

Stelios Sidiroglou · Angelos D. Keromytis

Execution transactions for defending against software failures: use and evaluation

Published online: 8 March 2006
© Springer-Verlag 2006

Abstract We examine the problem of containing buffer overflow attacks in a *safe* and *efficient* manner. Briefly, we automatically augment source code to dynamically catch stack and heap-based buffer overflow and underflow attacks, and recover from them by allowing the program to continue execution. Our hypothesis is that we can treat each code function as a transaction that can be aborted when an attack is detected, without affecting the application's ability to correctly execute. Our approach allows us to enable selectively or disable components of this defensive mechanism in response to external events, allowing for a direct tradeoff between security and performance. We combine our defensive mechanism with a honeypot-like configuration to detect previously unknown attacks, automatically adapt an application's defensive posture at a negligible performance cost, and help determine worm signatures. Our scheme provides low impact on application performance, the ability to respond to attacks without human intervention, the capacity to handle previously unknown vulnerabilities, and the preservation of service availability. We implement a stand-alone tool, DYBOC, which we use to instrument a number of vulnerable applications. Our performance benchmarks indicate a slow-down of 20% for Apache in full-protection mode, and 1.2% with selective protection. We provide preliminary evidence toward the validity of our transactional hypothesis via two experiments: first, by applying our scheme to 17 vulnerable applications, successfully fixing 14 of them; second, by examining the behavior of Apache when each of 154 potentially vulnerable routines are made to fail, resulting in correct behavior in 139 cases (90%), with similar results for *sshd* (89%) and *Bind* (88%).

Keywords Buffer overflows · Exception · Transactional Execution

S. Sidiroglou (✉) · A. D. Keromytis
Department of Computer Science, Columbia University,
1214 Amsterdam Avenue, 450 Computer Science Building,
New York, NY 10027, USA
E-mail: {stelios, ig2111, angelos}@cs.columbia.edu

1 Introduction

The prevalence of buffer overflow attacks as a preferred intrusion mechanism, accounting for approximately half the CERT advisories in the past few years [89], has elevated them into a first-order security concern. Such attacks exploit software vulnerabilities related to input (and input length) validation, and allow attackers to inject code of their choice into an already running program. The ability to launch such attacks over a network has resulted in their use by a number of highly publicized computer worms.

In their original form [2], such attacks seek to overflow a buffer in the program stack and cause control to be transferred to the injected code. Similar attacks overflow buffers in the program heap, virtual functions and handlers [4, 62], or use other injection vectors such as format strings. Due to the impact of these attacks, a variety of techniques for removing, containing, or mitigating buffer overflows have been developed over the years. Although bug elimination during development is the most desirable solution, this is a difficult problem with only partial solutions. These techniques suffer from at least one of the following problems:

- There is a poor tradeoff between security and availability: once an attack has been detected, the only option available is to terminate program execution [22, 31], since the stack has already been overwritten. Although this is arguably better than allowing arbitrary code to execute, program termination is not always a desirable alternative (particularly for critical services). *Automated, high-volume attacks, e.g., a worm outbreak, can exacerbate the problem by suppressing a server that is safe from infection but is being constantly probed and thus crashes.*
- Severe impact in the performance of the protected application: dynamic techniques that seek to detect and avoid buffer overflow attacks during program execution by instrumenting memory accesses, the performance degradation can be significant. Hardware features such as the

NoExecute (NX) flag in recent Pentium-class processors [31] address the performance issue, but cover a subset of exploitation methods (e.g., jump-into-libc attacks remain possible).

- Ease of use: especially as it applies to translating applications to a safe language such as Java or using a new library that implements safe versions of commonly abused routines.

An ideal solution uses a comprehensive, perhaps “expensive” protection mechanism only where needed and allows applications to gracefully recover from such attacks, in conjunction with a low-impact protection mechanism that prevents intrusions at the expense of service disruption.

1.1 Our contribution

We have developed such a mechanism that automatically instruments all statically and dynamically allocated buffers in an application so that any buffer overflow or underflow attack will cause transfer of the execution flow to a specified location in the code, from which the application can resume execution. *Our hypothesis is that function calls can be treated as transactions that can be aborted when a buffer overflow is detected, without impacting the application’s ability to execute correctly.* Nested function calls are treated as subtransactions, whose failure is handled independently. Our mechanism takes advantage of standard memory-protection features available in all modern operating systems and is highly portable. We implement our scheme as a stand-alone tool, named DYBOC (DYnamic Buffer Overflow Containment), which simply needs to be run against the source code of the target application. Previous research [70, 71] has applied a similar idea in the context of a safe language runtime (Java); we extend and modify that approach for use with unsafe languages, focusing on single-threaded applications. Because we instrument memory regions and not accesses to these, our approach does not run into any problems with pointer aliasing, as is common with static analysis and some dynamic code instrumentation techniques.

We apply DYBOC to 17 open-source applications with known buffer overflow exploits, correctly mitigating the effects of these attacks (allowing the program to continue execution without any harmful side effects) for 14 of the applications. In the remaining three cases, the program terminated; in no case did the attack succeed. Although a contrived micro-benchmark shows a performance degradation of up to 440%, measuring the ability of an instrumented instance of the Apache web server indicates a performance penalty of only 20%. We provide some preliminary experimental validation of our hypothesis on the recovery of execution transactions by examining its effects on program execution on the Apache web server. We show that when each of the 154 potentially vulnerable routines are forced to fail, 139 result in correct behavior, with similar results for *sshd*

and *Bind*. Our approach can also protect against heap overflows.

Although we believe this performance penalty (as the price for security and service availability) to be generally acceptable, we provide further extensions to our scheme to protect only against specific exploits that are detected dynamically. This approach lends itself well to defending against scanning worms. Briefly, we use an instrumented version of the application (e.g., web server) in a sandboxed environment, with all protection checks enabled. This environment operates *in parallel with* the production servers, but is not used to serve actual requests nor are requests delayed. Rather, it is used to detect “blind” attacks, such as when a worm or an attacker is randomly scanning and attacking IP addresses. We use this environment as a “clean room” to test the effects of “suspicious” requests, such as potential worm infection vectors. A request that causes a buffer overflow on the production server will have the same effect on the sandboxed version of the application. The instrumentation allows us to determine the buffers and functions involved in a buffer overflow attack. This information is then passed on to the production server, which enables that subset of the defenses that is necessary to protect against the detected exploit. In contrast with our previous work, where patches were dynamically generated “on the fly” [76, 77], DYBOC allows administrators to test the functionality and performance of the software with all protection components enabled. Even by itself, the honeypot mode of operation can significantly accelerate the identification of new attacks and the generation of patches or the invocation of other protection mechanisms, improving on the current state-of-the-art in attack detection [37, 65].

We describe our approach and the prototype implementation in Sect. 2. We then evaluate its performance and effectiveness in Sect. 3, and give a brief overview of related work in Sect. 4.

2 Our approach

The core of our approach is to automatically instrument parts of the application source code¹ that may be vulnerable to buffer overflow attacks (i.e., buffers declared in the stack or the heap) such that overflow or underflow attacks cause an exception. We then catch these exceptions and recover the program execution from a suitable location.

This description raises several questions: Which buffers are instrumented? What is the nature of the instrumentation? How can we recover from an attack, once it has been detected? Can all this be done efficiently and effectively? In the following subsections we answer these questions and describe the main components of our system. The question of efficiency and effectiveness is addressed in the next section.

¹ Binary rewriting techniques may be applicable, but we do not further consider them due to their significant complexity.

2.1 Instrumentation

Since our goal is to contain buffer overflow attacks, our system instruments all statically and dynamically allocated buffers, and all read and writes to these buffers. In principle, we could combine our system with a static analysis tool to identify those buffers (and uses of buffers) that are provably safe from exploitation. Although such an approach would be an integral part of a complete system, we do not examine it further here; we focus on the details of the dynamic protection mechanism. Likewise, we expect that our system would be used in conjunction with a mechanism like StackGuard [22] or ProPolice to prevent successful intrusions against attacks we are not yet aware of; following such an attack, we can enable the dynamic protection mechanism to prevent service disruption. We should also note the “prove and check” approach has been used in the context of software security in the past, most notably in CCured [57]. In the remainder of this paper, we will focus on stack-based attacks, although our technique can equally easily defend against heap-based ones.

For the code transformations we use TXL [40], a hybrid functional and rule-based language which is well-suited for performing source-to-source transformation and for rapidly prototyping new languages and language processors.

