

## **SIMULATING NON-STATIONARY CONGESTION SYSTEMS USING SPLITTING WITH APPLICATIONS TO CYBER SECURITY**

Martin J. Fischer  
Denise M.B. Masi

John F. Shortle  
Chun-Hung Chen

Noblis, Inc.  
3150 Fairview Park Drive  
Falls Church, VA 22042, USA

Dept of Systems Eng. & Operations Research  
George Mason University  
4400 University Dr., Fairfax, VA 22032, USA

### **ABSTRACT**

According to the former counterterrorism czar, Richard A. Clarke (2010), our national infrastructure could be severely damaged in 15 minutes by a cyber attack. A worm attack on an Internet Protocol (IP) network is one type of attack that is possible. Such an attack would result in a non-stationary arrival process of packets on a link in the network. In this paper we present an initial use of our Optimal Splitting Technique for Rare Events (OSTRE) to simulate the congestion imposed by the worm on the link. This initial application is oriented to testing the technique in this dynamic environment and report on its use as compared with conventional simulations.

### **1 INTRODUCTION**

According to the former counterterrorism czar, Richard A. Clarke (2010), our national infrastructure could be severely damaged in 15 minutes by a cyber attack. A worm attack on an Internet Protocol (IP) network is one type of attack that is possible. Such an attack would result in a non-stationary arrival process of packets on a link in the network. In this paper we present an initial use of our Optimal Splitting Technique for Rare Events (OSTRE) to efficiently simulate the congestion imposed by the worm on the link. This initial application is oriented to testing the technique in this dynamic environment and report on its use as compared with conventional simulations.

Cyber security has become a national priority. In the fall of 2003, the Department of Homeland Security (DHS) National Cyber Security Division created a new government organization called the U.S. Computer Emergency Response Team (US-CERT) to lead all cyber attack incident prevention and response efforts across the country. The Carnegie Mellon University CERT<sup>®</sup> Coordination Center (CERT/CC) reports that “along with the rapid increase in the size of the Internet and its use for critical functions, there have been progressive changes in intruder techniques, increased amounts of damage, increased difficulty of detecting an attack, and increased difficulty of catching the attackers” (Allen 2004). Figure 1 shows the number of vulnerabilities reported to CERT/CC from 1995 to 2007 (CERT/CC “CERT Statistics (Historical)”); CERT/CC stopped collecting this data in 2008.

While private IP networks, which are closed networks based on IP technology, can overcome the performance issues of the Internet and improve security somewhat, there are still security concerns for these private networks. The National Communications System (NCS) runs several emergency telecommunications programs for federal government users using the public Internet. Their mission is to ensure that priority traffic can be delivered during emergencies or cyber attacks.




Figure 1: Number of Vulnerabilities Reported to CERT/CC

Liu and Cheng (2009) give an excellent overview of cyber attacks. They discuss some causes of cyber security problems. These include bugs, defects, flaws, and vulnerabilities. They also examine the number of reported vulnerabilities to CERT/CC from 1995 to 2008 and draw the conclusion that it is very difficult for enterprises to keep up with the number of software patches to combat the attacks. They discuss how a cyber attack occurs and the four steps in the process: conducting reconnaissance, scanning targets, exploiting systems and establishing footholds. The Winter 2004 edition of *Sigma*, titled “The Impact of Cyber Attacks on Internet Protocol (IP) Networks,” discusses the basics of cyber attacks on IP networks and gives an overview of the relevant cyber attack problem on the Internet, discusses possible attacks such as worms and distributed denial of service attacks, and other critical Internet vulnerabilities such as the Border Gateway Protocol (BGP) and the Domain Name Servers (DNS).

While the focus of many efforts is on the detection and prevention of cyber attacks, a major gap is that little research and tools are devoted to the modeling of the performance impact on networks as a result of different kinds of cyber attacks. Thus, there is a recognized need for the modeling of the performance of cyber networks. This need is of continual interest to the national government. One particular concern is ensuring that critical traffic is delivered in the network during times of crisis or during a serious worm attack. Here we report on the use of our splitting methodology to simulate the link performance of an IP network under a worm attack. In Shortle et al. (2010), we introduce the use of this technique on cyber and power grid problems.

