

AVOID FAILURES! TESTED SUCCESS TIPS FOR SIMULATION PROJECT EXCELLENCE

David T. Sturrock

Simio LLC
504 Beaver St
Sewickley, PA 15143, USA

ABSTRACT

How can you make your projects successful? Modeling can certainly be fun, but it can also be quite challenging. With the new demands of Smart Factories, Digital Twins, and Digital Transformation, the challenges multiply. You want your first and every project to be successful, so you can justify continued work. Unfortunately, a simulation project is much more than simply building a model -- the skills required for success go well beyond knowing a particular simulation tool. A 35-year veteran who has done hundreds of successful projects shares important insights to enable project success. He also shares some cautions and tips to help avoid common traps leading to failure and frustration.

1 INTRODUCTION

This paper discusses many aspects of modeling that are often missed by new and aspiring simulationists. In particular, tips and advice are provided to help you avoid some common traps and help ensure that your first project is successful. The first four topics dealing with defining project objectives, understanding the system, creating a functional specification, and managing the project are often given inadequate attention by beginning modelers. The latter sections dealing with building, verifying, validating, and presenting the model offer some insight into some proven approaches.

Market demands are steadily increasing in volume and complexity. In particular the *Industry 4.0* and *Smart Factory* trends in automation and data exchange have caused a growing demand for a *Digital Twin* that often goes well beyond typical model demands. It is important to consider the topics of data-driven modeling and even data-generated models in your project objectives and tool selection.

2 DEFINE THE OBJECTIVES OF THE STUDY

When you first think about conducting a simulation study, one of the earliest things to consider is the project objectives. Why does someone want to simulate this system and what do they expect to get out of it? To be more specific, you must determine who your stakeholders are and how they define success.

2.1 Who Are Your Stakeholders?

A *stakeholder* is someone who has an interest in the outcome of the project, someone who cares. It seems like “Who are your stakeholders?” has an obvious answer – your manager or your client. But if you explore why someone would want to see the results of this study, you will probably discover additional stakeholders. Are you trying to improve plant productivity? If so, the manager in charge of day-to-day system operations will want to be sure it is accurate. Executives responsible for the bottom line will want to see the financial results. Worker representatives may be interested in work content changes. If staff changes are likely, human resources personnel may be interested in the study. Various other operations

(maintenance) and staff (process engineering) functions might also be interested. Even the Marketing department may be interested in using the animation for promotion.

Every project will have different stakeholders and obviously some stakeholders will be more interested than others. And some stakeholders may be more important than others – like those who control the budget. While it is obvious that the most important stakeholders must be satisfied, do not ignore the others. Many times, the cooperation and satisfaction of the less important stakeholders can make or break your project.

2.2 How Do Your Stakeholders Define Success?

The Pragmatic Marketing group has coined a phrase “Your opinion, while interesting, is irrelevant.” This is basically saying that the customer’s (or in this case, your stakeholder’s) opinion about project success counts much more than your own. Even if you personally consider the project to have been an overwhelming success, if your most important stakeholders consider it to be a failure, your project is a failure.

It is important to probe your stakeholders to find out what their needs and expectations really are. Do they want to reduce headcount or expenses? Improve profits? Improve system predictability or reliability? Increase output? Improve customer service? In all cases, you need to find out not only what they value, but how they measure it and how they want to see the results.

It is wise to also be aware of any “hidden agendas.” Is the real reason for performing simulation analysis that someone required them to build a model? Sometimes a customer or source of funding will require a simulation model be built as a condition of a contract. In this case, the stakeholder’s main objective may be to have a model that supports what they intend to do anyway. To quote the “Lost in Space” robot ... “Danger, Will Robinson!” Starting out with the answer you must “prove” is a situation to be avoided at all cost. Sometimes stakeholders are secretly advocating a project failure – getting to know all the stakeholders on a more personal level can help unmask their motives.

Knowing how your most important stakeholders define (and hopefully even measure) success, now you are ready to write your high-level objectives. This will be the starting point for further project discussions so that everyone has a shared vision. This information also provides a good start for the detailed functional specification you will be doing at a later point.

