Proceedings of the 2008 Winter Simulation Conference
S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, J. W. Fowler eds.

PANEL DISCUSSION: WHAT MAKES GOOD RESEARCH IN MODELING AND SIMULATION: ASSESSING THE QUALITY, SUCCESS, AND UTILITY OF M&S RESEARCH

Jeffrey S. Smith
(Panel Chair)
Industrial & Systems Engineering
3301 Shelby Center for Technology
Auburn University, AL 36849, USA

John A. "Drew" Hamilton, Jr.

Computer Science and Software Engineering
3101 Shelby Center for Technology
Auburn University
Auburn University, AL 36849, USA

Richard E. Nance

Orca Computer, Inc.
1800 Kraft Drive, Suite 111
Blacksburg, VA 24060, USA

Barry L. Nelson

Industrial Engineering and Management Sciences
2145 Sheridan Road, Room C250
Northwestern University
Evanston, IL 60208, USA

George F. Riley

Computer Engineering
266 Ferst Drive
Georgia Institute of Technology
Atlanta, GA 30332, USA

Lee W. Schruben

Industrial Engineering and Operations Research
4131 Etcheverry Hall
University of California-Berkeley
Berkeley, CA 94720, USA

ABSTRACT

This paper presents the "position papers" contributed by the participants of a panel at the 2008 Winter Simulation Conference. As the paper pre-dates the actual panel, the purpose of the paper is to provide some background information about the views of the individual panelists prior to the actual panel. Each panelist was asked to submit a position paper addressing the general question of "What makes good Modeling and Simulation research?" This paper presents a summary of these position papers along with an introduction and conclusion aimed at identifying the common themes to setup the conference panel.

1 PANEL DESCRIPTION

This panel was initiated by a discussion with Levent Yilmaz (Auburn University) regarding the general field of Modeling and Simulation (M&S) and, more specifically, the research areas in the field. We sent the following question to several researchers involved in various M&S-related research activities:

"What makes good and creative M&S Research?"

We asked the recipients to respond with a position paper discussing this general question. Interestingly, there was a very natural division in the responses and this division led to two separate panels, each addressing a slightly different question. This panel addresses the general question of what makes good research and the panel will be comprised of the following individuals:

- Drew Hamilton, Auburn University
- Richard Nance, Orca Computer, Inc.
- Barry Nelson, Northwestern University

- George Riley, Georgia Institute of Technology
- Lee Schruben, University of California, Berkeley

The companion panel (organized by Levent Yilmz) focuses on M&S as a stand-alone academic discipline and is entitled "What Makes Good Research in Modeling and Simulation? Sustaining the Growth and Vitality of the M& S Discipline". Both of these panels are directly related to the paper entitled "Prelude to the Panel on What Makes Good Research in Modeling & Simulation" that provides an overall framework for both of the panels.

2 INTRODUCTION

The panel participants' submissions were quite interesting and emphasized several common themes. As one would expect, one of the common themes among the position papers relates to the general issue of "Application vs. Theory" or "general purpose vs. special purpose" in simulation-related research. In this regard, Richard Nance conjectures that good M&S research should be foundational – the objectives should be to contribute to increased knowledge. Barry Nelson addresses this question along two dimensions – pay-as-you-go research vs. investment research, in general; and general purpose vs. special purpose research in terms of research methodologies. George Riley addresses these two basic types of M&S research completely separately.

Lee Schruben interjects in interesting twist by addressing *successful* research rather than "good" research, rightly pointing out that research quality (i.e. "goodness") is highly subjective, whereas "successful" is at least, somewhat less subjective. He describes the evolution of the definition of successful M&S research over the years and interplay between academic researchers and simulation software developers. The issues of commercial simulation software and simulation software development are also addressed to differing degrees by Drew Hamilton and Barry Nelson.

Richard Nance and Barry Nelson both discuss some of the difficulties in obtaining research funding and publishing the in M&S area. While both offer opinions on these issues, these questions are very much open for debate and should be topics for significant discussion of the panel.

At the risk of creating the perception of disagreement to liven up the panel discussion, there does seem to be some differing opinion when it comes to research focusing on strict applications of simulation modeling. Lee Shruben indicates that using well-known (simulation) methodologies to significant improve the design or operation of a real system is a valuable research contribution. I tend to agree with this assessment. Richard Nance's position paper seems to disagree with this and George Riley's paper seems to indicate that more is needed (than the application). While Barry Nelson's submission does not dispute the "goodness" of applications, he does discuss the difficulties associated with publishing this type of work. Certainly difficulty in publishing can be viewed by the academic community as an indicator of the lack of recognition that application-only papers provide significant contributions. However, Lee Schruben clearly opines that publication is not necessary for M&S-related work to be considered successful.

