

WEB-BASED SIMULATION: SOME PERSONAL OBSERVATIONS

Paul A. Fishwick

Department of Computer and Information Science and Engineering
University of Florida
Building CSE, Room 301
Gainesville, FL 32611

ABSTRACT

The web has a short history, but has grown at an exponential pace since its introduction six years ago. *Web-based simulation* represents the connection between the web and the field of simulation. Web-based simulation is not an existing field but rather an idea which represents an interest on the part of simulationists to exploit web technology. To further this cause, I present some issues and concepts in web-based simulation to serve as a back drop for a more formal discussion, and potentially the formation of a new simulation sub-area.

1 INTRODUCTION TO THE WEB

What is web-based simulation? First, we need to discuss a few terms before we can adequately address this question. But, before doing that, we need to say something about the nature of this topic. The web represents a rapid growth of Internet-based interaction over the past six years. Yet, the topic reflects more of a *technology-transfer* area than one with a solid theoretic underpinning. Rapid advances have taken place through practical application, engineering and entrepreneurial activity. A new web advancement will appear, almost overnight, as a working product to be tried by eager users. This situation—not unlike the gold rush in the mid-nineteenth century—suggests that, by default, an article on “web-based simulation” needs to focus on technology. Yet, as researchers, we collectively require a substantial knowledge of this technology if we are to apply the web to our disciplines. It is not healthy to ignore the web and remain fixated to either purely academic tomes or yesterday’s approaches to simulation. My purpose in this article is to inform the reader about some of the possibilities with an applications and technology transfer focus. I have been involved as a heavy user of web technology and will impart some of my personal experiences with the web

as they relate to simulation. I will unabashedly refer to specific packages and methods, not with the deliberate intention of presenting a biased view, but because this is not a survey paper and there are too many web-based methods to be able to develop a timely paper without making the content obsolete once it is in print. Indeed, novel web-based applications seem to crop up so frequently that information on web-based simulation will likely change as rapidly as these applications. In this sense, the article is somewhat anecdotal, and is meant to help simulation-interested readers who are beginning to explore the web, and possibly even simulation. Furthermore, in the spirit of the subject matter, all of our references will be universal resource locators (URLs) and a hypermedia copy of this document should be accessed to be able to link appropriately where required. The hypermedia document is available by pointing your browser at: <http://www.cise.ufl.edu/~fishwick/websim.html> as soon as you are near a PC or workstation. Beyond this unique URL, which serves as the gateway to all further citations and references, all other URLs will not be visible in this article for the sole reason that their employment is not possible without a browser. Nothing is gained from reading the printed URL, and with a browser, you will not require this document. So, I encourage you at the next opportunity to go to a browser at locate the previous URL. It is assumed that you will have access to a browser, as there are several free browsers available and some (such as Microsoft’s Internet Navigator) come bundled with their operating systems (Windows 95 and NT), which in turn are bundled with all PCs. The PC is the most popular platform for web-based applications due to the huge number of PCs and the PC open architecture. Other supported platforms include Apple’s MacIntosh and Unix workstations such as the ones made by Sun, Hewlett-Packard and Silicon Graphics. Most web software is currently packaged for Windows 95 first, then Windows NT and Windows 3.1, followed by Unix versions.

The Internet is a loosely connected world-wide network of heterogeneous computers. The World Wide Web (WWW) is a set of on-line hypermedia documents that reside on the Internet. Hypermedia documents are built using a language called hypertext markup language (HTML). HTML started as a subset of SGML (Standard Graphics Markup Language) but grew to absorb the needs of web authors. Access to the web is afforded using a web *browser* such as Netscape, Mosaic or Internet Navigator. When you use the browser, you are presented with a default *home page* which presents you with your first online document. This document will include *hyperlinks* which, when clicked using a mouse, link to another document. In this way, a document is connected to other documents in a sprawling network: there is no centralized control—just documents loosely connected to one another through links. Documents may contain more than the usual printed document:

- Text: various natural languages are supported.
- Images: typically GIF files seem to be the best supported but JPEG is also supported. Image converters exist for most platforms.
- Video: supported video types are MPEG and AVI, with the additional GIF89 that supports inline small video clips meant for small icon display.
- Audio: AU and WAV are supported formats, as well as digital standards such as MIDI.