The instrumentation is fairly straightforward: we move static buffers to the heap, by dynamically allocating the buffer upon entering the function in which it was previously declared; we deallocate these buffers upon exiting the function, whether implicitly (by reaching the end of the function body) or explicitly (through a *return* statement).

For memory allocation we use *pmalloc()*, our own version of *malloc()*, which allocates two zero-filled, write-protected pages surrounding the requested buffer (Fig. 1).

The guard pages are *mmap()*’ed from */dev/zero* as read-only. As *mmap()* operates at memory page granularity, every memory request is rounded up to the nearest page. The




Fig. 1 Example of *pmalloc()*-based memory allocation: the trailer and edge regions (above and below the write-protected pages) indicate “waste” memory allocated by *malloc()*

Original code	Modified code
<pre>int func() { char buf[100]; ... other_func(buf); ... return 0; }</pre>	<pre>int func() { char *buf = pmalloc(100); ... other_func(buf); ... pfree(buf); return 0; }</pre>

Fig. 2 First-stage transformation, moving buffers from the stack to the heap with *pmalloc()*

pointer that is returned by *pmalloc()* can be adjusted to immediately catch any buffer overflow or underflow depending on where attention is focused. This functionality is similar to that offered by the *ElectricFence* memory-debugging library, the difference being that *pmalloc()* catches both buffer overflow and underflow attacks. Because we *mmap()* pages from */dev/zero*, we do not waste physical memory for the guards (just page-table entries). Some memory is wasted, however, for each allocated buffer, since we round to the next closest page. While this could lead to considerable memory waste, we note that in our experiments the overhead has proven manageable.

Figure 2 shows an example of such a translation. Buffers that are already allocated via *malloc()* are simply switched to *pmalloc()*. This is achieved by examining declarations in the source and transforming them to pointers where the size is allocated with a *malloc()* function call. Furthermore, we adjust the C grammar to free the variables before the function returns. After making changes to the standard ANSI C grammar that allow entries such as *malloc()* to be inserted between declarations and statements, the transformation step is trivial. For single-threaded, nonreentrant code, it is possible to use *pmalloc()* once for each previously-allocated static buffer. Generally, however, this allocation needs to be done each time the function is invoked. We discuss how to minimize this cost in Sect. 2.3.

Any overflow or underflow attack to a *pmalloc()*-allocated buffer will cause the process to receive a Segmentation Violation (SEGV) signal, which is caught by a signal handler we have added to the source code. It is then the responsibility of the signal handler to recover from such failures.

2.2 Recovery: execution transactions

In determining how to recover from such exception, we introduce the hypothesis of an *execution transaction*. Very simply, we posit that for the majority of code (and for the purposes of defending against buffer overflow attacks), we can treat each function execution as a transaction (in a manner similar to a sequence of operations in a database) that can be aborted without adversely affecting the graceful termination of the computation. Each function call from inside that function can itself be treated as a transaction, whose success (or failure) does not contribute to the success or failure of its enclosing transaction. Under this hypothesis, it is sufficient to snapshot the state of the program execution when a new

transaction begins, detect a failure per our previous discussion, and recover by aborting this transaction and continuing the execution of its enclosing transaction. Currently, we focus our efforts inside the process address space, and do not deal with rolling back I/O. For this purpose, a virtual file system approach can be employed to roll back any I/O that is associated with a process. We plan to address this further in future work, by adopting the techniques described in [83]. However, there are limitations to what can be done, e.g., network traffic.

Note that our hypothesis does not imply anything about the correctness of the resulting computation, when a failure occurs. Rather, it merely states that if a function is prevented from overflowing a buffer, it is sufficient to continue execution at its enclosing function, “pretending” the aborted function returned an error. Depending on the return type of the function, a set of heuristics are employed so as to determine an appropriate error return value that is, in turn, used by the program to handle error conditions. Details of this approach are described in Sect. 2.3. Our underlying assumption is that the remainder of the program can handle truncated data in a buffer in a graceful manner. For example, consider the case of a buffer overflow vulnerability in a web server, whereby extremely long URLs cause the server to be subverted: when DYBOC catches such an overflow, the web server will simply try to process the truncated URL (which may simply be garbage, or may point to a legitimate page).

A secondary assumption is that most functions that are thusly aborted do not have other side effects (e.g., touch global state), or that such side effects can be ignored. *Obviously, neither of these two conditions can be proven, and examples where they do not hold can be trivially constructed, e.g., an mmap()'ed file shared with other applications.* Since we are interested in the actual behavior of real software, we experimentally evaluate our hypothesis in Sect. 3. Note that, in principle, we could checkpoint and recover from each instruction (line of code) that “touches” a buffer; doing so, however, would be prohibitively expensive.

To implement recovery we use `sigsetjmp()` to snapshot the location to which we want to return once an attack has been detected. The effect of this operation is to save the stack pointers, registers, and program counter, such that the program can later restore their state. We also inject a signal handler (initialized early in `main()`) that catches `SIGSEGV`² and calls `siglongjmp()`, restoring the stack pointers and registers (including the program counter) to their values prior to the call of the offending function (in fact, they are restored to their values as of the call to `sigsetjmp()`):

```
void sigsegv_handler() {
    /* transaction(TRANS_ABORT); */
    siglongjmp(global_env, 1);
}
```

(We explain the meaning of the `transaction()` call later in this section.) The program will then re-evaluate the injected

conditional statement that includes the `sigsetjmp()` call. This time, however, the return value will cause the conditional to evaluate to false, thereby skipping execution of the offending function. Note that the targeted buffer will contain exactly the amount of data (infection vector) it would if the offending function performed correct data-truncation. In our example, after a fault, execution will return to the conditional statement just prior to the call to `other_func()`, which will cause execution to skip another invocation of `other_func()`. If `other_func()` is a function such as `strcpy()`, or `sprintf()` (i.e., code with no side effects), the result is similar to a situation where these functions correctly handled array-bounds checking.

There are two benefits to this approach. First, objects in the heap are protected from being overwritten by an attack on the specified variable since there is a signal violation when data is written beyond the allocated space. Second, we can recover gracefully from an overflow attempt, since we can recover the stack context environment prior to the offending function’s call, and effectively `siglongjmp()` to the code immediately following the routine that caused the overflow or underflow. While the contents of the stack can be recovered by restoring the stack pointer, special care must be placed in handling the state of the heap. To deal with data corruption in the heap, we can employ data structure consistency constraints, as described in [24], to detect and recover from such errors. Thus, the code in our example from Fig. 2 will be transformed as shown in Fig. 3 (grayed lines indicate changes from the previous example).

To accommodate multiple functions checkpointing different locations during program execution, a globally defined `sigjmp_buf` structure always points to the latest snapshot to recover from. Each function is responsible for saving and restoring this information before and after invoking a subroutine respectively, as shown in Fig. 4.

Functions may also refer to global variables; ideally, we would like to unroll any changes made to them by an aborted transaction. The use of such variables can be determined fairly easily via lexical analysis of the instrumented function: any *l-values* not defined in the function are assumed to be global variables (globals used as *r-values* do not cause any changes to their values, and can thus be safely ignored).

```
int func()
{
    char *buf;
    buf = pmalloc(100);
    ...
    if (sigsetjmp(global_env, 1) == 0) {
        other_func(buf); /* Indented */
    }
    ...
    pfree(buf);
    return 0;
}

/* Global definitions */
sigjmp_buf global_env;
```

Fig. 3 Saving state for recovery

² Care must be taken to avoid an endless loop on the signal handler if another such signal is raised while in the handler. We apply our approach on OpenBSD and Linux RedHat.

```

int func()
{
    char *buf;
    sigjmp_buf curr_env;
    sigjmp_buf *prev_env;
    buf = pmalloc(100);
    ...
    if (sigsetjmp(curr_env, 1) == 0) {
        prev_env = global_env;
        global_env = &curr_env;
        other_func(buf); /* Indented */
        global_env = prev_env;
    }
    ...
    pfree(buf);
    return 0;
}

```

Fig. 4 Saving previous recovery context

```

/* Global variables */
int global;

int func()
{
    char *buf;
    sigjmp_buf curr_env;
    sigjmp_buf *prev_env;
    buf = pmalloc(100);
    int temp_dyboe_global;
    ...
    if (sigsetjmp(curr_env, 1) == 0) {
        temp_dyboe_global = global;
        /* OR: transaction(TRANS.START); */
        prev_env = global_env;
        global_env = &curr_env;
        other_func(buf); /* Indented */
        global_env = prev_env;
    } else {
        global = temp_dyboe_global;
        /* OR: transaction(TRANS.END); */
    }
    ...
    pfree(buf);
    return 0;
}

```

Fig. 5 Saving global variable

Once the name of the global variable has been determined, we scan the code to determine its type. If it is a basic type (e.g., integer, float, character), a fixed-size array of a basic type, or a statically allocated structure, we define a temporary variable of the same type in the enclosing function and save/restore its original value as needed. In the example shown in Fig. 5, variable “global” is used in *other_func()*.

