

Practical Fine-Grained Information Flow Control Using Laminar

DONALD E. PORTER, Stony Brook University
 MICHAEL D. BOND, Ohio State University
 INDRAJIT ROY, Hewlett Packard Labs
 KATHRYN S. MCKINLEY, Microsoft Research
 EMMETT WITCHEL, The University of Texas at Austin

Decentralized Information Flow Control (DIFC) is a promising model for writing programs with powerful, end-to-end security guarantees. Current DIFC systems that run on commodity hardware can be broadly categorized into two types: language-level and operating system-level DIFC. Language solutions provide no guarantees against security violations on system resources such as files and sockets. Operating system solutions mediate accesses to system resources but are either inefficient or imprecise at monitoring the flow of information through fine-grained program data structures. This article describes Laminar, the first system to implement DIFC using a unified set of abstractions for OS resources and heap-allocated objects. Programmers express security policies by labeling data with secrecy and integrity labels and access the labeled data in *security methods*. Laminar enforces the security policies specified by the labels at runtime. Laminar is implemented using a modified Java virtual machine and a new Linux security module. This article shows that security methods ease incremental deployment and limit dynamic security checks by retrofitting DIFC policies on four application case studies. Replacing the applications' ad hoc security policies changes less than 10% of the code and incurs performance overheads from 5% to 56%. Compared to prior DIFC systems, Laminar supports a more general class of multithreaded DIFC programs efficiently and integrates language and OS abstractions.

Categories and Subject Descriptors: D.2.4 [Software Engineering]: Software/Program Verification; D.3.3 [Programming Languages]: Language Constructs and Features; D.4.6 [Operating Systems]: Security and Protection—*Information flow controls*

General Terms: Languages, Performance, Security

Additional Key Words and Phrases: Information flow control, Java virtual machine, operating systems, security method

ACM Reference Format:

Donald E. Porter, Michael D. Bond, Indrajit Roy, Kathryn S. McKinley, and Emmett Witchel. 2014. Practical fine-grained information flow control using Laminar. *ACM Trans. Program. Lang. Syst.* 37, 1, Article 4 (October 2014), 50 pages.
 DOI: <http://dx.doi.org/10.1145/2638548>

This research was supported in part by NSF grants CAREER-0644205, CNS-0905602, CAREER CNS-1149229, CNS-1161541, CNS-1228839, CNS-1228843, CAREER-1253703, CSR-1218695, SHF-0910818, NIH LM011028-01, and the Office of the Vice President for Research at Stony Brook University.

Authors' addresses: D. E. Porter, Computer Science Department, Stony Brook University; M. D. Bond, Department of Computer Science and Engineering, Ohio State University; I. Roy, Hewlett Packard Labs; K. S. McKinley, Microsoft Research, Redmond; E. Witchel, Department of Computer Science, The University of Texas at Austin.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2014 ACM 0164-0925/2014/10-ART4 \$15.00
 DOI: <http://dx.doi.org/10.1145/2638548>

	1
	2
	3
	4
	5
	6
4	7
	8
	9
	10
	11
	12
	13
	14
	15
	16
	17
	18
	19
	20
	21
	22
	23
	24
	25
	26
	27
	28
	29
	30
	31
	32

Q1

33 **1. INTRODUCTION**

34 As computer systems support more aspects of modern life, from finance to health
35 care to energy, the security of these systems becomes increasingly important. Current
36 security policies and enforcement mechanisms are typically sprinkled throughout an
37 application, making security policies difficult to express, change, and audit. Operating
38 system security abstractions, such as file permissions and user IDs, are too coarse
39 to express many desirable policies, such as protecting a user's financial data from a
40 program downloaded from the Internet.

41 Furthermore, poor integration of Programming Language (PL) constructs and Oper-
42 ating System (OS) security mechanisms complicates the expression and enforcement
43 of security policies. For example, a policy against sending a user's credit card number
44 on the network should be enforced whether the number originates from a file or an
45 application data structure. In current systems, the OS governs the security of files,
46 and application-specific logic governs the security of data structures; because these
47 mechanisms are completely distinct, developers must understand both mechanisms
48 and ensure that they interoperate correctly. This article describes Laminar, which
49 integrates PL and OS security mechanisms under a common set of programmer ab-
50 stractions and uniformly enforces programmer-specified security policies at all levels
51 of the software stack.

52 Laminar builds on the Decentralized Label Model (DLM) [Myers and Liskov 1997],
53 which expresses more powerful, sophisticated, and intuitive security policies than tra-
54 ditional security models. The enforcement of DLM restrictions is called Decentralized
55 Information Flow Control (DIFC). DIFC is more expressive than traditional access
56 control. For instance, traditional access control models are all-or-nothing; once an ap-
57 plication has the right to read a file, it can do anything with that file's data. In contrast,
58 a DIFC policy may give an application the right to read a file and simultaneously for-
59 bid it to broadcast the contents of the file over an unsecured network channel. A DIFC
60 implementation dynamically or statically enforces user-specified security policies by
61 tracking information flow throughout the system.

62 In the decentralized label model, users create *tags*, which represent secrecy or in-
63 tegrity concerns. A set of tags is called a *label*, and all data and application threads
64 have an associated secrecy label and an integrity label. The system restricts the flow
65 of information according to these labels. Secrecy guarantees prevent sensitive infor-
66 mation from escaping the system (no illegal reads),¹ and integrity guarantees prevent
67 external information from corrupting the system (no illegal writes).

68 As an example, suppose Alice and Bob want to schedule a meeting without disclosing
69 other appointments on their calendars. In the DLM model, Alice and Bob each place
70 a tag in the *secrecy label* on their calendar files. Alice and Bob can give the calendar
71 application permission to read these files but only if the application *taints* its own
72 secrecy label with the secrecy tags of each file. A tainted application thread may no
73 longer write to less-secret outputs, such as the terminal or the network. In our example,
74 the tainted thread may read each calendar file and select an agreeable meeting time,
75 but the thread can only write output to a file or data structure labeled with both Alice's
76 and Bob's secrecy tags. In order for the calendar application to output a nonsecret
77 meeting time, Alice and Bob must provide a *declassifier* with the capability of removing
78 their tags from a datum's secrecy label. The declassifier is a piece of code responsible
79 for checking that its output conforms to a secrecy policy associated with a tag; the
80 declassifier may write acceptable data to a less-secret output. In the calendar example,
81 Alice and Bob might both provide declassifiers; each declassifier can generate output

¹The literature uses both *confidentiality* and *secrecy* for this guarantee. We use *S* for secrecy, *I* for integrity, and *C* for capabilities to avoid ambiguity.

without that user’s tag in the secrecy label. For instance, Bob’s declassifier might read the labeled meeting time and check that the output is simply a date and does not include mention of his upcoming vacation to Las Vegas. Note that DIFC exists in addition to traditional access control; for example, a web server would not be allowed to open either calendar file due to standard OS-level permission checks.

Similarly, Alice may use an *integrity label* on her calendar file to ensure that any updates to the file respect certain invariants. Suppose Alice’s calendar is stored as a chronologically sorted list of appointments. Untrusted code that adds appointments to Alice’s calendar might serialize her appointments into the on-disk format and store the pending data in a memory buffer. Alice could then run this buffer through an *endorser*, which ensures that the pending data write meets the specifications of her calendar format, such as checking that all appointments are sorted chronologically. Just as secrecy labels can be removed from the output of a computation by a declassifier, an endorser is trusted with the capability to add a tag to the integrity label of inputs that it validates. Once the endorser has validated that its input is trustworthy, the endorser adds Alice’s integrity tag to its integrity label and writes a new version of her calendar file.

DIFC provides two key advantages: precise rules for the legal propagation of data through an application and the ability to localize security policy decisions. In the calendar example, the secrecy labels ensure that any program that can read the data cannot leak the data, whether accidentally or intentionally. The label is tied to the data, and the label modulates how data can flow through threads and data containers (e.g., files and data structures). The decision to declassify data is localized to small pieces of code that programmers may closely audit. The result is a system where security policies are easier to express, maintain, and modify than with traditional security models.

Combining the Strengths of Language and Operating System Enforcement. Laminar is a new DIFC system design that features a common security abstraction and labeling scheme for program objects and OS resources such as files and sockets. The Java Virtual Machine (VM) and OS coordinate to comprehensively enforce rules within an application, among applications, and through OS resources.

Prior DIFC systems are implemented at the language level [Chandra and Franz 2007; Myers and Liskov 1997; Myers et al. 2001; Nair et al. 2008] in the operating system [Krohn et al. 2007; Vandebogart et al. 2007; Zeldovich et al. 2006], or in the architecture [Tiwari et al. 2009a; Vachharajani et al. 2004; Zeldovich et al. 2008]. Each approach has strengths and limitations. Language-based DIFC systems can track information flow through data structures within a program but have little visibility into OS-managed resources, such as files and pipes. In contrast, OS-based DIFC systems track labels at the coarse granularity of pages or a process’s virtual address space rather than on individual data structures. Information flow rules are enforced on OS-level abstractions, such as sockets and files. For many simple applications, these coarse-grained rules simplify DIFC adoption. However, OS protection mechanisms are not a good fit for managing information flow on data structures within an application because the OS’s primary tool is page-level protections. Although an application developer could group objects with similar labels on similarly labeled pages, this undermines developer productivity and application efficiency. Thus, we believe that coordinating language and OS mechanisms will maximize security and programmability.

We limit the scope of this article to DIFC implementations on commodity hardware. Architecture-based solutions track data labels on various low-level hardware features, such as CPU registers, memory, cache lines, or even gates, but require similar coordination with trusted software to manage the labels.

Language-based DIFC systems can be further categorized by how they enforce DIFC rules: static analysis [Myers and Liskov 1997; Myers et al. 2001], dynamic analysis

133 [Shroff et al. 2007], or a hybrid [Chandra and Franz 2007; Nair et al. 2008]. Static
 134 systems generally introduce a type system that is expressive and powerful but difficult
 135 to program or retrofit onto existing code. Because static systems do most security anal-
 136 ysis at compile time, they introduce little runtime overhead; static systems may insert
 137 dynamic checks for properties that cannot be established at compile time. Dynamic sys-
 138 tems generally enforce information flow rules by mediating every operation at runtime
 139 but with relatively high performance overheads. Purely dynamic systems also struggle
 140 to regulate implicit flows (discussed further in Section 6.4) and can ultimately reject
 141 safe programs or leak sensitive data [Russo and Sabelfeld 2010].

142 Most language-level systems are actually a hybrid of static and dynamic analysis.
 143 Each design strikes a balance among changes to the programming language to facili-
 144 tate static analysis, runtime overheads, and security guarantees. The Laminar design
 145 restricts the programming model slightly, ensuring that all security properties can be
 146 checked dynamically. Laminar does employ intraprocedural static analysis at Just-in-
 147 Time (JIT) compilation time to optimize security checks.

148 *Limiting the Scope of Analysis.* A second key contribution of Laminar is the design of
 149 a language-level feature, called a *security method*, which strikes a unique balance be-
 150 tween programmability and efficiency. Developers place all security-sensitive program
 151 logic in security methods. The Laminar VM requires that all operations on labeled data
 152 or system resources occur within security methods, according to developer-specified
 153 policies. In addition, all methods dynamically invoked by a security method, directly or
 154 transitively, are security methods. Code that attempts to manipulate security-sensitive
 155 data outside of a security method will fail.

156 Laminar enforces stringent requirements on transitions to and from security meth-
 157 ods, restricting both control and data flow. These restrictions are enforced dynamically
 158 by VM instrumentation. Security methods reduce the overhead of dynamic security
 159 checks because only code within security methods requires complex DIFC checks.

160 Security methods also minimize the code changes required to adopt DIFC. In our
 161 case studies, changes to adopt security methods account for 10% or fewer of the total
 162 lines of code, which suggests that pervasive program modifications are unnecessary to
 163 use DIFC with Laminar.

164 *Contributions.* The contributions of this article are as follows:

- 165 (1) We present the design and implementation of Laminar, the first system to unify
 166 PL and OS mechanisms for enforcing DIFC. Laminar features a novel division of
 167 responsibilities between the VM and OS.
- 168 Q2 (2) We introduce security methods, an intuitive security primitive that reduces the
 169 work required to convert an application to use DIFC, makes code auditing easier,
 170 and makes the DIFC implementation simpler and more efficient.
- 171 (3) We present the design and implementation of Laminar in the Linux OS and Jikes
 172 RVM, a Java research VM.
- 173 (4) We evaluate four case studies that retrofit security policies onto existing code.
 174 These case studies require modification of less than 10% of the total code base and
 175 incur overheads from Laminar ranging from 5% to 56%.
- 176 (5) Based on our experiences, we substantially modified the conference publication
 177 that introduced this research [Roy et al. 2009]. We replace security regions with
 178 security methods to simplify our implementation. We use only dynamic analy-
 179 sis to simplify the enforcement security policies. We identify and fix a covert
 180 channel bug arising from the interaction of termination and concurrency. Fur-
 181 thermore, we improve the programming model for initializing and using security
 182 labels.

Practical Fine-Grained Information Flow Control Using Laminar	4:5
(6) We describe strengths and limitations of the Laminar model, its open challenges, and potential solutions. In particular, Laminar is one of the few DIFC systems to attempt multithreading support, which is prone to high-bandwidth timing channels. Laminar cannot prevent all of these timing channels, but we outline how subsequent work by others [Askarov et al. 2010; Askarov and Myers 2012; Zhang et al. 2011] could strengthen the Laminar threading model.	183 184 185 186 187 188
Initial results suggest that integrating PL and OS DIFC enforcement is practical and incurs low overheads. Our experience with Laminar shows that it prevents some termination information channels, but it cannot yet make guarantees on some timing channels. We believe restrictions on the programming model within security methods can solve some of these problems, but this research leaves open the definition of such a formalism and accompanying proofs. Laminar provides a first step for application developers to write expressive abstractions with fine-grained, powerful, and useful security policies that span program data structures and system resources.	189 190 191 192 193 194 195 196
2. DIFC MODEL	197
This section describes how the DIFC model specifies and enforces safe information flows and how Laminar embodies the DIFC model. All DIFC systems denote the sensitivity of information and the privileges of the participating users, as well as describe application-specific policies that map between users and sensitive information. The security policy is defined in terms of <i>principals</i> that read and write data in the system. Examples of principals in DIFC systems are users [Myers et al. 2001], processes [Krohn et al. 2007], and kernel threads [Zeldovich et al. 2006]. Principals in Laminar are kernel threads, which ultimately work on behalf of human users or other application-level actors.	198 199 200 201 202 203 204 205 206
2.1. DIFC Abstractions	207
Standard DIFC security abstractions include tags, labels, and capabilities. Tags are short, arbitrary tokens drawn from a large universe of possible values (\mathcal{T}) [Krohn et al. 2007]. Programmers use tags to denote a unique secrecy or integrity property, but a tag has no inherent meaning. Programmers may create tags dynamically and may persist tags beyond execution of an application. A set of tags is called a label. In a DIFC system, any principal can create a new tag for secrecy or integrity. For example, a web application might create one secrecy tag for its user database and a separate secrecy tag for each user’s data. The secrecy tag on the user database will prevent authentication information from leaking to the network. The tags on user data will prevent a malicious user from writing another user’s secret data to an untrusted network connection.	208 209 210 211 212 213 214 215 216 217
Principals assign labels to data objects. Data objects include program data structures (e.g., individual objects, arrays, lists, and hash tables) and system resources (e.g., files and sockets). Previous OS-based systems limit principals to the granularity of a process or support threads by enforcing DIFC rules at the granularity of a page. Laminar is the first to support threads as principals and enforce DIFC at object granularity.	218 219 220 221 222
Each data object and principal x has two labels, S_x for secrecy and I_x for integrity. A tag t in the secrecy label S_x of a data object denotes that it may contain information private to principals with tag t . Similarly, a tag u in I_x indicates that the owner of integrity tag u endorses the data. Data integrity is a guarantee that data exist in the same state as when they were endorsed by a principal. For example, if Microsoft endorses a data file, then a user can choose to trust the file’s contents if she trusts Microsoft. With integrity enforcement, only Microsoft may modify the integrity-labeled file. However, Microsoft may choose to remove the integrity label, or some other application may write the file, but without the Microsoft integrity label. In either case, the file’s consumer will	223 224 225 226 227 228 229 230 231

232 no longer trust the contents as coming from Microsoft. In general, a principal's labels
233 restrict the interaction that the principal has with other principals and data objects.

234 A partial ordering of labels imposed by the subset relation forms a lattice [Denning
235 1976]. Secrecy and integrity may be treated separately, as asymmetric duals. The
236 bottom of the secrecy lattice is the least restricted label (public): Any principal can
237 read it. The bottom of the integrity lattice is the most secure (trusted): All principals
238 can trust it. Adding secrecy tags to a label restricts the use of the data, moving higher in
239 the lattice. The most restricted data at the top of the secrecy lattice includes all secrecy
240 tags. The bottom of the integrity lattice is the most secure (trusted) and includes all
241 integrity tags. Removing integrity tags moves the label higher in the lattice and the
242 data are less trusted. The top of the integrity lattice has no integrity tags—no principal
243 endorses it.

244 Some other DIFC explanations put an empty integrity label at the bottom of the
245 lattice so that adding tags moves up the lattice, as opposed to the preceding description
246 that places the label with all integrity tags at the bottom, so that moving up the lattice
247 adds restrictions. Both representations are functionally equivalent. For clarity, this
248 article generally discusses the secrecy and integrity labels separately, but occasionally
249 some explanations treat principals and data as having a single label and capability set
250 for ease of exposition.

251 Because Laminar's threat model includes code that may be contributed to the appli-
252 cation by an adversary, all application data are assigned an empty label (public and
253 untrusted) by default. Data from the JVM itself are public and trusted. The program-
254 mer need not label every data structure, nor does the OS need to label every file in the
255 file system. Code that executes with a nonempty integrity label must sanitize untrusted
256 data before a read. Default empty labels make Laminar easier to deploy incrementally,
257 but introduce some asymmetry in the treatment of secrecy and integrity tags.

258 A principal may change the label of a data object or principal if and only if the
259 principal has the appropriate capabilities, which generalize ownership of tags [Myers
260 and Liskov 1997]. A principal p has a capability set C_p that defines whether the
261 principal has the privilege to add or remove a tag. For each tag t , let t^+ and t^- denote
262 the capabilities to add and remove the tag t .

263 If tag t is used for secrecy, a principal with the capability t^+ may *classify* data
264 with secrecy tag t . Classification raises data to a higher secrecy level. Given the t^-
265 capability, a principal may *declassify* these data. Declassification lowers the secrecy
266 level. Principals may add t to their secrecy label if they have the t^+ capability. If the
267 principal adds t , then we call it *tainted* with the tag t . A principal taints itself when it
268 wants to read secret data. To communicate with unlabeled devices and files, a tainted
269 principal must use the t^- capability to untaint itself and to declassify the data it wants
270 to write. Note that DIFC capabilities are not pointers with access control information,
271 which is how they are commonly defined in capability-based operating systems [Levy
272 1984; Shapiro et al. 1999].

273 DIFC handles integrity similarly to secrecy. A principal with integrity tag t is claim-
274 ing to represent a certain level of integrity; the system prevents the principal from
275 reading data with a lower integrity label, which could undermine the integrity of the
276 computation. Given the t^+ capability, a principal may *endorse* data with integrity tag
277 t , generally after validating that the input data meet some requirements. Given the t^-
278 capability, a principal may drop the endorsement and read untrusted data. For exam-
279 ple, code and data signed by a software vendor could run with that vendor's integrity
280 tag. If the program wants to load an unlabeled, third-party extension, the principal
281 drops the endorsement of the tag.

282 Note that the capability set C_p is defined on tags. A tag can be assigned to a secrecy
283 or integrity label. In practice, a tag is rarely used for both purposes. C_p^- is the set of

tags that principal p may declassify (drop endorsements), and C_p^+ is the set of tags that p may classify (endorse). Principals and data objects have both a secrecy and integrity label; a data object with secrecy label s and integrity label i is written: $\langle S(s), I(i) \rangle$. An empty label set is written: $\langle S(), I() \rangle$. The capability set of a principal that can add both s and i but can drop only i is written: $\langle C(s^+, i^+, i^-) \rangle$.