Although web browsers support many exciting features, our purpose is to say something about how the web and the field of simulation relate to one another: web-based simulation. Given a browser and the ability to write new web pages, what are the issues for simulationists? How can we exploit the new technologies and how will our field change as a result of the web?

2 HOW SIMULATION WILL BE AFFECTED

There are many aspects to simulation as there are for any scientific discipline. We will focus on the following aspects: 1) Education and Training, 2) Publications, 3) Simulation Programs. All three of these topics are interrelated. For the academic readers, the topics of education and publication are probably on the top of the list, and industrial simulationists will want to focus more on training and the practical methods for building simulations that use the web.

2.1 Education and Training

How do we currently educate our students using simulation? Most modern simulation packages today have graphical front ends (so called graphical user interfaces—GUIs). This allows a student to learn through virtual world exploration. The web will affect this process in several ways. First, a diskette or CD-ROM has a limited amount of local storage but the web provides us with effectively infinite storage. The “storage” is on the Internet and not limited to one’s own machine. Therefore simulation packages which include help text and information about the pieces of the model, can have these pieces on the web so that they do not require local storage. This encourages some collaboration in that a company marketing their simulation tool, conveniently priced for educators, will have to think more *globally* about information and resources. It is not necessary for a company to re-build every piece of information about a manufacturing process, for instance, from scratch. Information on automated guided vehicles (AGVs), machine specifications, and automated conveyance mechanisms may already be located somewhere on the web. These sorts of devices are common in simulation programs built for manufacturing systems analysis. Why re-invent the wheel by building our own information databases on these manufacturing devices? The web encourages this sort of global view with a *re-use* of knowledge and information. While academics may revel in this kind of reusable knowledge, industry will have to rethink certain proprietary concerns: a company needs to make a profit to survive. They also need to produce quality simulation software, but this does not negate the use of the information of other web sites. Credit must be given where it is due, as usual, and it is not clear exactly how copyright, trademarks and patents will change as a result of the web phenomenon. It is too early to tell, but we need to address these issues now.

We have discussed the use of web sites in user documentation for a simulation program, but the program could just as easily be placed on the web as well. After all, is the program really that different from hypermedia “documentation?” We discuss this in more detail within this section. The most profound affect the web may have on our method of teaching simulation lies with the use of multimedia, since it is multimedia that we find in every web link. The web encourages distance learning more so than the typical simulation textbook. On any web page it is possible to include images and video of the instructor along with synchronized slides or overheads. This immerses the student in a synthetic learning environment that is more congenial than one they would get

simply by reading a book or watching a videotape. Various technologies are around to aid us in creating multimedia web pages for simulation education. Before mentioning a few of them, it is worth defining the concepts of *client* and *server* since the web encapsulates these two roles. A server provides a client with information and services. The daily metaphors of servers and clients (whether going to a restaurant or visiting your doctor or lawyer) are the best way of understanding these roles. If you enter your browser and access a travel agency in Connecticut, you are able to do this because the travel agency has a computer which contains a web server. This web server services your needs by delivering multimedia documents to you via your browser. The relationship is directional: the agency is the server and you (or your machine) are the client. Every machine, potentially, has the ability to be both server and client so there is no reason why the agency might not use your machine as a server if you have it set up as such, delivering information required by the travel agency. Any machine can serve, then, as both client and server.

A *plug-in* is a software program to allow third party vendors (with respect to the browser designer) to provide multimedia service for web pages. For example, if I install the PDF reader from Adobe, Inc. into Netscape, I will then be able to include PDF files in one of my web pages, so when I access someone else's web page which contains a server capability delivering PDF files, my PDF plug-in will activate and I will be able to see PDF files. Plug-ins exist for numerous functions. A recent trend has been for commercial vendors to provide plug-ins for free to anyone, but to charge a fee for those wanting to serve clients. If you want, for example, to design your own PDF files so that others can access them over the web, this will involve a cost *to you*, and this is where the company generates its required revenue. If web software is free, it is free only in the sense that someone else is paying for you to use it in the same way that TV and Radio advertisers pay for the programming which you enjoy.

