

# Software Decoys: Intrusion Detection and Countermeasures

James Bret Michael, *Senior Member, IEEE*, Mikhail Auguston, Neil C. Rowe, and Richard D. Riehle

**Abstract**—We introduce the notion of an intelligent software decoy, and provide both an architecture and event-based language for automatic implementation of them. Our decoys detect and respond to patterns of suspicious behavior, and maintain a repository of rules for behavior patterns and decoying actions. As an example, we construct a model of system behavior from an initial list of event types and their attributes in the interaction between computer worms and an operating system. The model represents patterns of suspicious or malicious events that the software decoy should detect, and specific actions to be taken in response. Our approach explicitly treats both standard and nonstandard invocations of components, with the latter representing an attempt to circumvent the public interface of the component.

**Index terms**—Behavior modeling, computer security, computer worm, deception, event trace, software decoy, intrusion detection, information operations, intrusion tolerance

## I. INTRODUCTION

Software components with poorly designed interfaces are susceptible to misuse or modification by rogue programs. Consider, for example, the result of unleashing the Morris Internet worm (*vid.* Spafford [1]): Some integral components of the UNIX operating system permitted the worm to propagate itself over the Internet. What were some of shortcomings of some versions of Unix in 1988? One was that some components permitted incorrect argument types to be executed, such as the *sendmail* program accepting commands instead of user addresses. Another weakness was that some components allowed for erroneous argument values to be passed to functions, such as a string of 536 bytes passed to the input buffer of the *fingerd* program: That string exceeds the size of the buffer, but *fingerd* and the functions it calls in the C language I/O library do not check, resulting in a buffer overflow.

Software patches were made to *sendmail* and *fingerd* in response to the Morris worm. The patches consisted of

exception-handling routines for catching errors in input arguments to these programs. The effectiveness of the patches was low at first because of delays between the release and installation.

History continues to repeat itself. Distributed systems are plagued by worms descended from the Morris worm, such as the well-publicized “Code Red” worm [2] that exploits weaknesses in the public interfaces of components of the Microsoft Windows NT operating system that manipulate the registry.<sup>1</sup> The “patch-and-pray” approach to dealing with such weaknesses in interfaces has not been effective. Shortly after the introduction of the original, a variant of the Code Red worm appeared which took advantage of other weaknesses in the interfaces of those same components of the operating system [3].

In this paper we explore a new way of thinking about the challenges associated with protecting components within distributed systems from the effects of attacks. In particular, we develop the mechanisms needed to realize a modified version of what Michael and Riehle call an *intelligent software decoy* [4]. Intelligent software decoys adapt in order to reveal and tolerate both intrusions into systems by unauthorized users and misuse of software components, instead of either indicating to the intruder or offending process that a violation of security policy has been detected or terminating the interaction with the process. Intelligent software decoys perform what could be termed cyber *Akido*, in which the decoy tolerates an attack and learns about the opponent’s weaknesses (*e.g.*, that the opponent relies on the use of buffer-overflow techniques) while simultaneously trying to neutralize the opponent by, in the following order, reducing or eliminating the will of the attacker (*e.g.*, by inserting delays into responses to method calls made by the opponent), changing the proximity of attack (*e.g.*, directing the attention of the opponent to a honeypot), and as a last resort reducing or eliminating the ability of the opponent to attack (*e.g.*, closing a communication port, killing a process, or launching a counter denial-of-service attack), at which point the effectiveness of deception techniques will likely be reduced because the opponent becomes aware or

---

Manuscript submitted on December 28, 2001; accepted for publication on April 18, 2002.

J. B. Michael, N. C. Rowe, and R. D. Riehle are with the Department of Computer Science, Naval Postgraduate School, 833 Dyer Rd., Monterey, Calif. 93943-5118 (e-mail: {bmichael, ncrowe, rdriehle}@nps.navy.mil).

M. Auguston is with the Department of Computer Science, New Mexico State University, Las Cruces, N.M. 88003-0001 (e-mail: mikau@cs.nmsu.edu).

---

<sup>1</sup> The registry is the repository of the local security policy used by the security reference manager and monitor to control access by threads to objects.

suspicious that the software decoy knows of the presence or weaknesses of the attacker.

In contrast, most intrusion-detection and response systems apply a form of cyber *Karate*: “if you try to kill me, I will try to kill you” by immediately terminating interaction with the attacker and thus indicating to the attacker that the attack was noticed and that countermeasures are already being applied. A counterexample is that of pH [5] which tries to initially tolerate potentially malicious interaction between a calling process and a wrapped component of the operating system by inserting delays into the responses to such calls so as to provide the intrusion-detection system with as much time as possible to distinguish between real attacks and false positives. The termination of a process based on a false positive would result in a legitimate process being denied service.

