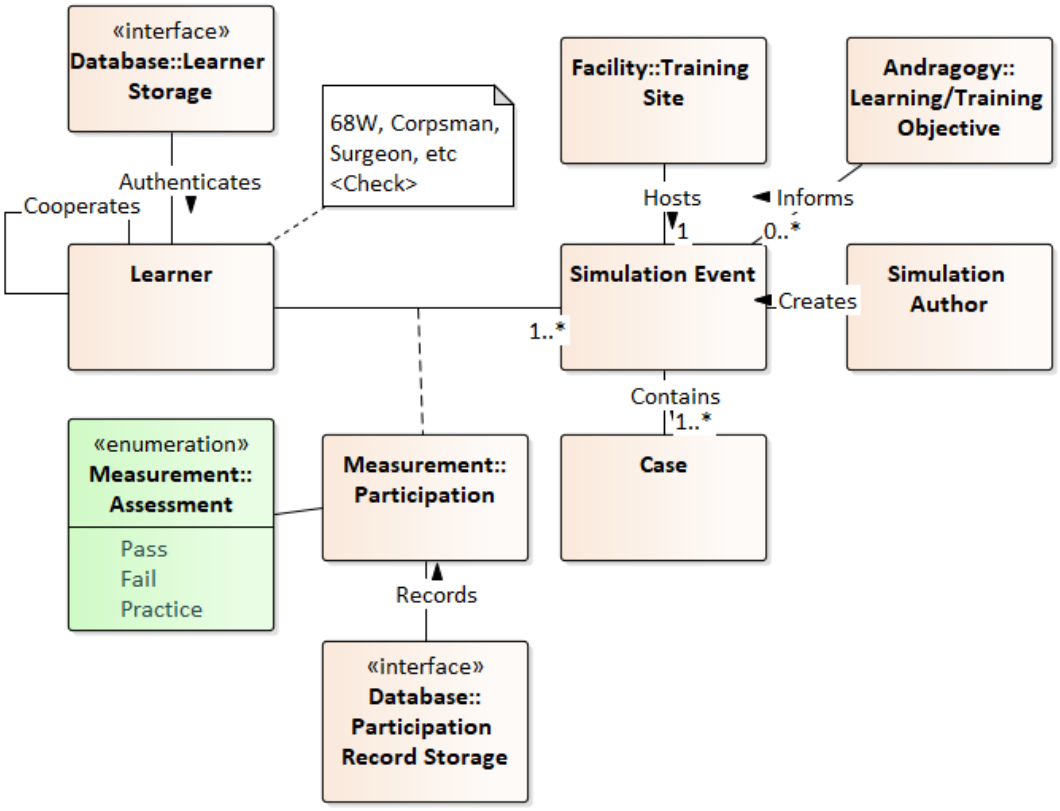

Figure 1. One of the 21 different views of the system, as depicted in the domain model. In this particular view, the learner completes a simulation-based training event at a dedicated training facility. The learner's identity is first validated from a centralized database. While performing the simulation, the learner's performance is assessed and the resulting learner records are stored in a dedicated Learner Record Store (LRS) for long-term tracking.

CURRENT STATUS

The purpose of this two-year research and development (R&D) effort is to develop a Product Line Architecture (PLA) that will support the next generation of medical simulation. Funding has been provided by the Program Executive Office for Simulation Training, Research, and Instrumentation and the (PEO-STRI), and program guidance has been provided by the newly-formed Joint Program Manager for Medical Modeling and Simulation (JPM-MMS). When the current effort was launched, there were multiple, partially-overlapping R&D efforts that each addressed different components of the larger problem space. Specifically, there were efforts to develop forward casualty care simulation systems that focus exclusively on point-of-injury care; Joint evacuation and transport simulation systems that focus on the aeromedical evacuation of wounded Warfighters from the front lines to Forward Surgical Hospitals (FSHs); hospital training simulation systems that focus on definitive care, and; learning system tactics and training systems that focus on the science of learning. The architecture is intended to be an interoperability framework that will permit these various systems to be networked together within a larger system-of-systems. Returning to the example of a wounded Warfighter with a GSW to the shoulder, it is envisioned that the architecture will serve as the "glue" that connects these various training systems, so that DoD medical personnel can practice these different components of care using highly realistic simulated patient avatars who are handed off from one care provider to another.

