

A Security Domain Model for Implementing Trusted Subject Behaviors

Alan Shaffer, Mikhail Auguston, Cynthia Irvine, Timothy Levin

Naval Postgraduate School, Computer Science Department, 1411 Cunningham Rd.,
Monterey, California, USA 93943
{abshaffe, maugusto, irvine, levin}@nps.edu

Abstract. Within a multilevel secure (MLS) system, trusted subjects are granted privileges to perform operations that are not possible by ordinary subjects controlled by mandatory access control (MAC) policy enforcement mechanisms. These subjects are trusted not to conduct malicious activity or degrade system security. We present a formal definition for trusted subject behaviors, which depends upon a representation of information flow and control dependencies generated during a program execution. We describe a security Domain Model (DM) designed in the Alloy specification language for conducting static analysis of programs to identify illicit information flows, access control flaws and covert channel vulnerabilities. The DM is compiled from a representation of a target program, written in an intermediate Implementation Modeling Language (IML), and a specification of the security policy written in Alloy. The Alloy Analyzer tool is used to perform static analysis of the DM to detect potential security policy violations in the target program. In particular, since the operating system upon which the trusted subject runs has limited ability to control its actions, static analysis of trusted subject operations can contribute to the security of the system.

Keywords: Security domain model, trusted subjects, static analysis, automated program verification, specification language.

1 Introduction

Within a multilevel secure (MLS) system, certain *trusted subjects* may be granted privileges to perform operations, in some cases within prescribed limits [17], not normally allowed for ordinary subjects controlled by mandatory access control (MAC) policy enforcement mechanisms. Granting of such privileges is predicated on the idea that trusted subjects will not conduct malicious activity or degrade the system's overall security. This paper presents a formal definition for trusted subject behaviors in certain program implementations. These behaviors depend upon a representation of information flow and control dependencies generated during a target program execution, thus extending classic work in this area [6][15][23]. We describe a security domain model to formally represent trusted subject behaviors, information flow tracing through program execution, various types of covert channels, and a means for conducting static analysis of certain program implementations.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE SEP 2008	2. REPORT TYPE	3. DATES COVERED 00-00-2008 to 00-00-2008			
4. TITLE AND SUBTITLE A Security Domain Model for Implementing Trusted Subject Behaviors		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School, Department of Computer Science, Monterey, CA, 93943		8. PERFORMING ORGANIZATION REPORT NUMBER ; NPS-CS-08-015			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S) NPS-CS-08-015			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 12	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Widely accepted evaluation standards [4][5][14] require that high assurance secure systems be designed, developed, verified and tested using rigorous processes and formal methods. This evaluation process must include demonstration of correct correspondence between system representations at various levels of abstraction, e.g., security policy objectives, security specifications, and program implementation. Formal security models are often based on concepts of program secure state and state transitions. Our approach analyzes programs for preservation of security properties through state transitions, and advances the concepts of secure information flow in classic work by Denning and others [6][23] by describing automated techniques for information flow static analysis. Our previous work has demonstrated the ability to detect illicit information flow security violations [19], and covert channel and overt access control flaws based on control dependency analysis [20].

The *Implementation Modeling Language* (IML), the first novel element in this approach, is a language that supports basic information processing via assignment statements, conditional and loop statements, read/write statements, file random access, and access to a system clock. The *target program* is an original high-level language program from which we extract a *base program*, the IML abstraction that provides a basis for analysis of the target program for adherence to a security policy.

The second novel element in this work is the definition of a security *Domain Model* (DM), represented as an Alloy [1][7] specification. The DM provides a framework for specifying program state and state transitions, as well as security-related concepts such as security policy, information flow, access control, and covert channel vulnerabilities. The DM is comprised of an *Invariant Model*, which defines the generic concepts of program state, information flow, and security policy; and an *Implementation Model*, which specifies the behavior of the base program. A specialized *DM-Compiler* was developed to translate a base program in IML into an Implementation Model, and to integrate it with the Invariant Model to form a complete DM specification.