What plug-ins are useful to a simulation educator? When you have a need to deliver real-time video and audio to your students, then you will need to purchase the server software, while your students can freely obtain the client-side plug-in. For non-real time video and audio, many free players exist (e.g., MPEG and AVI for video and AU, MIDI for audio), and so most companies seem to be making a profit for providing a much-needed real-time audio-video stream, especially over low bandwidth connections such as via 28.8K modems. VDOLIVE and VDOPHONE are especially good at providing close-to real time audio and video. VDOLIVE is a method of "streaming" video and au-

dio in the same way that a TV does. VDOPHONE is an Internet substitute for a videophone. Both of these packages employ wavelet compression to achieve excellent performance. Video frame rate is generally sacrificed where necessary to maintain uninterrupted audio. This is as it should be since "broken" video is acceptable, but this is not true for audio. Occasionally, another marketing technique is to provide a "personal server" for free. This provides you with the ability to use one stream, which is fine as long as two clients (i.e., students) do not attempt to access the video stream simultaneously. Many other plugins are useful as well, such as one that can display color slides, possibly in synchronization with the audio and video streaming plugins.

In the sense that our conferences represent "education and training," we need to make major overhauls of the way we do business. The use of teleconferencing will help us to interact prior to meeting at a conference. Personal meetings with each other cannot be replaced by electronic means, but new web tools such as CU-SEEME and VDOPHONE (section 2.2) have the means to allow us to communicate more effectively. E-mail, mailing lists and USENET news groups (including the new web-based news groups) are useful technologies. The USENET news group `comp.simulation` has been in operation for almost a decade and has helped to grow the discipline. The ELECSIM conference has been successful in forging ahead to permit us to communicate with one another without forcing cross-country trips to conferences.

2.2 Publications

We publish for one reason: to inform others of new or existing simulation research and applications. We write articles in everything from conference proceedings such as this one, and trade magazines to academic archival journals. The web is dramatically changing the way we publish, and it is changing so fast that many are rethinking how publications are delivered and purchased. Let's consider archival simulation journals first. If you publish an article in an archival journal, there is an accepted practice which is defined as follows:

1. *Refereeing*: you submit a manuscript to the journal's Editor in Chief. The editor then has an editorial board which is used to manage the review process for the paper. The idea is that, for your paper to be accepted, it should be sent to qualified reviewers who are experts in the sub-simulation area that represents your paper's focus. Moreover, these reviewers are

anonymous so as not to bias the review. Sometimes, conferences try to snip part of an author's submission so that the reviewer doesn't know who wrote the article. This procedure tends not to be practiced in the simulation community, and the procedure has flaws since self-reference (to establish background) in a paper generally links a specific person to the paper. Once the paper is reviewed, it can be rejected, accepted or subject to revision with provided comments and critique by the reviewers. A review process from time of submission to a decision can take anywhere from three months to six months.

2. *Publication*: The review process can be lengthy, and assuming that you have made the proper revisions to your paper, your paper may be accepted by the Editor in Chief on recommendation of a member of the board. The time from acceptance to the time your paper appears depends on two things: the periodicity of the journal and the current queue (backlog) of already accepted papers. Times from six months to two years are common.
3. *Archival*: Libraries typically archive journals so that others can search for your article (using a scientific index) for their own research.

This entire process is lengthy but there is good reason for much of the time taken to evaluate the paper. We, as scientists, also accept this method since it has proven to be the way to introduce original simulation knowledge to grow our information base. The web has changed the playing field - many new ideas are surfacing. Most of the simulation magazines and journals are making new inroads here. For example, both the ACM Transactions on Modeling and Computer Simulation (TOMACS) and the SCS Transactions on Simulation are rapidly transitioning to electronic form. Let's consider some issues:

1. *Literature Search*: a publication begins with a thorough literature search. This search is made easier on the web through search tools, but searching for the right information is still difficult. There is so much information that searches can become unwieldy. Often, searches on the web result in "missing link" problems: all information is distributed and so links change dynamically. Only evolved and maintained links remain in way not unlike the processes of genetic variation and fitness in living organisms. Libraries are changing their functions to help with information access and retrieval. There

is still an art to locating just the right information and this is the task of an information specialist—the modern librarian.

