

Verifiable Functional Purity in Java

Matthew Finifter Adrian Mettler Naveen Sastry David Wagner
Department of Electrical Engineering and Computer Science
University of California, Berkeley, USA

{finifter, amettler}@cs.berkeley.edu, naveen@ksastry.com, daw@cs.berkeley.edu

ABSTRACT

Proving that particular methods within a code base are *functionally pure*—deterministic and side-effect free—would aid verification of security properties including function invertibility, reproducibility of computation, and safety of untrusted code execution. Until now it has not been possible to automatically prove a method is functionally pure within a high-level imperative language in wide use, such as Java. We discuss a technique to prove that methods are functionally pure by writing programs in a subset of Java called Joe-E; a static verifier ensures that programs fall within the subset. In Joe-E, pure methods can be trivially recognized from their method signature. To demonstrate the practicality of our approach, we refactor an AES library, an experimental voting machine implementation, and an HTML parser to use our techniques. We prove that their top-level methods are verifiably pure and show how this provides high-level security guarantees about these routines. Our approach to verifiable purity is an attractive way to permit functional-style reasoning about security properties while leveraging the familiarity, convenience, and legacy code of imperative languages.

Categories and Subject Descriptors

D.2.3 [Coding Tools and Techniques]; D.2.4 [Software/Program Verification]

General Terms

Security, Languages, Verification

Keywords

Pure functions, determinism, static analysis, object-capabilities

1. INTRODUCTION

Critical real-world programs often have high-level security and privacy requirements expressed in terms of reproducibility, invertibility, non-interference, or containment of untrusted code. We would like to verify these properties given the programs' source

Permission to make digital or hard copies of all or part 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 bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CCS'08, October 27–31, 2008, Alexandria, Virginia, USA.
Copyright 2008 ACM 978-1-59593-810-7/08/10 ...\$5.00.

code, but this task is difficult in the languages commonly used to write real-world programs. These imperative languages permit side effects and data dependencies that are difficult to reason about. Purely functional languages, in which methods obey the semantics of mathematical functions, make reasoning about effects and information flow easier, but have not gained the popularity and code base of more traditional imperative languages. We present a technique for implementing verifiably *functionally pure* methods in imperative languages. To be functionally pure, a method must satisfy two critical properties¹:

First, it must have no side effects. For a computational method to be free of side effects, its execution must not have any visible effect other than to generate a result. A method that modifies its arguments or global variables, or that causes an external effect like writing to disk or printing to the console, is not side-effect free.

The second property is functional determinism: the method's behavior must depend only on the arguments provided to the method. The method must return the same answer every time it is invoked on equivalent arguments, even across different executions of the program. A simple example would be a method to upper-case a string: every time it is given a string containing the word "foo", it will return a string containing "FOO". Many methods do not satisfy this criterion, including ones whose behavior depends on the time of day, the amount of free memory, or whether a specific flag was present on the command line.

Electronic voting machines are one important application with a number of security requirements amenable to enforcement using functional purity. These machines are single-purpose computers running custom software designed to allow the voter to select his or her preferred candidates and to record the selections. Given the importance of these machines to our democracy and concerns over their trustworthiness, it would be useful if we could prove aspects of their operation correct.

For example, we argue that voting machines should be designed to ensure that each voter's voting experience will be a deterministic function of the ballot definition and that voter's actions. For a particular set of voter actions, the system should always present the same screens and record the same selections, independent of previous voters' interactions with the voting machine. Leaking any information about previous sessions could violate earlier voters' privacy and could create a conduit for a malicious voter to interfere with subsequent voters. Also, voting sessions should have no side effects; their only legitimate effect should be to return the voted ballot. Functional purity can help verify these security properties.

As another example, voting machines must serialize and possibly encrypt the voter's selections when writing them to stable storage.

¹More formal definitions of these two properties are provided in Section 3.

This data will be read and tallied at a future date, likely on a different machine. In order for the voter’s choices to be counted as they were cast, we must be certain that the reconstituted votes will match the originals. We propose a fail-stop check on the encoding process: the machine writing the data should decode the serialized output and verify that it matches the original vote selection data structure. If the decode method is deterministic, this check ensures that this data structure will be correctly reconstructed later when the votes are counted. If the serialization and deserialization routines are also side-effect free, they can be removed from the trusted computing base, as the check verifies their correctness as needed.

In general, verifying that a computation will be deterministic and free of side effects is a difficult task that typically requires careful examination of a program’s entire source code. Verifying side-effect freeness requires verifying that the computation does not *modify* the state of any parameters or global state and does not affect the outside world in any observable way (e.g., writing to an I/O device). Verifying determinism requires ensuring that the method does not *read* any information that may differ between different calls. Checking the latter property first requires ensuring that anything that is read by the method isn’t changed elsewhere in the program. Also, we must ensure that any value read by the method doesn’t depend on environmental factors that could differ between executions of the program.

We can see that the concepts of determinism and side-effect freeness are related, in that they both restrict access to state created outside the method. We use a unified approach to achieving both goals, based on object capabilities [18]. Specifically, we introduce and define the concept of *deterministic object-capability languages*, in which the ability to cause side effects and to observe data that varies between executions is conveyed by explicit object references that are propagated only by explicit program statements.

A key advantage of our approach is that it supports modular reasoning about purity, side effects, and determinism. In particular, a programmer can tell whether a particular method is pure simply by looking at its type signature. In our system, if all parameter types are immutable, then the method can be guaranteed to be pure. This allows purity specifications to be part of the contract of a method and simplifies the task of reasoning about program behavior. The *body* of a pure method has no additional constraints, permitting wide flexibility in how it is implemented. In particular, pure methods can call impure methods, and vice versa. In short, pure and impure code can easily be mixed; the majority of a program can be imperative, with purity still being enforced where needed.

We briefly describe how the Joe-E subset of Java satisfies the requirements of a deterministic object-capability language, and how it can be used to write methods that can be easily recognized as verifiably pure. In order to evaluate our approach to verifiable purity, we ported three legacy libraries (an AES implementation, serialization logic from an experimental voting machine implementation, and an HTML parser) to the Joe-E subset, and refactored them so that their top-level methods could be verified as pure.

As Joe-E was not explicitly designed to ease migration of legacy code, we found that the task of modifying existing code to satisfy the Joe-E restrictions was at times difficult. Certain recurring patterns account for much of this difficulty; code that avoids these patterns is much easier to port. Refactoring methods so they could be verified pure was generally harder than just porting to the Joe-E subset, and sometimes required changes to data structures and interfaces. We therefore recommend our approach primarily for use with new code that is designed with this approach to purity in mind.

We view the contributions of this work as follows:

- We enumerate several applications where the ability to verify that particular blocks of code are pure makes it easy to verify interesting high-level application-specific properties.
- We describe a class of imperative programming languages in which it is easy to verify purity.
- We introduce Joe-E’s enforcement of determinism and we show how this enables verifiable purity in Java.
- We share our experience refactoring legacy codebases so that they can be verified as pure, thus attaining useful security guarantees.
- Based on this practical experience, we identify programming patterns that are well-suited to writing verifiably pure systems as well as anti-patterns that make this task difficult.

2. APPLICATIONS

We argue that functional purity has many applications in security and reliability. Purity is a helpful tool for building more modular programs that are easier to reason about, and this makes it easier to verify many kinds of security properties. Languages and programming idioms that make this property easy to achieve and verify may be of benefit to programmers, especially those aiming to write maintainable, auditable, and understandable code.

2.1 Reproducibility

Consider the following scenario, inspired by [1]: Mallory generates a PDF file containing a contract for Alice to electronically sign. Mallory constructs this PDF file so that its displayed content depends on the system date. When viewed in January, the contract says that Mallory will pay Alice \$100; in any other month, the contract says that Alice will pay Mallory \$1,000. Suppose Alice reads and electronically signs the contract on January 1, and returns the signed contract to Mallory. On February 1, Mallory presents the signed contract to a judge, and the judge orders Alice to pay Mallory \$1,000.

The problem is that the computation that renders the text is not deterministic. The behavior of the PDF viewer depends on other factors aside from its input, the bits of the document file. This attack could not succeed if the PDF viewer’s computation was a pure function of the input file. If we could verify the purity of the viewer, we would be assured that Mallory’s attack will fail.

This is an example of a TOCTTOU vulnerability. Whenever we compute a result that is checked, and then recompute it later when it is used, we must be careful to ensure that the computation is reproducible. Pure functions are useful for this, because determinism ensures reproducibility and makes explicit the inputs a computation may depend upon.

