

# Paragon for Practical Programming with Information-Flow Control

Niklas Broberg, Bart van Delft, and David Sands

Chalmers University of Technology, Sweden

**Abstract.** Conventional security policies for software applications are adequate for managing concerns on the level of access control. But standard abstraction mechanisms of mainstream programming languages are not sufficient to express how information is allowed to flow between resources once access to them has been obtained. In practice we believe that such control - information flow control - is needed to manage the end-to-end security properties of applications.

In this paper we present Paragon, a Java-based language with first-class support for static checking of information flow control policies. Paragon policies are specified in a logic-based policy language. By virtue of their explicitly stateful nature, these policies appear to be more expressive and flexible than those used in previous languages with information-flow support.

Our contribution is to present the design and implementation of Paragon, which smoothly integrates the policy language with Java's object-oriented setting, and reaps the benefits of the marriage with a fully fledged programming language.

**Keywords:** information flow, static enforcement

## 1 Introduction

The general goal of this work is to construct innovative design methods for the construction of secure systems that put security requirements at the heart of the construction process, namely *security by design*. To do this we must (i) understand how we can unambiguously formulate the policy aims for secure systems, and (ii) develop technology to integrate these goals into design mechanisms and technologies that enable an efficient construction or verification of systems with respect to those policies.

We address this challenge using a programming language-centric approach, presenting a full-fledged security-typed programming language that allows the programmer to specify how data may be used in the system. These security policies are then enforced by compile-time type checking, thus requiring little run-time overhead. Through this we can guarantee that well-typed programs are secure by construction.

But which security policies might we want for our data, and why do we need special support to express them? Certain security policies, for example access control, are relatively easy to express in many modern programming languages. This is because limiting access to resources is something that good programming language abstraction mechanisms are designed to handle. However, access control mechanisms are often a poor tool to express the intended end-to-end security properties that we wish to impose on our applications.

Consider a travel planner “app” which permits you to plan a bus journey, and even add your planned trip to your calendar<sup>1</sup>. In order to function, the app must have access to the network to fetch the latest bus times, and must have access to your calendar in order to add or remove schedules. But an app with these permissions can, for example, send your whole calendar to anywhere on the net.

What we want is to grant necessary access, but limit the *information flows*. In this case we want to at least limit the information flows from the calendar to the network while retaining the app’s ability to read and write to both.

Research on controlling these information flows has progressed over the last decades. In this paper we identify three generations of security properties and control mechanisms for information-flow:

**Information-Flow Control** In the 70’s, Denning & Denning pioneered the idea of certifying programs for compliance with information flow policies based on military-style secrecy classifications [7,8]. They used program analysis to validate that information labelled with a given clearance level could never influence output at any levels lower in the hierarchy – so for example a certified program could never leak top-secret information over a channel labelled as public.

The language FlowCAML [24], a variant of ML with information-flow types and type inference, represents the state-of-the-art in support for static Denning-style confidentiality policies.

**Beyond Mandatory Information Flow Control** Although a rigid, static hierarchy of security levels may be appropriate in a military message-passing scenario, it became quickly apparent that such a strict and static information flow policy is too rigid for modern software requirements. In practice we need a finer-grained and more dynamic view of information flow.

The concept of *declassification* – the act of deliberately releasing (or leaking) sensitive information – is an important example of such a requirement. Without a possibility to leak secrets, some systems would be of no practical use.

For example an information purchase protocol reveals the secret information once a condition (such as “payment transferred”) has been fulfilled. Yet another example is a password checking program that inevitably leaks some information: even when a login attempt fails the attacker learns that the guess is *not* the password.

With this in mind, the *Jif* programming language [17,20] can be seen as the next milestone after the pure Denning-style approach. Jif is a subset of Java extended with information flow labels.

As well as implementing an important distributed view of data ownership, the so-called *Decentralised label model* [18,19], Jif included the possibility of *declassification*, which provides a liberal information flow escape hatch for programs which would otherwise be rejected.

---

<sup>1</sup> The example is based on a family of actual Android apps (e.g. de.hafas.android.vasttrafik).

**Paragon, a Third-Generation IF Language** Declassification, in many shapes and forms, has been widely studied in the research community in recent years [23]. The large variety of declassification concepts is testament to the fact that there is simply no single right way to control the flow of information that goes against the grain. Moreover, it is not always natural to view information flow policies as consisting of “good flows plus exceptional cases” at all; in some situations there is no obvious base-line policy, and the flows which are deemed acceptable may depend on the state of the system at any given moment.

In earlier work [5] we introduced a new highly versatile policy language, *Paralocks*, based around the idea of *Flow Locks*. We demonstrated its ability to model a wide variety of policy paradigms, from classical Denning-style policies to Jif’s decentralised label model, as well as the capability to model *stateful* information flow policies. But the idea of using Paralocks as types in a statically-checked programming language was only demonstrated for a toy language. The question whether a Flow Locks-based policy language could feasibly scale to inclusion in a full-fledged programming language, to allow practical programming with information flow control, was left open.