2. *Electronic Access*: make everything electronic. The author submits the paper using an electronic medium such as Postscript, which is fast becoming a kind of pseudo-standard for page layout among journals. Referees are sent (either via mail attachments or using file transfer-ftp) a compressed Postscript file. Referees submit their reviews by e-mail and the author is informed of the decision by e-mail as well. All of this makes sense and is happening now. All citations use URL links instead of printing a bibliography. This allows the reader to jump directly to a cited paper, avoiding lengthy and sometimes unsuccessful literature searches.
3. *Refereeing*: more comments can be made available to an author if anyone is permitted to attach comments to an article being considered for publication. Why limit the comments to two or three people? We call this "reader refereeing": let those who are interested in reading an article make comments on it. One possibility is to have a combination of traditional and reader refereeing.
4. *Revenue*: who pays for the article? Before, everything was so simple - the entity which published the hardcopy of the journal received basically all the revenue. However, if the volunteers (editorial board) are doing all the work and do not need the publisher, then should simulation journals be free? Well, there are costs lurking everywhere: the cost of the server equipment and the cost of archiving journal articles. Someone has to pay for these costs. Costs have been translated from paper-printing into digital service and storage.
5. *Preprints*: remember that the goal of publishing is for you to impart information to others on your new ideas. If you can put your papers immediately online in your web page, and have these web pages indexed by web robots and *intelligent agents*, then you don't have to wait any longer to see your work "in print." It is in print immediately. The generally accepted practice is for authors to place preprints on the web, either in a home page or in a preprint area maintained by someone else, and to simultaneously submit the article for consideration by a journal. Whereas before, only a handful of people might ever thoroughly read your journal article, now anyone with a web browser can find it

using a search engine and read it. Yet, there is so much information on the web, that we want to narrow our search. As societies abandon paper in favor of web-publishing, there still will be *value* in doing your search for simulation papers via society servers first since that is where papers have been checked for quality.

2.3 Simulation Programs

The most exciting part of web-based simulation is in the simulations themselves. Operating on the web gives new meaning to the word *document*. A web document can contain videos, interactive simulations, images and audio in addition to the usual text that traditional documents contain. How do we embed simulations in documents? A simulation requires a model and a model requires computing hardware to execute the model. Some parts of a model can be executed on another machine and other parts are executed on your machine where your web browser is located. For simplicity, we'll assume that you are the one building the simulation and others want to execute it. To build a simulation which executes on your machine requires that someone else (the client) access your simulation (the server) and cause a program to execute. Typically scripting languages such as Perl are used to build a high level program that runs your simulation code once you have obtained all required simulation inputs and parameters, generally using a "form" capability. If many users are running your simulation simultaneously, you will need to make sure your Perl script is set up to handle this. Your computer resources may be taxed, depending on the number of readers wanting to run the simulation. The idea of Java and Javascript is to bundle a piece of code (called an applet) and send it to the client's computer to execute, not yours. This way, it makes no difference how many clients want to run your simulation at least from the perspective of model execution.

In our discussion so far, you have constructed a simulation which runs on your computer or sends applets to a client's computer. This is fine for single-user simulations, but what if you need to have a simulation where multiple users *interact* with each other? The best example of user interaction is found in multi-user gaming simulations such as DOOM and QUAKE. In such a simulation, a global server (one person's machine) is used to capture and maintain the state of the world. Each player makes moves and information is moved from client to server. Multi-User Dungeons (MUD) operate in a similar fashion, the difference being that the interaction among users may extend toward a cooperative relation to other players

rather than a strictly competitive one. The MUD is also a basis for applying education through simulation within a virtual simulated "world." This represents a shared solution to multi-user simulation.

If the simulation model is spread over a network of computers, then we have a distributed solution. Each computer is associated with either a single physical entity or a set of entities. The distributed interactive simulation (DIS) initiative begun by the Department of Defense (DoD) is a good example of distributed model execution. The DoD focus is to have an interactive simulation capability mainly for training people, but increasingly other aspects of the military (such as "acquisition") are finding uses for DIS. Some of the computers in a DIS network may be running semi-automated forces (SAFOR) and others will have humans attached to often sophisticated human-computer interfaces so that the training exercise is made as realistic as possible. If there is no real-time human-computer interaction (HCI) necessary, we may still split our models and do parallel and distributed simulation (using conservative or optimistic time advance).

How does the web relate to all of this? Since the web is a collection of documents, an interactive simulation could ostensibly be entered from within a web document. A user locates a document, presses the mouse button on an image portraying a virtual world and then is immersed in a simulation. Most DIS simulations do not currently operate in this fashion, but there is no reason why this could not be done. After all, the web is all about storing and retrieving multimedia information, and simulation is a key *kind of information*. The virtual reality markup language (VRML) suggests a way to do this. VRML plugins allow one to construct 3D geometry and have it rendered in real time. A document contains a URL which points to a VRML file. The browser recognizes this particular information type and launches the VRML plugin. Then the user browses a 3D scene using whatever HCI devices are available.