Unfortunately, dynamically allocated global data structures (such as hash tables or linked lists) are not as straightforward to handle in this manner, since their size may be determined at run time and thus be indeterminate to a static lexical analyzer. Thus, when we cannot determine the side-effects of a function, we use a different mechanism, assisted by the operating system: we added a new system call, named *transaction()*. This is conditionally invoked (as directed by the *dyboe_flag()* macro) at three locations in the code, as shown in Fig. 5.

First, prior to invoking a function that may be aborted, to indicate to the operating system that a new transaction has

begun. The OS makes a backup of all memory page permissions, and marks all heap memory pages as read-only. As the process executes and modifies these pages, the OS maintains a copy of the original page and allocates a new page (which is given the permissions the original page had, from the backup) for the process to use, in exactly the same way copy-on-write works in modern operating systems. Both copies of the page are kept until *transaction()* is called again. Second, after the end of a transaction (execution of a vulnerable function), to indicate to the operating system that a transaction has successfully completed. The OS then discards all original copies of memory pages that have been modified during processing this request. Third, in the signal handler, to indicate to the OS that an exception (attack) has been detected. The OS then discards all dirty pages by restoring the original pages.

A similar mechanism could be built around the filesystem by using a private copy of the buffer cache for the process executing in shadow mode, although we have not implemented it. The only difficulty arises when the process must itself communicate with another process while servicing a request; unless the second process is also included in the transaction definition (which may be impossible, if it is a remote process on another system), overall system state may change without the ability to roll it back. For example, this may happen when a web server communicates with a back-end database. Our system does not currently address this, i.e., we assume that any such state changes are benign or irrelevant (e.g., a DNS query). Back-end databases inherently support the concept of a transaction rollback, so it is (in theory) possible to undo any changes.

The signal handler may also notify external logic to indicate that an attack associated with a particular input from a specific source has been detected. The external logic may then instantiate a filter, either based on the network source of the request or the contents of the payload.

2.3 Dynamic defensive postures

“Eternal vigilance is the price of liberty.” – Wendell Phillips, 1852

Unfortunately, when it comes to security mechanisms, vigilance takes a back seat to performance. Thus, although our mechanism can defend against all buffer overflow attacks and (as we shall see in Sect. 3) maintains service availability in the majority of cases, this comes at the cost of performance degradation. Although such degradation seems to be modest for some applications (about 20% for Apache, see Sect. 3), it is conceivable that other applications may suffer a significant performance penalty if all buffers are instrumented with our system (for example, a worst-case microbenchmark measurement indicates a 440% slowdown). One possibility we already mentioned is the use of static analysis tools to reduce the number of buffers that need to be instrumented; however, it is very likely that a significant number of these will remain unresolved, requiring further protection.

```

int func()
{
    char *buf;
    sigjmp_buf curr_env, *prev_env;
    char _buf[100];
    if (dyboc_flag(827))
        buf = pmalloc(100); /* Indented */
    else
        buf = _buf;
    ...
    if (dyboc_flag(1821)) {
        if (sigsetjmp(curr_env, 1) == 0) {
            prev_env = global_env;
            global_env = &curr_env;
            other_func(buf);
            global_env = prev_env;
        } else {
            other_func(buf);
        }
    }
    ...
    if (dyboc_flag(827)) {
        pfree(buf); /* Indented */
    }
    return 0;
}

```

Fig. 6 Enabling DYBOC conditionally

Our scheme makes it possible to *selectively* enable or disable protection for specific buffers in functions, in response to external events (e.g., an administrator command, or an automated intrusion detection system). In the simplest case, an application may execute with all protection disabled, only to assume a more defensive posture as a result of increased network scanning and probing activity. This allows us to avoid paying the cost of instrumentation most of the time, while retaining the ability to protect against attacks quickly. Although this strategy entails some risks (exposure to a successful directed attack with no prior warning), it may be the only alternative when we wish to achieve security, availability, and performance.

The basic idea is to only use *pmalloc()* and *pfree()* if a flag instructs the application to do so; otherwise, the transformed buffer is made to point to a statically allocated buffer. Similarly, the *sigsetjmp()* operation is performed only when the relevant flag indicates so. This flagging mechanism is implemented through the *dyboc_flag()* macro, which takes as argument an identifier for the current allocation or checkpoint, and returns true if the appropriate action needs to be taken. Continuing with our previous example, the code will be transformed as shown in Fig. 6. Note that there are three invocations of *dyboc_flag()*, using two different identifiers: the first and last use the same identifier, which indicates whether a particular buffer should be *pmalloc()*’ed or be statically allocated; the second invocation, with a different identifier, indicates whether a particular transaction (function call) should be checkpointed.

To implement the signaling mechanism, we use a shared memory segment of sufficient size to hold all identifiers (1 bit per flag). *dyboc_flag()* then simply tests the appropriate flag. A second process, acting as the *notification monitor* is responsible for setting the appropriate flag, when notified through a command-line tool or an automated mecha-

nism. Turning off a flag requires manual intervention by the administrator. We do not address memory leaks due to the obvious race condition (turning off the flag while a buffer is already allocated), since we currently only examine single threaded cases and we expect the administrator to restart the service under such rare circumstances, although these can be addressed with additional checking code. Other mechanisms that can be used to address memory leaks and inconsistent data structures are recursive restartability [12] and micro-booting [11]. We intend to examine these in future work.

2.4 Worm containment

Recent incidents have demonstrated the ability of self-propagating code, also known as “network worms,” to infect large numbers of hosts, exploiting vulnerabilities in the largely homogeneous deployed software base (or even a small homogeneous base [75]), often affecting the offline world in the process [48]. Even when a worm carries no malicious payload, the direct cost of recovering from the side effects of an infection epidemic can be tremendous. Countering worms has recently become the focus of increased research, generally focusing on content-filtering mechanisms.

Despite some promising early results, we believe that in the future this approach will be insufficient. We base this primarily on two observations. First, to achieve coverage, such mechanisms are intended for use by routers (e.g., Cisco’s NBAR); given the routers’ limited budget in terms of processing cycles per packet, even mildly polymorphic worms (mirroring the evolution of polymorphic viruses, more than a decade ago [84]) are likely to evade such filtering, as demonstrated recently in [18]. Network-based intrusion detection systems (NIDS) have encountered similar problems, requiring fairly invasive packet processing and queuing at the router or firewall. When placed in the application’s critical path, as such filtering mechanisms must, they will have an adverse impact on performance, as well as cause a large number of false positive alerts [60]. Second, end-to-end “opportunistic” encryption in the form of TLS/SSL or IPsec is being used by an increasing number of hosts and applications. We believe that it is only a matter of time until worms start using such encrypted channels to cover their tracks. These trends argue for an end-point worm-countering mechanism. Mechanisms detecting misbehavior [94] are more promising in that respect.

The mechanism we have described allows us to create an autonomous mechanism for combating a scanning (but not hit-list) worm that does not require snooping the network. We use two instances of the application to be protected (e.g., a web server), both instrumented as described above. The production server (which handles actual requests) is operating with all security disabled; the second server, which runs in honeypot mode [65], is listening on an unadvertised address. A scanning worm such as Blaster, CodeRed, or Slammer (or an automated exploit toolkit that scans and attacks any vulnerable services) will trigger an exploit on the honeypot server; our instrumentation will allow us to determine

which buffer and function are being exploited by the particular worm or attack. This information will then be conveyed to the production server notification monitor, which will set the appropriate flags. A service restart may be necessary, to ensure that no instance of the production server has been infected while the honeypot was detecting the attack. The payload that triggered the buffer overflow, the first part of which can be found on the instrumented buffer, may also be used for content-based filtering at the border router (with the caveats described above). Thus, our system can be used in quickly deriving content-filter rules for use by other mechanisms. Active honeypot techniques such as those proposed in [96] can make it more difficult for attackers to discriminate between the honeypot and the production server.

Thus, targeted services can automatically enable those parts of their defenses that are necessary to defend against a particular attack, without incurring the performance penalty at other times, and cause the worm to slow down. There is no dependency on some critical mass of collaborating entities, as with some other schemes: defenses are engaged in a completely decentralized manner, independent of other organizations' actions. Wide-spread deployment would cause worm outbreaks to subside relatively quickly, as vulnerable services become immune after being exploited. This system can protect against zero-day attacks, for which no patch or signature is available.