Some key ways in which intelligent software decoys differ from *honeypots* are that software decoys actively defend components rather than just collect data, are not intended to be used solely to attract the attention of adversaries, and do not rely on the concept of a perimeter of defenses (*i.e.*, any component within a distributed system can be decoy-enabled and either operate autonomously or in a cooperative mode with other software decoys). Moreover, we envision software decoys as being integrated into a larger joint (across the Armed Services) war-fighting and enforcement arsenal used to shape the cyber environment in real-time. Although software decoys could be used to help thwart the actions of “script kiddies” and hackers, the motivation for work reported here is to provide enabling technology for the warfighter to conduct automated information operations against sophisticated terrorist cells and nation states; deceiving, deterring, and dissuading such adversaries requires embedding into systems the ability for software components to adapt, through learning, their strategies for conducting information operations as feedback is received by those components on the effectiveness of such operations.

Each decoy-enabled software component has a *contract* consisting of a class invariant and one or more preconditions and postconditions; communication with a component is only permitted via its contract interface. Class invariants govern the nature and extent of any change to objects within a component by a nonstandard invocation (*i.e.*, circumvention) of the component’s public interface.

When a calling process passes arguments—via local or remote procedure call, or remote method invocation—that violate the contract of a component, the component transitions from its nominal operating mode to a deception mode, in which it attempts to both deceive the calling process into concluding that its violation of the contract has been successful, while simultaneously both containing the interaction and assessing the nature of the violation.

For example, suppose a precondition in a contract asserts that input strings shall not exceed 512 bytes. If the component receives a string that violates this condition, then the component could respond by simulating a buffer overflow with the aim of keeping the process engaged in interacting with the component.

Here we describe decoy strategy along with an architecture and event-based language for automatic implementation of software decoys.

## II. APPROACH

We start with the introduction of a precise *behavior model* for the system under consideration. This model is specified in terms of events and two binary relations over those events: precedence and inclusion.

We provide a formalism to specify rules for runtime intrusion detection and corresponding countermeasures based on behavior patterns over event traces and a catalog of decoy actions, such as blocking or substituting certain system commands. This implies an implementation based on automatic instrumentation for event detection derived from the behavior model. Intrusion detection rules are textually separated from the source code of the system, which allows for accumulating and formalizing knowledge of typical intrusion patterns and decoy strategies.

### A. Event Traces

An *event* is an abstraction of any detectable action performed at runtime. An event has a beginning, end, duration, and some other attributes, such as program states at the beginning and end of the event, source code associated with the event, and so on. Two binary relations are defined for events. One event may precede another event. For example, one statement execution may precede another. In addition, one event may be included in another. For instance, a statement-execution event may appear inside a procedure-call or method-invocation event. Each of these binary relations defines a partial ordering of events. System execution may be represented as a set of events with the two basic relations between them—an *event trace*. An *event grammar* [6] is a set of axioms that determines possible configurations of events of different types within the event trace. We use the notion of computations over event traces to specify behavior patterns and decoy actions. This mechanism is a basis for automatic instrumentation of the source code.

### B. Specification of Events

The behavioral model consists of the definition of event types and the attributes of the events. Our event grammar is used to describe the structure of events. The behavioral model consists of two parts: (i) a specification of *axioms*

that define constraints on the behavior of components; this specification is given in terms of inclusion and precedence relations, and (ii) a description of *patterns of behavior* expressed in terms of event patterns.

The event grammar is not intended for actual parsing of an event trace. Each event defined in the behavior model should be detectable by some independent means, for example, by proper instrumentation of the source code. There always is a main event `execute-program`, which encompasses all other events. Event type and event attribute declarations, and event grammar rules (*i.e.*, axioms) constitute the behavior model, or “lightweight” semantics specification for the system under consideration. The following is an example of an event-grammar rule that specifies that an event of type `execute-assignment` always contain (*i.e.*, the inclusion relation) two ordered events: `evaluate-right-hand-part` and `perform-destination`.

```
execute-assignment::
    (evaluate-right-hand-part
     perform-destination)
```

Intrusion detection and other monitoring activities are defined in terms of this behavior model. This model opens the way for automatic source program instrumentation, and provides a formalism for describing different kinds of behavior patterns, for example, typical intrusion patterns.

### C. Computations over Event Traces

*Event pattern A* matches successfully any event of the type *A*. The event pattern can contain additional context conditions, which typically involve event attributes, for instance, `X: perform-destination & source-code(X) = 'V'`. This event pattern matches an event of type `perform-destination`, such that source-code attribute of this event is 'V'. Notice the use of auxiliary variable *X* to denote the event under consideration.