During the first year of this two-year R&D effort, the research team focused primarily on designing the system architecture and Concept of Operations (CONOPS). We began by assembling a team with a wide representation of perspectives, including the developers of patient manikins and physiological models, scholars from the learning sciences, and representatives from academic medical centers. We then developed a detailed model of the current state-of-the-art in medical simulation which was represented in 21 different views, including elements of: simulation in general, simulation events, simulation modalities, simulator states, clinical training scenarios, clinical case data, human performance measurement, physiology engines, and the visualization of learner records, among others (Scheirich, 2017). To develop the system architecture, the research team also developed a series of 11 detailed technical documents – which we refer to as MSTA Requests for Comment (MRFCs) – that were disseminated to a select group of engineers and researchers to more fully explore specific topics within the domain space: lexicons, levels of patient representation, interoperability architectures, simulation control data models, physiology data models, human performance measurement, After Action Reviews (AARs), and the like. At the start of the development process, the purpose of these review documents was to summarize the current state-of-the-art within each specific topic area, with the intention of illuminating the likely path forward. For example, the patient physiology MRFC noted that many current human patient simulators are "symptomatic" simulations. By that we mean that the instructor can manually increase or decrease the patient's heart rate (for example). However, because there is no underlying physiological model that controls all of the physiological systems, the expected "ripple effects" – such as changes in blood pressure, respiration rate, and the like – do not always occur when the heart rate is manipulated. As a result, the learner can be provided with misleading or contradictory fidelity cues. Therefore, looking to the future, there is a long-term need for integration with comprehensive physiological models that can "drive" the behavior of multiple, different simulators. As our ideas matured, the documents' purpose shifted towards being the basis for forming a standards documentation that guides the process of maturing the technology into a useful product.

Our first year architecture and design activities – particularly the writing of the MRFCs and the resulting group discussions that they evoked – provided critical insights that were eventually incorporated into the final domain model. For example, one of the critical insights was the lack of standardization within the medical community. Unlike the tactical community, there is no single set of Mission Essential Task Lists (METLs) that clearly identifies the tasks that a team of medical professionals must be able to perform. Similarly, there are no universally accepted set of Tactics, Techniques, and Procedures (TTPs) that describe expected behaviors either in the abstract or in specific situational contexts. With few exceptions, such as emerging standards developed by the Committee on Tactical Combat Casualty Care (Co-TCCC), different providers perform the various medical tasks based on how and where they were trained. At best, there have been some attempts to provide a common lexicon and taxonomy of the human anatomy, physiological parameters, medical procedures, and medical tools, such as the open-source SNOMED CT database (SNOMED International, 2019). Another one of the critical insights from our Year 1 activities concerns the role of measurement during training. It is a truism that "you can't train what you can't measure," if for no other reason than you will never know if the learner has mastered the instructional content (National Research Council, 1996). Therefore, the research team has intentionally referenced emerging standards for measuring learning such as the Experience Application Program Interface (xAPI) and Human Performance Measurement Language (HPML) (P. S. Gallagher, Folsom-Kovarik, Schatz, Barr, & Turkaly, 2017; Wheeler Atkinson, Tindall, Killilea, Tolland, & Dean,

2017). We believe that the measurement of human performance data can serve multiple uses within the larger ecosystem, such as facilitating the post-training AAR, and providing a corpus of data with which to make data-driven decisions about training Return on Investment (ROI), and the acquisition of next generation simulators, among others. As part of our overall design philosophy, we have intentionally referenced these types of emerging standards – including SNOMED, xAPI, and HPML, among others – in our architectural and design documents. While imperfect, they minimize the need for rework and represent a critical step in the right direction when designing the next generation of medical simulation.

During the second year of this two-year R&D effort, the research team focused primarily on developing a reference implementation of the architecture, which is undergoing technical demonstration in June 2019. The reference implementation links together disparate simulators and training tools, which include both currently-deployed and next-generation training technologies (see Figure 2). Specifically, the following technologies are being networked as part of the reference implementation:

- A current generation patient manikin (Laerdal SimMan 3G)
- A constructive representation of the patient physiology (the PULSE open source physiology engine)
- An independent Simulation Control Panel (custom-developed Qt application),
- A part-task trainer (MGH Tension Pneumothorax Simulator),
- A virtual TCCC (Tactical Combat Casualty Care) Card (custom-developed Qt application), and
- A human performance data recording and review system (Aptima's Performance Measurement Engine).