3 Experimental evaluation

To test the capabilities of our system, we conducted a series of experiments and performance measurements. Results were acquired through the examination of the applications provided by the Code Security Analysis Kit (CoSAK) project.

3.1 Security analysis

To determine the validity of our execution transactions hypothesis, we examined a number of vulnerable open-source software products. This data was made available through the Code Security Analysis Kit (CoSAK) project from the software engineering research group at Drexel university. CoSAK is a DARPA-funded project that is developing a toolkit for software auditors to assist with the development of high-assurance and secure software systems. They have compiled a database of thirty open source products along with their known vulnerabilities and respective patches. This database is comprised of general vulnerabilities, with a large number listed as susceptible to buffer overflow attacks. We applied DYBOC against this data set.

Our tests resulted in fixing 14 out of 17 "fixable" buffer overflow vulnerabilities, a 82% success rate. The remaining 14 packages in the CoSAK suite were not tested because their vulnerabilities were unrelated (non buffer-overflow). In the remaining three cases (those for which our hypothesis

appeared not to hold), we manually inspected the vulnerabilities and determined that what would be required to provide an appropriate fix are adjustments to the DYBOC tool to cover special cases, such as handling multidimensional buffers and preinitialized arrays; although these are important in a complete system, we feel that our initial results were encouraging.

3.2 Execution transaction validation

In order to evaluate the validity of our hypothesis on the recovery of execution transactions, we experimentally evaluate its effects on program execution on the Apache web server. We run a profiled version of ApacheBench against a set of concurrent requests generated by ApacheBench and examine the subsequent call-graph generated by these requests with *gprof*.

The call tree is analyzed in order to determine which functions are used. These functions are, in turn, employed as potentially vulnerable transactions. As mentioned previously, we treat each function execution as a transaction that can be aborted without incongruously affecting the normal termination of computation. Armed with the information provided by the call-graph, we run a TXL script that inserts an early return in all the functions, simulating an aborted transaction.

This TXL script operates on a set of heuristics that were devised for the purpose of this experiment. Briefly, depending on the return type of the function, an appropriate value is returned. For example, if the return type is an *int*, a -1 is returned; if the value is *unsigned int*, we return 0, *etc.* A special case is used when the function returns a pointer. Specifically, instead of blindly returning a *NULL*, we examine if the pointer returned is dereferenced later by the calling function. In this case, we issue an early return immediately before the terminal function is called. For each simulated aborted transaction, we monitor the program execution of Apache by running *httperf*, a web server performance measurement tool. Specifically, we examined 154 functions.

The results from these tests were very encouraging; 139 of the 154 functions completed the *httperf* tests successfully: program execution was not interrupted. What we found to be surprising, was that not only did the program not crash but in some cases all the pages were served correctly. This is probably due to the fact a large number of the functions are used for statistical and logging purposes. Out of the 15 functions that produced segmentation faults, four did so at startup.

Similarly for *sshd*, we iterate through each aborted function while examining program execution during an *scp* transfer. In the case of *sshd*, we examined 81 functions. Again, the results were encouraging: 72 of the 81 functions maintained program execution. Furthermore, only four functions caused segmentation faults; the rest simply did not allow the program to start.

For *Bind*, we examined the program execution of *named* during the execution of a set of queries; 67 functions were

tested. In this case, 59 of the 67 functions maintained the proper execution state. Similar to `sshd`, only four functions caused segmentation faults.

Naturally, it is possible that Apache, Bind, and `sshd` will exhibit long-term side effects, e.g., through data structure corruption. Our experimental evaluation through a benchmark suite, which issues many thousand requests to the same application, gives us some confidence that their internal state does not “decay” quickly. To address longer-term deterioration, we can use either micro-rebooting (software rejuvenation) [11] or automated data-structure repair [24]. We intend to examine the combination of our approach with either of these techniques in future work.

3.3 Performance overheads

To understand the performance implications of our protection mechanism, we run a set of performance benchmarks. We first measure the worst-case performance impact of DYBOC in a contrived program; we then run DYBOC against the Apache web server and measure the overhead of full protection.

The first benchmark is aimed at helping us understand the performance implications of our DYBOC engine. For this purpose, we use an austere C program that makes an `strcpy()` call using a statically allocated buffer as the basis of our experiment.

After patching the program with DYBOC, we compare the performance of the patched version to that of the original version by examining the difference in processor cycles using the Read Time Stamp Counter (RDTSC), found in Pentium class processors. The results illustrated by Fig. 7 indicate the mean time, in microseconds (adjusted from the processor cycles), for 100,000 iterations. The performance overhead for the patched, protected version is 440%, which is expected given the complexity of the `pmalloc()` rou-




Fig. 7 Micro-benchmark results




Fig. 8 Apache benchmark results

tine relative to the simplicity of calling `strcpy()` for small strings.

We also used DYBOC on the Apache web server, version 2.0.49. Apache was chosen due to its popularity and source-code availability. Basic Apache functionality was tested, omitting additional modules. Our goal was to examine the overhead of preemptive patching of a software system. The tests were conducted on a PC with a 2 GHz Intel P4 processor and 1 GB of RAM, running Debian Linux (2.6.5-1 kernel).

We used ApacheBench, a complete benchmarking and regression testing suite. Examination of application response is preferable to explicit measurements in the case of complex systems, as we seek to understand the effect on overall system performance.

Figure 8 illustrates the requests per second that Apache can handle. There is a 20.1% overhead for the patched version of Apache over the original, which is expected since the majority of the patched buffers belong to utility functions that are not heavily used. This result is an indication of the worst-case analysis, since all the protection flags were enabled; although the performance penalty is high, it is not outright prohibitive for some applications. For the instrumentation of a single buffer and a vulnerable function that is invoked once per HTTP transaction, the overhead is 1.18%.

3.4 Space overheads

The line count for the server files in Apache is 226,647, while the patched version is 258,061 lines long, representing an increase of 13.86%. Note that buffers that are already being allocated with `malloc()` (and de-allocated with `free()`) are simply translated to `pmalloc()` and `pfree()` respectively, and thus do not contribute to an increase in the line count. The binary size of the original version was 2,231,922 bytes, while the patched version of the binary was 2,259,243 bytes, an increase of 1.22%. Similar results are obtained with OpenSSH




Fig. 9 Simulating the effects of using DYBOC on worm propagation. We vary the percentage of protected hosts and examine the effects on infection rate and overall susceptible population

3.7.1. Thus, the impact of our approach in terms of additional required memory or disk storage is minimal.

3.5 Effectiveness as a worm containment strategy

In order to analyze the effectiveness of our technique as a worm containment strategy, we examine its effects on worm propagation. In particular, we model the impact of our approach on a scanning worm with propagation characteristics similar to Code Red. Worm containment is comprised of three steps: worm detection, patch generation, and patch dissemination. In this work, we address the effects of detection and patch generation on worm propagation.

Modelling worm propagation has been addressed extensively in the literature [53, 94, 95, 97] and is often described using variations on the susceptible-infected-susceptible (SIS) and susceptible-infected-recovered (SIR) models borrowed from classic epidemiology research. The SIR model is more suitable for worm modelling since it takes into account recovered (patched) systems. We model our system as a variation to the SIR model where infected hosts that are protected using DYBOC recover after the time period required to detect a vulnerability and patch a system. Specifically, we use a model similar to the one proposed by Wong et al. [95] where the population size is 10000 and the probability of selecting a host out of the network's address space is 10000/65535.

As illustrated in Fig. 9, we examine the effects of using DYBOC on worm propagation by varying the percentage of

susceptible hosts that are protected with our tool. Specifically, we vary the percentage of DYBOC protected hosts from 0% to 20% and inspect the effects on overall infected population and the rate of infection. When a worm infects a node that is protected with DYBOC, the node continues as an infected host for the time required to generate and test a patch, effectively participating in worm propagation during this period. We model this period as a random value between one and three minutes. After this period, the node joins the group of patched (recovered) hosts.

Upon first glance, it is evident that although DYBOC mitigates a worm's effect on the susceptible population, protection of individual hosts will not suffice as an attempt at curtailing worm propagation (unless a very large portion of the population is protected) without further collaboration between hosts. Collaboration techniques may include patch dissemination and collaborative filtering. Patch dissemination techniques have been addressed in [20] and analyzed by Vojnovic and Ganesh [53] where the authors show that given significant participation and swift patch generation a worm's propagation rate can be restricted if the patch spreading rate is higher than the worm's infection rate. In more detail, they explore patch dissemination using a hierarchical system where hosts responsible for regions can also act as filters (or firewalls). Nicol et al. [59] also examine models of active worm defenses; they investigate the effects of filtering in conjunction with patch dissemination. Unfortunately, their analysis only holds for scanning worms. Very fast worms (flash, hitlist) tend to have infection rates that are as high, if not higher, as patch dissemina-

tion rates. Furthermore, fast scanning worms do not follow traditional epidemic models since their propagation is often shaped by the affect they have on the underlying network infrastructure.