Our approach uses the Alloy Analyzer tool [1] to perform static analysis of base programs to identify execution paths that might violate the security policy rules. The Alloy Analyzer performs symbolic execution of all base program execution paths within a defined *scope* (the upper limit of the size of models considered); the scope is generated heuristically, based on the total number of statements in the base program. It is assumed that the Alloy *small scope hypothesis*, which states that most flaws in models can be revealed on small instances [7], holds for information flow tracing. A description of the DM structure, and examples of refinement of a base program to Alloy can be found in [18][20].

Section 2 of this paper provides background discussion on trusted subjects and processes. Section 3 presents an overview of the Security DM methodology for modeling programs and security policies. Sections 4 and 5 describe analyses of several example base programs using the DM, and summarize our test results with these examples. Sections 6 and 7 discuss related work, and planned future work in this research.

2 Trusted Subjects

Users in a multilevel secure (MLS) environment are assigned a clearance level based on the relative level of trust placed in them by security administrators. A user is allowed to log into a system at any level that is at or below (*dominated by*) his assigned clearance, which creates a *session* at that level. Subjects that act on behalf of a given user are labeled with an access class that is at the same level as the user's session. A subject is allowed to read information (*objects*) whose sensitivity level is up to the subject's clearance level (*access class*), and write to objects at or above their access class.

In contrast to this, a *trusted subject* is one that is allowed to read and write within a *range* of access classes [12], which limits the authority of the trusted subject to "read-up" and "write-down" [3]. MLS systems with trusted subjects defined this way do not require a separate access control lattice or special rules specifically for their actions [12]. As a result, a trusted subject does not need to be given a privilege to bypass or violate the security rules.

Since trusted subjects are allowed to interact with (read & write) information across access classes, they must be trusted not to abuse these special privileges. The existence of trusted subjects is generally required for certain services provided in MLS systems, such as login, information downgrading, and data backup utilities across multiple access levels. MLS system administration may also require a trusted subject to interact with and manage regular user accounts and information across multiple access levels [22]. Such actions represent a good target for trusted subject implementation, however research has not always focused on the design principle that trusted subjects should be small and minimized within an MLS system.

With respect to security policies, a trusted subject should not move data between sensitivity levels, other than in constrained, explicitly defined ways [21]. The specification of a trusted subject must explicitly define how it can do this. Levin et al. [11] point out that trusted subjects do not violate a general policy in place, but their behavior must be a defined part of a policy. Such a policy for trusted subjects, referred to as a "relaxed MLS policy," must be integrated with the general MLS policy, such that the resultant union of the two can allow trusted subjects to effectively operate, while ensuring that non-trusted subjects cannot conduct malicious activity. In a "downgrader" role, for example, a trusted subject may essentially change the label of information from *high* to *low* by reading the information from a *high* object and moving it into another *low* object.

Trusted subjects can be defined by their behaviors in an MLS system. Steffan & Clow [21] described examples of trusted subject actions, including the ability to process information across multiple access control levels, e.g., to view (read) a highly sensitive document in order to comment (write) on its existence at a lower level; and the ability to modify the labels of sensitive information. In the latter case, a trusted subject may *regrade* a classified document, which would seem to require overriding the tranquility principle that a subject or object's label will not change while being referenced [3]. We maintain the tranquility of object labels, abstracting the idea of downgrading information by changing variable labels from the viewpoint of, i.e.,

internal to, the trusted subject. Allowing movement of information within a range of access classes represents the trusted subject actions we model in our DM approach.

3 Security Domain Model Methodology

Our approach to program security verification using the Security Domain Model has been described in [19][20], and an overview is provided below.

A *base program* represents an abstraction of a target program implementation, and is written in Implementation Modeling Language (IML) notation. The IML defines a simple imperative language that captures the basic capabilities and constructs, with respect to security, of high-level programming languages, such as Java or C++. The IML was motivated by a requirement to represent the information flow properties in target program implementations. A complete IML and DM reference manual is available online at [18].