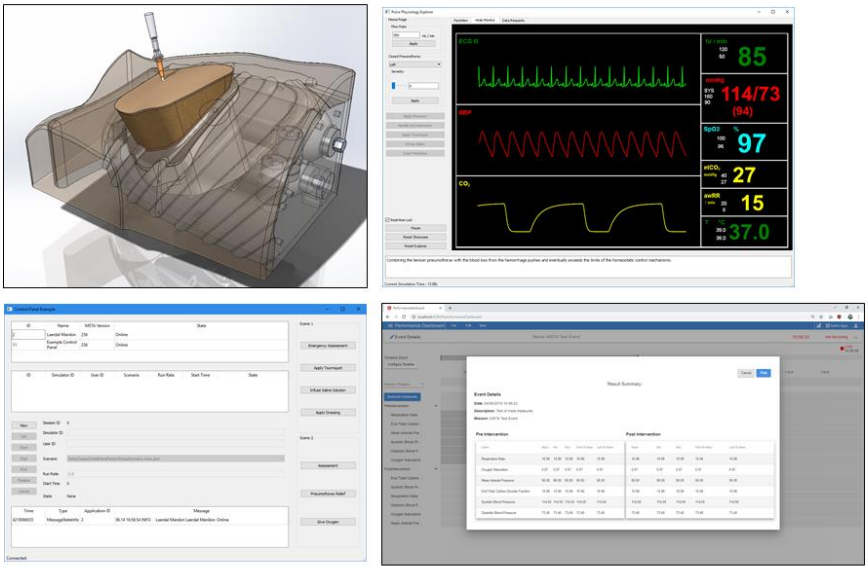

Figure 2. Images depicting several of the simulation and training technologies that are being networked together during the reference implementation. These include: the MGH tension pneumothorax simulator (upper left), the Kitware PULSE physiology engine (upper right), the independent simulation control panel (lower left), and Aptima's Performance Measurement Engine (lower right).

The June 2019 technical demonstration involves the following scenario: the Warfighter receives a GSW to the leg, thereby requiring the Combat Medic to control the bleeding by applying a tourniquet. After the Warfighter has been stabilized, he is then transported via aeromedical ambulance to a Role 2 care facility. At altitude, the Warfighter develops a tension pneumothorax, which is resolved by performing a needle decompression, thereby releasing the pressure from the pleural cavity and allowing to heart and lung function to return to normal. During the technical demonstration, the control panel will initiate the scenario, mark the exact times when the tourniquet and needle decompression procedures occur, and signal the scenario's conclusion. The PULSE physiology engine will then control the symptomology of both manikins. Finally, the PM Engine will visualize the patient's physiology parameters in real time, and provide summary statistics at the scenario's conclusion.

The design allows the whole chain to be exercised either as part of a large training event, or in separate steps where each part of the care chain can be trained independent of the others. In the latter situation, the simulation starts from the saved state of the prior simulation; learners care for the patient as per normal and then save their state to a library. That saved state can then be accessed and used as the starting point further along the care chain.

The design meets the current medical simulation needs and is flexible to adapt to changing technology both in the underlying transport layers and in the higher level application layers. Each of these has and continues to undergo change.

Network Transportation Layer: Three years ago, when the architecture was first being conceptualized, HLA appeared to be the choice as a transport layer. As we progressed, however, it became clear the newer DDS (Data Distribution Layer) offered some significant advantages in flexibility and openness, which represent two of our overarching design goals. The resultant protocol, however, is adaptable enough to ride on either of these, or even on DIS or some yet-to-be-defined new protocol. By restricting the features required of the transport to just those essential for medical training simulation, the protocol is able to avoid some of the complexity that a general purpose solution for real-time simulation would need. For example, medical/surgical simulation does not require millisecond coordination of timing among simulators to ensure that a fair, consistent, and accurate outcome is achieved. In addition, human patients are independent of each other, while a medical simulation architecture needs to support team training and learners need to interact with each other, the simulated patients don't. The architecture can provide a thin API which decouples applications from the chosen transport layer.

A second factor in the API design is the avoidance of encoding medical information directly into the API. the architecture does not have embedded within its code a notion of what information will be transported. That comes exclusively from participants connected to the network. Participants are free to add, delete, or change what information they present. This gives compliant simulations the ability to evolve with the changing participants' needs without requiring that every participant resynchronize on a new definition of the network data model. As new data is added, either new downstream participants will be added to consume the new data, or existing participants will use the data discovery mechanisms to detect and respond to the newly-available data. As a simple example of how this might work, consider a network which includes a physiology engine producing a heartbeat to drive a manikin, and the manikin reporting chest compression as a metric. An AAR tool connected to the network might originally pick up the heartbeat data and the chest compression data for an instructor to use in debriefing the learner. If the physiology engine is enhanced to now also product blood pressure information, the AAR tool, which is using the data discovery, can without any changes to its code or operation automatically detect the new data and display it along with the previous data. If the manikin is capable of using that information, it could similarly pick up the data.