Drew Hamilton discusses the multidisciplinary nature of M&S and stresses that thorough understanding of the software engineering, statistical, and modeling aspects of simulation are critical in order to conduct successful research. He also makes the interesting point that due to its multidisciplinary nature, successful M&S research often involves teams rather than individuals.

3 PARTICIPANTS' SYNOPSES

The following sections present the individual panel participants position papers which were submitted in response to the original question of "What makes good and creative M&S research?".

3.1 Drew Hamilton

Creative M&S Research requires a thorough understanding of the multiple disciplines which under pin modeling and simulation, specifically, software engineering, statistics and systems analysis as well as access to significant domain knowledge. For example to use M&S to study missile defense, requires expertise in physics in order to design the appropriate experiments.

What of research into M&S techniques themselves? A successful researcher needs depth as well as breadth and typically that depth can only be achieved by experience in developing and employing non-trivial simulations. Critical to successful research in simulation is having enough background in software development to understand the underlying execution of the simulation as well as a strong enough background in statistical analysis to understand the results. One major implication of this is that successful M&S research efforts are likely to be multidisciplinary team efforts.

We are seeing some simulation degree programs beginning to emerge. In many cases these are existing science or engineering degrees with some simulation courses grafted on to them. The detailed multi-disciplinary education requirements make for some lengthy prerequisite trees from many different departments and therefore make credible undergraduate modeling & simulation degree programs challenging to design.

3.2 Richard Nance

The interrogative title can be interpreted to produce a broad, general response or a very specific answer. The general response could take the form of characteristics of research topics or proposed studies that are indicative of worthy subjects and innovative or creative approaches. The more narrow answer could identify specific topics that represent challenging and thought-provoking inquiries. While both interpretations can provide useful guidance to the modeling and simulation (M&S) community, my response focuses on the general without totally ignoring the more specific.

Good research in M&S should first and foremost be foundational; i.e. the objectives should be to contribute to increased knowledge that is fundamental to modeling and simulation. A test of "fundamental" is to ask if the proposed work is confined to a specific application domain. A positive answer categorizes the work as a M&S application; and, despite the importance of the application domain, research support should be viewed as directed toward advancement of knowledge in that domain. To avoid misinterpretation, my argument is not that application domain research is less important for that application area, but it does not advance the state of knowledge in M&S. For all too long the scientific funding sources have not grasped the reality that studies providing answers to questions in physics, chemistry, material sciences, and biology do not contribute to answering key questions in systems modeling, model verification and validation, representational approaches to abstraction resolution, statistical analysis of model output, and a host of others. The recent report from the government Blue Ribbon Panel on Simulation-Based Engineering and Science: http://www.nsf.gov/pubs/reports/sbes_final_report.pdf is a patent example of how application domain specialists appear to be guiding major research decisions that portend to improve M&S knowledge, technology and practice.

Why does the M&S community, or the discrete event simulation community, lack recognition in some government funding agencies? Can the definition of what constitutes good, creative research assist in overcoming this identity problem? How can the uninformed be educated? Hopefully, the perceptions and opinions of others in this panel can assist in the formulation of convincing arguments for external observers.

3.3 Barry Nelson

As Editor-in-Chief of *Naval Research Logistics*, I distinguish between two types of research: *investment research* and *pay-as-you-go research*. Pay as you go is mostly what we do in engineering: The basic question here is "Is there someone who would really be interested in using this work if it were available in an easy-to-apply form?" Investment research advances the field in a meaningful, but not necessarily immediately useful, way by deepening our understanding. As academics, we tend to take more chances on investment research because it is hard to predict what will be useful.

In simulation analysis methodology research, I distinguish between *general purpose* and *special purpose* research. For general purpose research an acid test is would it add value if it were included in a commercially available simulation package such as Arena, ProModel, Automod, etc. For special purpose research the test is does it make solvable (or provide a much better solutions to) a problem that was not solvable (or solved well) before and the problem is one that someone is really interested in solving.