In this work the IML language syntax has been expanded to enable the representation of trusted subject behaviors. Specifically, the IML allows trusted subject to modify variable labels (“regrading”), and to “filter” variable values and labels based on existing content and/or label. To enable regrading, the IML provides a Trusted Assignment statement, which allows a trusted subject to assign a value to a *destination* variable with a specifically assigned security label. When an IML base program is translated, it is under the context that only a trusted subject may perform trusted assignment. In this operation, the *destination* variable takes on the data value of the source variable, however it does not automatically take its label, as would normally be the case for an assignment statement in IML. Instead, the *destination* is explicitly assigned a label, as determined by a filter function that is automatically invoked with each trusted assignment. The trusted assignment syntax follows:

```
Assign destination from source1 as source2;
```

The trusted assignment *source1* can be either a variable or constant, and the *source2* can be either a variable (in which case the access label currently assigned to the value stored in this variable is used) or an explicitly defined access label. The new content and access label of the *destination* variable are defined by an Alloy function `TS_filter` in the DM Invariant Model:

```
destination(value', label') =  
  TS_filter ( destination(value, label), source1(value, label), source2(value, label) )
```

This function specifies the behavior of trusted subjects in our model, and an example is described in detail in section 4.1.

The DM *Invariant Model* defines security rules that have the Bell & LaPadula security model [3] as their basis, i.e., flows from higher to lower secrecy levels are not allowed by either writing down or reading up. The general DM security policy defines a lattice with flows allowed from lower to higher (or equal) secrecy levels, represented by access labels, for instance:

```
one sig Policy {  
  ord: AccessLabel -> AccessLabel }
```

```

{ ord = ^( (SysLow -> SysMid)
           + (SysMid -> SysHigh) )
  + (iden & (AccessLabel -> AccessLabel) ) }

```

In Alloy notation this defines a recursive closure of the access label relations (*SysLow* -> *SysMid*) and (*SysMid* -> *SysHigh*). The “basic” security policy is defined in the DM Invariant Model by reads and writes of external I/O devices that conform to this policy lattice. The trusted policy is defined such that trusted subjects are allowed to change labels and data within the constraints of the *TS_filter*.

4 Testing and Analysis of Trusted Subject Behaviors in the DM

This section presents examples of program security vulnerabilities that illustrate how trusted subjects are constrained by both the basic security policy and the trusted policy (as implemented in the *TS_filter*). In these examples, the security rules for discovering information flow errors, overt access control flaws and covert channels, are described using Alloy notation, and a base program written in IML is presented to illustrate the particular security violation. The complete Alloy models for these and other examples can be found online at [18].

4.1 Flow Violation Caused by a Trusted Subject Operation (Example 1)

The first example illustrates a trusted subject regrade operation that, based on allowed trusted subject behavior, leads to an information flow violation. In the example, an attempt is made by a trusted subject to downgrade a destination variable label from *SysHigh* to *SysLow*. Here, trusted subjects are allowed to perform downgrading of information from *SysHigh* to *SysMid*. To support the policy, a *TS_filter* function is defined (below in Alloy notation) to ensure that “downward” info flows are allowed only from *SysHigh* to *SysMid*. The function takes as input parameters three *Values* and three *AccessLabels*, specifically, the data values and labels of the *destination*, *source1* and *source2* variables in the Trusted Assignment (see Section 3 for trusted assignment IML syntax), and returns an instance of *FTuple* (i.e., a filtered *Value* and *AccessLabel*). In essence, the policy for trusted subject behaviors is captured in the semantics of this filter function.

For example purposes here, this *TS_filter* function returns the greater of constant 0 and the *source1* *Value* (*s1v*), and the higher of *SysMid* and the *source2* *AccessLabel* (*s2a*). As shown in this example *TS_filter*, it is not necessary to use all of the parameters passed into the function to generate a resulting *FTuple*. Note that a different DM Invariant Model could define a *TS_filter* function that would return different results based on the specific input parameters, and thus define a different security policy for trusted subject behaviors.

```

sig FTuple {
  val: Value,
  label: AccessLabel
}

```

```

fun TS_filter[dv, s1v, s2v: Value,
             da, s1a, s2a: AccessLabel]: FTuple {
{ result: FTuple | {
  result.val = (((s1v->const0) in LT.lt)
=> const0 else s1v)
  result.label = (((da->s2a) in Policy.ord)
=> s2a else
  ((s2a->SysMid) in Policy.ord)
=> SysMid else s2a)) }
} }