2.3 How Will This Model Be Used?

While you will eventually need detailed functional specifications (more on that below), at this point it is important to understand high-level project objectives. Is this a one-of model or will this model form the base for use in other sites or answering ongoing questions? Are you evaluating a potential change to an existing operation or is this a greenfield operation with perhaps very limited information available? Is this a traditional design evaluation or are your stakeholders interested in using the same model for planning and scheduling as well? Is the customer looking for a Digital Twin to represent the physical assets and interact with multiple data sources to provide dynamic system behavior?

Answers to the above questions can have major impact on the project scope, tool selection, modeling approach, and even the background of the people or team required for project implementation. For example, where the model will have a longer life you might implement a *data-driven approach* where much of the model behavior is represented in data tables for easier understanding and model maintenance. If you are interfacing to existing systems (like an MES), perhaps the basic model structure can be imported from those existing systems – using a *data-generated model approach* to provide a quick project start. If you are extracting data from complex sources (like ERP or MES), you may need to supplement the background of the modeler with someone with additional database or programming background.

3 UNDERSTAND THE SYSTEM

If you are lucky, this is your system that you are modeling and you know it well. More typically, even if the system is owned by your company, you do not know it well enough to accurately model it. Every system

has subtleties which are often important. While it is not reasonable to expect a simulationist to know every system, a good simulationist should know the important questions to ask and be able to understand the answers.

One good way to start is to review the process so that you understand the key aspects. What are the entities? How are they being transformed? What is the process flow? What are the constraints? If possible, take advantage of the opportunity to literally walk through the actual or similar facility to discover things that might be missed in a discussion or diagram review.

Ask questions. Ask more questions. Ask different people the same questions and don't be surprised that you get different answers. Your goal at this stage is not to solve the problem, but to understand the problem and the system well enough that you can describe and estimate the work. Part of this stage is to identify what you don't know so that you can allow time and risk in the project for that enlightenment.

4 CREATE A FUNCTIONAL SPECIFICATION

There is an old adage that says "If you don't know where you are going, how will you know when you get there?" That is especially true in simulation projects. A functional specification clarifies the model scope and level of detail. And most importantly, it clearly defines the deliverables. It defines the objectives as well as the deliverables and determines how everyone will know when you are done.

A functional specification should clarify the project and bring everyone into a common understanding of the deliverables. Topics should include:

- Objectives – Summarize from your initial high-level objectives what you are intending to solve and what you are not intending to solve.
- Level of Detail - A model is always just an approximation of reality and can always be improved. It is important to define the limits of this model. For example, the level of detail for a particular model might be suitable for comparing the relative productivity of alternate designs, but might have insufficient detail to provide a reliable prediction of absolute system productivity.
- Data Requirements – Identify what data will be necessary to support the agreed level of detail. Where will this data come from? Who will be responsible for providing it? When will it be provided?
- Assumptions and Control Logic – Summarize your understanding of the logic in various points in the system. List any assumptions that you will be making so that you and all stakeholders have a common understanding of how much detail will be modeled for each part of the system. For example, details of dispatching, queue priority, and resource allocation should be agreed upon before modeling begins.
- Analysis and Reports – Determine who will be involved in the analysis phase of the project. Define the form and content of the results to be delivered. A mock-up of a final report is an important part of a functional specification. On review of the mockup, the stakeholders will almost certainly identify things that are missing and things that are unnecessary. It is much better to identify such items at this point than at the final project presentation.
- Animations – A certain level of animation is generally necessary for model development and validation. How important is animation to the stakeholders? In many cases stakeholders initially may indicate that animation has little importance to them. My general experience is that once stakeholders have seen the 2D or 3D animation done in development, they appreciate its value for communication and later demand it as part of the deliverable.
- Due Date and Agility – Simulation is often a process of discovery. As you model and learn about the system you will find new alternatives to explore and possibly areas of the model requiring more detail. Adequately exploring those areas can potentially make the project much more valuable. But the best results possible have no value if they are delivered after the decision has been made. When are results expected? When is the absolute "drop-dead" date after which the results will have no value?

You might think that your project doesn't need a functional specification or that it is too much formality for a small project or an internal project. It does not necessarily have to be formal. But every project needs a functional specification and it should take about 5-10% of the total project time to finalize. Even a project that is expected to complete in a single day should devote perhaps 30-60 minutes of time defining scope and detail. This time spent thinking ahead will more than pay itself back later in the project. In fact, it is best not to think of this as *extra* time spent at the beginning of a project, but rather *moving* selected tasks from late in the project up to the earliest phase. For example, at some point you will certainly need to determine what data you need and where it will come from – moving that step to the beginning of the project has significant benefits.

