

# Inoculating Software for Survivability

Anup K. Ghosh and Jeffrey M. Voas  
Reliable Software Technologies  
21515 Ridgetop Circle, #250  
Sterling, VA 20166  
{aghosh, jmvoas}@rstcorp.com  
www.rstcorp.com

## 1 Introduction

In early 1998, several dozen computer systems in U.S. military installations and government facilities were broken into via the Internet. The attacks successfully broke into systems belonging to the Navy and Air Force as well as at federally funded research laboratories including Oak Ridge National Laboratory, Brookhaven National Laboratories, U.C. Berkeley, and MIT. Although no classified systems allegedly were compromised, the attackers were able to obtain system privileges that could be used to read password files, delete files, or create back doors for later re-entry. Despite being called “the most organized and systematic attack” to date against DoD systems by the U.S. Deputy Defense Secretary John Hamre, these attacks were not the work of an organized terrorist group or nation; rather, authorities believe two northern California teenagers were responsible for hacking into these systems simply because they could.

Foretelling these attacks, the President’s Commission on Critical Infrastructure Protection (PCCIP) announced in October of 1997 that the increasing dependence of U.S. critical infrastructures on information and communications has made them vulnerable to information warfare attacks. The commission found that while the resources needed to conduct a physical attack against these infrastructures have not changed dramatically, the resources necessary to launch a comparable scale attack via information warfare are commonplace and consist of a personal computer and Internet connection. Furthermore, the ubiquity of Internet access and the ease of availability to cracker tools on underground Internet sites have significantly reduced both financial and intellectual barriers to launching effective attacks against critical systems. With roughly 95% of Defense Department communications relying on commercial infrastructure, the government finds itself as a major stake-holder in the security of commercial systems and is now proposing to spend U.S. \$1.46 billion to address the threat against critical infrastructures.

The Federal sector is not alone, however, in its concerns over information warfare. Wholesale payment systems such as the Federal Reserve’s FedWire and automated clearing houses (ACHs) move trillions of dollars over electronic networks. A compromise of the Federal Reserve system could dissolve trust in the electronic payments and clearing system, on which all banking transactions rely. As more corporations get connected online to inherently insecure public networks such as the Internet, their trade and financial secrets are being exposed and placed at risk. As societies transition to paperless commerce, individual privacy is threatened with each transaction. In short, as society becomes more “wired”, the security, privacy, and integrity of information becomes paramount. Likewise, the threat of information warfare looms ever larger.

At the heart of our national information infrastructure (NII) is software. Software is used to

enable the entire information infrastructure from the ubiquitous desktop Web browser to telecommunications switching software to the network servers and back-end databases. Software is pervasive in every component that enables the information economy. The greatest risk to our NII is failing software, be it from inadvertent flaws or from malicious attack. The most dangerous attacks against the information infrastructure are attacks against the software that composes it.

In this paper, we are concerned with the survivability of the infrastructure to software flaws, anomalous events, and malicious attack. In the past, finding and removing software flaws has traditionally been the realm of software testing. Software testing has largely concerned itself with ensuring that software behaves correctly — an intractable problem for any non-trivial piece of software. In this paper, we present “off-nominal” testing techniques that are not concerned with the correctness of the software, but with the survivability of the software in the face of anomalous events and malicious attack. Where software testing is focused on ensuring that the software computes the specified function correctly, we are concerned that the software continues to operate in the presence of unusual system events or malicious attacks.

The off-nominal testing approach uses fault injection analysis to determine the effect of unusual or malicious attacks against software. Fault injection is the process of perturbing or corrupting a data state during program execution. Fault injection analysis is the process of determining the effect of that perturbation. The analysis may consist of simply measuring whether the perturbed state affected a particular output, or the analysis may determine whether system attributes such as safety, security, or survivability have been affected [12].