4 Related work

Modeling executing software as a transaction that can be aborted has been examined in the context of language-based runtime systems (specifically, Java) in [70, 71]. That work focused on safely terminating misbehaving threads, introducing the concept of “soft termination”. Soft termination allows threads to be terminated while preserving the stability of the language runtime, without imposing unreasonable performance overheads. In that approach, threads (or *codelets*) are each executed in their own transaction, applying standard ACID semantics. This allows changes to the runtime’s (and other threads’) state made by the terminated codelet to be rolled back. The performance overhead of their system can range from 200 up to 2,300%. Relative to that work, our contribution is twofold. First, we apply the transactional model to an unsafe language such as C, addressing several (but not all) challenges presented by that environment. Second, by selectively applying transactional processing, we substantially reduce the performance overhead of the application. However, there is no free lunch: this reduction comes at the cost of allowing failures to occur. Our system aims to automatically evolve code such that it *eventually* (i.e., after an attack has been observed) does not succumb to attacks.

Some interesting work has been done to deal with memory errors at runtime. For example, Rinard et al. [69] have developed a compiler that inserts code to deal with writes to unallocated memory by virtually expanding the target buffer. Such a capability aims toward the same goal our system does: provide a more robust fault response rather than simply crashing. The technique presented in [69] is modified in [68] and introduced as *failure-oblivious computing*. Because the program code is extensively re-written to include the necessary check for *every* memory access, their system incurs overheads ranging from 80 up to 500% for a variety of different applications.

Modeling executing software as a transaction that can be aborted has been examined in the context of language-based runtime systems (specifically, Java) in [70, 71]. That work focused on safely terminating misbehaving threads, introducing the concept of “soft termination.” Soft termination allows threads to be terminated while preserving the stability of the language runtime, without imposing unreasonable performance overheads. In that approach, threads (or *codelets*) are each executed in their own transaction, applying standard ACID semantics. This allows changes to the runtime’s (and other threads’) state made by the terminated codelet to be rolled back. The performance overhead of their system can range from 200 up to 2,300%. Relative to that

work, our contribution is twofold. First, we apply the transactional model to an unsafe language such as C, addressing several (but not all) challenges presented by that environment. Second, by selectively applying transactional processing, we substantially reduce the performance overhead of the application. However, there is no free lunch: this reduction comes at the cost of allowing failures to occur. Our system aims to automatically evolve a piece of code such that it *eventually* (i.e., after an attack has been observed) does not succumb to attacks.

Some interesting work has been done to deal with memory errors at runtime. For example, Rinard et al. [69] have developed a compiler that inserts code to deal with writes to unallocated memory by virtually expanding the target buffer. Such a capability aims toward the same goal our system does: provide a more robust fault response rather than simply crashing. The technique presented in [69] is modified in [68] and introduced as *failure-oblivious computing*. Because the program code is extensively re-written to include the necessary check for *every* memory access, their system incurs overheads ranging from 80 up to 500% for a variety of different applications.

Suh et al. [82], propose a hardware based solution that can be used to thwart control-transfer attacks and restrict executable instructions by monitoring “tainted” input data. In order to identify “tainted” data, they rely on the operating system. If the processor detects the use of this tainted data as a jump address or an executed instruction, it raises an exception that can be handled by the operating system. The authors do not address the issue of recovering program execution and suggest the immediate termination of the offending process. DIRA [78] is a technique for automatic detection, identification and repair of control-hijacking attacks. This solution is implemented as a GCC compiler extension that transforms a program’s source code adding heavy instrumentation so that the resulting program can perform these tasks. The use of checkpoints throughout the program ensures that corruption of state can be detected if control sensitive data structures are overwritten. Unfortunately, the performance implications of the system make it unusable as a front line defense mechanism. Song and Newsome [58] propose dynamic taint analysis for automatic detection of overwrite attacks. Tainted data is monitored throughout the program execution and modified buffers with tainted information will result in protection faults. Once an attack has been identified, signatures are generated using automatic semantic analysis. The technique is implemented as an extension to Valgrind and does not require any modifications to the program’s source code but suffers from severe performance degradation.

One of the most critical concerns with recovering from software faults and vulnerability exploits is ensuring the consistency and correctness of program data and state. An important contribution in this area is [25], which discusses mechanisms for detecting corrupted data structures and fixing them to match some pre-specified constraints. While the precision of the fixes with respect to the semantics of the

program is not guaranteed, their test cases continued to operate when faults were randomly injected. Similar results are shown in [92]: when program execution is forced to take the “wrong” path at a branch instruction, program behavior remains the same in over half the times.

4.1 Source code analysis

Increasingly, source code analysis techniques are brought to bear on the problem of detecting potential code vulnerabilities [17]. The simplest approach has been that of the compiler warning on the use of certain unsafe functions, e.g., *gets()*. More recent approaches [26, 29, 41, 46, 74, 89] have focused on detecting specific types of problems, rather than try to solve the general “bad code” issue, with considerable success. While such tools can greatly help programmers ensure the safety of their code, especially when used in conjunction with other protection techniques, they (as well as dynamic analysis tools such as [47, 49]) offer incomplete protection, as they can only protect against and detect *known* classes of vulnerabilities.

MOPS [15, 16] is an automated formal-methods framework for finding bugs in security-relevant software, or verifying their absence. They model programs as pushdown automata, represent security properties as finite state automata, and use model-checking techniques to identify whether any state violating the desired security goal is reachable in the program. While this is a powerful and scalable (in terms of performance and size of program to be verified) technique, it does not help against buffer overflow or other code-injection attacks. RacerX [27] uses flow-sensitive, inter-procedural analysis to detect race conditions and deadlocks, geared towards debugging of large multithreaded systems.

4.2 Process sandboxing

Process sandboxing [61] is perhaps the best understood and widely researched area of containing bad code, as evidenced by the plethora of available systems like SubOS [39, 38], Janus [35], Consh [3], Mapbox [1], OpenBSD’s *sysrace* [66], and the Mediating Connectors [6]. These operate at user level and confine applications by filtering access to system calls. To accomplish this, they rely on *ptrace(2)*, the */proc* file system, and/or special shared libraries. Another category of systems, such as Tron [8], SubDomain [21] and others [30, 34, 52, 54, 63, 90, 93], go a step further. They intercept system calls inside the kernel, and use policy engines to decide whether to permit the call or not, logically extending the model of standard access control lists, which typically operate only on filesystem objects [36]. The main problem with all these is that the attack is not prevented: rather, the system tries to limit the damage such code can do, such as obtain super-user privileges. Thus, the system does not protect against attacks that use the subverted process’ privileges to bypass application-specific access control checks (e.g., read all the pages a web server has access

to), nor does it protect against attacks that simply use the subverted program as a stepping stone, as is the case with network worms. [33] identifies several common security-related problems with such systems, such as their susceptibility to various types of race conditions. In [32], the authors propose a delegating architecture that uses privileged agents that handle security-critical operations and a process’s interactions with the rest of the system.

User-level sandboxing approaches [64] using a second monitor process seem more promising, especially with the appearance of automated tools [9] and libraries [44] to assist the programmer; however, these require access and significant modifications to the source code of the application.

Model-carrying code [73] augments untrusted code with a concise high-level model of its security-relevant behavior. This model helps bridge the gap between high-level security policies and the binary code itself. The authors show that this approach can be used to verify whether the code satisfied a security policy, or to determine which policy (from a catalogue of pre-defined ones) is the most appropriate to use.

4.3 Compiler tricks

Perhaps the best-known approach to countering buffer overflows is StackGuard [22]. This is a patch to the popular *gcc* compiler that inserts a *canary* word right before the return address in a function’s activation record on the stack. The canary is checked just before the function returns, and execution is halted if it is not the correct value, which would be the case if a stack-smashing attack had overwritten it. This protects against simple stack-based attacks, although some attacks were demonstrated against the original approach [10], which has since been amended to address the problem. ProPolice [28] also uses a canary value to detect stack-based attacks. It also rearranges variables in the stack so that *char* buffers are always allocated at the bottom of the record. Thus, overflowing these buffers cannot harm other local variables, especially function-pointer variables, avoiding attacks that modify the values of local variables without overwriting the canary. However, attacks against both ProPolicy and Stackguard have been demonstrated recently [81], taking advantage of the linear encryption algorithm (XOR), used to obfuscate/protect pointers, to make a brute-force attack against the randomizing key feasible.

