Proceedings of the 2011 Winter Simulation Conference
S. Jain, R.R. Creasey, J. Himmelspach, K.P. White, and M. Fu, eds.

SIMULATING NETWORK CYBER ATTACKS USING SPLITTING TECHNIQUES

Denise M.B. Masi
Martin J. Fischer

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

John F. Shortle

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

Chun-Hung Chen

Department of Electrical Engineering
National Taiwan University
Taipei, TAIWAN

ABSTRACT

As a result of potential damage to our national infrastructure due to cyber attacks, a number of cyber-security bills have been introduced in Congress and a National Strategy for Trusted Identities in Cyber-space has been developed by the White House; a component of this strategy is the development of models to assess risks due to cyber incidents. A worm attack on a network is one type of attack that is possible. The simulation of rare events, such as the occurrence of a catastrophic worm attack, is impractical without special simulation techniques. In this paper we present an application of splitting methods to estimate rare-event probabilities associated with the propagation of a worm through a network. We explore the sensitivity of the benefits of splitting methods, as compared to standard simulation, to the rarity of the event and the level function used.

1 INTRODUCTION

Cybersecurity has become a national priority. 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 2011 survey by Unisys Corp. showed that Americans are feeling less secure about Internet security versus fall 2010, with an increase by 35 percent in the Internet Security Index, which was a greater increase by far than the other security areas measured (national, financial, and personal) (Jackson 2011). In addition to private citizens, the federal government has concerns about Internet security, and legislation to ensure more effective methods of securing federal networks has been in discussion since 2010 (Kash 2010). There are several versions of proposed cybersecurity legislation being considered, but with a number of controversial aspects (e.g., a "kill switch" allowing the President to shut down the Internet during emergencies, and imposition of controls on privately owned infrastructure), it has been said that "chances for passage of comprehensive federal cybersecurity legislation appear to be fading" (Jackson 2011). Another bill in the House would allow the Department of Defense to conduct clandestine operations in cyberspace against another country (Hardy 2011). Although the Obama administration has concerns about this aspect of the bill, this is an example where modeling cyber operations to predict the impact would be essential. The National Strategy for Trusted Identities in Cyberspace was developed by the White House in April

2011; a component of this strategy is the development of models to assess risks due to cyber incidents (The White House 2011).

A worm attack on an Internet Protocol (IP) network is one type of attack that is possible. Worms are "malicious programs that exploit some defect in a computer's software to implant a copy of the worm, and use the newly infected host as a platform to seek and infect other victims" (Nicol 2008). Two major classes of worms, scan-based worms, and email worms have frequently attacked computer networks. Well-known worms such as Code Red, Code Red II, Slammer, Blaster, and Sasser, were all scan-based worms. We focus on modeling scan-based worms in this paper. Scan-based worms spread through computer networks by searching, attacking, and infecting remote computers automatically (Zou et al. 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. Scans can be performed either randomly, as was the case in many of the worms mentioned above, or in a preferential fashion whereby a worm will choose targets close to itself with higher probability than distant targets. We focus on random scanning in this paper.

Because computers involved in a worm attack generate a large number of scanning packets, discrete-event simulation of such attacks can be difficult due to the demands imposed with generating many events. There is a large body of literature that uses traditional epidemic models based on differential equations to model worms (e.g., Liljenstam et al. 2002; Staniford et al. 2002; Zou et al. 2002). However, it is important to be able to simulate worms, as analytic approaches do not capture the variability that is a feature of worms. This paper focuses on estimating the probabilities of rare events pertaining to worm infections, such as a probability that a high percentage of network nodes become infected. Nicol (2008) describes efficient hybrid discrete-continuous approaches to simulating worms. However, unlike this paper, Nicol does not address estimation of rare-event probabilities related to worms, and does not incorporate possible repair of infected host computers.

In this paper we apply a splitting approach to more efficiently simulate the propagation of the worm through a network and estimate rare event probabilities associated with the worm. When evaluating rare events, the number of simulation runs required to achieve a reasonable confidence interval can be prohibitively high, requiring use of variance reduction techniques. We report on the use of splitting as compared with conventional simulations. This paper builds on our initial application of splitting to worm simulation modeling in Fischer et al. (2010). In Shortle and Chen (2010) we utilize splitting techniques to analyze power grid problems.