Event patterns are used in *aggregate operations* to select sequences of events from other event sets. For example, `[X:utility-call & name(X) = 'sendmail' FROM execute-program]` yields a sequence of events matching the pattern `X: utility-call & name(X) = 'sendmail'` selected from the whole event trace. An aggregate operation may produce a result different from just a sequence of events. For instance, `SUM/[X:utility-call & name(X) = 'sendmail' FROM execute-program APPLY duration(X)]` yields a total duration of events selected.

Event patterns can be combined in order to specify *path expressions*: patterns of event sequences. For instance,

the following path expression specifies a sequence of events starting with event of type *A* followed by zero or more events of types *B* or *C*:  $A (B \mid C)^*$

A *probe* is a Boolean expression involving event attributes and standard arithmetic and logic operations. It might also contain calls of subroutines written in some general programming language. Probes may be included in path expressions interlaced with event patterns. A probe is evaluated immediately after the preceding event pattern has been matched successfully. A probe is successful if it evaluates to the value True. For example, the following path expression specifies a utility call that does not change the value of variable *V*.

```
X: utility-call
  Probe
    (value(V at begin X) =
     value (V at end X))
```

These operations provide a basis for computational tasks such as filtering events to create views of the event trace subspace, evaluating the truth of assertions, and computing specific values (*e.g.*, counting the number of events in a sequence or the maximum elapsed time for an event pattern). An important feature of this formalism is that computations over event traces are separated from the source code itself, and the instrumentation of the source code can be automated, based on the behavior model (*i.e.*, event grammar) and the text of the trace computation.

### D. Code Instrumentation

Code instrumentation is needed to recognize patterns of events and to perform computations over the event trace. Our approach is to selectively specify and automatically instrument, on a component-by-component basis, only those event types and event attributes that are of interest. For example, one would likely specify events corresponding to popping (*i.e.*, deleting an item from) the stack, reading, writing, or handling threads (*e.g.*, create, reschedule, delete) as events of interest in the context of the behavior of components comprising an operating system; only the source code corresponding to these and other events of interest would be instrumented. For other components within the operating system, there may not be any events of interest and therefore the source code of those components would not be instrumented.

An automated tool for generating the instrumentation would accept as input the program's source code and the text of trace computation. We define the event pattern language and aggregate operations in such a way that event detection, event pattern matching, and other computations over event traces can be performed at runtime. Based on similar levels of instrumentation as reported by Sekar and Uppuluri [7] and Ko *et al.* [8], we anticipate runtime overhead associated with optimized instrumenta-

tion to fall within the range of four to seven percent. This effect of event-trace computations is important for high-performance systems, especially for components that are called frequently, such as those comprising either the scheduler or memory manager of an operating system.

As an example of how instrumentation would be selectively applied, consider a component that implements the Transmission Control Protocol (TCP); TCP is a frequently called component and known to be susceptible to certain types of attacks, such as that of the "Cheese" worm [9] that attempts to execute commands on a specific TCP port of a system and then masquerade as the *httpd* program on that system. The change to the *httpd* program via a Trojan horse is an example of an attack for which a class invariant could be used to guard against unauthorized changes to the behavior of this component. One of the events of interest would be the invocation of shell commands on a TCP port. Due to the fact that there is a sequence of calls resulting from the shell commands, the events at other components may need to be instrumented; this is the case for the Cheese worm. Some of the events of interest would be changing process names, overwriting system files (e.g., */etc/inetd.conf*), and repeated attempts to restart a system command (e.g., *inetd*). The compiler used for generating instrumentation would only do so for these types of events. Specific examples of how to specify events, their attributes, and actions (i.e., responses to events) are given in the remainder of the paper for the Morris worm.

### III. DOMAIN MODELS OF SYSTEM BEHAVIOR

*Domain models* of system behavior serve as the foundation from which to specify both patterns of suspicious behavior (i.e., slices of event histories) and the corresponding responses (i.e., actions) of software decoys. A domain model consists of a list of generic event types, along with their attributes, both of which are associated with a particular class of behavior and detectable. In this section, we report our work to date to develop a domain model of detectable events that can be triggered within an operating system by computer worms: that is, rogue programs that propagate themselves across computer networks. In other words, we treat the interaction between a worm and a particular type of software component as a unique domain for modeling purposes. Our motivation here is to both describe an event-trace language for implementing software decoys and give specific examples of how one would instrument code for runtime checking. Our examples are based on the behaviors of components of the Unix operating system with the Morris worm. More recent worms, such as "Code Red" and "Cheese," are variants of the Morris worm in at least two ways: (i) they are composed from the same set of generic event types, or "building blocks," and (ii) share the same basic

structure, in some cases initiating similar patterns of events.