Another approach implemented as a *gcc* patch [42] adds bounds-checking for pointers and arrays without changing the memory model used for representing pointers. This prevents buffer overflow exploits, but at a high performance cost, since all indirect memory accesses are checked, greatly slowing program execution. A somewhat more efficient version is described in [72]. Stack Shield [88] is another *gcc* extension with an activation record-based approach. The technique involves saving the return address to a write-protected memory area, which is impervious to buffer overflows, when the function is entered. Before returning from the function, the system restores the proper return address value. MemGuard [22] makes the location of the return address in the

function prologue read-only and restores it upon function return, disallowing any writes to the section of memory containing the return address. It permits writes to other locations in the same virtual memory page, but slows them down considerably because they must be handled by kernel code.

Program shepherding [45] uses an interpreter to verify the source and target of any branch instruction, based on some security policy. To avoid the performance penalty of interpretation, their system caches verified code segments and reuses them as needed. Despite this, there is a considerable performance penalty for some applications. *libverify* [7] dynamically re-writes executed binaries to add run-time return-address checks.

4.4 Worms

Computer viruses are not a new phenomenon, and they have been studied extensively over the last several decades. Cohen was the first to define and describe computer viruses in their present form. In [19], he gave a theoretical basis for the spread of computer viruses. In [80], the authors describe the risk to the Internet due to the ability of attackers to quickly gain control of vast numbers of hosts. Their analysis shows how quickly attackers can compromise hosts using “dumb” worms and how “better” worms can spread even faster. The strong analogy between biological and computer viruses led Kephart et al. to investigate the propagation of computer viruses based on epidemiological models. In [43], they extend the standard epidemiological model by placing it on a directed graph, and use a combination of analysis and simulation to study its behavior. They conclude that if the rate at which defense mechanisms detect and remove viruses is sufficiently high, relative to the rate at which viruses spread, they are adequate for preventing widespread propagation of viruses.

The CodeRed worm [13] was analyzed extensively in [98]. The authors of that work conclude that even though epidemic models can be used to study the behavior of Internet worms, they are not accurate enough because they cannot capture some specific properties of the environment these operate in: the effect of human countermeasures against worm spreading (i.e., cleaning, patching, filtering, disconnecting, *etc.*), and the slowing down of the worm infection rate due to the worm’s impact on Internet traffic and infrastructure. They derive a new general Internet worm model called *two-factor worm* model, which they then validate in simulations that match the observed Code Red data available to them. Their analysis seems to be supported by the data on Code Red propagation in [55] and [79] (the latter distinguishes between different worms simultaneously active). A similar analysis on the SQL “Slammer” worm [14] can be found in [85]. More recent analyses [97] show that it is possible to predict the overall vulnerable population size using Kalman filters early in the propagation cycle of a worm, allowing for detection of a fast-spreading worm when only 1%–2% of vulnerable computers on the network have been infected.

CodeRed inspired several countermeasures, such as La Brea [51], which attempts to slow the growth of TCP-based worms by accepting connections and then blocking on them indefinitely, causing the corresponding worm thread to block. Under the connection-throttling approach [87], each host restricts the rate at which connections may be initiated. If adopted universally, such an approach can reduce the spreading rate of a worm by up to an order of magnitude, without affecting legitimate communications.

[56] describes a design space of worm containment systems using three parameters: reaction time, containment strategy, and deployment scenario. The authors use a combination of analytic modeling and simulation to describe how these design factors impact the dynamics of a worm epidemic. Their analysis suggests that there are significant gaps in containment defense mechanisms that can be employed, and that considerable more research, and better coordination between ISPs and other entities, is needed when employing network-based defenses. In [91], the authors describe a mechanism for pushing to workstations vulnerability-specific, application-aware filters expressed as programs in a simple language. These programs roughly mirror the state of the protected service, allowing for more intelligent application of content filters, as opposed to simplistic payload string matching.

The HACQIT architecture [67] uses various sensors to detect new types of attacks against secure servers, access to which is limited to small numbers of users at a time. Any deviation from expected or known behavior results in the possibly subverted server to be taken off-line. Similar to our approach, a sandboxed instance of the server is used to conduct “clean room” analysis, comparing the outputs from two different implementations of the service (in their prototype, the Microsoft IIS and Apache web servers were used to provide application diversity). Machine-learning techniques are used to generalize attack features from observed instances of the attack. Content-based filtering is then used, either at the firewall or the end host, to block inputs that may have resulted in attacks, and the infected servers are restarted. Due to the feature-generalization approach, trivial variants of the attack will also be caught by the filter. [86] takes a roughly similar approach, although filtering is done based on port numbers, which can affect service availability. Cisco’s Network-Based Application Recognition (NBAR) allows routers to block TCP sessions based on the presence of specific strings in the TCP stream. This feature was used to block Code-Red probes, without affecting regular web-server access.

HoneyStat [23] runs sacrificial services inside a virtual machine, and monitors memory, disk, and network events to detect abnormal behavior. For some classes of attacks (e.g., buffer overflows), this can produce highly accurate alerts with relatively few false positives, and can detect zero-day worms. Although the system only protects against scanning worms, “active honeypot” techniques [96] may be used to make it more difficult for an automated attacker to differen-

tiate between HoneyStats and real servers. The Internet Motion Sensor [5] is a distributed blackhole monitoring system aimed at measuring, characterizing, and tracking Internet-based threats, including worms.

Lin, Ricciardi, and Marzullo study how computer worms affect the availability of services. In [50], they study the fault tolerance of multicast protocols under self-propagating virus attacks.

5 Conclusion

The main contribution of this paper is the introduction and validation of the *execution transaction* hypothesis, which states that every function execution can be treated as a transaction (similar to a sequence of database operations) that can be allowed to fail, or forced to abort, without affecting the graceful termination of the computation. We provide some preliminary evidence on the validity of this hypothesis by examining a number of open-source software packages with known vulnerabilities.

For that purpose, we developed DYBOC, a tool for instrumenting C source code such that buffer overflow attacks can be caught, and program execution continue without any adverse side effects (such as forced program termination). DYBOC allows a system to dynamically enable or disable specific protection checks in running software, potentially as a result of input from external sources (e.g., an IDS engine), at an very high level of granularity. This enables the system to implement policies that trade off between performance and risk, retaining the capability to re-evaluate this trade-off very quickly. This makes DYBOC-enhanced services highly responsive to automated indiscriminate attacks, such as scanning worms. Finally, our preliminary performance experiments indicate that: (a) the performance impact of DYBOC in contrived examples can be significant, but (b) the impact in performance is significantly lessened (less than 2%) in real applications, and (c) this performance impact is further lessened by utilizing the dynamic nature of our scheme.

Our plans for future work include enhancing the capabilities of DYBOC by combining it with a static source-code analysis tool, extending the performance evaluation, and further validating our hypothesis by examining a larger number of applications.