In Section 2 we formally define the problem and discuss the 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

We model the propagation of a worm through a network. Using the traditional epidemic model notation (Hethcote 2000), the possible states for a host computer are assumed to be:

- Susceptible (S) – a host computer that is susceptible to infection;
- Infectious (I) – a host that is infected and can infect other hosts;
- Removed (R) – a host that is removed from the infectious populations by repair; a host in state R cannot leave it.

Figure 1 depicts the possible states and flows between states that are typically assumed in the SIR model for host computers infected by a worm. Susceptible hosts can become infected when scanned by an infectious computer, and then become infectious themselves. Infectious host computers can be repaired, and are then removed from the infectious population and cannot become infected again. Our assumption is

that only the infectious computers are repaired, although one could also repair and remove the susceptible computers if desired.

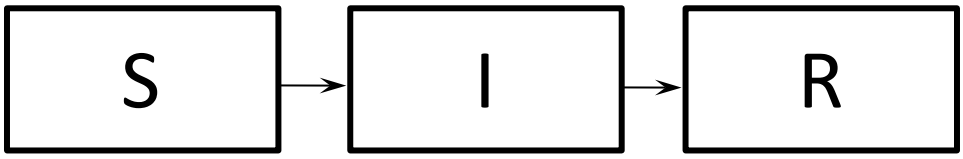

Figure 1: SIR state diagram

Our assumptions on the worm propagation are as follows. We assume that there are N host computers in the network, and initially the number of susceptible hosts is $S_0 < N$, due to the susceptibility of some operating systems and not others, or patching in advance that may have occurred on some host computers. Each infectious host scans with rate λ, and picks a host at random to scan. It might scan a repaired host in which case nothing happens, or a susceptible host in which case it becomes infected, or an infected host in which case nothing happens. Each infectious host gets repaired with rate μ. All event times are exponentially distributed (see Nicol 2006 for justification), so this is a continuous time Markov chain. The state space is defined as $X(t) = (R(t), I(t))$, where $R(t)$ is the number of repaired hosts at time t, $I(t)$ is the number of infected hosts at time t, and $S(t) \equiv N - R(t) - I(t)$ is the number of susceptible hosts at time t. Initially, one host computer is infectious ($I(0) = 1$). The exponential assumption is not required for the simulation, but it is helpful in checking the simulation against analytical results that can be derived via analysis of continuous time Markov chains.

We are interested in the rare event that x% of the susceptible network nodes become infected (where x% could be 100%) before the worm propagation terminates. Effectively, once more than 100-x% of the susceptible nodes are removed, the rare event can no longer occur. Our simulation model estimates the expected value and variance of this probability.

3 SIMULATION METHODOLOGY

We employ a splitting approach to improve the efficiency of rare-event simulation. Importance sampling is another technique used to more efficiency simulate rare events. The basic idea of splitting is to create separate copies of the simulation whenever it gets close to the rare event (Figure 2). 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 et al. (2011) referred to as Optimal Splitting Technique for Rare-Event simulation (OSTRE). Similar to the idea of Optimal Computing Budget Allocation (e.g., Chen et al. 2000, 2008, 2010), OSTRE intends to maximize the efficiency of splitting for rare-event simulation. We also investigate an equal-allocation implementation (e.g., L'Ecuyer et al. 2006).

Our initial examination in Fischer et al. (2010) focused on simulation of the congestion on a network link as a worm spreads attack traffic through the network and an initial simulation of worm propagation through a network, both using OSTRE. This paper extends that work by studying the sensitivity of the benefit of splitting to the number and location of the levels and also examining equal-allocation splitting.

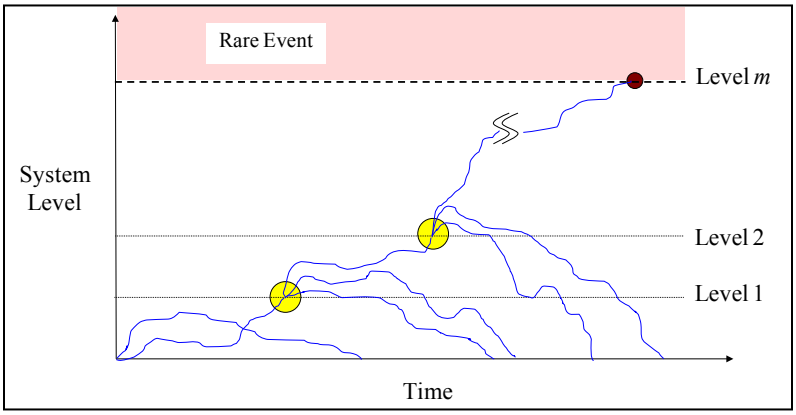

Figure 2: Level splitting

Figure 2 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. The total number of levels is denoted by *m*. 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*:

$$\min_{N_1,N_2,..,N_m} \text{Var}(\hat{\gamma}) \quad \text{s.t.} \quad b_1N_1 + b_2 N_2 + \ldots + b_m N_m = T.$$