#### A. Interaction between the Worm and Operating System

Although it is still debated as to whether the intent of the author of the worm was to just experiment out of academic interest and the worm "got loose" due to an egregious use of contracts at the interfaces of the components it called, the interaction of the worm with software components resulted in a temporary denial of service across much of the Internet.

The Morris worm interacts with the components of the operating system by repeatedly accessing a data structure (e.g., the system password and network configuration files), calling a program to obtain state information (e.g., *netstat*), or repeatedly invoking a program and permuting the arguments before each invocation (e.g., *rsh*). The worm also invokes sequences of functions and procedures, such as issuing the *DEBUG* command to the *sendmail* program followed by a sequence of commands (i.e., the vector portion of the worm program) to be acted on by procedures available in that operating mode.

In addition to accessing components, the Morris worm attempts to write to files (e.g., set a flag) and execute copies of itself in the form of Trojan horses. It also attempts to propagate itself while simultaneously eluding detection or spoofing by forking itself after a pre-specified number of tries at infecting other computers or elapse of time, and then deleting both the parent process and all of the temporary files created by the parent process.

#### B. Description of the Domain Model

Here we provide a first approximation of the description of the model. We begin by specifying five event types and their attributes:

```
fingerd_call
  Attributes:  caller_id
               begin_time

param_pass
  Attributes:  length

read
utility_call
  Attributes:  begin_time
               caller_id
               caller_process_type

sendmail_call
  Attributes:  in_debug_mode
```

Next we introduce an axiom that specifies the event *fingerd\_call* contains another event, *param\_pass*:

```
Axiom: fingerd_call:: ( param_pass )
```

Based on the foregoing specification of events and axioms, we can start to build a behavioral specification for the Morris worm as follows:

```
fingerd_call:: ( x: param_pass )  
  & length(x) > max_buffer_size  
  read +  
  probe( buffer_overflow )
```

This behavior pattern, described in terms of a violation of the contract for *fingerd* calls, states that first an event of the type *fingerd\_call* is detected, such that the value of attribute *length* of the included event *param\_pass* exceeds a given constant *max\_buffer\_size*; in this case, *x* is a tentative name associated with the event *param\_pass*. Next, one or more events of the type *read* should appear, and finally a probe (*i.e.*, an evaluation of a Boolean expression) specifying buffer overflow should evaluate to the value *True*. Successful detection of those events and conditions in the specified order indicates that the system might be under attack by the Morris worm or some other worm. Now let us specify yet another behavior pattern:

```
x: sendmail_call::  
  ( [ utility_call + ]  
    && CONST(caller_id)  
    && CONST(caller_process_type)  
    && FREQUENCY(time_interval)  
  )  
  & in_debug_mode(x) = True
```

Here *sendmail\_call* is specified with the attribute *in\_debug\_mode* equal to the value *True*. This event contains a sequence of *utility\_call* events with the same *caller\_id* and *caller\_process\_type* attribute values, such that the time between those *utility\_call* events does not exceed *time\_interval*.

### C. Summary

In summary, our event-trace language integrates the language constructs that are necessary for specifying event patterns, probes, and actions. An event pattern is a particular interval-based signature constructed from a structural model of events of interest (*e.g.*, system calls to write to a file). The probe is a Boolean expression; for instance, a probe can be used to compare the values of two attributes over the same interval of time, or the value of the same attribute at different points in time. An action is a message that is generated based on the recognition that a particular pattern of behavior has been observed. The pattern of behavior can be composed from one or more slices of the same or different event traces. The messages are used to trigger rules that embody the target component's decoy-mode response to its interaction with

the calling process whose interaction with the component may be malicious or benign.

## IV. DECOY STRATEGY

Decoy methods can show a wide spectrum of complexity. A key dimension is information in the mathematical sense they provide to the attacker [10]. One might ask: Why provide any information to one's adversary? This is the nature of deception, such as providing misinformation or misleading information to one's adversary. Our decoy strategy borrows from military tactics for deceiving one's adversary (*e.g.*, *viz.* [11]), which is only one of the five pillars of information operations, the others being psychological operations, electronic warfare, computer-network attack, and operational security.