The main contribution of this paper is to answer that question with an emphatic yes. We present the new programming language Paragon, which extends the Java programming language with information flow policy specifications based on an object-oriented generalisation of Paralocks. Not only does it turn out to be feasible, but the marriage of our stateful policy mechanism and Java’s encapsulation facilities yields a whole that is greater than the sum of its parts: it allows for the creation of complex policy mechanisms as libraries, giving even stronger control over flows and declassification than the policy language alone.

The remainder of the paper is structured as follows. Section 2 presents the language Paragon: its policy language and integration in Java; followed by a simple showcase of Paragon in Section 3. In Section 4 we give an overview of the enforcement mechanism and details of our implementation. Section 5 discusses our experience from two larger case studies. Related work in Section 6 and conclusions in Section 7 round out the paper.

## 2 The Language Paragon

Paragon is largely an extension to the Java language and type system. Our choice for Java is motivated by its relatively clear semantics and the wide adoption of Java in both commercial and academical settings. In addition, it allows us to reuse existing ideas from, and simultaneously compare Paragon with, Jif [17,20], the only (other) java-based full-fledged security-typed programming language to date. We discuss Paragon’s relation to Jif in more detail in § 6. We do not, however, rely on any particular features of Java for the integration of our policy language to work, and posit that it would be equally feasible to do this for other statically typed languages with safe memory management, e.g. ML or Scala.

In this section we give a high-level overview of Paragon and its various components, leaving more technical features, such as extending Java Generics, to the technical report version of this paper [1].

## 2.1 The Paragon Policy Language

The subjects of Paragon policies are the information-flow relevant entities, which we refer to as *actors*. An actor could be a user, a resource, a system component, an information source or sink, etc.; any entity that has an information-flow concern.

In Paragon, these entities are represented by *object references*. For instance, the code fragment below creates regular instances of the `User` and `File` class, where `alice` and `f1` can play a dual role; both as program variables, and as actors in Paragon policies.

---

```
User alice = new User();
User bob   = new User();
File f1   = new File();
File f2   = new File();
```

---

A *paragon policy* is used to label information containers in the program (fields, local variables), and specifies to which actors the information in that container is allowed to flow. A policy consists of a set of *clauses*, each specifying one particular actor, or one group of actors of a particular type.

For example, the policy `p1` states that information may flow to the specific users Alice and Bob, while the policy `p2` states that information may flow to any file:

---

```
policy p1 = { alice: ; bob: };
policy p2 = { File f: };
```

---

This makes the policy `{ Object o: }` the most permissive, and the policy with no clauses (denoted `{ : }`) the most restrictive paragon policy.

A clause may have a *body* that constrains the states in which the information may flow to the actors specified in the head. These constraints come in the form of *locks*; typed predicates representing the policy-relevant state of the system.

A lock can be opened or closed for given actor arguments. Viewing a lock as a predicate, opening a lock corresponds to assigning it the value true. Below we define two locks, one for modelling the ownership of files, and another for the organisational hierarchy among users.

---

```
lock Owns(File, User);
lock ActsFor(User, User);
policy p3 = { File f: Owns(f, alice) } ;
policy p4 = { (User u) File f: Owns(f, u), ActsFor(u, alice) };
```

---

The policy `p3` expresses that information can flow to any file owned by Alice, while the policy `p4` states that `u` ranges over users, and that information having this policy may flow to any file `f` for which `f` is owned by some `u` such that `u` acts for `alice`. Note that variables that are mentioned in the head of a clause are universally quantified, whereas those only appearing in the body are existentially quantified.

A lock can be declared to have *properties*. A property specifies conditions under which some locks are *implicitly* open. For example, we might want to express that the acts-for relation is transitive and reflexive.

This requirement can be stated at the point of declaration by replacing line 2 in the above with:

---

```
lock ActsFor(User, User)
{ (User x) ActsFor(x, x) :
; (User x y z) ActsFor(x, y) : ActsFor(x, z), ActsFor(z, y) }
```

---

Transitivity and reflexivity properties (as well as symmetry) are a common pattern, so Paragon provides syntactic sugar for these:

---

```
reflexive transitive lock ActsFor(User, User);
```

---

The Paragon Policy Language is an object-based generalisation of Paralocks [5,29], and is described in more detail in the technical report [1].

## 2.2 Information-flow policies in Paragon

The various flows of information that need to be controlled in Paragon are essentially the same as the ones occurring in Java. As is common in information-flow analysis we make a distinction between *direct* and *indirect* information flows.

**Direct flows** The typical direct flow is an assignment, where information flows directly from one location to another. Direct flows also happen at method calls (arguments flow to parameters), returns (the return value flows to the caller) and exception throws (an exception value flows to its enclosing catch clause).