From a research reviewing standpoint, it is much easier to assess *mathematical novelty* than it is to assess *usefulness*. This makes reviewing "applied" papers more difficult. Papers that crack a tough mathematical nut are often worth publishing based on the investment argument, but just having tough math is not enough.

There are significant reasons why it is difficult to publish empirically based research: We don't have a good community sense of what constitutes a thorough/useful/controlled/compelling empirical study. We are overly worried about the exceptions (math can prove truth for ALL cases, while generalizing from an experiment requires a certain leap of faith). Simulation can be thrown at virtually any problem, so it really is tough to construct empirical studies that cover the waterfront. GIVEN THAT SIMULATION IS AN EMPIRICAL TECHNIQUE, IT IS AMAZING WE HAVE THIS PROBLEM.

Another significant difficulty in publishing M&S related research is that Referees, Area Editors, and Chief Editors often have a general approach of looking for something wrong rather than looking for something right.

3.4 George Riley

In order to address the posed question, we must first differentiate between research into basic modeling and simulation techniques and research that uses modeling and simulation to make some claims about a particular problem domain. We will address each of these issues separately.

Good research in basic modeling and simulation techniques should present some new methodology that can be universally applied to all, or nearly all, application domains. Further, the research should describe clearly the strengths and weaknesses of the approach, and present some conclusive evidence that the technique preforms as claimed. For example, suppose the research is proposing a new method for time synchronization in a conservative simulation system. The research should explain and demonstrate:

1. What is fundamentally different between the new method and all prior approaches for time synchronization? A small tweak on an existing algorithm or approach is not good research, even if

empirical measurements show a decent performance improvement.
2. Under what conditions were the performance results measured?
3. Is the approach applicable to a variety of platforms and environments, such as tightly–coupled shared memory systems, networks of workstations on low–cost networks, and supercomputer platforms with a variety of high– speed interconnects and shared memory interfaces?
4. Did the measured performance include all overhead associated with the approach?
5. Under what conditions (if any) will the proposed method perform poorly, perhaps worse than existing known methods. This is extremely important for the reader, and is often lacking in published works. All too often, a new method is demonstrated to exhibit considerable improvement under conditions carefully chosen by the researcher.
6. Explain clearly what about the new method causes the improved performance. Simply presenting empirical results is not sufficient to convey to the reader the fundamental reasons the approach works better.

Research that uses modeling and simulation methodology to demonstrate the effectiveness of some approach in some problem domain also has a number of common problems, not unlike those mentioned above. Suppose that a researcher has developed a new method for multi–hop wireless routing in a mobile ac–hoc network. A good research paper in this scenario should:

1. If possible, give an analytical discussion and analysis of the method before presenting simulation–based results. Try to explain to the reader the fundamental basis for the improvement in performance.
2. Under what scenarios was the approach simulated? Is there something about the particular scenario that highlights the new method's performance improvement?
3. What simulator was used and what parameters (both default and non–default) were used in the simulator. This will allow others to re–create the simulation results and further analyze the approach if they so desire.
4. Is there a particular scenario that will demonstrate any weaknesses in the proposed method? It is not necessary for any improved method to outperform all others in all possible cases. Clearly state which scenarios, if any, would lead to worse performance that other methods and explain why this occurred.
5. Any anomalies in the measured must be understood and explained. Unexplained dips or valleys in performance graphs, or numbers that go down when they should go up are often clear indications of flaws in the simulation methodology. Few things will turn off a reader of a paper faster than statements to the effect of "we do not understand why X occurred".

To summarize, good research in modeling and simulation is in many ways the same as good research in other domains. You must understand and explain clearly the environment in which the experiments were performed, include all experimental results even if they do not always perform as expected, and be able to convince the reader that the approach is indeed valid and the reported results are meaningful.

3.5 Lee Schruben

Not everyone will agree on whether a particular modeling and simulation research contribution is good or not, nor should we. If we ever reach complete consensus, we will stop learning. Most researchers find papers in our own area of research more interesting than those on other topics. The more flattering the references to our own work, the higher our opinions of a research paper tend to be, and vice versa. Quality is a largely a matter of individual taste. Recognizing the subjective nature of any discussion of quality, I will try to focus instead on what I believe makes modeling and simulation research *successful*.