Developing a prototype during the functional specification phase can be enlightening to all parties. You might find that the system is easier or harder to model than you thought. The prototype does not have to be large – it should be just big enough to meet the needs of the moment, often illustrating potential simulation benefits to that particular application and identifying sources of confusion or modeling complexity. Even on the smallest project it is generally worth showing a quick model and asking the question: Is this what you mean? You might find that you have a totally different understanding than the stakeholders. Often a prototype model can be made that addresses a large percentage of what the stakeholders *say* they need. As soon as they see the prototype, they remember the complex situations and all the *other* needs that they neglected to identify earlier.

The final part of the functional specification phase is the sign-off. It should be made clear to everyone that this functional specification defines the project and that the project will be considered complete and successful when all of the aspects of the functional specification are delivered. Ideally, the final specification should be formally approved by at least the primary stakeholders to avoid later controversies.

5 MANAGE THE PROJECT

While the best time to start a simulation study is very early in the associated project's lifecycle, that is unfortunately not the most common situation. It is far more common that simulation is first considered when problems are encountered late in the cycle; perhaps a short time before the final decisions must be made. At this point, everything becomes urgent, and you may even be "late" before you have started.

In such a situation, the temptation is to go into reactive mode, letting the urgency pull you in first one direction and then another. And there is always pressure to skip important steps like deciding exactly what you want to accomplish (the functional specification phase). This tends to result in less than optimal work flow and even an incomplete project.

Manage the project, don't let it manage you. A project that is completed just *after* the decision is made is of little value. It is part of your job to manage the simulation project so that you provide valuable insight in a timely fashion. Note the words "valuable insight". All simulations are an approximation. Although a close approximation has more value, a rougher approximation can still provide valuable insight. If there is insufficient time to do the entire project well, then select a subset or a rougher approximation that you can do well in the time allotted. This should be reflected in the assumptions of the functional specification.

Simulation is often a process of discovery. You will gain knowledge as you go from the effort to accurately describe the system to the early simulation results. Often this new information may move the study in new directions. A certain amount of agility is appropriate in responding to such needs; however, too much agility can prevent project completion. At such times, you must take the difficult step of telling your stakeholders "no" and deferring such requests to a later project phase. While no one likes to hear the word no, most stakeholders would prefer an honest no to a misleading yes which basically says "Yes, I will do what you request, but as a result the project may not return any useful results within your deadline." Budget your time so that the important tasks will be completed and only then allow the project to explore some unanticipated directions.

6 COLLECT INPUT DATA

The topic of input data often catches simulationists by surprise. And it can easily be cause for project failure. In the days before the prevalence of computers and automation, it was typically the case that little or no data was available. Now, it is much more likely that you will be overwhelmed by data. Organizing and making sense of that data is often the challenge.

The first challenge is to know your data. Here is a simple, but fairly common example: Perhaps you collect some machine downtime data and when you analyze it you find that it has a minimum repair time of 8 minutes, a mode of 32 minutes and a maximum of 9.5 hours. Without additional study you might not discover that the maximum repair time also included an 8 hour off-shift time when the repair started near the end of a shift. It would be easy to use such data incorrectly in the model and generate bad results. It is important to know your data and how good it is, “scrub” it clean of any invalid data, and perform appropriate input analysis.

Since collecting data can be expensive, the objectives of your simulation study should be evaluated to determine where you need the most accurate data. For example, if you are evaluating operator utilization, it is important to have enough data related to the specific tasks for which the operators are responsible. However, the data related to another area of the system with no impact on operators may be able to be approximated. Some software contains features to help you evaluate the impact of input parameters on your output responses and recommend where to put additional data collection efforts.

You can also use your model and some pilot runs to help determine where you need better data by determining how sensitive the model is to different data values. You should check sensitivity to both the magnitude (e.g. the mean) and the variability (e.g. the range) – if the model results have little change when you use other reasonable input data, then your present numbers may be good enough. However if you notice a significant change in results, with a relative minor change to magnitude or variability, then that may be an indication that you should spend more time and effort in assuring that you have the best data possible for that parameter.

You have already specified in your functional specification who is responsible for providing data and when. It is prudent to let people know well ahead of time when you need the data and at what point the project will be delayed without it. While you may be able to place blame on someone else for causing a late project, it is far better to work together to ensure that the project is on-time and successful.