Continuing in the style of the examples from the previous section, let  $x$  be a variable with the policy `{File f:Owns(f, alice)}` and  $y$  a variable with the policy `{f1:}` in the assignment `y = x;`. Whether or not the assignment will be flagged as an error by the Paragon compiler depends on the lock state in which the direct flow occurs.

If the `Owns(f1, alice)` lock is statically determined to be closed the compiler raises an error, since the information stored in  $x$ , according to its policy, should only flow into file `f1` when the file is owned by Alice, whereas the information in  $y$  can always flow to `f1`. In other words, the assignment has insecure information flow because it moves information to a place where it becomes visible to more actors than its policy declares. If, however, the lock is determined to be open, i.e. declaring that `alice` owns `f1`, the assignment occurs in a state where `f1` can already read the information in  $x$ , and so the program compiles successfully.

**Indirect flows** An indirect flow is one where the effect of evaluating one term reveals information about a completely different term that was evaluated previously. The typical indirect flow is a side-effect happening in a branch that reveals which branch was chosen, which in turn reveals the value of the conditional expression that was branched on. Indirect flows also arise from other control flow constructions (including loops and exceptions), and field updates or instance method-calls (possibly revealing the object they belonged to).

Due to the delayed nature of these information flows, the lock state in effect at the time of the indirect flow might be different to that in effect at the point at which it is revealed. Therefore, indirect flows are handled conservatively, by not allowing the lock state to affect which of these flows are considered secure.

### 2.3 Policy annotations

When integrating the policy language into Java, the two core design issues are (i) how policies are to be associated to data, and (ii) how the lock state is specified, updated, and queried.

**Policies as modifiers** In Paragon every information container (field, variable, parameter, lock) has a policy detailing how the information contained therein may be used. Every expression has an effective policy which is (an upper bound on) the conjunction of all policies on all containers whose contents affect its resulting value – we refer to this as the expression’s *read effect*.

Paragon separates policies from base (Java) types syntactically by having all policy annotations as modifiers. A modifier `?pol` denotes a policy on an information container, and the read effect of accessing that container. When used on a method we refer to it as the *return policy*, as it is the read effect on the value returned by the method. Using modifiers for policies allows for a clean separation of concerns, allowing us to analyse base types and policies separately.

Similarly, every expression (and statement) has a *write effect*, which is (a lower bound on) the disjunction of all policies on all containers whose contents are modified by the expression. Write effects allow us to control implicit information flows, by limiting the contexts in which expressions with side-effects may occur. A modifier `!pol` denotes a write effect, and is used to annotate methods.

Policy modifiers are also placed on exceptions declared to be thrown by a method. A read effect modifier on an exception denotes the read effect of inspecting the thrown and caught exception object. More interesting is the write effect modifier, which serves two purposes in relation to indirect flows. First, it restricts the contexts in which the exception may be thrown within the method. Second, it imposes a restriction of its own on all subsequent side-effects until the point where the exception has been caught and handled. Together, these two restrictions ensure that no information leaks can occur by observing whether or not an exception has been thrown.

All exceptions in Paragon must be *checked*, i.e. declared to be thrown by methods that may terminate with such exceptions. This implies the need for analyses that can rule out the possibility of exceptions, in particular for null pointers, to avoid a massive blow-up in the number of potential exceptions that must be declared. Paragon adds the modifier `nonnull` for fields, variables and method parameters that may never be null, to aid the null-pointer analysis.

To reduce the burden on the programmer to put in policy annotations, Paragon attempts either to infer, or to supply clever defaults for, policies on variables, fields and methods. We omit the details of policy defaulting, and discuss the inference mechanism in § 4.

**Lock state analysis** Manipulation of the lock state is done programmatically through the use of the Paragon-specific statements `open` and `close`. The compiler performs a *lock state analysis* which conservatively approximates the set of locks guaranteed to be open at any given program point.

In cases where we cannot know statically that a lock is open, we allow runtime *lock queries* to guide the analysis: A lock can be used syntactically as an expression of type

`boolean`, with the value `true` if the lock is currently open. If a lock query appears as the condition to e.g. an `if` statement, the analysis can include the knowledge of the lock's status when checking the respective branches.

To facilitate modularity, Paragon introduces three modifiers, used on methods and constructors, to specify their interaction with the lock state:

- `+locks` says that the annotated method *will* open the specified lock(s), for every execution in which the method returns normally. We call this the *opens* modifier.
- `-locks`, dubbed the *closes* modifier, says that the method *may* close the specified lock(s), for *some* execution.
- `~locks`, the *expects* modifier, says that the specified lock(s) must be open whenever the method is called.