A decoy that provides the same canned response to a user command in all circumstances transmits zero bits of information. A decoy that simulates an arbitrary protected file by generating a fake file with random choices transmits information equal to the logarithm of the number of possible fake files it can create for a given request. A decoy that simulates a denial-of-service attack, providing additional operating-system delays monotonically increasing with the number of requests in a recent time interval, transmits information equal to the logarithm of the number of distinguishable time intervals it can provide. A decoy that simulates attempts to change a module of the operating system transmits information equal to the logarithm of the size of that module (or perhaps the part of it relevant to a particular attacker activity). It might seem that the ultimate decoy would simulate all aspects of a computer and its software as proposed by Thimbley *et al.* [12], requiring information equal to the logarithm of the size of the complete description of the computer and programming environment, and thus be equivalent to a universal Turing machine. However, we claim a truly ultimate decoy would be "intelligent" and need to do more: It would need to analyze properties of arbitrary code of viruses and Trojan horses to predict their conceptual consequences. But such ultimate decoying appears impossible in light of the undecidability of nontrivial properties of recursive functions (*vid.* Rice's Theorem [13]).

Applying these ideas to the Morris worm, the propagation of the worm could be decoyed by having the operating system count the number of recent buffer-overflowing *fingerd* calls (requiring a single integer of memory) and provide additional process-suspension time in all operating-system activities; this could most easily be done by just inserting delays into a frequently used utility like file reading, delays that would only occur for the process threads created by the perceived attacker and would be monotonic with the number of ongoing threads they have created. The *sendmail* trapdoor could be decoyed by per-

mitting normal debugging behavior except for commands writing or executing parts of memory, which would appear to execute but would be prevented from changing memory.

Consider some examples of decoy strategies for the Morris worm. The propagation of the worm could be decoyed by providing additional process-suspension time for all utility calls within the same process, creating an illusion to the intruder that a delay has been introduced (assuming that the intruder monitors response times) while the other non-offending processes will continue to execute unhindered.

```
detect [x: fingerd_call::
  ( y: param_pass )
  & length(y) > max_buffer_size
  read + ] && CONST(process_id)
probe (buffer_overflow)
from execute-program
do enable delay (t) to z: utility_call
  & process_id(x) = process_id(z)
```

The *sendmail* trapdoor could be decoyed by directing its invoker to a decoy function `protected_read()` that simulates the behavior with correct read commands and by deactivating write commands via the decoy function `dummy()` to substituting write operations:

```
detect x: sendmail_call::
  ( [ utility_call + ]
  && CONST(caller_id)
  && CONST(caller_process_type)
  && FREQUENCY(time_interval)
  ) & in_debug_mode(x) = True
from execute-program
do substitute z: write
  & process_id(x) = process_id(z)
  by dummy()
  substitutue w: read
  & process_id(x) = process_id(w)
  by protected_read()
```

Finally, the *fingerd* buffer-overflow to change the operating system's stack could be simulated via the instrumentation of the pop operations on the stack to simulate a jump to the new designated area of memory and execution of its commands with the current execution environment, requiring a full simulation of what happens there but on a copy of the operating system designated only for the offending process. Decoy methods like these can be generalized for broader applicability (by inserting code to call a general decoy monitor into many executables and parts of the operating system) or specialized for efficiency (by compiling separate decoy-enabled versions of executables that can be called when suspicious activities occur) depending on design priorities.

A software decoy should try to fool an intruder into thinking it is real. A necessary condition is that it returns most of the bits of information that would be returned by the real component corresponding to the simulated component. (This is not a sufficient condition because the type of information returned may be important even for the same number of bits.) A necessary information-theoretic condition for this is that the size of the decoy in bits must be most of the size of the amount of information it must transmit—and in the case of complex decoys that are too hard to summarize in an input-response table, perhaps larger still to model the decoying behavior in an easy-to-debug and easy-to-update way. However, in many practical situations it is sufficient that the decoy provide only a sample of all possible behaviors in response to only a sample of possible attacker inputs. For instance, a software decoy need only provide convincing fakes for the protected files that an attacker asks to see. For such situations, a good measure of decoy quality is the difference between the number of bits expected by the attacker and the number of bits transmitted. So if the same response is given to each of three actions, the number of bits transmitted remains the same as if one response is given to one action, but the number of bits expected is tripled.

Our behavior model provides means for specifying a broad spectrum of decoying strategies from specific ones for a particular intrusion type to a more generic when it is possible to come up with a generic description of intrusion patterns common to a whole class of attacks.

## V. SOFTWARE-DECOY ARCHITECTURE

We turn now to the implementation of the event-trace language to support software decoys. The structural (*i.e.*, domain behavior) model will consist of event types, event attributes, and axioms. Event types and attributes should be detectable independently.

We envision the possible implementation as an instrumentation of kernel libraries and log monitors available on the chosen platform. The basic instrumentation provides for the detection of events and event attributes specified in the generic behavior model and the interface for executing monitoring programs. A specialized compiler will generate monitoring programs from our high-level language for specifying both deception strategies and rules to be used by the software decoys. An option that permits the user to store event traces for debugging purposes may also be included in the general monitoring framework.