Shortle et al. (2011) give an asymptotically optimal solution (as $T\to\infty$) to this computing budget allocation problem. The result holds under the condition that the success probability of a stage-*j* run is independent of the starting state from level *j*-1:

$$N_1\sqrt{\frac{b_1p_1}{1-p_1}} = N_2\sqrt{\frac{b_2p_2}{1-p_2}} = \cdots = \; N_m\sqrt{\frac{b_mp_m}{1-p_m}}$$

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 $N_j\sqrt{b_jp_j/(1-p_j)}$ (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).

The assumption made to derive this result – namely that the success probability of a stage-*j* run is independent of the starting state – does not generally hold for the worm-attack model. This is because the probability of reaching a certain number of infected computers depends on both the number of currently infected computers (the number in the set *I*) as well as the number of currently repaired computers (the number in the set *R*). In the special case when the rare event of interest is that *all* nodes become infected, this assumption does hold. This is because the problem state space becomes one dimensional (the number in the set *I*). The number in the set *R* essentially does not matter since once any computer is repaired (the number in *R* is non-zero), the rare event can no longer occur, so the simulation can terminate at that point. All the results in this paper are for this simpler case, so it should be noted that the OSTRE allocation may

not work as well when the problem of interest is to determine the probability that x% become infected where $x\% < 100\%$.

When $b_j = 1$ and $p_j = \gamma^{1/m}$, equal-allocation splitting and optimal splitting are the same (Shortle et al. 2011). The variance is reduced by a factor $m^2 \gamma^{1-1/m}$, as compared with standard simulation. This results in an optimal approximate number of levels: $m = -\ln(\gamma)/2$ (e.g., L'Ecuyer et al. 2006).

4 SIMULATION STUDY DESCRIPTION

In this section we describe the details of the simulation study. Three simulation methods are compared: standard simulation, equal-allocation splitting, and optimal splitting.

The simulations are 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) is implemented in VB and used for this simulation study.

The rare-event set is defined as $K = \{(r, i)$ such that $i \geq a$, where $a \leq S_0\}$, and S_0 is the initial number of susceptible nodes. We define the level function $h(r, i) = i$. The rare event probability we wish to estimate is $\gamma = \Pr\{h(X(t)) \geq a$ before $h(X(t)) = 0\}$. For the case of $a = S_0$, which is the primary focus in this paper, the system can never reach the rare event once a single susceptible node is repaired ($r > 0$).

In the level-splitting implementations, initially five intermediate levels are used, where the levels are evenly spaced on $(0, S_0)$. The sensitivity of the benefits of splitting methods as compared to standard simulation to the rarity of the event is first explored by varying the repair rate, μ.

Secondly, optimally selecting the number and location of levels based on the theory provided in Section 3 is investigated. The number of levels is selected equal to $m = -\ln(\gamma)/2$, and the location of the levels is selected so that the probability of each level was approximately equal to e^{-2}.

To assist in this, we can formulate the model as a continuous time Markov chain, since all event times are exponentially distributed. The continuous time Markov chain can be used to obtain the rare event probability of interest (γ) directly. However, as analytic approaches do not capture the variability that is a feature of worms, we use the Markov chain to facilitate in the selection of the optimal number and location of levels for the simulation. The continuous time Markov chain is used to quickly obtain estimates of the probability of i infected before returning to zero infected given that i-1 was infected, for all $i = 1, 2, \ldots, S_0$. That is, $q_i = \Pr\{h(X(t)) \geq i$ before $h(X(t)) = 0) \mid h(X(t)) \geq i$-1 before $h(X(t)) = 0)\}$, for all $i = 1, 2, \ldots, S_0$. These probabilities were computed before simulating the system. These are similar to the p_j (the probability that a stage-j run reaches level j before returning to 0), but are the probabilities for each of the $i = 1, 2, \ldots, S_0$ rather than the probabilities for the levels. Alternatively, these estimates could be obtained with initial simulation runs, but the analytic model can be run more quickly. The product of these q_i across all i gives an estimate of γ, which is used to compute the optimal number of levels ($-\ln(\gamma)/2$). Then, the q_i are used to determine the location of the levels so that the probabilities of each of the levels is approximately e^{-2}. The simulation is then run using these specifications.