In Section 2 we formally define the problem and discuss the initial worm simulation study. Section 3 presents our simulation methodology. Section 4 discusses how the splitting technique is applied to this cyber problem. The results of our simulation study are presented in Section 5. Section 6 contains our conclusions and next steps.

## 2 PROBLEM DEFINITION AND SIMULATION STUDY

We have developed an analytic performance model of an IP network for analysis of cyber attacks, called the IP Network Performance and Analysis Tool (IP-NPAT) (Masi, Fischer, and Garbin 2004a; Masi, Fischer, and Garbin 2004b; Masi and Fischer 2005). For worm attacks, it uses the epidemic approach to predict the number of hosts infected at discrete points in time. At each time step, the origin-destination network traffic is increased corresponding to the scanning traffic generated by the number of infected hosts, and the network performance is estimated. The worm propagation and traffic adjustment equations can be modified to assume multiple hosts behind each router. Efficiency becomes a big concern when si-

mutating rare events that large-scale attacks may cause. A detailed description of the model we developed of a worm attack is given in “Internet Protocol (IP) Network Performance and Analysis Tool (IP-NPAT),” Scientific and Technical Report, 2005.

A worm spreads through computer networks by searching, attacking, and infecting remote computers automatically (for example, see Zou, Gong, and Towsley 2002, 2003, 2004 and 2006). Based on the vulnerability of the targeted host, some hosts become infected and others do not. As more hosts are infected, the network sees a significant increase in packet volumes due to scanning. In fact, “the biggest impact of these worms is that their propagation effectively creates a denial of service in many parts of the Internet because of the huge amounts of scan traffic generated” (CERT/CC “Overview of Attack Trends”). As the worm propagates through the network, the traffic loading on the network is changing.

Simulation is a powerful tool that can be used to analyze a wide variety of systems. In principle, given an accurate model and ample computer time, simulation can provide answers to many important problems. However, when the problems require the evaluation of rare events, such as packet performance under a worm attack, the number of simulation runs (or replications) required to achieve a reasonable confidence interval can be prohibitively high; especially considering the non-stationarity of the network traffic.

Here, we apply an effective simulation methodology, called the Optimal Splitting Technique for Rare Events (OSTRE) simulation technique, to simulate the link performance. The key idea is to integrate the notion of optimal computing budget allocation (Chen et al. 2005, 2008, 2010) into splitting to optimally allocate the limited computing budget so that the overall simulation efficiency can be maximized. This technique provides an optimal choice for the numbers of simulation runs at each level, regardless of the choice of the importance function or the choice of the number and locations of the levels or the simulation costs at all levels.

Our initial application of OSTRE compares it with conventional simulations. We use the worm model in IP-NPAT to generate the attack traffic on the Atlanta Provider to Customer Edge T1 (1,536 kbps) access link of the 7 router, 14 link network shown Figure 2. This access link is the focus of our simulation. Future work will extend the application to the entire network, rather than focusing on a single link.

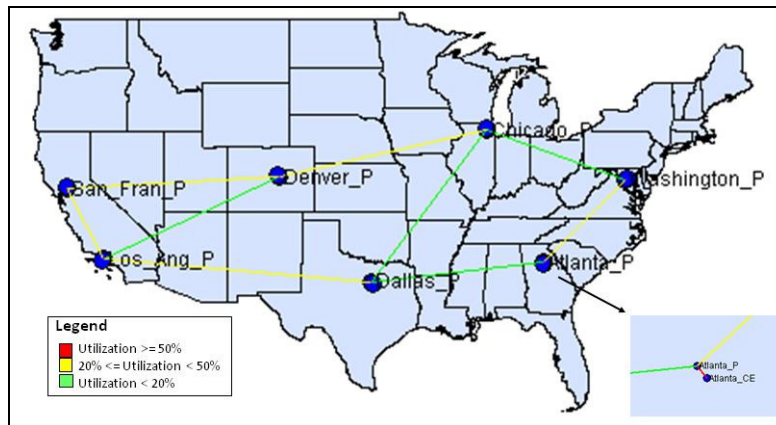


Figure 2: Seven Router, Fourteen Link Network

The resultant link loading in packets per millisecond (ppms) is given in Figure 3.