Another application is in transactional systems. Suppose we take periodic checkpoints of an application and log all its inputs. If the application is deterministic, then we can recover from crashes: reincarnating the application and replaying from an old snapshot and input trace will always reproduce the same behavior that the previous incarnation of the application followed. This eliminates the need to checkpoint every intermediate state. It also allows a replicated system to transparently fail over to a backup system that is receiving the same stream of input events.

2.2 Invertibility

The serialization example given in the introduction is representative of a class of applications that have a matched pair of algorithms (Encode, Decode) for which it is intended that Decode is an inverse of Encode. Specifically, the *inverse property* should hold: for all x , $\text{Decode}(\text{Encode}(x))$ should yield some output x' that is functionally equivalent to x . To ensure the original x will be re-

coverable in the future, this has to hold even if the invocation of Decode takes place at some later time on a different machine.

Purity helps support *fail-stop* enforcement of this property, in which errors are detected at runtime but before any harmful consequences have taken place. One can test `Encode(x)` at runtime to ensure that it will be decoded correctly by Decode:

```
y := Encode(x)
abort if x != Decode(y)
```

If Decode is purely functional, its determinism ensures that the check can be performed at any time and will accurately reflect whether the message can be correctly decoded in the future. Also, if Decode is side-effect free, adding this check to existing code won't break the program.

This approach applies to, e.g., serialization and deserialization, encryption and decryption, and compression and decompression. In many such applications, it is better to fail and warn the user than it is to proceed and lose data. If this pattern is used to ensure that all data that is encoded can be recovered, neither the encoder nor decoder need to be trusted correct in order to establish the property that data is never lost or corrupted.

Formally verifying the correctness of serialization and deserialization with static analysis is a difficult task. Serialization and deserialization typically involve walking a (potentially cyclic) object graph, and thus inevitably implicate complex aliasing issues, which is known to make static analysis difficult. Therefore, purity seems better-suited to this task than classical approaches.

Deterministic functions can also be used for enforcement of more complex functional relations than invertibility. The exokernel Xok's stable storage system uses what the authors call a UDF (untrusted deterministic function) for each type of metadata disk block (e.g., inodes) to translate the set of blocks referenced by the metadata into a form recognized by the kernel [13]. The determinism of this function allows Xok to verify that metadata can only claim ownership of the correct set of disk blocks. This is done by verifying, when the metadata is updated, that the set of blocks claimed by the new metadata is the same as the set claimed by the old metadata with the intended change applied. This mechanism is only sound if the metadata decoding function is known to be deterministic.

2.3 Untrusted code execution

Purity gives us a way to execute untrusted code safely: we first verify that the untrusted code is pure, and then many useful privacy and security properties will follow. In particular, the lack of side effects means that the pure, untrusted computation cannot violate the integrity of the rest of the program it interacts with², so pure code inherently executes in a sandbox.

Purity can also be used to structure programs in a way that reduces our reliance upon the correctness of some subset of the code. If we use a pure method to process (possibly malicious) data from an untrusted source, and if the output from the pure method is no more trusted than its input, the method doesn't need to be trusted to defend itself from malicious data successfully. Even if a malicious input is able to somehow subvert the proper operation of that method, at worst it can only influence the result of the pure computation; it cannot harm the proper operation of the rest of the program.

Bernstein's discussion of address-extraction code in `sendmail` [4] illustrates these ideas well. The address-extraction code is responsible for parsing an email message and extracting an email ad-

²Untrusted code can still deplete resources or fail to terminate. Limits on resource usage or looping would be needed if denial of service is a concern [22], but that is beyond the scope of this paper.

dress from a particular header. At one point, this code contained a remotely exploitable vulnerability that allowed an attacker to gain root by taking control of `sendmail`. Bernstein proposed an alternate architecture:

Suppose that the same address-extraction code is run under an interpreter enforcing two simple data-flow rules:

- the only way that the code can see the rest of the system is by reading this mail message;
- the only way that the code can affect the rest of the system is by printing one string determined by that mail message.

The code is then incapable of violating the user's security requirements. An attacker who supplies a message that seizes complete control of the code can control the address printed by the code—but the attacker could have done this anyway without exploiting any bugs in the code.

We note that Bernstein's two conditions are exactly determinism and side-effect-freeness, so implementing the address-extraction code as a pure method would provide the desired security benefits.

Determinism allows us to bound what information a pure method can read—in particular, the method can only observe the value of objects that are reachable from one of its arguments, but cannot gain any information about any other data in the program. Moreover, deterministic code cannot listen on covert channels: for instance, any differences in behavior due to timing information or resource limits would violate the determinism properties. This ensures that the untrusted method cannot spy on any sensitive program state that was not explicitly provided to it.

Purity also limits the untrusted code's ability to leak sensitive information to others through overt channels. It can communicate to others only through its return value (or thrown exceptions) and its resource consumption. However, it can transmit over a timing- or resource-based covert channel to a receiver that is not pure. For instance, we might download an untrusted tax calculator and verify that it is pure before executing it. Then even if we type our salary into it, it cannot leak our salary to others directly, though it may be able to leak our salary through a covert channel.

Purity may also be useful for application extensions and plugins. For example, consider an image viewer that, out of the box, supports only a handful of image formats. It might allow installation of a plugin for viewing images in a different format only if that plugin is written as a verifiably pure function that, given the contents of an image file, returns a bitmap to be displayed by the image viewer. Once verified as pure, any such plug-in could be downloaded and executed safely; it cannot gain any information about other private information stored on the system, nor can it corrupt the state of any other part of the program.³

2.4 Building robust systems

Pure methods are also helpful for writing trustworthy security-critical code that mediates between untrusted components.

For the purposes of preserving application integrity, pure methods are always safe to expose to untrusted code. Their functionality could always be duplicated by the untrusted code itself, so they cannot pose an additional threat. Pure methods may still be part of the TCB, but only if their behavior is trusted for semantic correctness, not because the method is granted privileged access to program internal state. This is a consequence of the lack of side effects. It

³It should be noted that purity does not eliminate all threats that the plugin could pose to the program. The invoker of a pure plugin method must still ensure that only appropriate data is passed to the plugin and defend against unexpected return values from the plugin.

is possible, however, that specific instances of immutable objects, and thus their associated pure methods, might convey confidential or malicious information. One must still be careful about data flow.

A pure method is automatically “defensively consistent” [18, § 5.6], provided that in the absence of malice it provides correct service to individual clients and provided that each invocation of the pure method processes information from only a single client. (An object that serves multiple clients with independent interests is defensively consistent if, even when one client violates its preconditions, it continues to comply with its specification for other clients that satisfy its preconditions.) Defensive consistency ensures that one malicious client cannot attack other clients who may rely upon the same pure method.

Bernstein presents an example in which a pure `jpegtopnm` converter receives a JPEG image from the network, decompresses it, and outputs a raw bitmap of the image to a user [4]. The “client” here is the sender of the image, who as an arbitrary remote party, is untrustworthy and may wish to corrupt or disrupt processing of images from other sources. A defensively consistent implementation would thwart such attacks by continuing to provide correct conversion of all images originating from other senders, even if one sender sends malformed data with the intent to exploit the converter. The purity of the converter ensures defensive consistency because it allows one to know that each image is converted independently from any others, preventing a malformed image from affecting the processing of other images.

2.5 Bug reduction

Pure functions can help us eliminate certain classes of bugs. Of course, anything that reduces the number of bugs in security-critical code helps security.

A pure computation is automatically thread-safe, requiring no locks, and can always be run in parallel with other computations without risk of interference. Determinism guarantees that concurrent operations cannot disrupt its correct execution, and the lack of side effects means that it cannot disrupt other computations.

Reproducibility is particularly useful when debugging and testing applications. It is often the case with modern applications that bugs are discovered in the wild rather than during testing, due to novel configurations that were not considered during testing. In many cases, it can be difficult to reproduce the bug as there are a number of hidden variables that cause the behavior of a program to differ between runs. If a method is pure, any failure of the method will be reproducible given the same well-defined, bounded set of inputs. This known set of data can be collected and used to reproduce the bug for the developer, who can then fix the program.

Deterministic functions can also make testing more effective. If the computation is deterministic, we only need to cover any particular input once; on the other hand, if the computation is nondeterministic, it may conceal bugs that trigger nondeterministically, so it is difficult to know whether we have tested all possible behaviors. For instance, Bernstein cites dealing with nondeterministically triggered error cases as one challenge in testing `qmail`, and proposes that testing would have been easier if the code had been structured as a purely functional computation plus a simple wrapper that interacts with the environment (so that the wrapper can be easily mocked in testing) [4]. Verifiable purity would enable developers to check that this discipline was followed correctly and preserve it under maintenance.

