# Model-Based Evaluation: From Dependability to Security

David M. Nicol, *Fellow*, *IEEE*, William H. Sanders, *Fellow*, *IEEE*, and Kishor S. Trivedi, *Fellow*, *IEEE*

**Abstract**—The development of techniques for quantitative, model-based evaluation of computer system dependability has a long and rich history. A wide array of model-based evaluation techniques is now available, ranging from combinatorial methods, which are useful for quick, rough-cut analyses, to state-based methods, such as Markov reward models, and detailed, discrete-event simulation. The use of quantitative techniques for security evaluation is much less common, and has typically taken the form of formal analysis of small parts of an overall design, or experimental *red team*-based approaches. Alone, neither of these approaches is fully satisfactory, and we argue that there is much to be gained through the development of a sound model-based methodology for quantifying the security one can expect from a particular design. In this work, we survey existing model-based techniques for evaluating system dependability, and summarize how they are now being extended to evaluate system security. We find that many techniques from dependability evaluation can be applied in the security domain, but that significant challenges remain, largely due to fundamental differences between the accidental nature of the faults commonly assumed in dependability evaluation, and the intentional, human nature of cyber attacks.

**Index Terms**—Dependability evaluation, security evaluation, performability evaluation, stochastic modeling.

## 1 Introduction

Computer system and network security is an issue of increasing practical concern and research attention. As a research discipline, computer security is a venerable one with its own culture, assumptions, and language. The increased emphasis on system security has brought new researchers from different backgrounds to the field, bringing different perspectives and different skillsets. All of this is for the good since new viewpoints can lead to new insights.

Most of the older work in computer security focused on details in complex protocols or details in complex systems, for the simple reason that the root causes of security gaps are often found in the failures associated with such details. Later work expanded the attention to system-level security, that is, to the study of how systems can be designed to be secure in the sense that they perform their intended function in spite of possible malicious attacks. New work is examining intrusion tolerance (e.g., [1], [2], [3], [4]), which is a means of designing systems to continue to perform their intended function in spite of partially successful attacks.

No system-level methodology currently exists that can quantify the amount of security provided by a particular system-level approach. So far, most attempts at validation of security have been qualitative, focusing more on the process used to build a system that should be secure. Since it is impossible in practice to build a perfectly secure system, it is important to be able to quantitatively validate the efficacy of systems intended to be secure. Efforts aimed at quantitative validation of security have usually been based on formal methods (e.g., [5]), or have been informal using "red teams" to try to compromise a system (e.g., [6]). Both approaches, while being valuable in identifying system vulnerabilities, have their limitations, especially when they are applied to large systems. Only recently, due to efforts led by researchers accustomed to quantifying other system measures, have attempts been made to quantify measures associated with system security.

This paper surveys concepts and methodologies for the evaluation of system dependability [7], [8], [9] and summarizes how these are now being extended to evaluate system security. These techniques differ from most other treatments of system security in that they frequently use stochastic modeling. Stochastic assumptions are needed to describe systems that have yet to be built and for systems whose specific vulnerabilities remain unknown. In such cases, it is appropriate to make stochastic assumptions about the introduction and discovery of vulnerabilities, about attacker behavior, about system behavior (in terms of the effects the exploited vulnerabilities have on it, and in terms of the system's responses to attacks), and about transient periods of vulnerability, and to solve (or simulate) the model for stochastic measures. Stochastic models are also very useful for sensitivity analysis.

In this survey, we find that there is much in classical dependability analysis that can be transferred to security analysis. However, we also find that there are attributes of security that cannot be integrated naturally into a dependability framework. In addition, we point out that the root causes of system failure in the context of classical dependability are fundamentally different from the root causes of security violations in ways that impact the usefulness of the models we develop to describe those failures. We conclude the paper with ideas for future work in the area.

- *D.M. Nicol and W.H. Sanders are with the Coordinated Science Lab and Department of Electrical and Computer Engineering, University of Illinois, Urbana, IL, 61801. E-mail: nicol@crhc.uiuc.edu and whs@uiuc.edu.*
- *K.S. Trivedi is with the Electrical and Computer Engineering Department, Duke University, Durham, NC 27708. E-mail: kst@ee.duke.edu.*

*Manuscript received 21 June 2004; revised 24 Aug. 2004; accepted 25 Aug. 2004.*
*For information on obtaining reprints of this article, please send e-mail to: tdsc@computer.org, and reference IEEECS Log Number TDSC-0094-0604.*

## 2 Measures of Dependability and Security

Engineers have long used models to evaluate system designs. The models employed typically focus on the questions that are most pressing to an engineer (e.g., Will the bridge collapse? Are structural reinforcements needed to meet stress tolerances? Will the computer system provide a specified response time? Is the flight control system able to tolerate three simultaneous equipment failures?). Today, we are faced with pressing questions about computer system and network *security*, which is defined informally as the resilience of a computer system or network to malicious attacks. Our thesis is that modeling can help one design and evaluate systems or networks that are intended to be secure, just as modeling has helped us provide and evaluate other system properties. Furthermore, we see interesting similarities between computer system failures due to intentional attacks and system failures due to accidental component failures; we thus see value in exploring how evaluation techniques developed to quantify system dependability might be extended to quantify system security.

However, we also see that security brings new issues to be considered, most notably related to causes of component and system failure. Dependability analysis to date usually assumes that failures are caused by random events in hardware or rare events in software, and that this randomness can be quantified (even if correlated) in a way that permits the determination of system-level properties. Security analysis must assume that failures are caused by human intent, resulting in security failures that are definitely correlated, that depend in subtle ways on system state, and that attackers learn over time. While such events might appear to be random, as perceived by an outside observer, they tend to depend on each other in subtle ways that make them difficult to represent accurately using classical stochastic models. To highlight these differences, we review classical dependability measures, pointing out challenges that arise when they are applied to systems that fail due to malicious attacks.

*Reliability* is the probability that a system performs a specified service throughout a specified interval of time. Reliability analysis therefore depends on stochastic models of the frequency, duration, and intensity of faults in hardware and software. While one can certainly *assume* some probabilistic structure when modeling cyber attacks, the problem of developing and validating good stochastic models is very much an open issue. However, it is one we must solve if we are to use classical reliability analysis to predict reliability in the face of security breaches. Having said that, it may be possible to use a conjectured probabilistic specification of the occurrence of cyber attacks to support sensitivity analysis of a system's reliability or availability. Indeed, several initial attempts have been made to quantify system security using ideas developed to quantify the effect of accidental failures (see, for example, [10], [11], [12], [13], [14]).

*Availability* is a quantification of the alternation between proper and improper service, and is often expressed as the fraction of time that a system can be used for its intended purpose during a specified interval of time or in steady state. Challenges similar to those described above in the context of reliability analysis apply when evaluating a system's availability under malicious attack. Furthermore, as we consider the impact of security on availability, we see that a system's availability may be affected in several ways by a cyber attack; for example, it may be affected by the attack's own impact on the system and by the efforts to diagnose the attack and restore system service following the attack. Therefore, availability analysis of a system under malicious attack needs to specify explicitly how (and how long) a system remains unavailable following a successful attack.

*Safety* is the probability that a system does not fail in a manner that causes catastrophic damage during a specified period of time. Since system safety depends on the effect of a system failure rather than on the cause of the failure, one can easily imagine quantifying system safety in the context of cyber attack. While the safety of data from accidental erasure has certainly been a consideration in safety analysis of information systems, when we add security considerations, we will also need to consider the security of sensitive data in the event of a security breach as a component of safety. For example, the exposure of very sensitive data (e.g., private keys) might enable an attacker to cause catastrophic damage.

*Performability* [15] quantifies system performance in the presence of failures (either component or system). Performability analysis is often carried out by specifying a set of *structural* states for a system, each state corresponding to a configuration that results in a particular system performance, and specifying how the system changes state (often in the form of transition rates). Once these states and transitions have been identified, one then quantifies the amount of performance obtained in each state by specifying the rate at which *reward* is obtained in each state, and the amount (*impulse*) of reward that is obtained when a particular state transition is taken.

Using that mathematical structure, it is possible to specify performability measures in terms of the amount of reward accumulated during a specified interval of time (where the start or length of the interval can tend to infinity), or the rate of accumulation of reward at a specified instant of time or in steady state. (More details can be found in the next section.) Note that this method of specifying reward is general enough to support a wide variety of dependability and performability measures [16], [8], [17]. For example, if the rate at which reward is earned in an operational state is 1, while the reward rate earned in a nonoperational state is 0, the expected reward accumulated over a specified interval of time is just the expected amount of time the system was available over that epoch; in other words, it is the system's expected interval availability. If, instead, we define the rate reward for each state as the rate at which work is accomplished in that state and define state in the model in a manner that captures changes in the rate at which work can be done that are due to component failures, then the reward accumulated during an interval specifies a measure of the quality of service provided to a user.

The ease with which measures can be specified in the framework described above is quite useful for the analysis of security breaches. For example, consider a denial-of-

service cyber attack. The impact of that type of cyber attack and the system's attempts to cope with it can be reflected in the reward accumulated while spending time in states that reflect the attack. Another application of performability concepts in the security context follows from the fact that security measures may make a system harder to use. For example, security measures can slow the inherent rate at which work may be accomplished. Performability measures and analysis techniques thus provide a framework for considering the impact cyber attacks have on overall system performance. However, performability analysis faces the same principal challenge faced by reliability and availability analyses, namely, the development of meaningful stochastic descriptions of events that occur during a cyber attack.

System security includes attributes in addition to those we have described above in the context of dependability. In particular, *data confidentiality* means that a system does not allow protected data to be read in an unauthorized fashion, while *data integrity* means that a system does not allow protected data to be modified in an unauthorized fashion. A breakdown in confidentiality or integrity need not imply a failure of reliability or availability, at least not in their usual senses. It may qualify as a safety issue and could, depending on the structure of the associated measure, be expressed as a performability measure.

However, confidentiality and integrity are arguably different measures from those normally considered by dependability analysis, insofar as they are not concerned with system behavior so much as certain system properties. Likewise, *nonrepudiation* is a system property that prevents future false denial of involvement by either party in a transaction. *Authentication* is a related property, by which the claimed identity of a party to a transaction can be independently verified. There are no obvious connections of authentication or nonrepudiation to classical dependability measures.

Another important property in the security domain is survivability. As defined in [18], *survivability* is the capability of a system to fulfill its mission, in a timely manner, in the presence of attacks, failures, or accidents. Since survivability measures quantify the ability of a system to perform an intended function, modeling approaches that are applicable to availability and performability evaluation can be adapted for survivability evaluation. Recent work in quantifying survivability is found in [19], [20], [21], among other places.

Models for security analysis must describe how and when security breaches occur; they must describe the impact on the system when they do, as well as the mechanisms, effects, and costs of system recovery, system maintenance, and defenses. Stated as such, those are identical to the requirements of dependability models. However, there need to be significant differences in the nature and details of security models. This is most pointedly seen when we consider the introduction of security failures. A leading source of security vulnerability is misconfiguration. Failures due to misconfiguration can of course happen in other contexts, but a distinguishing feature in the security context is that some external agent must deliberately exercise the vulnerability in order for the failure to occur.

Latent software faults (e.g., buffer overflow problems) are another cause of security failure. Any given fault has specific idiosyncratic behaviors and requirements for accessing and exploiting it. Like misconfiguration, a security penetration made possible by a latent software fault does not occur accidently, but is actively induced by an attacker. Furthermore, a security penetration may require an attacker to exercise several vulnerabilities before compromising a prized asset (such as root access). This coupling of system vulnerabilities and attackers' exploitation of them distinguishes security failures from the types of failures traditionally considered by dependability analysis. The key issue is that of how to characterize attacker behavior.

## 3 Model Representation/Analysis Techniques

Research in dependability analysis has led to a variety of models, each focusing on particular levels of abstraction and/or system characteristics. As we extend that type of analysis into the security domain, we again find utility in diverse model types. We now review important classes of model representation and report on how they are being extended.

### 3.1 Combinatorial Methods

In contrast with state-space models, combinatorial models do not enumerate all possible system states to obtain a solution. Instead, simpler approaches are used to compute system dependability measures. Despite several extensions that have been made to combinatorial models, they do not easily capture certain features, such as stochastic dependence and imperfect fault coverage. We present a brief overview of combinatorial models.

#### 3.1.1 Reliability Block Diagrams (RBD)

An RBD is a graphical structure with two types of nodes: blocks representing system components and dummy nodes for connections between the components. Edges and dummy nodes model the operational dependency of a system on its components. At any instant of time, if there exists a path in the system from the start dummy node to the end dummy node, then the system is considered operational; otherwise, the system is considered failed. A failed component blocks all the paths on which it appears. RBDs thus map the operational dependency of a system on its components and not the actual physical structure of the system.