The generated monitors will include mechanisms for monitoring behavior-pattern implementation, via the use of state-transition diagrams—similar in some respects to

the NetSTAT [14] approach to network monitoring—and buffers to maintain event-attribute values.

Fig. 1 represents one possible architectural solution for implementing intelligent software decoys, in particular, for software components within an operating system. The separation of the repository of rules and behavior patterns from the supervisor supports the maintenance of the repository without changing the implementation (*i.e.*, component wrappers and supervisor); the repository and behavior model are confidential within the system under protection.

Execution of system commands contributes the main part of the event types in our behavior model. The basic instrumentation assumes that each system command participating in the model is enclosed in a wrapper, which is responsible for sending to the supervisor messages about event beginning and end and event attributes. The wrapper also contains a lightweight interface for implementing decoy actions, for example, inserting a delay or substituting the execution of the command by another subroutine execution. Behavior pattern detection is based on state transition diagrams maintained in the monitor. The interpreter within the monitoring kernel is responsible for rule interpretation and triggering decoying actions.

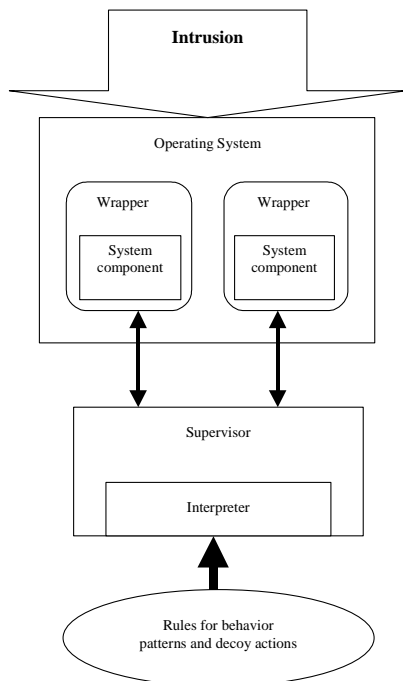


Fig. 1. Software decoy architecture

In addition to the goal of supporting software-based deception, our architecture is founded on two of the goals driving the specification and implementation of the Survivable Autonomic Response Architecture (SARA)

[15]: provision for both fast and coordinated responses to patterns of observed behavior. In the software-decoy architecture, the supervisor coordinates the actions among the decoyed-enabled components in order to both make decisions about how best to instrument components given the latest information about component interaction and select responses—that conform to information operations doctrine and policy—to effect the system-wide deception strategy.

In contrast to SARA, our architectural framework provides for performing machine learning at both the local (*i.e.*, decoy) and global (*i.e.*, supervisory) level, although we realize the need to evaluate the consequences of this decision in terms of requirements for the architecture such as adaptability, extensibility, and scalability.

## VI. RELATED WORK

The notions of event, path expression, and assertion language are well known in testing and debugging automation research. An event-based debugger for the C programming language called Dalek [16] provides user-defined events that typically are points within a program execution trace. A target program has to be instrumented manually in order to collect values of event attributes. The path expression technique was introduced for specifying parallel programs in [17]. This technique has been used in several projects on high-level debugging tools (*e.g.*, in [18]) where path rules are suggested as a kind of debugger command. Assertions (or annotations) currently in use are mostly based on Boolean expressions attached to selected points of the target program (*e.g.*, the *assert* macro in C). The ANNA [19] annotation language for Ada supports assertions on variable and type declarations. In the TSL [20] annotation language for Ada, the notion of event is introduced to describe the behavior of tasks. The RAPIDE project [21] provides a rich event-based assertion language for software architecture description. In [22], a practical approach to assertions for the C language is advocated. Our language for computations over traces provides a flexible means for writing both local and global assertions, including those about temporal relations between events. It supports all kinds of assertions used in the systems above in a uniform framework.

Sekar and Uppuluri [7] describe a high-level formalism for specifying intended program behaviors using patterns over sequences of system calls. The patterns also capture conditions on the values of system-call arguments. Their intrusion prevention and detection system is based on intercepting system calls, comparing them against the specifications of valid behaviors. In their approach, it is possible to modify a system call before it is delivered to the operating system's kernel, permitting the system to react before the execution by a process of any damage-

causing system call. The paper also presents a low-overhead algorithm for matching runtime behaviors against specifications.

We suggest a more powerful and expressive pattern specification language based on precise behavior models (event grammars) that will be able to capture a broader class of behaviors than those described by Sekar and Uppuluri. We also expect that our implementation strategy will be close to the guidelines presented by Sekar and Uppuluri, thus keeping the runtime overhead at or below four percent.