We describe two applications of fault injection analysis: one to improve the survivability of software before release and one to test the survivability of software once deployed in a fielded system. The former approach is aimed at software vendors to provide additional assurance prior to releasing the software (and after performing traditional testing) that the software has been exercised under unusual conditions that might be otherwise unattainable via standard testing. Fault injection analysis is performed in the software source code in order to identify vulnerabilities in the source that can be potentially exploited to compromise system security and survivability. The results from the analysis can be used to harden the software against anomalous events or malicious attack. This approach is developed in detail with case studies to illustrate in Section 2.

The second approach of using fault injection analysis addresses the growing need to provide assurance of survivability in systems composed of Commercial Off The Shelf (COTS) software. The purpose of using fault injection analysis is, like before, to simulate anomalous conditions that would otherwise be difficult to obtain via standard testing, and to observe the resulting effect on system survivability. Today, few systems are ever built from the ground up, using custom-written software components. Instead, today’s software systems are a mixture of COTS software, legacy software, and custom-written software. Therefore, we must develop techniques that can provide assurance of survivability without requiring access to source code.

We describe an approach and tool that permits assessment for how robust a software program is under anomalous system resource conditions. For instance, if the operating system throws an exception during operation, the analysis can determine *a priori* how robust the software is to these anomalous conditions. The discussion presents a prototype tool for testing the robustness of Windows NT application software under anomalous operating system conditions.

## 2 Developing More Survivable Systems

Two approaches to improving the survivability of the NII are to: (1) develop more survivable systems, and (2) add survivability to fielded systems. In an ideal world, software development firms

would spend an appropriate amount of time and resources to developing more survivable systems. In reality, however, the commercial pressures to bring a product to market usually override concerns over providing rigorous assurance of security or survivability. As a result, little security/survivability testing is performed in software products prior to release, in spite of the potential for software flaws to adversely affect system security and survivability. However, even with market pressures set aside, there is little tool support for security and survivability oriented testing that even the best intentioned software development firms can use.

In this section we present an approach and tool that supports the first approach for improving the survivability of the NII, *i.e.*, enabling the development of more survivable system by providing security and survivability assurance technologies during software development. See [3, 1, 8, 7, 6] for other related security-oriented testing technologies. In the next section, we describe an approach and tool for assessing the survivability of fielded COTS systems. The approach recognizes that no matter how good (or inadequate) the efforts made to develop more survivable systems are, the complexity of today’s systems combined with different operational environments in which software is deployed makes the deployment of perfectly survivable systems a practical impossibility. Thus, technologies to assess the survivability of fielded systems in a particular environment to anomalous conditions are essential for finding vulnerabilities and retrofitting software with survivable mechanisms such as software wrappers.

Fault injection originated out of testing of integrated circuits, but recent advances have allowed it be applied to testing the safety properties of safety-critical systems. In [12], case studies of fault injection analysis are described that detect a serious flaw in the safety-monitoring routine of a nuclear control application, potential safety-critical hazards in a computer-controlled surgical device, and safety-critical flaws in a metropolitan subway control system.

Having demonstrated value to safety-critical systems, fault injection analysis has since been developed to analyze security properties in security-critical software systems. Specifically, it was theorized in [11] that fault injection can be used to find locations in source code where security-related vulnerabilities might exist. The idea is simple: perform fault injection in locations throughout the software source code, and then observe whether the program exhibits insecure or non-robust behavior. Those locations where fault injection resulted in undesirable behavior would require strengthening (or fault tolerance) to ensure that the corrupted internal states will never manifest during the actual use of the program. An application of fault injection analysis to security-critical software is described in detail in [4].

By discovering where critical flaws may be during product development using automated fault injection analysis, the opportunity to develop more survivable software systems—before damage has occurred—is afforded to the software vendor. If this proactive approach is employed by vendors of critical software within the information infrastructure (*e.g.*, operating systems, system utilities, network servers and clients) one result is that on the whole, the information infrastructure will be more survivable. Next, we describe a tool and case studies of using fault injection analysis.

