Testing for Software Vulnerability Using Environment Perturbation [*][†]

WENLIANG DU

Center for Systems Assurance
Department of Electrical Engineering and Computer Science
Syracuse University, 121 Link Hall, Syracuse, NY 13244
Email: wedu@ecs.syr.edu. Tel: (315)443-9180.

ADITYA P. MATHUR

Department of Computer Sciences
Purdue University, West Lafayette, IN 47907

Abstract

We describe an methodology for testing a software system for possible security flaws. Based on the observation that most security flaws are caused by the program's inappropriate interactions with the environment, and triggered by user's malicious perturbation on the environment (which we call an environment fault), we view the security testing problem as the problem of testing for the fault-tolerance properties of a software system. We consider each environment perturbation as a fault and the resulting security compromise a failure in the toleration of such faults. Our approach is based on the well known technique of fault-injection. Environment faults are injected into the system under test and system behavior observed. The failure to tolerate faults is an indicator of a potential security flaw in the system. An Environment-Application Interaction (EAI) fault model is proposed which guides us to decide what faults to inject. Based on EAI, we have developed a security testing methodology, and apply it to several applications. We successfully identified a number of vulnerabilities including vulnerabilities in Windows NT operating system.

Keywords: Security testing, security flaws, fault injection, environment perturbation.

1 INTRODUCTION

Security testing

Reports of security violations due to software errors are becoming increasingly common. We refer to such errors as "security errors" or "security flaws." This has resulted in security related concerns among software developers and users. All stages of software development are motivated by the desire to make the product secure and invulnerable to malicious intentions of some users. Our work is concerned with the testing of software with the goal of detecting errors that might lead to security violations.

Traditional methods for detecting security flaws include penetration analysis and formal verification of security kernels [18, 20]. Penetration analysis relies on known security flaws in software systems other than the one being tested. A team of individuals is given the responsibility of penetrating the system using this knowledge. Formal methods use a mathematical description of the security requirements and that of the system that implements the requirements. The goal of these methods is to show formally that the requirements are indeed met by the system.

A weakness of penetration analysis is that it requires one either to know or be able to postulate the nature of flaws that might exist in a system. Further, the effectiveness of penetration analysis is believed to be as good as that of the team that performs the analysis. A lack of an objective criterion to measure the adequacy of penetration analysis leads to uncertainty in the reliability of the software system for which penetration analysis did not reveal any security flaws.

Attractive due to the precision they provide, formal

methods suffer from the inherent difficulty in specifying the requirements, the system, and then applying the process of checking the requirements specification against system specification.

Recently, several specific security testing techniques have been developed [4, 8, 19, 24, 22, 29]. Adaptive Vulnerability Analysis (AVA) is designed by Ghosh et al. to quantitatively assess information system security and survivability. This approach exercises software in source-code form by simulating incoming malicious and non-malicious attacks that fall under various threat classes [22, 23, 28, 29]. Fuzz [9, 24] is a black-box testing method designed by Miller et al, which feeds randomly generated input stream to system utilities in order to test how reliable they are. Bishop and Dilger studied one class of the time-of-check-to-time-of-use (TOCTTOU) flaws [4], and investigated using static analysis method to identify such type of flaws. Fink and Levitt employ application-slicing technique to test privileged applications [8]; Gligor has also proposed a security testing method using control synthesis graphs [19]. As discussed in section 5, these techniques are either restricted to some specific security flaws or limited by the underlying testing techniques.

Another alternative for security testing is to use general testing techniques, such as path testing, data-flow testing, domain testing, and syntax testing [2]. However, the effectiveness of these techniques in revealing security flaws is still unknown and more studies are needed to justify their use in testing for security flaws.

Outline of our approach

Our approach for security testing employs a well known technique in the testing of fault-tolerant systems, namely fault injection. This approach has drawn upon years of research and experience in vulnerability analysis [1, 3, 6, 17, 21]. Our approach relies on an empirically supported belief that the environment plays a significant role in triggering security flaws that lead to security violations [10, 17].

The problem

For the purpose of our discussion, we assume that a "system" is composed of an "application" and its "environment." Thus, potentially, all codes that are not considered as belonging to the application belong to the environment. However, we can reduce the size of the environment, by considering only those portions of the code that have a direct or indirect coupling with the application code. Such coupling might arise, for example, due to the application's use of global variables declared in the environment or the use of common resources such as files and network elements.

For various reasons, programmers tend to make assumptions about the environment in which their application will function. When these assumptions hold, the application is likely to behave appropriately. But, because the environment, as a shared resource, can often be perturbed by other subjects, especially malicious users, these assumptions might not be true. How to know whether a program can tolerate the environment perturbation is the key problem that we want to solve in this paper.

If we consider environment perturbations, especially malicious perturbations to be (malicious) faults, the above problem is considered as whether a program is able to tolerate various environment faults (not leading to security violations is considered toleration of such faults). In the remainder of this paper, we will use the terms "environment perturbation" and "environment fault" interchangeably where there is no confusion.

Fault injection–the deliberate insertion of faults into an operational system to determine its response–offers an effective solution to validate the dependability of fault-tolerant computer and software systems [5]. In our approach, faults are injected into environment thereby perturbing it. In other words, we perturb the application environment during testing to see how the program responds and whether there will be a security violation under this perturbation. If not then the system is considered secure.