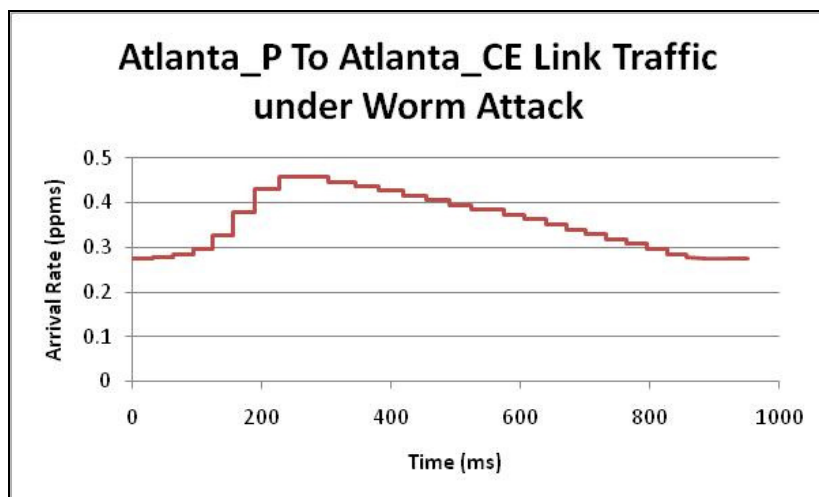


Figure 3: The Non-Stationary Worm Link Loading

From Figure 3 we see that the traffic on the Atlanta access link increases from a steady state value up to a maximum value as the worm propagates through the network, and then decreases back down to the steady state value as host computers are removed from the infectious or susceptible states through patching and cleaning. The length of time at each rate is based on the network congestion, and, thus, varies. This non-stationary nature of the traffic puts steady-state results in question. Here we investigate the use of the OSTRE simulation technique. In general, we use a discrete packet size distribution; sizes are 320, 4608, and 12000 bytes with probabilities 0.615, 0.165 and 0.23. This data is based on IP packet size distribution generated by Merit Network, Inc., a consortium of Michigan-based Internet Service Providers (ISPs). This traffic is the steady state distribution; we assume that the worm scanning traffic is also using this same packet size distribution.

We also assume Poisson internal and external arrivals to the network links, including the Atlanta access link that is simulated in this study. The arrival rates vary as shown in Figure 3. Thus we are considering a non-stationary  $M(t)/D_3/1/K$  queueing system. Technically, the network model does not have internal Poisson flows. However, previous analysis (Masi et al. 2004a) compared the analytic modeling assuming Poisson external and internal arrivals with a simulation which does not assume Poisson internal arrivals. The good agreement in the simulation and analytical results indicates that the assumption of internal Poisson flows due to superposition of the arrival processes is likely valid.

We assume the buffer size  $K$  is 80 and are interested in the rare event of packet buffer overflow. We compare the simulation variance of the buffer overflow using the splitting technique and using standard simulation. Future work will investigate this problem at the network level.

### 3 SIMULATION METHODOLOGY

We employ a splitting approach to improve the efficiency of rare-event simulation. The basic idea of splitting is to create separate copies of the simulation whenever it gets close to the rare event (Figure 4). Effectively, this multiplies promising runs that are “near” the rare event, thus improving the efficiency of the simulation. There are many variations of the basic splitting concept. Here, we apply an implementation presented in Shortle and Chen (2008). The key idea is to apply the notion of optimal computing budget allocation (Chen et al. 2005, 2008, 2010) to determine a good allocation of simulation runs at the intermediate levels. The following is a rough description of the method, omitting some details. For further details, see Shortle and Chen (2008).