## 2.1 A tool for fault injection security analysis

Fault injection analysis for identifying potential vulnerabilities in software has been implemented in a working tool named the Fault Injection Security Tool (FIST). The tool automates fault injection analysis of software using program inputs, fault injection functions, and assertions in programs written in C and C++.

A schematic diagram of FIST is shown in Figure 1. The fault injection engine provides the analyst the ability to instrument data variables with fault injection functions. The security policy assertion component provides the ability to capture security constraints on the software and to

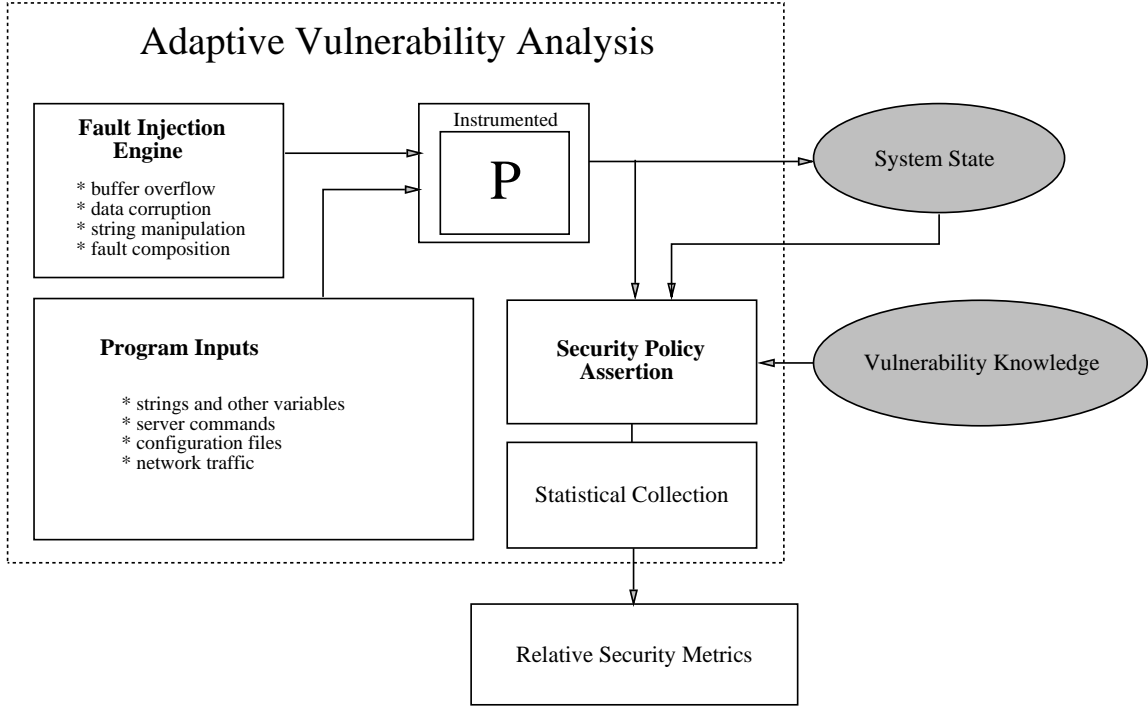


Figure 1: Overview of the Fault Injection Security Tool. A program,  $P$ , is instrumented with fault injection functions and assertions of its security policy (based on the vulnerability knowledge of the program). The program is exercised using program inputs. The security policy is evaluated online by examining program and system states. If a security policy assertion is violated during the dynamic analysis, the specific input and fault injection function that triggered the violation is identified.

determine if a security violation occurs during testing.

When a program is loaded into FIST, it is automatically instrumented with fault injection functions to perturb all possible data variables. Fault injection functions are instrumented by default for each different data type. For example, Booleans are perturbed to their opposite value during execution, integers are perturbed using a random function with a uniform distribution centered around their current value, character strings are perturbed using random values. In addition to these default settings, the user can programmatically instrument functions such as appending a particular string command to the end of a random string or instrumenting buffer overrun functions.