2.2. Restricting Information Flow

Programs implement policies to control access and propagation of data by using labels to limit the interaction among principals and data objects. Information flow is defined in terms of data moving from a source x to a destination y , at least one of which is a principal. For example, principal x writing to file y , principal x sending a message to principal y , and principal y reading a file x are all information flows from x to y . If principal x writes to a file y , then we say information flows from source x to destination y . Laminar enforces the following information flow rules for x to y :

Secrecy rule. Bell and LaPadula introduced the simple security property and the *-property for secrecy [Bell and LaPadula 1973]. The simple security property states that no principal may read data at a higher level (*no read up*), and the *-property states that a principal may not write data to a lower level (*no write down*). Expressed formally, information flow from x to y preserves secrecy if:

$$S_x \subseteq S_y$$

Note that x or y may make a flow feasible by using their capabilities to explicitly drop or add a label. For example, x may make a flow feasible by removing a tag t from its label S_x if it has the declassification capability for t (i.e., $t^- \in C_x^-$). Similarly, y may use its capabilities in C_y^+ to extend its secrecy label and receive information.

Integrity rule. The integrity rule constrains who can alter information and restricts reads from lower integrity (*no read down*) and writes to higher integrity (*no write up*) [Biba 1977]. Laminar enforces the following rule:

$$I_y \subseteq I_x$$

Intuitively, the integrity label of x should be at least as strong as destination y . Just like the secrecy rule, x may make a flow feasible by endorsing information sent to a higher integrity destination, which is allowed if x has the appropriate capability in C_x^+ . Similarly, y may need to reduce its integrity level, using C_y^- , to receive information from a lower integrity source.

Label changes. According to the previous two rules, a principal can enable information flow by using its current capabilities to drop or add tags from its label. Laminar requires that the principal must *explicitly* change its current labels. Zeldovich et al. show that automatic, or implicit, label changes can form a covert storage channel [Zeldovich et al. 2006].

In Laminar, a principal p may change its label from L_1 to L_2 if it has the capability to add tags present in L_2 but not in L_1 , and can drop the tags that are in L_1 but not in L_2 . This is formally stated as:

$$(L_2 - L_1) \subseteq C_p^+ \text{ and } (L_1 - L_2) \subseteq C_p^-.$$

2.3. Calendar Example

Again, consider scheduling a meeting between Bob and Alice using a calendar server that is not administered by either Alice or Bob. Alice’s calendar file has a secrecy tag, a , and integrity tag i ; Bob’s calendar file has a secrecy tag, b .

326 *Ensuring secrecy.* Focusing on Alice, she gives a^+ to the scheduling server to let it
 327 read her secret calendar file, which has label $\langle S(a) \rangle$. A thread in the server uses the
 328 a^+ capability to start a security method with secrecy tag a that reads Alice's calendar
 329 file. Once the server's thread has the label $\langle S(a) \rangle$, it can no longer return to the empty
 330 label because it lacks the declassification capability, a^- . As a result, the server thread
 331 can read Alice's secret file, but it can never write to an unlabeled device like the disk,
 332 network, or display. If the server thread creates a new file, it must have label $\langle S(a) \rangle$,
 333 which is unreadable to its other threads. Before the server thread can communicate
 334 information derived from Alice's secret file to another thread, the other thread must
 335 add the a tag, and it also becomes unable to write to unlabeled channels.

336 *Ensuring integrity.* Alice also places an integrity label on her calendar file, which
 337 is propagated to the heap data structures representing her calendar. In order for any
 338 thread to update Alice's calendar, the thread must add the i integrity tag to its label.
 339 In general, the capability to add this tag would be restricted to code that is trusted to
 340 check that inputs or updates uphold application invariants. In this example, much of
 341 the calendar code may run with an empty integrity label, but once a meeting request
 342 is ready to be added to Alice's calendar, the meeting request is checked by Alice's
 343 endorser. If the checks pass, Alice's endorser adds the i tag to the meeting request data
 344 structure. The code that writes the updated calendar to disk must also run with the i
 345 tag, preventing data from untrusted heap objects from inadvertently being written to
 346 the calendar file.

347 *Sharing secrets with trusted partners.* Alice and Bob collaborate to schedule a meeting
 348 while both retain fine-grained control over what information is exposed. After the
 349 scheduler has read Alice's and Bob's calendar files, the output data are labeled with the
 350 a and b secrecy tags. Alice's module has access to her a^- capability, so the server calls
 351 her code, which validates that the output does not disclose unintended information
 352 to Bob. Alice's module then removes the a tag from the output data, publishing the
 353 meeting time to Bob. Alice controls which of her data flow out of the scheduler. Bob
 354 does the same, and the scheduler can communicate with both of them and coordinate
 355 their possible meeting times.

356 *Discussion.* In this example, Alice specifies a declassifier as a small code module
 357 that can be loaded into a larger server application, which can be completely ignorant
 358 of DIFC and requires no modifications to work with Alice's DIFC-aware module. For
 359 previous DIFC systems, this example is more cumbersome. OS-based DIFC systems
 360 require the declassifier to run as a separate process. Language-based DIFC systems re-
 361 quire programmers to annotate the entire application. By integrating OS and language
 362 techniques, Laminar simplifies incremental DIFC adoption.

363 2.4. Goals and Threat Model

364 This subsection describes our threat model and its rationale at a high level. We revisit
 365 these security properties in Section 6, after describing our system design and imple-
 366 mentation. Section 10 surveys related work in more detail, but here we summarize key
 367 categories of DIFC systems and challenges in DIFC adoption. DIFC systems can be
 368 roughly categorized by how they enforce flows: static analysis, dynamic language-level
 369 analysis, or OS-level enforcement.

370 *Incremental Adoption.* A key design goal of Laminar is facilitating incremental adop-
 371 tion of DIFC on a large body of code. The ease with which a programmer can adopt
 372 DIFC is an issue for most DIFC designs. DIFC based on static analysis often re-
 373 quires substantial annotations of the program with a new type system. OS-based DIFC

requires substantial reorganization of the application code in order to segregate data pages and code by label. It is unclear whether a language-level dynamic analysis is any easier to adopt. Although there has been some work in this area, it has generally enforced only simple policies on outputs [Chandra and Franz 2007] or had problems with “label creep,” which requires error-prone, manual analysis by the programmer [Nair et al. 2008].

The insight underlying Laminar’s security-method-based design is that many applications already handle sensitive data only in relatively small portions of their code. For instance, web server authentication code is generally small relative to all of the code that generates and transmits web content. Thus, Laminar is designed so that the programmer audits only these relatively small portions of preexisting code for correct handling of sensitive data. Sensitive code is placed in security methods, and the system dynamically checks that all information flows according to the restrictions imposed by the developer and end users.

Laminar enforces DIFC rules using a combination of dynamic analysis and programmer annotations (i.e., security methods). Compared to other dynamic or hybrid language-level systems, Laminar is generally more efficient than previous systems because of careful implementation choices and limited scope of analysis. As discussed earlier, all DIFC systems require some measure of work to adopt, and our experience is that security methods minimize the effort without sacrificing functionality.

Integration of OS and PL Abstractions. Laminar integrates OS and PL DIFC abstractions to implement uniform policies and label management across resources. Existing systems cannot easily integrate these abstractions. For instance, a PL system might enforce all-or-nothing policies about output or might make educated guesses about information flow through OS abstractions, but it cannot ensure that these rules are followed once data leaves the application.

Threading. A key aspect of incremental deployability is tracking information flow through a program, including with multiple, concurrent threads. OS systems mediate multiprocess concurrency through explicit channels and at page granularity. In practice, these systems cannot track fine-grained information flow through traditional thread packages without major modifications to the application. PL systems have generally avoided multithreading because it increases the risks of covert channels. Laminar does permit multithreading, but cannot prevent all timing channels attacks. This article identifies some threats and points to solutions developed after the initial publication of this work [Roy et al. 2009] that could be integrated into security methods to mitigate these channels.

Threat Model. In a DIFC system, the primary concern is limiting the ability of one principal to access another principal’s data. So, in our threat model, the attacker may have contributed code to the application and is executing as principal (thread) A. Laminar does not allow principal A to explicitly read or write another principal B’s data (e.g., by explicit assignment in the program) without acquiring appropriate secrecy and integrity labels. Any other user controls access to her data by controlling which principals she gives the capabilities to add and remove tags associated with her data.

Limitations. Like most DIFC systems, the Laminar VM and OS mediate all explicit assignments of labeled data, as described in Sections 4 and 5. Laminar prevents implicit information flows by restricting the visibility that untrusted code has into the control flow of a security method, including restrictions on input and output variables (discussed further in Section 6.4).

Eliminating all timing, termination, and other covert channels are open problems [Denning and Denning 1977; Lampson 1973] and beyond the scope of this article.

424 In particular, it is well established that preventing all these channels on a general-
425 purpose programming model is tantamount to solving the halting problem [Denning
426 and Denning 1977]. In order to eliminate information leaks due to unbounded execu-
427 tion, more recent work has investigated highly restricted programming models (e.g.,
428 without unbounded loops [Tiwari et al. 2009b]) or bounding the execution time of code
429 that manipulates sensitive data [Tiwari et al. 2009a].

430 To facilitate incremental adoption, Laminar places capabilities in threads, rather
431 than statically mapping them to functions. The underlying tradeoff is that the program-
432 mer can more easily invoke standard libraries from a security method. For example,
433 programmers may therefore manipulate secure objects using standard implementa-
434 tions of arrays, lists, and sets. Code invoked from a security method executes as if it were
435 in the security method. This choice introduces some risk for a confused deputy prob-
436 lem and requires trusting the caller of a security method to manage capabilities. The
437 Laminar design mitigates the risk of capability management errors by requiring that
438 all endorsers and declassifiers be declared `final` and that non-endorser/declassifier
439 security methods do not accept capabilities as arguments. These issues are discussed
440 further in Section 6.6.

441 Two key innovations of Laminar are support for multiple threads and the ability of
442 a single thread to transition between different trust levels—facilitating incremental
443 adoption but also introducing new opportunities for covert channels based on the timing
444 of these transitions. Section 6 describes the new classes of timing and termination
445 channels that these features could introduce and how Laminar mitigates them. To
446 summarize, Laminar restricts the ability to create a channel based on control flow
447 within a security method by requiring a single exit point from a security method
448 and carefully mediating any OS- or VM-level storage channel, such as the thread's
449 capabilities. The article also discusses how more recent work, such as predictive timing
450 models [Askarov et al. 2010], could be applied to the Laminar prototype to further
451 reduce covert channels, especially through thread synchronization. These issues are
452 discussed further in Section 6.

453 3. DESIGN

454 This section describes the Laminar programming model and how Laminar enforces
455 DIFC in an enhanced VM and OS.

456 3.1. Overview

457 Figure 1 illustrates the Laminar architecture. The OS kernel reference monitor medi-
458 ates accesses to system resources. The VM enforces DIFC rules within the application's
459 address space. Only the OS kernel and VM are in the Laminar trusted computing base.
460 The OS kernel and VM trust each other as well.

461 The Laminar OS kernel extends a standard OS kernel with a Laminar security mod-
462 ule for information flow control. Users and programmers invoke the Laminar kernel
463 security APIs to create tags, store capabilities for their tags, and label their data in
464 files. Users launch processes with a subset of their tags and capabilities. The Laminar
465 OS kernel governs information flows through all standard kernel interfaces, including
466 through devices, files, pipes, and sockets. DIFC rules are enforced by the kernel on
467 all threads, whether the threads are of the same or different processes. Resources and
468 principals without an explicit label have empty secrecy and integrity labels, facilitating
469 incremental adoption. Our prototype uses the Linux Security Modules [Wright et al.
470 2002] framework, although the design could be extended to any OS that provides simi-
471 lar hooks to an in-kernel reference monitor to label kernel objects and mediate system
472 calls that could create an information flow.

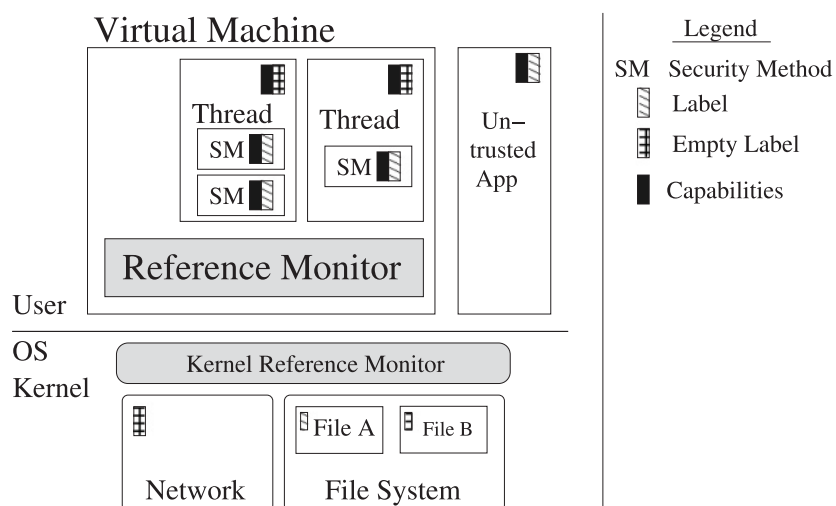


Fig. 1. Design of Laminar. An OS kernel reference monitor and VM reference monitor enforce information flow control. All data is labeled, including objects in the VM, as well as OS abstractions, such as files and sockets. Objects without explicit labels default to empty secrecy (public) and integrity (untrusted) labels. Threads have capabilities and an empty label. A thread enters a Security Method (SM) to acquire a nonempty label. A security method may optionally receive capabilities from the calling thread.

To regulate information flows within an application, Laminar extends the runtime system of a standard Java VM. By default, the Laminar OS kernel requires all threads within a process to have the same secrecy and integrity labels. The OS relaxes this restriction for threads running on the trusted Laminar VM. The Laminar VM binary is labeled with a special TCB integrity tag, which indicates to the OS that this application is trusted to control information flows within its address space. Although the kernel trusts the Laminar VM to regulate flows within the address space, the kernel still checks all accesses to system resources.

The Laminar VM regulates information flow between heap objects within a thread and between threads of the same process via these objects. The Laminar VM inserts dynamic DIFC checks to regulate DIFC flows.

The key abstraction for manipulating labeled data is the *security method*. Programmers explicitly declare security methods. In addition, any method invoked directly or transitively from a declared security method is implicitly defined as a security method. Outside of a security method, a thread has empty secrecy and integrity labels and may only read or write data with empty secrecy and integrity labels. The VM terminates the program if it attempts to read or write any labeled data outside a security method. If a thread has the capability to add a tag to its secrecy or integrity labels, the thread may change its labels by entering a security method. Within a security method, a thread may read or write data with nonempty labels as long as the reads and writes constitute a legal information flow according to the capabilities specified by the parameters. A thread typically runs with a subset of the user’s capabilities, and a security method specifies a subset of the thread’s capabilities. Security methods may nest. Each security method may only have a subset of the parent security method’s capabilities and may only change its labels as permitted by the parent’s capabilities.

For example, Alice writes a program in Java with the Laminar programming model and uses the Laminar API (see next paragraphs and Table I). Alice compiles the code using a standard, untrusted bytecode generator such as `javac`. The Laminar JIT compiler and VM execute the bytecode, and the Laminar OS kernel executes the Laminar

4:12

D. E. Porter et al.

Table I. Laminar API

These functions manipulate labels and capabilities. `LabelType` denotes whether a request is for a secrecy label or integrity label. `CapType` denotes plus, minus, or both capabilities for a tag. The `getCurrentLabel` and `copyAndLabel` functions may be invoked inside of a security method, but tags and capabilities may only be created and destroyed outside of a security method, using `createAndAddCapability` and `removeCapability`, respectively. The API has wrapper functions (not shown) for the new Laminar system calls. `Label` stores a set of secrecy or integrity tags (Section 5.1).

Function	Description
<code>Label getCurrentLabel(LabelType t)</code>	Return the current secrecy or integrity label of the security method as an opaque object. Label objects cannot be enumerated.
<code>Object copyAndLabel(Object o, Label S, Label I)</code>	Return a copy of the object <code>o</code> with new secrecy label <code>S</code> and integrity label <code>I</code> .
<code>CapSet getCurrentCapabilities()</code>	Return the current capability set of the thread as an opaque object.
<code>Label createAndAddCapability()</code>	Create a new tag and add capabilities to the current thread. Must be used outside of a security method.
<code>void removeCapability(CapType c, Label l)</code>	Drop the capabilities listed in <code>c</code> (plus and/or minus) associated with the tags in <code>l</code> from the current thread. Must be used outside of a security method.

502 VM. During execution, the program labels data with security and integrity tags that it
 503 obtains from the kernel API. The OS kernel and VM thus use the same tag namespace
 504 for the system resources and objects. For example, the application reads data from a
 505 labeled file into a data structure with the same labels. The Laminar VM ensures that
 506 any accesses or modifications to labeled data follow the DIFC rules and occur in a
 507 security method, a labeled method specified by the programmer.

508 Restricting security policies to security methods makes it easier to add security
 509 policies to existing programs. Furthermore, auditing security methods will generally be
 510 easier than auditing the entire program. These features should facilitate incremental
 511 deployment of Laminar in existing systems. Security methods also decrease the cost of
 512 performing dynamic security checks in the Laminar runtime.

513 Laminar does not track information flows through local variables. Because labeled
 514 data are manipulated in security methods, locals in an untrusted parent are out of
 515 scope inside the security method and vice versa. With additional static analysis on
 516 information flow through locals, one might be able to safely implement security methods
 517 as arbitrary, lexically scoped regions, as originally proposed [Roy et al. 2009]. We
 518 expect that the additional static analysis required to support lexically scoped regions
 519 would be easiest to implement in the Java compiler (`javac`), but these properties might
 520 also be checked by the JVM during bytecode verification. We found mediating flows
 521 through locals at method boundaries to strike a good balance between implementation
 522 complexity for the application programmer and JVM developer.

523 3.2. Programming Model, VM, and OS Interaction

524 Laminar provides language extensions, a security library, and security system calls.
 525 Table I depicts the Java APIs, which include methods that perform tag creation, de-
 526 classification, label queries, and capability queries. The Laminar OS kernel exports
 527 security system calls to the trusted VM for capability and label management, as shown
 528 in Table II. These system calls are used by the Laminar VM internally to implement
 529 security methods and are not directly exposed to Laminar applications. An application
 530 not running on the Laminar VM may directly use these system calls to manage its
 531 capabilities and labels, excluding `drop_label_tcb`, which can only be issued by the
 532 trusted Laminar VM. The Laminar OS securely stores all of the persistent capabilities
 533 of a user so that these capabilities can be used across user sessions. On login, the OS
 534 kernel gives the capabilities of the user to the login shell. Laminar does not innovate in

Table II. Laminar System Calls

The `tag_t` and `capability_t` types represent a single tag or capability, respectively. The `struct label` type represents a set of tags that compose a label, and the `capList_t` type is a list of capabilities.