Advantages of our approach

The use of environment fault injection technique leads to several advantages. First, in practice, it is hard to trigger certain anomalies in the environment, and knowing how to trigger them depends on the tester's knowledge of the environment. Therefore, testing software security under those environment anomalies becomes difficult. Fault injection technique provides a way of emulating the environment anomalies without having to be concerned with how they could occur in practice. Second, our approach provides a systematic way of deciding when to emulate environment faults. If we want to test whether a system will behave appropriately under certain environment anomalies, we need to set up those environments. However, the set up time is often difficult to control. If the set-up is too early, the condition might change during the test, and the environment state might not be what is expected when an interaction between the application and the environment takes place. If the environment is set up too late, the effect it has on the application's behav-

ior might not serve the purpose for which it was set up. By exploiting static information in the application and the environment's source code, our approach can, however, decide deterministically when to trigger environment faults. Third, unlike penetration analysis, where the procedure is difficult to automate and quantify, fault injection technique provides a capability of automating the testing procedure. In addition, we adopt a two-dimensional metrics to quantify the quality of our testing procedure.

Research issues

Fault injection requires the selection of a fault model [5]. The choice of this model depends on the nature of faults. Software errors arising from hardware faults, for instance, are often modeled via bits of zeroes and ones written into a data structure or a portion of the memory [15, 26], while protocol implementation errors arising from communication are often modeled via message dropping, duplication, reordering, delaying etc. [14]. Understanding the nature of security faults provides a basis for the application of fault injection. Several studies have been concerned with the nature of security faults [1, 3, 6, 17, 21].) However, we are not aware of any study that classifies security flaws from the point of view of environment perturbation. Some general fault models have also been widely used [13, 27, 22, 29]. The semantic gap between these models and the environment faults that lead to security violations is wide and the relationship between faults injected and faults leading to security violations is still unknown. We have developed an Environment-Application Interaction (EAI) fault model which serves as the basis for the fault injection technique described here. The advantage of the EAI model is in its capability of emulating environment faults that are likely to cause security violations.

Another issue in fault injection technique is the location, within the system under test, where faults are to be injected. In the current stage of our research, we inject environment faults at the points where the environment and the application interact. In future work, we plan to exploit static analysis to further reduce the number of fault injection locations by finding the equivalence relationship among those locations. The motivation for using static analysis method is that we can reduce the testing efforts by utilizing static information from the program.

A general issue about software testing is "what is an acceptable test adequacy criterion?" [11]. We adopt a two-dimensional coverage metric (code coverage and fault coverage) to measure test adequacy.

The remainder of this paper is organized as follows: section 2 presents the fault model. A methodology for security testing is presented in section 3. In section 4 we will show the results of using this methodology in detecting real world programs. Finally a brief overview of related studies is presented in section 5 followed by summary of this research and the potential for future work in section 6.

2 AN ENVIRONMENT FAULT MODEL

In order to determine system behavior under various environment conditions, an engineer must be able to determine the effects of environment perturbation on a given system. Therefore, it is useful to inject faults that manifest themselves as errors in systems at the environment-application interaction level. To maintain confidence in the validity of the errors, the model used for these injections should be drawn from actual environment faults, while faults injected into the system should be able to emulate those environment faults appropriately. One assumption behind this requirement is that a security violation resulting due to the injected fault is similar to one that results due to an environment fault that arises during the intended use of the system.

2.1 Terminology

Definition 2.1. *(Internal State and Internal Entity)* Any element in an application's code and data space is considered an internal entity. A state consisting of the status of these entities is called an internal state.

Variable i in a application, for example, is an internal entity. The value of i is a part of an internal state. The size of a buffer used in the application is also a part of its internal state. In general, all information in this application's data space, stack space, and heap space are part of its internal state.

Definition 2.2. *(Environment Entity and Environment State)* Any element that is external to an application's code and data space is called an environment entity. A state that consists of the status of these entities is called an environment state.

For instance, file and network are treated as environment entities. The permission of a file, existence of a file, ownership of a file, real user-id of a process, and the effective user-id of process are different parts of an environment state.

It should be noted that the distinction between the internal entity and environment entity always depends

on the scope of the program. If the program under test is the whole operating system, then many usual environment entities, such as files, belong to the internal entities; however, if the program under test is an application of a component of the operating system, entities like files now become environment entities. In this study, we focus on testing applications and components of an operating system.

A key difference between an environment and an internal entity, which makes implementation of a secure system difficult and error-prone, is the shared nature of the environment entity. An application is not the only one that can access and change an environment entity. Other objects, such as other users, may access and change the environment entity as well. Internal entity, on the other hand, is private to an application in the sense that only the application can modify and access them, assuming that the underlying operating system provides protected process space.