2.6 Assertions and Specifications

It is widely accepted that assertions should be side-effect free. If evaluating the assertion condition causes no side effects, a pro-

gram that always satisfies the assertion will behave the same way whether the assertion is enabled or disabled. This restriction could be checked by a lint-type tool that would warn about potentially impure assertions.

In applications where assertions are used for debugging, it is also helpful to know that the assertion condition is deterministic. If a deterministic assertion succeeds, we know that it will not fail on another run of the program due to dependence on seemingly unrelated state or nondeterministic behavior of the underlying platform. Sometimes, programmers use assertions specifically to check for and abort in the face of incompatible platform configurations. In deterministic languages like Joe-E, however, platform-specific behavior is mostly hidden from the program, which would reduce the need for this pattern.

Some specification languages allow methods to have pre- and post-conditions that are defined using the same language as the code, and these conditions may call other methods. For instance, in JML, a specification language for Java, specifications are only supposed to call methods that are “pure.” JML’s notion of purity forbids side effects but does not require determinism and places no restrictions on what state the method’s behavior may depend upon [15]. Since JML specifications can be compiled to assertions and checked at runtime [8], the purity requirement is intended to ensure that these assertions do not change the program’s semantics.

We argue that pre-conditions, post-conditions, and object invariants should be deterministic as well as side-effect free. When methods are used in specifications, the specification cannot be considered fully defined unless the method is deterministic. In particular, if the requirements on the method are predicated on external state whose value changes from invocation to invocation, it will not be possible to statically verify that the method satisfies its contract. While JML’s restrictions on side effects in specifications may suffice to prevent runtime enforcement from changing the semantics of the program, static checking is more difficult than it would be if specifications were functionally pure.

We have not implemented a tool for checking the purity of assertions, but verifying the purity of some assertions would be a straightforward extension of our techniques for methods. There are some common patterns for assertions and specifications (such as those that require executing methods on non-immutable objects) for which our approach may not be applicable.

3. DEFINITIONS

Our definition of functional purity derives from the concept of a mathematical function, a well-defined one-to-one mapping from inputs to outputs. We consider a method in a program to be functionally pure if and only if it is both side-effect free and deterministic.

3.1 Side-effect freeness

A method is *side-effect free* if the only objects that the method ever modifies are created as part of the execution of the method. This definition permits the method to create and modify new objects, any subset of which may be reachable from its return value, but does not permit it to make any change that would be observable from outside the method.

In addition to this linguistic notion of side-effect freeness, we also require that a pure method not cause any side effects outside the language environment, with the exception of resource consumption (memory and CPU cycles). For example, it must not write to files, communicate over the network, or print to the console. This is necessary to soundly constrain the effects of untrusted code.

3.2 Determinism

A mathematical relation is considered a function if each distinct input is associated with a single specific output. Any two evaluations of a mathematical function with the same inputs will give the same result. This result depends deterministically on the inputs, and nothing else. Our determinism requirement for functions in a program is analogous. We want any two calls with equivalent arguments to a pure function to give the same result. The result must depend only on the arguments and not on other global or thread-local state such as the current time or the stack trace.

For numbers and other mathematical constructs, there is a well-accepted canonical definition of what it means for two sets of function inputs to be the same, namely mathematical equality for each argument. We need a similarly precise definition for equivalence of arguments in programs. Element-by-element equality works for value types, but pointers and references raise questions. Should two calls of a method be equivalent if their arguments have the same numeric values but different aliasing relationships? If they have the same aliasing but reside at different addresses in memory?

There is not a single obviously right answer to these questions. Determinism is thus a *parameterized* property: given a definition of what it means for arguments to be equivalent, a method is deterministic if all calls with equivalent arguments return results that are indistinguishable from within the language. The determinism guarantee is only useful for calls whose arguments are equivalent according to the definition of equivalence. The definition should make semantically equivalent invocations (those that look equivalent to the programmer or auditor) have equivalent arguments.

If the criteria for equivalence include memory layout information such as the concrete addresses of pointers, invocations will essentially never be equivalent, and the determinism guarantee will be meaningless. For types that serve purely as collections of immutable data, we can avoid addresses completely by comparing purely by value, not by reference. This prevents pointer aliasing from causing seemingly equivalent invocations to be distinct. For example, consider a Java method that concatenates two strings. If aliasing information is included, `concat(str, str)` will not be equivalent to `concat(str, str.clone())` because in the first case both arguments refer to the same object, and in the second case they refer to different objects. For other types (e.g., graph nodes) that are generally compared by reference, aliasing information may be important to include. Our equality definition (§4.1) excludes memory layout details but includes aliasing relationships for types where identity provides the semantic notion of equality.

4. APPROACH

Our approach to purity is based on leveraging the properties of a *deterministic object-capability language*, i.e., an object-capability language that (a) has no nondeterministic language primitives and (b) requires a restricted capability for any access to nondeterminism.

An object-capability language [18] is one with the following properties:

- all state that can be communicated between methods is stored in objects
- all objects can only be accessed by references
- references can only propagate by being passed as arguments or being stored in a shared object
- references are unforgeable (for instance, the language must be memory-safe, and it must not permit unsafe casts)
- access to references is strictly limited by lexical scoping of variables and transitive reachability of references

In such a language, references serve as capabilities, and capabilities can be granted only by explicitly passing references. For these properties to be effective in restricting code's effects, the global scope must not contain any capabilities to affect program state or the outside world. In other words, using methods and objects in the global scope it must not be possible to have any effect on mutable program state or external effects aside from resource consumption.

In a deterministic object-capability language, the observable global state must never change and must be the same on every execution. Any data returned from global methods is considered to be part of the observable global state. This means, for example, that no globally accessible methods can provide the time of day or state of the filesystem, even though this data is not explicitly stored anywhere in the global scope. A method's view of global state will thus be the same every time it is invoked, so globals can effectively hold only compile-time constants. Then, since the only variables in a method's scope are globals and arguments, any variation in the method's behavior can always be attributed to differences in its arguments.

From §3.2, we must specify when function arguments are equivalent. We do so in the following section, in which the concrete arguments to a method invocation are considered a set of named references. For an instance method, the implicit reference to the target object is treated as an argument named "this".

4.1 Equivalence of reference lists

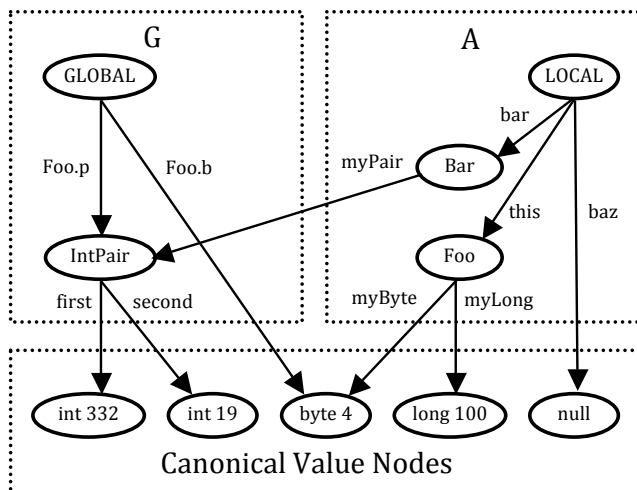
At a high level, we consider two sets of named references equivalent if their reachable object graphs (including values, types, and aliasing relationships) are isomorphic. Figure 1 gives an example; the rules for its construction follow.

Many object-oriented languages include both *reference types* and *value types*. Objects of reference type have an identity distinct from their value. The language distinguishes between references that point to the same object and those pointing to different objects with identical contents. Any type with mutable fields is by necessity a reference type, as changes to one instance will not affect another, but immutable objects can also be reference types if the language provides a way to test for object identity (such as Java's `==` operator). In contrast, value types can be compared only by value; there is no other notion of identity for these types. In Java, the primitive types (`boolean`, `char`, and the integer and floating point types) are value types, whereas `Object` and all its subtypes (including arrays) are reference types.

For simplicity in the following formal definition, values are represented in the object graph as references to canonical instances. One such instance exists for each distinct value of each value type. Similarly, we treat null pointers as references to a single canonical null object that belongs to every non-primitive type.