The following parameter values are used in these experiments:

- Number of host computers in the network $N = 100$
- Initial number of susceptible hosts $S_0 = 75$
- Worm scan rate $\lambda = .1$
- Repair rate is varied: $\mu = .0005$, .00275, and .005
- Twenty replications of runs of five minutes each are performed.

5 SIMULATION RESULTS

Simulation results using five levels that are evenly spaced on $(0, S_0)$ are shown in Table 1 for each of the three simulation methods and repair rates. The measures of interest are the sample mean, sample variance,

and relative error of the rare-event probability estimator $\hat{\gamma}$. The relative error is the sample standard deviation of the estimator divided by its sample mean. The simulation results in Table 1 for the estimated gamma were found to be consistent with the Markov chain results. Table 1 shows that for the cases of μ = .005 and .00275, where the probability that all hosts become infected is rare, both splitting methods result in a much lower estimate for the variance of the rare event probability than standard simulation. In fact, for μ = .005, there are no observations of the rare event with standard simulation. For μ = .0005, the probability that all hosts become infected is much less rare (about .1), and splitting is less beneficial as compared to standard simulation. Figures 3 and 4 display these results graphically for μ = .00275, where the much smaller variance with splitting is evident as well as the resulting narrower confidence intervals. These confidence intervals are approximate, as they are based on the normal distribution, which may not apply to rare events.

Table 1: Simulation estimates with levels evenly spaced on $(0, S_0)$.

	Standard Simulation			**Equal-Allocation Splitting**			**Optimal Level Splitting**		
μ	**Sample Mean**	**Sample Var**	**Relative Error**	**Sample Mean**	**Sample Var**	**Relative Error**	**Sample Mean**	**Sample Var**	**Relative Error**
0.0005	1.05E-01	8.76E-06	2.81E-02	1.04E-01	1.97E-06	1.35E-02	1.05E-01	3.95E-06	1.90E-02
0.00275	2.91E-05	1.12E-09	1.15E+00	3.16E-05	6.78E-11	2.60E-01	3.40E-05	2.00E-11	1.32E-01
0.005	No Observations			4.98E-08	1.98E-15	8.95E-01	4.99E-08	7.22E-16	5.39E-01

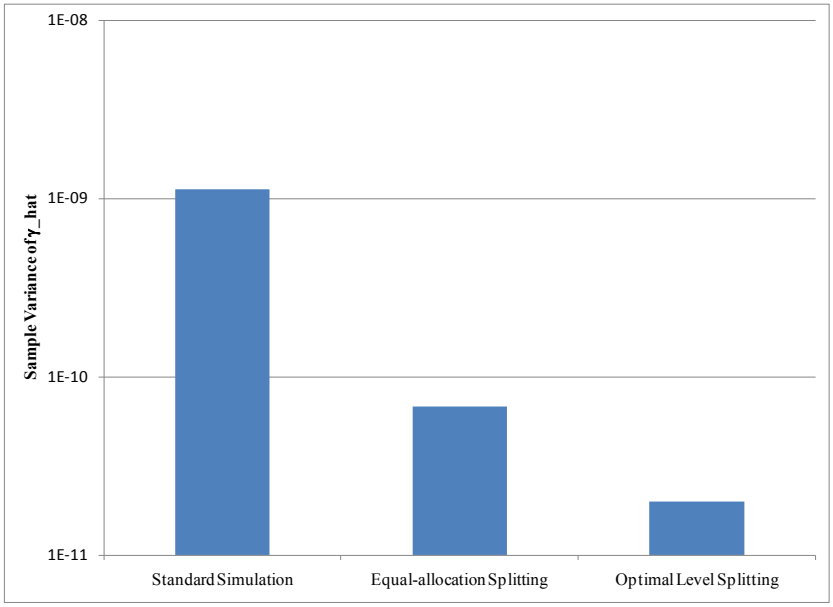

Figure 3: Variance of probability that all susceptible nodes get infected, μ=.00275