Series-Parallel RBDs are useful not only because they are very intuitive, but also because they can be solved in linear time [22]. Such RBDs are quite frequently used in reliability and availability modeling [8], [9], and many software packages exist that support construction and solution of RBD models (e.g., [23], [22]). We have yet to see an application of RBDs in security modeling. For such an application to be possible, one would need to create a compositional theory of security. At a glance, it seems that such a theory ought to have different semantics; in

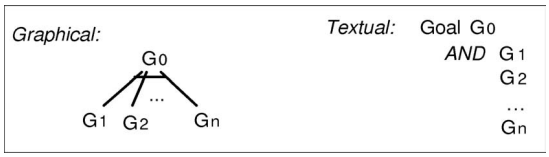


Fig. 1. AND node.

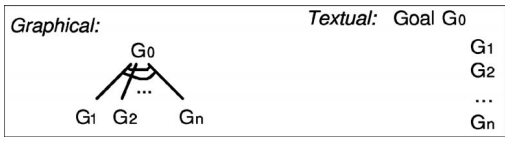


Fig. 2. OR node.

particular, an architecture needs to isolate an insecure component, not just provide a replicate in parallel.

### 3.1.2 Fault Trees (FTs)

A fault tree is an acyclic graph with internal nodes that are logic gates (e.g., AND, OR, k-of-n) and external nodes (leaves or basic events) that represent system components. The edges represent the flow of failure information in terms of Boolean entities (TRUE and FALSE or 0s and 1s). Typically, if a component has failed, a TRUE is transmitted; otherwise, a FALSE is transmitted. The edge connections determine the operational dependency of the system on the components. At any instant of time, the logic value at the root node determines whether or not the system is operational. If shared (repeated) nodes (nodes that share a common input) are not allowed, then the acyclic structure is a rooted tree.

Fault trees without shared nodes are equivalent to series-parallel RBDs [9], but when shared nodes (or repeated events) are allowed, fault trees are more powerful [24]. Many solution algorithms exist for fault trees with repeated events, including those based on sums of disjoint products (see [25]) and binary decision diagrams (e.g., [26], [27]). Fault trees have been extensively used in reliability and availability modeling (e.g., [28], [29], [30], [31]), safety modeling (see [32]), and modeling of software fault tolerance (e.g., [33]).

Fault trees have been extended to include various types of gates, such as priority AND gates, sequence dependency gates, exclusive OR gates, and inhibitor gates [34]. There have been extensions to include imperfect coverage [35], multistate systems [36], and phased mission systems [37]. A large number of software packages that support the construction and solution of fault trees are available (e.g., [38], [39], [22]). The main difficulty of using combinatorial methods in practice is the common assumption that all basic events must be statistically independent.

### 3.1.3 Attack Trees

Attack trees are closely related to fault trees, in that they consider a security breach as a system failure, and describe sets of events that can lead to system failure in a combinatorial way. An attack tree thus models all possible attacks against a system, just as a fault tree models all failures. They provide a formal, methodical way to describe the security of systems and subsystems based on various types of attacks (originally described in [40]) using graphics that are somewhat different from those that have become standard for fault trees). In an attack tree, the attacks to a system are represented in a tree structure, with the goal as the root node and the different ways to achieve that goal as leaf nodes. The security of a large system can be modeled with a set of attack trees, where the root of each tree represents an attack that can significantly damage the system's operation.

**Structures and Semantics**. In an attack tree, each nonleaf node represents an attack goal (or subgoal), and leaf nodes are atomic attacks. There are two kinds of nonleaf nodes: AND nodes and OR nodes. An AND node represents an attack goal for which a set of subgoals must be achieved in order for the attack to succeed. These attack subgoals are represented by the AND node's children. An OR node represents an attack goal that can be achieved in several ways, which are represented by the OR node's children.

Attack trees can be represented graphically or textually. A representation of an AND node is shown in Fig. 1. The figure shows a goal $G_0$ that can be achieved if the attacker achieves each of $G_1$ through $G_n$.

A representation of an OR node is shown in Fig. 2. It shows a goal $G_0$ that can be achieved if the attacker achieves any one of $G_1$ through $G_n$.

**Assigning Node Values**. Once the attack tree has been created, different values can be assigned to the leaf nodes. These values can be:

1. *Boolean (e.g., possible versus impossible):* A possible node represents a feasible attack scenario; an impossible value means the attack cannot be carried out in the current situation. Some other Boolean values include easy versus not easy, expensive versus not expensive, intrusive versus nonintrusive, legal versus illegal, and special equipment required versus no special equipment required.
2. *Continuous:* Sometimes it is not enough to use Boolean values to describe the attacks. For example, we may want to know the probability that an attack goal can be achieved. We can do so by assigning continuous values to each leaf node. These values may include the cost in dollars to attack/defend, the effort spent to achieve/repulse, the probability that the attack will succeed/fail, and the likelihood that the attacker will try the attack.

**Evaluating Attack Trees**. After assigning values to each leaf node, it is possible to propagate the node value up to the root of the tree. A node's value is a function of its children's values. Depending on whether a node is an AND node or an OR node and the nature of the assigned values, the calculation rules may differ. For example, if the possible/impossible values are under consideration, the AND node's value is the Boolean *and* of all values of its children, while the OR node's value is the Boolean *or* of all values of its children. When cost value is considered, the value of the AND node is the sum of the values of its children, and the value of the OR node is the minimum of the values of its children.

The attack tree can be used to evaluate different aspects of the system security, depending on the kind of value that is assigned to the leaf nodes. If a possible/impossible value is assigned, one can enumerate all sets of possible atomic attacks that achieve the attack goal; if a probability value is assigned, one can use an attack tree to evaluate the probability that the attack goal can be achieved. If a cost value is assigned, an attack tree can be used to evaluate the minimum cost needed to reach an attack goal.

Since atomic attacks can have multiple attributes, each leaf node can have several different value types. Therefore, an attack tree can be used to combine these values and help users learn more about a system's vulnerabilities. For example, if one assigns a possible/impossible value as well as a cost-in-dollars value to each node, one can use the attack tree to find the lowest-possible-cost attack sets for the system; if a probability value as well as a special equipment value is assigned, one can obtain the most-probable attack sets with no special equipment required.

Although fault trees are used primarily in dependability analysis and attack trees are used primarily in the security context, they share much in common. They have the same tree structures; they both contain AND nodes and OR nodes, which are the two most commonly used node types in both trees; and they have similar calculation rules to propagate node values up to the root. Therefore, many techniques used in fault tree analysis can be applied to analyze attack trees. For example, the SDP, BDD, or factoring methods used to compute system reliability can also be used in attack trees to obtain the probability that the attack goal will be reached; the minimum cut-set and minimum path sets analysis from fault trees can be used to find all sets of atomic attacks that achieve the goal; and similar concepts and computations of importance measures in fault tree analysis can be applied to attack trees to evaluate the impact of certain atomic attacks on the overall system security.

Attack trees thus provide a systematic way to describe the security vulnerabilities, thus making it possible to assess risks and make security decisions. They capture knowledge and expertise in a reusable form; once the attack tree for a certain security feature has been built, it can be included as part of a larger attack tree for a system that uses the security feature.

### 3.2 Model Checking

Another important type of dependability and security analysis is based on a reachability analysis of the model state space, an activity sometimes known as *model checking*. The general approach has a long history in hardware verification. The idea is to analyze the state space implied by some formal expression of the system. Certain states reflect deleterious conditions; model-checking algorithms explore the entire state space, report states of interest as they are uncovered, and for each one give an example sequence of state transitions that reaches it.

An early paper [41] described how to model the behavior of certain types of public key protocols in terms of the action of the protocol plus participant and intruder knowledge, specify precisely the meaning of "security fault," and search the implicit state-space for such faults. Limitations on protocols that could be so analyzed have been relaxed (e.g., [42], [43], [44]).

Model-checking is also being used to analyze computer programs for security flaws [45]. The fundamental data structure is the program's control flow graph, and the fundamental concept being analyzed is a set of program properties of an execution path, as it evolves in accordance with the control flow graph. Chen et al. [46] report success in analyzing large-sized well-known software packages (Apache HTTPD, BIND, Postfix, OpenSSH, Samba, Sendmail) for a set of particular security flaws. Actions a program execution might take to exercise a flaw are described with finite state automata; the model-checking involves analyzing the control-flow graph to determine whether the program satisfies the transitions described by any fault automata.

Model-checking approaches are also finding application in modeling attacks on systems (e.g., [47], [12]), where the network state includes a description of hosts and their vulnerabilities and a description of connectivity. In that approach, an attacker's state includes capabilities and access gained so far in the course of an attack. A state transition occurs when there is a match between a capability in the attacker's state and a vulnerability in the network state, resulting (usually) in increased access somewhere in the network. States that represent an attacker's access to network assets (e.g., gaining root access on a host, or an ability to retrieve critical information from a database) reflect successful exploits. For every asset, one can ask whether it can be compromised (a precise definition of which depends on the model) and, in principle, determine the number of paths to states in which the asset is first compromised.

The difficulty, of course, is the size of the state space. Advances in state representation have led to an ability to represent very large state spaces; for example, [48] indicates that extremely large state-transition-rate diagrams can be represented compactly using symbolic data structures such as Binary Decision Diagrams (BDDs) and Multiterminal Decision Diagrams (MDDs) (e.g., [26], [49]). These techniques have been shown to be able to represent state-transition-rate diagrams with $10^{20}$ or more states. Sheyner et al. [12] report a model build-time of two hours on a model that has 229 bits of state. From the point of view of tractable formal models, 229 bits is remarkable; however, this serves only to blunt the onset of the curse of dimensionality. With respect to the need to represent the complexities of real systems, it is still very small. A nascent effort at finessing the state space issue by *sampling* paths through the state space is reported in [50], in which a model with 1,700 bits of state is analyzed. The fundamental idea is to use importance sampling to guide the path sampling strategy, in order to estimate inherent system security metrics such as the number of unique exploits that compromise a given asset.

### 3.3 State-Based Stochastic Methods

Combinatorial methods are quite limited in the stochastic behavior that they can express. While attack-tree analysis has become a staple in the diet of system security analysts, the classical formulation does not capture the dependence

of security vulnerabilities on *sequencing* of events; to be successful, a buffer overflow attack that gains root access must precede the attack `rm -r *`.

State-space methods are much more comprehensive. They allow explicit modeling of complex relationships (e.g., [51]), and their transition structure encodes important sequencing information. Historically, state-space methods have been explored in the context of mathematical models that specify probabilistic assumptions about time durations and transition behavior. We now review those models and comment on how they are being applied in the security context.

#### 3.3.1 Markov Reward Models [17], [8], [52]

Let $\{X(t), t \geq 0\}$ be a homogeneous finite state continuous time Markov chain (CTMC) with state space $S$ and infinitesimal generator matrix $\mathbf{Q} = [q_{ij}]$. Let $P_i(t) = P\{X(t) = i\}$ denote the unconditional probability that the CTMC will be in state $i$ at time $t$, and the row vector $\mathbf{P}(t) = [P_1(t), P_2(t), \ldots, P_n(t)]$ represent the transient state probability vector of the CTMC. The transient behavior of the CTMC can be described by the Kolmogorov differential equation:

$$\frac{d\mathbf{P}(t)}{dt} = \mathbf{P}(t)\mathbf{Q}, \quad \text{given } \mathbf{P}(0), \tag{1}$$

where $\mathbf{P}(0)$ represents the initial probability vector (at time $t = 0$). The steady-state probability vector $\pi = \lim_{t\to\infty} \mathbf{P}(t)$ satisfies:

$$\pi\mathbf{Q} = 0, \quad \sum_{i\in S} \pi_i = 1. \tag{2}$$

In addition to transient state probabilities, cumulative probabilities are sometimes of interest. Define $\mathbf{L}(t) = \int_0^t \mathbf{P}(u)du$; then, $L_i(t)$ denotes the expected total time the CTMC spends in state $i$ during the interval $[0, t)$. $\mathbf{L}(t)$ satisfies the differential equation:

$$\frac{d\mathbf{L}(t)}{dt} = \mathbf{L}(t)\mathbf{Q} + \mathbf{P}(0), \quad \mathbf{L}(0) = 0. \tag{3}$$

With these definitions, most interesting dependability measures can be defined.

CTMCs with absorbing states deserve additional attention. Here, the measures of interest are based on the time a CTMC spends in nonabsorbing states before an absorbing state is ultimately reached. To compute this measure, the state space $S = A \cup T$ is partitioned into the set $A$ of absorbing states and the set $T$ of nonabsorbing (transient) states. Let $\mathbf{Q}_T$ be the submatrix of $\mathbf{Q}$ corresponding to the transitions between transient states. Then, the time spent in transient states before absorption can be calculated by $\mathbf{L}_T(\infty) = \lim_{t\to\infty} \mathbf{L}_T(t)$ restricted to the states of the set $T$. The mean time to absorption (MTTA) can be written as $MTTA = \sum_{i\in T} L_i(\infty)$.