The *opens* and *closes* modifiers are also used to annotate each exception type thrown by a method, to signal to the analysis what changes to the lock state can be assumed if the method terminates by throwing an exception of that type.

As a middleground between private and public locks, Paragon introduces the modifier `readonly` for locks, indicating that outside the class the lock can be queried, but not opened or closed.

### 3 Brief Example

To illustrate the language features of Paragon we present the scenario of a simple social network. In the network, users can befriend each other and share messages in the form of posts that can be read by their friends. The scenario contains two information flow policies that we want Paragon to enforce.

First, posts can only be read by a direct friend of the poster or, if the poster so indicates, by friends of friends of the poster. A user can decide, per post, whether it should be shared with friends-of-friends or not. Paragon should thus enforce that the network properly checks the friendship relations before allowing a user to read a post.

Second, to prevent injection or scripting attacks, a message should be properly *sanitized* before it is stored in the network. That is, we want to enforce the policy that all posted messages first pass through a sanitising function.

The Paragon implementation of this network is shown in Figure 1. Some policy annotations are omitted in the implementation, since Paragon provides default policies in these cases. For example, all fields that do not specify a read effect automatically get the least restrictive policy `{Object x:}`.

To establish the first policy we define the `Friend` lock to model friendships. Similarly we create a lock `FoFriend` to model friend-of-friend relations. Since the `User` class does not explicitly open or close this lock and exports it as `readonly` we know that it models a purely derived property of the `Friend` lock, and thus one that will evolve correctly as the friendship status changes over time.

With the locks in place we can now create the desired policy as `messagePol`, which we use for the read-effect on a post's content. We assume that the correct `Friend` instances are opened elsewhere in the program. Turning sharing with friends-of-friends on per post is handled in the `post` method by opening the `ShareFoF` lock for that post.

```

1 public class User {
2     public reflexive symmetric lock Friend(User, User);
3     public readonly lock FoFriend(User, User)
4         { (User x y z) FoFriend(x,y) : Friend(x,z), Friend(z,y) };
5     public void receive(?{this:} String data) {
6         ... // User receives provided data
7     }
8 }
9
10
11 public class Post {
12     public lock ShareFoF(Post);
13     public final User poster;
14     public static final policy messagePol =
15         { User x : User.Friend(x, poster)
16           ; User x : User.FoFriend(x, poster), ShareFoF(this) };
17     public final ?messagePol String message;
18     public Post(?{Object x:} User p, ?messagePol String m) {
19         this.poster = p;
20         this.message = m;
21     }
22 }
23
24
25 public class Sanitiser {
26     private lock Sanitised;
27     public static final policy unsanitised = {Object x : Sanitised};
28     public static ?{Object x:} String sanitise (?un sanitised String s) {
29         open Sanitised {
30             return /* Sanitised string */ ;
31         }
32     }
33 }
34
35
36 public class Network {
37     private static Post[] posts = new Post[10]; // Shifting list of posts
38     private static int index = 0; // Where to place the next post
39
40     !{Object x:} static void post ( ?{Object x:} User user
41                                   , ?Sanitiser.unsanitised String message
42                                   , ?{Object x:} boolean shareFoF ) {
43         String sM = Sanitiser.sanitise(message);
44         Post p = new Post(user, sM);
45         if (shareFoF)
46             open Post.ShareFoF(p);
47         posts[index] = p;
48         index = (index + 1) % posts.length; // Next time overwrite oldest post
49     }
50
51     static void read(?{Object x:} User user, ?{Object x:} int i) {
52         ?{user:} String res = null;
53         Post p = posts[i];
54         if (p != null) {
55             if (User.Friend(user, p.poster))
56                 res = p.message;
57             if (Post.ShareFoF(p))
58                 if (User.FoFriend(user, p.poster))
59                     res = p.message;
60         }
61         user.receive(res);
62     }
63 }

```

Fig. 1. A simple social network application written in Paragon.



As an effect of calling this method the array `posts` is changed (among others). Any observer that may notice this change (i.e. of level `{Object x:}` and above) may thus notice that this method has been called. To prevent this method from being called in a context where these side-effects result in implicit flows, we are required to annotate the method with the corresponding write effect.

The user's `receive` method lets the user read the provided information, therefore arguments to this method should be allowed to flow to that user. All combined, we get Paragon's enforcement ensuring that the policy-relevant state is properly checked before sharing a post with another user.

Leveraging on Java's encapsulation mechanism we are able to provide the ingredients for the sanitisation policy entirely as a separate library. The lock `Sanitised` is private to the class, meaning that no code outside the class is able to open, close or even mention the lock. Therefore, any data labelled with the `unsanitised` policy cannot lose its `Sanitised` constraint, other than by actually sanitising the data by calling the exported `sanitise` method. With this library we can thus easily enforce our second policy by labelling each newly incoming message as `unsanitised`.

The example demonstrates the three different generations of information-flow control policies and how Paragon models them.