The buffer overrun function overwrites the return address of the stack frame in which the buffer variable is allocated with the address of the buffer itself. By tracing the frame pointer back through the stack, the fault injection function is able to determine where to overwrite the return address. The opcodes for machine instructions are written into the buffer being perturbed. If the program is vulnerable to a buffer overrun attack, the activation record containing the modified return address will be popped off the program stack and the program will jump to the machine instructions embedded by the fault injection function. These instructions will be executed as if they were a part of the normal operation of the program. This fault injection function can determine whether buffer variables are susceptible to buffer overrun attacks.

Once instrumented, the program is iteratively run, where for each test case, a different fault injection function is triggered on each run until all test cases and all feasible fault injection functions

Program	Instrumented Locations	Successful Simple Perturbations	Successful Buffer Overruns	Function Coverage
Samba v1.9.17p3	1264	12	15	45.5%
NCSA httpd v1.5.2a	463	27	3	40.14%
wu-ftp v2.4	476	11	3	58.62%
pop3d v1.005h	73	2	1	63.64%
kfingerd v0.07	146	12	5	38.1%

Table 1: **Results from fault injection analysis of network daemons.**

are executed (see [4] for the algorithm). This process is automated in an iterative execution environment. The effect of a single injected fault on program security is assessed by determining which assertions fire. The specific fault injection function that triggers a particular security assertion is identified. As a result, the analyst can tie the violation of security policy to a specific line of source code that when perturbed violates security. This information allows the developer to harden the code with fault-tolerant or survivable mechanisms such as assertions or stack guards [2].

## 2.2 Case studies of fault injection analysis

In a case study of fault injection analysis for software security, five common network services were analyzed [4]. Network daemons are interesting from a security standpoint because they provide services to untrusted users. Most network daemons allow connections from anywhere on the Internet, opening them up to attack from malicious users anywhere. Network daemons sometimes run with super-user, or `root`, privilege levels in order to bind to sockets on reserved ports, or to navigate the entire file system without being denied access. Successfully exploiting a weakness in a daemon running with high privileges can allow the attacker complete access to the server. Therefore, it is imperative that network daemons be free from security-related flaws that could permit untrusted users access to high privilege accounts on the server.

The programs analyzed were NCSA `httpd` version 1.5.2.a, the Washington University `wu-ftp` version 2.4, `kfingerd` version 0.07, the Samba daemon version 1.9.17p3, and `pop3d` version 1.005h. The source code for these programs is publicly available on the Internet. Samba, `httpd`, and `wu-ftp` are popular programs and can be found running on many sites on the Internet. The analysis of those programs was performed on a Sparc machine running SunOS 4.1.3-U. The other programs, `pop3d` and `kfingerd`, are Linux programs found in public repositories for Linux source code on the Internet. The analysis of those programs was performed on a Linux 2.0.0 kernel. The programs were instrumented with both simple fault injection functions as well as the buffer overrun functions where applicable.

A summary of results from the analysis is shown in Table 1. The table shows the total number of instrumented locations together with the number of simple perturbations and buffer overrun perturbations that resulted in security violations. Clearly, the automated analysis shows a number of “trouble” spots where fault injection functions violated the security policy of the software. In the case of buffer overruns, the security policy was simply that the program did not allow the buffer overrun function to execute its own code. In the case of the simple perturbations, the security policy involved illegal accesses to protected files.

The last column shows the percentage of the functions in the source code that were executed as a result of the test cases employed. Higher coverage results can be achieved through more testing and may result in more potential security hazards flushed out through the analysis.

### 3 Assessing the Survivability of COTS-based Systems

The preceding section described fault injection analysis as a viable technique for improving the survivability of software before it is released. The approach, however, is not a silver bullet solution for survivability, and market pressures tend to reward quicker release cycles of products with more and more features, *i.e.*, complexity, than stronger security and survivability. As a result, we cannot depend on software development and testing to produce survivable systems. Furthermore, even with the open source software movement afoot, most commercial software firms are loathe to release source code, which would permit peer review to identify bugs and vulnerabilities. As a result, we are bound by the practical constraints of commercial software releases to develop assurance technologies that can work with COTS software.