Let G be the set of named global (static) variables in the program. Let A be a set of named object references (such as the arguments to a method). We then define the *reachable object set* A^* corresponding to the set A as the transitive closure of objects reachable by following references from all fields of the objects pointed to by $A \cup G$. (The global state of the program will be the same across all executions, so the set of objects reachable from G won't change. This portion of the graph is only included in order to represent all observable aliasing relationships.)

The *object graph* for the reference set A is then constructed as follows: We create special nodes labeled *Global* and *Local*. We construct a canonical node for each primitive value in the program, in addition to a canonical node for `null`. For each reference-type object in A^* , we construct a node labeled with its concrete type. For variables in G or A , we add edges originating in *Global* and



```

public class Foo {
    public static byte b;
    public static IntPair p;

    private long myLong;
    private byte myByte;

    public void foo(Bar bar, Baz baz) {
        // ...
    }
}

public class Bar {
    private IntPair myPair;
    // ...
}

public class IntPair {
    public int first, second;
}

```

Figure 1: An example of an object graph and corresponding Java class definitions.

Local respectively. These edges point to the node representing the object pointed to by the variable and are labeled with the variable’s name. For each field of each object in A^* , we draw a directed edge from the node holding the reference to the node representing the referenced object (or canonical value), labeled with the field name.

Two sets of object references are considered equivalent if they result in identical object graphs (nodes and edges with the same labels). Note that the object graph reflects all aliasing relationships between non-value objects reachable from the set A and from global variables.

4.2 Immutability

We also need the language to provide support for types that are verifiably immutable. An object is immutable if its state, and the state of all objects reachable from it, can never change during its lifetime. Immutability is transitive: all objects reachable from an immutable object must themselves be immutable. To statically verify purity, we need the language to provide some way to verify that a type T is immutable, i.e., that every instance of T will be immutable.

Objects that represent a capability to affect or observe external state must not be considered immutable, even if their explicit state in the language is immutable.

5. PURE METHODS

Languages that meet our requirements (§ 4) make verification of purity easy. The condition is simple: If all parameters to a method (including the implicit `this` parameter) are statically declared to be of an immutable type, then the method is pure.

A pure method cannot cause any side effects, because it will never be able to obtain a reference to a shared mutable object. The only mutable objects it can access are those that it creates itself, but any changes it makes to them are not side effects since they will not be visible outside the method’s execution. The method is also prevented from causing any effects external to the language runtime, as all such effects would require an appropriate capability. Such capabilities are not considered immutable and thus are unavailable to methods with only immutable arguments.

The only state observable to the method is that in the global scope and reachable from its arguments. Since the global scope is constant (essentially determined at compile time), the only varying

state it can observe is from its arguments. As immutable arguments cannot provide a view of external state, the only observable state that immutable arguments can provide is captured in the definition of reference-list equivalence above (§ 4.1). For a language without any inherently nondeterministic constructs, ensuring that two invocations of a method start in indistinguishable states is sufficient to ensure that they will terminate identically.

In our approach, purity is part of the contract of a method: we can verify that a method is pure simply by examining its type signature. This is powerful, because it means that we do not need to inspect the implementation of the method or of other code that it might call. Internally, it may make use of mutating operations and impure methods. Any operation expressible in the language, no matter how much internal mutation it contains, can be verifiably pure if a wrapper method is written with a pure interface. We do not need an automated tool to identify for us which methods are pure; instead, programmers can recognize pure methods from their type signatures. An annotation may be used on such methods in order to ensure that they remain pure under maintenance, i.e., that their argument types remain immutable in later versions of the program.

6. IMPLEMENTATION

The object-capability language Joe-E was designed by Mettler et al. and specified as a technical report [16]. The contribution of this paper is not the design of the object-capability language Joe-E, which will be the subject of a forthcoming paper. Instead, this paper covers our subsequent work to make Joe-E a *deterministic* object-capability language and how this enables verification of purity.

Joe-E is a subset of Java; the Joe-E verifier, implemented as an Eclipse 3.2 plug-in, checks Java source code to confirm that it falls within the Joe-E subset. Any Java program that is accepted by the verifier is also a Joe-E program with the same semantics, but not all Java programs pass the verifier.

6.1 Side effects and Nondeterminism

The restrictions on globally-available side effects and nondeterminism are accomplished by exposing only a subset of the fields and methods defined in the Java libraries to Joe-E code. The Joe-E language defines a whitelist of fields and methods from the Java libraries that Joe-E code is allowed to use; the Joe-E verifier will reject programs that make reference to any field or method not on

```
Integer.getInteger(SYSTEM_PROPERTY)
System.identityHashCode(new String())
string.intern() == string
Integer.newInteger(integer) == integer
```

Figure 2: Several methods from the Java library are nondeterministic and callable by all Java code. These are not exposed to Joe-E code.

```
boolean randomBit() {
    return (new Object().hashCode() % 2) == 1;
}
```

Figure 3: This Java method is not deterministic. It is not legal Joe-E code, because Joe-E forbids calls to `Object.hashCode()`.

the list. We use this mechanism to prevent Joe-E code from calling any method that exposes the ability to observe or modify the environment outside the JVM, provides access to nondeterminism, or allows reading or writing of global mutable state. See Figure 2 for examples.

6.1.1 Object identity

In Java, objects have identity: conceptually, they have an “address”, and we can compare whether two object references point to the same “address” using the `==` operator. This notion of object identity can expose nondeterminism. As depicted in Figure 3, `Object.hashCode()` exposes nondeterminism, which is incompatible with reasoning about the purity of a method solely by examining its type signature. We have the Joe-E verifier forbid calls to methods such as `hashCode()` that expose nondeterministic representations of object identity.

Also, `String.intern()`’s static object cache contains global mutable state that is visible to Java code as a result of object identity (see Figure 4). For some types, such as `String`, we sidestep the problems caused by object identity and simplify reasoning by making the type a value type in Joe-E. This is accomplished by prohibiting Java’s object identity comparisons on these types: e.g., Joe-E code is not allowed to use the `==` operator on `Strings` or other value types. This makes our determinism guarantees more meaningful and more useful, as the programmer does not need to worry about aliasing relationships between variables of these types. Otherwise, a pure method could act differently on two invocations that the programmer might see as identical, but that differed only in aliasing relationships. For example, a string upper-casing method could behave differently if its argument aliases a global variable. In Joe-E, most immutable types are also value types, allowing more intuitive reasoning about equivalence of invocations.

6.1.2 Exceptions

We treat a thrown exception as a form of return value from a method, since the caller can catch and obtain a reference to the thrown object. This makes it challenging to reason about determinism in Java, for two reasons.

First, every `Throwable` object contains a stack trace generated when the throwable is constructed. This operation is not deterministic in the arguments to the throwable’s constructor, because the stack trace stored in the throwable depends on the calling context, which is not a function of the constructor’s arguments. Any construction of an exception (which can happen implicitly) thus results in an object which in Java would be a source of nondeterminism. We solve this problem by preventing programmatic access to the stack trace. Specifically, Joe-E code is not allowed to call the `getStackTrace()` and `printStackTrace()` methods.

```
boolean previouslyInterned(String s) {
    String t = new String(s);
    return t.intern() != t;
}
```

Figure 4: This Java method’s return value is not a deterministic function of its input. `previouslyInterned("foo")` returns true iff some other code has previously called `s.intern()` on a string `s` such that `s.equals("foo")`. This is not legal Joe-E code, because Joe-E does not allow using `!=` on `Strings`.

```
int estimateAvailStackSpace() {
    try { return estimateAvailStackSpace() + 1; }
    catch (VirtualMachineError e) { return 1; }
}
```

Figure 5: The JVM throws a `StackOverflowError` when the available stack space is exhausted, so this recursive method’s return value is nondeterministic.

Second, the Java Virtual Machine exposes true nondeterminism when it throws a `VirtualMachineError`. This occurs under a number of exceptional conditions, some of which (like running out of memory) can be triggered by the application. Consequently, virtual machine errors are a source of nondeterministic behavior; see, e.g., Figure 5.

This limits the guarantees we can provide: we cannot promise that whether or not a method terminates successfully will be a deterministic function of its arguments, since other conditions (e.g., the amount of free memory) might influence whether it aborts with a `VirtualMachineError`. Instead, we provide the following guarantee: any two calls to a method with equivalent sets of arguments will yield equivalent results, as long as neither method aborts with a `VirtualMachineError`. On the other hand, if one or both calls throw a `VirtualMachineError`, then we promise nothing.