```

The base program example below demonstrates a security violation based on the trusted subject filter and security policy. Initially, values are read into two variables with security labels *SysHigh* and *SysMid*, respectively (s1-s2). A trusted assignment operation is then performed (s3), in which the data value stored in x2 is copied into variable x1, and x1 is assigned a *SysLow* label. During this statement operation, the *TS_filter* function is applied to the parameters of the trusted assignment, "filtering" the label assignment to *SysMid*, which results in x1 being assigned a higher label than was intended by the trusted assignment operation (s3).

```

(s1) Read_dev (SysHigh, x1);
(s2) Read_dev (SysMid, x2);
(s3) Assign x1 from x2 as SysLow;
(s4) Write_dev (SysLow, x1);
(s5) Stop;

```

When the next statement (s4) attempts to write the value held in x1 to a *SysLow* external device, an illicit flow results since x1 is labeled as *SysMid*. The Alloy Analyzer detects this situation as a violation of the Alloy security predicate below, and correctly reports an illicit information flow, tracing execution through statements (s1)(s2)(s3)(s4). The same base program, under a DM Invariant Model with a different policy and filter function, would not necessarily result in this flow violation.

```

pred consistent_with_FlowPolicy [current: State] {
  let stm = current.stmt | {
    ( stm.type in (Write_dev + PutDirectFile) &&
      stm.source in Variable )
    => (current.access_label[stm.source] ->
        stm.subject_label) in Policy.ord
  } }

```

4.2 Dual Trusted Subject Flow Violation & Overt Flaw (Example 2)

The second example base program illustrates two different security violations that may result from a trusted subject operation. In the program, a successful trusted subject regrade creates an overt control dependency flaw, however when the trusted subject regrade fails to occur, illegal information flow results. For purposes of this example, the security policy and *TS_filter* function described above apply.

In the base program, values are initially read into three variables, with assigned security labels *SysHigh*, *SysMid* and *SysLow*, respectively (s1-s3). Depending on the value stored in *x1* (s4), a trusted assignment statement is performed (s5) in which the value of *x1* is modified to that of *x2*, and the label of *x1* is downgraded to that of *x3*, *SysMid* in this case. Since a regrade from *SysHigh* to *SysMid* is allowed by the security policy (as reflected in the *TS_filter* function), *x1* is assigned the *SysMid* label.

```
(s1) Read_dev (SysHigh, x1);
(s2) Read_dev (SysLow, x2);
(s3) Read_dev (SysMid, x3);
(s4) if x1 < 0 then {
(s5)   Assign x1 from x2 as x3;
(s6)   Write_dev (SysMid, x1); }
      else
(s7)   Write_dev (SysMid, x1);
(s8) Stop;
```

The next statement (s6) attempts to write the value of *x1* to a *SysMid* external device, a seemingly legal flow. However, since this operation occurs within the if-block, it creates a control dependency from *SysHigh* (*x1* label when it was examined in s4) to *SysMid*, representing an overt access control flaw (i.e., in the *SysHigh* context, a write to *SysMid* violates the security policy). Based on the Alloy security rule predicate below, the Alloy Analyzer properly detects this violation, tracing execution through statements (s1)(s2)(s3)(s4)(s5)(s6).

```
pred dependency_flaw_found [current: State] {
  let stm = current.stmt,
      pre = current.influenced_by[stm.source] |
  {
    stm.type = Write_dev &&
    stm.source in Variable &&
    not ((pre.access_label[pre.stmt.source] ->
          stm.subject_label) in Policy.ord)
  } }
```

An additional violation occurs when the conditional check (s4) evaluates to false, and the else-branch is executed. In this case, an attempt is made to write the value stored in *x1* (still assigned its original *SysHigh* label) to a *SysMid* external device (s7). Since this represents an overt illegal flow from *SysHigh* to *SysMid*, the Alloy Analyzer properly identifies and reports the error, tracing execution through statements (s1)(s2)(s3)(s4)(s7).