Assigning rewards to states or to transitions between states of a CTMC defines a Markov reward model (MRM). Rewards are referred to as *rate rewards* in the former case, and as *impulse rewards* in the latter case. If we consider rate rewards only, let the reward rate $r_i$ be assigned to state $i$. Then, the random variable $Z(t) = r_{X(t)}$ refers to the instantaneous reward rate of the MRM at time $t$. The accumulated reward over the interval $[0, t)$ is given by

$$Y(t) = \int_0^t Z(u)du = \int_0^t r_{X(u)}du. \tag{4}$$

Based on the definitions of $X(t)$, $Z(t)$, and $Y(t)$, which are nonindependent random variables, various measures can be defined. The most general is the distribution of the accumulated reward over time $[0, t)$, that is, $P\{Y(t) \leq y\}$, which is difficult to compute for unrestricted models and reward structures (see [52], [53] for a survey of methods to compute the distribution of reward accumulated over a finite interval).

The problem is considerably simplified if we restrict ourselves to the expectations and other moments of random variables. In that case, the expected instantaneous reward rate can be computed from

$$E[Z(t)] = \sum_{i\in S} r_i P_i(t) \tag{5}$$

and the expected reward rate in steady state (when the underlying CTMC is ergodic) is

$$E[Z] = \sum_{i\in S} r_i \pi_i. \tag{6}$$

To compute the expected accumulated reward over $(0, t)$, we use

$$E[Y(t)] = \sum_{i\in S} r_i L_i(t). \tag{7}$$

For models with absorbing states, the limit as $t \to \infty$ of the expected accumulated reward is called the *expected accumulated reward until absorption*, which is

$$E[Y(\infty)] = \sum_{i\in T} r_i L_i(\infty). \tag{8}$$

Note that the reward rate assignments (or reward structure) clearly depend on which attribute we are interested in with respect to a system's dependability and security. Markov and Markov reward models have been extensively used for dependability analysis of hardware systems (see, for example, [8], [9], [22]), real-time system performance in the presence of failures [54], [55], architecture-based analysis software systems (e.g., [56]), combined analysis of hardware-software reliability (e.g., [57], [58]), system performance analysis (e.g., [17]), and performability analysis (e.g., [8], [52], [59], [60], [61], [53]). Many sources for solution algorithms are available (e.g., [62], [63]), and many software packages exist (e.g., [62], [22], [64], [65], [66]). Several models constructed from and validated against measurement data have also been published [67], [68].

For complex systems with large numbers of components, the number of system states can grow prohibitively large. This is called the *largeness* problem for MRM models. Thus, significant work is being done to reduce the size of the Markov chain required for realistic system models. There are two general approaches for dealing with the size problem: *largeness avoidance* and *largeness tolerance* [22].

### 3.3.2 Largeness Avoidance Techniques

When largeness avoidance is employed, the size of a model is reduced and, therefore, a large model is not generated. State truncation methods [69], [70], hierarchical model solution [22], fixed point iteration [71], [72], and hybrid models that judiciously combine different model types [22] are examples of largeness avoidance.

State lumping (e.g., [73], [74]) has also been used extensively as an avoidance approach. Lumping reduces the size of a CTMC by considering the quotient of the CTMC with respect to an equivalence relation (i.e., replaces a set of states with a single lumped state) that preserves the Markov property and supports the desired performance measures defined on the CTMC. By solving the smaller CTMC, it is possible to compute exact results for the larger CTMC and, therefore, measures of interest for the original model.

A *state-level* lumping technique is a lumping technique that exploits the lumping properties at the CTMC level. The main advantage of state-level lumping techniques is that they generate the optimal (i.e., smallest possible) lumped CTMC. However, since they can perform efficiently only on a sparse matrix representation of a CTMC, they have prohibitive space requirements for very large CTMCs; therefore, they are usually used along with other CTMC solution techniques. Buchholz [75] gives a state-level lumping algorithm with $O(mn)$ time complexity and $O(m + n)$ space complexity for computing the optimal (i.e., coarsest) lumping of a CTMC represented as a sparse matrix, where $n$ is the number of states and $m$ is the number of transitions of the CTMC.

Several authors have also addressed the problem of computing bisimilarity [76], which is, in some ways, similar to the problem of state-level CTMC lumping. Kanellakis and Smolka gave a partition refinement algorithm with time complexity $O(mn)$ [77]. They conjectured that an algorithm exists that reduces the time complexity to $O(m \log n)$. A few years later, Paige and Tarjan designed such an algorithm [78]. An implementation of Paige and Tarjan's algorithm can be found in [79]. $O(m \log n)$ complexity has been claimed without formal proof by Bernardo and Gorrieri [80] for CTMCs and by Huynh and Tian [81] for discrete-time Markov chains (DTMCs). Derisavi et al. [82] recently provided an $O(m \log n)$ lumping algorithm, along with a rigorous proof. The approach they take is based on Tarjan and Paige, and uses new data structures to obtain the bound.

In contrast, *model-level* lumping techniques identify appropriate lumping properties by operating on a higher-level formalism (see the following section for a description of one such formalism) and directly constructing a lumped CTMC, rather than by constructing the unlumped CTMC and then operating on it. The lumping equivalence relation is established by the modeling formalism itself in some model-level lumping techniques. That holds for stochastic well-formed nets (SWNs) [83] and replicate/join operators in stochastic activity network-based composed models (SANs) [84], in which the lumping results from equivalence of the replicas of a particular submodel. Extending the work of [84], Obal developed [85] a graph composition formalism and used symmetry detection, a type of model-level lumping technique. The technique automatically identifies and exploits all the structural symmetries due to the interaction between submodels of a state-sharing composed model, that is, a model consisting of submodels that share a subset of their state variables. Restricted versions of a symmetry detection technique similar to the one described in [84] have also been used for process algebras in [86], [87]. Unlike state-level lumping, model-level lumping techniques do not always find the optimal lumping because they do not operate at the CTMC level.

Other lumping techniques, which we call *compositional lumping* techniques, can be applied to composed models provided that the specific high-level formalism satisfies a particular set of assumptions. In these techniques, each of the individual interacting submodels is lumped separately from the others using a state-level lumping algorithm, and is then replaced in the overall model by its lumped version. For example, based on the fact that lumping is a congruence with respect to parallel composition in a number of process algebra formalisms and stochastic automata networks, compositional lumping can be used in those formalisms to generate lumped state spaces (e.g., [88], [89], [90], [91], [92], [93]). Most of the work on compositional lumping applies only to state-level lumping inside the submodels. In some cases, in addition to lumpability in each of the submodels, the structural symmetry of the interaction among the submodels may also be exploited to achieve even smaller CTMCs. In other words, for some composed models, a compositional lumping algorithm that applies the state-level and model-level lumping techniques at the same time could give an extra opportunity to shrink the CTMC. Therefore, a fairly general algorithm that integrates the two techniques for a compositional formalism is desirable.

Another largeness avoidance technique is called *aggregation*.[1] In this approach, as in lumping, a set of conditions for partitioning the set of states of a CTMC is given such that a smaller CTMC is constructed by replacing the set of states in each block of the partition with a single state. The aggregation differs from lumping in that the solution of the aggregated CTMC gives approximate results (with or without bounds) on the original CTMC, as opposed to the exact results that would be obtained from the lumped CTMC. Moreover, some aggregation techniques are only applicable to CTMCs that satisfy a strict set of conditions. However, aggregation conditions could result in a coarser partition and, therefore, a smaller CTMC, compared to the lumping conditions. An aggregation technique for steady-state analysis of a general CTMC has been proposed in [94], [95]. It gives the best known bounds on the result but is computationally costly, and the bounds are tight only if the matrix satisfies some strict constraints. Bobbio and Trivedi [96], [97] extended Courtois's technique to the domain of transient analysis of CTMCs. Daly et al. [98] give a more general aggregation technique that can solve for both steady-state and transient measures of general CTMCs. It introduces a new partial order on the set of states of a CTMC that is a generalization of the concept of lumping.

1. Note that different authors use the terms *lumping* and *aggregation* differently.

### 3.3.3 Largeness Tolerance Techniques

Even a lumped CTMC can be extremely large and, hence, *largeness tolerance* techniques are needed to provide practical modeling support for large CTMCs. In those techniques, one starts with a concise high-level representation of the system being modeled (usually a variant of stochastic Petri net or stochastic process algebra; see a following section). Then, new algorithms are designed to manipulate the large underlying CTMCs, and special data structures and/or representations are utilized to reduce the space requirements of the state space, the generator matrix, and the iteration vectors. Those techniques are usually, but not always, associated with compositional modeling.

Binary [26] and multivalued decision diagram [99] (BDD and MDD) data structures have been successfully applied to efficiently explore and represent large unlumped state spaces. The key idea is to encode states as paths in a directed acyclic graph. Techniques that generate state spaces using decision diagrams are referred to as *symbolic* state-space exploration and representation techniques (e.g., [49], [100]).

MDD data structures have been used in [48] to explore large state spaces of models built using an *action synchronization* (also known as *action-sharing*) high-level compositional formalism in which submodels interact by synchronized firing of a subset of their actions. Saturation, a state-space exploration technique that was introduced in [101], improved the running time of the algorithm given in [48] by up to a few orders of magnitude, thus enabling the exploration of even larger state spaces. In [48], [101], it was assumed that the state spaces of individual submodels were known a priori, i.e., the state spaces of the submodels computed by exploring the submodels in isolation are finite and result in the same state spaces they would have if they had been explored in interaction with the rest of the model. This assumption was relaxed later in [102]. Those state-space exploration techniques are applicable to action-synchronization composed models that conform to a set of particular structural restrictions called the *logical product form* property [48].

One approach in space-efficient representation of generator matrices is to follow a divide-and-conquer strategy and represent the matrix with a set of relatively small component matrices that are appropriately combined. The earliest attempt using that approach was made by Plateau [103] and Plateau and Atif [104], who proposed a technique in which the generator matrix of a CTMC generated from a specific compositional high-level formalism need not be explicitly stored. Instead, the matrix is implicitly represented as a mathematical expression consisting of *Kronecker* operators and a number of relatively small matrices derived from the structure of submodels. Later, the "Kronecker representation" technique was extended to more general formalisms, and a number of its shortcomings were resolved [105], [106], [107], [108], [109], [110]. The approach is applicable only to action-synchronization models that satisfy certain structural constraints.

A parallel effort was undertaken by Ciardo and Miner [111], who proposed the matrix diagram (MD) data structure to store the generator matrix of action-synchronization composed models. An MD is structurally similar to an MDD and, along with an MDD, represents the set of states and transitions of a very large CTMC. Efficient algorithms to manipulate MDDs and MDs have been given in [111]. In [111], the algorithm that generates the MD data structure is time-efficient, but works only for composed models that hold the logical product form property. Later, Miner [112] developed canonical MDs (CMDs), a proper subset of MDs, and presented an algorithm to store virtually any matrix in the form of a CMD. In particular, he used the algorithm to generate the CMD representation of the generator matrix of models based on the generalized stochastic Petri net (GSPN) formalism, which is a fairly general Markov modeling formalism. Since the algorithm added matrix elements to the CMD data structure one by one, and without exploiting any structural information, its running time was prohibitive. As mentioned above, both CMDs and MDs can represent virtually any generator matrix regardless of the modeling formalism from which it was generated. The challenge is then to develop algorithms that build (C)MD representations of generator matrices of other formalisms in a time-efficient manner.

In contrast, the *disk-based* approach first introduced in [113] performs steady-state solution by storing the generator matrix of the CTMC in the disk instead of the memory while using a variant of block Gauss-Seidel as the iterative solution algorithm. To increase the utilization of the CPU, the algorithm implementation concurrently fetches parts of the matrix from the disk and performs computation on other in-memory parts of the matrix. By using disk instead of memory to store the matrix, the technique enables the solution of CTMCs that are one or two orders of magnitude larger than what would be possible if using only memory.

Likewise, the "on-the-fly" technique of [114] completely avoids the storage of the generator matrix by (re)generating the elements of the matrix as they are needed in an iterative solution algorithm (steady state or transient solution). The elements are computed on-the-fly from the model, which is given in a high-level formalism. Repetitive calculations of the elements incur a substantial computational overhead.

The path-based approach is yet another largeness tolerance technique for performing transient analysis of CTMCs while avoiding the storage of the CTMC and possibly the iteration vector. In this approach, a limited number of paths (i.e., sequences of transitions) of the CTMC that make a major contribution toward the measures of interest are enumerated. Then, the reward is computed only for those paths. The first notable work based on that approach was given by de Souza e Silva and Gail [115], and they later improved it in [116]. Later, Qureshi and Sanders improved the numerical stability and computational complexity of [116] in [117]. Most recently, Lam et al. [118] use Kronecker operators to represent both the CTMC and the iteration vector to compute approximate results for transient analysis of an action-synchronization composed model.

### 3.3.4 Other Challenges

When there is a large difference between failure and repair rates or failure and job arrival rates in the model, it leads to

the *stiffness* problem for MRM models. Stiffness can be reduced by separating the performance and availability models, but the stiffness within the availability model remains. Stiffness may be avoided by using aggregation [96], [97] that yields approximate solutions. To tolerate stiffness, special stable stiff solvers may be used (e.g., [119], [120], [121], [122]).