As per traditional non-interference, some flows are never allowed in the network. For example, Paragon enforces that a posted message can only flow to users in the network, and not to any other channel. We see an example of the exceptional information declassification pattern in the sanitiser library: the `sanitise` function serves as a declassifier, deliberately allowing the provided argument to flow to more actors. Finally, the locks used to model friendships exemplify third-generation information-flow policies. There is no explicit declassification of information, rather flows are allowed or not depending on the state of the system – in this case the state of the social network.

## 4 Enforcement of Paragon Policies

Enforcement of information flow policies in Paragon is no small task, and presenting the information flow type system in its entirety is beyond the scope of this paper. Instead, we sketch a high-level overview of the most important analyses involved, presented as a sequence of *phases*, and focus on the last phase in which information flows are tracked.

**Phase 1: Type checking** The first phase roughly corresponds to ordinary Java type checking, albeit with some additions for Paragon-specific constructs. Particularly, we must assure that arguments to locks are type correct, and that policy expressions used in modifiers are indeed of type `policy`. This phase also checks that potential (runtime) exceptions are properly handled.

**Phase 2: Policy type evaluation** Locks, policies, and object references all play a dual role, both as type-level and value-level entities. In this phase the values of each of these entities are statically approximated. For *locks* we ensure that, whenever a lock is queried, the information in the query is propagated to the respective branch (or loop body).

For fields and variables holding *actors*, i.e. object references, approximating their runtime values means performing an alias analysis. Our present analysis is simple but has performed well enough in practice. However, work is in progress to improve its precision by adapting the work by Whaley and Rinard [30].

Since policies can be used as values at runtime, and dynamically hoisted to the type level, our analysis needs to approximate policies as *singleton types*, similar to the analysis of actors. For each field or variable storing a *policy*, and for each policy expression appearing in a modifier, we thus calculate upper and lower bounds on the policy held by that variable at runtime.

Further, we need ways to relate policies that are not known statically to other (static or dynamic) policies, to improve precision. Similar to runtime lock queries, we thus let our policy analysis be guided by inequality constraints between policies appearing as the condition in `if` statements and conditional `?:` expressions. This problem has been studied by Zheng and Myers in the context of Jif [32], and our solution closely follows theirs.

**Phase 3: Lock state evaluation** The next sub-phase approximates the *lock state*, i.e. it calculates the set of locks which we can statically know to be open, at each program point. This amounts to a dataflow analysis over the control flow graph, to properly capture the influence of method calls and exceptions, and to handle loops. Each program point where a direct flow takes place is annotated with the lock state in effect at that point.

**Phase 4: Policy constraint generation** The constraint generation phase will result in a set of constraints on the form  $p \sqsubseteq_{LS} q$  where  $p$  and  $q$  are policy expressions and  $LS$  is the lock state (calculated in Phase 3) at the program point where the constraint was generated (omitted if empty). As argued in §2.2 the lock state is only taken into account for direct flows. Policy expressions possibly contain meta-variables, for which the constraint solving phase then solves.

**Phase 5: Policy constraint solving** The last phase solves the generated constraints, on a per-method basis. A solution to a set of constraints is an assignment of policies to constraint variables that satisfies all the policy comparison constraints. The algorithm needs only determine whether there exists a solution, and does not need to actually produce one. The constraint solver is based on the algorithm presented by Rehof and Mogensen [21].

#### 4.1 Paragon implementation

We have implemented Paragon in a compiler that performs type checking for policies, and compiles policy-compliant programs into vanilla Java code. Once we know that a given program satisfies the intended information flow properties, we can safely remove all Paragon-specific type-level aspects of policies, locks and actors.

We must still retain the runtime aspects, such as querying the lock state and performing inequality comparisons between policies. The Paragon runtime library provides Java implementations for locks and properties, including operations for opening, closing and

querying locks to which the Paragon `open`, `close` and query statements are compiled. Similarly, the library provides Java implementations for policies and operations for performing runtime inequality comparisons between them.

**Compiler statistics** Our Paragon compiler is written in Haskell and comprises roughly 16k lines of code, including comments. Approximately half of that code is due to our policy type checker, and only a small fraction, just over 600 lines of code, deals with generation of Java code and the Paragon interface files needed for modular compilation. On top of that, some 1500 lines of Java code are written for our runtime representations of Paragon entites. The compiler can be downloaded from our Paragon website [1], or from the central Haskell “hackage” repository using the command `cabal install paragon`.

**Runtime overhead** Supporting lock queries and policy comparisons at runtime yields a negligible overhead on Paragon programs. Most of the additional generated code handles the initialisation of policies and locks upon class or object instantiation, as well as the opening and closing of locks, which should not give any significant performance penalty. More involved are the lock queries and policy comparisons themselves since they resemble essentially Datalog program evaluation and respectively containment [29]. However, our experience shows that clause bodies consist of just a few atoms, and have yet to find an example involving locks with arity higher than two, so in practice we posit that this overhead is negligible as well.