References

1. Acharya, A., Raje, M.: Mapbox: Using parameterized behavior classes to confine applications. In: Proceedings of the 9th USENIX Security Symposium, pp. 1–17 (2000)
2. Aleph One: Smashing the stack for fun and profit. Phrack 7 (1996)
3. Alexandrov, A., Kmiec, P., Schauser, K.: Consh: A confined execution environment for Internet computations (1998)
4. Arce, I.: The Shellcode generation. IEEE Secur. Priv. 2, 72–76 (2004)
5. Bailey, M., Cooke, E., Jahanian, F., Nazario, J., Watson, D.: The Internet motion sensor: a distributed blackhole monitoring system. In: Proceedings of the ISOC Symposium on Network and Distributed Systems Security (SNDSS), pp. 167–179 (2005)
6. Balzer, R., Goldman, N.: Mediating connectors: A non-bypassable process wrapping technology. In: Proceedings of the 19th IEEE International Conference on Distributed Computing Systems (1999)
7. Baratloo, A., Singh, N., Tsai, T.: Transparent run-time defense against stack smashing attacks. In: Proceedings of the USENIX Annual Technical Conference (2000)
8. Berman, A., Bourassa, V., Selberg, E.: TRON: Process-specific file protection for the UNIX operating system. In: Proceedings of the USENIX Technical Conference (1995)
9. Brumley, D., Song, D.: Privtrans: Automatically partitioning programs for privilege separation. In: Proceedings of the 13th USENIX Security Symposium, pp. 57–71 (2004)
10. Bulba, K.: Bypassing StackGuard and StackShield. Phrack 5 (2000)
11. Candea, G., Fox, A.: Crash-only software. In: Proceedings of the 9th Workshop on Hot Topics in Operating Systems (2003)
12. Candea, G., Fox, A.: Recursive restartability: Turning the reboot sledgehammer into a scalpel. In: Proceedings of the 8th Workshop on Hot Topics in Operating Systems (HotOS-VIII), pp. 110–115 Schloss Elmau, Germany, IEEE Computer Society, (2001)
13. CERT Advisory CA-2001-19: ‘Code Red’ Worm Exploiting Buffer Overflow in IIS Indexing Service DLL. <http://www.cert.org/advisories/CA-2001-19.html> (2001)
14. Cert Advisory CA-2003-04: MS-SQL Server Worm. <http://www.cert.org/advisories/CA-2003-04.html> (2003)
15. Chen, H., Dean, D., Wagner, D.: Model checking one million lines of C code. In: Proceedings of the Network and Distributed System Security (NDSS) Symposium, pp. 171–185 (2004)
16. Chen, H., Wagner, D.: MOPS: an infrastructure for examining security properties of software. In: Proceedings of the ACM Computer and Communications Security (CCS) Conference, pp. 235–244 (2002)
17. Chess, B., McGraw, G.: Static analysis for security. IEEE Secur. Priv. 2, 76–79 (2004)
18. Christodorescu, M., Jha, S.: Static analysis of executables to detect malicious patterns. In: Proceedings of the 12th USENIX Security Symposium, pp. 169–186 (2003)
19. Cohen, F.: Computer viruses: theory and practice. Comput. Secur. 6, 22–35 (1987)
20. Costa, M., Crowcroft, J., Castro, M., Rowstron, A., Zhou, L., Zhang, L., Barham, P.: Vigilante: end-to-end containment of Internet worms. In: SOSP ’05: Proceedings of the Twentieth ACM symposium on Operating systems principles, vol. 39, pp. 133–147. ACM, New York (2005)
21. Cowan, C., Beattie, S., Pu, C., Wagle, P., Gligor, V.: SubDomain: parsimonious security for server appliances. In: Proceedings of the 14th USENIX System Administration Conference (LISA 2000) (2000)
22. Cowan, C., Pu, C., Maier, D., Hinton, H., Walpole, J., Bakke, P., Beattie, S., Grier, A., Wagle, P., Zhang, Q.: StackGuard: automatic adaptive detection and prevention of buffer-overflow attacks. In: Proceedings of the 7th USENIX Security Symposium (1998)
23. Dagon, D., Qin, X., Gu, G., Lee, W., Grizzard, J., Levine, J., Owen, H.: HoneyStat: Local worm detection using Honeypots. In: Proceedings of the 7th International Symposium on Recent Advances in Intrusion Detection (RAID), pp. 39–58 (2004)
24. Demsky, B., Rinard, M.C.: Automatic detection and repair of errors in data structures. In: Proceedings of the 18th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Application (OOPSLA) (2003)
25. Demsky, B., Rinard, M.C.: Automatic detection and repair of errors in data structures. In: Proceedings of the 18th Annual ACM SIGPLAN Conference on Object Oriented Programming, Systems, Languages, and Applications (2003)

26. Dor, N., Rodeh, M., Sagiv, M.: CSSV: Towards a realistic tool for statically detecting all buffer overflows in C. In: Proceedings of the ACM Conference on Programming Language Design and Implementation (PLDI) (2003)
27. Engler, D., Ashcraft, K.: RacerX: Effective, static detection of race conditions and deadlocks. In: Proceedings of ACM SOSP (2003)
28. Etoh, J.: GCC extension for protecting applications from stack-smashing attacks. <http://www.trl.ibm.com/projects/security/ssp/> (2000)
29. Foster, J., Fahndrich, M., Aiken, A.: A theory of type qualifiers. In: Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI) (1999)
30. Fraser, T., Badger, L., Feldman, M.: Hardening COTS software with generic software wrappers. In: Proceedings of the IEEE Symposium on Security and Privacy. Oakland, CA (1999)
31. Garber, L.: New chips stop buffer overflow attacks. *IEEE Comput.* **37**, 28 (2004)
32. Garfinkel, T., Pfaff, B., Rosenblum, M.: Ostia: A delegating architecture for secure system call interposition. In: Proceedings of the Network and Distributed System Security (NDSS) Symposium, pp. 187–201 (2004)
33. Garfinkel, T.: Traps and pitfalls: Practical problems in system call interposition based security tools. In: Proceedings of the Symposium on Network and Distributed Systems Security (SNDSS), pp. 163–176 (2003)
34. Ghormley, D.P., Petrou, D., Rodrigues, S.H., Anderson, T.E.: SLIC: an extensibility system for commodity operating systems. In: Proceedings of the 1998 USENIX Annual Technical Conference, pp. 39–52 (1998)
35. Goldberg, I., Wagner, D., Thomas, R., Brewer, E.A.: A secure environment for untrusted helper applications. In: Proceedings of the 1996 USENIX Annual Technical Conference (1996)
36. Grünbacher, A.: POSIX access control lists on Linux. In: Proceedings of the USENIX Technical Conference, Freenix Track, pp. 259–272 (2003)
37. Hernacki, B., Bennett, J., Lofgren, T.: Symantec deception server experience with a commercial deception system. In: Proceedings of the 7th International Symposium on Recent Advances in Intrusion Detection (RAID), pp. 188–202 (2004)
38. Ioannidis, S., Bellovin, S., Smith, J.M.: Sub-operating systems: a new approach to application security. In: Proceedings 10th SIGOPS European Workshop, pp. 108–115 (2002)
39. Ioannidis, S., Bellovin, S.M.: Building a secure browser. In: Proceedings of the Annual USENIX Technical Conference, Freenix Track, pp. 127–134 (2001)
40. Cordy, J.R., Dean, T.R., Schneider, K.: Source transformation in software engineering using the TXL transformation system. *J. Inf. Software Technol.* **44**, 827–837 (2002)
41. Johnson, R., Wagner, D.: Finding user/kernel pointer bugs with type inference. In: Proceedings for the 13th USENIX Security Symposium, pp. 119–134 (2004)
42. Jones, R.W.M., Kelly, P.H.J.: Backwards-compatible bounds checking for arrays and pointers in C programs. In: Third International Workshop on Automated Debugging (1997)
43. Kephart, J.O.: A biologically inspired immune system for computers. In: Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems, pp. 130–139. MIT Press, Cambridge, MA (1994)
44. Kilpatrick, D.: Privman: A library for partitioning applications. In: Proceedings of the USENIX Technical Conference, Freenix Track, pp. 273–284 (2003)
45. Kiriansky, V., Bruening, D., Amarasinghe, S.: Secure execution via program shepherding. In: Proceedings of the 11th USENIX Security Symposium, pp. 191–205 (2002)
46. Larochelle, D., Evans, D.: Statically detecting likely buffer overflow vulnerabilities. In: Proceedings of the 10th USENIX Security Symposium, pp. 177–190 (2001)
47. Larson, E., Austin, T.: High coverage detection of input-related security faults. In: Proceedings of the 12th USENIX Security Symposium, pp. 121–136 (2003)
48. Levy, E.: Crossover: Online pests plaguing the offline world. *IEEE Secur. Priv.* **1**, 71–73 (2003)
49. Lhee, K., Chapin, S.J.: Type-assisted dynamic buffer overflow detection. In: Proceedings of the 11th USENIX Security Symposium, pp. 81–90 (2002)
50. Lin, M.J., Ricciardi, A., Marzullo, K.: A new model for availability in the face of self-propagating attacks. In: Proceedings of the New Security Paradigms Workshop (1998)
51. Liston, T.: Welcome To My Tarpit: The Tactical and Strategic Use of LaBrea. <http://www.threenorth.com/LaBrea/LaBrea.txt> (2001)
52. Loscocco, P., Smalley, S.: Integrating flexible support for security policies into the Linux operating system. In: Proceedings of the USENIX Annual Technical Conference, Freenix Track, pp. 29–40 (2001)
53. Vojnovic, M., Ganesh, A.: On the effectiveness of automatic patching. In: WORM '05 (2005)
54. Mitchem, T., Lu, R., O'Brien, R.: Using kernel hypervisors to secure applications. In: Proceedings of the Annual Computer Security Applications Conference (1997)
55. Moore, D., Shanning, C., Claffy, K.: Code-Red: a case study on the spread and victims of an Internet worm. In: Proceedings of the 2nd Internet Measurement Workshop (IMW), pp. 273–284 (2002)
56. Moore, D., Shannon, C., Voelker, G., Savage, S.: Internet quarantine: requirements for containing self-propagating code. In: Proceedings of the IEEE Infocom Conference (2003)
57. Necula, G.C., McPeak, S., Weimer, W.: CCured: type-safe retrofitting of legacy code. In: Proceedings of the Principles of Programming Languages (PoPL) (2002)
58. Newsome, J., Song, D.: Dynamic taint analysis for automatic detection, analysis, and signature generation of exploits on commodity software. In: The 12th Annual Network and Distributed System Security Symposium (2005)
59. Nicol, D., Liljenstam, M.: Models of active worm defenses. In: Proceedings of the Measurement, Modeling and Analysis of the Internet Workshop (IMA) (2004)
60. Pasupulati, A., Coit, J., Levitt, K., Wu, S., Li, S., Kuo, J., Fan, K.: Buttercup: On network-based detection of polymorphic buffer overflow vulnerabilities. In: Proceedings of the Network Operations and Management Symposium (NOMS), vol. 1, pp. 235–248 (2004)
61. Peterson, D.S., Bishop, M., Pandey, R.: A flexible containment mechanism for executing untrusted code. In: Proceedings of the 11th USENIX Security Symposium, pp. 207–225 (2002)
62. Pincus, J., Baker, B.: Beyond stack smashing: recent advances in exploiting buffer overflows. *IEEE Secur. Priv.* **2**, 20–27 (2004)
63. Prevelakis, V., Spinellis, D.: Sandboxing applications. In: Proceedings of the USENIX Technical Annual Conference, Freenix Track, pp. 119–126 (2001)
64. Provos, N., Friedl, M., Honeyman, P.: Preventing privilege escalation. In: Proceedings of the 12th USENIX Security Symposium, pp. 231–242 (2003)
65. Provos, N.: A virtual Honeypot framework. In: Proceedings of the 13th USENIX Security Symposium, pp. 1–14 (2004)
66. Provos, N.: Improving host security with system call policies. In: Proceedings of the 12th USENIX Security Symposium, pp. 257–272 (2003)
67. Reynolds, J.C., Just, J., Clough, L., Maglich, R.: On-line intrusion detection and attack prevention using diversity, generate-and-test, and generalization. In: Proceedings of the 36th Annual Hawaii International Conference on System Sciences (HICSS) (2003)
68. Rinard, M., Cadar, C., Dumitran, D., Roy, D., Leu, T., Beebe, J.: Enhancing server availability and security through failure-oblivious computing. In: Proceedings 6th Symposium on Operating Systems Design and Implementation (OSDI) (2004)
69. Rinard, M., Cadar, C., Dumitran, D., Roy, D., Leu, T.: A dynamic technique for eliminating buffer overflow vulnerabilities (and other memory errors). In: Proceedings 20th Annual Computer Security Applications Conference (ACSAC) (2004)