Largeness and stiffness problems may also be caused by combining the performance and availability models in a single monolithic performability model. Solving an overall model of system behavior can potentially yield more accurate results than solving two smaller, less stiff models that only lead to approximate solutions. However, we should note that numerical difficulties arising from largeness and stiffness may very well negate this gain [52], [60].

A major objection to the use of homogeneous Markov models in the evaluation of performance and dependability behavior of systems is the assumption that the sojourn (holding) time in any state is exponentially distributed.

The exponential distribution has many useful properties that lead to analytic tractability, but it does not always realistically represent the observed distribution functions. One way to deal with nonexponential distributions is the phase-type approximation, which consists of modeling a distribution with a set of states and transitions between those states such that the holding time in each state is exponentially distributed [8], [123]. The simplest examples of phase approximation are the hyperexponential distribution with a coefficient of variation larger than 1, and hypoexponential distribution with a coefficient of variation less than 1. Although the method of phase-type approximation enables us to use MRMs, its major drawback is that it usually results in a large state space.

If transition rates in a CTMC are allowed to be time-dependent, where time is measured from the beginning of system operation, the model becomes a nonhomogeneous CTMC. Such models are used in software reliability modeling (e.g., [124], [125]) and in hardware reliability models of nonrepairable systems (e.g., [126]).

Due to the assumptions that holding times in the state are exponentially distributed and that past behavior of the process is completely summarized by the current state of the process, any observation instant in a homogeneous CTMC acts as a regeneration point for the process. The first assumption can be relaxed by allowing the holding time to have any distribution, thus resulting in a *semi-Markov process* (SMP), where the epoch of each state transition is a regeneration point [127]. This assumption can be relaxed by assuming that not all state transitions are regeneration points, thereby resulting in a *Markov regenerative process* (MRGP) [128].

### 3.3.5 Higher-Level Model Representations

CTMCs are rarely used directly to specify a system's model in a typical modeling process. Many high-level modeling formalisms have been created to fill the gap between CTMC specification and system design specification. Examples of those formalisms include variants of stochastic Petri nets (e.g., [129], [64], [130], [131]), variants of stochastic process algebras [132], [90], [80], [133], and interactive Markov chains (IMC) [93]. To illustrate the usefulness of these model types, we describe stochastic Petri nets and extensions in more detail, and illustrate their use in the context of security.

Stochastic Petri nets (SPNs) and extensions have been developed as extensions to untimed Petri nets (originally introduced by C.A. Petri in 1962) with timed transitions for which the firing time distributions are assumed to be exponential. SPNs have been extensively used in the area of dependability evaluation (e.g., [8], [134]) due to the small size of their descriptions and their visual/conceptual clarity. They allow the designer to focus more on the system being modeled rather than on error-prone and tedious manual creation of a lower-level MRM.

To specify an SPN, one has to define a set of places $P$, a set of transitions $T$, and a set $A$ of arcs between transitions and places: $A \subseteq (P \times T) \cup (T \times P)$. Each place can contain zero or more tokens. Graphically, places are depicted as circles, transitions as bars, tokens as dots in circles, and arcs as arrows.

The distribution of tokens over the places is called a *marking* and corresponds to the notion of state in a Markov chain. All places from which arcs go to a particular transition are called the *input places* of that transition. All places to which arcs go from a particular transition are called the *output places* of the transition. A transition is said to be *enabled* when all of its input places contain at least one token. If a transition is enabled it may fire. Upon firing, a transition removes one token from each of its input places and puts one token in each of its output places, possibly causing a change of marking, i.e., a change of state.

The firing of transitions is assumed to take an exponentially distributed amount of time. Given the initial marking of an SPN, all the markings as well as the transition rates can be derived, under the condition that the number of tokens in every place is bounded. Thus, a Markov chain is obtained.

Classically, SPNs and extensions are solved via an underlying MRM that can be automatically derived, thereby making it possible to use the wide variety of available techniques for MRMs.

In the last two decades, many extensions to the basic SPN model have been proposed to enhance its modeling power and flexibility of use. They include arcs with multiplicity, a shorthand notation for multiple arcs between a place and a transition, immediate or instantaneous transitions that fire in zero time, and inhibitor arcs from places to transitions that prevent a transition from firing as long as there are tokens in the place. The most popular model of this type is called generalized stochastic Petri nets (GSPNs) [129]. More flexible firing rules have also been proposed, most notably the introduction of gates in stochastic activity networks (SANs) [130], [131] and guards or enabling functions in stochastic reward nets (SRNs) [64]. The extended stochastic Petri net (ESPN), in which general firing time distributions are allowed, has a semi-Markov process [135] as the underlying process, when certain restrictions are met.

Deterministic stochastic Petri nets (DSPNs) allow the definition of immediate, exponential, and deterministic transitions [136]. The stochastic process underlying a DSPN

is a Markov regenerative process. Markov regenerative stochastic Petri nets (MRSPNs) generalize DSPNs and still have MRGP as an underlying stochastic process [137]. Concurrent generalized Petri nets (CGPNs) allow simultaneous enabling of any number of immediate, exponentially timed, and generally distributed timed transitions, provided that all generally distributed transitions are enabled at the same instant. The stochastic process underlying a CGPN is also an MRGP [138]. Fluid stochastic Petri nets allow for continuous state variables [139], [140].

### 3.4 Applications to Security Modeling

As described earlier, state-based techniques have been extensively developed and used in classical dependability contexts. They are now also beginning to be used in security analysis. For example, Ortalo et al. [141] have proposed modeling of known system vulnerabilities using "privilege graphs," followed by a combination of the privilege graphs with simple assumptions about attacker behavior to obtain "attack-state graphs." The latter can be analyzed using the Markovian reward models described above to obtain probabilistic measures of security. An interesting definition of reward used here is the "effort" needed to make a transition (which usually represents some sort of system compromise).

The types of Markovian models described above ascribe distributional properties to high-level system and attacker behaviors. Accepting this, we have still the issue of quantifying the scale of these stochastically modeled activities, particularly those related to attacks. Early groundbreaking work in this regard was done by Littlewood et al. [142]. Their work was exploratory in nature and identified "effort" made by an attacker as an appropriate measure of the security of the system. Effectively, the model consisted of only two states, viz. "working" and "security failed state." The latter state was assumed to be an absorbing state. With respect to the above discussion, the relevant security measure turns out to be "mean effort to (security) failure." Jonsson and Olovsson [143] attempted to build a quantitative Markov model of attacker behavior using data from several experiments they conducted over a two-year period. They postulated that the process representing an attacker may be broken into multiple phases, each of which has an exponential time distribution. The overall attacker behavior therefore requires nonexponential characterization.

More recently, Singh et al. [13] have used stochastic activity networks to validate an intrusion-tolerant system, emphasizing the effects of intrusions on the system behavior and the ability of the intrusion-tolerant mechanisms to handle those effects, while using very simple assumptions about the discovery and exploitation of vulnerabilities by the attackers to achieve those intrusions. Gupta et al. [14] have used a similar approach to evaluate the security and performance of several intrusion-tolerant server architectures. Madan et al. [10] have used a semi-Markov model to evaluate the security properties of an intrusion-tolerant system. Depending on the particular attack scenario, various states may be associated with failure of availability, integrity, and confidentiality. The security may then be quantified in terms of the mean time to security failure and in terms of the absorption probabilities.

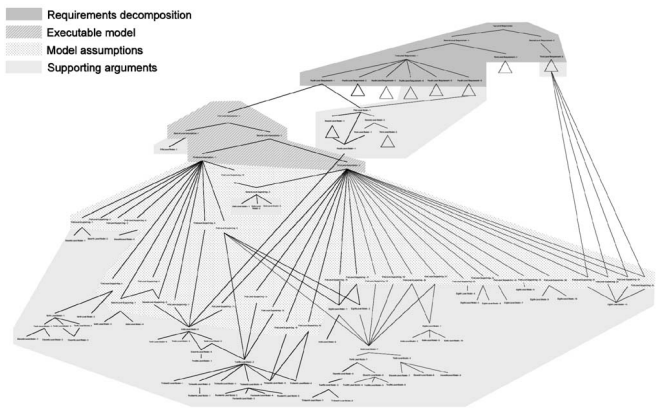


Fig. 3. DPASA argument graph structure.

Furthermore, Dacier et al. [144] have proposed a two-stage technique that starts by converting a privilege graph into an SPN by treating each atomic attack as a transition in a stochastic Petri net. In the second stage, the markings of this SPN generate a continuous-time Markov chain. Making simple assumptions about the attacker, the chain may be analyzed using the Markov reward techniques to obtain probabilistic measures of security in terms of the mean time to reach "security failed" states and other related measures. In [145], Wang et al. use stochastic reward nets to model both attacker and system behavior for an intrusion-tolerant architecture named SITAR [4]. Likewise, probabilistic methods have been used to model the DPASA [21] architecture. In the DPASA project, system validators combined structured requirement specification, probabilistic modeling, experimental evaluation, logical arguments, and formal methods to build an overall survivability argument for the architecture. Fig. 3 shows an example of an "argument graph" that links together arguments that make use of these methods.

A promising application of stochastic state-based methods in the security context is to quantify direct, high-level measures of the *service* that one can expect from a computer system or network in spite of cyber attacks that may occur. In order to do so, one needs a notion of time in the system model and a probabilistic notion of system and attacker behavior. High-level service measures can be defined to quantify system performance under cyber attack, so that 1) systems can be represented as state-level models in a way that captures either known or unknown vulnerabilities, 2) attacker behavior can be modeled in such frameworks, and 3) measurement can be used to quantify parameters describing known vulnerabilities.

A model for probabilistic validation of security with respect to high-level system properties (e.g., availability, privacy, and integrity) should have several components. It should contain representations of attackers, the system, and the assets, resources, and privileges associated with the system. It should represent attacker decision-making and all temporal aspects of his, and the system's, behavior, and the application (or applications) that provide services of interest [146].

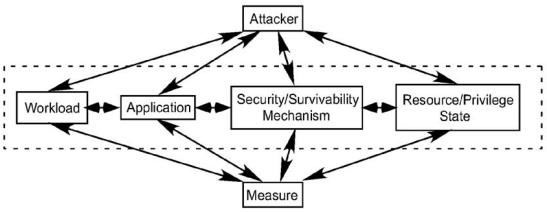


Fig. 4. Probabilistic security model structure.

Model specification of behaviors needs to focus on the appropriate stochastic measures and their relationship to services, and determine how they are reflected in the model. Such a model ought to capture relevant *attacker behavior*, the *workload* demanded of the system, the intended *application* (which defines the service that must be provided in spite of cyber attacks that may occur), a model of *security and survivability mechanisms* employed, and a model of the *resources/privileges* that are needed by the application to provide its services. Fig. 4 shows the relationship of these parts of the model to one another. The arcs connecting submodels in the figure represent possible interactions between submodels that can change their state. For example, the attacker may be able to change the state of a resource or amount of privilege granted to him (as represented by the directed arc from the attacker to the Resource/Privilege State submodel), or the attacker may change his state (and, hence, change his behavior) by using knowledge he has gained by observing the state of the system (represented by the directed arcs from the system submodels to the attacker). This general framework [146] is one we are using to develop models of specific applications on specific architectures.

Development of measurement techniques is just as important as the development of models to quantify security. Measurements should be developed for two purposes in this regard: 1) to guide construction of and provide input parameter values for models, and 2) to validate the correctness of models. In particular, with regard to the first objective, the appropriate level of detail/abstraction depends on the input parameter values (obtained from measurement data) available for each model. The type and accuracy of input parameter values available will depend on the stage of development of the system that is being validated. For existing systems, one could use methods similar to Jonsson's as a starting point to obtain values that quantify the behavior of an attacker. A recent study that placed significant emphasis on the development of an attacker effect model is [21]. The attacker model developed there is quite detailed, but tailored to the system being studied. It is not clear if it will be possible to build attacker models that are more generic, but still have the detail needed for meaningful representation.

Concerning the Resource/Privilege State model, one could use a security scanner like Nessus [147] combined with a network exploration and security auditing tool like nmap [148] to obtain model parameter values. Likewise, tools like COPS [149], Tiger [150], or Ferret [151] could be used to quantify host security vulnerabilities. The advantage of these tools is that the list of vulnerabilities is updated as new vulnerabilities are discovered. The application, security and survivability mechanism, and workload models are more case-specific, and require further study to determine appropriate input parameters and their values. It appears that the best route forward is from specific to general. We and others are currently building models of this type for specific system designs, and investigating how to obtain parameter values for them. It is our hope that these experiences will guide us in constructing more general models and approaches for parameter value estimation.

The second objective, model validation, is a more difficult issue, both for classical dependability models and for security models whose goal is quantification. We are not aware of any "silver bullet," other than the hard and meticulous work and significant time required to collect data on attackers and systems that are intended to be secure. Once that has been done, one can compare the insights gained from the data collected to those obtained from quantitative models. Note, however, that the value of most quantitative security models, like most classical dependability models, will be in gaining insight and making decisions about how best to make a system secure/dependable, rather than making a precise statement about a specific system's absolute dependability or security.