The definition of success in simulation and modeling research has evolved in an interesting manner. Early on, the success of academic simulation methodology researchers, to a large degree, was dependent on the judgment of simulation software developers. Research results that were automated and included as part of a commercial simulation language were considered successful. This was, and still is, the only way that the results of simulation methodology research can have a broad impact. Also, simulation modeling and application papers submitted to the top academic journals were mostly refereed by analysis methodology researchers. Acceptance of research for publication was mostly determined by academics, many of whom valued mathematical rigor over modeling innovation or study impact. What was done, or why, was not nearly as important to successfully getting a paper published as how it was done. We had simulation software developers judging the success of academic researchers, and academic researchers in turn judging the success of simulation applications. This value chain implied that success across a broad range of simulation modeling and analysis research (as well as, by the transitive property, simulation applications) came primarily from the commercial simulation software developers. This situation, in my opinion, has improved greatly with more open-source simulation software available and outlets for research communication beyond the printed word. For research to be successful, people have to know about it, and care. It has to be motivated, explained, and communicated. Publication helps, but this alone does not ensure success. Sometimes publishing papers appears to be the final result of a research project, if that were true, then I would not feel the research was much of a success. Successful research needs to be understood, appreciated, and used.

Nor is journal publication necessary for research success. Successful research can be communicated through conversations, group discussions, public panels, blogs – many valuable research contributions have been acknowledged as "private communications". When evaluating the success of a piece of research, I like to judge it by looking at it through both "I's", Innovation and Impact. Successful letters I have written supporting tenure and promotion always include one or both of these "I's". If I cannot do that, I do not write the letter. Successful research does not have to have both innovation and impact, but it should be able to make a supportable claim to at least one of these. Of the two, I feel impact is the more important.

Innovative research is not merely applying conventional methodologies to domains for which they really are not appropriate. This leads only to methodology bloat (Petri Nets come to mind here). A completely new approach in a new domain is preferred to shoehorning old shoes onto new feet.

Innovative research is doing something in a new way. Whether it has impact or not depends on the "something". Modeling and simulation by its nature has a conceptual, real, proposed, or at least hypothetical system as its context. Simulation research done out of a systems context is pure art. That does not mean it is uninteresting. Innovative art can be fascinating and entertaining, and includes some excellent mathematical research. But this should not be confused with successful modeling and simulation research. For that, there needs to be a system that is being simulated. There may not even be an immediate application, but successful research needs to address important questions and problems about systems.

Successful simulation and modeling research can have a huge impact without inventing new methodologies. Research that takes current, well known, methodologies and uses them appropriately to make significant improvements in the design or operations of a real system is a valuable simulation and modeling research contribution. Solid examples of simulation success stories in the open literature are inspirational to other simulation practitioners as well as to simulation modeling and analysis methodology researchers. They are critical to our field's growth, indeed, its survival.

My personal opinion is that successful modeling and simulation research should not be constrained by current commercial simulation software designs or limited by current computational capabilities, but rather motivates extensions to these. Just because something cannot be done easily now, does not mean we should not start trying to figure out how to do it. Successful modeling and simulation research should explicitly acknowledge that dynamic systems are non-stationary and good simulation models are often stochastic. What simulation is ultimately trying to do is forecast how a system will behave in different possible future environments. Modeling a current system using current (and necessarily old) information and data offers a useful framework for retrospection, but with simulation we can do much better than "what if" hindsight. I feel that successful modeling and simulation research is abstract (general) but applicable (relevant) to solving important problems in engineering, science, and society. Finally, the most successful modeling and simulation research should provide new intuition and insights making a qualitative contribution to human knowledge. This is probably more important and lasting than a new quantitative analysis algorithm or modeling methodology.

4 SUMMARY

The position papers submitted by the panelists have raised some very interesting and important points regarding the assessment of quality, usefulness, and utility in M&S related research. While there are several points of agreement, there are also issues for which there does not appear to be universal agreement.