4.3 Covert Channel Resulting from a Trusted Subject Operation (Example 3)

The third scenario describes execution of a trusted assignment that could produce a covert storage channel [10]. Our earlier paper [20] describes in detail how the DM formalizes the notion of covert channels, and defines a security rule to identify a class of covert storage channel vulnerabilities in a base program execution.

In the base program below, we assume a direct file with a maximum capacity of two records, initially empty. To begin, *SysLow* values are read into variables $x1$ and $x2$ (s1-s2). A trusted assignment is then performed (s3) in which $x1$ is assigned the value of $x2$, and upgraded to a *SysHigh* label. Next, the value of $x1$ is examined to verify whether it is non-negative (s4). Because the *TS_filter* function (see section 4.1) only returns filtered values of 0 or greater, $x1$ holding a non-negative value is an indication that the trusted assignment resulted in the assignment of source data to the destination variable. When this check evaluates to true, the values of $x1$ and $x2$ are stored into the direct file by the *SysHigh* sender, resulting in the internal *full* direct file flag being set.

```
(s1) Read_dev (SysLow, x1);
(s2) Read_dev (SysLow, x2);
(s3) Assign x1 from x2 as SysHigh;
(s4) if x1 > const_minus_1 then {
(s5)   PutDirectFile (SysHigh, 1, x1);
(s6)   PutDirectFile (SysHigh, 2, x2); }
```

The next sequence of program statements represent execution by a *SysLow* covert channel receiver. When the *SysLow* subject attempts to store a value into the direct file using a new key 3 (s7), the system issues a failure indication since the direct file is *full* (note that in the translation to an base program, the internal system flag translates to an explicit flag, accessible in IML as in statement (s8)). Depending on the success or failure of the direct file store (s8), *SysLow* writes a constant '1' or a '0' to an external device (s9 & s10) to exploit the storage channel.

```
(s7) PutDirectFile (SysLow, 3, 1);
(s8) if full = True then
(s9)   Write_dev (SysLow, 1);
(s10) else Write_dev (SysLow, 0);
(s11) Stop;
```

Because a higher-labeled subject caused the direct file to become full, the Alloy Analyzer detects and reports this violation of the below Alloy security predicate, tracing the flow of execution through statements (s1)(s2)(s3)(s4)(s5)(s6)(s7). The actions of two regular subjects at *SysHigh* and *SysLow*, acting in collusion to exploit the direct file, could bring about the same security violation (i.e., a storage channel).

```
pred storage_channel_found [current: State] {
  let stm = current.stmt | {
    stm.type = PutDirectFile &&
    current.direct_file.full = const1 &&
    not (current.direct_file.last_written ->
        stm.subject_label) in Policy.ord
  } }
```

5 Testing Results

The base program examples presented above were evaluated using Alloy Analyzer 4.1.7, running under Mac OS X™ 10.5.4 on a 2.5 GHz Intel Core 2 Duo processor, with 2 GB of memory. In test runs, the Alloy Analyzer successfully found valid counterexamples for violations of each security rule assertion described above. Test run times, in *ms*, were as follows; *total time (time to generate model, time to find counterexample)*:

- Example 1 (scope = 7): 1516 (1277, 239)
- Example 2, overt control dependency flaw (scope = 9): 3335 (2290, 1045)
- Example 2, information flow violation (scope = 9): 2692 (2236, 456)
- Example 3 (scope = 12): 48631 (9852, 38779)

6 Related Work

Earlier work in trusted subject implementation [24] developed a framework for running a trusted multi-level database management system (DBMS), referred to as a “trusted subject,” to be hosted on any trusted operating system. This work established a layered policy, with a general policy for the *trusted computing base* (TCB) layer of the operating system, and a separate policy for the DBMS TCB layer. Their premise was that, for a DBMS hosted on a known secure operating system, only the DBMS TCB layer must be subjected to security analysis to ensure that it meets all access control requirements. This concept provided “portability” of the DBMS trusted subject, and negated the time and expense of testing every system on which the DBMS may be targeted. This work did not appear to outline a traditional concrete security policy for trusted subjects, and only used them in the context of a trusted DBMS. This paper focused on modification of object tranquility as a valid action for trusted subjects.