7 BUILD AND VERIFY THE MODEL (ITERATIVE)

Building a model is the process of creating a representation of the real system adequate to support meeting the stated objectives. Verifying the model is the process of ensuring that the model really does what you think it is doing. While building and verifying the model are two different tasks, they are covered under a single topic to emphasize the importance of always doing them iteratively.

7.1 Building The Model

Novices will sometimes build a large part of the model, or perhaps even the entire model, before starting verification. This is a significant cause for project failure. When you start verifying a large model, there is so much going on that understanding the detailed interactions becomes difficult or impossible. It is much more effective to instead take an iterative approach – build a piece of the model, verify it, then continue adding additional pieces of logic to the model. Two very effective approaches to model building can be summarized as ‘breadth first’ or ‘depth first’.

In ‘*breadth first*’ modeling, you might build the entire model or a major section of it with a minimal level of detail. You can then verify the model works before continuing on. This has the advantage of immediately generating a potentially useful model. Your first pass could actually be the prototype used in the functional specification. Another advantage is that you can more easily get stakeholder feedback from a complete (albeit not fully detailed) model, and get regular feedback on where more detail is required. You can sometimes even do some measure of validation (discussed later) as part of the iterative cycle.

In ‘*depth first*’ modeling, you select one small section of the system and model it in the full detail required. You can verify this model section completely and in the extreme case never have to review it again. An advantage of this approach is the ability to modularize the model – particularly important if several people could be working on the model at once. A novice might choose to build an easy section of the model first to gain experience. A more experienced simulationist might implement the hardest or trickiest sections first to eliminate some project risk early on. A modeler with some “agile” background might do the highest priority or most important sections first. With this latter approach, at any stage the most important aspects of the model have been completed. This helps reduce the risk of running out of time or budget without being able to produce any meaningful results.

‘Breadth first’ and ‘depth first’ approaches can also be combined by alternately adding some detail at the entire model level, then adding some detail to (or completing) a particular subsection. But the most important aspect is to add relatively small sections of model logic and then verify each section before adding more logic.

In each cycle of verification, you want to definitively answer two questions: Does the section of model I just built perform as I intended (e.g., are there bugs in the logic of this new section)? When this new section interacts with previously built sections of the model, does the entire model still perform as intended (e.g., are there bugs in the interactions between sections)? As your model gets larger, you might want to make your new sections smaller to make answering the second question easier.

7.2 Hey, this is hard...

Many simulation vendors would like you to think that modeling is easy, if only you choose their products. To be fair, some tools are easier to use than others, and some tools are easier to use in specific targeted applications than others. But it is rarely *easy* to build a complex model in adequate detail to effectively solve the problem. Even the most experienced simulationist will often struggle to solve some problems. *A significant portion of modeling effort is often spent resolving modeling issues.* But that’s why modelers are so highly paid (or at least we modelers wish they were.)

But “forewarned is forearmed.” Plan time for things to go wrong during modeling, because they often will. One obvious advantage of an expert user is knowledge of the tools – the ability to build a model quickly and accurately. A related, more subtle advantage is knowing how to use the tools to identify, isolate, and eradicate model bugs. Most products have some level of debugging tools. When you have a choice, select a product with the best debugging tools possible. Then take the time to learn to use those tools effectively (see next section).

7.3 How Do You Verify A Model And How Do You Isolate A Problem When You Find It?

The most obvious ways to find and diagnose model problems are to *watch the animation* and to *carefully examine the output results*. Most products also have a variety of other tools to support model verification. *Model trace* is often available that can provide great detail on exactly what is happening step by step in your model. You may want to start by watching a single entity go through the entire process. Typically there will be controls in your software to allow you to *step* through a model or to *break execution* at a particular location, time, or condition. Often there will be a *watch* window that allows you to explore the detailed system state at any time or for any object to help further clarify what is happening. And certainly take advantage of any *dashboards* or other *interactive statistics and graphics* offered by your software.

The verification process is certain to be an enlightening and quite necessary part of the project. Unexpected results are not a problem – they indicate the learning that is a primary reason for doing a simulation. Unexplainable results are a problem. When the model generates an unexpected result, you need to use all your available tools to find the explanation. In some cases that might lead to discovery of a model bug that must be fixed. In other cases it leads to an “ah-ha” moment – a flash of enlightenment about how a complex system works.