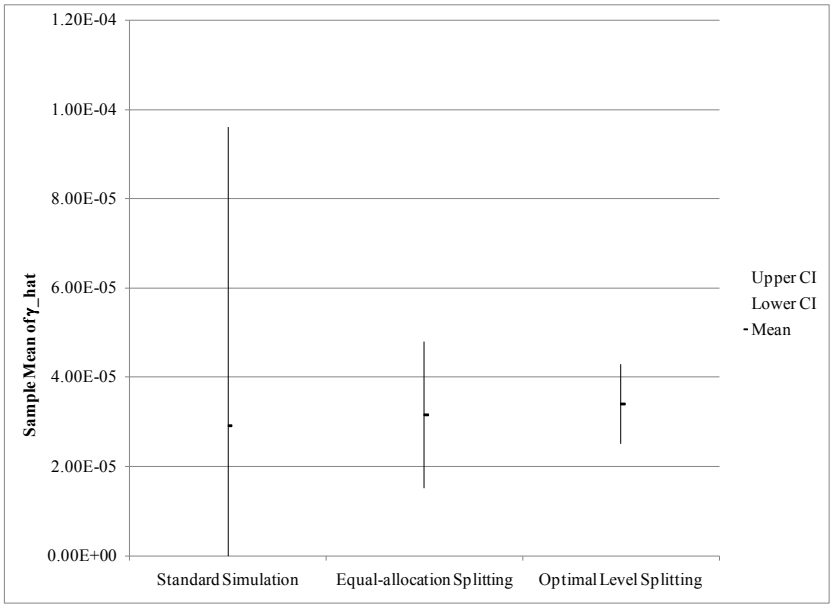

Figure 4: Probability that all susceptible nodes get infected, μ=.00275

Figure 5 shows the allocation of runs to each of the levels for the three methods. Optimal level splitting results in a much greater allocation of the runs to the highest simulation level closest to the rare event, as it essentially multiplies promising runs that are more likely to reach the rare event.

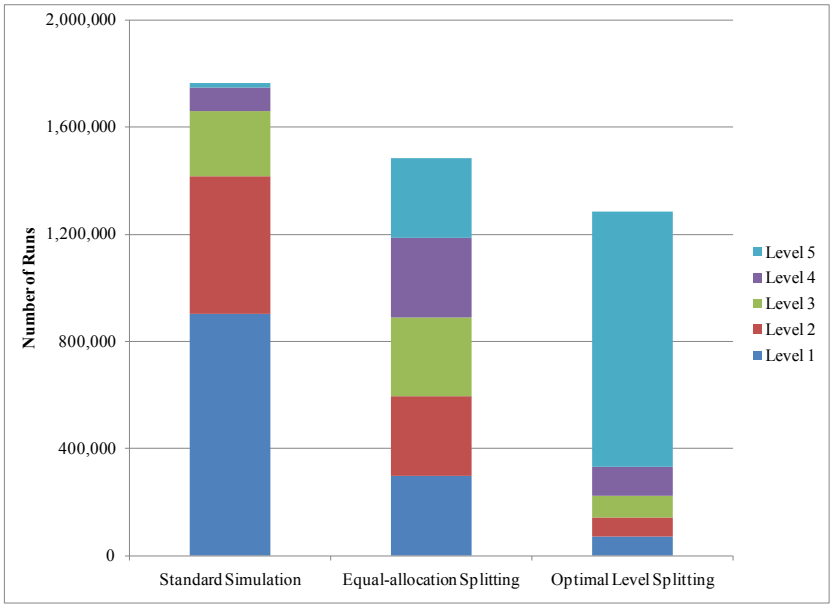

Figure 5: Run allocation, μ=.00275

Next, selecting the number and location of levels based on the theory provided in Section 3 is investigated. We use the continuous time Markov chain model to obtain initial estimates of the rare event probability and optimal number of levels. The estimate of the probability that all S_0 nodes get infected from the analytic model is used to compute the optimal number of levels, m, as seen in Table 2. Scenarios where the event of interest is more rare optimally would have a greater number of levels. Figure 6 shows the q_i for $i = 1, 2, \ldots S_0$ (solid lines) for each of the three repair rates that we studied. The locations of the levels

chosen so that the probability of each level (determined by taking the product of the q_i up to potential level boundaries) is close to e^{-2} is displayed with the dashed lines in Figure 6. Figure 7 shows the differences in the probabilities of the levels in the two configurations, for the μ=.00275 case. Figure 8 shows that the benefits of splitting are greatly increased when the levels are chosen so that the probabilities of the levels are close to e^{-2} for the μ=.00275 case. Figure 9, for the case of μ=.005, also shows that the splitting is more beneficial when the levels are more carefully chosen so that the probabilities of the levels are close to e^{-2}. In either method of selecting the levels, however, using splitting is more efficient than standard simulation.