## 5 Case Studies

We put the compiler to the test with two case studies, both based on applications written in the Jif programming language, to which we further relate in §6.

**Mental Poker** In [3], a non-trivial cryptographic protocol for playing online poker without a trusted third party is implemented in Jif. During the distribution of the cards, players communicate cards encrypted with a per-player, per-game symmetric key. That is, the receiving player cannot decrypt the received card. At the end of the game the players reveal their symmetric key such that the other player may verify the outcome of the card distribution. For the purpose of non-repudiation each player signs outgoing messages with her private key.

From an information-flow perspective we desire an implementation of this protocol to satisfy various policies. The public key of a player is visible to everyone, as it is used to verify the player’s signatures, but the private key should never leave the player’s client. The cards to be communicated should not be sent before they are encrypted with the symmetric key and then signed. The symmetric key should remain confidential to the player until the end of the game.

The value of the symmetric key leaks partially when performing encryption. In our Paragon implementation (6.5k lines) this leak is controlled by a lock private to the class performing the encryption, similar to the approach taken in the sanitiser class from the example in §3. That is, the class ensures that only the result of the operation is released

and not the value of the key involved. The symmetric key is protected with a policy guarded by this private lock. A similar approach is used to protect the private key, to only reveal the outcome of the signing operation in which it was involved. The cards to be encrypted are protected with both the private locks of the encryption and the signing operation, indicating that they have to go through both declassifiers before they can be sent to the other player. The symmetric key is also allowed to be released when the `EndGame` lock is open. That is, this lock is used to represent a policy-relevant state of the application.

By contrast, Jif uses owner-based policies. The Jif policies here can simply state whether the data is owned by a given player or not, and cannot, in an obvious way, express anything beyond that. The fact that a Jif program has access to exactly one declassification mechanism prevents it from distinguishing or controlling different forms of declassification. In this case study it cannot make a distinction between declassifications that are allowed due to encryption, and those due to signing. In addition, Jif does not provide a means to write temporal policies and needs to rely on programming patterns to prevent declassifications occurring in a state where they are not supposed to be allowed.

**JPMail** The second case study implements a functional e-mail client based on JPMail [11]. In JPMail the user can specify a mail-policy file, partly dictating the information-flow policies that the mail client has to enforce.

JPMail ensures that an e-mail is only sent if its body has been encrypted under an algorithm that is trusted by the receiver of the e-mail. Which encryption algorithms are trusted by what JPMail users is specified in the mail-policy file. In addition JPMail needs to enforce more static policies, e.g. preventing the login credentials from flowing anywhere else than to the e-mail server.

In the Paragon implementation (2.6k lines) these latter, static policies are easily modelled as specifying only the e-mail server as a receiving actor. The partly dynamic policy on the e-mail body is represented by a set of clauses of the form:

---

```
(User u) server: Receiver(mail, u), AESEncrypted(mail), TrustsAES(u)
```

---

That is, the e-mail can be sent to the mail server only if it has been encrypted under AES and the receiver of that e-mail trusts AES encryption. The `TrustsAES` and similar locks representing the user-specific policies are opened after parsing the mail-policy file, during initialisation of the client. The `Receiver` lock is opened based on the To-field information, and the `AESEncrypted` lock is encapsulated analogous to the encryption / signing locks of the previous case study.

The issues for the Jif implementation in the mental poker case study show up in the JPMail example [11] as well. Moreover, stateful policies are central to this example and are challenging to model in Jif; Hicks *et al's* solution involves generating the policy part of the Jif source code from the mail-policy file, hardcoding the user-specific policies in the client. This implies that if a mail-policy file changes, the only way for the Jif solution to handle it is by recompilation of the code. By contrast, Paragon handles policy change mechanisms naturally (by opening and closing locks) without stopping the code or recompiling.

## 6 Related Work

In this section we consider the related work on languages and language support for expressive information flow policies. We focus on actual systems rather than theoretical studies on policy mechanisms and formalisms. We note, however, that there are several policy languages in the access control and authorisation area which have some superficial similarity with the Paragon Policy Language, since they are based on datalog-like clauses to express properties like delegation and roles, see e.g. [4,9,13,14]. Key differences are (i) the information flow semantics that lies at the heart of Paragon, and (ii) the fact that the principal operation in Paragon is comparison and combination of policies, whereas in the aforementioned works the only operation of interest is querying of rules.

**Languages with explicit information-flow tracking** Two “real-sized” languages stand out as providing information-flow primitives as types, namely FlowCAML and Jif – as discussed in the introduction.