In this section, we describe an approach to software assurance that is designed for COTS software where the source code is not available, but executable binaries and application programming interfaces (APIs) are. The approach leverages fault injection analysis of software interfaces to analyze a critical attribute of survivability — robustness of software to anomalous events. Robustness is defined by the IEEE Standard Glossary of Software Engineering Terminology as “the degree to which a system or component can function correctly in the presence of invalid inputs or stressful environmental conditions”. In order to assure survivability of deployed software systems, we are concerned precisely with the unusual or stressful conditions that often arise in the field that are rarely tested by the software vendor.

In a break with traditional dependability research, we have applied this approach to the Microsoft Windows platform. The Windows platform (Windows 95, Windows 98, Windows NT, Windows CE, and Windows 2000) represents the most popular commercial platform and it is increasingly be used in critical applications. For example, under the Information Technology in the 21st century (IT-21) directive, the U.S. Navy requires its ships to migrate to Windows NT workstations and servers. While modernizing the fleets’ technology base is appropriate, the risks of migrating to new platforms are great, particularly in mission-critical applications. A stark example of the risks is illustrated by the saga of the USS Yorktown, a U.S. Navy Aegis missile cruiser, which suffered a significant software problem in the Windows NT systems that control the “smart ship”. Reportedly, an exception thrown by the Windows NT platform crashed the ship’s propulsion system software [10]. The end result: the ship had to be towed back to the Norfolk Naval shipyard.

In order to assess the survivability of COTS-based systems, we employ fault injection analysis on the interfaces between the software application and the operating system (OS). The fault injection functions simulate the effect of failing system resources, such as memory allocation errors, network failures, file input/output (I/O) problems, as well as the range of exceptions that can be thrown by OS functions when improperly used. The fault injection analysis tests the robustness of the application to unusual, anomalous, potentially malicious, and stressful environment conditions. An application is considered robust when it does not hang, crash, or disrupt the system in the presence of anomalous or invalid inputs, or stressful environmental conditions.

In our previous studies of the Windows platform, we analyzed the robustness of Windows NT OS functions to unexpected or anomalous inputs [5, 9]. We developed test harnesses and test data generators for testing OS functions with combinations of valid and anomalous inputs. Results from these studies show non-robust behavior from a large percentage of tested DLL functions when presented anomalous inputs. This information is particularly relevant to application developers that use these functions. That is, unless application developers are building in robustness to handle exceptions thrown by these functions, their applications may crash if they use these functions in unexpected ways.

Using traditional testing techniques to trigger non-robust operating system behavior may be

very difficult in practice. As a result, we employ fault injection functions at the software interface between the application and the operating system to artificially trigger non-robust operating system behavior in order to assess the robustness of the application.

### 3.1 A Failure Simulation Tool for Windows Applications

The approach to fault injection analysis of binary executables involves “wrapping” the software’s interface to the Win32 API with our own functions. The Win32 API is a set of functions standard on Windows NT, Windows 95/98, and Windows CE platforms. These functions exist in Dynamically Linked Libraries (DLLs) and represent programmer’s interface to the Windows operating system. The application’s Import Address Table (IAT), which is used to look up the address of imported DLL functions, is modified to point to our own wrapper DLL. When a target function is called by the application, the wrapper DLL is called instead. The wrapper DLL, in turn, executes, providing the ability to replace the value returned by the requested DLL function with an exception. Only exceptions that have been documented as part of the function’s interface or verified through actual testing are returned by the wrapper.