System Call	Description
<code>tag_t alloc_tag(capList_t &caps)</code>	Return a new tag, add plus and minus capabilities to the calling principal, and write new capabilities into <code>caps</code> .
<code>int add_task_tag(tag_t t, int type)</code>	Add a tag to the current principal's secrecy or integrity label (secrecy or integrity selected by <code>type</code>), as allowed by the principal's capabilities.
<code>int remove_task_tag(tag_t t, int type)</code>	Remove a tag from the current principal's secrecy or integrity label (secrecy or integrity selected by <code>type</code>).
<code>int drop_label_tcb(pid_t tid)</code>	Drop the current temporary labels of the thread without capability checks, invoked only by threads with the special integrity tag.
<code>int drop_capabilities(capList_t *caps, int tmp)</code>	Drop the given capabilities from the current principal. <code>tmp</code> is a flag used by the VM to suspend a capability only for a security method or during a <code>fork()</code> .
<code>int write_capability(capability_t cap, int fd)</code>	Send a capability to another thread via a pipe.
<code>int create_file_labeled(char* name, mode_t m, struct label *S, struct label *I)</code>	Create a labeled file with labels <code>S</code> and <code>I</code> .
<code>int mkdir_labeled(char* name, mode_t m, struct label *S, struct label *I)</code>	Create a labeled directory with labels <code>S</code> and <code>I</code> .

capability persistence but rather adopts a simple and stylized model. Asbestos develops a more robust model for persistent storage of tags [Vandebogart et al. 2007]. 535

The secure keyword applies to methods used as security methods. The VM and kernel enforce the rule that the program may only access labeled data objects (e.g., files, heap-allocated objects, arrays) inside security methods, which includes all methods directly or transitively invoked from a declared security method. Outside security methods, threads always have empty labels but may hold capabilities that determine whether the thread may enter a security method. Threads are the only principals in Laminar, and the VM modifies the thread's labels and capabilities when it enters and exits a security method. When a thread enters a security method, it dynamically passes the desired secrecy label and integrity label as arguments to the method, using the opaque `Label` object. If the security method endorses or declassifies data, it may also accept the necessary capabilities as an argument, as a `CapSet` object. During the execution of a security method, the VM internally uses these labels and capabilities for DIFC enforcement, and the kernel mediates thread accesses to system resources according to the security method's labels. Because security methods are not visible to the kernel, the VM proxies the security method by tainting the thread with the correct labels and capabilities. At the end of the security method, the VM restores the thread's original capabilities and labels. 536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553

3.3. Security Methods 554

A *security method* is a special method type that has parameters for a secrecy and integrity label. A security method that can endorse or declassify data also has a parameter for a capability set. The labels dictate which data the program may touch inside the security method. Labels on secure method parameters must satisfy the data flow constraints of the labels on the security method, and the label on the returned data must satisfy the data flow equations for the labels of the calling context. In the Laminar implementation, these labels and capabilities are represented as sets that can be variably sized and assigned at runtime. Label and capability sets are stored as opaque objects, which cannot be enumerated (see Section 5.1). 555
556
557
558
559
560
561
562
563

564 Only code within a security method can access data with nonempty labels. Security
565 methods demarcate the methods that are security sensitive, thus easing the program-
566 mer's burden when adding security policies to existing programs. The programmer
567 must place all code that references labeled data in a security method, such as a routine
568 that reads a sensitive file into a data structure. In our experience, only a small portion
569 of code and data in a program is security sensitive and will belong in a security method,
570 which simplifies the task of auditing security-sensitive code. This design also limits the
571 amount of work the VM must do to enforce DIFC.

572 *Dynamic DIFC Enforcement with Barriers.* The VM inserts *barriers* that ensure no
573 reference outside a security method reads or writes labeled data, and all references
574 inside a security method follow the DIFC rules. A barrier is a snippet of code the VM
575 executes before every read and write to an object and is a standard implementation
576 feature in VMs for garbage-collected languages [Blackburn and Hosking 2004]. The
577 Laminar VM inserts barriers at every object read and write. Outside the security
578 method, the barrier throws an exception if the program tries to read or write data with
579 a nonempty secrecy or integrity label. Inside a security method, every time the program
580 reads or writes objects or kernel resources, the barrier checks that the information flow
581 follows the policies specified by the *current* labels of the security method.

582 For example, an assignment $w=r$ inside a security method M is safe if and only if the
583 information flow from r to w is legal for the thread inside M . Note that the Laminar
584 library API (Table I) does not include a routine for adding labels to a thread. In order
585 to add labels, threads must start a security method.

586 Security methods have the added benefit that they make the DIFC implementation
587 more efficient because the barrier checks outside a security method are simpler than
588 checks inside a security method. Outside a security method, the barrier simply checks
589 if data have nonempty labels and throws an exception if they do, since any access to
590 secure data is forbidden. Inside a security method, the DIFC barriers must compare
591 DIFC labels of references to ascertain if the information flow is legal.

592 In summary, security methods make it easier for programmers to add security policies
593 to existing programs. They make it easier for programmers to audit security code.
594 They limit the effects of implicit information flows (Section 6.4), and they make the
595 implementation more efficient.

596 *Inputs, Return Values, and Container Objects.* Security methods have two default
597 parameters: a secrecy label and an integrity label. A declassifier or endorser includes
598 a third parameter: the capability set. Security methods may take other parameters
599 as inputs and/or return an output value so long as the input or return is a valid
600 information flow. These input and output values may be primitives (`int`, `boolean`, etc.)
601 or object references. Generally speaking, security methods with a nonempty secrecy
602 label cannot return a value, and security methods with a nonempty integrity label
603 cannot read inputs without being wrapped in an endorser. Section 5.2 details the
604 specific rules.

605 A key abstraction in Laminar that improves programmability is the stylized use
606 of *container* objects. The programmer allocates objects with secret labels outside of a
607 secret security method by invoking `new` and passing the appropriate labels. Labeled
608 object creation is explained further in Section 5.2. The program then passes this object
609 to security methods. Each security method may update the contents of the object, but
610 outside of the security method, the object's contents and any modifications are opaque.
611 Code outside of a security method may not dereference references to objects with non-
612 empty labels.

613 *Example.* Figure 2 presents the calendar example from the introduction. A calendar
614 server calls code provided by Bob that reads the `Calendar` object belonging to Alice,

```

public class MeetingScheduler {
    Calendar AliceCal; // has labels ⟨S(a), I()⟩
    Calendar BobCal; // has labels ⟨S(b), I()⟩
    // Alice Calendar file "alice.cal" has labels ⟨S(a), I(i)⟩

    void ScheduleMeeting() {
        CapSet C = getCurrentCapabilities(); // C = (a+, b+, b-, i+)
        CapSet CapDeclassify = C.minus(i+);
        CapSet CapEndorse = C.minus(b+, b-);

        // Create a Meeting object with a secrecy label, to act as a security container
[L1] Meeting Mtg = new Meeting (S(a), I(), C(a+)) ();

        // Bob computes a mutually agreeable meeting time, declassifies to Alice
[L2] Mtg.BobFindMeetingTime(S(a,b), I(), CapDeclassify) (AliceCal, BobCal);

        // Alice's endorser
[L7] Mtg.AliceCheckMeeting (S(a), I(), CapEndorse) ();
    }
}

public class Meeting {
    MeetingTime val; // class variable

    secure (Label S, Label I, CapSet c)
    final void BobFindMeetingTime (Calendar Other, Calendar Bob) {
[L3] MeetingTime tmp = getMeetingTime(Other, Bob);

        // tmp now has labels ⟨S(a,b), I()⟩
[L4] Label newS = S.minus(S(b));
[L5] if ( /* bob's tests */ ) { BobDeclassify (newS, I, c) (tmp); }
    }

    secure (Label S, Label I, CapSet c)
    final void BobDeclassify (MeetingTime t) {
[L6] this.val = Laminar.copyAndLabel(t, newS, I);
    }

    secure (Label S, Label I, CapSet c)
    final void AliceCheckMeeting () {
[L8] Label AliceFileIntLabel = I.union(I(i));

        // Check that the MeetingTime is well-formed
[L9] if (AliceIntegrityCheck(this.val) {
[L10] this.val = Laminar.copyAndLabel(this.val, S, AliceFileIntLabel);
        }

        // this.val now has labels ⟨S(a), I(i)⟩
[L11] AliceWriteCalendar(S, AliceFileIntLabel) ();
    }

    secure (Label S, Label I)
    void AliceWriteCalendar () {
[L12] FileOutputStream AliceCalFile = new FileOutputStream("alice.cal");
        // Calculate offset into the file
[L13] AliceCalFile.Write(this.val, offset, this.val.length);
[L14] AliceCalFile.Close();
    }
}

```

Fig. 2. Example security methods that read and write a secure calendar. Bob provides the first two security methods. `BobFindMeetingTime` executes with both Alice and Bob's secrecy tags (a and b , respectively). This method selects a meeting time such that Alice and Bob are both available and places it in a `MeetingTime` object with label $\langle S(a, b), I() \rangle$. `BobDeclassify` then removes Bob's secrecy tag (b). The calendar application then executes Alice's endorser (`AliceCheckMeeting`), which checks that the `MeetingTime` object is well-formed, and then adds the i integrity tag and writes the meeting time to a file with label $\langle S(a), I(i) \rangle$. The label on Alice's calendar ensures both secrecy of her calendar data, as well as that all updates have been checked by trusted code. Execution order is indicated with LX, where X is the line number if the code were inlined into `ScheduleMeeting`.

615 creates a Meeting object, exports the meeting time to Bob, and writes the meeting
 616 into Alice’s calendar file. The code begins at L1 by allocating a container object, Mtg,
 617 which is passed to subsequent security methods to store the secret meeting time.
 618 When the thread enters the security method BobFindMeetingTime to read the calendars
 619 (line L2), the VM sets the thread’s secrecy label to $S(a, b)$ and therefore the thread
 620 can read secret Calendar data guarded by tags a and b . The code that reads each
 621 calendar (line L3) is a valid information flow because the thread’s labels are more
 622 restrictive than either Calendar’s labels (e.g., $\langle S(a, b), I() \rangle \supseteq \langle S(a), I() \rangle$). The thread
 623 has the capability $C(b^-)$ to declassify tag b , which is used to enter the security method
 624 BobDeclassify, at line L5. Entering the nested declassifier at line L5 may be conditioned
 625 on additional checks to prevent information leaks. Before writing the meeting time into
 626 Alice’s calendar file, the thread must acquire integrity label $I(i)$ by calling an endorser
 627 function, AliceCheckMeeting, which checks that the data to be written meet Alice’s
 628 invariants, such as prohibiting conflicting meeting times.

629 The VM inserts barriers that check the information flow safety at every object read
 630 and write. Locals are limited to method scope and implicitly have the same label as the
 631 security method. The code at Line L3 computes the common meeting time and stores
 632 it in the container object referred to by Mtg. For instance, the read barrier code tests
 633 if reading fields of objects Other and Bob are valid information flows and whether the
 634 writes into the newly created Meeting object are legal. The writes into the Meeting object
 635 are legal because the object has the same secrecy label as the thread in the security
 636 method at that point. A nested security method declassifies the meeting to Alice (L5/L6),
 637 updating the meeting time in this.val. By replacing the object referenced by this.val
 638 with a copy with a lower secrecy level (Line L6), this code effectively removes the tag b
 639 from the output. Copying and relabeling tmp at L6 is legal because the method has the
 640 b^- capability and can declassify data protected by the secrecy tag b . Notice that if line
 641 L6 performed copyAndLabel(tmp, S(), I()) to remove all tags from the secrecy label,
 642 the VM would throw an exception because when the thread is in the security method,
 643 it does not have the a^- capability and cannot remove the a tag from the data. In this
 644 example, the kernel checks the file operations in line L12 and L13 that write to Alice’s
 645 calendar file, and the VM checks the other operations on application data structures.

646 *Security Method Initialization.* Laminar enforces the following rules when a thread
 647 enters a security method. Let S_R , I_R , and C_R be the secrecy label, integrity label,
 648 and capability sets of a security method, R . Similarly, let S_P , I_P , and C_P be the sets
 649 associated with a kernel thread P that enters and then leaves R . Laminar supports
 650 arbitrary *nesting* of security methods. P could, therefore, already be inside a security
 651 method when it enters R . When the thread P enters the security method R , the VM
 652 ensures that the following rules hold:

$$(S_R - S_P) \subseteq C_P^+ \text{ and } (S_P - S_R) \subseteq C_P^- \quad (1)$$

653
$$(I_R - I_P) \subseteq C_P^+ \text{ and } (I_P - I_R) \subseteq C_P^- \quad (2)$$

654
$$C_R \subseteq C_P \quad (3)$$

655 The first two rules state that, in order for principal P to enter method R , P must have
 656 the required capabilities to change its labels to R ’s labels. The third rule states that the
 657 principal P can only retain a subset of its current capabilities when it enters a security
 658 method. While the security method executes, the sets associated with P change to S_R ,
 659 I_R , and C_R .

660 These rules encapsulate the common-sense understanding that a parent principal,
 661 P , has control over the labels and capabilities it passes to a security method and

that the VM will prevent the principal from creating a security method with security properties that the principal itself lacks. The rules also state that security methods nest in the natural way based on the labels and capabilities of the thread entering the nested method.

3.4. VM and OS Interface

Our design trusts the Laminar VM and OS kernel for DIFC enforcement. The VM is trusted to enforce DIFC policies on application data structures and implements security methods without kernel involvement. The kernel is responsible for enforcing DIFC rules on OS kernel abstractions, such as files and pipes. If a security method does not perform a system call, the VM does all the enforcement and does not involve the kernel. For example, if code inside a security method with secrecy tags tries to write to a public object, the VM will throw an exception that will end the security method. As an optimization, the VM does not notify the kernel of changes to the thread's labels until the VM needs to issue a system call on behalf of the application. The kernel enforces DIFC rules on each system call according to the thread's labels and the labels of any other objects involved (e.g., writing data to a file or the network). For standard system calls, such as read, the labels of the thread and file handle are implicit system call arguments. The VM communicates security metadata to the kernel via the Laminar system calls (Table II). For instance, the VM changes the labels on the current application thread (embodied as a kernel thread) executing within a security method using the `add_task_tag` system call. The kernel ensures that the labels are legal given the thread's capabilities.

Acquiring Tags and Capabilities. Principals (threads) in Laminar acquire capabilities in three ways. They allocate a new tag, they inherit them through `fork()`, or they perform interprocess communication. A thread working on behalf of one user may call security methods provided by another user; for instance, Alice's thread may call Bob's declassifier with the capability to read Alice's calendar. Another thread running on Bob's behalf can only acquire Alice's capability if Alice shares it over an IPC channel. The system carefully mediates capability acquisition lest a principal incorrectly declassify or endorse data. Laminar assumes a one-to-one correspondence between application and kernel threads. Application threads use the Laminar language API, which in turn invokes the system calls for managing tags and capabilities.

A principal allocates a new tag in the kernel via the `alloc_tag` system call, which is used to implement the language API function `createAndAddCapability`. As a result of the system call, the kernel security module will create and return a new, unique tag. The principal that allocates a tag becomes the owner of the new tag. The owner can give the plus and minus capabilities for the new tag to any other principal with whom it can legally communicate. A thread explicitly selects which capabilities it will pass to a security method, and the trusted VM can temporarily remove the capability from the thread using the `drop_capabilities` system call.

Threads and security methods form a natural hierarchy of principals. When a kernel thread forks off a new thread, it can initialize the new thread with a subset of its capabilities. Similarly, when a thread enters a security method, the thread retains only the subset of its capabilities specified by the method. In general, when a new principal is created, its capabilities are a subset of its immediate parent, which the VM and kernel enforce.

The passing of all interthread and interprocess capabilities is mediated by the kernel, specifically with the `write_capability` kernel call. This system call checks that the labels of the sender and receiver allow communication.

712 *Removing Tags and Capabilities.* The Laminar VM is responsible for correctly setting
713 thread labels and capabilities inside security methods. When a thread enters a security
714 method, the VM first makes sure that the thread has sufficient capabilities to enter
715 the method. If the thread may enter the security method, the VM sets the labels and
716 capabilities of the thread to equal those specified by the security method. The VM sets
717 the thread's capabilities to the empty set when it enters a security method that is not
718 passed capabilities (i.e., not a declassifier or endorser). Similarly, when the thread exits
719 the security method, the VM restores the labels and capabilities the thread had just
720 before it entered the method. On exiting a nested security method, the VM restores the
721 labels and capabilities of the thread to those of the parent security method.

722 The Laminar language API provides a method, `removeCapability`, that removes
723 a thread's capability in the VM, preventing use as an argument to a later security
724 method. To prevent threads from using the capability set as a covert channel, capa-
725 bilities must be created and removed outside of a security method (Section 6.6). The
726 `removeCapability` VM call uses the `drop_capability` system call to drop the capability
727 from the OS kernel thread.

728 Similarly, if a security method issues a system call, the VM first invokes
729 the `add_task_tag` or `remove_task_tag` system calls to change the thread's labels in
730 the OS kernel. As an optimization, the VM postpones setting the thread's labels in the
731 kernel until just before the first system call and at the end of the security method. This
732 system call has no user API; it is used solely by the VM.

733 The Laminar VM prohibits security methods from changing their labels; labels stay
734 the same throughout the security method to prevent leaks through local variables
735 (Section 5.2). Labels are stored as opaque objects that cannot be enumerated. To change
736 labels in the middle of a security method, a thread begins a nested security method.

737 Consider an example when a thread only has the a^+ capability and starts a security
738 method with secrecy label $\langle S(a) \rangle$. The Laminar VM sets the secrecy label of the thread
739 to $\langle S(a) \rangle$ when the security method begins. When the security method ends, the VM
740 forces the thread to drop the secrecy label, even if it does not have the a^- capability.
741 To drop $\langle S(a) \rangle$ from a thread, the VM contains a high-integrity thread, running with a
742 special integrity tag called `tcb` that is trusted by the kernel. Using the `drop_label_tcb`
743 system call, this trusted thread may drop all current labels for a thread without having
744 the appropriate capabilities.

745 A single, high-integrity thread in the VM limits exposure to bugs because the kernel
746 enforces that only the thread with the `tcb` tag may drop labels within a single address
747 space. The VM cannot drop the labels on other applications. Only a small, auditable
748 portion of the VM is trusted to run with this special label.

749 *Capability Persistence and Revocation.* Capability persistence and revocation are
750 always issues for capability-based systems, and Laminar does not innovate any solu-
751 tions. However, its use of capabilities is simple and stylized. The OS kernel stores the
752 persistent capabilities for each user in a file. On login, the OS gives the login shell all
753 of the user's persistent capabilities, just as it gives the shell access to the controlling
754 terminal. If a user wishes to revoke access to a resource for which she has already
755 shared a capability, she must allocate a new capability and relabel the data. Because
756 tags are drawn from a 64-bit identifier space, tag exhaustion is not a concern.

757 3.5. Security Discussion

758 The Laminar OS mediates information flow on OS resources, such as files and pipes.
759 The Laminar JVM mediates information flows within the application using barriers,
760 by restricting the programming model, and constraining how data enter and leave a
761 security method. Implicit flows are mediated by masking the control flow within the

security method (Section 6.4). Updating the capability set of a thread is treated as a public write, preventing covert channels through this abstraction. We allow security methods to execute concurrently. Our threat model assumes that security methods will terminate and will not leak information through timing channels, including the execution time of a security method; Sections 6.4 and 6.5 discuss techniques that could be adopted in a production Laminar deployment to uphold these assumptions.

Although the security method design facilitates incremental adoption because threads manage capabilities, this choice unfortunately places a measure of trust in the code that calls security methods. To limit the risk of capability mismanagement, only security methods that endorse or declassify data are passed capabilities, and these methods must be declared `final`. Section 6.6 discusses this issue in more detail.

3.6. Labeling Data

The VM labels data objects at allocation time to avoid races between creation and labeling. The VM labels objects allocated within a security method with the secrecy and integrity labels of that method. The `create_file_labeled` and `mkdir_labeled` kernel calls create labeled files and directories. Other system resources use the labels of their creating thread.

Similar to most other DIFC systems, Laminar uses immutable labels. To change a label, the user must copy the data object. Section 5.4 discusses implementation details and the interaction of object labels with the Java memory model. Dynamic relabeling in a multithreaded environment requires additional synchronization to ensure that a label check on a data object and its subsequent use by principal *A* are atomic with respect to relabeling by principal *B*. Without atomicity, an information flow rule may be violated. For example, *A* checks the label, *B* changes the label to be more secret, *B* writes secret data, and then *A* uses the data. Atomic relabeling can prevent this unauthorized flow from *B* to *A*. Laminar currently supports only immutable labels on files. It may be possible to safely relabel files using additional synchronization.

3.7. Compatibility Challenges

Although Laminar is designed to be incrementally deployed, some implementation techniques are incompatible with any DIFC system. For instance, a library might memoize results without regard for labels. If a function memoizes its result in a security method with one label, a later call with a different label may attempt to return the memoized value. Because the memoized result is secret, Laminar will prevent the attempt to return it. Programmers or the VM must modify such code to work in any DIFC system.