Table 2. Optimal number of levels.

μ	Optimal Number of Levels
0.0005	1
0.00275	5
0.005	8

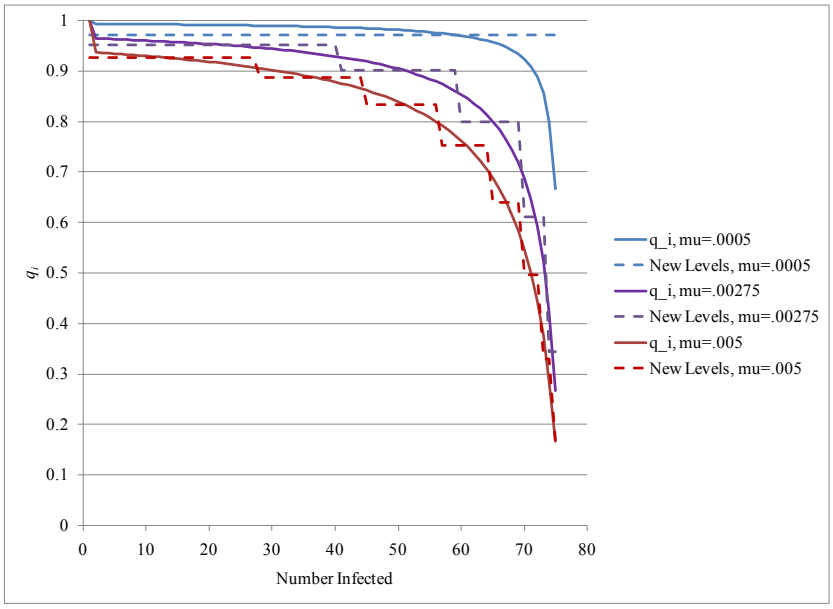

Figure 6: q_i for $i = 1, \ldots, S_0$, from continuous time Markov chain

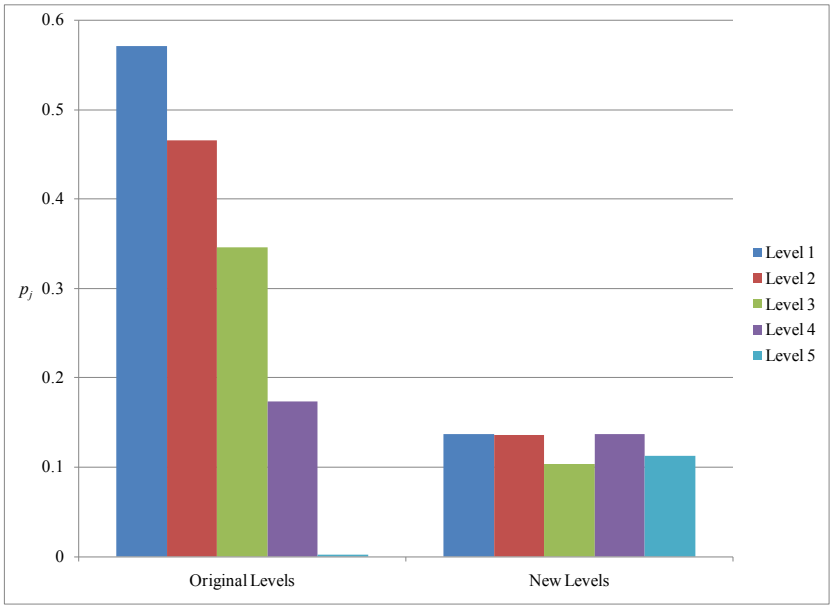

Figure 7: Probabilities p_j that a run reaches level j (before returning to 0), $\mu = .00275$