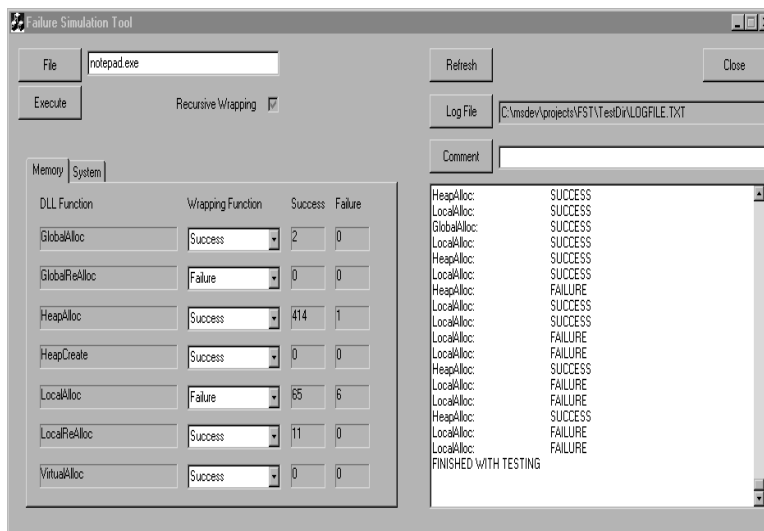


Figure 2: The graphical interface to the failure simulation tool that wraps Windows NT executable programs. The tool allows the user to interactively fail OS functions when they are called. The right panel shows a log of the calls made and the failures simulated, if any, for each call.

A failure simulation tool has been written to enable the user to interactively fail system functions during testing. If the application crashes, then we know that the application is non-robust to these exceptions thrown by the OS function.

Figure 2 shows the graphical user interface to the failure simulation tool that allows selective failing of operating system resources. The window shows the example memory functions that can be wrapped with failure or success functions. Other system functions such as file I/O functions are available for instrumentation via the **System** tab shown in the window in Figure 2.

The tool can be applied to any Win32 program. Testing of common desktop applications has already shown differences in robustness. For example, Microsoft’s Powerpoint 97 is more robust to exceptions thrown by the `CreateFile()` function than Microsoft’s Word 97. PowerPoint gives the user the opportunity to save work and exit when the `CreateFile()` function throws an exception. Word, on the other hand, immediately crashes when this function throws an exception. This

capability can be applied to any number of applications to test the robustness to failing operating system functions.

## 4 Retrofitting Survivability into COTS-based Systems

By using the failure simulation tool described in the preceding section, we can determine how robust or survivable a given application is to unusual or stressful environmental conditions (such as those one might encounter in a mission-critical environment). In order to increase the survivability of the vulnerable software, we have two options: (1) inform the software vendor of robustness/survivability problems and hope for a patch, or (2) harden the application with software wrappers. The former option is attractive in order to fix the problem at its source, however, the response might be less than desired. Unless it can be demonstrated that the failure of the application occurred in a non-simulated environment, *i.e.*, a *real* mission-critical failure, *and* that the non-robustness will impact a significant number of users, it is unlikely that the vendor will rectify the problem. Of course, waiting for a mission-critical failure to occur before complaining about a problem is too late to ensure survivability. Thus, the second option is attractive and we pursue it briefly in this section.

The approach leverages the wrapping method described in the preceding section. In the case that an application is non-robust to an exception thrown by an operating system function, the program can be wrapped to make it more robust. Instead of throwing exceptions to test robustness, the wrapper will catch any exceptions thrown by OS functions and return them as a specified error code. While many programmers will not handle exceptions, the return value from a function is almost always checked for specified error values. Thus, handling an exception thrown by an OS function at the wrapper and returning a specified error value is a robust way of dealing with a non-robust OS function.

Using the wrapping approach to testing for robustness described in the preceding section, the robustness of an application to error codes can be verified *a priori*. Therefore, the robustness of the wrapping approach to handling exceptions can be known before online deployment. The wrapper approach is particularly useful for mission-critical COTS software, where access to the source code is not available, but where robustness is important. The wrapper can be deployed with the application such that whenever the application is started, it is started with the wrapper in place.