We mitigate this shortcoming by ensuring that if the JVM does throw a `VirtualMachineError`, then the program will terminate immediately—the error will propagate to the top level and no Joe-E code will execute after the error occurs. To enforce this property, the Joe-E verifier prohibits catching `Error` or any subtype of `Error`. In addition, Joe-E must also prohibit the use of `finally` clauses, as they allow code to execute after (and in response to) a `VirtualMachineError`. See Appendix 10 for details.

6.2 Immutability

To support reasoning about purity, we want to allow the programmer to write user-defined classes that are verifiably immutable. The programmer can communicate to the Joe-E verifier that class `C` is intended to be immutable by declaring it to implement the interface `org.joe_e.Immutable`. The verifier then confirms that `C` truly is immutable by checking that all its fields are `final` and have a static type that is a primitive or immutable.

Unfortunately, this is not quite sufficient. In Java, code that has access to a partially constructed object `O` can read `final` fields of `O` before they have been initialized, so reading the same object’s fields twice might yield two different answers. To prevent this anomaly, Joe-E places several restrictions on all constructors to ensure that a reference to the object being constructed cannot escape from the constructor [16]. The most relevant is that constructors are prohibited from calling instance methods on the object being constructed.

Some reference types in the Java library, such as strings, are observationally immutable but are not declared to implement the `Immutable` interface. The Joe-E verifier handles these classes specially, treating them as if they did implement `Immutable`.

Java arrays are mutable, but are often used in situations where their mutability is not useful or desirable, as they are the simplest representation of a collection of objects. Therefore, Joe-E introduces a class, `ImmutableArray`, for storing an immutable sequence of immutable objects. Specialized subtypes are provided for holding primitive values without having to “box” them into objects; e.g., `ByteArray` holds an immutable sequence of bytes.

6.3 Verifying Purity

With our extensions, Joe-E makes it easy to reason about purity. In particular, in a Joe-E program, if all of the parameters to a method (including the implicit `this` parameter) are immutable, then the method is pure. For instance, suppose that we are decompressing a compressed file from an untrusted source. We might design the decompression interface as follows:

```
ByteArray decompress(ByteArray compressed)
```

This function will be pure, so even if the decompression code is buggy or insecure, a malicious compressed file cannot cause it to corrupt other application data structures. This allows us to contain the effect of any security holes in the decompression code.

As another example, suppose that we are building election tabulation software, which reads the contents of memory cards from the voting machines in the field, parses that data, accumulates the votes, and produces a report summarizing the tallies and winners. The parsing, accumulation, and report-generation code might be implemented following this interface:

```
String tabulate(ImmutableArray<CardData> cards)
```

If `CardData` is an immutable data structure holding the data read from a single memory card, this function will be pure. Hence we can be confident tabulation will be deterministically repeatable, and that the tabulation operation cannot (even if it is buggy or insecure) corrupt other election data.

Some caveats apply. The soundness of our determinism guarantee depends on consistent behavior from the portion of the Java libraries that Joe-E programs are allowed to call. While it is straightforward to block truly nondeterministic library methods, the semantics of some methods differs between Java library releases, including useful methods that are fully deterministic within a particular version. For example, the behavior of many string and character routines has changed between versions to reflect characters added to the Unicode specification. For this reason, our implementation only guarantees reproducibility between executions using the same library version. Since Joe-E does not currently require all assertions to be pure, our determinism guarantee is also predicated on whether or not assertions are enabled. In Java, results of floating point operations may differ between JVMs, which may cause code to have platform-dependent behavior.

7. EVALUATION AND EXPERIENCE

Our approach is intended primarily for programmers developing new code in Joe-E with verifiable purity in mind. Since Joe-E is intended to be as familiar as possible to Java programmers, we wanted to understand to what extent our approach would require Java programmers to change the coding style they are used to. We chose three Java libraries and retrofitted them (a) to pass the Joe-E verifier and (b) to have verifiably pure methods and resulting security properties. The refactoring was performed by a programmer who had no prior experience using Joe-E or any other object-capability language.

We give a detailed account of our experience, for three purposes: (1) to give the reader a sense of the type and magnitude of changes

	Source lines of code		Num. classes		Num. methods	
	Before	After	Before	After	Before	After
AES	319	276	1	1	9	9
Voting	688	692	25	25	80	79
HTML	12,652	10,848	94	99	965	947

Table 1: Basic code metrics for the three libraries used for evaluation, as measured both before and after refactoring.

that were necessary, (2) to understand the programming patterns that could potentially act as a barrier to the adoption of our system, and (3) to evaluate the strengths and limitations of our approach to verifiable purity. See Table 1 for the three applications we analyze. (We used the Eclipse Metrics Plugin [25] for all code metrics.)

7.1 AES library

7.1.1 Motivation

We started with an open-source AES implementation written in Java [6]. We sought to prove that the `encrypt` and `decrypt` methods are pure. This would then enable us to check at runtime that these methods satisfy the inverse property, as described in Section 2.2.

7.1.2 Changes to the codebase

First, we refactored the code to pass the Joe-E verifier. The AES library initially contained mutable static state: it used static variables of array type to hold the S-box tables. We replaced these with `ImmutableArrays`, to meet Joe-E’s requirement that all static variables be immutable.

Second, we refactored the class to provide verifiably pure methods. Originally, the AES library’s interface had this type signature:

```
public AES()
public void setKey(byte[] key)
public byte[] encrypt(byte[] plain)
public byte[] decrypt(byte[] cipher)
```

After refactoring, the signatures for the relevant methods and constructors became:

```
public AES(ByteArray key)
public ByteArray encrypt(ByteArray plain)
public ByteArray decrypt(ByteArray cipher)
```

Method signatures for `encrypt` and `decrypt` were changed so that all parameters would have an immutable type, thus making the methods verifiably pure. This was accomplished by replacing each `byte[]` array with a `ByteArray`. Also, because `encrypt` and `decrypt` are instance methods on the AES class, we had to make the AES class immutable. As an immutable class, it can no longer have its key specified using a setter method that mutates its state. Instead, the key is specified as an argument to the constructor. Immutability of AES also required making all instance variables `final` and immutable. To accomplish this, we had to remove debugging trace information from the class. In a case where preserving such information is important, a suitable solution would be to return from the top-level methods an object containing both the original return value and the debugging trace for that method call.

Notice that we also changed the return type of the `encrypt` and `decrypt` methods to an immutable type. This was not strictly necessary for verifying the purity of the AES library, but it helps clients of the AES library write their own pure methods that manipulate data returned from the AES library. In general, returning an immutable data structure helps verify purity of other parts of the code.

After our refactoring, clients of the AES library are able to check that decryption is the inverse of encryption by inserting

```
k.check(x);
```

before every call to `k.encrypt(x)`, where instance `k` is of type AES. The `check` method can be defined as follows:

```
public void check(ByteArray x) {
    assert(decrypt(encrypt(x)).equals(x));
}
```

Since this method is pure, inserting the call to `check` cannot change the program's behavior. Moreover, this call ensures that `encrypt` and `decrypt` satisfy the inverse property for every value `x` that is ever encrypted by any client of the AES library.

7.2 Voting machine

7.2.1 Motivation

Next, we examined the serialization and deserialization code of an experimental voting machine implementation [24]. We refactored the code to make serialization and deserialization pure. Our goal was to confirm at runtime that deserialization is the inverse of serialization, following the pattern described in Section 2.2. This ensures that all votes that are successfully recorded will be read back correctly during vote tallying.

7.2.2 Changes to the codebase

Nearly all changes made were simply replacing standard Java arrays with Joe-E immutable arrays. Another common change was adding the `Immutable` interface to classes that were already observationally immutable. (This required nothing more than adding `implements Immutable` to the class declaration.)

Another modification involved the use of a monotonically increasing serial number to filter out duplicate ballots. The last received serial number was stored as a static field inside the `BallotMessage` class. Inside the deserialization method, the serial number of the ballot was compared with the static value of the most recently received serial number; if the serial number was already received, the deserialization method would return null. With detection of duplicate ballots written in this way, the deserialization is not deterministic. To fix this, we separated the deserialization functionality from duplicate ballot suppression.

The only other significant change necessary was to require that a `Ballot` received all of its `Races` upon construction. Prior to refactoring, the `Ballot` class exposed a method `addRace(Race r)`. This method had to be removed in order to make the `Ballot` class immutable.