In concurrent programming, shared resources are handled by using the mutual exclusion and the semaphore mechanism to guarantee assumptions about the state of shared resources. However, we believe that few programmers use a similar mechanism to guarantee their assumption about the state of the environment. There are several reasons for this. First, programmers might not have recognized that the environment entities are shared resources. When, for example, an application writes to a file, it checks that it has the permission to write to that file, and then assumes that right in subsequent operations to that file without noticing that a malicious attacker could have change the environment thereby rendering the assumption false. Most security flaws resulting from race conditions [4] are caused by such dubious assumptions. Second, although some mechanisms, such as file locking, guarantee that a programmer's assumption hold on some part of the environment state, there is no general mechanism to do the same because the environment entity has various attributes and the mutual exclusion and semaphore mechanisms could not handle them easily. As a result, programmers often use *ad hoc* mechanisms to guarantee the correctness of their assumptions. This can lead to errors more readily than would be the case when a standard mechanism is used.

2.2 Developing a fault model

In order to provide high confidence in the validity of the security flaws caused by environment faults, the methodology described here models systems at a high level. We refer to this level as the Environment-Application Interaction (EAI) level. Fault injection

at the interaction level attempts to emulate what a "real" attacker does. Since most of the vulnerability databases record the way attackers exploit a vulnerability, we transform these exploits to environment faults to be injected with little analysis on those records thereby narrowing the semantic gap between faults injected at the interaction level and faults that really occur during the intended use of the system. In contrast, other studies [22, 29] inject faults at the program statement level thereby leaving a large semantic gap between faults injected and those that might arise during the intended use of the application.

2.3 An EAI fault model

In general, environment faults affect an application in two different ways. First, an application receives inputs from its environment. The environment faults now become faults in the input, which is then inherited by an internal entity of the application. From this point onwards the environment faults propagate through the application via the internal entities. If the application does not handle the faults correctly, a security violation might occur. The direct reason for this violation appear to be faults in the internal entity. However, this violation is due to the propagation of environment faults. Stated differently, the environment indirectly causes a security violation, through the medium of the internal entity. Figure 1(a) shows this indirect way in which the environment faults affect an application.

Consider the following example. Suppose that an application receives its input from the network. Any fault in the network message related to this input is inherited by an internal entity. When the application does a memory copy from this message to an internal buffer without checking the buffer's boundaries, the fault in the network message, the fault being "message too long," now triggers a violation of security policy.

A second way in which an environment fault affects the application is that the fault does not propagate via the internal entity; instead, it stays within the environment entity and when the application interacts with the environment without correctly dealing with these faults, security policy will be violated. In this case, the environment faults are the direct cause of security violation and the medium for environment faults is the environment entity itself. Figure 1(b) shows this direct way in which the environment faults affect an application.

Let us now consider an example to illustrate this second kind of interaction. Suppose that an application needs to execute a file. There are two possibilities: one is that the file belongs to the user who runs
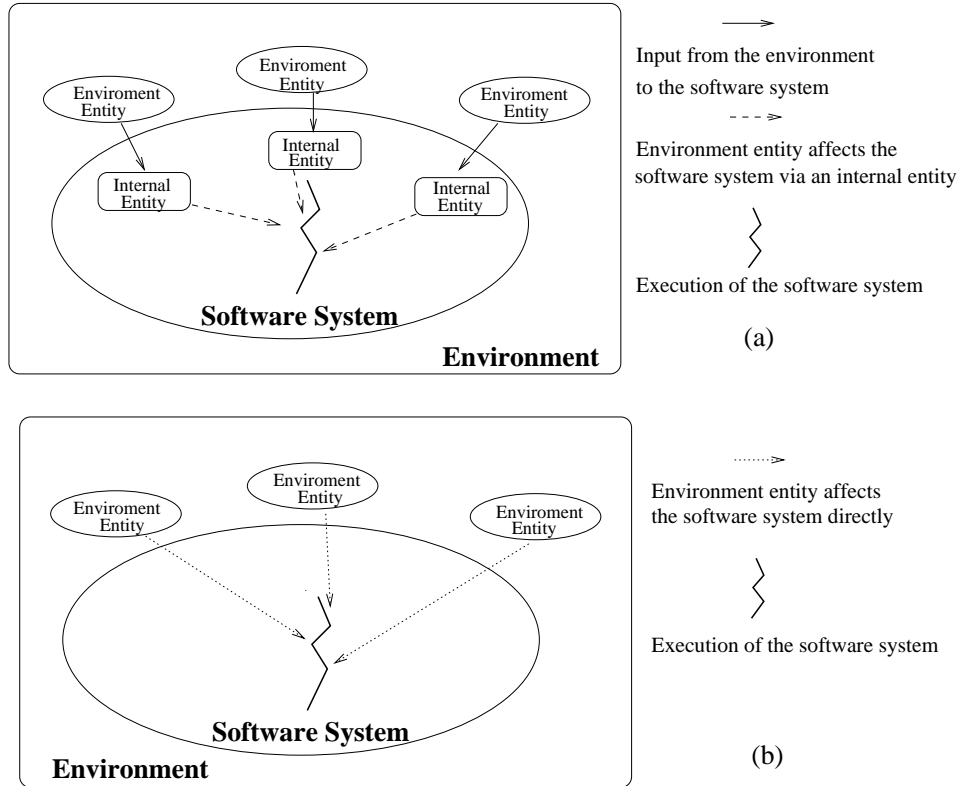
Figure 1. Interaction Model

the application. Here the environment attribute is the file's ownership. In this case the execution is safe. The other possibility is that the file belongs to some malicious user. This is an environment fault created by the malicious user. Now the individual who runs the application assumes that the file belongs to the application. If the application does not deal with this environment fault, it might execute arbitrary commands in that file thereby resulting in a security violation.