The data layer draws on the National Institute of Health's (NIH) Unified Medical Language System (UMLS). In particular, one of the constituent libraries, SNOMED CT, is used to provide a common set of terms to describe what is happening medically within a simulation. ULMS and SNOMED CT have grown and matured over 30 years to provide international standards for expression of medical information. In addition to the well-founded basics for representation of information, they are actively maintained to stay current with medical developments. SNOMED has the advantage of combining an exhaustive list of what they call concepts with the ability to create expressions for further refinement of these concepts. In SNOMED, even terms are defined by using this expression syntax which lets computerized systems reason about the terms even if those terms are not known specifically. But while SNOMED CT provides a broad set of terms, it was not designed with trauma and simulation specifically in mind. As a result, there are concepts which arise in simulation that are not naturally encoded using SNOMED CT directly. In the demo scenario we were able to convey the pressure of a specific side of the lung by using a custom expression.
With the help of SNOMEDs syntax we were able to address specific physiological structures and convey information about them. On the other hand for example while there is the concept of a tourniquet, a junctional tourniquet is not available in the current identifier set and while we could extend the definition of a tourniquet introducing a new concept for a junctional tourniquet would be more appropriate.
In future efforts extensions which are compatible with SNOMED CT will be developed, those will be supplied with the source code. A parallel effort is required to then take new terms to the maintainers of SNOMED CT for potential inclusion in the standard.

FUTURE PLANS

There are several potential directions that future development could take. Eventually, a fully-developed architecture will encompass all of these areas and more. These are mutually-supportive directions that can be taken to build on the base that has already been established.

- *Linkage to Synthetic Training Environment (STE)*. Linking with STE would involve being able to send and receive information from STE and even use STE resources to develop joint tactical/medical simulations, such that the medical activities are inherently coupled with STE and thus any exercise developed in STE has the capability to send information to the medical simulators; loading them with casualty data and launching them as appropriate. Similarly, the results of the these simulations would be reflected back to the STE so that resources consumed, or personnel taken out of action, can leave the field as the outcome dictates.
- *Development of a portal, library, and simulation control features*. The current implementation is a library devoid of user interface other than those used to test and demonstrate its functionality. Moving forward, any system or systems built on this architecture will share certain attributes. One of these is the need to access a library of simulations. These simulations will arise in multiple ways. Some will be instructor-designed simulations that teach specific skills. Others will come from the STE environment arising from simulation events there. Still others will be results of simulations which are moving along the continuum of care being handed-off from one simulator to the next.
- *Encoding of tactical medical training information in the lexicon*. This lexicon draws from the SNOMED CT ontology, which is mapped to the National Institutes of Health (NIH) Unified Medical Language System (UMLS), as it basis for medical knowledge. This international standard for the encoding and exchange of medical data comprises over 1 million biomedical concepts and 5 million concept names, all of which stem from the over 100 incorporated controlled vocabularies and classification systems. To ensure that the simulations involved all choose the same terms to encode casualty care events, a down selection from the million concepts to an "accepted" set must be made.
- *Introduction of representational state transfer (REST) style interface to the protocol*. RESTful interfaces are an enabling feature for Web services. Web services in turn make data access and control information from the rest of the network available to web applications and mobile apps. RESTful interfaces have a number of defining characteristics. Pertinent to the architecture are the *stateless interface*, such that requests provide all the information required to perform the requested service; *layerability*, so that the clients do not know if they are directly connected to the service or if there are intermediate layers that provide features such as scalability and security; and *uniform interface*, that provides a simple and complete method for accessing and manipulating the service. Adding a RESTful API will provide a standard mechanism for any web service or tool to access and manipulate medical simulation.
- *Automated verification/certification tools*. To ensure that all parts of the ecosystem interact with each other as planned, and that information flows seamlessly from point to point, it is vital to have verification tools which can automatically detect that the encoding of events are correct and that protocol usage is also correct. These checks will range from correct negotiation of the communication protocols and interactions to higher-level checks that medical terms and concepts are properly expressed in the lexicon.
- *Adapting legacy training systems/devices to be compliant*. Creating adaptors for legacy training systems which have a planned operational life of more than a few years will speed the transition from the currently deployed non-integrated training systems to the planned vision of the future. Starting with major elements, such as the most commonly used patient manikins, will give the developing standard a jump-start in this direction.