Neumann and Porras describe the architecture for EMERALD—an environment for anomaly and misuse detection—and subsequent analysis of the behavior of the systems and networks [23]. The architecture of EMERALD is based on component principles and can be adjusted to various platforms and configurations. A production-based expert system is used to analyze behavior patterns and signatures.

The use of event traces for reasoning about the behavior of information systems is not new. Event traces have been used to reason about the correctness of the execution of hardware-level instructions. For instance, Bressoud and Schneider [24] present an approach to fault tolerance based on the use of hypervisors that emulate the target hardware architecture: Each hypervisor serves as a virtual machine. The hypervisor is then instrumented to watch for sequences of instructions, within an epoch (*i.e.*, interval of time), that are critical to the correct operation of replicated hardware within a system.

Working at a higher level of abstraction, Erlingsson and Schneider [25] introduce an approach known as Security Automata SFI Implementation (SASI). In this approach, the object code is instrumented so that as sequences of events occur, the trusted computing base, which in this case includes the software that analyzes and modifies object code, can use the event traces to reason about whether the traces are conformant to security policy (*e.g.*, memory-protection policy).

Bressoud and Schneider [24] discuss some of the difficulties in instrumentation of virtual machines and reasoning over sequences of hardware instructions. Likewise, Erlingsson and Schneider admit that it is difficult to recognize patterns of interest from specifications of object and security policy represented in their security automaton language (SAL), and that the development of such specifications is an “awkward and error-prone” process. Their new approach is to rely on the use of “annotations of the object code that are easily checked and that expose application-level abstractions,” rather than checking every machine-level instruction. Our formalism of decoys relies on specifying events for which the source code of components can be instrumented.

Sekar and Uppuluri [7] attempt to defend against buffer-overflow attacks by injecting code and addresses into memory, and transferring control flow to the injected code. They manually instrument system calls and their arguments for each component to be protected against buffer-overflow attacks. A run-time monitor then compares these patterns of nominal and anomalous behavior with that of the runtime behavior of the protected component. Our approach differs from that of Sekar and Uppuluri, especially in terms of the method of instrumentation: that is, in our approach, the software components are automatically instrumented.

Liu and Jajodia [26] introduce specialized decoying methods and strategies for transaction-level intrusion tolerance and damage confinement in database management systems (DBMS). The approach is bound to transaction events within the DBMS and relies on the use of filtering strategies. In contrast, our approach can be specialized to different types of events and countermeasures.

## VII. CONCLUSION

Our approach explicitly treats the two types of invocations of components: standard and nonstandard, with the latter representing an attempt by a process to circumvent the public interface of the component. For standard types of invocations or components of legacy systems for which the source code is not available, then certain types of events of interest will not be detectable.

The behavior model based on event grammars provides a uniform framework for behavior-pattern recognition and decoying actions. In our opinion, this is a practical road to formalizing knowledge about typical intrusion patterns and constructing a flexible system of countermeasures.

Our next steps are to implement the architecture, develop measures of effective, and conduct experiments (similar in nature to those used by Turing in his experiments with the “imitation game” [27]) with the decoys embedded in real systems operating in test-and-evaluation mode (*e.g.*, weapon control systems with live-fire disabled).

## VIII. ACKNOWLEDGEMENTS

This research is supported by the Naval Research Laboratory under contract no. N41756-01-WR-10433, Office of Naval Research under contract N00014-01-10746, the U.S. Army Research Laboratory under contract no. 40473-MA, and the National Research Council. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright annotations thereon.