3.8. Trusted Computing Base

To implement Laminar, we added approximately 2,000 lines of code to Jikes RVM [Alpern et al. 2000],² added a 1,000-line Linux security module, and modified 500 lines of the Linux kernel. This relatively small amount of code means that Laminar can be audited.

The Laminar design does not trust `javac` to enforce information flow rules but does trust `javac` to provide valid bytecode that faithfully represents the Java source. Jikes RVM does not include a bytecode verifier—a feature of a secure, production VM that should reject malformed bytecode.

We rely on the standardization of the VM and the OS as the basis of Laminar’s trust. In addition to trusting the base VM, Laminar requires that the VM correctly inserts the appropriate read and write barriers (Section 3.3) for all accesses and optimizes them

²<http://www.jikesrvm.org>.

809 correctly. Read and write barrier insertion is standard in many VMs [Blackburn and
 810 Hosking 2004]. In Linux, Laminar assumes that the kernel has the proper mechanisms
 811 to call into Linux Security Modules (LSM) [Wright et al. 2002]. Because many projects
 812 rely on LSMs, the Linux code base is under constant audit to make sure all necessary
 813 calls are made.

814 4. OS SUPPORT TO CONTROL INFORMATION FLOW

815 We have implemented support for DIFC in Linux version 2.6.22.6 as an LSM [Wright
 816 et al. 2002]. LSM provides hooks into the kernel that customize authorization rules.
 817 We added a set of system calls to manage labels and capabilities, as listed in Table II.
 818 Some LSM systems, such as SELinux [Loscocco and Smalley 2001], manage access
 819 control settings through a custom filesystem similar to `/proc`. A custom filesystem is
 820 isomorphic to adding new system calls. The Laminar security module contains about
 821 1,000 lines of new code, and we modified about 500 lines of existing kernel code to
 822 implement the Laminar OS.

823 *Tags, Labels, and Capabilities.* Tags are represented by 64-bit integers and allocated
 824 via the `alloc_tag()` system call. The OS stores labels and capabilities for system re-
 825 sources in the opaque security field of the appropriate Linux objects (e.g., `task_struct`,
 826 `inode`, and `file`). The OS persistently stores secrecy and integrity labels for files in the
 827 files' extended attributes. Most of the standard local filesystems for Linux support ex-
 828 tended attributes, including `ext2`, `ext3`, `xfs`, and `reiserfs`. A mature implementation
 829 of Laminar could adopt a similar strategy to Flume for filesystems without extended
 830 attributes, encoding a label identifier in the extra bits of the user and group identifier
 831 fields of a file's `inode` [Krohn et al. 2007].

832 *Files.* Using LSM, Laminar intercepts `inode` and `file` accesses, which perform all
 833 operations on unopened files and file handles, respectively. The `inode` and `file` data
 834 structures are used to implement a variety of abstractions, such as sockets and pipes.
 835 The Laminar security hooks perform a straightforward check of the rules listed in
 836 Section 2.2. The secrecy and integrity labels of an `inode` protect its contents and its
 837 metadata, except for the name and labels, which are protected by the labels of the
 838 parent directory.

839 For instance, if a process with secrecy label $\langle S(a) \rangle$ tries to read directory `foo` with the
 840 same secrecy label, the process will be able to see the names and labels of all files in
 841 `foo`. If file `foo/bar` has secrecy label $\langle S(a, b) \rangle$, any attempt to read the file's attributes,
 842 such as its size, will fail, as size of the file could otherwise be used to leak information
 843 about the file's contents.

844 In a typical filesystem tree, secrecy increases from the root to the leaves. Creating
 845 labeled files in a DIFC system is tricky because it involves writing a new entry in
 846 a parent directory, which can disclose secret information. For example, we prevent
 847 a principal with secrecy label $\langle S(a) \rangle$ from creating a file with secrecy label $\langle S(a) \rangle$ in
 848 an unlabeled directory because it can leak information through the file name. In-
 849 stead, the principal should pre-create the file before tainting itself with the secrecy
 850 label.

851 A principal may use the newly introduced `create_labeled` and `mkdir_labeled` sys-
 852 tem calls to create a file or directory with secrecy and integrity labels different from the
 853 principal's current labels. Informally, a principal may create a differently labeled file if
 854 its current labels permit reading and writing the parent directory, and it has capabili-
 855 ties such that it can change its labels to match the new file. More formally, we allow a
 856 principal with labels $\langle S_p, I_p \rangle$ to create a labeled file or directory with labels $\langle S_f, I_f \rangle$ if
 857 (1) $S_p \subseteq S_f$ and $I_f \subseteq I_p$, (2) the principal has capabilities to acquire labels $\langle S_f, I_f \rangle$, and

(3) the principal can read and write the parent directory with its current secrecy and integrity label. This approach prevents information leaks during file creation while maintaining a logical and useful interface. 858
859
860

Applying integrity labels to a filesystem tree is more complex than secrecy. The intuitive reason for integrity labels on directories is to prevent an attacker from tricking a program into opening the wrong file, for instance using symbolic links. The practical difficulty with integrity for directories is that a task with integrity label I_A cannot read any files or directories without this label, potentially including $/$. If system directories, such as $/home$, have the union of all integrity labels, then an administrator cannot add home directories for new users without being given the integrity labels of all existing users. Flume solves this problem by providing a flat namespace that elides this problem with hierarchical directory traversal and simplifies application-level data storage with integrity labels [Krohn et al. 2007]. 861
862
863
864
865
866
867
868
869
870

Applying integrity labels to a traditional Unix directory structure brings out a fundamental design tension in DIFC OSes between usability and minimizing trust in the administrator. Laminar finds a middle ground by labeling system directories (e.g., $/$, $/etc$, $/home$) with a system administrator integrity label when the system is installed. A user may choose to trust the system administrator's integrity label and read absolute paths to files, or she may eschew trust in the system administrator by exclusively opening relative paths. In the worst case, she creates her own `chroot` environment. Simple relative paths were sufficient for all of the case studies in this article. Laminar's approach supports incremental deployability by allowing users to choose whether to trust the system administrator at the cost of extra work for stronger integrity guarantees. 871
872
873
874
875
876
877
878
879
880

Pipes. Laminar mediates Interprocess Communication (IPC) over pipes by labeling the `inode` associated with the pipe message buffer. A process may read or write to a pipe so long as its labels are compatible with the labels of the pipe. Message delivery over a pipe in Laminar is unreliable. An error code due to an incompatible label or a full pipe buffer can leak information, so messages that cannot be delivered are silently dropped. Unreliable pipes are common in OS DIFC implementations [Krohn et al. 2007; Vandebogart et al. 2007]. Linux does not include LSM hooks in the pipe implementation; Laminar adds LSM hooks to the pipe implementation in order to mediate reads and writes to pipes. 881
882
883
884
885
886
887
888
889

To prevent illegal information flows in Laminar, a pipe does not deliver an end-of-file (EOF) notification when the writer exits or closes the pipe if the writing thread cannot write to the pipe at the time it exits. This lack of termination implies that, if a process exits inside of a security method, the JVM must ensure that the thread's label is visible to the kernel (Section 3.4) before issuing an `exit` system call, so that the appropriate policies are applied when the OS closes the open file descriptors. Thus, Laminar, like many OS DIFC implementations, only delivers EOF notifications if writing the notification constitutes a legal flow. 890
891
892
893
894
895
896
897

Thus, the practical implication of unreliable delivery and eliminating EOF notification is that reads from a pipe should be nonblocking. Otherwise, an application may hang waiting for an EOF notification. In the common case where all applications in a pipeline have the same labels, traditional Unix pipe behavior can be approximated with a timeout. Using pipes in programs with heterogeneous, dynamic labels may require modification for a DIFC environment. 898
899
900
901
902
903

Network Sockets and Other IPC. The Laminar OS prototype treats network sockets and other IPC channels as having empty secrecy and integrity labels. Thus, input from the network must be read by code with empty secrecy and integrity labels, and the data must be labeled in a security method that validates the input. Managing information 904
905
906
907

4:22

D. E. Porter et al.

908 flows across systems is beyond the scope of this work, but has been addressed in other
909 systems including DStar [Zeldovich et al. 2008]. The inodes associated with other
910 Linux IPC abstractions, such as System V IPC, could be labeled similarly to pipes but
911 would likely require additional analysis of any potential information flows resulting
912 from idiosyncratic behavior.

913 5. JAVA VM SUPPORT TO CONTROL INFORMATION FLOW

914 We implement Laminar’s trusted VM in Jikes RVM 3.0.0,¹ a well performing Java-in-
915 Java VM [Alpern et al. 2000]. Our Laminar implementation is publicly available on
916 the Jikes RVM Research Archive and on GitHub.³ All subsequent uses of the term VM
917 refer to the Laminar-enhanced implementation in Jikes RVM.

918 When a thread starts a security method, the VM inserts a check that determines
919 if the thread has the capabilities to initialize the security method with the specified
920 labels and capabilities, as described in Section 3.3. Thread capabilities are stored in
921 the kernel. The VM caches a copy of the current capabilities of each thread to make
922 the checks efficient inside the security method.

923 The VM enforces information flow control for accesses to three types of application
924 data: objects, which reside in the heap; locals, which reside on the stack and in registers;
925 and statics, which reside in a global table.⁴ This section describes how the VM enforces
926 the DIFC rules on objects, local variables, and static variables.

927 5.1. Controlling Information Flow on Objects

928 The VM interposes on every read and write to an object or static by transparently
929 adding *barriers* before the operation. Barriers are not visible to the programmer and
930 cannot be avoided, thus creating a natural point to mediate explicit data flows. The
931 VM uses barriers to ensure that all accesses to data with nonempty labels occur within
932 a security method and that references inside a security method conform to the DIFC
933 rules in Section 2.

934 *Heap Objects.* The VM tracks information flow for labeled heap objects. When an
935 object is allocated, the VM assigns immutable secrecy and integrity labels to the object.
936 We modify the allocator to take secrecy and integrity labels as parameters; the allocator
937 adds two words to each object’s header, which point to secrecy and integrity labels.
938 The VM assigns objects allocated inside security methods the labels of the method
939 at the allocation point. To change an object’s labels, our implementation provides an
940 API call, `copyAndLabel`, that clones an object with specified labels. The label change
941 must conform to the label change rule (Section 2). The VM allocates labeled objects
942 into a separate *labeled object space* in the heap, which we exploit to optimize the
943 instrumentation that checks whether an object is labeled or not.

944 Each object acts as a *security container* for its fields, and the object’s labels protect
945 the fields from illegal access. The Laminar prototype requires that all fields of an object
946 have the same labels. For example, consider an object pointed to by the reference `o`.
947 The object has two fields, primitive integer `x` and reference `y`. When the program reads
948 or writes `o.x` or `o.y`, the VM enforces DIFC rules based on the labels of the object
949 referenced by `o`. If the program has labels that allow it to read the object referenced by
950 `o`, then it may read or copy `o.x` and `o.y`. However, the object that `o.y` references may
951 have the same or different labels. Thus, the programmer may assign distinct labels

³<http://www.jikesrvm.org/Research+Archive> and <https://github.com/ut-osa/laminar>.

⁴Although objects and nonvolatile statics may be register-allocated and nonescaping objects can be scalar-replaced, objects appear to be in the heap and statics appear to be in the global table when the VM compiler adds the needed barriers.

to the object referenced by `o.y`. If the application performs the dereference `o.y.foo`, the VM must verify that the security method may read and dereference the reference `o.y` based on the labels of the object referenced by `o` and then separately check the read of reference `foo` based on the labels of the object referenced by `o.y`. The security container model simplifies the task of labeling objects at allocation time, which is easier for programmers to reason about and cheaper for the VM to enforce compared to labeling individual object fields.

Labels. Applications do not have direct access to labels on data or principals, which are used internally by the VM to enforce DIFC rules. Recall that a label may contain one or more tags.

The Laminar API provides two functions that return a label. The functions return the label in an immutable, opaque object of type `Label`. The instantiations of `Label` support operations such as `isSubsetOf()`, `minus()`, and `union()`. The function `createAndAddCapability` invokes the `alloc_tag` Laminar OS system call, which creates a new tag and adds the associated capabilities to the current thread, and returns a `Label` object containing the single new tag to the application. The `getCurrentLabel()` function returns the secrecy or integrity label of the enclosing security method.

For efficiency, `Label` objects may be safely shared by objects, security methods, and threads because they are immutable; operations such as `minus()` and `union()` return a new object instead of modifying an existing `Label`. `Label` objects are not used internally by the VM for DIFC enforcement. Internally, the VM implements `Label` as a sorted array of 64-bit integers to hold tags. Because a `Label` object is opaque, applications cannot observe the individual values of the tags. Moreover, because object labels are immutable, any attempt to change the labels on an object requires writing a reference to the new object somewhere, which is an explicit, regulated information flow. Thus, a program cannot create a covert channel by creating a `Label` with irrelevant tags.

Similar to any other object, the VM associates secrecy and integrity labels with the instances of `Label`. An application may create a `Label` object using the new keyword or by using trusted Laminar API functions. When a `Label` object is created, it has the secrecy level of the thread at the time it was created. The integrity level of the `Label` object depends on which function created it: `Label` objects created by `new` also have the integrity of the thread at the time of creation, whereas `Label` objects created by the Laminar API functions are given the highest integrity (\top , representing the set of all possible integrity tags) because we trust the API and the VM. In general, `Label` objects have high integrity and empty secrecy and can be used as parameters to any security method.

VM Instrumentation. To enforce DIFC rules, the VM's compiler inserts barrier instrumentation just prior to every read and write in the application (Section 3.3). Inside security methods, the compiler inserts barriers at a *labeled object allocation* (before the compiler invokes the application's constructor) that sets the labels. It inserts barriers at every *read from* and *write to* an object field or array element. Inside security methods, barriers load the accessed objects' secrecy and integrity `Labels` and check that they conform to the current security method's labels and capabilities. Outside security methods, read and write barriers check that the accessed objects are *unlabeled* (i.e., have empty secrecy and integrity labels).

The compiler inserts different barriers depending on whether the access occurs inside or outside a security method. If a method is called both from inside and outside security method contexts, the compiler will produce two versions of the method. In our prototype implementation, when a method first executes, the JVM invokes the compiler, and it checks whether the thread is executing a security method and inserts barriers accordingly. Subsequent recompilation at higher optimization levels reuses

1002 this decision. This approach, which we call *static barriers*, fails if a method is called
1003 from both within and outside a security method. Thus, we also support a configuration
1004 in which the compiler adds *dynamic barriers*. The barriers check whether the current
1005 thread is in a security method or not and then execute the correct barrier. A produc-
1006 tion implementation should use cloning to compile two versions of methods called from
1007 both contexts, and each call site can call the appropriate version based on context.
1008 (Some software transactional memory implementations use a similar approach [Ni
1009 et al. 2008].) Because which version to call is statically knowable at each compiled
1010 call site, the overhead one would attain with a method cloning implementation should
1011 match what we measure for static barriers.

1012 Because object labels are immutable, and security methods cannot change their
1013 labels, repeating barriers on the same object is redundant. We implemented an in-
1014 traprocedural, flow-sensitive data-flow analysis that identifies redundant barriers and
1015 removes them. A read (or write) barrier is redundant if the object has been read
1016 (written), or if the object was allocated, along every incoming path. Although this opti-
1017 mization is intraprocedural, the VM's dynamic optimizing compiler inlines small and
1018 hot methods by default, thus increasing the scope of redundancy elimination.

1019 *Example.* Figure 3 computes the sum of the grades obtained by two different stu-
1020 dents. The `student1` and `student2` objects are labeled and have different secrecy values
1021 associated with them. Once the security method starts, the VM assigns the thread the
1022 secrecy and integrity labels specified by `S` and `I`, respectively. Lines L2 and L3 read
1023 labeled objects and result in a security exception if the flow from `student1.grades` or
1024 `student2.grades` to the thread in the security method is not allowed. Line L4 stores the
1025 value in a new labeled Integer object and stores the reference in the labeled `avgHolder`
1026 object. At lines L5–L6, the thread calls a declassifying security method, passing it the
1027 capability to add and remove the secrecy tags by making an unlabeled copy of the
1028 `avgHolder.value` object. If the `CapSet` passed to the security method is not a subset of
1029 the current thread's capabilities, then the program throws a security exception at L5,
1030 which the end of the security method may catch; this is followed by returning from the
1031 security method. Security exceptions are a category of Java language exceptions and
1032 may be caught by the security method author. The VM does not propagate exceptions
1033 out of a security method (Section 6.2). Because the declassifier runs with an empty
1034 label, it may assign the new reference into the unlabeled `outputHolder.value` field. In
1035 practice, a declassifier such as `declassifyAverage` would be nested inside a security
1036 method with a nonempty secrecy label that first checked the potential output, as in
1037 Figure 2, and the application of rules in the VM would be similar.

1038 5.2. Restricting Information Flow for Locals and Parameters

1039 Laminar does not track labels on local variables because they cannot be used outside
1040 the scope of the current method, thus precluding an information flow to or from a
1041 security method. Laminar assumes that locals have the secrecy and integrity labels
1042 of the enclosing security method or empty labels outside of a security method. All
1043 security methods take as input two parameters: the secrecy label and integrity label.
1044 Declassifiers and endorsers may take a third parameter: the capability set. For clarity,
1045 these are indicated in examples with separate argument parentheses on the secure
1046 keyword.

1047 If the explicit flow is legal, security methods in Laminar can accept additional inputs
1048 and return outputs of primitive values (`int`, `boolean`, etc.) and references, which are
1049 passed-by-value. A security method with a nonempty integrity label may only accept
1050 input if the calling function is also in a security method with higher integrity or the
1051 capability to add all missing integrity tags (i.e., an endorser). A security method with


```

// threadCaps = C(s1+, s1-, s2+, s2-)
// ST is the thread's current secrecy label, initially empty
// IT is the thread's current integrity label, initially empty
// Label S_Empty = S(), I_Empty = I()

[L1] secure (Label S, Label I)
void computeAverage
(IntHolder avgHolder) {
    {VM operations}
    {? changeLabel(ST, S, threadCaps) }
    {? changeLabel(IT, I, threadCaps) }
    {save ST, IT, threadCaps}
    {ST = S}
    {IT = I}
    {threadCaps = C()}

[L2] int m1 = student1.grades;
    {? secrecy-label(student1) ⊆ (ST = (s1, s2))}
    {? (IT = ()) ⊆ integrity-label(student1)}

[L3] int m2 = student2.grades;
    {? secrecy-label(student2) ⊆ (ST = (s1, s2))}
    {? (IT = ()) ⊆ integrity-label(student2)}

    int avg = (m1 + m2) / 2;

[L4] avgHolder.value = new Integer(S, I) (avg);
    {? ST ⊆ secrecy-label(new Integer)}
    {? integrity-label(new Integer) ⊆ IT}
    {? ST ⊆ secrecy-label(avgHolder)}
    {? integrity-label(avgHolder) ⊆ IT}

    return;
    {restore ST, IT, threadCaps}
}

[L5] secure (Label S, Label I, CapSet C) {? C ⊆ threadCaps}
final void declassifyAverage
(IntHolder avgHolder,
 IntHolder outputHolder) {
    {? changeLabel(ST, S, C) }
    {? changeLabel(IT, I, C) }
    {save ST, IT, threadCaps}
    {ST = S}
    {IT = I}
    {threadCaps = C}

[L6] outputHolder.value =
    Laminar.copyAndLabel
    (AvgHolder.value,
     S_Empty, I_Empty);
    {? ST ⊆ secrecy-label(outputHolder)}
    {? integrity-label(outputHolder) ⊆ IT}
    {? changeLabel(ST,
                    secrecy-label(avgHolder.value),
                    C)}
    {? changeLabel(IT,
                    integrity-label(avgHolder.value),
                    C)}
    {? changeLabel(secrecy-label(avgHolder.value)
                    S_Empty, C)}
    {? changeLabel(integrity-label(avgHolder.value)
                    I_Empty, C)}

    return;
    {restore ST, IT, threadCaps}
}

// Label S == S(s1, s2), I == I();
// CapSet C_Declassify = C(s1+, s1-, s2+, s2-)
IntHolder avgHolder = new IntHolder (S, I) (); // labeled
IntHolder outputHolder = new IntHolder (); // unlabeled
[L1-L4] computeAverage(S, I) (avgHolder);
[L5-L6] declassifyAverage(S_Empty, I, C_Declassify) (avgHolder, outputHolder);
System.out.println("Average is " + outputHolder.value.toString());

```

Fig. 3. Securely computing the average grades of two students. The student1 and student2 objects are labeled. The object credentials contains the secrecy, integrity, and capabilities sets with which the security method is initialized. The statements on the right side are the checks that are performed by the VM. The symbol ? indicates an assertion, ⊆ indicates an information flow check, and the internal function change-Label(Label to, Label from, CapSet caps) checks whether a label change would be permitted given the input capabilities (caps).

an empty integrity label may read any input. Similarly, a security method may only return a value if it has an empty secrecy label or the value is returned to a more secret security method. Because declassifiers tend to be nested, most declassification examples write the output to an object or security container with an empty secrecy label that is passed as input to the method. 1052 1053 1054 1055 1056

4:26

D. E. Porter et al.

1057 To enable containers for secure data, Laminar permits creation of objects with an
1058 empty integrity label and nonempty secrecy label outside of a security method. By
1059 creating objects with a nonempty secrecy label outside of a security method, the cre-
1060 ation itself and the return of the reference cannot be dependent upon any secret in-
1061 formation and thus cannot create any information flow. Once the initial secret object
1062 reference is created, it can be passed to multiple security methods that operate on
1063 secret data, acting as a container for secret data. Note that the first security method
1064 the reference is passed to is the object's *constructor*. If a labeled constructor throws
1065 an exception, *new* must still return the labeled but uninitialized object. Because crit-
1066 ical regions with an empty secrecy label but a nonempty integrity label can return a
1067 value, allocation of objects with integrity tags can always be wrapped in an endorser
1068 without any secrecy tags. The endorser may allocate an object with both secrecy and
1069 integrity tags in its label so long as it drops its secrecy label before returning the
1070 object.