TECHNICAL CHALLENGES AND LESSONS LEARNED

In the following paragraphs, we review several technical challenges that we experienced, and present some guidelines that may benefit other system developers, researchers, and instructors who are attempting to conduct their own networked medical simulation-based training, regardless of whether or not they choose to adopt the architecture.

- *Guideline #1: Be mindful of the past, but design with the future in mind.* During the course of our background research, we became aware of two small-scale demonstration efforts that are attempting to network a handful of

patient manikins using the High Level Architecture (HLA) communications protocol (Petty & Windyga, 1999). Our understanding is that these two efforts specifically chose HLA because it is widely used among the tactical simulation community, thereby paving the way for potential joint tactical-medical simulation exercises. At one point, we also considered using HLA for the same reason. However, we eventually abandoned the idea in favor of a middleware-based approach that leverages the Data Distribution Service (DDS) protocol instead. HLA's open source infrastructure did not look well supported to us, thereby imposing an additional cost on the simulation vendors to incorporate it into their designs. Another limitation is that different vendor libraries don't align perfectly with one another. While this can be resolved using network bridges, doing so represents yet another cost to simulation developers or host facilities that seek to work within the architecture. Ultimately we chose a middleware-based approach that isolates the upper layers from the actual DDS protocol. This way, if there is ever a decision to use a different protocol moving forward, only the supporting parts need to be re-written. We feel that this approach provides continuity with the past and flexibility toward the future.

- *Guideline #2: Never underestimate the power of language as a potential unifying (or dividing) force.* One of the advantages of working in the tactical domain is that there are clearly-worded terms, definitions, and standards. For example, Mission Essential Task Lists (METLs) clearly identify the tasks that a unit must be able to perform; similarly, Tactics, Techniques, and Procedures (TTPs) describe expected unit behaviors, both in the abstract and in specific situational contexts. Unfortunately, this level of standardization is simply not present in the medical and healthcare domains. Instead, there are countless different ways to define the human anatomy, physiological parameters, medical procedures, medical tools, and the like. The lack of standardization applies to the medical modeling and simulation community as well. For example, many human patient manikins only represent the surface symptoms such as bleeding and heart attacks, while others actually model the underlying physiological systems that produce these symptoms. The closest thing that we have found to a common vocabulary/lexicon is the SNOMED CT database of clinical terminology (SNOMED International, 2019). SNOMED CT is an international, copyright-free ontology of several million clinical terms, each of which is also mapped to the NIH Unified Medical Language System (UMLS). To facilitate understanding among simulation designers and other potential users, all of the concepts referenced in the architecture have been mapped to their respective SNOMED codes. While imperfect, this is one initial step at attempting to provide some sense of unity to this otherwise disjointed domain.
- *Guideline #3: Leverage agile development methods to test out your ideas, collect feedback, and revise as appropriate.* When we first started this effort, the research team developed a number of knowledge products, the MRFCs, that sought to define critical components of the system architecture and how it might be used in practice. So for example, we developed MRFCs regarding: levels of patient representation, system interoperability, simulation control, human physiology, conceptual and data models, After Action Reviews (AARs), and human performance measurement, among others. One of the challenges is that while each knowledge product was internally consistent, it was difficult to align them because they were all written in the abstract. Therefore, we organized a reference implementation in which a single constructive patient simulator would generate the physiological parameters that drive the physical representation of two different patient manikins. Moreover, while a combat medic performed two different medical procedures (applying a tourniquet to control a hemorrhage, and a needle decompression to reduce the pressure within the pleural cavity), we are visualizing the physiological parameters in real time on a large-screen dashboard. While at some level this technical demonstration does divert some resources away from the ongoing architectural design efforts, we believe that it is critical to the project's long-term success because it will help us to better understand how the various components work (or don't work) together. The lessons learned in that practical work will be integrated back into the MRFCs. A similar agile development approach is being used in the Total Learning Architecture community (P. S. Gallagher et al., 2017). We believe that such demonstrations can also be a useful vehicle for socializing the concept to interested parties within government, industry, and academia.
- *Guideline #4: Think beyond a single training event (followed by a debriefing), and instead design the system with lifelong learning in mind.* One of the key design concepts is that training and human performance assessment are "two sides of the same coin" (National Research Council, 1996). Simply put, if you can't measure the learner's proficiency, you will never know if the learner has mastered the instructional content. Therefore, one of our design considerations was to leverage existing standards of human performance measurement, including Human Performance Measurement Language (Wheeler Atkinson et al., 2017) and the Experience Application Program Interface (P. S. Gallagher et al., 2017), respectively. Our long-term vision is to support end-users by measuring learner proficiency on every training trial, and archiving these measures in a LRS for long-term tracking and trending purposes. Aggregate learner records can then be mined for a variety of purposes, such as personalized learning, modeling skill acquisition and decay rates, calculating training Return on Investment (ROI), and making