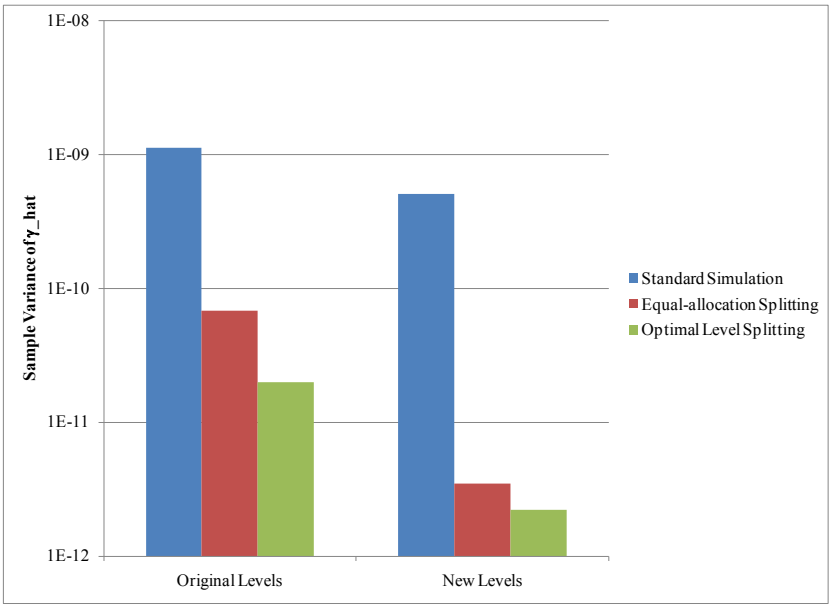

Figure 8: Sample variance of rare event probability with different level locations, μ=.00275

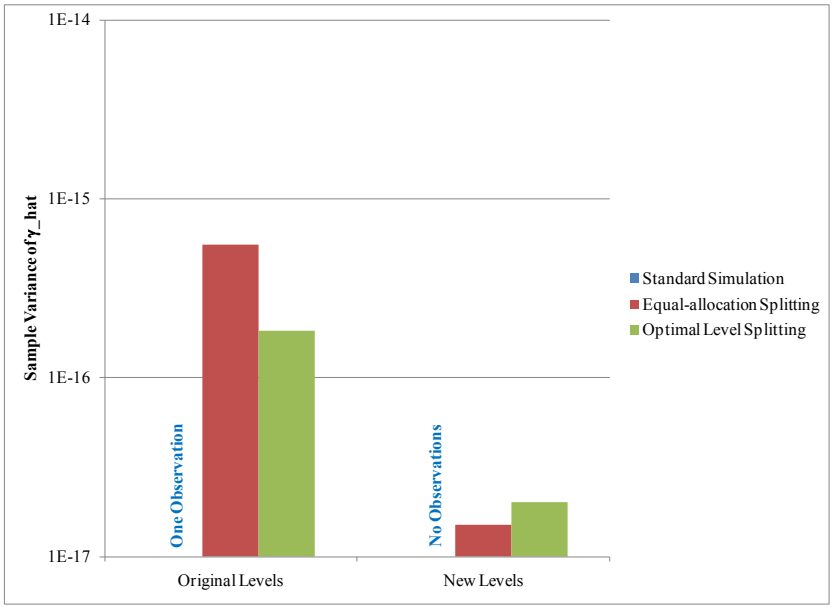

Figure 9: Sample variance of rare event probability with different level locations, μ=.005

6 CONCLUSIONS AND NEXT STEPS

Cyber attacks are of increasing national concern. In this paper we have applied two splitting simulation methods — equal-allocation and optimal splitting —to evaluate the propagation of a worm in a network during a cyber attack. The rare event that all susceptible host computers become infected was estimated. Both splitting methods are more efficient than standard simulation at estimating rare event probabilities, although the optimal splitting method (OSTRE) is superior to equal-allocation splitting in all cases investigated. Methods of selecting the levels were investigated. When analytics (continuous time Markov chain) can be employed to get estimates of the number of hosts infected, that provides a method of easily selecting the best level locations resulting in a low variance estimator of the rare event.

Future work will also explore the efficiency of alternate level functions when simulating worms. Additional research on the use of the Markov chain approach for worm attack modeling, for instance to compute additional measures of interest, will be explored. Lastly, incorporation of a network topology that will enable modeling other types of worms utilizing preferential rather than random scanning will be investigated.