Figure 4: Level Splitting

Figure 4 shows sample paths of the “system level” as a function of time. System level measures proximity to the rare event and can be represented as a function from the (possibly multi-dimensional) system state to a non-negative real number. Define a “stage- $j$  run” to be a simulation run that starts at level  $(j-1)$  and proceeds until either hitting level  $j$  or returning to level 0. Let  $N_j$  be the number of stage- $j$  runs (a decision variable). Let  $p_j$  be the probability that a stage- $j$  run reaches level  $j$  (before returning to 0). Let  $b_j$  be the average computing time to conduct a stage- $j$  run. The objective is to minimize the variance of the rare-event estimator  $\hat{\gamma}$  (defined as the product of individual estimators for  $p_j$ ) subject to a computing budget  $T$ :

$$\min_{N_1, N_2, \dots, N_m} \text{Var}(\hat{\gamma}) \quad \text{s.t.} \quad b_1 N_1 + b_2 N_2 + \dots + b_m N_m = T,$$

Shortle and Chen (2008) give an asymptotically optimal solution (as  $T \rightarrow \infty$ ) to this computing budget allocation problem:

$$N_1 \sqrt{\frac{1-p_1}{b_1 p_1}} = N_2 \sqrt{\frac{1-p_2}{b_2 p_2}} = \dots = N_m \sqrt{\frac{1-p_m}{b_m p_m}}.$$

The optimal allocation suggests that more replications should be made at stages that (a) are less expensive to run, and (b) have a lower probability of advancing to the next stage. This can be implemented in a sequential manner – after each replication, simulate the stage  $j$  with the lowest estimate for  $N_j \sqrt{(1-p_j)/b_j p_j}$  (where  $N_j$  is the number of stage- $j$  runs conducted so far in the simulation). This method is called Optimal Splitting Technique for Rare-Event simulation (OSTRE).

#### 4 SIMULATION STUDY DESCRIPTION

In this section we describe the details of the simulation study for the Atlanta Provider to Customer Edge T1 link mentioned previously. Two simulation methods are compared: standard simulation and optimal level splitting.

The simulations were performed in Visual Basic (VB). The linear congruential random number generator available in VB is not very robust (L’Ecuyer 2001). Instead of that VB generator, the random number generator package with multiple streams described in L’Ecuyer (2001) and L’Ecuyer et al. (2002) was implemented in VB and used for this simulation study.

The rare-event probability we wish to estimate is  $\gamma \equiv P(T_K < T_W)$ , where  $T_K$  is the first time there are  $K$  in the system, and  $T_W$  is the duration of the worm attack. The worm continues to infect the network,

while host computers in the network remain in the susceptible or infectious states. That is,  $\{T_R < T_W\}$  is the event that a buffer overflow occurs before the end of the attack, while host computers in the network remain in the susceptible or infectious states. A value of  $K = 80$  was used for the  $M(t)/D_3/1/K$  system. The value of  $K = 80$  was selected so the probability of a buffer overflow would be small. The simulations initially start at time zero, using the initial arrival rate (.275 ppms), shown earlier in Figure 2.

In the level-splitting implementation, four intermediate levels were used. Sensitivity of the results to evenly spaced levels versus unevenly spaced levels was investigated. Each simulation run continues while the number in the system is less than the current level  $i$ , or the end of the worm attack is reached.

For the cyber security application, the arrival process is non-stationary. OSTRE was previously applied only to stationary processes (Shortle and Chen 2008, Shortle et al. 2010). To extend it to non-stationary processes, we expand the state space by adding the time  $t$  as a new dimension. The state space is defined on  $(n, t)$ , where  $n$  is the number in the system. When each run terminates, the current state  $(n, t)$  is saved as a candidate starting state for future runs of the next level. This is a change from the application to a queue with stationary arrival rates, where the state is one-dimensional and only consists of the number in the system. For OSTRE, the starting state (both  $n$  and  $t$ ) is selected according to the distribution of saved starting states for the selected optimal level. For standard simulation, the levels are simulated sequentially, and the starting state (both  $n$  and  $t$ ) for the next level is the end state from the previous level.