The most error-prone interactions between an application and the environment are those that involving files. Programmers tend to use an abstraction of a file that includes only a subset of the file attributes. A file name with a location or file content, for example, is a commonly used abstraction of a file. The environment faults, such as a long file name or a file name with special characters, associated with this abstraction will propagate via the internal entity. If the application does not place appropriate checks on these internal entities, such environment faults will cause security violations such as those due to buffer overflow and the execution of an unintended command. The environment faults associated with the remaining file attributes, such as whether the file is a symbolic link, the ownership of the file, existence of the file, and the

permissions associated with the file, will not propagate via an internal entity. Although these attributes are extrinsic to the application, if not dealt correctly, they are likely to directly affect the interaction between application and environment.

In summary, we have categorized the environment faults according to the way they affect applications. Environment faults which affect programs via internal entities are called *indirect environment faults*. Environment faults which affect programs via environment entities are called *direct environment faults*.

Indirect environment faults are further divided into five sub-categories according to their origins: (1) user input, (2) environment variable, (3) file system input, (4) network input, (5) process input. These five categories are the primary sources where a program get its inputs from.

A number of environment entities can affect a program's behavior directly. We categorize them according to traditional operating system categorization, namely, we further classify direct environment faults into three sub-categories: (1) file system, (2) process, (3) network. Although there are other type of environment entities in an operating system, we decide not to adopt them in our direct environment fault model

5

because we do not see significant amount of vulnerabilities caused by those types of environment entities. However, we could easily extend our model to accommodate future new type of vulnerabilities.

3 ENVIRONMENT FAULT INJECTION METHODOLOGY

3.1 Fault injection

Like the EAI model, which models the environment faults at the interaction level, fault injections are also done at the interaction level. The previous section classifies the environment faults into direct and indirect environment faults. These faults are injected using the following mechanisms:

1. **Indirect Environment Fault Injections:** An indirect environment fault occurs at the interaction point where an application requests its environment for an input. The input that the environment provides to the application will most likely affect the application's behavior. A secure application should tolerate an unexpected anomaly in the environment input. One way to perturb the input is to use random input as in Fuzz [9, 24]. However, this approach dramatically increases the testing space, which and calls for a significantly large amount of testing effort. The Fuzz approach does not exploit the semantics of each input. Our vulnerability analysis, however, has shown that inputs most likely to cause security violations tend to have patterns according to their semantics. If, for instance, the input is a list of paths used to search for a command, then security failure will most likely occur when the order of these paths is altered, a new path is inserted or deleted, or the length of the list is increased. Other kinds of perturbations are less likely to cause security failure. Thus, by an examination of rare cases and by concentrating instead on fault patterns already observed, we reduce the testing space considerably.

   Faults injected into the application are based on patterns that are likely to cause security faults. These patterns come from our investigation of a vulnerability database and other studies reported in the literature. The faults are summarized in Table 1.

2. **Direct Environment Faults Injections:** A direct environment fault occurs at the interaction point where the application accesses an environment entity for creation, modification, reading or execution of an environment entity. Different status of environment entity attributes will affect the consequences of those interactions. Thus, the environment fault injections are used to perturb the attributes of an environment entity at points of interaction to see how the application responds to the perturbation. For example, before an application executes an `open` operation to a named `file`, several perturbations are performed on this file by changing its attributes such as its existence, permissions, ownership, and the type of the file since failure to handle these attributes is likely to cause security violations. These attributes are and their their perturbation are presented in Table 2.

3.2 Test adequacy criterion

An important issue in the management of software testing is to "ensure that prior to the start of testing the objectives of testing are known and agreed upon and that the objectives are set in terms that can be measured." Such objectives "should be quantified, reasonable, and achievable" [12].

We use *fault coverage* and *interaction coverage* measure test adequacy. Fault coverage is defined as the percentage of the number of faults tolerated with respect to that of the faults injected. Our conjecture is that the higher the fault coverage the more secure the application is. In addition to fault coverage, an additional measurement of the testing effort is the interaction coverage. Interaction coverage is defined as the percentage of the number of interaction points where we injected faults with respect to the total number of interaction points. Once again, we conjecture that the higher the interaction coverage, the more dependable the testing result are. Of course we assume that faults found during testing are removed. These two coverage criteria lead to a 2-dimensional metric for measuring test adequacy.

Figure 2 shows the 2-dimensional metric and four sample points of significance. The metric serves as a quantitative evaluation of a test set. Point 1 is representative of the region where testing resulted in low interaction and fault coverage. In this case testing is considered inadequate. Point 2 is representative of the region where the fault coverage is high but interaction coverage is low. The test is considered inadequate since in this test, only a few interactions are perturbed, how the system behaves under perturbation of other interactions is still unknown.