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

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

Chen, C. H., J. Lin, E. Yücesan, and S. E. Chick. 2000. "Simulation Budget Allocation for Further Enhancing the Efficiency of Ordinal Optimization," *Journal of Discrete Event Dynamic Systems: Theory and Applications*, 10, 251-270.

Chen, C. H., D. He, M. Fu, and L. H. Lee. 2008. "Efficient Simulation Budget Allocation for Selecting an Optimal Subset," *Journal on Computing*, 20(4), 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*, 42(1), 60-70.

Clarke, R. 2010. *Cyber War: The Next Threat to National Security and What to Do About It*, Ecco.

Fischer, M.J., Masi, D.M. B., Shortle, J.F., and Chen, C.H. 2010. "Simulating Non-Stationary Congestion Systems Using Splitting with Applications to Cyber Security." In *Proceedings of the 2010 Winter Simulation Conference,* Edited by B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yücesan, Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Hardy, Michael. 2011. "Administration Questions Military Role in Cyberspace". *Federal Computer Week*, May 24.

Hethcote, H.W. 2000. "The Mathematics of Infectious Diseases." *SIAM Review*, 42(4), 599 – 653.

Jackson, William. 2011. "Fear Factor: Americans Feeling Less Secure About Nearly Everything, Survey Shows." *Government Computer News*, May 4.

Kash, Wyatt. 2010. "A Cyber Bill Worth Enacting." *Government Computer News*, June 18.

L'Ecuyer, P. 2001. "Software for Uniform Random Number Generation: Distinguishing the Good and the Bad." In *Proceedings of the 2001 Winter Simulation Conference*, edited by B. A. Peters, J. S. Smith, D. J. Medeiros, and 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.

L'Ecuyer, P., V. Demers and B. Tuffin. 2006. "Splitting for rare-event simulation."In *Proceedings of the 2006 Winter Simulation Conference,* edited by L. R. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, 137-148. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.. IEEE, Piscataway, NJ.

Liljenstam, M., Y. Yuan, B. Premore, and D. Nicol. 2002. "A Mixed Abstraction Level Simulation Model of Large-Scale Internet Worm Infestations", *in Proceedings of the Tenth IEEE/ACM Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)*, IEEE Computer Society Press, Fort Worth, TX.

Nicol, D. M. 2006. "The Impact of Stochastic Variability on Worm Detection". In *Proceedings of ACM WORM 2006* (Fairfax, VA). 57–64.

Nicol, D.M. 2008. "Efficient simulation of internet worms." *ACM Transactions on Modeling and Computer Simulation*, 18(2):5:1–5:32.

Shortle, J., and C.H. Chen. 2008. "A Preliminary Study of Optimal Splitting for Rare-Event Simulation", *Proceedings of 2008 Winter Simulation Conference*, edited by 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. and C.H. Chen**,** 2010. "Optimal Level Splitting for Rare-event Simulation," Presentation given at the 2010 Department of Energy Applied Mathematics Program Meeting, May 3-5, Berkeley, CA.

Shortle, J.F., C.H. Chen, A. Brodsky, and D. Brod. 2011. "Optimal Level Splitting for Rare-Event Simulation". To appear in *IIE Transactions*.

Staniford, S., V. Paxson, and N. Weaver. 2002. "How to Own the Internet in Your Spare Time". *Proceedings of the USENIX Security Symposium*, 2002.

The White House, 2011. "National Strategy for Trusted Identities in Cyberspace", April.

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.

Zou, C.C., D. Towsley, and W. Gong. 2006. "On the Performance of Internet Worm Scanning Strategies", *Performance Evaluation*, 63(7), p.700-723, July.

AUTHOR BIOGRAPHIES

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*.

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*.

JOHN F. SHORTLE is an Associate Professor of Systems Engineering and Operations Research at George Mason University. His research interests include simulation and queueing applications in air transportation, telecommunications, and energy. Previously, he worked at US WEST Advanced Technologies. He received his doctorate degree in industrial engineering and operations research from UC Berkeley in 1996. His email address is *jshortle@gmu.edu*.

CHUN-HUNG CHEN is a Professor of Electrical Engineering at National Taiwan 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. He served as Co-Editor of the Proceedings of the 2002 Winter Simulation Conference and Program Co-Chair for 2007 INFORMS Simulation Society Workshop. 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 in 1994. His email address is *cchen9@gmu.edu*.