As time progresses in the simulation, the appropriate arrival rate is used (see Figure 2), depending on the current simulation time. A time-based computing budget of five minutes per replication was used for both the standard simulation and optimal level splitting. The VB Timer function was used to track the usage of the computing budget, and the average simulation time for each level, which is needed to determine the asymptotically optimal allocation of runs for each level. Thirty replications of 5 minutes per replication were performed.

## 5 SIMULATION RESULTS

Simulation results from the case of evenly spaced levels (0, 20, 40, 60, and 80) are shown in Figures 5 and 6, and Table 1. The measures of interest are the expected value, variance, and relative error of the rare-event probability estimator  $\hat{\gamma}$ , that is  $E[\hat{\gamma}]$ ,  $\text{Var}[\hat{\gamma}]$ , and  $\text{RE}[\hat{\gamma}]$ . The relative error is the standard deviation of the estimator divided by its mean. Figure 5 shows that optimal level splitting results in a much lower estimate for the variance of the rare event probability than standard simulation.




Figure 5: Variance of Rare-Event Probability, with Evenly Spaced Levels

Figure 6 shows the allocation of the runs among the various levels for both methods. Optimal level splitting results in a much greater allocation of the runs to the highest simulation level (closest to the rare

event). Table 1 gives the expected value, variance, and relative error of the rare event probability for both methods. The optimal level splitting is clearly superior to the standard simulation.

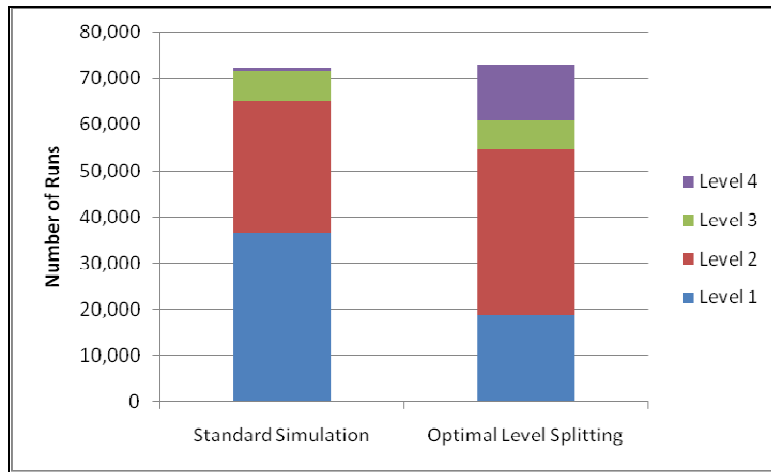


Figure 6: Run Allocation, with Evenly Spaced Levels

Table 1: Simulation Estimates, with Evenly Spaced Levels

Method	$E[\hat{\gamma}]$	$Var[\hat{\gamma}]$	$RE[\hat{\gamma}]$
Standard Simulation	8.70E-04	9.07E-07	1.09E+00
Optimal Level Splitting	6.82E-04	1.89E-07	6.37E-01

Simulation results from the case of unevenly spaced levels (0, 65, 70, 75, 80) are shown in Figures 7 and 8, and Table 2. Figure 7 shows that with the unevenly spaced levels, the variance of the rare event probability is lower for the optimal level splitting, showing that the OSTRE method will still do well even if poor selections are made for the levels. Figure 8 shows the allocation of the runs among the various levels for both methods. Standard simulation results in very few observations in levels 2, 3, and 4. Table 2 gives the expected value, variance, and relative error of the rare event probability for both methods. The optimal level splitting method is superior to the standard simulation in the case of poorly selected levels.