## 5 Conclusions

The fault injection approaches we described in this paper as off-nominal testing can be thought of as a means of inoculating software for survivability. The analogy to vaccinations is apt. People are inoculated against disease by injecting infectious matter into the body in non-lethal forms. The body builds appropriate antibodies to the infectious matter in order to combat future infections of a more lethal instance of the disease. In the same way, fault injection analysis injects faults into an executing program in order to determine where it is vulnerable. Unfortunately, today we do not have automatic learning systems for protecting software states, though it is an endeavor of on-going research. Instead, once the program is found to be vulnerable through fault injection analysis, it can be retrofitted with fault-tolerant mechanisms in order to increase its likelihood of survivability.

## Acknowledgments

This work was sponsored by the Defense Advanced Research Projects Agency (DARPA) under Contracts F30602-95-C-0282 and F30602-97-C-0117. THE VIEWS AND CONCLUSIONS CONTAINED IN THIS DOCUMENT ARE THOSE OF THE AUTHORS AND SHOULD NOT BE INTERPRETED AS REPRESENTING THE OFFICIAL POLICIES, EITHER EXPRESSED OR IMPLIED, OF THE DEFENSE ADVANCED RESEARCH PROJECTS AGENCY OR THE U.S. GOVERNMENT. We also wish to acknowledge Tom O'Connor, Gary McGraw, Matt Schmid, and Frank Hill from Reliable Software Technologies for contributions to the work described in this article.

## References

- [1] M. Bishop and M. Dilger. Checking for race conditions in file accesses. In *The USENIX Association, Computing Systems*, pages 131–152, Spring 1996.
- [2] C. Cowan, C. Pu, D. Maier, H. Hinton, J. Walpole, P. Bakke, S. Beattie, A. Grier, P. Wagle, and Q. Zhang. Stackguard: Automatic adaptive detection and prevention of buffer-overflow attacks. In *Proceedings of the 7th USENIX Security Symposium*, pages 63–78, San Antonio, TX, January 1998.
- [3] G. Fink and M. Bishop. Property-based testing: A new approach to testing for assurance. *ACM SIGSOFT Software Engineering Notes*, 22(4), July 1997.
- [4] A.K. Ghosh, T. O'Connor, and G. McGraw. An automated approach for identifying potential vulnerabilities in software. In *Proceedings of the 1998 IEEE Symposium on Security and Privacy*, pages 104–114, Oakland, CA, May 3-6 1998.
- [5] A.K. Ghosh, M. Schmid, and V. Shah. Testing the robustness of Windows NT software. In *Proceedings of the Ninth International Symposium on Software Reliability Engineering (IS-SRE'98)*, pages 231–235, Los Alamitos, CA, November 1998. IEEE Computer Society.
- [6] R. Hamlet. Testing programs to detect malicious faults. In *Proceedings of the IFIP Working Conference on Dependable Computing*, pages 162–169, February 1991.
- [7] C. Ko, G. Fink, and K. Levitt. Automated detection of vulnerabilities in privileged programs by execution monitoring. In *10th Annual Computer Security Application Conference*, pages 134–144, December 1994. Orlando, FL.
- [8] R. Lo, K. Levitt, and R. Olsson. MCF: A malicious code filter. *Computers and Security*, 14(6):541–566, 1995.
- [9] M. Schmid and F. Hill. Data generation techniques for automated software robustness testing. In *Proceedings of the International Conference on Testing Computer Software*, 1999. To appear.
- [10] G. Slabodkin. Software glitches leave navy smart ship dead in the water, July 13 1998. Available online: [www.gcn.com/gcn/1998/July13/cov2.htm](http://www.gcn.com/gcn/1998/July13/cov2.htm).
- [11] J. Voas, A. Ghosh, G. McGraw, F. Charron, and K. Miller. Defining an adaptive software security metric from a dynamic software failure tolerance measure. In *Proceedings of the 11th Annual Conference on Computer Assurance*, pages 250–263, June 1996.
- [12] J.M. Voas and G. McGraw. *Software Fault Injection: Inoculating Programs Against Errors*. John Wiley and Sons, New York, 1998.