### 3.5 Simulation

As just argued, the state space of a system model may be too large to be analyzed in its entirety. Nevertheless, in principle, we can construct statistical estimators of all the system measures computed by the quantitative techniques we have described. Rather than generate and analyze the entire state space, we generate and analyze randomly chosen paths through the state space; typically, we sample many paths independently of each other (in a strict probabilistic sense of "independent"). This focus on evaluating a system by individual trajectories of the system is commonly known as *simulation*.

#### 3.5.1 Statistical Issues

A major issue when using simulation to estimate system measures is that of how to do so in a way that ensures the statistical quality of the estimates. The first concern is that the estimator be *unbiased*, which is a technical term that means that the mean value of the estimator (itself a random variable) is the same as the mean value of the random system measure being estimated. Proof that an estimator is unbiased is powerful insofar as it implies (by the law of large numbers) that the arithmetic average of many independent unbiased estimates converges to the mean of the random samples, which is the same as the mean of the system measure being estimated.

A more subtle statistical issue is the variance exhibited by the estimator. To illustrate this issue, imagine a system measure whose mean is a small probability, say $1e-5$. Further imagine that one way of estimating this probability creates estimates that fall almost entirely in the range $0.95e-5$ to $1.05e-5$, and another generates estimates that fall almost entirely in the range $1e-7$ to $1e-3$. The first method of estimation is clearly better because its values are almost always closer to the true mean than the values of the second method are. The practical impact of using a low

variance estimator is that, for a given level of accuracy, lower variance implies that fewer samples are needed to achieve that accuracy.

The issues of bias and variance are particularly significant when the measure of interest is a small probability and, therefore, are significant in the context of dependability and security. So-called *variance reduction techniques* [152] have been developed in reliability and performance analysis, but are characterized by the need to exploit a model's structure in order to be provably effective. Work is needed to explore how variance reduction technology can be applied to models focused on the idiosyncrasies of security issues.

The technique of simulating a system independently many times is necessary when one is interested in transient measures; the measures of interest can be repeatedly sampled at the given instants of interest. This same technique is used to estimate asymptotic measures. For example, consider the estimation of a queue's asymptotic average length. Using independent replications one runs $N$ statistically independent simulations, and generates $N$ independent samples of the average queue length. Each sample is the queue length averaged over an interval of simulation time $[s, T]$, where the system is deemed to be in "steady state" by time $s$, and $T$ is the simulation termination time. Standard statistical methods can be used to compute confidence intervals around a sample mean. However, replications per se are sometimes not needed to estimate asymptotic measures. One can instead push the model along a single sample path, but for a period of simulation time long enough to force the system into steady state, and then use samples taken across time, rather than samples taken across replications. Applied to the example above, one runs the simulation to time $T'$. The method of batch means [153] partitions the interval $[s, T']$ into successive epochs of duration $\Delta$, and measures average queue length over each. For many types of models, when $\Delta$ is large, the measurements are nearly independent, so confidence intervals can be constructed. Depending on 1) how long the simulation must be run until reaching steady state (i.e., time $s$), 2) the length of simulation time needed for a "good" sample ($T - s$), and 3) the size of the epoch needed to give near-statistical independence to successive estimates (i.e., $\Delta$), the computational cost of doing a long run can be smaller than the computational cost of $N$ independent replications. If we take the total length of simulation time traversed as an indicator of computational work, the cost of $N$ independent replications is $N \times (s + (T - s))$, while the cost of a long run is $s + N \times \Delta$. A little algebra shows that independent replications cost more when $(N - 1)s > N(\Delta - (T - s))$.

The point of selecting a large $\Delta$ is to try to heuristically create statistical independence between successive samples in a long run. It is possible to bring more rigor to this intuition, using the notion of regeneration points [154]. The idea is that certain systems sometimes enter states from which they probabilistically start over in the sense that the future behavior is independent of the past. The simplest example of such a regeneration point is the beginning of an idle period for an M/G/1 queue; the Poisson arrivals imply that there is no probabilistic memory of the time since the last arrival, and the fact that the queue is empty implies that there are no complications owing to interrupted service times. At such an instant, the future behavior of the queue is in a probabilistic sense completely independent of anything that has happened in the past. The implication is that epochs separated by regeneration points can, for the purposes of sampling measures, serve as independent replications of system behavior. Just as discovery of sampling schemes that lead to variance reduction depends very much on the particulars of the model, so likewise does identification of regeneration points. To make this technology work in the context of security analysis, we must identify or construct models in which regeneration points can be readily identified.

The challenges of using stochastic simulation models on system models in a security context partition along the lines of attributes of interest. Those attributes described by numbers, e.g., availability or performability, can be treated using known quantitative techniques; however, optimizations such as variance reduction require domain-specific insight that has yet to be generally developed in this context. New attributes more closely tied to security (e.g., confidentiality, nonrepudiation, and authentication) are fundamentally different in that they are system properties, not system measures. If classical stochastic simulation output analysis is to be used to help evaluate these attributes, then numeric measures of some sort will have to be developed to quantify these properties.

### 3.5.2 Simulation Model Representation

Model formalisms are often developed to expose underlying structure common to different models. The beauty of a stochastic Petri net formalism is that one may use it equally well to describe a communication synchronization protocol and a machine maintenance and repair schedule since the underlying mathematics and analysis algorithms remain the same. The flip side of the coin, however, is that such formalisms have built-in constraints whose effects limit how a model can be expressed. Some models just do not fit the formalisms. A common use of simulation is to express systems using formalisms that are not as mathematically tractable as Markov chains or SPNs, but which allow greater freedom of expression and "look" more like the systems they represent. Simulation languages have been developed to this end (for a comprehensive overview of simulation languages, see [155]); general-purpose programming languages are also used, e.g., [156], [157]. Generality of expression is extremely important if we are to capture the complicated unexpected interactions that trigger security failures.

On the other hand, certain types of simulation models demand less detail, rather than more. A good example of this is Internet worm propagation, for which differential equations have been used [158], [159], [160], [161], [162]. The need for such high-level abstractions is computational; worms propagate over the entire Internet, and it is computationally infeasible to model individually infected hosts, sending individual probes and infection packets.

#### 3.5.3 Simulation and Security Analysis Today

By far, the most common use of simulation today in a security context is the use of normal network simulation tools to model a system and then model traffic representing attacks. Such models are useful for quantifying diverse system measures in the midst of attacks and countermeasures, e.g., see [163], [164], [165]. Equally common, though, is the use of simulation as a means of education and discovery. For example, there is an effort underway at the Naval Postgraduate School [166], [167] to build a simulator designed to get students engaged in the process of system configuration and cyber attack defenses, in addition to playing the role of cyber attackers. The interest here is less on the specifics of quantifying system behavior, and more on providing enough realism for students to learn about security issues. Similar efforts are underway elsewhere [168], [169].

Simulation is also being used to help government agencies practice appropriate responses to cyber attacks on their information technology infrastructure. Two notable examples are a cyber attack exercise conducted in Seattle in conjunction with the May 2003 TOPOFF dirty-bomb exercise [170] and the October Livewire cyber war exercise [171]. In the Livewire exercise, simulated attacks on a simulated network caused disruption of simulated services. Exercise players then worked through their responses (e.g., who to call, appropriate network reconfiguration, and so forth) to degraded capability. Again, the users looked to the simulation to provide a "realistic" simulated network behavior, not to quantify system metrics precisely (although capturing the general trend of system metrics under cyber attack is critical to the whole approach).

Despite the seeming divergence of this style of simulation from very mathematical quantitative analyses, there is potential for closer linkage. Loose requirements on the accuracy of quantitative measures open the possibility for other types of models and analyses to play an important role as fast approximations, albeit hidden from the user. We can look forward to work that incorporates diverse modeling methodologies in such applications of simulation.

## 4 Challenges and Conclusions

In the previous sections, we have reviewed measures that are pertinent to dependability and security evaluation, surveyed existing techniques for dependability evaluation, and given examples of how those techniques are currently being applied to the evaluation of certain security properties. While these applications suggest that there is merit to using stochastic techniques to evaluate security properties, they also suggest that significant new work is necessary to create a sound, model-based framework for quantifying system security.

That goal is clearly important since history suggests that it will be difficult, if not impossible, to build systems that can be shown to be perfectly secure. Hence, in order to have confidence that a given design will perform its intended function, we must be able to quantify its security. At the highest level, we believe that this work falls into two categories: 1) modeling attacker behavior and 2) creating a single, comprehensive methodology for evaluating whether a design meets one or more high-level requirements related to security. We outline the issues and challenges related to each of these needs in the following.

The first challenge is related to appropriate modeling of the behavior of cyber attackers. Just as appropriate fault models are critical to dependability evaluation, appropriate attacker models are critical to quantitative security evaluation. Determining the appropriate level of detail/abstraction in an attacker model is very important and depends on the scope and purpose of the model. Different attacker models will be needed for different purposes and different attack classes. For a given model, the level of detail/abstraction that is appropriate will depend on many factors. For example, the system submodels should represent the parts of a system that are important, relative to the types of attacks considered and the expression of a particular security measure. In particular, they must be detailed enough to support the expression of those parts of state that an attacker may change and those parts that may change his behavior.

Depending on the nature of the attack, the attacker model may either represent details of the attack or intrusion itself (corresponding to explicit representation of a fault in a dependability model) or represent the effect of the intrusion (corresponding to the representation of the error in a dependability model). We believe that by representing attacker behavior in terms of effects, rather than attacks/intrusions, we can cover a large class of attacks/intrusions (including unknown ("zero day") attacks) in a model.

Development of a comprehensive methodology for system-level security quantification is also a significant challenge. As described earlier in the paper, stochastic evaluation techniques originally intended for use in dependability evaluation have been successfully used to evaluate certain security attributes, including availability and survivability. However, other attributes, such as confidentiality and nonrepudiation, are more difficult to evaluate using standard, stochastic, techniques.

These measures may be better validated via so-called "formal" methods. The different natures of these multiple security measures suggest that the individual application of any of the techniques we have described is insufficient to validate large systems that are intended to be secure. While each of those techniques has the ability to evaluate certain kinds of security measures, such abilities may have limitations when one is attempting to use a single approach to validate the system with respect to a fairly high-level security requirement.

What is needed is an integrated validation framework that permits the use of multiple evaluation techniques in an organized manner. Starting with a system and a high-level set of security requirements, the framework should provide a top-down approach to methodically break the problem of validating the system with respect to its security requirements into manageable tasks, and provide steps that deal with each of those tasks. Each step would use one or more individual evaluation techniques. A symbiotic relationship should be established among the various techniques such that they complement and supplement each other to build the overall argument.

Such an approach would make it possible to handle the validation of very large-scale systems, producing systematic and well-documented arguments about their security. Such integrated approaches to evaluation have been applied in the safety community, resulting in so-called "safety cases," suggesting that a similar approach might be useable for security quantification.

In summary, stochastic evaluation techniques inspired by dependability evaluation methods have the potential to be used, with appropriate extension, for security evaluation. Several studies that take this approach have already been made, indicating the promise of this approach. However, there are still significant obstacles to the creation of a comprehensive, integrated approach to the evaluation of multiple security properties, as outlined above. There are ample opportunities for further research.

## ACKNOWLEDGMENTS

The authors would like to thank their past and current research sponsors for supporting this work, and their colleagues and students, both past and present, who also contributed to the research described herein. They would particularly like to thank Mr. Salem Derisavi, Mr. Yun Liu, Dr. Bharat Madan, and Mr. Dazhi Wang in this regard. They would also like to thank Ms. Jenny Applequist for her editorial assistance. Their sponsors include the US National Science Foundation (CCR-0209144, EIA-99-75019, INT-0233490, CCR-0311616, CNS-0406351), AFOSR MURI (F49620-1-0327), the Defense Advanced Research Projects Agency, Motorola, and Pioneer Hi-Bred. Their work was also supported under Award No. 2000-DT-CX-K001 from the Office for Domestic Preparedness, US Department of Homeland Security. Points of view in this document are those of the author(s) and do not necessarily represent the official position of the US Department of Homeland Security, or any of the other sponsors. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the US National Science Foundation.

## REFERENCES

[1] Y. Deswarte, L. Blain, and J.C. Fabre, "Intrusion Tolerance in Distributed Computing Systems," *Proc. IEEE Symp. Research in Security and Privacy,* pp. 110-121, May 1991.

[2] B. Dutertre, V. Crettaz, and V. Stavridou, "Intrusion-Tolerant Enclaves," *Proc. IEEE Int'l Symp. Security and Privacy,* pp. 216-224, May 2002.

[3] M. Cukier, J. Lyons, P. Pandey, H.V. Ramasamy, W.H. Sanders, P. Pal, F. Webber, R. Schantz, J. Loyall, R. Watro, M. Atighetchi, and J. Gossett, "Intrusion Tolerance Approaches in ITUA," *Supplement of the Proc. 2001 Int'l Conf. Dependable Systems and Networks,* pp. B-64-B-65, July 2001.