With DIS and MUDs, we are only scratching the surface of what is possible. You have probably noticed that most companies currently have web sites. What if they were to take the products that they manufacture and put digital equivalents of these products on their web site? This would radically change the way in which we do modeling and simulation. One of the biggest reasons why simulation of large systems is expensive is that each one of us rebuilds model components instead of re-using what could be out there on the web. Customers will benefit from such web-based models since they can try out a product before purchasing it. This also goes for industrial customers as much as it does for indi-

viduals. For the original equipment manufacturers, this scheme will allow for them to sell more equipment since potential buyers can try out the equipment simply by linking the part of their simulation model which models the equipment to the manufacturer's web site. Manufacturers will have to spend more time building digital web-based models for their products, but this will be time well spent once readers and users locate the models.

We have done some preliminary studies of web-based simulation at the University of Florida. Figure 1 displays the home web page for a simulation of four disks and one CPU. A job enters the CPU and then is assigned a disk unit for file access. Then, this job loops a certain number of times. This looping simulates the time taken by an average program (i.e., job) which has need of CPU and file storage resources. The web page is organized in an overall linear fashion, starting with a paper title, author list and then proceeding with a picture of the system to be simulated. The square area of the screen uses an "ISMAP" feature which allows us (the authors of the web page) to perform a given action if the reader clicks on a certain region within this square. At the bottom of Figure 1, the user clicks on this to begin entering simulation input information. All input is entered using the FORMS feature of HTML. Once the the information has been entered, the reader can execute the simulation. Since our simulation uses the "server side approach," using a Perl script which invokes an executable on a SPARC Unix workstation, the simulation executes and produces time-dependent graphs for CPU and Disk resource queue sizes. The graphs are produced by dynamically creating an HTML web page from the simulation output, which is then available to the reader.

3 REDUCING THE COST OF SIMULATION

Simulation has always been a costly enterprise, and even though more inexpensive computer equipment is driving costs down, the ever-lurking cost makes itself known in all of our concerns: education, publication and simulation programs. We have just discussed reducing this cost by the re-use of objects within simulation models but there is a broader strategy: advertising. Everyone has questions about how money is to be effectively exchanged over the web. Revenue issues will affect the simulation community as much as any other, so we need to consider the transfer of money. Simulation products and services can be bought through traditional credit card means, and we have seen that many services (in the form of web products) are free for the client side of information

delivery. The cost is on the server side. Many simulation services may be possible through inventive advertising methods such as including an advertiser's message in a web page accessed by simulation practitioners. Most web-search engines contain this feature. The web user benefits by using the search engine and the search engine service is paid for by the advertiser. Everyone wins. In this way, education, publication and simulation software may become cheaper as a result of embedded advertisements. We should be concerned about over-commercializing simulation knowledge; however, web advertisements are *less obtrusive* to the reader and provide direct links to the advertiser. Furthermore, by the web's nature, advertiser information tends to be more information-rich compared to other types of media such as television, radio and magazines. The advertising links are also interactive so there is the potential for providing quality information to readers. This means that simulation publications, educational products and interactive simulation programs of the near future might incorporate direct advertising to cover costs. The advertising is less blatant, more accommodating, and provides for a three-way benefit to the server, client and advertiser.

4 SUMMARY

We are in a kind of twenty-first century gold rush since everyone is rushing to put multimedia information on the web, but there is great uncertainty about how things such look and feel. How will simulation-related societies change and deliver their publications? Will companies expand their sales by offering simulation models of their products? One thing is for certain. We will never know any of these answers unless we experiment and take risks if necessary. The "right way to do simulation" will naturally emerge and evolve as we try various web-based simulation approaches. The worst possible approach is to sit back and adopt a "let's wait" attitude: simulation will be worse off as a discipline unless we move forward now to incorporate the web-based technologies.