Point 3 is representative of an insecure region because the fault coverage is so low that we consider the application is likely to be vulnerable to the perturba-

Table 1. Indirect Environment Faults and Environment Perturbations

| Internal Entity	Semantic Attribute	Fault Injections
User Input	file name + directory name	change length, use relative path, use absolute path, insert special characters such as "..", "/" in the name
	command	change length, use relative path, use absolute path, insert special characters such as "|", "&", ">" or newline in the command
Environment Variable	file name + directory name	change length, use relative path, use absolute path, use special characters, such as "|", "&" or ">" in the name
	execution path + library path	change length, rearrange order of path, insert a untrusted path, use incorrect path, use recursive path
	permission mask	change mask to 0 so it will not mask any permission bit
File System Input	file name + directory name	change length, use relative path, use absolute path, use special characters in the name such as "|", "&" or ">" in name
	file extension	change to other file extensions like ".exe" in Windows system; change length of file extension
Network Input	IP address	change length of the address, use bad-formatted address
	packet	change size of the packet, use bad-formatted packet
	host name	change length of host name, use bad-formatted host name
	DNS reply	change length of the DNS reply, use bad-formatted reply
Process Input	message	change length of the message, use bad-formatted message

Figure 2. Test Adequacy Metric

tion of the environment. The safest region is indicated by point 4 which corresponds to a high interaction and fault coverage.

3.3 Procedure

The procedure of our Environment Fault Injection Methodology consists of the following steps:

1. Set **count** and **n** to 0.

2. For each test case, do step 3 to 9.

3. For each interaction point in the execution trace, decide if the application asks for an input. If there is no input, only inject direct environment faults; if there is an input, inject both direct and indirect environment faults.

4. Decide the object where faults will be injected.

5. Establish a fault list corresponding to this object using Table1 and Table 2.

6. For each fault in the list, inject it before the interaction point for the direct environment faults; inject each fault after the interaction point for the indirect environment faults since in this case, we want to change the value the internal entity receives from the input.

7. Increase **n** by 1.

8. Detect if security policy is violated. If so, increase **count** by 1.

9. Calculate interaction coverage. If the test adequacy criteria for interaction coverage is satisfied then stop else repeat steps 3-9 until the adequacy criteria for interaction coverage is achieved.

10. Divide **count** by **n** yielding α to obtain the vulnerability assessment score (fault coverage) for the application.

3.4 Multiple Perturbations

Although in theory, a combination of k multiple perturbations can cause security breaches while each single perturbation cannot, our experience shows that this theoretic situation is rarely usual in practice. Therefore our testing procedure only deals with a single perturbation. To deal with multiple perturbations, many challenging issues still exist. For example, an important issue is how to form a multiple perturbation scenario? We cannot afford to use an arbitrary combination of all the perturbation candidates. We will continue evaluate the validity of our experience about the multiple perturbations, and if necessary, to develop a procedure that can handle multiple perturbations.

3.5 Example

To illustrate the steps shown above, we consider an example of fault injection. The following code is taken from `BSD` version of `lpr.c`. Notice that `lpr` is a privileged application. It is a `set-UID` application which means that it runs in the root's privilege even when it is invoked by a user who does not have the same privilege as the root.

```
f = create(n, 0660);
if (f<0) {
    printf(''%s: cannot create %s'', name, n);
    cleanup();
}
... (code skipped here)
if (write(f, buf, i)!=i) {
    printf(''%s: %s: temp file write error\n'',
            name, n);
    break;
}
```

Suppose that we have decided to perturb the environment at a place where the `create` system call is issued. This is an interaction point where `lpr` interacts with the `file system`. There is no input in this case and hence we simply carry out direct environment fault injections.

The next step is to identify the object. Here, `n` is a file name, and hence the object is the file referred to using this file name. Then we refer to Table 2 and get a list of attributes that need to be perturbed. This list includes 1) file existence, 2) file ownership, 3) file permission, 4) symbolic link, 5) file content invariance,

6) file name invariance and 7) working directory. A further analysis shows that attributes 5 and 6 are not applicable in this case as this is supposed to be the first time the file is encountered.

We then perturb the remaining four attributes of the file and inject the faults into the application. For example, the perturbation of the "existence" means that we make the file exist or not exist before the application creates it. The perturbation of "symbolic link" means that we make the file link to some other file, such as the password file, before the application creates it.

After fault injection, we execute the application and detect if there is any violation of the security policy. In this case the violation is detected when we perturb attributes 1, 2, 3 and 4. Doing so causes `lpr` to write to a file even when the user who runs it does not have the appropriate ownership and file permissions. Thus when the file is linked to the password file, the password file is be modified by `lpr`. The problem here is that the application assumes that the file does not exist before the creation or assumes that the file belongs to the user who runs the application. In a real environment, this assumption could easily be false and the fault injection test points out a security vulnerability.

4 RESULT

4.1 Turnin

`Turnin` is a program used in Purdue for electronically submitting files for grading. Before students in a class can use this program, the teaching assistant (TA) for this class should have set up his account (or a dedicated course account) correspondingly. This includes creating a `submit` directory under the home directory of this account, creating a `Projlist` file under `submit` directory, which specifies a list of projects students could be able to turnin. Students can type "turnin -c coursename -l" to view the list of projects; students can type "turnin -c coursename -p projectname files" to turnin their project files. After submission, the submitted files will be copied to TA's `submit` directory.