1071 Because a method with a nonempty secrecy label cannot return a value, the security
1072 container abstraction serves as a means to facilitate passing secret, intermediate values
1073 among security methods. The security container abstraction also neatly integrates with
1074 common Java patterns of using the implicit *this* input parameter. In other words,
1075 security methods may construct container objects and operate on them, as illustrated
1076 in Figure 2.

1077 5.3. Static Variables

1078 Static (global) variables in the Laminar prototype have empty secrecy and integrity
1079 labels. By inserting barriers at static variable accesses inside security methods, security
1080 methods with an empty secrecy label may write static variables, and security methods
1081 with an empty integrity label may read static variables.

1082 We expect that a production implementation could support nonempty labels on statics
1083 with modest overhead because static accesses are relatively infrequent compared to
1084 field and array element accesses. Good security programming practices, like general-
1085 purpose programming practices, recommend sparse use, if any, of statics. We did not
1086 find this functionality necessary, and none of the applications in Section 9 needed
1087 labeled static variables.

1088 5.4. Instantiating Labels

1089 Some care must be taken when creating objects to prevent race conditions between
1090 assigning the object label and concurrent attempts to dereference the object. In the
1091 Laminar implementation, the label fields of each object are hidden from the program-
1092 mer (VM-internal) and are assigned between object allocation and calling the object's
1093 constructor. From the perspective of the Java memory model [Manson et al. 2005; Pugh
1094 2005], the label fields should be treated similar to *final* fields. In the Java memory
1095 model, *final* fields are visible to all threads before the constructor returns. The VM
1096 must prevent reordering these assignments outside of the constructor, and the construc-
1097 tor writer must not make external assignments of the *this* object. In order to protect
1098 against a malicious constructor writer, a production Laminar VM would strengthen
1099 this requirement slightly: The label assignments must be visible to all threads before
1100 the constructor is called.

1101 6. SECURITY IMPLICATIONS AND INFORMATION FLOW ENFORCEMENT MECHANISMS

1102 This section summarizes the major classes of information flows that Laminar mediates
1103 and the security implications of the Laminar design. This section pays particular atten-
1104 tion to changes in the programming model introduced by Laminar, including security
1105 methods and thread capabilities. This section also discusses security issues that are

Table III. Laminar's Programming Requirements and the Attacks They Prevent

Requirement	Attack Prevented
Explicitly labeled objects in the JVM and OS.	Illegal explicit information flow through objects.
Restrict information flow through explicit function arguments and return values.	Illegal explicit information flow through arguments and return values. Prevents information flow through locals, which are out of scope in a security method.
Static fields have empty security and integrity labels.	Illegal explicit information flows through static fields.
A security method may only have one exit point, including exceptions. All exceptions will be caught at the end of a security method.	Implicit information flows based on security method control flow.
A security method will execute for a fixed amount of time (not implemented).	Limits the bandwidth of timing and termination channels, which would otherwise be increased by multithreaded synchronization.
Dropping or creating a capability is treated as a write to the thread's capability set and requires an empty secrecy label.	Prevents information flow through the thread's capability set.
A security method that takes a capability set as a third parameter must be declared <code>final</code> .	Prevents passing capabilities to unintended functions via inheritance.

not addressed in the Laminar prototype and how subsequent research could mitigate these concerns. This section connects implementation details described previously in Sections 4 and 5 with the system's security properties. Table III summarizes the key programming abstractions and requirements that Laminar places on the programmer and the attacks they prevent, all of which are discussed later in more detail.

In each example and figure in this section, we use the following notation for labels. The value of a secrecy label with tags a and b is represented as $S(a, b)$. In Java, this label is stored in a `Label` object. Similarly, an integrity label with tag i is represented $I(i)$. Finally, a capability set with the ability to add a and remove i is represented as $C(a^+, i^-)$.

6.1. Explicit Information Flows

An explicit information flow occurs when a program moves data from one variable to another or from program memory into an OS-managed data sink, such as a file. The Laminar JVM and OS kernel collaborate to track explicit information flows and prevent illegal information flows. This subsection reviews the strategy for each major programming abstraction and provides backward references for the implementation details.

OS abstractions (Section 4). Laminar extends the Linux 2.6.22.6 kernel with an LSM that adds secrecy and integrity labels to a task (OS-visible thread) and file inodes, which include most IPC abstractions, such as pipes. The Laminar LSM interposes on all file handle reads and writes to validate the flow, as well as other system calls such as creating files and directories. We extend the Linux kernel with a few additional system calls and security hooks.

Java objects (Section 5.1). Objects in the Laminar VM are explicitly labeled, and the VM checks the labels of an object before all reads from and writes to an object field. The Laminar VM extends Jikes RVM, which provides *barriers* that interpose each object read and write.

Local variables (Section 5.2). Laminar does not label or track information flows through local variables. Because labeled data are accessed in security methods, locals

```

// Object o has labels ⟨S(h), I()⟩
// and members
// L has labels ⟨S(), I()⟩
// Invariant: y == 2/x
static boolean L = false;
secure (Label S = S(h), Label I = I())
void explicit (Object o) {
    o.x++;
    L = o.H;
    o.y = 2 / o.x;
    ...
} catch (ArithmeticException e) {
    o.y = 2;
    o.x = 1;
} catch (...) {
    o.y = 2 / o.x;
}

```

Fig. 4. Catch blocks handle illegal flows. Programmer may handle security exceptions separately from other runtime errors (e.g., divide by zero). Label and capability values are inlined for clarity.

1135 in an untrusted parent are out of scope inside the security method and vice versa.
 1136 With additional static analysis on information flow through locals, one could safely
 1137 implement security methods as arbitrary, lexically scoped blocks within a method, as
 1138 originally proposed [Roy et al. 2009], but we found the implementation was much more
 1139 complex.

1140 *Arguments and return values (Section 5.2).* Laminar permits primitives and refer-
 1141 ences as input values to a security method as long as reading the input values would
 1142 not violate an integrity rule (e.g., no read down). Note that even object references are
 1143 passed by value in Java, so manipulating any input variables will not affect a local in
 1144 the calling frame. The VM will mediate all accesses to an object with barriers. Simi-
 1145 larly, the programming model is restricted such that a security method may only return
 1146 a value in the calling context if the write would not violate a secrecy rule (no write
 1147 down).

1148 In the case of nested security methods, a more secret calling method may pass
 1149 an input to a less secret inner security method if the outer method has appropriate
 1150 declassification capability. Similarly, a higher integrity method may accept input from a
 1151 lower integrity parent if the outer method has the appropriate endorsement capability.
 1152 These rules for nested security methods are necessary to facilitate declassification and
 1153 endorsement.

1154 *Static variables (Section 5.3).* The Laminar prototype treats all static fields as having
 1155 empty labels. The Laminar VM interposes on all static field accesses and prevents
 1156 illegal information flows to statics. In general, static fields are used infrequently, and
 1157 our application case studies did not require nonempty labels for static variables.

1158 6.2. Handling Illegal Flows

1159 When code in a security method attempts an illegal explicit information flow, the VM
 1160 creates an exception that transfers control to the end of the security method. As a
 1161 programmer convenience, the security method may catch exceptions in order to restore
 1162 program invariants. Any exceptions uncaught by the programmer will be caught by
 1163 the VM before the security method ends, thus hiding the control flow of the security
 1164 method from the caller.

1165 For example, the code in Figure 4 shows an illegal explicit flow. The code attempts to
 1166 copy and thus leak the value of secret variable `o.H`, which it may not declassify, to the

static, nonsecret variable `L`. Laminar raises an exception because the security method does not have the right to declassify `o.H`. The value of `L` does not change. The catch block gives the programmer a chance to restore program invariants before exiting the security method.

If a thread tries to enter a security method for which it lacks appropriate capabilities, or if the thread passes illegal inputs to the security method, the VM raises an exception and transfers control to the security method's terminating catch block. Essentially, entering a security method with invalid credentials will effectively skip execution of the security method without revealing any information to the calling thread.

If a system call is attempted that would generate an illegal information flow, the OS returns a unique error code to the VM. The VM treats this error as a security exception; that is, the same way as an illegal flow through application-level variables.

An attempt to access labeled data outside of a security method will terminate the application. To prevent covert channels by testing whether an object is labeled at all, assignments to references must be treated as explicit information flows, described in the next subsection.

6.3. Information Flow through Object References

When a labeled object is created in a security method, Laminar restricts how the object stores references in order to prevent information leaks. One option is that an object reference can be written to a `static` variable, which must have empty secrecy and integrity labels. Therefore, only a method with the capability to drop its labels (i.e., a declassifier) can store a labeled object reference in a `static`. Similarly, security methods with an empty secrecy label can return an object reference to the caller.

A security method may store a reference to one object inside of another. For instance, suppose a security method writes a reference to newly created object `x` into object `o`'s field `o.p`. This assignment is an explicit flow from the security method into object `o`, and the VM-inserted barriers check the information flow. If `o`'s labels are $\langle S(o), I() \rangle$ and the security method's labels are $\langle S(o, x), I() \rangle$, this assignment is an illegal flow that would violate the secrecy rule, and it triggers a security exception.

Programmers may find it helpful to pass an object reference as input to multiple security methods. This convention does not leak data because object references are passed by value in the Java calling convention. As discussed earlier, returning an initial reference to an object or storing the reference in a `static` requires the capability to declassify the secret. Subsequent reads of the reference will not leak secret data. A subsequent security method cannot update a `static` reference unless it can declassify all of its secret data. Similarly, overwriting an input parameter in a security method does not propagate information to the caller because object references are passed by value in Java.

This pattern for passing secret data among security methods can be generalized by creating *security container* objects—an object whose reference is public which stores a set of secret data or object references. Security methods with the same secrecy label as the security container may conveniently write to the object and accept its reference as input. This convention does not leak any information because the public reference is never changed, and the contents of the container are protected by VM barriers.

To facilitate this pattern, we permit new to operate as a security method that can return a newly constructed object. Because the object is actually allocated from the heap and labeled before the constructor is called, a labeled object can always be returned without leaking secret information. If the constructor fails or throws an exception, the exception is masked, just as with any other security method, and a partially initialized, but labeled, object is returned. Objects with integrity tags must be allocated inside of

4:30

D. E. Porter et al.

```

public static void main (..) {
    MyObj m, k;
    // Label S = S(a), I_b = I(b), I_Empty = I();
    // CapSet C = C(b+)

    // Create an object with labels ⟨S(a), I(b)⟩
[L1] m = new MyObj(S, I_b) ();
[L2-4] manipulateObj(S, I_b, C) (m);
[L5] updateSecret(S, I_empty) (m);

[L6] k = m; /* Legal copy of a local object reference */
[L7] k.val = 0; /* Runtime error, since k points to a labeled object */

}

secure (Label S = S(a), Label I = I(b), CapSet C = C(b+))
void manipulateObj(m) {
    // Endorse reference m
[L2] manipulateObjInternal(S, I) (m);
}

// Use m in a security method
secure (Label S = S(a), Label I = I(b))
void manipulateObjInternal(m) {
[L3] MyObj n = new MyObj();
[L4] m.val = n.val + 5;
}

secure (Label S = S(a), Label I = I())
void updateSecret(m) {
[L5] m.secret.x = computeNewSecret(); // Internal function, not shown
}

```

Fig. 5. Allocating and passing objects among security methods using local references. Runtime values of labels and capability sets are inlined for clarity.

1218 an endorser security method. Nested security methods can allow an endorser with the
 1219 capability to add a secrecy tag to create an object with both secrecy and integrity tags
 1220 in its label. This approach makes it easier for the programmer to create a security
 1221 container and pass it among security methods, without creating data leaks.

1222 *Local reference example.* Figure 5 shows an example in which a local object reference
 1223 m is passed among security methods. The constructor for the new `MyObj` creates a labeled
 1224 object at line L1. This object is assigned to local reference m and passed to the security
 1225 method `manipulateObj`, where it is modified (L4). Outside of the security method, the
 1226 reference m may be safely copied to another reference k . An attempt to dereference
 1227 either reference outside of a security method will result in a runtime exception, since
 1228 both point to a labeled object.

1229 *Integrity example.* Figure 5 also illustrates how Laminar guarantees integrity. In
 1230 line L1 we create an object and label it with integrity label b . This object is returned to
 1231 the calling thread and assigned to m . This reference m is passed to a security method
 1232 (`manipulateObj`) but because the local reference itself is untrusted, the reference must
 1233 be endorsed (L2) and then passed to the nested, high-integrity security method. The
 1234 reference k is also assigned outside the security method to a high-integrity object at
 1235 line L6. Since Laminar does not track labels of references, such an assignment outside
 1236 the security method is valid. However, Laminar would prevent low-integrity code from
 1237 modifying the high-integrity object. For example, the VM will raise an exception at line

```

// Object o has labels ⟨S(h), I()⟩
// and members
// L has labels ⟨S(), I()⟩
static boolean L = false;
secure (Label S = S(h), Label I = I())
void implicit (Object o) {
    if (o.H) L = true;
    ...
} catch (...) {
    ...
}

```

Fig. 6. An example implicit flow. Label and capability values are inlined for clarity.

L7 when the value of the object pointed to by *k* is dereferenced outside of a security method. 1238
1239

Secrecy example. Figure 5 illustrates how an object can also be used as a security 1240
container. As discussed earlier, reference *m* points to a secret object, which cannot be 1241
dereferenced outside of a security method. This object may store other secrets, such 1242
as object reference *x*, which can be read and modified inside high-secrecy security 1243
methods, illustrated in Line L5. 1244

6.4. Implicit Information Flows 1245

Security methods limit implicit flows by hiding the control flow within the security 1246
method and preventing exceptional control flow from leaving the method. An implicit 1247
information flow leaks secret data through control flow decisions [Denning and Denning 1248
1977]. To deal with implicit flows due to exceptional control flow, the VM requires 1249
every security method to have a catch block, as shown in Figure 4. The catch block 1250
executes with the same labels and capabilities as the security method. A security 1251
method may explicitly catch specific exception types (e.g., an arithmetic exception 1252
caused by a potential divide by zero in Figure 4) and use the ellipsis syntax to catch 1253
all other exceptions (equivalent to a catch block that catches any `Throwable`). The VM 1254
suppresses other types of exceptions inside a security method that are not explicitly 1255
caught inside the security method, including exceptions within the catch block. Thus, 1256
exceptions cannot escape a security method. The VM continues execution after the 1257
security method. 1258

A major benefit of security methods is that they limit the amount of analysis neces- 1259
sary to restrict implicit information flows. Figure 6 includes an attempt to create an 1260
implicit flow. This security method code tries to leak the value of secret variable *o.H*, 1261
which it may not declassify, by deliberately creating an exception when it attempts 1262
an illegal explicit flow to the variable *L*. A thread might attempt to register an excep- 1263
tion handler outside of the security method that would learn the value of *o.H* based on 1264
whether an exception occurred. This attack will not work because the VM will suppress 1265
any exceptions from leaving a security method. 1266

To prevent information leaks, recall that security methods also cannot return a value 1267
unless they have an empty secrecy label (Section 5.2). Thus, security exceptions inside 1268
a secret security method cannot be reflected in the return value. A security method 1269
that has an empty secrecy label but a nonempty integrity label may return a value. 1270
If such a nonsecret security method incurs a security exception, the return value will 1271
either be set by the catch block or will be the default value for the return type. 1272

Alternatively, a VM prototype could permit security methods to be simple blocks, as 1273
we proposed initially, called *security regions* [Roy et al. 2009]. Security regions must exit 1274
via fall-through control flow. Security regions cannot use `break`, `return`, or `continue` to 1275

4:32

D. E. Porter et al.

```

// Object o has labels (S(h), I()),
// contains boolean H.
secure (Label S = S(h),
        Label I = I())
void termination (Object o) {
    if (o.H) while (true) { }
} catch (...) { }

```

Fig. 7. Leaking data via a termination channel. Runtime values of labels and capability sets are inlined for clarity.

```

// Object o has labels (S(h), I()),
// contains boolean H.
secure (Label S = S(h),
        Label I = I())
void simpleTiming (Object o) {
    if (o.H) Thread.sleep (5000) { }
} catch (...) { }

void untrustedCode() {
    // Label S = S(h), I = I()
    long start = System.currentTimeMillis();
    simpleTiming (S, I) (o);
    long end = System.currentTimeMillis();

    if (end - start > 4000) {
        System.out.println("`o.H is true`");
    } else {
        System.out.println("`o.H is false`");
    }
}

```

Fig. 8. Leaking data via a single-threaded timing channel. Runtime values of labels and capability sets are inlined for clarity.

1276 exit, except in the trivial case where the control flow will continue at the statement
 1277 that immediately follows the security region.

1278 Laminar thus eliminates implicit flows by hiding the control flow of a security method
 1279 from code outside of the security method. In Figure 4, code outside the security method
 1280 cannot distinguish an execution where $o.H$ is true from one where it is false. In contrast,
 1281 DIFC systems that rely on static analysis prevent these flows by detecting them during
 1282 compilation [Myers and Liskov 1997]. To prevent implicit flows, dynamic DIFC systems
 1283 generally either restrict the programming model, which we have done, or adopt a
 1284 hybrid of static and dynamic analysis [Chandra and Franz 2007; Nair et al. 2008;
 1285 Venkatakrishnan et al. 2006].

1286 6.5. Timing and Termination Channels

1287 In addition to explicit and implicit flows, an adversary may try to leak information
 1288 covertly through timing and termination channels [Lampson 1973]. A *timing channel*
 1289 attempts to leak information based on how long a piece of code executes. A *termination*
 1290 *channel* is a special timing channel that leaks information by executing in an infinite
 1291 loop depending on a secret value. We do not eliminate all timing and termination
 1292 channels for multithreaded programs, but we discuss potential solutions that minimize
 1293 their bandwidth.

1294 *Termination Channels.* Figure 7 shows an example of a termination channel that
 1295 attempts to leak secret information based on whether the application terminates. If
 1296 control returns from this security method, then unprivileged code can learn that $o.H$
 1297 is false. Similarly, a colluding application might learn that $o.H$ is true if the application
 1298 appears to hang.