data-based simulator acquisition decisions. We believe that this emphasis on "big data" and lifelong learning represents the next evolution of learning science, and have intentionally positioned the architecture to be ready for it.

- *Guideline #5: To encourage use, minimize barriers to entry.* One of our overarching design philosophies is to purposely minimizing barriers to entry. That is one of the reasons why we have heavily leveraged open source technologies (DDS, xAPI, HPML) and knowledge products (SNOMED CT), whenever possible. We realize that nudging the medical simulation community, away from its current reliance on closed architectures and proprietary systems, towards an open architecture will take many years to accomplish. However, there are other barriers to entry besides technology and cost. Another key barrier is time. While we look very highly on the good work performed by standards organizations such as the Simulator Interoperability and Standards Office (SISO), we also realize that the development of standards can take years because it requires building consensus among government, industry, and academia. Looking toward the future, we do envision some form of certification process for a technology to be considered "compliant," but we envision this process being more akin to the one used by DDS, which relies primarily on automated system compliance tests. Simply put, the medical modeling and simulation community is nearly 10 years behind the tactical community. In order to catch up, we will need to streamline the process whenever possible, and learn from our mistakes as we go.

ACKNOWLEDGEMENTS

This work was supported by the Program Executive Office for Simulation Training, Research, and Instrumentation and the (PEO-STRI) and the Joint Program Manager for Medical Modeling and Simulation (JPM-MMS) under Contract No. W900KK-17-C-0041 to SimQuest, LLC. The views, opinions and/or findings contained in this report are those of the authors and should not be construed as an official Department of Defense position, policy or decision unless so designated by other documentation.

REFERENCES

Advanced Modular Manikin. (2015). Advanced Modular Manikin™ Phase 2 Standards: A Common Platform Towards a New Paradigm in Health Care Simulation. Retrieved from https://www.advancedmodularmanikin.com/

Bray, A., Webb, J. B., Enquobahrie, A., Vicory, J., Heneghan, J., Hubal, R., . . . Clipp, R. B. (2019). Pulse physiology engine: An open-source software platform for computational modeling of human medical simulation. *SN Comprehensive Clinical Medicine, 1*, 362-377.

Gallagher, A. G., & O'Sullivan, G. C. (2011). Simulation in and for medicine: Where next? In *Fundamentals of Surgical Simulation* (pp. 297-323). New York: Springer.

Gallagher, P. S., Folsom-Kovarik, J. T., Schatz, S., Barr, A., & Turkaly, S. (2017). Total learning architecture development: A design-based research approach. In *Proceedings of the 2017 Interservice/Industry Training, Simulation, and Education Conference. Paper No. 17117*. Arlington, VA: National Training and Simulation Association.

Kelliher, T. P., Lenoir, J., Novotny, P., & Scheirich, H. (2014). Open Surgical Simulation (OSS): A community resource. In J. D. Westwood (Ed.), *Medicine Meets Virtual Reality (MMVR) 21* (pp. 197-203). Amsterdam: IOS Press.

National Research Council. (1996). *National science education standards*. Washington, DC: National Academies Press.

Petty, M. D., & Windyga, P. S. (1999). A high level architecture-based medical simulation system. *Simulation, 73*(5), 281-287.

Scheirich, H. (2017). *Medical Simulation Training Architecture (MSTA) Domain Model*. Annapolis ,MD: SimQuest, LLC.

SNOMED International. (2019). SNOMED CT International Edition 20190131. Retrieved from https://browser.ihtsdotools.org

Wheeler Atkinson, B., Tindall, M., Killilea, J., Tolland, M., & Dean, C. (2017). Standardizing performance measurement for ease of data analytics. In *Proceedings of the 2017 Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC)* (pp. 2187-2195). Arlington, VA: National Training and Simulation Association.