Landauer et al. [8] introduced a formal model for managing trusted processes, by defining a state machine whose state space can be locked, or isolated, in order to allow privileged actions to overlap. The authors described a trusted process as possessing special privileges to alter operating system kernel access control decisions, or other security critical operations; their definition provides a good description of a trusted subject. This paper provided an in-depth mathematical analysis of the security policy derived from trusted process principles, and is valuable as a source of background discussion on security policy issues for trusted subjects.

Steffan & Clow [21] defined a set of trusted process classes, to identify their relative privileged status. These classes correspond to combinations of override privileges in the areas of Tranquility (labels), MAC (content) and DAC (privileges). As the class numbers increase, so do the privileges granted, and the risk associated with using a trusted process in that class. In contrast to this paper, our work characterizes trusted subjects without violating tranquility of object labels.

Thomas & Sandhu [22] presented three architectures for trusted object-oriented databases, based on: a kernel, a replicated DBMS, and a trusted subject architecture. The last of these was the focus of their paper. Their architecture was composed of a

session manager, which was trusted and running across multiple layers; several message managers, which were untrusted, and operated within a single layer; and service read/write requests to the DB from a client. The trusted session manager has the advantage that it can always maintain a global snapshot of the system for a given session, across all access control layers, to allow it to coordinate message requests and scheduling. Clearly, the session manager must be a trusted subject for this architecture to maintain security of messages across layers. This paper did not appear to focus on security policies for trusted subjects, i.e., how separate policies might be treated in the session manager and message managers and how such policies could effectively coexist.

Levin et al, [9] discussed trusted subject actions within a security kernel architecture. With respect to the principle of least privilege [16], they described how a trusted subject in a downgrader role must be able to perform only the minimum required operations, namely, downgrading of labels in this case. Other operations such as “dirty word search” (DWS) of a document for specific words or phrases prior to downgrade, must be handled by other system processes to prevent unintended or malicious consequences. They defined a framework for performing filtering and downgrade of information, separating tasks between users and processes, both untrusted and trusted. We believe our model is in line with this thinking, when one considers that if our trusted subject acts as a downgrader, the Invariant Model filter function can reflect a separate untrusted process in the target program that performs DWS. We generalize this concept by allowing the trusted subject to downgrade based on content or label information. In our model, the DWS might represent examination of a highly classified document for specific references to some classified topic, with subsequent removal of these references prior to downgrading the document. Alternately, the DWS could represent filtering of a document by its creation date, where downgrading of the document will occur only if this information is older than some predetermined date.

7 Discussion and Future Work

This paper has provided a survey of ongoing research to develop a formal security domain model that formalizes security policies for both regular and trusted subjects. The model formalizes trusted subject behaviors, using the specialized imperative language. Our approach defines a formal security Domain Model (DM) that facilitates specification of security vulnerabilities and trusted subject behaviors, independent of program implementation.

By the *small scope hypothesis* [7] it is assumed that most program errors may be revealed by relatively small counterexamples. Using the Analyzer to perform static analysis of the DM provides assurance that, within a specified search scope, a counterexample will be found when one exists, and that false negatives and false positives are eliminated within the defined scope. This assumption necessitates our implementation of a relatively small trusted subject, which is inline with the Reference Monitor Concept principle that a reference validation mechanism “must be small enough to be subject to analysis and tests” [2] to ensure its correctness.

Future work will expand the DM to enable dynamic security policies [10]. This concept would allow the DM to support a sequence of policies during program execution, and support the ability of a system to adapt to a dynamically changing security environment by using different policies [13]. We could extend this by adding functionality for multiple trusted subjects. By defining multiple filter functions within a DM Invariant Model, and modifying the IML syntax to support this, the model could represent separate trusted subjects, each governed by a different policy as defined by its own filter function.