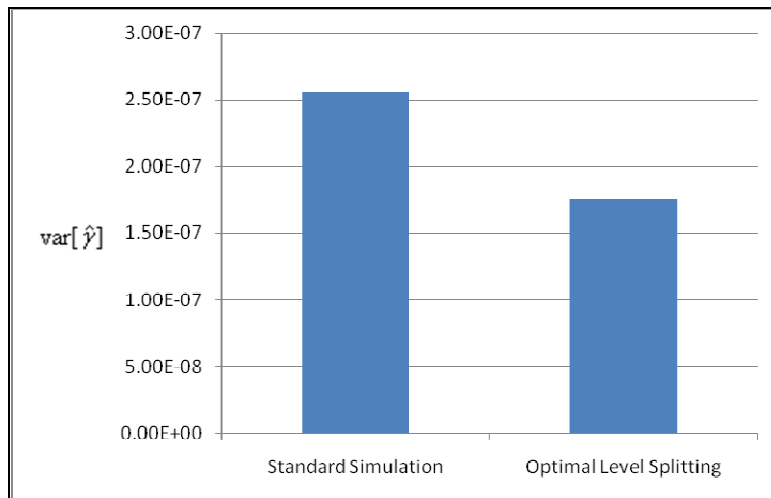


Figure 7: Variance of Rare-Event Probability, with Unevenly Spaced Levels

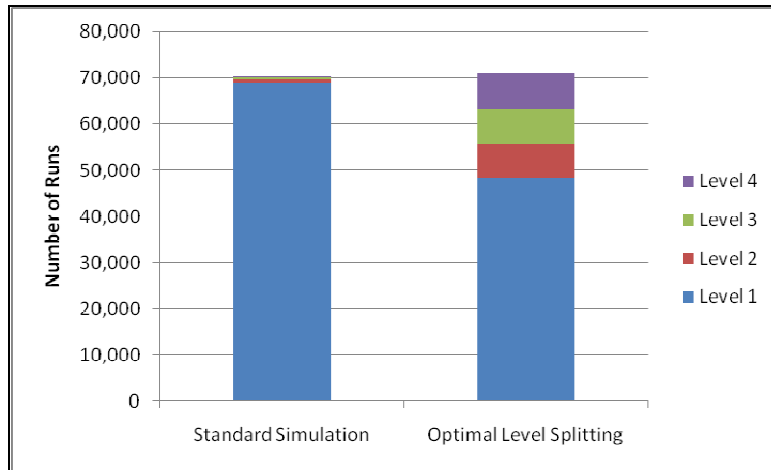


Figure 8: Run Allocation, with Unevenly Spaced Levels

Table 2: Simulation Estimates, with Unevenly Spaced Levels

Method	$E[\hat{\gamma}]$	$Var[\hat{\gamma}]$	$RE[\hat{\gamma}]$
Standard Simulation	1.25E-03	2.55E-07	4.05E-01
Optimal Level Splitting	1.11E-03	1.75E-07	3.76E-01

## 6 CONCLUSIONS AND NEXT STEPS

In this paper we have applied a splitting simulation method—OSTRE—to evaluate the performance of a link in an IP network during a cyber attack. To generate the traffic on the link, we used the worm model in IP-NPAT. The resultant link loading is non-stationary. Our work demonstrates that the optimal splitting method can be applied to systems with non-stationary arrival processes, such as this cyber attack scenario. The initial results indicate the optimal level splitting yields superior results to standard simulation in this application. Even when the intermediate levels are poorly selected, the method still provides a benefit compared with standard simulation.

This initial work is a proof of concept of the application of optimal splitting to systems with non-stationary arrival processes. One extension is to estimate the probability of being in the rare event state, rather than the probability of hitting the rare event. Future work will also focus on estimating rare event probabilities in a subset of the network consisting of a multi-link path between a particular origin and destination node pair, rather than focusing on a single link. Our longer term goal is to improve rare-event simulation efficiency in an entire network under cyber attack

## ACKNOWLEDGMENTS

This work has been supported in part by Department of Energy under Award DE-SC0002223. This research has also been funded by Noblis and is part of an ongoing research partnership between George Mason University and Noblis known as the Center for Network Based Systems, see <http://www.noblis.org/cnbs/index.htm>.



## DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

## REFERENCES

- Allen, J. 2004. Information Security as an Institutional Priority, Software Engineering Institute, Carnegie Mellon University, <[www.cert.org/archive/pdf/info-sec-ip.pdf](http://www.cert.org/archive/pdf/info-sec-ip.pdf)>.
- CERT/CC, CERT Statistics (Historical), <<http://www.cert.org/stats/>>.
- CERT/CC, Overview of Attack Trends, <[www.arccert.gov.ar/webs/textos/attack\\_trends.pdf](http://www.arccert.gov.ar/webs/textos/attack_trends.pdf)>.
- Chen, C. H., and D. He. 2005. Intelligent Simulation for Alternatives Comparison and Application to Air Traffic Management, *Journal of Systems Science and Systems Engineering*, Vol. 14, No. 1, pp. 37-51.
- Chen, C. H., D. He, M. Fu, and L. H. Lee. 2008. Efficient Simulation Budget Allocation for Selecting an Optimal Subset, *Infornis Journal on Computing*, Vol. 20, No. 4, pp. 579-595.
- Chen, C. H., E. Yücesan, L. Dai, and H. C. Chen. 2010. Efficient Computation of Optimal Budget Allocation for Discrete Event Simulation Experiment, *IIE Transactions*, Vol. 42, No. 1, pp. 60-70.
- Clarke, R. 2010. *Cyber War: The Next Threat to National Security and What to Do About It*, Ecco.
- The Impact of Cyber Attacks on Internet Protocol (IP) Networks. 2004. *Sigma*, Noblis, Falls Church, VA, Winter.
- Internet Protocol (IP) Network Performance and Analysis Tool (IP-NPAT), Scientific and Technical Report, Contract Number: NBCH-D-02-0039 CDRL #: 007 / DID#: 80711, Technology and Programs Division (N2), 701 South Courthouse Road Arlington, VA 22204, January 2005.
- Kleinrock, L 1976. *Queueing Systems, Volume II: Computer Applications*, Wiley.
- L'Ecuyer, P. 2001. Software for Uniform Random Number Generation: Distinguishing the Good and the Bad. In *Proceedings of the 2001 Winter Simulation Conference*, ed. B. A. Peters, J. S. Smith, D. J. Medeiros, M. W. Rohrer, 95-105. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- L'Ecuyer, P., R. Simard, E.J. Chen, and W.D. Kelton. 2002. An Object-Oriented Random-Number Package with Many Long Streams and Substreams, *Operations Research* 50, 1073-1074.
- Liu, S. and B. Cheng. 2009. Cyber attacks: Why, What, Who, and How, *IT Pro*, May/June 2009.
- Masi, D. M. B., M. J. Fischer, and D. A. Garbin. 2004a. Modeling Internet Protocol Networks: OPNET and Analytics, *Sigma: Noblis Technology Summaries*, pp. 52-63, Winter. <<http://www.noblis.org/NewsPublications/Publications/TechnicalPublications/SigmaJournal/Pages/SigmaWinter2004.aspx>>.
- Masi, D. M. B., M. J. Fischer, and D. A. Garbin. 2004b. Analytically Modeling Cyber Attacks in Internet Protocol Networks, *2004 Advanced Simulation Technologies Conference*, pp. 69-74, Arlington, VA, April 18-22.
- Masi, D. M. B. and M. J. Fischer. 2005. Analytically Modeling Worm Attacks in Internet Protocol Networks, *Ninth INFORMS Computing Society (ICS) Conference*, Annapolis, MD, January 5-7.
- Mathematical Challenges in Cyber security. 2009. SANDIA Report 2009-0805, February.