Since `turnin` program allows students to copy their files to TA's protected directory, the program is running as `SUID`, which means its effective user is `root`. The program consists of 1310 lines of code.

Following our method, we have identified 8 interaction places where programmers could possibly have made assumptions about the environment. We make 41 environment perturbation to check whether programmers indeed made the assumptions, and whether the failure of these assumptions can affect program's security. Among those perturbations, 9 perturbation lead

Table 2. Direct Environment Faults and Environment Perturbations

Environment Entity	Attribute	Fault Injections
File System	file existence	delete an existing file or make a non-existing file exist
	file ownership	change ownership to the owner of the process, other normal users, or root
	file permission	flip the permission bit
	symbolic link	if the file is a symbolic link, change the target it links to; if the file is not a symbolic link, change it to a symbolic link
	file content invariance	modify file
	file name invariance	change file name
	working directory	start application in different directory
Network	message authenticity	make the message come from other network entity instead of where it is expected to come from
	protocol	purposely violates underlying protocol by omitting a protocol step, adding an extra step, reordering steps
	socket	share the socket with another process
	service availability	deny the service that application is asking for
	entity trustability	change the entity with which the application interacts to a untrusted one
Process	message authenticity	make the message come from other process instead of where it is expected to come from
	process trustability	change the entity with which the application interacts to a untrusted one
	service availability	deny the service that application is asking for

to security violation, which means the failure of assumptions on these 9 situation could lead to a vulnerability in the program. Then we investigated each assumptions by asking whether they are reasonable. For example, programmers obviously made an assumption that `/usr/local/lib/turnin.cf` file is trusted. Our perturbation testing found out that if this assumptions is false, the system's security will be violated. Since the `turnin.cf` will always be protected, so is its directory, we believe the assumption is quite reasonable, there is no vulnerability regarding to this assumption.

However, one assumption seems unreasonable to us, it turns out to be a vulnerability, and is hence exploited by us after we have known the assumption. The problematic code is list in the following:

```
if ((FILE *)0 == (fp = fopen(pcFile, "r"))) {
    printf("can not find project list file\n");
    exit(9);
}
```

Since `fopen` is an interaction point where potential assumption might be made, we perturb the environment status of `pcFile`, making it only readable by root, not by the people who is running the `turnin` program. The result is that by running "turnin -c coursename -l", we can successfully read the contents of the file that we are not supposed to be able to read. So, here the programmers have made an assumption that people are allowed to read file pointed by `pcFile` using `turnin` program, and its failure can cause security violation. Now, the question is: is this assumption reasonable? The result turns out to be NO since TA can make `pcFile` point to any file he wants, then using `turnin` program to read the contents of that file.

Knowing this fact, we designed a following scenario: a TA makes the `Projlist` a symbolic link to `/etc/shadow`, which is not readable by anyone except `root`. Then the TA runs "turnin -c coursename -l", Voila, the program prints out the content of `/etc/shadow`!

Another perturbation we have done is perturbing the attributes of the argument in the following code:

```
execve (acTar, nargv, environ);
```

Since *nargv* contains file names, according to table 1, we have inserted special characters, such as "/", "../", in front of the file names. The program does a good job in forbidding the "/" character, however, it does not resist the perturbation of inserting "../" in the front. Knowing this fact, a student can submit several ".login" files with different number of "../" in front of the ".login" file, such that when his TA unpacks the submitted file, the TA's ".login" will be overwritten by the

student's malicious ".login" file, which can do anything evil to the TA.

The `turnin` program has been used in Purdue University widely since 1993, and we became the first to identify these vulnerabilities. After our discovery, the university quickly verified the problem and patched the `turnin` program.

4.2 Windows NT Registry

In Windows NT operation system, registry directory is a critical part to the system security. Registry directory is essentially an organized storage for operating system's and application's data which are globally shared by different applications and components. An appropriate configuration on each registry key in the registry directory is a key factor for security. Many security vulnerabilities have been reported due to inappropriate configuration of the registry keys. In the Windows NT 4.0 (SP3), there are still keys that are not protected. Our task is to test the related modules of the operating system, and find whether it is secure to leave those registry keys unprotected.

First of all, we use static analysis technique to find out where these unprotected keys are used [7], then we apply the EPA method to find if programmers have made assumptions that can fail.

We have identified 9 unprotected registry keys that could be exploited to break the system security, and indeed we came up with test cases to actually exploit the vulnerabilities. Furthermore, based on the similarities of these 9 registry keys and other 20 unprotected keys, we speculate that the same vulnerabilities exist for those 20 keys as well. However, we have not been able to perturb the modules that used the other 20 keys yet due to the lack of knowledge of how those modules work. The 9 registry keys that we have exploited are the results of applying our perturbation technique.

Due to the agreement with Microsoft, we are not revealing the exact keys and source codes that have the vulnerabilities. So, in the next discussion, we will not refer to any specific key, except the purpose of the key and the problem with the key.