[4] F. Wang, F. Gong, C. Sargor, K. Goševa-Popstojanova, K.S. Trivedi, and F. Jou, "SITAR: A Scalable Intrusion Tolerance Architecture for Distributed Services," *Proc. IEEE Second SMC Information Assurance Workshop,* pp. 38-45, June 2001.

[5] C. Landwehr, "Formal Models for Computer Security," *Computer Surveys,* vol. 13, no. 3, pp. 247-278, Sept. 1981.

[6] J. Lowry, "An Initial Foray into Understanding Adversary Planning and Courses of Action," *Proc. DARPA Information Survivability Conf. and Exposition II (DISCEX '01),* pp. 123-133, 2001.

[7] A. Avizienis, J. Laprie, and B. Randell, "Fundamental Concepts of Dependability," LAAS-CNRS, Technical Report N01145, Apr. 2001.

[8] K.S. Trivedi, *Probability and Statistics with Reliability, Queuing, and Computer Science Applications,* second ed. New York: John Wiley and Sons, 2001.

[9] M.L. Shooman, *Probabilistic Reliability: An Engineering Approach,* second ed. Malabar, Fla.: R.E. Krieger Publishing Co., 1990.

[10] B. Madan, K. Goševa-Popstojanova, K. Vaidyanathan, and K. Trivedi, "Modeling and Quantification of Security Attributes of Software Systems," *Proc. Int'l Conf. Dependable Systems and Networks,* pp. 505-514, 2002.

[11] S. Jha, O. Sheyner, and J. Wing, "Minimization and Reliability Analysis of Attack Graphs," Technical Report CMU-CS-2-109, Carnegie Mellon Univ., May 2002.

[12] O. Sheyner, J. Haines, S. Jha, R. Lippmann, and J. Wing, "Automated Generation and Analysis of Attack Graphs," *Proc. 2002 IEEE Symp. Security and Privacy,* pp. 273-284, May 2002.

[13] S. Singh, M. Cukier, and W.H. Sanders, "Probabilistic Validation of an Intrusion-Tolerant Replication System," *Proc. Int'l Conf. Dependable Systems and Networks,* pp. 616-624, June 2003.

[14] V. Gupta, V.V. Lam, H.V. Ramasamy, W.H. Sanders, and S. Singh, "Dependability and Performance Evaluation of Intrusion-Tolerant Server Architectures," *Dependable Computing: Proc. First Latin-Am. Symp. (LADC 2003),* pp. 81-101, 2003.

[15] J.F. Meyer, "On Evaluating the Performability of Degradable Computing Systems," *IEEE Trans. Computers,* vol. 29, no. 8, pp. 720-731, Aug. 1980.

[16] W.H. Sanders and J.F. Meyer, "A Unified Approach for Specifying Measures of Performance, Dependability, and Performability," *Dependable Computing for Critical Applications, Vol. 4 of Dependable Computing and Fault-Tolerant Systems,* A. Avizienis, H. Kopetz, and J. Laprie, eds., Springer-Verlag, pp. 215-237, 1991.

[17] G. Bolch, S. Greiner, H. de Meer, and K.S. Trivedi, *Queueing Networks and Markov Chains.* New York: John Wiley & Sons, 1998.

[18] R.J. Ellison, D.A. Fisher, R.C. Linger, H.F. Lipson, T. Longstaff, and N.R. Mead, "Survivable Network Systems: An Emerging Discipline," Technical Report CMU/SEI-97-TR-013, CMU Software Engineering Institute, Nov. 1997.

[19] Y. Liu and K.S. Trivedi, "A General Framework for Network Survivability Quantification," *Proc. 12th GI/ITG Conf. Measuring, Modelling and Evaluation of Computer and Comm. Systems (MMB) together with Third Polish-German Teletraffic Symp. (PGTS),* 2004.

[20] Y. Liu, V.B. Mendiratta, and K.S. Trivedi, "Survivability Analysis of Telephone Access Network," *Proc. IEEE Int'l. Symp. Software Eng. (ISSRE '04),* 2004.

[21] F. Stevens, T. Courtney, S. Singh, A. Agbaria, J.F. Meyer, W.H. Sanders, and P. Pal, "Model-Based Validation of an Intrusion-Tolerant Information System," *Proc. 23rd Symp. Reliable Distributed Systems (SRDS 2004),* Oct. 2004.

[22] R.A. Sahner, K.S. Trivedi, and A. Puliafito, *Performance and Reliability Analysis of Computer Systems: An Example-Based Approach Using the SHARPE Software Package.* Kluwer Academic Publishers, 1996.

[23] http://www.relexsoftware.com/products/relanalysissoft.asp, 2004.

[24] M. Malhotra and K. Trivedi, "Power-Hierarchy of Dependability Model Types," *IEEE Trans. Reliability,* vol. 43, no. 2, pp. 493-502, Sept. 1994.

[25] S. Rai, M. Veeraraghavan, and K. Trivedi, "A Survey on Efficient Computation of Reliability Using Disjoint Products Approach," *Networks,* vol. 25, no. 3, pp. 147-163, 1995.

[26] R.E. Bryant, "Graph Based Algorithms for Boolean Function Manipulation," *IEEE Trans. Computers,* vol. 35, no. 8, pp. 677-691, Aug. 1986.

[27] X. Zang, H. Sun, and K. Trivedi, "A BDD-Based Algorithm for Reliability Analysis of Phased-Mission Systems," *IEEE Trans. Reliability,* vol. 48, no. 1, pp. 50-60, Mar. 1999.

[28] J.E. Arsenault and J.A. Roberts, *Reliability and Maintainability of Electronic Systems.* Rockville, MD: Computer Science Press, 1980.

[29] R.E. Barlow and F. Proschan, *Statistical Theory of Reliability and Life Testing.* New York: Holt, Rinehart and Winston, 1975.

[30] B.S. Dhillon and C. Singh, *Engineering Reliability: New Techniques and Applications.* New York: Wiley, 1981.

[31] E. Henley and H. Kumamoto, *Reliability Engineering and Risk Assessment.* Englewood Cliffs, N.J.: Prentice-Hall, 1981.

[32] N.G. Leveson, *Safeware: System Safety and Computers.* Addison-Wesley Publishing Co., 1995.

[33] J.B. Dugan and M.R. Lyu, "Dependability Modeling for Fault-Tolerant Software and Systems," *Software Fault Tolerance,* M.R. Lyu, ed., Chichester: John Wiley & Sons, pp. 109-138, 1995.

[34] J.B. Dugan, S.J. Bavuso, and M.A. Boyd, "Fault Trees and Sequence Dependencies," *Proc. Reliability and Maintainability Symp.,* pp. 286-293, 1990.

[35] J.B. Dugan, "Fault Trees and Imperfect Coverage," *IEEE Trans. Reliability,* vol. 38, no. 2, pp. 177-185, 1989.

[36] X. Zang, D. Wang, H. Sun, and K. Trivedi, "A BDD-Based Algorithm for Analysis of Multistate Systems with Multistate Components," *IEEE Trans. Computers,* vol. 52, no. 12, pp. 1608-1618, Dec. 2003.

[37] Y. Ma and K. Trivedi, "An Algorithm for Reliability Analysis of Phased-Mission Systems," *Reliability Eng. and System Safety,* vol. 66, no. 2, pp. 157-170, 1999.

[38] "CAFTA: A Fault Tree Analysis Tool Designed for PSA," *Proc. Probabilistic Risk Assessment and Risk Management Conf. (PSA '87),* vol. 2, pp. 588-592, 1987.

[39] http://www.ds-s.com/risk_and_reliability_tools.asp, 2004.

[40] B. Schneier, *Secrets and Lies: Digital Security in a Networked World.* John Wiley & Sons, Aug. 2000.

[41] D. Dolev and A. Yao, "On the Security of Public Key Protocols," *IEEE Trans. Information Theory,* vol. 29, no. 2, pp. 198-208, 1983.

[42] C. Meadows, "Applying Formal Methods to the Analysis of a Key Management Protocol," *J. Computer Security,* vol. 1, no. 1, pp. 5-36, 1992.

[43] T. Woo and S. Lam, "A Semantic Model for Authentication Protocols," *Proc. 1993 IEEE Symp. Security and Privacy,* pp. 178-195, 1993.

[44] W. Marrero, E. Clark, and S. Jha, "Modeling Checking for Security Protocols," Technical Report CMU-SCS-97-139, Carnegie Mellon Univ., May 1997.

[45] F. Besson, J. Jensen, D.L. Métayer, and T. Thorn, "Model Checking Security Properties of Control Flow Graphs," *J. Computer Security,* vol. 9, no. 3, pp. 217-250, 2001.

[46] H. Chen, D. Dean, and D. Wagner, "Model Checking One Million Lines of C Code," *Proc. 11th Ann. Network and Distributed System Security Symp.,* 2004.

[47] R.W. Ritchey and P. Ammann, "Using Model Checking to Analyze Network Vulnerabilities," *Proc. IEEE Symp. Security and Privacy,* pp. 156-165, May 2000.

[48] G. Ciardo and A.S. Miner, "Efficient Reachability Set Generation and Storage Using Decision Diagrams," *Proc. 20th Int'l Conf. Application and Theory of Petri Nets,* pp. 6-25, 1999.

[49] J.R. Burch, E.M. Clarke, K.L. McMillan, D.L. Dill, and L.J. Hwang, "Symbolic Model Checking: $10^{20}$ States and Beyond," *Information and Computation,* vol. 98, no. 2, pp. 142-170, June 1992.

[50] S. Singh, J. Lyons, and D. Nicol, "Fast Model-Based Penetration Testing," *Proc. 2004 Winter Simulation Conf.,* Dec. 2004.

[51] J.K. Muppala, M. Malhotra, and K.S. Trivedi, "Markov Dependability Models of Complex Systems: Analysis Techniques," *Reliability and Maintenance of Complex Systems,* S. Ozekici, ed., Germany: Springer, pp. 442-486, 1996.

[52] B. Haverkort, R. Marie, G. Rubino, and K.S. Trivedi, *Performability Modeling Tools and Techniques.* Chichester, England: John Wiley & Sons, 2001.

[53] K.S. Trivedi, J.K. Muppala, S.P. Woolet, and B.R. Haverkort, "Composite Performance and Dependability Analysis," *Performance Evaluation,* vol. 14, no. 3-4, pp. 197-215, 1992.

[54] J.K. Muppala, S.P. Woolet, and K.S. Trivedi, "Real-Time Systems Performance in the Presence of Failures," *Computer,* vol. 24, no. 5, pp. 37-47, May 1991.

[55] K.S. Trivedi, S. Ramani, and R.M. Fricks, "Recent Advances in Modeling Response-Time Distributions in Real-Time Systems," *Proc. IEEE,* vol. 91, no. 7, pp. 1023-1037, 2003.

[56] K.G. Popstojanova and K. Trivedi, "Architecture Based Approach to Reliability Assessment of Software Systems," *Performance Evaluation,* vol. 45, no. 2-3, pp. 179-204, 2001.

[57] S. Garg, Y. Huang, C.M.R. Kintala, S. Yajnik, and K. Trivedi, "Performance and Reliability Evaluation of Passive Replication Schemes in Application Level Fault Tolerance," *Proc. 29th Int'l Symp. Fault-Tolerant Computing,* pp. 322-328, June 1999.

[58] J.-C. Laprie and K. Kanoun, "X-Ware Reliability and Availability Modeling," *IEEE Trans. Software Eng.,* vol. 18, no. 2, pp. 130-147, 1992.

[59] W.H. Sanders and J.F. Meyer, "Performability Evaluation of Distributed Systems Using Stochastic Activity Networks," *Proc. Int'l Conf. Petri Nets and Performance Models,* pp. 111-120, 1987.

[60] Y. Ma, J. Han, and K. Trivedi, "Composite Performance & Availability Analysis of Wireless Communication Networks," *IEEE Trans. Vehicular Technology,* vol. 50, no. 5, pp. 1216-1223, Sept. 2001.

[61] R.M. Smith, K.S. Trivedi, and A.V. Ramesh, "Performability Analysis: Measures, an Algorithm, and a Case Study," *IEEE Trans. Computers,* vol. 37, no. 4, pp. 406-417, Apr. 1988.

[62] W.J. Stewart, *Introduction to the Numerical Solution of Markov Chains.* Princeton, 1994.

[63] A. Reibman, R.M. Smith, and K. Trivedi, "Markov and Markov Reward Models: A Survey of Numerical Approaches," *European J. Operations Research,* pp. 257-267, 1989.

[64] G. Ciardo, J. Muppala, and K. Trivedi, "SPNP: Stochastic Petri Net Package," *Proc. Third Int'l Workshop Petri Nets and Performance Models,* pp. 142-151, 1989.

[65] W.H. Sanders, W.D. Obal, M.A. Qureshi, and F.K. Widjanarko, "The *UltraSAN* Modeling Environment," *Performance Evaluation,* vol. 24, no. 1, pp. 89-115, Oct.-Nov. 1995.