The method that we wished to make verifiably pure serves to check the serialization. The method is called after the object has been serialized to an array of bytes, and tests that the serialized form deserializes to match the original ballot. Its signature was initially:

```
public static boolean
    deserializesTo (byte[] serialized,
                  BallotMessage bm)
```

After refactoring, the method was changed to use a `ByteArray` instead of `byte[]`. The actual deserialization, which is performed by a constructor that takes a `ByteArray`, is also verifiably pure.

7.3 HTML parser

Our third application, an HTML parser [20], was a much larger and more instructive undertaking. Since our modifications to this

library were significant, we ensured that it retained its functionality by verifying that our modified version and the original version produced the same results when run on a corpus of HTML test cases [11].

7.3.1 Motivation

Our primary goal was to refactor the code to make the top-level `parse` method pure. From a security perspective, a pure parse method is valuable for any system in which parses need to be performed on behalf of different users or using data from different sources. An example of this is on a web forum, where posts to the forum must be sanitized to prevent cross-site scripting attacks. A pure parse method together with a pure sanitization routine ensures that there can be no accidental data contamination between different posts, and that no private information about a user can be accidentally leaked into another post or to another user. Additionally, a pure parse method guarantees that a given parse is reproducible on any machine under virtually any circumstances.

Before refactoring, the top-level method signature, which resides in the `Parser` class, was the following:

```
public NodeList parse (NodeFilter filter)
    throws ParseException
```

Originally, neither the `Parser` class nor the `NodeFilter` class was immutable, and hence this method was not verifiably pure.

7.3.2 Mutable static state

We removed several instances of mutable static state from the HTML library, so that the code would pass the Joe-E verifier. For example, originally the only way to pass options to the parser was to set a global flag, `parse`, and then restore the flag, as follows:

```
boolean oldValue = SomeClass.SOMEFLAG;
SomeClass.SOMEFLAG = true;
try { parser.parse(); }
finally { SomeClass.SOMEFLAG = oldValue; }
```

This pattern seems to have been used to avoid propagating a configuration parameter through several levels in the call hierarchy. However, this use of global variables makes it harder to see how the flag is specified, renders the code thread-unsafe, and violates Joe-E's prohibition on shared mutable state.

We eliminated this pattern by augmenting the API with a top-level `parse` method that takes an extra argument and passes it as necessary to other parts of the program. The original top-level `parse` method remains, using a default value for the flag.

The original codebase also violated Joe-E restrictions by printing to `System.out` for debugging and reading the default locale using `java.util.Locale.getDefault()`. Using `Locale` objects in Joe-E code is not problematic, but the *default* locale is system-dependent and therefore non-deterministic. We instead modified the API to require that a `Locale` be passed as a parameter to objects needing access to the locale.

7.3.3 Instance method calls in constructors

We found many constructors that called other instance methods during their execution. As discussed earlier, Joe-E prohibits this (see § 6.2), so we had to eliminate all calls to instance methods from within constructors. This was bothersome—it was one of the few changes we had to make that did not reflect poor or nonstandard style in the original code—but fortunately we were able to work around the problem in every case by inlining the instance method, replacing the instance method with a static method, or using a factory method instead of a constructor. Nonetheless, this

```
String html = getHtmlStringFromSomewhere();
Parser p = new Parser(html);
NodeList list = p.parse(null); // null NodeFilter
// do something with the parse "tree" in list
```

Figure 6: A typical use of the `Parser` class. The HTML document is supplied to the constructor as a string. Then, the `parse` method is called with a `NodeFilter` as a parameter. A `NodeList` is returned, which contains a list of the top-level nodes from the HTML document.

experience suggests that the restriction on calling instance methods from constructors may place an undue burden on Joe-E programmers. We are currently considering less restrictive alternatives for future Joe-E releases.

7.3.4 Immutable classes

Once the HTML parser's code passed the Joe-E verifier, we refactored the `Parser` and `NodeFilter` classes to be immutable.

The original `Parser` class contained a `Lexer` as an instance variable. In order to make the `Parser` class immutable, this instance variable had to be removed due to the fact that a `Lexer` is inherently mutable. We refactored the code to construct and use a `Lexer` inside the top-level `parse` method. A typical use of the `Parser` class can be seen in Figure 6.

Making the classes that implement the `NodeFilter` interface immutable was straightforward, except for the `IsEqualFilter` class. This required significant effort due to the fact that this class, which tests whether two nodes are equivalent to each other, contained an instance variable of type `Node`. As a result, all classes that implemented the `Node` interface had to be made immutable, which necessitated removing all setter methods from any `Node` subclass and requiring that all fields were set upon construction of any subclass of `Node`.

Refactoring the `Node` subclasses to be immutable proved difficult due to a nonstandard construction pattern. The library used a prototype construction pattern to support the creation of custom parsers that recognize varying sets of HTML tags⁴: before parsing, the caller could register a set of `Node` prototypes. When the `Lexer` needed to construct a new `Node`, it would clone a prototype and then overwrite the relevant fields of the clone using setter methods.

We refactored the code to use a more standard construction pattern in which `Nodes` are constructed using a constructor that takes an argument for each instance field that needs to be set. Minor functionality was lost with this change, as it is no longer possible to create a custom `NodeFactory` (without creating a custom class implementing the `NodeFactory` interface) to recognize a different set of nodes.

Also, to make `Nodes` immutable, we had to split the `Page` class into two classes. Before refactoring, the `Page` class conflated two distinct purposes. It was used by both the `Lexer` and by the `Node` classes. The `Lexer` used a `Page` instance during lexing to maintain information about the current position of the cursor in the page and to get and unget characters. This inherently requires a mutable class. On the other hand, the `Node` classes only used the `Page` object for finding out the line and column numbers for characters in the page. This information is fixed and will never change after construction of a `Node`. To reflect this fact, we created an immutable class called `PageInfo` to hold this information and extracted it from the mutable `Page` class. Now, when the `Lexer` creates a `Node`,

⁴For example, one could create a parser that recognizes only `img` tags, and treats all other tags as generic tags with no hierarchical structure.

it obtains the `PageInfo` from the `Page` and passes it to the `Node`'s constructor.

As illustrated above, immutability is a property that necessarily spreads through related classes. We noticed similarities between these immutable data structures and those used in functional programming. For instance, immutable data structures must be constructed from the bottom up and hence are necessarily acyclic.

We also refactored the `parse` method to make the parse tree it returns be immutable. This is not necessary for the purity of the `parse` method, but it aids the creation of other pure methods that use the data structure returned by the `parse` method, since callers of the `parse` method can directly pass the parse tree it returns as a parameter to other pure methods.

7.4 Summary of patterns

Using a strictly-functional style throughout a program is the most reliable pattern for attaining verifiable purity, as it ensures that every method will be pure. Such a strict approach is generally not necessary to achieve useful purity guarantees. None of the three applications that we refactored were written in an exclusively functional style, either before or after our modifications. Our approach to purity requires only that immutable types (and thus functional programming style) be used for the interface of a pure function, allowing its internal algorithms to be written in an imperative fashion if the programmer so desires.

Objects that have cycles (for example, a tree with parent pointers or a doubly-linked list) pose a challenge for our approach. A cyclic object graph, even if it is observationally immutable once fully constructed, cannot be statically verified as immutable in our system. We may therefore be unable to verify purity for methods that use such objects.

Joe-E required us to eliminate the use of mutable static state and pass parameters explicitly as arguments instead of using mutable global variables. We found that this brought our code closer to a functional style and had benefits of its own.

We believe that new code can take better advantage of Joe-E's guarantees if the class hierarchy is designed with immutability in mind. If part of a class is immutable but the rest of the class is not, the entire class must be treated as mutable. Consequently, if a concept has separable mutable and immutable aspects, it may be helpful to represent it as two separate classes.

7.5 Waterken Server

The Waterken server is an extensible web server designed for building distributed web services [9]. Waterken is implemented in a mixture of Joe-E and Java. The Joe-E code was not retrofitted from Java to Joe-E, as in our previous examples, but was designed and implemented following object-capability principles. The Joe-E portion is substantial, comprising 8,246 source lines of code and 132 classes.

We counted the number of pure methods in the Waterken Joe-E code. (See Table 2.) Our results are somewhat surprising: a large fraction of methods (17%) and an even larger fraction of constructors (29%) are verifiably pure. While the code was written in an object-capability style, verifiable purity was not an explicit goal. This suggests that verifiable purity can (and does) occur as a natural consequence of object-capability discipline.