1299 No general-purpose DIFC system can ensure termination of a program (or, in Lam-
 1300 inar's case, a security method). The primary goal in dealing with termination chan-
 1301 nels is preventing a deterministic or high-bandwidth channel. OS-based systems can
 1302 suppress termination notification [Efsthathopoulos 2008; Krohn et al. 2007; Zeldovich
 1303 et al. 2006] and thereby eliminate termination channels. Even this approach arguably


```

// Object o has labels ⟨S(h), I()⟩
//   and contains boolean H
// L has labels ⟨S(), I()⟩
static boolean L = true;

/* Thread 1 */
secure (Label S = S(h),
        Label I = I())
void timing1 (o) {
    if (o.H==true)
        Thread.sleep(5000);
} catch (...) { }
L=false;

/* Thread 2 */
secure (Label S = S(h),
        Label I = I())
void timing2 () {
    Thread.sleep(1000);
} catch (...) { }
System.out.println(L);

```

Fig. 9. A timing channel attack that with high probability prints L with the same value as the secret H. Runtime values of labels and capability sets are inlined for clarity.

creates some disruption in the CPU scheduling that might create a channel that is noisy and thus difficult to exploit.

In a language-level system like Laminar, untrusted code placed after a security method can detect whether a security method has terminated. A number of solutions have been explored in the literature, surveyed by Kashyap et al. [2011]. One option is to use static analysis to identify the labels of all variables used to make control flow decisions and only permit the code to execute if it can declassify these labels [Chandra and Franz 2007; Liu et al. 2009] or to restrict the programming model to forbid using a secret value as a conditional variable [Volpano and Smith 1999]. Another option is to partition and schedule the code based on labels [Kashyap et al. 2011]. A final option is to bound the maximum execution time of sensitive code [Askarov et al. 2010; Tiwari et al. 2009a] and return control to the untrusted code even if the sensitive code has not completed.

For Laminar, the most attractive approach to termination channels is simply bounding the execution time of a security method. If the maximum execution time (perhaps specified by the programmer) is exceeded, a security exception would be generated. Control would be transferred to the catch block, permitting the secure code to clean up (again for a bounded period) and then return to the unlabeled thread. This approach would prevent security methods from leaking data based on termination by artificially forcing all security methods to terminate.

Timing Channels. Similarly, a timing channel can leak information based on the execution time of a security method (or other privileged code in a different DIFC system). Figure 8 shows a timing channel that artificially delays execution based on the value of secret variable `o.H`. This sort of channel can be created even with a single thread by recording the time before and after execution.

In practice, these timing and termination channels have been low bandwidth and are difficult to exploit—especially in single-threaded applications. However, multithreaded applications are more vulnerable to these exploits because more threads can synchronize the order in which they execute a security method, which is then visible outside the security methods. Figure 9 illustrates a timing channel where threads artificially delay the length of security method execution based on the value of secret variable `o.H`. Even though neither security method explicitly leaks anything or fails to terminate, the execution time orders the execution of updates to the static variable `L`, thus leaking the value of `H` with high probability, but somewhat slowly.

Figure 10 shows a more subtle attack that efficiently and deterministically leaks one bit of information per security method execution. In this example, `signal` is a variable

```

// Object o has labels ⟨S(h), I()⟩
// and has boolean field H
// signal has labels ⟨S(), I()⟩
static int signal = -1;

/*Thread 1*/
secure (Label S = S(h),
        Label I = I())
void timing3 (o) {
    while (o.H==false && signal!=0);
} catch (...) { }
if(signal==-1) signal=1;
System.out.println(signal);

/*Thread 2*/
secure (Label S = S(h),
        Label I = I())
void timing4 (o) {
    while (o.H==true && signal!=1);
} catch (...) { }
if(signal==-1) signal=0;
System.out.println(signal);

```

Fig. 10. Leaking information via synchronization and timing. Runtime values of labels and capability sets are inlined for clarity.

1340 with empty labels. If the secret, H , is 1 then Thread 2 gets into a while loop until
 1341 Thread 1 exits its security method and sets the value of `signal` to 1. Thus, at the end
 1342 of the security method, the value of `signal` is the same as that of secret H . A variant
 1343 of this attack is also possible with only one thread in a security method and a second
 1344 thread sleeping and then writing to `signal`.

1345 We observe that the attacker in Figure 10 increases the bandwidth of timing chan-
 1346 nels by leveraging a data race on a `signal` variable. It is likely that the bandwidth of
 1347 some timing attacks in a dynamic DIFC system like Laminar could be reduced if the
 1348 program were known to be Data Race Free (DRF). Relying on programmers to write
 1349 DRF programs is straightforward but will fail to prevent attacks if programmers make
 1350 mistakes. Guaranteeing DRF through language design and type checking would pro-
 1351 hibit data races but requires programmer effort [Boyapati et al. 2002]. Alternatively,
 1352 the memory model could be strengthened so that synchronization-free regions appear
 1353 to execute atomically [Ouyang et al. 2013]. We note that even DRF programs can still
 1354 include timing channels, such as the one in Figure 8. Previous work has demonstrated
 1355 how a language-based DIFC system can reduce or eliminate timing channels in multi-
 1356 threaded programs by requiring data race freedom and that all traces of accesses to
 1357 public or low-secrecy variables are not influenced by secret inputs [Huisman et al. 2006;
 1358 Zdancewic and Myers 2003]. In general, locks for variables that are accessed across
 1359 multiple labels must be acquired in code with either the lowest secrecy and highest
 1360 integrity. Since correct lock acquisition complicates the programming model, we leave
 1361 further investigation to future work.

1362 Recent work [Askarov et al. 2010; Askarov and Myers 2012; Zhang et al. 2011]
 1363 provides an alternative promising approach to mitigating timing channels in language-
 1364 based systems by (1) predicting the expected runtime of a security-sensitive method,
 1365 (2) ensuring that every instance runs at least this long by delaying the return, and
 1366 (3) if the prediction is exceeded, increasing the prediction for future instances. This
 1367 predictive mitigation strategy substantially limits the ability of an attacker to create
 1368 a timing-based implicit flow.

1369 A variant of timing-based mitigation could be adopted by Laminar, in which program-
 1370 mers specify the execution time of a security method, plus some epsilon for imprecision
 1371 in the runtime system. Fixing execution time would address both timing and termina-
 1372 tion channels, and we expect that this would be robust to synchronization-based timing
 1373 attacks. We leave a formal treatment of this approach in the presence of concurrency
 1374 to future work.

```

// Object o has secrecy label S(a)
// and has a field boolean H
void thread() {
    String path;
    CapSet current = getCurrentCapabilities();
    Label L = createAndAddCapability();

    // Label S = S(a), I = I()

    leakH (S, I) (o);
    myMkdir (L, I) (path);

    File theDir = new File(path);
    if (theDir.exists())
        report("H was false");
    else
        report("H was true");
}

secure (Label S = S(a), Label I = I())
void leakH (Object o, Label L) {
    if (o.H) removeCapability(PLUS, L); // Runtime error in Laminar
}

secure (Label S = S(l), Label I = I())
void myMkdir (String name) {
    // thread capabilities would include l+ if o.H == false,
    // affecting whether mkdir_labeled succeeds
    mkdir_labeled(name, S, I);
}

```

Fig. 11. An attempt to use thread capabilities as a storage channel. Runtime values of labels and capability sets are inlined for clarity. Based on the value of H, the thread tries to permanently drop a capability. Laminar prevents this leak by ensuring programs only make permanent capability changes outside of a security method.

6.6. Capability Management

1375

Laminar adds a set of capabilities to each thread that persist across security methods. A critical concern is to ensure that the capability set not be used to create a storage channel to leak information. To avoid this, we treat a thread’s capability set as a nonsecret, trusted variable, and any tag creation or deletion is an explicit, mediated information flow. Because we trust the JVM to manipulate the capability set correctly, the capability set’s integrity label is treated as \top inside of a security method, and we permit threads to read the capability set outside of a security method.

1376

1377

1378

1379

1380

1381

1382

Figure 11 illustrates how such an attack might be attempted otherwise. The attacker thread initially creates a disposable capability for tag l , in Label L . Inside one security method, the thread drops l based on the value of `secret o.H` and later tries to use the capability (implicitly) in another security method to create a labeled directory. The thread does possess the $C(l^-)$ to declassify any data protected by l , which should create a public output if it is successful. In this attack, the value of `o.H` determines whether the thread drops the $C(l^-)$ capability, which determines whether the thread can create a directory with an irrelevant tag in its label. The untrusted code can see whether the directory exists and learn the value of `o.H`. In this example, the thread is essentially using the thread’s capability set as a storage channel to leak a secret value.

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

4:36

D. E. Porter et al.

```

// The thread takes in the reference to these caps as inputCap;
// 0 is alice, 1 is bob.
server_thread(Capability inputCap[2])

    // The thread has capabilities  $C(a^{+,-}, b^{+,-})$ 
    AliceCap = inputCap[0];
    BobCap = inputCap[0]; // Bug: Really AliceCap

    /* Schedule the Meeting Mtg, with labels Label S =  $S(A,B)$ , Label I =  $I()$  */

    BobDeclassify( S, I, BobCap)(Mtg);
    AliceDeclassify( S, I, AliceCap)(Mtg);

secure (Label S =  $S(A,B)$ , Label I =  $I()$ , Capability C = BobCap)
final void BobDeclassify() {
    // Bobcap is really alice cap.
    // Copy Alice's calendar to Bob with labels  $\langle S(B), I() \rangle$ 
}

```

Fig. 12. A potential “confused deputy” when managing capabilities in an untrusted thread. Runtime values of labels and capability sets are inlined for clarity.

1393 To prevent such a leak in Laminar, threads may only drop a capability either (1) out-
1394 side of a security method or (2) in a security method with an empty secrecy label.
1395 Essentially, dropping a capability is a write to the thread’s capability set, which has
1396 an empty secrecy label. Thus, this operation must be treated as an explicit write and
1397 mediated appropriately.

1398 *Capabilities and Confused Deputies.* One problem with threads that dynamically
1399 assign capabilities to security methods is that a bug in the untrusted thread code can
1400 inadvertently give a security method an inappropriate capability. Figure 12 shows a
1401 “confused deputy” [Hardy 1988] variant of the calendar example. The server thread
1402 accidentally gives Bob’s declassifier Alice’s declassification capability. Perhaps realizing
1403 the mistake, Bob copies Alice’s entire calendar into his calendar—a legal information
1404 flow.

1405 Dynamic capability management was a design decision made in the interest of pro-
1406 grammability. Unfortunately, as it stands, this choice increases the auditing burden on
1407 the security method developer. Not only must Alice audit her own security methods,
1408 she must audit the capability management code of threads that hold her declassifi-
1409 cation capability. Capabilities are Alice’s primary credentials in Laminar, so it is not
1410 surprising that capability management code requires a security audit. In some cases, it
1411 might be possible to trade auditing capability management code for auditing all secu-
1412 rity methods that a thread may call. However, that set might be difficult to determine
1413 statically, and it might include dynamically loaded methods and methods written by
1414 other users.

1415 To mitigate some of the risks of accidentally passing capabilities to the wrong secu-
1416 rity method, especially in the presence of inheritance of standard methods, capabilities
1417 must be explicitly passed to endorsing and declassifying security methods. Moreover,
1418 security methods receiving capabilities must be marked as `final`, thus disabling in-
1419 heritance. For security methods that manipulate labeled data without label changes,
1420 no capabilities need be passed to the method. When a security method is not explic-
1421 itly passed capabilities, the thread’s capability set will be temporarily assigned to the
1422 empty set for the duration of the security method.

1423 When a security method calls a function, its capabilities are not passed to this
1424 function unless the function is a nested security method that is explicitly passed

capabilities as arguments. This restriction reduces the risk of unexpected information flows through a third-party library. 1425
1426

An alternative design could allow Alice to map her capability to a hash of specific security methods, either in addition to or instead of thread capabilities. Such a mapping of capabilities to a security method alleviates the need to audit any code outside of the security method. We leave development of such a mechanism for future work. 1427
1428
1429
1430

7. LIMITATIONS 1431

Although Laminar regulates explicit information flows and hides control flow within a security method to prevent implicit flows, it is prone to attacks that exploit covert channels. For example, in the case of dynamic class loading, a user can query the VM to determine if a class has been loaded and use this additional information to leak sensitive data. In multithreaded programs, attackers may collude and use timing channels to leak information. We propose to mitigate these timing channels by fixing the execution time of a security method (Section 6.5). Such channels could also be mitigated by restricting the behavior of the scheduler [Sabelfeld and Myers 2003]. Laminar assumes that code blocks enclosed inside security methods always terminate. Otherwise, as explained in Section 6.4, information can leak through termination channels. 1432
1433
1434
1435
1436
1437
1438
1439
1440

The Laminar prototype trusts Java Native Interface (JNI) code that is included as part of the JVM. It does not track information flow through JNI code and does not allow third-party JNI modules. A production JVM could track the information flow through correct JNI code because the JNI specification requires C and Java to use separate heaps. The required copying of input and output data could serve as a natural point at which to check labels. We note that, as an optimization, many JVM implementations do give the C code pointers into the Java heap. This optimization must be disabled. A deeper concern is protecting against untrusted, user-provided JNI code. Because C is not memory safe, a malicious JNI module could guess or otherwise discover the location of JVM-internal bookkeeping. Protecting the JVM from untrusted JNI code would require a sandboxing technique, such as running the JNI in a separate address space, and is beyond the scope of our work. 1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452

In general, the implementation could handle Java reflection calls by intercepting them and handling them like normal calls for the purposes of Laminar’s security checks. The implementation could similarly handle calls to `sun.misc.Unsafe` methods, which perform raw memory accesses, by instrumenting the methods to perform Laminar’s checks. However, the prototype currently ignores these cases. 1453
1454
1455
1456
1457
1458

There is, however, a specific concern with combining reflection, multithreading, and file descriptors to create a covert channel. For instance, one could conditionally create a secret file inside of a security method, which influences the assigned file descriptor to file or socket creation outside of a security method, leading to a covert channel. This risk is only introduced when threads with different labels share a file descriptor table. The current Laminar prototype blocks this attack by relying on the fact that file descriptor values are hidden from the application in the `FileInputStream` and `FileOutputStream` classes without reflection or `sun.misc.Unsafe`. Thus, care would need to be taken in allowing an application to directly interact with the file descriptor table. 1459
1460
1461
1462
1463
1464
1465
1466
1467

The current implementation of Laminar does not allow application developers to read object labels, which may be useful for debugging. It is possible that some degree of visibility into object labels could be given to developers without creating new covert channels, but we leave this issue to future work. 1468
1469
1470
1471

The current implementation of Laminar treats static variables as unlabeled instead of associating labels with them (Section 5.3). Since most programs use statics infrequently, an improved implementation could track their labels without affecting the performance results. 1472
1473
1474
1475




Fig. 13. Laminar VM overhead on programs without security methods.

1476 As described in Section 3.8, the Laminar design does not trust javac to implement
1477 any DIFC enforcement but does trust javac to correctly compile the Java source to
1478 bytecode. Our prototype does not include a bytecode verifier, which could detect and
1479 reject invalid bytecode. Thus, the prototype Laminar VM trusts that bytecodes con-
1480 form to the specification. A production Laminar VM implementation would include a
1481 bytecode validator.

1482 The Laminar design requires security code to be written as methods, in order to
1483 simplify the enforcement of information flow rules on local variables (Section 5.2). With
1484 additional static analysis, it is possible that security code could be arbitrary, lexically
1485 scoped regions, as originally proposed [Roy et al. 2009]. After experimenting with a
1486 number of variations on the design, our experience is that restricting security code to
1487 methods strikes the best balance among programmability, security, and efficiency.

1488 Finally, several restrictions on the programming model are not currently checked in
1489 the Laminar VM runtime system but would be implemented in a production system.
1490 Rather, we require programmers to adhere to these restrictions and manually enforce
1491 them in our application studies. Specifically, the Laminar JVM prototype does not
1492 currently enforce input and output restrictions to security methods, enforce restrictions
1493 on labeled allocation, or restrict that security methods that accept capabilities are
1494 declared final. The VM could easily enforce all of these rules at runtime, and the JIT
1495 compiler could use static analysis to enforce some of them as well.

1496 8. LAMINAR PERFORMANCE

1497 This section reports the performance overheads incurred by adding Laminar to Jikes
1498 RVM and Linux. We conducted these experiments on a quad-core Intel Xeon 2.83GHz
1499 processor with 4GB of RAM. We configure Jikes RVM to run on four cores. The VM's
1500 heap is configured with a maximum size of 1,024MB. All results are normalized to
1501 values obtained on unmodified Linux 2.6.22.6 and Jikes RVM 3.0.0. We measured
1502 Laminar's overhead on standard Java benchmarks without security methods to be less
1503 than 10% using static barriers specific to code outside security methods. We measured
1504 Laminar OS overhead on lmbench, a standard OS benchmark, to be less than 8% on
1505 average.

1506 8.1. JVM Overhead

1507 Figure 13 shows the overhead of Jikes RVM with the Laminar enhancements on the
1508 DaCapo Java benchmarks [Blackburn et al. 2006], version 2006-10-MR2, and a fixed-
1509 workload version of SPECjbb2000 called pseudojbb [Standard Performance Evaluation
1510 Corporation 2001]. Because compilation decisions are nondeterministic, running times
1511 vary, so we execute 25 trials of each experiment and take the mean.

Table IV. Execution Time in Microseconds of Several lmbench OS Microbenchmarks on Linux with Laminar

Benchmark	Linux	Linux w/ Laminar	% Overhead
stat	0.92	0.94	2.0
fork	96.40	97.00	0.6
exec	300.00	302.00	0.6
0k file create	6.29	6.56	4.0
0k file delete	2.54	2.68	6.0
mmap latency	6,877.00	7,035.00	2.0
prot fault	0.24	0.26	7.0
null I/O	0.13	0.17	31.0

These bars represent two sets of runs, one with dynamic barriers and one with only static barriers. The darker bar shows the overhead of dynamic barriers, which check dynamically if they are in a security method as well as performing the secrecy and integrity checks as appropriate. Dynamic barriers add 23% overhead on average. The lighter bar is the overhead of using static barriers, which only do the appropriate per-object DIFC checks. This overhead is 9.7% on average. As discussed in Section 5.1, a mature implementation of Laminar would use method cloning to eliminate dynamic barriers. Because method cloning has comparable overheads to static barriers, code outside of a security method is expected to have an average overhead of 9.7%. This result is consistent with Blackburn and Hosking’s measurements of barriers [Blackburn and Hosking 2004].

8.2. OS Overhead

We use the standard lmbench [McVoy and Staelin 1996] system call microbenchmark suite to measure the overheads imposed on unlabeled applications when running on Laminar OS. A selection of the results is presented in Table IV.

In general, the overhead of the Laminar OS modifications are less than 8%, which is similar to previously reported overheads for Linux security modules [Wright et al. 2002]. The only performance outlier is the “null I/O” benchmark, which has an overhead of 31%. This benchmark represents the worst case for Laminar because the system call does very little work to amortize the cost of the label check. As a comparison, Flume adds a factor of 4-35× to the latency of system calls relative to unmodified Linux [Krohn et al. 2007].

9. APPLICATION CASE STUDIES

This section describes four case studies (GradeSheet, Battleship, Calendar, and FreeCS) and how we retrofitted these applications with DIFC security policies. GradeSheet implements a database with security policies for entering and reading grades by professors, TAs, and students. Battleship is a two-player game that keeps secrets about ship locations. Calendar manages multiple users calendars and arranges meeting times, similar to our running example. FreeCS is a chat server that implements security policies on group memberships and invitations. For each benchmark, we describe in more detail its functionality, modifying and retrofitting its security policies, and its performance.

Table V summarizes application details. All of the retrofitted applications implement more powerful security policies than their unmodified counterparts, yet all modifications add at most 10% to the source code. We list the lines of code statically within a security method, which is under 7% of the total lines of code. This count does not include code in methods called by a security method (e.g., library methods) that do not implement security policies.