References

1. The Alloy Analyzer, <http://alloy.mit.edu/>.
2. Anderson, J. (1972). *Computer Security Technology Planning Study*. Technical Report ESD-TR-73-51, Air Force Electronic Systems Division, Hanscom AFB, Bedford, MA.
3. Bell, D., & LaPadula, L. (1973). *Secure Computer Systems: Mathematical Foundations and Model*, MITRE Report. The MITRE Corp.
4. *Common Criteria for Information Technology Security Evaluation, Part 1: Introduction and General Model*, version 3.1. Document number CCMB-2006-09-001. September 2006.
5. *Department of Defense Trusted Computer Security Evaluation Criteria*, DOD 5200.28-STD, National Computer Security Center, December 1985.
6. Denning, D. E., & Denning, P. J. (1977). Certification of programs for secure information flow. *Communications of the ACM*, 20(7), 504-512. ACM Press.
7. Jackson, D. (2006). *Software Abstractions: Logic, Language, and Analysis*. Cambridge, MA, USA, and London, England: MIT Press.
8. Landauer, J., Redmond, T., & Benzel, T. (1989). Formal policies for trusted processes. *Proceedings of the Computer Security Foundations Workshop II*, 31-40.
9. Levin, T., Irvine, C., and Nguyen, T. (2006). Least Privilege in Separation Kernels. *Proceedings of the 2006 International Conference on Security and Cryptography*, 355-362.
10. Levin, T., Irvine, C., & Spyropoulou, E. (2006). *Quality of security service: Adaptive security*. Handbook of Information Security (H. Bidgoli, ed.), vol. 3, 1016-1025. Hoboken, NJ: John Wiley and Sons.
11. Levin, T., Irvine, C., Weissman, C., & Nguyen, T. (2007). Analysis of three multilevel security architectures. *Proceedings of the 2007 ACM Workshop on Computer Security Architecture*, 37-46. ACM Press, New York, NY.
12. Lunt, T., Denning, D., Schell, R., Heckman, M., & Shockley, W. (1990). The seaview security model. *IEEE Transactions on Software Engineering*, 16(6), 593-607.
13. National Security Agency IA Directorate. (2004). *Global Information Grid Information Assurance Reference Capability/Technology Roadmap*, Version 1.0.
14. National Security Agency. (2007). *U.S. Government Protection Profile for Separation Kernels in Environments Requiring High Robustness*, Version 1.03.
15. Sabelfeld, A., & Myers, A. (2003). Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1), 5-19. IEEE Press.
16. Saltzer, J. & Schroeder, M. (1975). The protection of information in computer systems. *Proceedings of the IEEE*, 63(9), 1278-1308.
17. Schell, R., Tao, T., & Heckman, M. (1985). Designing the GEMSOS Security Kernel for Security and Performance. *Proceedings of the 8th National Computer Security Conference*, 108 - 119.

18. Security Domain Model Project website, <http://cisr.nps.edu/projects/sdm.html>.
19. Shaffer, A., Auguston, M., Irvine, C. and Levin, T. (2007). Toward a security domain model for static analysis and verification of information systems. *Proceedings of the 7th OOPSLA Workshop on Domain-Specific Modeling*, 160-171. Montreal, Canada.
20. Shaffer, A., Auguston, M., Irvine, C. and Levin, T. (2008). A Security Domain Model to Assess Software for Exploitable Covert Channels. *Proceedings of the ACM SIGPLAN Third Workshop on Programming Languages and Analysis for Security (PLAS'08)*, 45-56. Tucson, Arizona. ACM Press.
21. Steffan, W., & Clow, J. (1996). Trusted process classes. *Proceedings of the 19th National Information Systems Security Conference*.
22. Thomas, R., & Sandhu, R. (1996). A trusted subject architecture for multilevel secure object-oriented databases. *IEEE Transactions on Knowledge and Data Engineering*, 8(1), 16-31.
23. Volpano, D., Smith, G., & Irvine, C. (1996). A sound type system for secure flow analysis. *Journal of Computer Security*, 4(3), 167-187.
24. Wilson, J. (1989). A security policy for an A1 DBMS (a trusted subject). *Proceedings of the 1989 IEEE Symposium on Security and Privacy*, 116-125.