One of the keys in the registry directory specifies a name for a font file. It seems pretty safe to give everybody the right to modify this registry key until we have found a module in the system that invokes a function call to actually delete this file. To know whether the program has done the correct checking before the delete or not, we did a perturbation on the properties of this file according to Table 2, making it writable only by administrator, and also making it point to a very important file (such as system configuration file)

instead of just a font file. It turns out that the program fails to respond securely under this environment perturbation - when administrators run this module, they will actually delete the file specified by this registry key regardless of whether this file is a font file or a security critical file. The assumption behind this "delete" environment interaction is that the programmers assume the file name always points to a font file or a unimportant file, however, since everybody has the right to modify the value of this registry key, the assumption fail to sustain.

Another vulnerability we have found is associated with user logon module. When a user logons, the module will find the user's profile from a directory specified in a registry key. Using our EAI model, we have managed to perturb the trustability attribute of the directory, and found out that the program does not deal with the situation when the directory is not trusted, which means, whenever a user logs in, the logon module will go to the untrusted directory, and grab a specified profile for you. Therefore, by the environment perturbation, we have found out that programmers have made a fatal assumption about the trustability of the profile directory. After knowing the fact, it becomes straightforward to design a test case and fail the programmers' assumptions.

5   COMPARISON WITH RELATED WORK

A significant amount of computer security testing is performed using penetration testing. Security is assessed by attempting to break into an installed system by exploiting well-known vulnerabilities. Several researchers, including Linde and Attanasio [18], Pfleeger [25], describes the process of penetration testing. As pointed out by Pfleeger, the success of penetration testing depends on testers' skill, experience, and familiarity with the system. Moreover, the lack of well defined and tested criteria to decide when to stop penetration testing causes penetration testing difficult to use.

Our research attempts to overcome the above mentioned difficulties. It has a deterministic procedure to conduct and test, a criterion to decide when testing should stop. It overcomes the limitation of the lack of knowledge of the environment by emulating possible attacks using the faults injection technique. Finally, our approach overcomes the limitation of testers' knowledge by offering a set of concrete faults that should be injected into application.

Adaptive Vulnerability Analysis (AVA) is designed by Ghosh et al. to quantitatively assess information system security and survivability. This approach exercises software in source-code form by simulating incoming malicious and non-malicious attacks that fall under various threat classes [22, 23, 28, 29]. In this respect, our own work parallels the AVA approach. A major divergence appears, however, with respect to how incoming attacks are simulated. AVA chooses to perturb the internal state of the executing application by corrupting the flow of data and the internal states assigned to application variables. Our approach chooses to perturb the environment state by changing the attributes of the environment entity and perturbing the input that an application receives from the environment. Our approach should be considered as complementary to AVA.

For attacks that do not affect the internal states of an application, AVA appears incapable of simulating them by only perturbing the internal states. For vulnerabilities that are caused purely by incorrect internal states, our approach cannot simulate them by only perturbing the environment. One disadvantage of the AVA is the semantic gap between the attacks during the use of an application and the perturbation AVA makes during testing. In other words, knowing that the application fails under certain perturbation, it is difficult to derive what kind of attacks correspond to this failure. This makes it difficult to assess the validity of the perturbation. Our approach narrows the semantic gap by perturbing at the environment-application level since most attacks really occur due to intentional perturbation of the environment.

Fuzz [9, 24] is a black-box testing method designed by Miller et al, which feeds randomly generated input stream to system utilities in order to test how reliable they are. The Ballista [16] testing methodology involves automatically generating combinations of exceptional and valid parameter values to be used in calling software modules. Both of these testing methods focus on the system reliability instead of security.

Bishop and Dilger studied one class of the time-of-check-to-time-of-use (TOCTTOU) flaws [4], and investigated using static analysis method to identify such type of flaws. Fink and Levitt employ application-slicing technique to test privileged applications [8]; Gligor has also proposed a security testing method using control synthesis graphs [19]. They both achieve a certain degree of success in security testing.

6   SUMMARY AND FUTURE WORK

We have presented a white-box security testing methodology using environment perturbation technique, a variant of the fault injection technique. The

methodology is based on the Environment-Application Interaction (EAI) model, which captures the properties of a family of software vulnerability. We have applied this methodology to several real-world systems and applications, and we have successfully identified a number of security flaws that exist for several years without being discovered.

Future work will concentrate on applying this methodology to more applications. We are in the progress of developing and conducting a set of experiments to evaluate the effectiveness of this methodology. In the future, we hope to be able to develop a prototype tool for security testing based on this methodology.

References

[1] T. Aslam. A taxonomy of security faults in the unix operation system. Master's thesis, Purdue University, August 1995.

[2] B. Beizer. *Software Testing Techniques*. Van Nostrand Reinhold, New York, 1990.

[3] M. Bishop. A taxonomy of unix system and network vulnerabilities. Technical Report CSE-95-10, Department of Computer Science, University of California at Davis, May 1995.

[4] M. Bishop and M. Dilger. Checking for race conditions in file acesses. *The USENIX Association Computing Systems*, 9(2):131–151, Spring 1996.

[5] J. Clark and D. Pradhan. Fault injection: A method for validating computer-system dependability. *IEEE Computer*, pages 47–56, June 1995.

[6] W. Du and A. Mathur. Categorization of software errors that led to security breaches. In *21st National Information Systems Security Conference*, Crystal City, VA, 1998.

[7] W. Du and A. Mathur. Security relevancy analysis on the registry of windows nt 4.0. In *ACSAC'99 15th Annual Computer Security Applications Conference*, Phoenix, Arizona, December 6-10 1999.