Table V. Application Characteristics

Application	LOC	Protected Data	Added LOC (%)	SM LOC (%)	% time in SMs
GradeSheet	900	Student grades	92 (10%)	62 (6.9%)	<1%
Battleship	1,700	Ship locations	95 (6%)	57 (3.4%)	18%
Calendar	6,200	Schedules	290 (5%)	189 (3.0%)	1%
FreeCS	22,000	Membership properties	1,200 (6%)	80 (0.4%)	<1%

Lines of code (LOC), security sensitive data, Laminar specific LOC we added, LOC inside a Security Method (SM) statically (excluding code called by an SM), and percent time in security methods.




Fig. 14. Overhead of executing applications retrofitted with Laminar.

1550 Figure 14 breaks down into four parts the overheads added by securing them using
 1551 Laminar. *Start/end SM* is the overhead of application modifications to support DIFC,
 1552 including starting and ending security methods and security operations, such as copy-
 1553 AndLabel. The *Alloc barriers* configuration denotes the time overhead for allocating
 1554 labeled objects and assigning their label sets. The *Static barriers* configuration is the
 1555 overhead from read and write barriers when the security context is known at compile
 1556 time. Finally, the *Dynamic barriers* configuration is the overhead from barriers
 1557 that check context at runtime. We note that GradeSheet and Battleship run correctly
 1558 with static barriers, but Calendar and FreeCS require dynamic barriers because some
 1559 methods are called from both inside and outside security methods. As discussed in
 1560 Section 5.1, method cloning would obviate the need for dynamic barriers, and we thus
 1561 expect that in practice overhead will match the overhead of *Static barriers*.

1562 In all our experiments, we disabled the GUI, as well as other I/O and network-
 1563 related operations, so that the Laminar overheads are not masked by them. Hence,
 1564 the slowdown in deployed applications would be less than what is reported in our
 1565 experiments. In particular, when we wait for the Battleship game to draw the GUI
 1566 between scripted moves in the test cases, the measured Laminar overhead drops to 1%.
 1567 For comparison, Flume [Krohn et al. 2007] adds 34–43% slowdown on the MoinMoin
 1568 wiki application. Flume labels data at the granularity of an address space and cannot
 1569 enforce DIFC rules on heterogeneously labeled objects in the same address space.

1570 **9.1. GradeSheet**

1571 GradeSheet is a small program that manages the grades of students [Birgisson et al.
 1572 2008]. GradeSheet has three types of end users: professors, TAs, and students. The
 1573 main data structure is a two-dimensional object array `GradeCell`. The $(i, j)^{th}$ object
 1574 of `GradeCell` stores the information about student i and her grades on project j . A
 1575 sample policy states that (1) the professor can read/write any cell, (2) the TA can read
 1576 the grades of all students but only modify the ones related to the project that she
 1577 graded, and (3) students can only view their own grades for all projects.

1578 Table VI shows how to express this policy by assigning labels and capabilities to the
 1579 data and the threads working on behalf of each type of user, respectively. Specifically,

Table VI. GradeSheet Security Sets for Objects and Threads Serving End Users, where *S* Is Secrecy, *I* Is Integrity, and *C* Is Capability

Name	Security Set
GradeCell (i,j)	$S = \langle s_i \rangle, I = \langle p_j \rangle$
Student (i)	$C = \langle s_i^+, s_i^- \rangle$
TA (j)	$C = \langle \bigcup_{i=1}^{i=n} s_i^+, p_j^+, p_j^- \rangle$
Professor	$C = \langle \bigcup_{i=1, j=1}^{i=n, j=m} (s_i^+, s_i^-, p_j^+, p_j^-) \rangle$

we guard the $(i, j)^{th}$ entry in the GradeCell with the secrecy tag s_i and the integrity tag p_j . Each student i has the capabilities to add or remove s_i , so students can read their own grades in any project. Each TA j has the capability to add tags s_i and the integrity tag for the project that she graded (p_j). This tag ensures that TAs can read the grades of all students, but the integrity constraint prevents them from modifying grades for projects that they did not grade.

Interestingly, Laminar found an information leak in the original policy. The policy allowed a student to calculate and read the average grades in a project, which leaks information about the grades of other students. Using Laminar, we specified that only the professor is allowed to calculate the average and declassify it.

Our experiments measure the time taken by the server to process a mix of queries by the TA. Overall, the queries are 72% writes and 28% reads, including reads of student ID and average grade, and reads and writes of student grades. The Laminar-enabled version has a 7% slowdown compared to the unmodified version.

9.2. Battleship

Battleship is a common board game played between two players. Each player secretly places her ships on the grid in her board. Play proceeds in rounds. In each round, a player shoots a location on the opponent's grid. The player who first sinks all the opponent's ships wins the game.

We modified JavaBattle,⁵ a 1,700-line Battleship program available on SourceForge. Each player P_i allocates a tag p_i and labels her board and the ships with it. The capability p_i^- is not given to anyone else, ensuring that only the player can declassify the locations of her ships. In the original implementation, players directly inspect the coordinates of a shot to determine whether it hit or missed an opponent's boat. Under Laminar, each player sends her guess to her opponent, who then updates his board inside a security method. The opponent then declassifies whether the guess was a hit or a miss and sends that information back to the first player. We added fewer than 100 lines of code to secure the program to run with Laminar.

In our experiments, the game is played between computers on a 15×15 grid without a GUI. Figure 14 shows that the secured version adds 56% overhead with static barriers. The overhead is high because the benchmark spends a substantial portion of its time of its time (18%) inside security methods. In a deployed Battleship, which would display the intermediate state of the board to the players, the overhead is significantly less. In an experiment where we displayed the shot location after each move, the runtime of the application increases significantly, and Laminar overhead drops to 1%.

9.3. Calendar

We modified k5nCa1,⁶ a multithreaded desktop calendar that provides a graphical interface and allows users to subscribe to multiple external iCalendar-based calendars. It

⁵<http://sourceforge.net/projects/javabattle/>.

⁶<http://k5ndesktopcal.sourceforge.net>.

1618 has different threads for rendering the GUI, importing calendar files, and periodically
1619 fetching updates from remote calendars. Our modifications provide similar functional-
1620 ity as in the examples from earlier in this article. We label all data structures and `.ics`
1621 files that store a user’s calendar information with the user’s secrecy tag. We wrap all
1622 functions that access private calendar data inside security methods, including a sched-
1623 uler that finds available meeting times for multiple users. In the original program, a
1624 user could view the calendar of other users, a feature we disabled.

1625 Our experiments measure the time to schedule a meeting, which includes reading the
1626 labeled calendars of Bob and Alice, finding a common meeting date, and then writing
1627 the date to another labeled file that Alice can read. The scheduling code executes in
1628 a thread that has the capabilities to read data for both Alice and Bob, but can only
1629 declassify Bob’s data. The output file is protected by Alice’s secrecy tag. Our experiment
1630 schedules 1,000 meetings. Figure 14 shows that the secured version of Calendar runs
1631 6% slower than unmodified Calendar.

1632 We note that a substantial portion of the time in the calendar application is spent
1633 on internal thread creation and management, and even more time would be spent
1634 rendering a GUI if we had not disabled this feature. For comparison, we lifted the
1635 scheduling code out of the rest of the application and wrote a microbenchmark that
1636 scheduled appointments in a tight loop on a single thread. In this case, the percentage
1637 of time in security methods increased to 71%, and the total overhead was 77%. In
1638 practice, we expect things like user interaction and thread management to dominate
1639 execution time, thus minimizing the impact of security methods.

1640 9.4. FreeCS Chat Server

1641 FreeCS⁷ is an open-source chat server written in Java. Multiple users connect to the
1642 server and communicate with each other. FreeCS supports 47 commands, such as creat-
1643 ing groups, inviting other users, and changing the theme of the chat room. The original
1644 security policy consists of an authorization framework that restricts what commands
1645 can be used by a user. All these policies are written in the form of `if...then` checks.
1646 These authorization checks are actually checks on the *role* of a user. For example, a
1647 user who is in the role of a VIP and has superuser power on a group can ban another
1648 user in the group.

1649 We improve the security code in FreeCS by labeling sensitive data structures and
1650 accessing them inside security methods. We made most of our modifications in two
1651 classes—Group and User. We localized all security checks by adding security methods
1652 to these classes. The abstraction of a role maps naturally onto integrity labels. For
1653 example, we protected the `banList` data structure with two tags, one that corresponds
1654 to the notion of VIP and the other for the group’s superuser. Now, only users who
1655 have the add capability for these two tags can use the ban command. We modified
1656 the authentication module to assign each user either the VIP capability, superuser
1657 capability, or no enhanced capability when she logs in. The authentication module is
1658 trusted to manage the VIP and superuser capabilities. Our experiments measure the
1659 time to process requests from 4,000 users, each invoking three different commands.
1660 Laminar’s overhead is 5% (Figure 14).

1661 9.5. Summary

1662 The four case studies reveal a pattern in the way applications are written. First, most
1663 applications have only a few key data structures that need to be secured, like the ar-
1664 ray of student grades in GradeSheet or the playing boards in Battleship. Second, the
1665 interface to access these data structures is quite narrow. For example, `InternalServer`

⁷<http://freecs.sourceforge.net>.

in `GradeSheet` and `DataFile` in `Calendar` contain the functions used to access the important data. These observations support our hypothesis that Laminar requires only localized and modest changes to add DIFC security to many types of applications. Third, most of the data structures require heterogeneous labeling—the single array data structure `GradeCell` has different labels corresponding to different students. Heterogeneous labeling is impractical in OS-based systems [Krohn et al. 2007; Vandebogart et al. 2007; Zeldovich et al. 2006] because they support a single label on the whole address space or require the programmer to map application data structures onto labeled pages. The Laminar VM easily solves this problem with fine-grained tracking of labels on the data structure, for example, individual array elements and objects in `GradeSheet`.

An open question is the how this approach will scale to larger applications. Our experience with Laminar is that the performance overheads are primarily determined by the amount of code that executes inside a security method and that developer effort is a function of how many declassification or endorsement points the code requires, rather than the amount of data the program secures. The case studies presented here had natural and simple endorsement and declassification points, which were close to the actual uses of labeled data—minimizing overheads and effort. For larger applications, this trend may continue. However, it is possible that larger applications may instead require a larger portion of code in security methods to manipulate labeled data or that more substantial refactoring may be required to minimize the code that must execute in a security method. We leave larger application case studies for future work.

10. RELATED WORK

Previous DIFC systems have either used only PL abstractions or OS abstractions. Laminar instead enforces DIFC rules for Java programs using an extended VM and OS. By unifying PL and OS abstractions for the first time with a seamless labeling model, Laminar combines the strengths of previous approaches and further improves the DIFC programming model. Table VII summarizes the taxonomy of design issues common to DIFC systems, ranging from the trusted code base, security guarantees, resource granularity, to threats, all addressed in more detail here.

From IFC to DIFC. Information Flow Control (IFC) stemmed from research in multi-level security for defense projects [Department of Defense 1985]. In the original military IFC systems, an administrator must allocate all labels and approve all declassification requests [Karger et al. 1991]. Modern Mandatory Access Control (MAC) systems, like security-enhanced Linux (SELinux), also limit declassification and require a static collection of labels and principals. DIFC systems provide a richer model for implementing security policies in which applications allocate labels and assign them to data and declassify [Myers and Liskov 1997].

Static DIFC analysis. Many language-based DIFC systems augment the type system to include secrecy and integrity constraints enforced by the bytecode generator [Myers 1999; Myers et al. 2001; Simonet and Rocquencourt 2003]. These systems label program data structures and objects at a fine granularity but require programming an intrusive type system or in an entirely new language. These language-based systems trust the whole OS and provide no guarantees against security violations on system resources, such as files and sockets.

Hybrid DIFC enforcement. A key strength of static analysis is that static analysis tends to be the most robust language-level defense against implicit channels (Section 6.4). Purely dynamic systems generally cannot effectively regulate implicit flows. As a result, a number of primarily dynamic JVM systems have augmented

Table VII. High-level Approaches to DIFC Implementation

Issue	PL solutions	OS solutions	PL & OS solution
	[Arden et al. 2012; Chandra and Franz 2007; Liu et al. 2009; Myers et al. 2001; Simonet and Rocquencourt 2003]	Asbestos [Efstathopoulos 2008; Vandebogart et al. 2007], HiStar [Zeldovich et al. 2006], Flume [Krohn et al. 2007]	Laminar
Modified	Compiler & type system ([Arden et al. 2012; Liu et al. 2009; Myers et al. 2001; Simonet and Rocquencourt 2003]) or JVM and bytecode compiler ([Chandra and Franz 2007])	(1) Complete OS or (2) User-level reference monitor & kernel module	VM and kernel module
Trusted	Compiler, VM, & OS	OS	VM & OS
Fine-grained information flow tracking?	Interprocedural static analysis or JVM instrumentation	Either not supported or inefficient because of page table mechanisms	Dynamic VM enforcement via inserted read/write barriers
Secure files & OS resources?	Can label file handles in the application and add dynamic checks to system calls, but limited visibility into OS to validate these assumptions.	(1) Modify entire OS or (2) User-level reference monitor & kernel	Kernel
Implicit flows?	Static analysis, combined with dynamic checks in some cases [Arden et al. 2012; Chandra and Franz 2007; Liu et al. 2009].	Not applicable—tracks information flow at thread or address space granularity	<i>Security method</i> design restricts visibility into control flow from outside the security method.
Termination, probabilistic, and timing channels?	Predictive Mitigation [Askarov et al. 2010; Askarov and Myers 2012; Zhang et al. 2011]	HiStar, Flume, & Asbestos suppress termination notification	Not handled

Laminar combines aspects of PL and OS solutions, and innovates in dynamic flow tracking.

1715 dynamic enforcement of explicit flows with static analysis for implicit flows (thus called
1716 a Hybrid DIFC system).

1717 Chandra and Franz develop a version of the JVM that enforces information flow control
1718 policies on unmodified Java programs [Chandra and Franz 2007]. Like Laminar,
1719 this JVM combines static analysis on Java bytecode with dynamic analysis. Security
1720 policies are expressed externally—such as restricting how sensitive data may exit the
1721 program. This system relies on whole-program, side-effect analysis to restrict implicit
1722 flows by labeling the program counter. Moreover, this system does not address threads
1723 and allows implicit flows through uncaught exceptions. Finally, the dynamic analysis
1724 in this system is relatively expensive, 23–159%, whereas Laminar’s reported applica-
1725 tion overheads are 5–56%. Laminar’s security methods instead strike a balance that
1726 minimizes programmer effort but substantially limits the scope and overhead of static
1727 and dynamic analysis.

1728 Trishul adopts a similar design as Chandra and Franz, but better handles implicit
1729 flows through caught exceptions via static analysis [Nair et al. 2008]. Trishul does not
1730 handle uncaught runtime exceptions, such as divide by zero. Trishul relies on a con-
1731 servative global program counter secrecy label when static analysis cannot prevent an
1732 implicit flow, such as when referencing certain object reference fields. This abstraction

is prone to “label creep,” and programmers must manually remove labels according to application security policies. The performance overhead of Trishul varies and tends to be highest on object manipulation and lowest for system calls. For a prime number benchmark, the overhead is 167% [Nair 2009]. A key contribution of Laminar is a highly optimized JVM design, as well as a judicious and programmable abstraction selection that keeps overheads low.

The Laminar VM prevents implicit flows instead by restricting how control can return from a security method—a property that can be checked dynamically. Arguably, some restrictions could be relaxed with additional static analysis. Although Laminar does not rely on static analysis for safety, it does employ some analysis during JIT compilation to optimize security checks (Section 5.1) and could be considered a hybrid DIFC system.

OS IFC. Asbestos [Vandebogart et al. 2007] and HiStar [Zeldovich et al. 2006] are new OSs that provide DIFC properties. Flume [Krohn et al. 2007] is a user-level reference monitor that provides DIFC guarantees without making extensive changes to the underlying OS. These OS DIFC systems provide little or no support for tracking information flow through application data structures with different labels. Flume tracks information flow at the granularity of an entire address space. HiStar enforces information flow at page granularity and supports a form of multithreading by forcing each thread to have a page mapping compatible with its label. Using page table protections to track information flow is expensive, both in execution time and space fragmentation, and complicates the programming model by tightly coupling memory management with DIFC enforcement. Laminar supports a richer, more natural programming model in which threads may have heterogeneous labels and access a variety of labeled data structures. For example, all of our application case studies use threads with different labels.

Laminar provides DIFC guarantees at the granularity of methods and data structures with modest changes to the VM. It also adds a security module to a standard operating system, as opposed to Asbestos and HiStar, which completely rewrite the OS. Most of Laminar’s OS DIFC enforcement occurs in a security module whose architecture is already present within Linux (LSMs) [Wright et al. 2002]). The Laminar OS does not need Flume’s *endpoint* abstraction to enforce security during operations on file descriptors (e.g., writes to a file or pipe) because the kernel-level reference monitor can check the information flow for each operation on a file descriptor.

Laminar adopts the label structure and the label/capability distinction derived from Jif and used by Flume. Capabilities in DIFC systems are distinct from *capability-based* operating systems, such as EROS [Shapiro et al. 1999]. These systems use pointers with access control information to combine system and language mechanisms for stronger security but use a centralized IFC model, rather than the richer DIFC model. Thus, capability systems cannot enforce DIFC rules, and programs must be completely rewritten to work with the capability programming model.

Integrating language and OS security. Hicks et al. observe that security-typed languages can ensure that OS security policies are not violated by trusted system applications, such as logrotate [Hicks et al. 2007]. Their framework, called SIESTA, extends Jif to enforce SELinux [Loscocco and Smalley 2001] MAC policies at the language level. The aims of Laminar and SIESTA are orthogonal. SIESTA provides developers with a mechanism to prove to the system that an application is trustworthy, whereas Laminar provides the developer a unified abstraction for specifying application security policies.

Implicit information flows. Implicit information flows can leak secret data based on program control flow, as when a conditional statement is based on the value of a secret

1783 variable. DIFC systems based on static analysis can identify when labeled variables
1784 can influence control flow and use this information to label the program counter of
1785 the function [Chandra and Franz 2007; Liu et al. 2009]. In other words, a function
1786 with a secret program counter label may not be called by a function with an unlabeled
1787 program counter.

1788 Venkatakrisnan et al. develop a framework that statically transforms program code
1789 in a simple procedural language into a form that can detect implicit flows at runtime
1790 [Venkatakrisnan et al. 2006]. Their model essentially adds explicit assignments to a
1791 program counter variable in the code at all conditional statements and procedure calls
1792 and catches illegal flows at runtime. This model is applied in the context of IFC and
1793 noninterference and has not been extended to DIFC or concurrency.

1794 Le Guernic proposes an automaton-based information flow model and type system
1795 that uses static analysis to identify potential implicit flows [Guernic 2007]. Unlike
1796 other systems, this model also identifies synchronization events in threaded systems
1797 and imposes additional restrictions at runtime around conditional statements. These
1798 restrictions include requiring that locks be acquired before any conditional is evaluated
1799 based on a secret variable and executing statements within certain conditionals atomi-
1800 cally. Unlike many IFC systems, this design avoids termination channels on failures by
1801 suppressing individual lines of code that might cause an implicit flow. Laminar adopts
1802 a similar approach to securing concurrency by limiting the possible interleavings of
1803 security methods.

1804 Shroff proposes a dynamic monitor and type system that can prevent implicit flows
1805 either with the help of a static type analysis, which can be overly conservative in some
1806 cases, or by learning the implicit flows in repeated executions [Shroff et al. 2007]. In the
1807 dynamic-only mode, the system records the explicit flows within all taken branches. In
1808 subsequent executions, if a different branch is taken, the recorded flows of previous ex-
1809 ecutions are used to identify potential implicit flows. In dynamic mode, this system can
1810 permit some number of leaks before converging on a tight approximation of secure rules.

1811 Fabric and Mobile Fabric add additional checks, both static and dynamic, to prevent
1812 additional implicit flows in distributed and federated systems, respectively [Arden et al.
1813 2012; Liu et al. 2009]. For example, loading a class from a remote server may indicate
1814 that a secret code took a certain execution path. The Fabric systems add additional
1815 labels and checks to prevent these flows.