# REFERENCES

- [1] Spafford, E. H. The Internet worm: Crisis and aftermath. *Comm. ACM* 32, 6 (June 1989), 678-687.
- [2] "Code Red" worm exploiting buffer overflow in IIS indexing service DLL. Incident Note IN-2001-08, CERT Coordination Center, Carnegie-Mellon Software Eng. Inst., Pittsburgh, Penn., July 19, 2001.
- [3] "Code Red II:" Another worm exploiting buffer overflow in IIS indexing service DLL. Incident Note IN-2001-09, CERT Coordination Center, Carnegie-Mellon Software Eng. Inst., Pittsburgh, Penn., Aug. 6, 2001.
- [4] Michael, J. B. and Riehle, R. D. Intelligent software decoys. In *Proc. Monterey Workshop: Eng. Automation for Software Intensive Syst. Integration*, n.p., (Monterey, Calif., June 2001), 178-187.
- [5] Somayaji, A. B. Operating System Stability and Security through Process Homeostasis. Ph.D. dissertation, University of New Mexico, 2002.
- [6] Auguston, M. Tools for program dynamic analysis, testing, and debugging based on event grammars. In *Proc. Twelfth Int. Conf. Software Eng. and Knowledge Eng.*, Skokie, Ill.: Knowledge Syst. Inst. (Chicago, Ill., July 2000), 159-166.
- [7] Sekar, R. and Uppuluri, P. Synthesizing fast intrusion detection/prevention systems from high-level specifications. In *Proc. USENIX Sec. Symp.*, Berkeley, Calif.: USENIX Assn. (Washington, D.C., Aug. 1999), 63-78.
- [8] Ko, C., Fraser, T., Badger, L., and Kilpatrick, D. Detecting and countering system intrusions using software wrappers. In *Proc. Ninth USENIX Sec. Symp.*, Berkeley, Calif.: USENIX Assn. (Denver, Col., Aug. 2000), 145-156.
- [9] The "cheese" worm. Incident Note IN-2001-05, CERT Coordination Center, Carnegie-Mellon Software Eng. Inst., Pittsburgh, Penn., May 17, 2001.
- [10] Evans, S., Bush, S. F., and Hershey, J. Information assurance through Kolmogorov complexity. In *Proc. DARPA Inf. Survivability Conf. and Exposition*, vol. 2, IEEE (Anaheim, Calif., June 2001), 322-331.
- [11] Fowler, C. A. and Nesbit, R. F. Tactical deception in air-land warfare. *J. Electronic Defense* 18, 6 (June 1995), 37-44, 76-79.
- [12] Thimbleby, H., Anderson, S., and Cairns, P. A framework for modelling Trojans and computer virus infection. *Comput. J.* 41, 7 (1998), 444-458.
- [13] Rice, H. G. Classes of recursively enumerable sets and their decision problems. *Trans. Amer. Math. Soc.* 89 (1953), 25-99.
- [14] Vigna, G. and Kemmerer, R. A. NetSTAT: A network-based intrusion detection approach. In *Proc. Fourteenth Annual Computer Sec. Applications Conf.*, ACM (Scottsdale, Ariz., Dec. 1998), 25-34.
- [15] Lewandowski, S. M., Van Hook, D. J., O'Leary, G. C., Haines, J. W., and Rossey, L. M. SARA: Survivable autonomic response architecture. In *Proc. DARPA Inf. Survivability Conf. and Expo.*, vol. 1, IEEE (Anaheim, Calif., June 2001), 77-88.
- [16] Olsson, R., Crawford, R., and Wilson, W. A dataflow approach to event-based debugging. *Software—Practice & Experience* 21, 2 (Feb. 1991), 19-31.
- [17] Campbell, R. H. and Habermann, A. N. The specification of process synchronization by path expressions. In Goos, G. and Hartmanis, J., eds., *Proc. Int. Symp. Operating Syst., Lect. Notes Comput. Sci.*, vol. 16, Berlin: Springer-Verlag (Rocquencourt, Fr., Apr. 1974), 89-102.
- [18] Bruegge, B., and Hibbard, P. Generalized path expressions: A high-level debugging mechanism. *J. Syst. Software* 3, 4 (1983), 265-276.
- [19] Luckham, D., Sankar, S., and Takahashi, S. Two-dimensional pinpointing: Debugging with formal specifications. *IEEE Software* (Jan. 1991), 74-84.
- [20] Luckham, D., Bryan, D., Mann, W., Meldal, S., and Helmbold, D. P. An introduction to task sequencing language, TSL version 1.5 (Preliminary version), Stanford Univ., Feb. 1, 1990, pp. 1-68.
- [21] Luckham, D. and Vera, J. An event-based architecture definition language. *IEEE Trans. Software Eng.* 21, 9 (Sept. 1995), 717-734.
- [22] Rosenblum, D. A practical approach to programming with assertions. *IEEE Trans. Software Eng.* 21, 1 (Jan. 1995), 19-31.
- [23] Neumann, P. G. and Porras, P. A. Experience with EMERALD to date. In *Proc. First USENIX Workshop Intrusion Detection and Network Monitoring* (Santa Clara, Calif., Apr. 1999), 73-80.
- [24] Bressoud, T. and Schneider, F. B. Hypervisor-based fault-tolerance. *ACM Trans. Comput. Syst.* 14, 1 (Feb. 1996), 80-107.
- [25] Erlingsson, Ú. and Schneider, F. B. SASI enforcement of security policies: A retrospective. In *Proc. New Sec. Paradigms Workshop*, ACM (Caledon Hills, Ont., Sept. 1999), 87-94.
- [26] Liu, P. and Jajodia, S. Multi-phase damage confinement in database systems for intrusion tolerance. In *Proc. Fourteenth Comput. Sec. Foundations Workshop*, IEEE (Cape Breton, N.S., June 2001), 191-205.
- [27] Turing, A. M. Computing machinery and intelligence. *Mind* 59, 236 (Oct. 1950), 433-460