[66] D.D. Deavours, G. Clark, T. Courtney, D. Daly, S. Derisavi, J.M. Doyle, W.H. Sanders, and P.G. Webster, "The Möbius Framework and Its Implementation," *IEEE Trans. Software Eng.,* vol. 28, no. 10, pp. 956-969, Oct. 2002.

[67] K. Vaidyanathan and K. Trivedi, "A Measurement-Based Model for Estimation of Resource Exhaustion in Operational Software Systems," *Proc. 10th Int'l Symp. Software Reliability Eng.,* pp. 84-93, Nov. 1999.

[68] M.C. Hsueh, R. Iyer, and K. Trivedi, "Performability Modeling Based on Real Data: A Case Study," *IEEE Trans. Computers,* vol. 37, no. 4, pp. 478-484, Apr. 1988.

[69] D. Chen, D. Selvamuthu, D. Chen, L. Li, R.R. Some, A.P. Nikora, and K. Trivedi, "Reliability and Availability Analysis for the JPL Remote Exploration and Experimentation System," *Proc. Int'l Conf. Dependable Systems and Networks,* pp. 337-344, June 2002.

[70] J.K. Muppala, A.S. Sathaye, R.C. Howe, and K.S. Trivedi, "Dependability Modeling of a Heterogenous VAXcluster System Using Stochastic Reward Nets," *Hardware and Software Fault Tolerance in Parallel Computing Systems,* Ellis Horwood Ltd., pp. 33-59, 1992.

[71] V. Mainkar and K. Trivedi, "Sufficient Conditions for Existence of a Fixed Point in Stochastic Reward Net-Based Iterative Models," *IEEE Trans. Software Eng.,* vol. 22, no. 9, pp. 640-653, Sept. 1996.

[72] L. Tomek and K. Trivedi, "Fixed-Point Iteration in Availability Modeling," *Informatik-Fachberichte, Vol. 283: Fehlertolerierende Rechensysteme,* Springer-Verlag, Berlin, pp. 229-240, 1991.

[73] J.G. Kemeney and J.L. Snell, *Finite Markov Chains.* D. Van Nostrand Company, Inc., 1960.

[74] P. Buchholz, "Exact and Ordinary Lumpability in Finite Markov Chains," *J. Applied Probability,* vol. 31, pp. 59-74, 1994.

[75] P. Buchholz, "Efficient Computation of Equivalent and Reduced Representations for Stochastic Automata," *Int'l J. Computer Systems Science & Eng.,* vol. 15, no. 2, pp. 93-103, 2000.

[76] R. Milner, *Communication and Concurrency.* London: Prentice Hall, 1989.

[77] P.C. Kanellakis and S.A. Smolka, "CCS Expressions, Finite State Processes, and Three Problems of Equivalence," *Proc. ACM Symp. Principles of Distributed Computing,* pp. 228-240, 1983.

[78] R. Paige and R.E. Tarjan, "Three Partition Refinement Algorithms," *SIAM J. Computing,* vol. 16, no. 6, pp. 973-989, 1987.

[79] J.C. Fernandez, "An Implementation of an Efficient Algorithm for Bisimulation Equivalence," *Science of Computer Programming,* vol. 13, no. 2-3, pp. 219-236, 1990.

[80] M. Bernardo and R. Gorrieri, "A Tutorial on EMPA: A Theory of Concurrent Processes with Nondeterminism, Priorities, Probabilities and Time," *Theoretical Computer Science,* vol. 202, pp. 1-54, 1998.

[81] D.T. Huynh and L. Tian, "On Some Equivalence Relations for Probabilistic Processes," *Fundamenta Informaticae,* vol. 17, pp. 211-234, 1992.

[82] S. Derisavi, H. Hermanns, and W.H. Sanders, "Optimal State-Space Lumping in Markov Chains," *Information Processing Letters,* vol. 87, no. 6, pp. 309-315, Sept. 2003.

[83] G. Chiola, C. Dutheillet, G. Franceschinis, and S. Haddad, "Stochastic Well-Formed Colored Nets and Symmetric Modeling Applications," *IEEE Trans. Computers,* vol. 42, no. 11, pp. 1343-1360, Nov. 1993.

[84] W.H. Sanders and J.F. Meyer, "Reduced Base Model Construction Methods for Stochastic Activity Networks," *IEEE J. Selected Areas in Comm.,* vol. 9, no. 1, pp. 25-36, Jan. 1991.

[85] W.D. Obal II, "Measure-Adaptive State-Space Construction Methods," PhD Dissertation, Univ. of Arizona, 1998.
[86] H. Hermanns and M. Ribaudo, "Exploiting Symmetries in Stochastic Process Algebras," *Proc. 12th European Simulation Multiconf. (ESM),* pp. 763-770, 1998.
[87] S. Gilmore, J. Hillston, and M. Ribaudo, "An Efficient Algorithm for Aggregating PEPA Models," *IEEE Trans. Software Eng.,* vol. 27, no. 5, pp. 449-464, May 2001.
[88] P. Buchholz, "Exact Performance Equivalence: An Equivalence Relation for Stochastic Automata," *Theoretical Computer Science,* vol. 215, no. 1/2, pp. 263-287, 1999.
[89] P. Buchholz, "Hierarchical Markovian Models: Symmetries and Reduction," *Performance Evaluation,* vol. 22, no. 1, pp. 93-110, Feb. 1995.
[90] P. Buchholz, "Markovian Process Algebra: Composition and Equivalence," *Proc. Second Workshop Process Algebras and Performance Modelling,* Arbeitsberichte des IMMD, vol. 27, no. 4, pp. 11-30, 1994.
[91] P. Buchholz, "Equivalence Relations for Stochastic Automata Networks," *Computation with Markov Chains,* W.J. Stewart, ed. Kluwer Int'l Publishers, pp. 197-216, 1995.
[92] P. Buchholz, "A Framework for the Hierarchical Analysis of Discrete Event Dynamic Systems (habilitations thesis)," PhD dissertation, Univ. Dortmund, Germany, 1996.
[93] H. Hermanns, *Interactive Markov Chains and the Quest for Quantified Quality.* Springer, LNCS vol. 2428, 2002.
[94] P.-J. Courtois and P. Semal, "Computable Bounds for Conditional Steady-State Probabilities in Large Markov Chains and Queueing Models," *IEEE J. Selected Areas in Comm.,* vol. 4, no. 6, pp. 926-937, Sept. 1986.
[95] P.J. Courtois, *Decomposability.* New York: Academic Press, 1977.
[96] A. Bobbio and K. Trivedi, "An Aggregation Technique for the Transient Analysis of Stiff Markov Chains," *IEEE Trans. Computers,* vol. 35, no. 9, pp. 803-814, Sept. 1986.
[97] A. Bobbio and K.S. Trivedi, "Computing Cumulative Measures of Stiff Markov Chains Using Aggregation," *IEEE Trans. Computers,* vol. 39, no. 10, pp. 1291-1297, 1990.
[98] D. Daly, P. Buchholz, and W.H. Sanders, "An Approach for Bounding Reward Measures in Markov Models Using Aggregation," Technical Report UILU-ENG-04-2206 (CRHC-04-06), Univ. of Illinois at Urbana-Champaign Coordinated Science Laboratory, July 2004.
[99] A. Srinivasan, T. Kam, S. Malik, and R.E. Brayton, "Algorithms for Discrete Function Manipulation," *Proc. Int'l Conf. CAD (ICCAD '90),* pp. 92-95, 1990.
[100] E.M. Clarke, O. Grumberg, and D. Peled, *Model Checking.* MIT Press, 1999.
[101] G. Ciardo, G. Lüttgen, and R. Siminiceanu, "Saturation: An Efficient Iteration Strategy for Symbolic State-Space Generation," *Proc. Int'l Conf. Tools and Algorithms for the Construction and Analysis of Systems,* pp. 328-342, 2001.
[102] G. Ciardo, R.M. Marmorstein, and R. Siminiceanu, "Saturation Unbound," *Proc. Int'l Conf. Tools and Algorithms for the Construction and Analysis of Systems,* pp. 379-393, 2003.
[103] B. Plateau, "On the Stochastic Structure of Parallelism and Synchronization Models for Distributed Algorithms," *Proc. ACM SIGMETRICS Conf. Measurement and Modeling of Computer Systems,* pp. 147-154, 1985.
[104] B. Plateau and K. Atif, "Stochastic Automata Network for Modeling Parallel Systems," *IEEE Trans. Software Eng.,* vol. 17, no. 10, pp. 1093-1108, Oct. 1991.
[105] P. Buchholz, "Numerical Solution Methods Based on Structured Descriptions of Markovian Models," *Computer Performance Evaluation,* Elsevier Science Publishers B.V. (North-Holland), pp. 251-267, 1991.
[106] P. Buchholz and P. Kemper, "Numerical Analysis of Stochastic Marked Graphs," *Proc. Sixth Int'l Workshop Petri Nets and Performance Models (PNPM '95),* pp. 32-41, Oct. 1995.
[107] S. Donatelli, "Superposed Stochastic Automata: A Class of Stochastic Petri Nets Amenable to Parallel Solution," *Proc. Fourth Int'l Workshop Petri Nets and Performance Models,* pp. 54-63, 1991.
[108] S. Donatelli, "Superposed Generalized Stochastic Petri Nets: Definition and Efficient Solution," *Proc. 15th Int'l Conf. Applications and Theory of Petri Nets,* pp. 258-277, 1994.
[109] P. Kemper, "Numerical Analysis of Superposed GSPNs," *Proc. Sixth Int'l Workshop Petri Nets and Performance Models (PNPM '95),* pp. 52-61, 1995.
[110] P. Buchholz, G. Ciardo, S. Donatelli, and P. Kemper, "Complexity of Memory-Efficient Kronecker Operations with Applications to the Solution of Markov Models," *INFORMS J. Computing,* vol. 12, no. 3, pp. 203-222, 2000.
[111] G. Ciardo and A. Miner, "A Data Structure for the Efficient Kronecker Solution of GSPNs," *Proc. Eighth Int'l Workshop Petri Nets and Performance Models,* pp. 22-31, 1999.
[112] A.S. Miner, "Efficient Solution of GSPNs Using Canonical Matrix Diagrams," *Proc. Ninth Int'l Workshop Petri Nets and Performance Models,* pp. 101-110, Sept. 2001.
[113] D.D. Deavours and W.H. Sanders, "An Efficient Disk-Based Tool for Solving Large Markov Models," *Performance Evaluation,* vol. 33, pp. 67-84, 1998.
[114] D.D. Deavours and W.H. Sanders, "On-the-Fly' Solution Techniques for Stochastic Petri Nets and Extensions," *IEEE Trans. Software Eng.,* vol. 24, no. 10, pp. 889-902, Oct. 1998.
[115] E. de Souza e Silva and H.R. Gail, "Calculating Availability and Performability Measures of Repairable Computer Systems," *J. ACM,* vol. 36, pp. 171-193, Jan. 1989.
[116] E. de Souza e Silva and H.R. Gail, "Calculating Transient Distributions of Cumulative Reward," *Proc. SIGMETRICS/Performance-95,* pp. 231-240, May 1995.
[117] M.A. Qureshi and W.H. Sanders, "A New Methodology for Calculating Distributions of Reward Accumulated During a Finite Interval," *Proc. 26th Int'l Symp. Fault-Tolerant Computing,* pp. 116-125, June 1996.
[118] V.V. Lam, P. Buchholz, and W.H. Sanders, "A Structured Path-Based Approach for Computing Transient Rewards of Large CTMCs," *Proc. First Int'l Conf. Quantitative Evaluation of Systems (QEST),* Sept. 2004.
[119] J. Muppala, M. Malhotra, and K. Trivedi, "Stiffness-Tolerant Methods for Transient Analysis of Stiff Markov Chains," *Microelectronics and Reliability,* vol. 34, no. 11, pp. 1825-1841, 1994.
[120] A. Reibman and K.S. Trivedi, "Numerical Transient Analysis of Markov Models," *Computers and Operations Research,* vol. 15, no. 1, pp. 19-36, 1988.
[121] A. van Moorsel and W.H. Sanders, "Adaptive Uniformization," *ORSA Comm. in Statistics: Stochastic Models,* vol. 10, no. 3, pp. 619-648, Aug. 1994.
[122] A.P.A. van Moorsel and W.H. Sanders, "Transient Solution of Markov Models by Combining Adaptive & Standard Uniformization," *IEEE Trans. Reliability,* vol. 46, no. 3, pp. 430-440, Sept. 1997.
[123] M. Malhotra and A. Reibman, "Selecting and Implementing Phase Approximations for Semi-Markov Models," *Comm. Statistical Stochastic Models,* vol. 9, no. 4, pp. 473-506, 1993.
[124] S. Gokhale and K. Trivedi, "A Time/Structure Based Software Reliability Model," *Annals of Software Eng.,* vol. 8, pp. 85-121, 1999.
[125] S. Gokhale, P.N. Marinos, M.R. Lyu, and K. Trivedi, "Effect of Repair Policies on Software Reliability," *Proc. 12th Ann. Conf. Computer Assurance (COMPASS),* pp. 105-116, June 1997.
[126] R. Geist, M. Smotherman, K.S. Trivedi, and J.B. Dugan, "Reliability Analysis of Life-Critical Systems," *Acta Informatica,* vol. 23, no. 6, pp. 621-642, 1986.
[127] G. Ciardo, R.A. Marie, B. Sericola, and K.S. Trivedi, "Performability Analysis Using Semi-Markov Reward Processes," *IEEE Trans. Computers,* vol. 39, no. 10, pp. 1251-1264, Oct. 1990.
[128] V. Kulkarni, *Modeling and Analysis of Stochastic Systems.* New York: Chapman Hall, 1995.
[129] M. Ajmone Marsan, G. Balbo, and G. Conte, "A Class of Generalized Stochastic Petri Nets for the Performance Evaluation of Multiprocessor Systems," *ACM Trans. Computer Systems,* vol. 2, no. 2, pp. 93-122, 1984.
[130] J.F. Meyer, A. Movaghar, and W.H. Sanders, "Stochastic Activity Networks: Structure, Behavior, and Application," *Proc. Int'l Workshop Timed Petri Nets,* pp. 106-115, July 1985.
[131] W.H. Sanders and J.F. Meyer, "Stochastic Activity Networks: Formal Definitions and Concepts," *Lectures on Formal Methods and Performance Analysis, First EEF/Euro Summer School on Trends in Computer Science,* LNCS, no. 2090, pp. 315-343, 2001.
[132] J. Hillston, *A Compositional Approach to Performance Modelling.* Cambridge Univ. Press, 1996.
[133] H. Hermanns and M. Rettelbach, "Syntax, Semantics, Equivalences, and Axioms for MTIPP," *Proc. Second Workshop Process Algebras and Performance Modelling,* Arbeitsberichte des IMMD, vol. 27, no. 4, pp. 71-87, 1994.
[134] M. Malhotra and K. Trivedi, "Dependability Modeling Using Petri Nets," *IEEE Trans. Reliability,* vol. 44, no. 3, pp. 428-440, Sept. 1995.