1816 Laminar restricts implicit information flows by restricting how exceptional control
1817 flow returns from a security method. OS DIFC systems generally do not need to address
1818 implicit flows because DIFC is enforced at process granularity, which hides control flow
1819 within the process by design.

1820 *Asymmetric behavior for secrecy and integrity.* In general, DIFC systems treat secrecy
1821 and integrity as duals. As a result, the bottom of the label lattice, or least-restricted
1822 data, is public and trusted. In Laminar, most application code and data are untrusted
1823 and have an empty label. We believe this choice is appropriate for a threat model
1824 where an adversary may have contributed code to the application, and any given policy
1825 concern applies to a small subset of the code. As a result, however, the measures taken
1826 to ensure secrecy and integrity are different. For instance, removing capabilities and
1827 creating security container objects must execute outside of a security method to ensure
1828 that the operation is public and does not leak secret information. In contrast, security
1829 methods trusted with an integrity tag must generally sanitize public data and endorse
1830 this input. In the worst cases, malformed public data can make the system unavailable.

1831 *Termination, timing, and probabilistic channels.* Implicit flows can be combined
1832 with termination, storage, and other features to create more powerful channels.

Vachharajani et al. argue that implementing DIFC with dynamic checking is as correct as static checking by showing that the program termination channels of static and dynamic DIFC systems leak an arbitrary number of bits [Vachharajani et al. 2004]. They prove that a correct dynamic DIFC system will overapproximate information flow, rejecting some programs that do not contain actual information flow violations. Russo and Sabelfeld similarly prove that a purely dynamic DIFC system will reject programs that a static analysis would not under a flow-sensitive analysis (i.e., when variables can change labels over the course of the computation) [Russo and Sabelfeld 2010]. Russo and Sabelfeld argue that these deficiencies can be recovered in a hybrid model, where some measure of static and dynamic analysis are combined. Laminar is a hybrid DIFC system but relies on dynamic checks and restricting the programming model to mitigate covert channels, and thus its security methods explicitly overapproximate information flow.

Recent work by Zhang, Askarov, and Myers developed a predictive mitigation strategy for timing channels [Askarov et al. 2010; Askarov and Myers 2012; Zhang et al. 2011]. Predictive mitigation essentially ensures that all instances of a sensitive method run for the same length of time. If a method runs longer than expected, all future instances run for the new maximum length. This strategy has been developed in static analysis systems but could be extended to dynamic DIFC systems such as Laminar.

In general, DIFC systems attempt to eliminate covert channels, which may be used to leak information, but do not eliminate timing channels [Lampson 1973] or probabilistic channels [Sabelfeld and Myers 2003]. DIFC systems can eliminate some implicit flows, as discussed in Section 6.4.

Formalizing information flow properties. Prior work has formally defined safety properties for information flow systems, primarily in the context of a type system. The most restrictive is *noninterference* [Goguen and Meseguer 1982], in which the output of a low-security computation cannot be influenced by the values of high-security computation. This definition precludes declassification and endorsement. In the case of our calendar example, a calendar application that enforced noninterference could not output a mutually agreeable meeting time. *Observational determinism* is a generalization of noninterference to concurrent programs [Huisman et al. 2006; Zdancewic and Myers 2003]. Observational determinism generally requires the program to be DRF, as well as requiring equivalent traces of possible accesses to nonsecret data.

An alternative safety condition is *robustness* [Chong and Myers 2005, 2006; Myers et al. 2004; Zdancewic and Myers 2001]. Within the lattice of labels in the decentralized label model, a robust system enforces boundaries on the ability of a principal to influence data or read data. In other words, a robust system would not allow a principal to expand its ability to read data based on the parts of the system it can influence. This model incorporates declassification, endorsement, multiple and mutually distrusting principals, and principals that can contribute and execute code.

Noninterference, observational determinism, and robustness have been applied primarily to DIFC-type systems. We leave adapting a property such as observational determinism to a dynamic DIFC system for future work.

In summary, Laminar combines the strengths of PL and OS DIFC systems. Laminar handles implicit flows and enforces the same fine-grained information flow control policies as performed by prior PL DIFC systems but without static analysis of all program components. Laminar enforces the same DIFC security policies on system resources as the OS DIFC systems enforce. Laminar, however, makes it easier to deploy and use information flow control systems by introducing security methods that encapsulate security code and are intuitive to program.

1884 **11. CONCLUSION**

1885 Laminar is the first DIFC system to unify PL and OS mechanisms for information flow
 1886 control. It provides a natural programming model to retrofit powerful and auditable
 1887 security policies onto existing, complex, multithreaded programs.

1888 Although abstractions such as the security method minimize the refactoring burden
 1889 on the programmer who wishes to adopt DIFC, the implementation mechanisms, such
 1890 as dynamic policy enforcement and allowing a thread to execute methods with different
 1891 labels, introduce additional opportunities for covert channels. To prevent some covert
 1892 channels, the current Laminar implementation imposes a number of modest restric-
 1893 tions that we would like to relax in future work, such as limiting the use of static
 1894 variables and forbidding file relabeling. This future work should be driven by a formal
 1895 model of security methods that facilitates careful reasoning about security properties,
 1896 especially about covert channels that arise due to concurrency.

1897 **ACKNOWLEDGMENTS**

1898 We thank the anonymous reviewers, David Mazières, Eddie Kohler, Maxwell Krohn, and Nikolai Zeldovich
 1899 for their feedback on earlier drafts of this document. Both David and the TOPLAS reviewers identified
 1900 timing and termination attacks that led to improvements in the Laminar model.

1901 **REFERENCES**

- 1902 B. Alpern, C. R. Attanasio, J. J. Barton, M. G. Burke, P. Cheng, J.-D. Choi, A. Cocchi, S. J. Fink, D. Grove,
 1903 M. Hind, S. F. Hummel, D. Lieber, V. Litvinov, M. Mergen, T. Ngo, J. R. Russell, V. Sarkar, M. J. Serrano,
 1904 J. Shepherd, S. Smith, V. C. Sreedhar, H. Srinivasan, and J. Whaley. 2000. The jalapeño virtual machine.
 1905 *IBM Systems Journal* 39, 1 (2000), 211–238.
- 1906 O. Arden, M. D. George, J. Liu, K. Vikram, A. Askarov, and A. C. Myers. 2012. Sharing mobile code securely
 1907 with information flow control. In *Proceedings of the IEEE Symposium on Security and Privacy*. 191–205.
- 1908 A. Askarov, D. Zhang, and A. C. Myers. 2010. Predictive black-box mitigation of timing channels. In *Proceed-
 1909 ings of the 17th ACM Conference on Computer and Communications Security (CCS'10)*. 297–307.
- 1910 D. E. Bell and L. J. LaPadula. 1973. *Secure Computer Systems: Mathematical Foundations*. Technical Report
 1911 MTR-2547, Vol. 1. MITRE Corp., Bedford, MA.
- 1912 K. J. Biba. 1977. *Integrity Considerations for Secure Computer Systems*. Technical Report ESD-TR-76-372.
 1913 USAF Electronic Systems Division, Bedford, MA.
- 1914 A. Birgisson, M. Dhawan, Úlfar Erlingsson, V. Ganapathy, and L. Iftode. 2008. Enforcing authorization poli-
 1915 cies using transactional memory introspection. In *Proceedings of the 15th ACM Conference on Computer
 1916 and Communications Security (CCS'08)*.
- 1917 S. M. Blackburn, R. Garner, C. Hoffman, A. M. Khan, K. S. McKinley, R. Bentzur, A. Diwan, D. Feinberg,
 1918 D. Frampton, S. Z. Guyer, M. Hirzel, A. Hosking, M. Jump, H. Lee, J. E. B. Moss, A. Phansalkar,
 1919 D. Stefanović, T. VanDrunen, D. von Dincklage, and B. Wiedermann. 2006. The DaCapo benchmarks:
 1920 Java benchmarking development and analysis. In *Proceedings of the 21th Annual ACM SIGPLAN
 1921 Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPLA'06)*. 169–
 1922 190.
- 1923 S. M. Blackburn and A. L. Hosking. 2004. Barriers: Friend or foe? In *Proceedings of the ACM International
 1924 Symposium on Memory Management*. DOI: <http://dx.doi.org/10.1145/1029873.1029891>
- 1925 C. Boyapati, R. Lee, and M. Rinard. 2002. Ownership types for safe programming: Preventing data races
 1926 and deadlocks. In *Proceedings of the 2013 ACM SIGPLAN International Conference on Object Ori-
 1927 Q4 ented Programming Systems Languages & Applications (OOPSLA'02)*. 211–230. DOI: [http://dx.doi.org/
 1928 10.1145/582419.582440](http://dx.doi.org/10.1145/582419.582440)
- 1929 D. Chandra and M. Franz. 2007. Fine-grained information flow analysis and enforcement in a java virtual
 1930 machine. In *Proceedings of the 23rd Annual Computer Security Applications Conference (ACSAC'07)*.
 1931 463–475.
- 1932 S. Chong and A. C. Myers. 2005. Language-based information erasure. In *Proceedings of the 15th IEEE
 1933 Computer Security Foundations Workshop (CSFW'05)*. 241–254.
- 1934 S. Chong and A. C. Myers. 2006. Decentralized robustness. In *Proceedings of the 16th IEEE Computer
 1935 Security Foundations Workshop (CSFW'06)*. 242–256.

Practical Fine-Grained Information Flow Control Using Laminar

4:49

- D. E. Denning. 1976. A lattice model of secure information flow. *Communications of the ACM* 19, 5 (May 1976), 236–243. 1936
1937
- D. E. Denning and P. J. Denning. 1977. Certification of programs for secure information flow. *Communications of the ACM* 20, 7 (July 1977), 504–513. 1938
1939
- Department of Defense. 1985. *Department of Defense Trusted Computer System Evaluation Criteria* (DOD 5200.28-STD (The Orange Book) ed.). Department of Defense. 1940
1941
- P. Efstathopoulos. 2008. *Policy Management and Decentralized Debugging in the Asbestos Operating System*. Ph.D. Dissertation. University of California, Los Angeles. 1942
1943
- J. A. Goguen and J. Meseguer. 1982. Security policies and security models. In *Proceedings of the 1982 Symposium on Security and Privacy (SSP'82)*. 11–20. 1944
1945
- G. Le Guernic. 2007. Automaton-based confidentiality monitoring of concurrent programs. In *Proceedings of the Computer Security Foundations Symposium (CSF'07)*. 218–232. 1946
1947
- N. Hardy. 1988. The confused deputy: (Or why capabilities might have been invented). *SIGOPS Operation Systems Review* 22, 4 (Oct. 1988), 36–38. 1948
1949
- B. Hicks, S. Rueda, T. Jaeger, and P. McDaniel. 2007. From trusted to secure: Building and executing applications that enforce system security. In *Proceedings of the 2007 USENIX Annual Technical Conference (USENIX'07)*. 1950
1951
1952
- M. Huisman, P. Worah, and K. Sunesen. 2006. A temporal logic characterisation of observational determinism. In *Proceedings of the 19th IEEE Computer Security Foundations Workshop (CSFW'06)*. 1953
1954
- P. A. Karger, M. E. Zurko, D. W. Bonin, A. H. Mason, and C. E. Kahn. 1991. A retrospective on the VAX VMM security kernel. *IEEE Transactions in Software Engineering* 17, 11 (1991). 1955
1956
- V. Kashyap, B. Wiedermann, and B. Hardekopf. 2011. Timing- and termination-sensitive secure information flow: Exploring a new approach. In *Proceedings of the 2011 IEEE Symposium on Security and Privacy (SP'11)*. 413–428. 1957
1958
1959
- M. Krohn, A. Yip, M. Brodsky, N. Cliffer, M. F. Kaashoek, E. Kohler, and R. Morris. 2007. Information flow control for standard OS abstractions. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP'07)*. 1960
1961
1962
- B. W. Lampson. 1973. A note on the confinement problem. *Communications of the ACM* 16, 10 (1973), 613–615. 1963
1964
- H. M. Levy. 1984. *Capability-Based Computer Systems*. Digital Press, Bedford, MA. 1965
- J. Liu, M. D. George, K. Vikram, X. Qi, L. Wayne, and A. C. Myers. 2009. Fabric: A platform for secure distributed computation and storage. In *Proceedings of the 21st ACM Symposium on Operating Systems Principles (SOSP'09)*. 321–334. 1966
1967
1968
- P. Loscocco and S. Smalley. 2001. Integrating flexible support for security policies into the Linux operating system. In *Proceedings of the 2001 USENIX Annual Technical Conference (USENIX'01)*. 1969
1970
- J. Manson, W. Pugh, and S. V. Adve. 2005. The Java Memory Model. Retrieved from <http://dl.dropbox.com/u/1011627/journal.pdf>. 1971
1972
- L. McVoy and C. Staelin. 1996. LMbench: Portable tools for performance analysis. In *Proceedings of the 1996 USENIX Annual Technical Conference (USENIX'96)*. 1973
1974
- A. C. Myers. 1999. JFlow: Practical mostly-static information flow control. In *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'99)*. ACM, New York, NY, 228–241. 1975
1976
1977
- A. C. Myers and B. Liskov. 1997. A decentralized model for information flow control. In *Proceedings of the ACM Symposium on Operating Systems Principles (SOSP'97)*. 129–142. 1978
1979
- A. C. Myers, N. Nystrom, L. Zheng, and S. Zdancewic. 2001. Jif: Java Information Flow. Retrieved from <http://www.cs.cornell.edu/jif>. 1980
1981
- A. C. Myers, A. Sabelfeld, and S. Zdancewic. 2004. Enforcing robust declassification. In *Proceedings of the 17th IEEE Computer Security Foundations Workshop (CSFW'04)*. 172–186. 1982
1983
- S. K. Nair. 2009. *Remote Policy Enforcement Using Java Virtual Machine*. Ph.D. Dissertation. Vrije University, Amsterdam. 1984
1985
- S. K. Nair, P. N. D. Simpson, B. Crispo, and A. S. Tanenbaum. 2008. A virtual machine based information flow control system for policy enforcement. *Electronic Notes on Theoretical Computer Science* 197, 1 (Feb. 2008), 3–16. 1986
1987
1988
- Y. Ni, A. Welc, A.-R. Adl-Tabatabai, M. Bach, S. Berkowits, J. Cownie, R. Geva, S. Kozhukow, R. Narayanaswamy, J. Olivier, S. Preis, B. Saha, A. Tal, and X. Tian. 2008. Design and implementation of transactional constructs for C/C++. In *Proceedings of the ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'08)*. 195–212. 1989
1990
1991
1992
1993
DOI: <http://dx.doi.org/10.1145/1449955.1449780>

4:50

D. E. Porter et al.

- 1994 J. Ouyang, P. M. Chen, J. Flinn, and S. Narayanasamy. 2013. ...and region serializability for all. In *Proceedings of the 5th USENIX Workshop on Hot Topics in Parallelism (HotPar'13)*.
- 1995 W. Pugh. 2005. May 12th Description of Final Fields. Retrieved from <http://www.cs.umd.edu/~pugh/java/memoryModel/may-12.pdf>.
- 1997 I. Roy, D. E. Porter, M. D. Bond, K. S. McKinley, and E. Witchel. 2009. Laminar: Practical fine-grained decentralized information flow control. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'09)*.
- 1998 A. Russo and A. Sabelfeld. 2010. Dynamic vs. static flow-sensitive security analysis. In *Proceedings of the 23rd IEEE Computer Security Foundations Symposium (CSF'10)*. 186–199.
- 1999 A. Sabelfeld and A. C. Myers. 2003. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications* 21 (2003).
- 2000 J. S. Shapiro, J. M. Smith, and D. J. Farber. 1999. EROS: A fast capability system. In *Proceedings of the ACM Symposium on Operating Systems Principles (SOSP'99)*.
- 2001 P. Shroff, S. Smith, and M. Thober. 2007. Dynamic dependency monitoring to secure information flow. In *Proceedings of the Computer Security Foundations Symposium (CSF'07)*. 203–217.
- 2002 V. Simonet and I. Rocquencourt. 2003. Flow Caml in a nutshell. In *Proceedings of the 1st APPSEM-II Workshop*. 152–165.
- 2003 Standard Performance Evaluation Corporation. 2001. *SPECjbb2000 Documentation* (release 1.01 ed.). Standard Performance Evaluation Corporation.
- 2004 M. Tiwari, X. Li, H. M. G. Wassel, F. T. Chong, and T. Sherwood. 2009a. Execution leases: A hardware-supported mechanism for enforcing strong non-interference. In *Proceedings of the Annual IEEE/ACM International Symposium on Microarchitecture (MICRO'09)*. 493–504.
- 2005 M. Tiwari, H. M. G. Wassel, B. Mazloom, S. Mysore, F. T. Chong, and T. Sherwood. 2009b. Complete information flow tracking from the gates up. In *Proceedings of the 14th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'09)*. 109–120.
- 2006 N. Vachharajani, M. J. Bridges, J. Chang, R. Rangan, G. Ottoni, J. A. Blome, G. A. Reis, M. Vachharajani, and D. I. August. 2004. RIFLE: An architectural framework for user-centric information-flow security. In *Proceedings of the 37th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO'04)*.
- 2007 S. Vandebogart, P. Efstathopoulos, E. Kohler, M. Krohn, C. Frey, D. Ziegler, F. Kaashoek, R. Morris, and D. Mazières. 2007. Labels and event processes in the Asbestos operating system. *ACM Transactions on Computer Systems* 25, 4 (2007), 11.
- 2008 V. N. Venkatakrishnan, W. Xu, D. I. C. DuVarney, and R. Sekar. 2006. Provably correct runtime enforcement of non-interference properties. In *Proceedings of the 2006 International Conference on Information and Communications Security (ICICS'06)*. 332–351.
- 2009 D. Volpano and G. Smith. 1999. Probabilistic noninterference in a concurrent language. *Journal of Computer Security* 7, 2–3 (Nov. 1999), 231–253.
- 2010 C. Wright, C. Cowan, S. Smalley, J. Morris, and G. K. Hartman. 2002. Linux security modules: General security support for the Linux kernel. In *Proceedings of the USENIX Security Symposium*.
- 2011 S. Zdancewic and A. C. Myers. 2003. Observational determinism for concurrent program security. In *Proceedings of the IEEE Computer Security Foundations Workshop (CSFW'03)*. 29–43.
- 2012 Steve Zdancewic and Andrew C. Myers. 2001. Robust declassification. In *CSFW*. IEEE Computer Society, Washington, DC, 15–23.
- 2013 N. Zeldovich, S. Boyd-Wickizer, E. Kohler, and D. Mazières. 2006. Making information flow explicit in HiStar. In *Proceedings of the USENIX Symposium on Operating Systems Design and Implementation (OSDI'06)*.
- 2014 N. Zeldovich, S. Boyd-Wickizer, and D. Mazières. 2008. Securing distributed systems with information flow control. In *Proceedings of the 5th USENIX Symposium on Networked Systems Design & Implementation (NSDI'08)*.
- 2015 N. Zeldovich, H. Kannan, M. Dalton, and C. Kozyrakis. 2008. Hardware enforcement of application security policies using tagged memory. In *Proceedings of the 8th USENIX Symposium on Operating Systems Design and Implementation (OSDI'08)*.
- 2016 Danfeng Zhang, Aslan Askarov, and Andrew C. Myers. 2011. Predictive mitigation of timing channels in interactive systems. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS'11)*. 563–574.
- 2017 D. Zhang, A. Askarov, and A. C. Myers. 2012. Language-based control and mitigation of timing channels. In *Proceedings of the 33rd ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'12)*.

Received March 2011; revised February 2014; accepted June 2014

QUERIES

- Q1:** AU: Please include full author addresses and emails.
- Q2:** AU: Is a term or word missing “we introduce security methods, an intuitive..”?
- Q3:** AU: Should this phrase be changed to “the lowest secrecy or highest integrity” as it follows “either”?
- Q4:** AU: For Boyapati et al. (and all refs) please supply the publisher and the page range.