AUTHOR BIOGRAPHIES

JEFFREY S. SMITH is a professor in the Industrial & Systems Engineering department at Auburn University. Prior to joining Auburn, Dr. Smith was an associate professor in the Industrial Engineering Department at Texas A&M University. He received the B.S. in Industrial Engineering from Auburn University in 1986 and the M.S. and Ph.D. degrees in Industrial Engineering from Penn State University in 1990 and 1992, respectively. In addition to his academic positions, Dr. Smith has held industrial positions at Electronic Data Systems and Philip Morris U.S.A. Dr. Smith is an active member of IIE and INFORMS. His email and web addresses are <jsmith@auburn.edu> and <http://sim.auburn.edu/~jsmith>

JOHN A. "DREW" HAMILTON JR. is an associate professor of computer science and software engineering at Auburn University and director of Auburn University's Information Assurance Center. Prior to his retirement from the U.S. Army, he served as the first director of the Joint Forces Program Office and on the electrical engineering and computer science faculty of the U.S. Military Academy, as well as chief of the Ada Joint Program Office. He has a bachelor of arts in journalism from Texas Tech University, masters degrees in systems management from the

University of Southern California and in computer science from Vanderbilt University, and a doctorate in computer science from Texas A&M University. He is currently President of SCS, the Society for Modeling & Simulation, International. His email and web addresses are <hamilton@auburn.edu> and <http://www.eng.auburn.edu/~hamilton/>

RICHARD E. NANCE is an Emeritus Professor of Computer Science at Virginia Tech and Chief Scientist, Orca Computer, Inc. He received B.S. and M.S. degrees from N.C. State University in 1962 and 1966, and the Ph.D. degree from Purdue University in 1968. He has served on the faculties of Southern Methodist University and Virginia Tech, where he was Department Head of Computer Science, 1973-1979. Dr. Nance held research appointments at the Naval Surface Weapons Center (1979-80) and at the Imperial College of Science and Technology (UK). Within ACM, he has chaired two special interest groups: Information Retrieval (SIGIR), 1970-71 and Simulation (SIGSIM), 1983-85. He is the author of over 150 papers on discrete event simulation, performance modeling and evaluation, computer networks, and software engineering. Dr. Nance has held several editorial positions and was the founding Editor-in-Chief of the ACM Transactions on Modeling and Computer Simulation, 1990-1995. He served as Program Chair for the 1990 Winter Simulation Conference. Dr. Nance has received several awards for his editorial and professional contributions, most recently the Lifetime Professional Achievement Award from the INFORMS Simulation Society. In 1996, he was named an ACM Fellow. In 2006 he was designated by the faculty as a Distinguished Alumnus of the Edward P. Fitts Department of Industrial & Systems Engineering, North Carolina State University, He is a member of Sigma Xi, Alpha Pi Mu, Upsilon Pi Epsilon, ACM, and INFORMS. His e-mail address is <nance@vt.edu>.

BARRY L. NELSON is the Charles Deering McCormick Professor of Industrial Engineering and Management Sciences at Northwestern University. His research centers on the design and analysis of computer simulation experiments on models of stochastic systems. He has published numerous papers and two books. He has served the profession as the Editor in Chief of *Naval Research Logistics*, the Simulation Area Editor of *Operations Research*, and President of the INFORMS (then TIMS) College on Simulation. He has held many positions for the Winter Simulation Conference, including Program Chair in 1997 and current membership on the Board of Directors. His e-mail and web addresses are <nelsonb@northwestern.edu> and <www.iems.northwestern.edu/~nelsonb/>.

GEORGE F. RILEY is an Associate Professor of Electrical and Computer Engineering at the Georgia Institute of Technology. He received his Ph.D. in computer science from the Georgia Institute of Technology, College of Computing, in August 2001. His research interests are large--scale simulation using distributed simulation methods. He is the developer of Parallel/Distributed ns2 (pdns), the Georgia Tech Network Simulator (GTNetS), and is co-PI on the ns3 development effort. Before turning to a career in academia, Dr. Riley spent 20 years as an independent consultant and business owner, primarily working at the Air Force Eastern Test Range, developing and deploying systems for real-time missile launch support. His email and web addresses are <riley@ece.gatech.edu> and < http://www.ece.gatech.edu/faculty-staff/fac_profiles/bio.php?id=86>.

LEE W. SCHRUBEN is a Chancellor's Professor and past Chairman in the Department of Industrial Engineering and Operations Research at the University of California at Berkeley. Prior to joining the Berkeley faculty, he was on the Operations Research and Industrial Engineering faculty at Cornell where he held the Schultz Endowed Professorship in Engineering. He received his PhD from Yale and is a Fellow of the Institute for Management Science and Operations Research. Professor Schruben's research interests are in simulation modeling and analysis methodologies with a broad range of applications, most recently focusing on biopharmaceutical production and supply chains. His email and web addresses are <lees@berkeley.edu> and <http://www.ieor.berkeley.edu/People/Faculty/schruben.htm>