7.4 Help From A Good Listener

Even with all of the above, you might find that you have a situation that just doesn't look right, but you cannot explain why. It's time for a model walk-through.

Find a good listener, ideally a simulationist or one of your stakeholders, and go through all of the relevant model sections and explain to them what is going on. If your listener has the ability to understand what you are explaining and ask questions, that's a bonus. But in a large percentage of the time, *you will find your own problem* by methodically walking through the interactions. Keeping this in mind opens up wide possibilities for a candidate listener. An uninvolved co-worker, a spouse, or even a pet are good candidates. While dogs and cats can sometimes be good listeners, nothing beats a pet goldfish for a captive audience. The key is that explaining your model out loud seems to open up a different part of your brain and allows you to solve your own problem.

7.5 How Do You Know When You Are Done?

As mentioned earlier, a model is just an approximation of a real system. Usually the modeler and the stakeholders want the model to be as accurate and comprehensive as possible. To avoid never-ending, late, and over-budget projects, you need to go back to your functional specification document. Your goal is to build a model with just enough detail to meet the stated objectives and no more!

Animation is an area where it is easy to "get lost." Animation can be the most fun and instantly gratifying work in the project. It is easy to let it take more time than it should. Most packages have some level of automatic animation. This is typically good enough for model verification. Likewise, many packages have some level of 2D or 3D animation that is very easy to generate. Some amount of this can make validation easier by providing an additional measure of reality and recognition by stakeholders. But again you must go back to that section of the functional specification. Your final animation should be just good enough to meet the previously identified customer objectives, and no more!

8 VALIDATE THE RESULTS

Model validation needs to be done to determine if the model represents reality to the extent necessary to meet objectives. You can sometimes complete some measure of validation as you do the model building and verification iterations and should take advantage of every opportunity to do so. But you will still need to do additional validation on the completed model. Perfect verification and validation is usually impossible because the only perfect model is the real system. But there are some ways that you can attempt to demonstrate that the model is valid enough for project purposes.

One common validation technique is to start with a model of the existing system (assuming that the real system exists). Compare the results of the "as is" model against the performance of the real system. A stochastic comparison might take a representative period (e.g. 30 days or 30 weeks) and compare the average results over that period. Another approach is to make the model as deterministic as is feasible (e.g. use exact entity arrival times, exact failure data, etc.) and compare the results for that shorter period. Each of these approaches is valuable in their own way. In both cases you strive to identify and explain any significant differences.

Another validation technique is to use the experience of your stakeholders. They know the system well and should be able to watch an animation and provide some measure of confidence. You should also give them the opportunity to see the model perform under a wide variety of situations, such as high volume, low volume, or recovering from a failure. Ideally stakeholders should even be able to create such situations themselves e.g. "I want to see Machine A200 fail ... now."

While a single stakeholder can provide valuable insight, a group of stakeholders from different backgrounds can provide even greater value. Perhaps an engineer might say "Yes, you captured the design exactly as I described it," to which an operator might respond, "Maybe so, but we would never actually do it that way. Here's how we would run it...". At that point the simulation is already providing significant

value as a communication tool. Your role in the remainder of that meeting is to facilitate the discussion and take notes.

9 EXPERIMENT, ANALYZE, AND PRESENT THE RESULTS

During the experimentation phase you will be generating the scenarios identified in the functional specification. Most likely, you will also need a few additional scenarios based on what you have learned as the project progressed. It is not unusual to start with the objective to evaluate four scenarios, but discover along the way that one of those no longer makes sense, but two additional scenarios are now worthy contenders. This is a natural result of the learning and improved understanding that is the outcome of most simulation studies.

The details of the statistical analysis are beyond the scope of this paper, but proper statistical analysis is critical. Since stochastic inputs are almost always involved, it is essential to do proper stochastic analysis of the output data including an adequate number of replications and appropriate depiction of the results. It is unfortunate, but true, that simulation models often rely heavily on input data estimates (e.g. “I think it usually requires 5-9 minutes.”) and inadequate samples (“based on the 12 occurrences we observed ...”). Look for software features that can help you understand how these estimates impact the accuracy of your results. See the additional reading section for some thorough treatment of appropriate experimentation and statistical analysis.