ACKNOWLEDGMENTS

I would like to acknowledge the following funding sources which have contributed towards our study of modeling and simulation. (1) Rome Laboratory, Griffiss Air Force Base, New York under contract F30602-95-C-0267 and grant F30602-95-1-0031; (2) Department of the Interior under grant 14-45-0009-1544-154 and the (3) National Science Foundation Engineering Research Center (ERC) in Particle Science and Technology at the University of Florida

An Interactive Web Simulation of CPU/Disk Performance

Paul Fishwick, Matthew Belk, Bradley Spatz

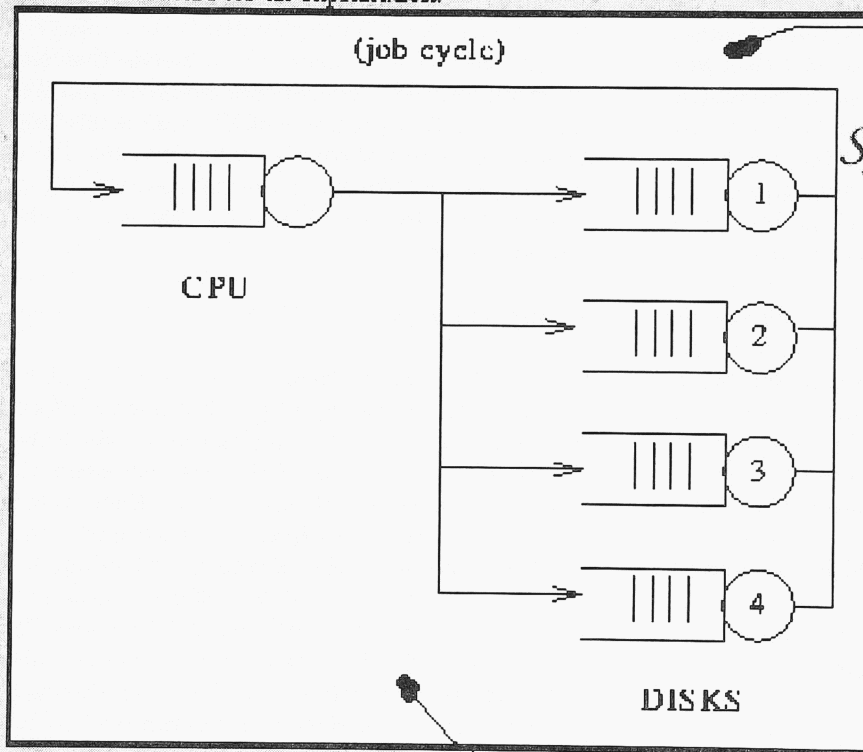
Home Pages

We present a functional model of a CPU/Disk subsystem with an interactive execution using SimPack. Aside from providing technical data on the model, the paper serves as an experiment in publishing an interactive simulation on the World Wide Web. See Citing and Referencing Web Publications for further information. Let's consider a combined CPU/Disk system (Audio Introduction):

Point to a feature for an explanation.

Underlined Hyperlinks

System to Simulate



Let's simulate this system!

ISMAP Region

Execute Perl Script (Run Simulation)

Figure 1: Web page of a CPU/Disk simulation.

(with Industrial Partners of the ERC) under grant EEC-94-02989.

AUTHOR BIOGRAPHY

PAUL A. FISHWICK is an Associate Professor in the Department of Computer and Information Science and Engineering at the University of Florida. He received the PhD in Computer and Information Science from the University of Pennsylvania in 1986. He also has six years of industrial/government production and research experience working at Newport News Shipbuilding and Dry Dock Co. (doing CAD/ CAM parts definition research) and at NASA Langley Research Center (studying engineering data base models for structural engineering). His research interests are in computer simulation modeling and analysis methods for complex systems. He is a senior member of the IEEE and the Society for Computer Simulation. He is also a member of the IEEE Society for Systems, Man and Cybernetics, ACM and AAAI. Dr. Fishwick founded the `comp.simulation` Internet news group (Simulation Digest) in 1987. He has chaired workshops and conferences in the area of computer simulation, and will serve as General Chair of the 2000 Winter Simulation Conference. He was chairman of the IEEE Computer Society technical committee on simulation (TCSIM) for two years (1988-1990) and he is on the editorial boards of several journals including the *ACM Transactions on Modeling and Computer Simulation*, *IEEE Transactions on Systems, Man and Cybernetics*, *The Transactions of the Society for Computer Simulation*, *International Journal of Computer Simulation*, and the *Journal of Systems Engineering*. Dr. Fishwick's WWW home page is <http://www.cise.ufl.edu/~fishwick> and his E-mail address is `fishwick@cise.ufl.edu`.