8. DISCUSSION

One advantage of our approach is that it can facilitate reasoning about side-effects and data dependencies for methods even if they do not strictly meet our requirements to be functionally pure. Since the accessible data and possible effects of a method are limited to

	Pure	Total	% pure
Methods	89	524	17%
Constructors	37	128	29%

Table 2: The number of pure and impure methods and constructors in the Waterken Server.

objects reachable from its arguments, these effects are still bounded even if some arguments are mutable. In particular, the method can only mutate objects that are reachable from its non-immutable arguments. Typing and capability reasoning can limit this set to a small portion of the in-memory objects in the program, e.g., the values in a single array of ints, or the private instance fields of an object.

One can sometimes use these bounded effects to achieve purity properties from methods that are not individually pure. For instance, consider the following set of operations on a non-immutable object `o`:

```
T o = new T(a);
o.f(b); o.g(c);
... // do something with o
```

If the constructor is pure and the arguments `a`, `b`, and `c` are all immutable, then the state of `o` after this sequence of operations will be a deterministic function of `a`, `b`, and `c` and no other side effects will occur. We will refer to a sequence of invocations with this property as a *functionally pure sequence*.

If all of the inputs are known in advance, the sequence above can be written as a single verifiably pure function; for this example, we would have:

```
T pure(A a, B b, C c) {
  T o = new T(a);
  o.f(b); o.g(c);
  return o;
}
```

The more interesting case is where the inputs are not known in advance, such as if some of them come from interactions with a user. In this case, some of the inputs depend on information received from the program, e.g., return values from invocations on `o`. This case can be expressed as a sequence of verifiably pure method calls if it is refactored to use a purely functional style. Specifically, we would refactor `T` to be immutable and replace each mutating instance method of `T` with one that returns both the original return value and a new object that has the modifications applied.

For cases in which making `T` immutable is impractical or cumbersome, we need a new set of rules sufficient to verify such a sequence is functionally pure. It is safe to add return values to the first scenario above, as long as they do not enable modifications to the object’s internal state. This limitation can easily be verified by requiring the return values to be immutable. (Thrown exceptions would also be a concern, but Joe-E already requires all throwables to be immutable). This set of restrictions is not as trivial to check as the ones needed for individual methods to be verifiably pure, but it allows for reasoning about useful properties of non-immutable objects.

The pattern allows for purity to be demonstrated in event-based interactive systems, such as a voting machine. Each voter’s actions constitute a stream of events that should be interpreted as they arrive to produce the voted ballot. Pure sequences can allow us to verify that the voted ballot and behavior of the voting machine are a deterministic function of the sequence of input events.

Functionally pure sequences also occur in Waterken’s implementation of deterministic server processes that react to input events. Each event causes mutations in the internal state of a handler object dedicated to that event’s connection. Because these mutations are local to the per-connection object, the behavior of each server process is a deterministic function of the sequence of input events it receives, even though each individual call to the event-processing method is not verifiably pure.

9. RELATED WORK

Object-capability languages have a long history [18]. Most of these languages are more functional in style than the Joe-E subset of Java, making it easier to limit side effects. In the E language, the `FUNCTIONAL` auditor examines an object to check that “every method on the object has no side effects and produces an immutable result depending solely on its arguments” [28]. The auditor verifies this property using runtime introspection on an object; in contrast, we verify it statically for individual methods. In this paper, we build upon Joe-E [16], another object-capability language designed by some of the authors of this paper in related work, but we expect that many of our techniques could be applied to many other object-capability languages as well.

We were inspired by the notion of “environment-freeness” [24], which is essentially the determinism portion of our notion of purity. Environment-freeness was used to verify determinism of a decoding operation and for fail-stop enforcement of the inverse property.

Most previous work on purity in imperative languages has focused on side-effect freeness and paid little attention to determinism. The definition of side-effect free used has generally been weaker than ours, as it has included only objects in memory and has excluded state external to the program. For legacy code, Rountev [21] and Salcianu and Rinard [23] provide pointer-based analyses that recognize side-effect free methods. Both address only in-language side effects; neither mentions any special treatment for native methods, which can cause external side effects. Analysis-based approaches have the advantage of being directly applicable to legacy code. Language-based approaches, on the other hand, provide more guidance for programmers in writing side-effect free methods.

In Joe-E, we make use of class immutability both in enforcement of determinism and for side-effect freeness. Specific classes in the standard Java type system are considered immutable; standard Java type safety and final field enforcement ensures that objects of such classes are never mutated after construction. An alternative is to use an extended type system that treats some references or instances as read-only while allowing others to be mutated. The C++ `const` qualifier for pointers is the most well-known example of this. It is a compilation error to assign the fields of, or invoke a non-`const` method on, a `const` reference. Its use in preventing side effects is limited because the restrictions are not transitive; it is possible to modify an object contained in a field reached via a `const` pointer. A transitive analogue of this was introduced by the KeyKOS operating system [10] as a “sensory key”; such a key prohibits writes and also causes all keys retrieved through it to be sensory. This concept is also found in the type system of a few programming languages to improve reasoning about immutability and side effects. Such types allow for documentation and modular checking of effect restrictions on a per-function basis.

Ieuralimschy and Rodriguez [12] use such a qualified type to enforce side-effect freeness in the SmallTalk-like language School. Methods annotated to be side-effect free are type-checked with all arguments and the instance pointer implicitly marked with an `old` qualifier. This qualifier prevents writes to the fields of `old` ob-

jects. The type checker only allows invocations of side-effect free methods on `old` objects, and treats the return values from all such invocations as `old`. This ensures that the only non-`old` (and thus mutable) objects that can be used by the method are ones it creates itself. The paper makes no mention of rules for dealing with mutable objects in the global scope or external side effects; their emphasis on soundness would suggest that School has neither. The Javari [26] type system provides similar qualifiers for Java, but instead of having a side-effect free annotation, Javari uses explicit `readonly` qualifiers as a transitive, sound version of C++’s `const`. Like C++, Javari provides a way for fields of a class to be declared as exempt from the `readonly` restrictions.

In addition to a sound, transitive version of `const` with no escape clauses for mutable fields, the D language [5] provides an instance-immutability qualifier `invariant` that can be used to achieve functional purity. Functions marked with the `pure` keyword must have only `invariant` arguments, can only read `invariant` global state, and can only call other `pure` methods. Their compiler restricts `invariant` variables in the global scope to constant values that can be computed by the compiler⁵, ensuring determinism. While this approach avoids the need to eliminate mutable state and determinism from the global scope, there is a substantial cost in expressivity as it prevents pure functions from making any use of impure functions and methods. The result is essentially of a partition of the program into imperative and purely functional portions, whereas our approach allows pure functions to make full use of the rest of the program, limited only by the references they hold.

The increased convenience of reference immutability (`const` or `readonly`) over class immutability is attractive, as one can just use the type qualifier with existing classes rather than modifying the type hierarchy. However, class or instance immutability is necessary to ensure determinism in a concurrent program, as otherwise a mutable alias can be used to concurrently modify the object. For non-concurrent programs, reference immutability would be adequate provided that the global scope can only store immutable references. As a general mechanism for defensive programming, reference immutability can only serve to protect the originator of a reference from unwanted modifications; the recipient of an immutable reference may still need to make a defensive copy.

Instance immutability, like provided in D, is an interesting alternative to class immutability that deserves further exploration. For Java, however, the lack of type safety for generics is likely to be an issue. For immutable classes, we perform runtime checks to ensure that the elements placed in an `ImmutableArray` are actually immutable; this would not be possible with a purely compile-time `invariant` type qualifier as would be required to preserve full compatibility with Java.

Spec# [3] and JML [7, 15] are extensions to C# and Java that allow the programmer to specify invariants on functions and classes. They follow Bertrand Meyer’s suggestion that classes and methods should have a contract specified by invariants [15]. They support annotating methods with the `pure` attribute, but purity as defined for JML includes only side-effect freeness and not determinism.

In strictly functional languages, like Haskell, nearly all functions are pure. Monads can be defined to allow writing in a more imperative style, in which each operation takes an input state and returns a monad instance that wraps the result along with auxiliary information such as side effects [27]. The monad type defines an operator for sequencing such invocations to obtain a final result; syntactic sugar makes this look like a sequence of imperative state-

⁵The D compiler can perform a substantial amount of computation to determine these values, unlike Java’s, which only pre-assigns literal constants.

ments. While some monads provide a way to retrieve a sequence’s final result integrated with any auxiliary information, other monads do not. They are “one-way”: once a value is wrapped with the monad, it never comes out. The I/O monad is an example. All functions that potentially expose nondeterminism or cause external side effects use this monad, which allows them to be recognized as potentially nondeterministic. All functions whose return type does not mention the I/O monad are functionally pure. While monads provided a means to use effects in Haskell, the language is primarily oriented at the functional style. In contrast to a mechanism for imperative patterns in a functional language, our approach is focused on being able to recognize pure methods in an otherwise imperative language. This reduces the changes needed to existing code and programming patterns.