- Mathematical Underpinnings for Science-Based Cybersecurity. 2008. *U.S. Department of Energy White Paper*,  
<<https://wiki.cac.washington.edu//download/attachments/7479040/doecybermath25feb08.doc>>.
- Nicol, D.M., M. Liljenstam, and J. Liu. 2003. Multiscale Modeling and Simulation of Worm Effects on the Internet Routing Infrastructure. *The 13th International Conference on Modelling Techniques and Tools for Computer Performance Evaluation (Performance TOOLS 2003)*, Urbana, IL, September 2-5.
- Shortle, J., C.H. Chen, M.J. Fischer and D.M. Masi. 2010. An Effective Rare-event Simulation Technique and Its Application to Cyber Security and Electric Grid, DHS Conference MSAHS 2010, Arlington, VA, March.
- Shortle, J., and C.H. Chen. 2008. A Preliminary Study of Optimal Splitting for Rare-Event Simulation In *Proceedings of the 2008 Winter Simulation Conference*, eds. S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, J. W. Fowler, 266-272. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Shortle, J.F., C.H. Chen, A. Brodsky, and D. Brod. 2010. Optimal Level Splitting for Rare-Event Simulation. Submitted to *IIE Transactions*.
- Zou, C.C., W. Gong, and D. Towsley 2002. Code Red Worm Propagation Modeling and Analysis. *9th ACM Conference on Computer and Communication Security (CCS)*, Nov. 18-22, Washington DC, USA.
- Zou, C.C., W. Gong, and D. Towsley. 2003. Worm Propagation Modeling and Analysis under Dynamic Quarantine Defense. *ACM CCS Workshop on Rapid Malcode (WORM'03)*, Oct. 27, Washington DC, USA.
- Zou C.C., D. Towsley and W. Gong. 2004. Email Virus Propagation Modeling and Analysis. *Umass ECE Technical Report TR-CSE-03-04*.
- Zou C.C., D. Towsley and W. Gong. 2006. On the Performance of Internet Worm Scanning Strategies. *Performance Evaluation*, 63, pages 700-723.

## AUTHOR BIOGRAPHIES

**MARTIN J. FISCHER** is a Senior Fellow in National Security and Intelligence at Noblis. He has 40 years of experience in the field of network design and performance analysis of telecommunications systems. This experience includes 25 years with the Defense Information Systems Agency and 15 years with Noblis. Until recently he was an adjunct professor at George Mason University and is a team member with faculty at George Mason University that has received two National Science Grants and one from the Department of Energy. Over his career he has published or presented over 200 papers, approximately 50 of which have appeared in refereed journals. He received a doctorate in Operations Research from Southern Methodist University. His email address is <[mfischer@noblis.org](mailto:mfischer@noblis.org)>.

**DENISE M. B. MASI** is a Fellow in National Security and Intelligence at Noblis. Her experience and research interests include queueing theory and simulation applied to telecommunications networks. Prior to joining Noblis in 1998, Dr. Masi worked in statistical analysis and modeling for the A.C. Nielsen Company. Dr. Masi received her B.S. in Industrial Engineering from Texas A&M University and an M.S. in Industrial Engineering from Purdue University. She received her Ph.D. in information technology and engineering, with a concentration in operations research, at George Mason University. Her email address is <[dmasi@noblis.org](mailto:dmasi@noblis.org)>.

**JOHN F. SHORTLE** is an Associate Professor of Systems Engineering at George Mason University. His experience includes developing stochastic, queueing, and simulation models to optimize networks and operations. His research interests include simulation and queueing applications in telecommunications, air

transportation, and energy. He received his doctorate degree in operations research from UC Berkeley. His email address is <[jshortle@gmu.edu](mailto:jshortle@gmu.edu)>.

**CHUN-HUNG CHEN** is a Professor of Systems Engineering & Operations Research at George Mason University. Dr. Chen has led research projects in stochastic simulation and optimization, systems design under uncertainty, and air traffic management, which are sponsored by NSF, FAA, and NASA. Dr. Chen received the Kayamori Best Automation Paper Award from the 2003 IEEE International Conference on Robotics and Automation, 1994 Eliahu I. Jury Award from Harvard University, and the 1992 MasPar Parallel Computer Challenge Award. He is serving on the editorial boards of IEEE Transactions on Automatic Control, IIE Transactions, *Journal of Simulation Modeling Practice and Theory*, and *International Journal of Simulation and Process Modeling*. He received his Ph.D. degree from Harvard University. His email address is <[cchen9@gmu.edu](mailto:cchen9@gmu.edu)>.