Comparing Paragon to Jif is inevitable, being at the same time a competitor and a source of inspiration. Due to the unique position Jif has enjoyed in the domain of information flow research over many years, much research has been done using Jif for context and examples. It is thus natural to ask how research done on or with Jif can carry over to Paragon.

The main advantage of Paragon over Jif is undoubtedly the flexibility of the concept of locks, including their stateful nature. Where Jif has a single declassify construct, Paragon can provide different declassifying methods to work on different data, as needed by the domain at hand, and relate that declassification to the state of the program. Jif rigidly builds in some stateful aspects in the form of authority and delegations, which in Paragon would just be special cases of working with locks.

In many aspects, our work on Paragon has greatly benefitted from Jif’s trailblazing, as well as research done in the context of Jif. Policy defaulting mechanisms, handling of runtime policies, and having all exceptions checked, are all features where we have been able to adopt Jif’s solution directly.

In separate work, as of yet unpublished, we have conducted a complete and in-depth comparison between the two languages and all their features, including a Paragon library that gives a complete implementation of Jif, but the full details of that comparison are out of scope for this paper.

**Compilers performing IF tracking** Information flow tracking can be performed in a language which has no inherent security policies, lattice-based or otherwise. In such a setting one tracks the way that information flows from e.g. method parameters to outputs. Examples of tools performing such analysis are the Spark Examiner operating over a safety-critical subset of Ada [6], and Hammer and Snelting [10] explain how state-of-the-art program slicing methods can support a more accurate analysis of such information flows in Java (e.g. both flow sensitive and object sensitive).

**Dynamic Information Flow Tracking with Expressive Policies** Runtime information flow tracking systems have experienced a recent surge of interest. The most relevant examples from the perspective of the present paper are those which perform full

information flow tracking (rather than the semantically incomplete “taint analysis”), and employ expressive policies. The first example is Stefan *et al*’s embedding of information flow in Haskell [25]. In principle one could also use Paralocks in a dynamic context, and we are currently investigating a stateful extension of their LIO framework which could be instantiated with the generalised Paralocks described in this work.

Yang *et al*’s *Jeeves* language [31] focusses on confidentiality properties of data expressed as context-dependent data visibility policies. The *Jeeves* approach is noteworthy in its novel implementation techniques and greater emphasis on the separation of policy and code.

**Encoding Information Flow Policies with Existing Type Systems** With suitably expressive type systems and abstraction mechanisms, static information flow constraints can be expressed via a library [15,16,22].

A number of recent expressive languages are aimed at expressing a variety of rich security policies, but do not have information flow control as a primitive notion (as Paragon or Jif) [12,28]. F\* [27], a full-fledged implementation of a dependently typed programming ML-style programming language, is perhaps the most successful in this class, with a large number of examples showing how security properties can be encoded and verified by typing.

**Typestate** The way that Paragon tracks locks is related to the concept of *typestate* [26]. Typestate acknowledges that the runtime state of e.g. an object often determines which methods are safe to call. For example, for a Java `File` object, the method `read()` can only be called if the file has first been opened with the `open()` method. Systems with typestate, such as *Plaid* [2], support formal specification of typestate properties, and enforce that programs correctly follow the specifications. In Paragon, typestate properties can be specified through the use of lock state modifiers. Paragon cannot express features that depend on *Plaid*’s first-class states, e.g. “an array of open files”, but can otherwise express solutions to their motivating examples.

## 7 Conclusions and Further Work

It is our expectation that one day programming languages with built-in support for expressing and enforcing information-flow concerns become widely deployed. Paragon’s strong integration with Java and its relatively natural yet expressive policy specification language lowers the threshold for adopting information-flow aware programming outside the research community. Still, much work is left to be done before Paragon can become a serious competitor to existing programming.

One notable direction for future work in the Paragon language is concurrency support. This direction requires both theoretical and practical work, in particular if declassification mechanisms are shared among threads. Another planned direction is to present a more substantial formalisation of Paragon’s type system, including a proof of soundness with respect to information flow security.

**Acknowledgments** Thanks to the reviewers and all members of the ProSec group at Chalmers for improving this paper with every iteration. We thank our colleagues Wolfgang Ahrendt, Pablo Buiras, Filippo Del Tedesco, Willard Rafnsson and Alejandro Russo for their valuable comments and feedback. Furthermore we would like to thank Jens Lideström, Shayan Najd Javadipour, Javed Nazir and Yannick Zakowski for assisting in the implementation of the Paragon compiler. This work is partly funded by the Swedish agencies SSF and VR, and Websand EC FP7-ICT-STREP.