Other systems have mixed imperative and functional programming styles to varying degrees. The Eiffel language [17] separates what it calls *commands* and *queries*. Commands may have side effects, while queries are supposed to be side-effect free. This is, however, only a convention; it is not enforced in any way. Similarly, both Euclid [14] and SPARK [2] define two distinct constructs for routines: *procedures* can have side effects, while *functions* are only able to compute a value (and thus are guaranteed to be free of side effects). In Euclid, functions can only import variables read-only, which prevents side effects and ensures determinism if the imported variables (which may be modified elsewhere) are treated as additional arguments. In SPARK, annotations on procedures specify exactly which variables can be modified by the procedure, and which variables their modifications are derived from. This and other information flow policies are verified by the SPARK Verifier.

Jif [19] extends Java with label-based information flow checking. Variable declarations are annotated with labels that indicate an owning principal and a policy for data stored in the variable. The policy can specify (a) the principals whose data the variable may *depend on* and (b) those whose data are allowed to *affect* the information stored in the variable. These restrictions are enforced statically, in cases where it is possible to statically guarantee that the policy is followed, and dynamically, in cases where it is not. As a special case, it is possible to specify data flow restrictions that ensure that a particular method is pure. In contrast, while our approach does not allow for the rich policies expressible in Jif, obtaining purity in Joe-E does not require the explicit specification of principals or policies.

10. CONCLUSIONS

Verifiable purity is useful for verifying many kinds of high-level security properties. A language with appropriate characteristics can greatly simplify the task of writing verifiably pure code. By combining determinism with object-capabilities, we describe a new class of languages that allow purity to be achieved in largely imperative programs. As such a language, Joe-E allows programmers to flexibly leverage verifiable purity while still using imperative algorithms in a familiar language.

Acknowledgments

We thank David Molnar, Karl Chen, Arel Cordero, Tyler Close, Toby Murray, Sandro Magi, Mike Samuel, our shepherd Sriram Rajamani, and our anonymous reviewers for helpful comments on earlier drafts of this paper. This research was supported by National Science Foundation grants CNS-0716715, CNS-0524745, CNS-0430585, and CCF-0424422.

11. REFERENCES

- [1] M. Backes, M. Dürmuth, and D. Unruh. Information flow in the peer-reviewing process (extended abstract). In *IEEE Symposium on Security and Privacy, Proceedings of SSP'07*, pages 187–191, May 2007.
- [2] J. Barnes. *High Integrity Software: The SPARK Approach to Safety and Security*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2003.
- [3] M. Barnett, K. R. Leino, and W. Schulte. The Spec# programming system: An overview. In *Proceedings of Construction and Analysis of Safe, Secure and Interoperable Smart Devices (CASSIS)*, 2004.
- [4] D. J. Bernstein. Some thoughts on security after ten years of qmail 1.0. In *CSAW '07: Proceedings of the 2007 ACM workshop on Computer security architecture*, pages 1–10, New York, NY, USA, 2007. ACM.
- [5] W. Bright. D language 2.0. <http://www.digitalmars.com/d/2.0/>.
- [6] L. Brown. AEScalc. <http://www.unsw.adfa.edu.au/~lpb/src/AEScalc/AEScalc.jar>.
- [7] L. Burdy, Y. Cheon, D. Cok, M. Ernst, J. Kiniry, G. Leavens, K. R. Leino, and E. Poll. An overview of JML tools and applications. *International Journal on Software Tools for Technology Transfer (STTT)*, 7(3):212–232, June 2005.
- [8] Y. Cheon and G. Leavens. A runtime assertion checker for the Java Modeling Language, 2002.
- [9] T. Close and S. Butler. Waterken server. <http://waterken.sourceforge.net/>.
- [10] N. Hardy. Keykos architecture. *SIGOPS Oper. Syst. Rev.*, 19(4):8–25, 1985.
- [11] HTML4 Test Suite. <http://www.w3.org/MarkUp/Test/HTML401/current/>.
- [12] R. Ierusalimsky and N. de La Rocque Rodriguez. Side-effect free functions in object-oriented languages. *Comput. Lang.*, 21(3/4):129–146, 1995.
- [13] M. F. Kaashoek, D. R. Engler, G. R. Ganger, n. Hector M. Brice R. Hunt, D. Mazières, T. Pinckney, R. Grimm, J. Jannotti, and K. Mackenzie. Application performance and flexibility on exokernel systems. In *SOSP '97: Proceedings of the sixteenth ACM symposium on Operating systems principles*, pages 52–65, New York, NY, USA, 1997. ACM.
- [14] B. W. Lampson, J. J. Horning, R. L. London, J. G. Mitchell, J. G. Mitchell, G. J. Popek, and G. J. Popek. Report on the programming language Euclid. *SIGPLAN Not.*, 12(2):1–79, 1977.
- [15] G. Leavens and Y. Cheon. Design by contract with JML, 2003.
- [16] A. Mettler and D. Wagner. The Joe-E language specification, version 1.0. Technical Report UCB/EECS-2008-91, EECS Department, University of California, Berkeley, August 7, 2008.
- [17] B. Meyer. *Eiffel: The Language*. Object-Oriented Series. Prentice Hall, Englewood Cliffs, NJ, USA, 1992.
- [18] M. S. Miller. *Robust Composition: Towards a Unified Approach to Access Control and Concurrency Control*. PhD thesis, Johns Hopkins University, Baltimore, Maryland, USA, May 2006.
- [19] A. C. Myers and B. Liskov. A decentralized model for information flow control. In *Symposium on Operating Systems Principles*, pages 129–142, 1997.
- [20] D. Oswald, S. Raha, I. Macfarlane, and D. Walters.

```
class IntException extends Exception {
    final int data;
    IntException(int data) { this.data = data; }
}

int nondet() {
    try { freemem(); return 0; }
    catch (IntException ie) { return ie.data; }
}

void freemem() throws IntException {
    int shift;
    try {
        for (shift = 0; ; ++shift) {
            new double[1 << shift];
        }
    } finally { throw new IntException(shift); }
}
```

Figure 7: finally clauses expose nondeterminism.

HTMLParser 1.6.

<http://htmlparser.sourceforge.net/>.

- [21] A. Rountev. Precise identification of side-effect-free methods in java. In *ICSM '04: Proceedings of the 20th IEEE International Conference on Software Maintenance*, pages 82–91, Washington, DC, USA, 2004. IEEE Computer Society.
- [22] A. Rudys and D. S. Wallach. Termination in language-based systems. *ACM Transactions on Information and System Security*, 5(2), May 2002.
- [23] A. Salcianu and M. C. Rinard. Purity and side effect analysis for java programs. In *VMCAI*, pages 199–215, 2005.
- [24] N. K. Sastry. *Verifying Security Properties in Electronic Voting Machines*. PhD thesis, University of California at Berkeley, 2007.
- [25] F. Sauer. Eclipse metrics plugin 1.3.6. <http://metrics.sourceforge.net/>.
- [26] M. S. Tschantz and M. D. Ernst. Javari: Adding reference immutability to Java. In *Object-Oriented Programming Systems, Languages, and Applications (OOPSLA 2005)*, pages 211–230, San Diego, CA, USA, October 18–20, 2005.
- [27] P. Wadler. The essence of functional programming. In *Conference Record of the Nineteenth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 1–14, Albuquerque, New Mexico, 1992.
- [28] K.-P. Yee and M. Miller. Auditors: An extensible, dynamic code verification mechanism, 2003. <http://www.erights.org/elang/kernel/auditors/index.html>.

APPENDIX: FINALLY CLAUSES

Figure 7 shows how to return a nondeterministic value without explicitly catching an `Error` by using a `finally` clause instead. The `freemem()` method tries to allocate larger and larger arrays of doubles until triggering an `OutOfMemoryError`. This causes the `finally` clause to execute, which then throws an `IntException` that hides the pending `Error`. The `IntException` contains nondeterministic state (how many arrays could be allocated before running out of memory), which is extracted from the exception and returned by `nondet()`.

Fortunately, Joe-E's prohibition of the use of `finally` does not reduce expressivity: Joe-E code can explicitly catch `Exception`, which allows the catching and appropriate handling of any non-`Error` throwable in the Java library.