As with all the other portions of the project, make sure you provide enough time in the schedule for experimentation and analysis. Many times, if you fall behind on the model building, verification or validation phases of the project, you may find yourself in a time crunch for the analysis. Keep in mind that the reason for doing the simulation project is typically to analyze various scenarios, so make sure to plan accordingly and leave plenty of scheduled time for the final analysis phase.

Your primary goal should be to help your stakeholders make the best decision possible given the time and resources allocated. While you might have other personal goals such as to build credibility or make a profit, it is likely those goals will be met if you concentrate on helping the stakeholders.

Consider the background and particular needs of each stakeholder before creating your report. Although you are probably proud of your model and the detailed way in which you solved complex problems, few stakeholders will share that interest. Most stakeholders are interested in three things. First, what alternatives were considered? Second, what are your conclusions or recommendations? Third, what supporting information can you provide to merit their confidence in your analysis?

Although you need to have data to support your conclusions, do not overwhelm your stakeholders with too many details. Try to provide information in the context needed. For example, instead of simply stating “Average driver utilization was 76%”, you might say “Since the average driver utilization is high (76%), there is inadequate slack time to catch up during peak periods without causing line delays.”

Don’t over-represent the accuracy of the output data. Acknowledge and even emphasize to the stakeholders that the model is an approximation and it will not generate exact answers. Display your data with appropriate precision based on the accuracy of your data and modeling assumptions (e.g. 76.2% not 76.2315738%). And display the accuracy of your numbers when possible. Most stakeholders can relate to a confidence interval like $76.2\% \pm 1.3\%$.

10 SUMMARY

In spite of what you might have heard, doing simulation projects well is not easy. There are many ways that even an experienced simulationist can fail. In this paper we have discussed some common traps and ways to avoid them. While following these suggestions will not guarantee a bulls eye, it will certainly improve your chance of hitting the target.

ADDITIONAL READING

- Banks, J., J. S. Carson, B. L. Nelson, and D. M. Nicol. 2010. *Discrete-Event System Simulation*. 5th ed. Upper Saddle River, New Jersey: Prentice-Hall, Inc.
- Law, A. M. 2015. *Simulation Modeling and Analysis*. 5th ed. New York, NY: McGraw-Hill, Inc.
- McHaney, R. and T. P. Cronan. 2000. "Toward an Empirical Understanding of Computer Simulation Implementation Success". *Information & Management* 37(3):135-151.
- McHaney, R., D. White, and G. Heilman. 2002. "Simulation Project Success and Failure: Survey Findings". *Simulation & Gaming* 33(1):49 – 66.
- Pegden, C. D. 2015. *Process Improvement Principles: A Concise Guide for Managers*. 1st ed. Pittsburgh, US: Simio LLC.
- Pegden, C. D. 2017. *Deliver On Your Promise – How Simulation-Based Scheduling Will Change Your Business*. 1st ed. Pittsburgh, US: Simio LLC.
- Sadowski, D. A. and M. R. Grabau. 1999. "Tips for Successful Practice of Simulation". In *Proceedings of the 1999 Winter Simulation Conference*, edited by P. A. Farrington, H. B. Nembhard, D. T. Sturrock, and G. W. Evans, 60-66. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Smith, J. S., D. T. Sturrock, and W. D. Kelton. 2019. *Simio and Simulation: Modeling, Analysis, Applications*. 5th ed. Pittsburgh, US: Simio LLC.
- Sturrock, D. T. 2011. "Simulation Success Skills". *Industrial Engineer Magazine*, December 2011, 43(12).
- Sturrock, D. T. n.d. Success in Simulation, Ongoing blog and discussion at <http://simio.com/blog/about-this-blog/>, accessed April 2019.

AUTHOR BIOGRAPHY

DAVID T. STURROCK is Vice President of Operations for Simio LLC www.simio.com. He graduated from The Pennsylvania State University in Industrial Engineering. He has over 35 years of experience in the simulation field and has applied simulation techniques in the areas of transportation systems, scheduling, plant layout, call centers, capacity analysis, process design, health care, packaging systems and real-time control. He is co-author of two leading simulation textbooks, contributed to several others, and teaches simulation at the University of Pittsburgh. In his present role for Simio he is responsible for development, support, and engineering services resources for the Simio simulation and scheduling product suite. His email is dsturrock@simio.com.