[8] G. Fink and K. Levitt. Property-based testing of privileged programs. In *Proceedings of the 10th Annual Computer Security Applications Conference; Orlando, FL, USA; 1994 Dec 5-9*, 1994.

[9] B. Miller, L. Fredriksen and B. So. An empirical study of the reliability of unix utilities. *Communications of the ACM*, 33(12):32–44, December 1990.

[10] S. Garfinkel and G. Spafford. *Practical UNIX & Internet Security*. O'Reilly & Associates, Inc., 1996.

[11] J. Goodenough and S. Gerhart. Toward a theory of testing: Data selection criteria. *current Trends in Programming Methodology*, 2:44–79, 1977.

[12] H. Zhu, P. Hall and J. May. Software unit test coverage and adequacy. *ACM Computing Surveys*, 29(4):366–427, December 1997.

[13] W. Kao, R. Iyer and D. Tang. FINE: A fault injection and monitoring environment for tracing the unix system behavior under faults. *IEEE Transactions on Software Envineering*, 19(11):1105–1118, November 1993.

[14] S. Dawson, F. Jahanian and T. Mitton. ORCHESTRA: A fault injection environment for distributed systems. In *26th International Symposium on Fault-Tolerant Computing (FTCS)*, pages 404–414, Sendai, Japan, June 1996.

[15] G. Kanawati, N. Kanawati and J. Abraham. FERRARI: A tool for the validation of system dependability properties. In *Proceedings 22nd International Symposium Fault Tolerant Computing*, pages 336–344, July 1992.

[16] N. Kropp, P. Koopman and D. Siewiorek. Automated robustness testing of off-the-shelf software components. In *28th Fault Tolerant Computing Symposium*, June 1998.

[17] I. Krsul. *Software Vulnerability Analysis*. PhD thesis, Purdue University, Department of Computer Sciences, West Lafayette, Indiana, 1998.

[18] R. R. Linde. Operating system penetration. In *AFIPS National Computer Conference*, pages pp. 361–368, 1975.

[19] V. D. Gligor, C. S. Chandersekaran, W. Jiang, A. Johri, G. L. Luchenbaugh and L. E. Reich. A new security testing method and its application to the secure xenix kernel. *IEEE Transactions on Software Engineering*, SE-13(2):169–183, February 1987.

[20] E. J. McCauley and P. J. Drongowski. The design of a secure operating system. In *National Computer Conference*, 1979.

[21] C. E. Landwehr, A. R. Bull, J. P. McDermott and W. S. Choi. A taxonomy of computer program security flaws. *ACM Computing Surveys*, 26(3), September 1994.

[22] A. Ghosh, T. O'Connor, G. McGraw. An automated approach for identifying potential vulnerabilities in software. In *IEEE Symposium on Security and Privacy*, Oakland, CA, 1998.

[23] J. Voas, F. Charron, G. McGraw, K. Miller and M.Friedman. Predicting how badly "good" software can behave. *IEEE Software*, 14(4):73–83, August 1997.

[24] B. Miller, D. Koski, C. Lee, V. Maganty, R. Murthy, A. Natarajan and J. Steidl. Fuzz revisited: A re-examination of the reliability of unix utilities and services. Technical report, Computer Sciences Department, University of Wisconsin, 1995.

[25] C. Pfleeger, S. Pfleeger and M. Theofanos. A methodology for penetration testing. *Computers and Security*, 8(7):613–620, 1989.

[26] S. Han, K. Shin and H. Rosenberg. Doctor: An integrated software fault injection environment for distributed real-time systems. Technical report, University of Michigan, Department of Elect. Engr. and Computer Science, 1995.

[27] M. Hsueh, T. Tsai and R. Iyer. Fault injection techniques and tools. *IEEE Computer*, pages 75–82, April 1997.

[28] J. Voas. Testing software for characteristics other than correctness: Safety, failure tolerance, and security. In *Proc. of the Int'l Conference on Testing Computer Software*, 1996.

[29] J. Voas and G. McGraw. *Software Fault Injection: Incoculating Programs Against Errors*. John Wiley & Sons, Inc., 1998.

Author's Biographies:

Wenliang Du is an Assistant Professor of Electrical Engineering and Computer Science at Syracuse University. He holds a PhD from the Computer Science Department at Purdue University, West Lafayette, Indiana. While he studied in Purdue, he worked in the Center for Education and Research in Information Assurance and Security (CERIAS). His research areas include computer and information security, software security testing, secure multiparty computation and applied cryptography.

Aditya P. Mathur is a Professor of Computer Science at Purdue University. He conducts research in the areas of software testing, reliability, formal approaches for software process control, and the management of smart spaces. Mathur has been a crusader for the use of code coverage criteria in the estimation of software reliability or as an orthogonal metric to assess confidence in the reliability estimates. He has proposed the "Saturation Effect" as a motivating device for the use of a sequence of testing techniques. This device is often used by vendors to enhance marketing of their test tools. Mathur's recent work deals with test and management of Internet services. His research group has developed homeWabash, a tool for the management of networked devices. In collaboration with Raymond DeCarlo, Mathur is also investigating the use of feedback control theory for the control of software development processes.