[135] J.B. Dugan, V. Nicola, R. Geist, and K. Trivedi, "Extended Stochastic Petri Nets: Applications and Analysis," *Proc. Conf. Performance '84,* pp. 507-519, 1985.
[136] M.A. Marsan and G. Chiola, "On Petri Nets with Deterministic and Exponentially Distributed Firing Times," *Advances in Petri Nets,* LNCS, vol. 266, Springer, pp. 132-145, 1987.
[137] H. Choi, V. Kulkarni, and K. Trivedi, "Markov Regenerative Stochastic Petri Nets," *Performance Evaluation,* vol. 20, pp. 337-357, 1994.
[138] V. Catania, A. Puliafito, M. Scarpa, and L. Vita, "Concurrent Generalized Petri Nets," *Proc. Conf. Numerical Solution of Markov Chains,* pp. 359-382, Jan. 1995.
[139] G. Horton, V. Kulkarni, D. Nicol, and K.S. Trivedi, "Fluid Stochastic Petri Nets: Theory, Application, and Solution Techniques," *European J. Operations Research,* vol. 105, no. 1, pp. 184-201, Feb. 1998.
[140] G. Ciardo, D.M. Nicol, and K.S. Trivedi, "Discrete-Event Simulation of Fluid Stochastic Petri Nets," *IEEE Trans. Software Eng.,* vol. 25, no. 2, pp. 207-217, 1999.
[141] R. Ortalo, Y. Deswarte, and M. Kaâniche, "Experimenting with Quantitative Evaluation Tools for Monitoring Operational Security," *IEEE Trans. Software Eng.,* vol. 25, pp. 633-650, Oct. 1999.
[142] B. Littlewood, S. Brocklehurst, N. Fenton, P. Mellor, S. Page, and D. Wright, "Towards Operational Measures of Computer Security," *J. Computer Security,* vol. 2, pp. 211-229, 1993.
[143] E. Jonsson and T. Olovsson, "A Quantitative Model of the Security Intrusion Process Based on Attacker Behavior," *IEEE Trans. Software Eng.,* vol. 23, no. 4, pp. 235-245, Apr. 1997.
[144] M. Dacier, Y. Deswarte, and M. Kaâniche, "Quantitative Assessment of Operational Security: Models and Tools," Technical Report 96493, Laboratory for Analysis and Architecture of Systems, May 1996.
[145] D. Wang, B. Madan, and K. Trivedi, "Security Analysis of SITAR Intrusion-Tolerant System," *Proc. ACM Workshop Survivable and Self-Regenerative Systems,* pp. 23-32, 2003.
[146] W.H. Sanders, M. Cukier, F. Webber, P. Pal, and R. Watro, "Probabilistic Validation of Intrusion Tolerance," *Supplemental Volume of the Proc. Int'l Conf. Dependable Systems & Networks (DSN-2002),* pp. B-78-B-79, June 2002.
[147] http://www.nessus.org/, 2004.
[148] http://www.insecure.org/nmap/, 2004.
[149] D. Farmer and E.H. Spafford, "The COPS Security Checker System," *Proc. Summer Usenix Conf.,* pp. 165-170, 1990.
[150] http://www.net.tamu.edu/network/tools/tiger.html, 2004.
[151] A. Sharma, J.R. Martin, N. Anand, M. Cukier, and W.H. Sanders, "Ferret: A Host Vulnerability Checking Tool," *Proc. 10th IEEE Pacific Rim Int'l Symp. Dependable Computing (PRDC-10),* pp. 389-394, Mar. 2004.
[152] P. Heidelberger, "Fast Simulation of Rare Events in Queueing and Reliability Models," *ACM Trans. Modeling and Computer Simulation,* vol. 1, no. 5, pp. 43-85, 1995.
[153] J. Banks, J. Carson, B. Nelson, and D. Nicol, *Discrete-Event System Simulation.* Upper Saddle River, N.J.: Prentice-Hall, 2000.
[154] G. Shedler, *Regenerative Stochastic Simulation.* Boston: Prentice-Hall, 1993.
[155] R.E. Nance, "A History of Discrete Event Simulation Programming Languages," *Proc. Second ACM SIGPLAN Conf. History of Programming Languages,* pp. 149-175, 1993.
[156] http://www.isi.edu/nsnam/ns/, 2004.
[157] http://www.ssfnet.org, 2004.
[158] D. Moore, C. Shannon, and K. Claffy, "Code-Red: A Case Study on the Spread and Victims of an Internet Worm," *Proc. Internet Measurement Workshop (IMW),* pp. 273-284, Nov. 2002.
[159] D. Nicol, M. Liljenstam, and J. Liu, "Multiscale Modeling and Simulation of Worm Effects on the Internet Routing Infrastructure," *Proc. 13th Int'l Conf. Modeling Techniques and Tools for Computer Performance Evaluation (Performance TOOLS 2003),* pp. 1-10, Sept. 2003.
[160] M. Liljenstam, Y. Yuan, B. Premore, and D. Nicol, "A Mixed Abstraction Level Model of Large-Scale Internet Worm Infestations," *Proc. 10th IEEE/ACM Symp. Modeling, Analysis and Simulation of Computer and Telecomm. Systems (MASCOTS),* pp. 109-116, Oct. 2002.
[161] C.C. Zou, L. Gao, W. Gong, and D. Towsley, "Monitoring and Early Warning for Internet Worms," *Proc. 10th ACM Conf. Computer and Comm. Security,* pp. 190-199, 2003.
[162] C.C. Zou, W. Gong, and D. Towsley, "Code Red Worm Propagation Modeling and Analysis," *Proc. Ninth ACM Conf. Computer and Comm. Security,* pp. 138-147, 2002.
[163] V. Venkataraghavan, S. Nair, and P.-M. Seidel, "Simulation-Based Validation of Security Protocols," *Proc. OPNETWORKS 2002 Conf.,* Aug. 2002.
[164] D. Apostal, T. Foote-Lennox, T. Markham, A. Dowd, R. Lu, and D. O'Brian, "Checkmate Network Security Modeling," *Proc. DARPA Information Survivability Conf. and Exposition,* pp. 214-226, June 2001.
[165] V. Gorodetski, I. Kotenko, and O. Karsaev, "Multi-Agent Technologies for Computer Network Security: Attack Simulation, Intrusion Detection and Intrusion Detection Learning," *Int'l J. Computer Systems Science and Eng.,* vol. 18, no. 4, pp. 191-200, July 2003.
[166] N. Falby, M. Thompson, and C. Irvine, "A Security Simulation Game Definition Language," *Innovative Program Abstracts—Colloquium on Information Systems Security Education,* June 2004.
[167] C. Irvine and M. Thompson, "Teaching Objectives of a Simulation Game for Computer Security," *Proc. Informing Science and Information Technology Joint Conf.,* June 2003.
[168] J. Drew, "Simulation to Support Security Issues Related to System Interoperability," *Proc. Summer Simulation Conf.,* pp. 14-18, 2002.
[169] S. Lathrop, J. Hill, and J. Surdu, "Modeling Network Attacks," *Proc. 12th Conf. Behavior Representation in Modeling and Simulation,* pp. 401-407, May 2003.
[170] W. Dizard III, "Seattle Cybergame Preceded Last Week's Drill and Simulated Reality," *Government Computer News,* vol. 22, no. 11, http://www.gcn.com/22_11/news/22099-1.html, 2003.
[171] T. Bridis, "Gov't Simulates Terrorist Cyberattack," Assoc. Press, http://www.zone-h.org/en/news/read/id=3728, Nov. 2003.

**David M. Nicol** is a professor of electrical and computer engineering at the University of Illinois at Urbana-Champaign. He is a long-time contributor in the field of parallel and distributed discrete-event simulations, having written one of the early PhD theses on the topic. He has also worked in parallel algorithms, algorithms for mapping workload in parallel architectures, performance analysis, and reliability modeling and analysis. His research contributions extend to approximately 150 articles in leading computer science journals and conferences, and he is coauthor of the textbook *Discrete-Event Systems Simulation*. His current interests lie in modeling and simulation of very large systems, particularly communications and other infrastructure, with applications in security. He was coarchitect of the *Scalable Simulation Framework* (SSF) and several of its implementations, now in wide use for network analysis in education, industry, and government. From 1997 to 2003, he was the editor-in-chief of the *ACM Transactions on Modeling and Computer Simulation*, and from 2002-2003 he served as the associate director for research, and then the acting director of the Institute for Security Technology Studies at Dartmouth College. Professor Nicol is a fellow of the IEEE.

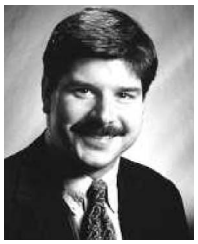

**William H. Sanders** is a professor in the Department of Electrical and Computer Engineering and the Coordinated Science Laboratory at the University of Illinois. He is vice-chair of IFIP Working Group 10.4 on Dependable Computing. In addition, he serves on the editorial board of *IEEE Transactions on Reliability* and is the area editor for simulation and modeling of computer systems for the *ACM Transactions on Modeling and Computer Simulation*. He is a past chair of the IEEE Technical Committee on Fault-Tolerant Computing. He is a fellow of the IEEE and the ACM. Dr. Sanders's research interests include performance/dependability/security evaluation and dependable and secure computing. He has published approximately 150 technical papers in these areas. He has served as an organizer and on the program committees of numerous conferences and workshops. He is a codeveloper of three tools for assessing the performability of systems represented as stochastic activity networks: METASAN, *UltraSAN*, and Möbius. Möbius and *UltraSAN* have been distributed widely to industry and academia; more than 300 licenses for the tools have been issued to universities, companies, and NASA for evaluating the performance, dependability, security, and performability of a variety of systems. He is also a codeveloper of the Loki distributed system fault injector and the AQuA/ITUA middlewares for providing dependability/security to distributed and networked applications.

**Kishor S. Trivedi** holds the Hudson Chair in the Department of Electrical and Computer Engineering at Duke University, Durham, North Carolina. He is the Duke-Site Director of a US National Science Foundation Industry-University Cooperative Research Center between North Carolina State University and Duke University for carrying out applied research in computing and communications. He has been on the Duke faculty since 1975. He is the author of a well-known text entitled *Probability and Statistics with Reliability, Queuing and Computer Science Applications,* with a thoroughly revised second edition being published by John Wiley. He has also published two other books, entitled *Performance and Reliability Analysis of Computer Systems* (Kluwer Academic Publishers) and *Queueing Networks and Markov Chains* (John Wiley). His research interests are in reliability and performance assessment of computer and communication systems. He has published more than 300 articles and lectured extensively on these topics. He has supervised 37 PhD dissertations. He is a fellow of the IEEE. He is a golden core member of the IEEE Computer Society. He is a codesigner of the HARP, SAVE, SHARPE, and SPNP software packages, which have been well-circulated.