70. Rudys, A., Wallach, D.S.: Termination in language-based systems. *ACM Trans. Inf. Syst. Secur.* **5** (2002)
71. Rudys, A., Wallach, D.S.: Transactional rollback for language-based systems. In: *ISOC Symposium on Network and Distributed Systems Security (SNDSS)* (2001)
72. Ruwase, O., Lam, M.S.: A practical dynamic buffer overflow detector. In: *Proceedings of the Network and Distributed System Security (NDSS) Symposium*, pp. 159–169 (2004)
73. Sekar, R., Venkatakrisnan, V.N., Basu, S., Bhatkar, S., DuVarney, D.C.: Model-carrying code: A practical approach for safe execution of untrusted applications. In: *Proceedings of ACM SOSP* (2003)
74. Shankar, U., Talwar, K., Foster, J.S., Wagner, D.: Detecting format string vulnerabilities with type qualifiers. In: *Proceedings of the 10th USENIX Security Symposium*, pp. 201–216 (2001)
75. Shannon, C., Moore, D.: The spread of the Witty worm. *IEEE Secur. Priv.* **2**, 46–50 (2004)
76. Sidiroglou, S., Keromytis, A.D.: A network worm vaccine architecture. In: *Proceedings of the IEEE Workshop on Enterprise Technologies: Infrastructure for Collaborative Enterprises (WET-ICE), Workshop on Enterprise Security*, pp. 220–225 (2003)
77. Sidiroglou, S., Keromytis, A.D.: Countering network worms through automatic patch generation. *IEEE Secur. Priv.* (2005) (to appear).
78. Smirnov, A., Chiueh, T.: DIRA: Automatic detection, identification, and repair of control-hijacking attacks. In: *Proceedings of the 12th ISOC Symposium on Network and Distributed System Security (SNDSS)* (2005)
79. Song, D., Malan, R., Stone, R.: A Snapshot of Global Internet Worm Activity. Technical report, Arbor Networks (2001)
80. Staniford, S., Paxson, V., Weaver, N.: How to own the Internet in your spare time. In: *Proceedings of the 11th USENIX Security Symposium*, pp. 149–167 (2002)
81. Steven, A.: Defeating compiler-level buffer overflow protection. *USENIX ;login* **30**, 59–71 (2005)
82. Suh, G.E., Lee, J.W., Zhang, D., Devadas, S.: Secure program execution via dynamic information flow tracking. *SIGOPS Oper. Syst. Rev.* **38**, 85–96 (2004)
83. Sun, W., Liang, Z., Sekar, R., Venkatakrisnan, V.N.: One-way isolation: an effective approach for realizing safe execution environments. In: *Proceedings of the 12th ISOC Symposium on Network and Distributed Systems Security (SNDSS)*, pp. 265–278 (2005)
84. Ször, P., Ferrie, P.: Hunting for Metamorphic. Technical report, Symantec (2003)
85. The Spread of the Sapphire/Slammer Worm. <http://www.silicondefense.com/research/worms/slammer.php> (2003)
86. Toth, T., Kruegel, C.: Connection-history based anomaly detection. In: *Proceedings of the IEEE Workshop on Information Assurance and Security* (2002)
87. Twycross, J., Williamson, M.M.: Implementing and testing a virus throttle. In: *Proceedings of the 12th USENIX Security Symposium*, pp. 285–294 (2003)
88. Vindicator: Stack shield. <http://www.angelfire.com/sk/stackshield/> (2000)
89. Wagner, D., Foster, J.S., Brewer, E.A., Aiken, A.: A first step towards automated detection of buffer overrun vulnerabilities. In: *Network and Distributed System Security Symposium*, pp. 3–17 (2000)
90. Walker, K.M., Stern, D.F., Badger, L., Oosendorp, K.A., Petkac, M.J., Sherman, D.L.: Confining root programs with domain and type enforcement. In: *Proceedings of the 1996 USENIX Security Symposium*, pp. 21–36 (1996)
91. Wang, H.J., Guo, C., Simon, D.R., Zugenmaier, A.: Shield: vulnerability-driven network filters for preventing known vulnerability exploits. In: *Proceedings of the ACM SIGCOMM Conference* (2004)
92. Wang, N., Fertig, M., Patel, S.: Y-Branches: When you come to a fork in the road, take it. In: *Proceedings of the 12th International Conference on Parallel Architectures and Compilation Techniques* (2003)
93. Watson, R.N.M.: TrustedBSD: adding trusted operating system features to FreeBSD. In: *Proceedings of the USENIX Annual Technical Conference, Freenix Track*, pp. 15–28 (2001)
94. Weaver, N., Staniford, S., Paxson, V.: Very fast containment of scanning worms. In: *Proceedings of the 13th USENIX Security Symposium*, pp. 29–44 (2004)
95. Wong, C., Wang, C., Song, D., Bielski, S., Ganger, G.R.: Dynamic quarantine of Internet worms. In: *DSN '04: Proceedings of the 2004 International Conference on Dependable Systems and Networks (DSN'04)*, p. 73. Washington, DC, IEEE Computer Society, (2004)
96. Yegneswaran, V., Barford, P., Plonka, D.: On the design and use of Internet sinks for network abuse monitoring. In: *Proceedings of the 7th International Symposium on Recent Advances in Intrusion Detection (RAID)*, pp. 146–165 (2004)
97. Zou, C.C., Gao, L., Gong, W., Towsley, D.: Monitoring and early warning for Internet worms. In: *Proceedings of the 10th ACM International Conference on Computer and Communications Security (CCS)*, pp. 190–199 (2003)
98. Zou, C.C., Gong, W., Towsley, D.: Code Red worm propagation modeling and analysis. In: *Proceedings of the 9th ACM Conference on Computer and Communications Security (CCS)*, pp. 138–147 (2002)