## References

1. Paragon website. <http://www.cse.chalmers.se/research/group/paragon>, July 2013.
2. J. Aldrich, J. Sunshine, D. Saini, and Z. Sparks. Tystate-oriented programming. In *OOP-SLA Companion*, pages 1015–1022, 2009.
3. A. Askarov and A. Sabelfeld. Security-typed languages for implementation of cryptographic protocols: A case study. In *Proc. European Symp. on Research in Computer Security*, volume 3679 of *LNCS*. Springer-Verlag, 2005.
4. M. Y. Becker, C. Fournet, and A. D. Gordon. Design and semantics of a decentralized authorization language. In *Proc. IEEE Computer Security Foundations Symposium*, pages 3–15. IEEE Computer Society, 2007.
5. N. Broberg and D. Sands. Paralocks – role-based information flow control and beyond. In *POPL’10, Proceedings of the 37th Annual ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, 2010.
6. R. Chapman and A. Hilton. Enforcing security and safety models with an information flow analysis tool. *ACM SIGAda Ada Letters*, 24(4):39–46, 2004.
7. D. E. Denning. A lattice model of secure information flow. *Comm. of the ACM*, 19(5):236–243, May 1976.
8. D. E. Denning and P. J. Denning. Certification of programs for secure information flow. *Comm. of the ACM*, 20(7):504–513, July 1977.
9. D. J. Dougherty, K. Fisler, and S. Krishnamurthi. Specifying and reasoning about dynamic access-control policies. In *IJCAR*, volume 4130 of *LNCS*. Springer, 2006.
10. C. Hammer and G. Snelling. Flow-sensitive, context-sensitive, and object-sensitive information flow control based on program dependence graphs. *International Journal of Information Security*, 8(6):399–422, 2009.
11. B. Hicks, K. Ahmadizadeh, and P. D. McDaniel. From languages to systems: Understanding practical application development in security-typed languages. In *ACSAC*. IEEE Computer Society, 2006.
12. L. Jia and S. Zdancewic. Encoding information flow in aura. In *Proceedings of the ACM SIGPLAN Fourth Workshop on Programming Languages and Analysis for Security*, 2009.
13. T. Jim. SD3: A trust management system with certified evaluation. In *Proc. IEEE Symp. on Security and Privacy*, 2001.
14. N. Li, J.C. Mitchell, and W.H. Winsborough. Design of a role-based trust-management framework. In *IEEE Symposium on Security and Privacy*, pages 114–130, 2002.
15. P. Li and S. Zdancewic. Arrows for secure information flow. *Theor. Comput. Sci*, 411(19), 2010.
16. J. Morgenstern and D. R. Licata. Security-typed programming within dependently-typed programming. In *Proceedings of the 15th ACM SIGPLAN international conference on Functional Programming*, 2010.

17. A. C. Myers. JFlow: Practical mostly-static information flow control. In *Proc. ACM Symp. on Principles of Programming Languages*, pages 228–241, January 1999.
18. A. C. Myers and B. Liskov. A decentralized model for information flow control. In *Proc. ACM Symp. on Operating System Principles*, pages 129–142, October 1997.
19. A. C. Myers and B. Liskov. Protecting privacy using the decentralized label model. *ACM Transactions on Software Engineering and Methodology*, 9(4):410–442, 2000.
20. A. C. Myers, L. Zheng, S. Zdancewic, S. Chong, and N. Nystrom. Jif: Java information flow. Software release. Located at <http://www.cs.cornell.edu/jif>, 2001–2013.
21. J. Rehof and T. Mogensen. Tractable constraints in finite semilattices. In Radhia Cousot and David A. Schmidt, editors, *Static Analysis*, volume 1145 of *Lecture Notes in Computer Science*, pages 285–300. Springer Berlin Heidelberg, 1996.
22. A. Russo, K. Claessen, and J. Hughes. A library for light-weight information-flow security in haskell. In *Proceedings of the 1st ACM SIGPLAN Symposium on Haskell*, 2008.
23. A. Sabelfeld and D. Sands. Declassification: Dimensions and principles. *Journal of Computer Security*, 15(5):517–548, 2009.
24. V. Simonet. The Flow Caml system. Software release. Located at <http://crystal.inria.fr/~simonet/soft/flowcaml>, July 2003.
25. D. Stefan, A. Russo, J. C. Mitchell, and David Mazières. Flexible dynamic information flow control in Haskell. In *Proceedings of the 4th ACM symposium on Haskell*, 2011.
26. R. E. Strom and S. Yemini. Typestate: A programming language concept for enhancing software reliability. *IEEE Trans. Software Eng.*, 12(1):157–171, 1986.
27. N. Swamy, J. Chen, C. Fournet, P. Strub, K. Bharagavan, and J. Yang. Secure distributed programming with value-dependent types. In *The 16th ACM SIGPLAN International Conference on Functional Programming*, 2011.
28. N. Swamy, B.J. Corcoran, and M. Hicks. Fable: A language for enforcing user-defined security policies. In *Proc. IEEE Symp. on Security and Privacy*, pages 369–383, 2008.
29. B. van Delft, N. Broberg, and D. Sands. A datalog semantics for paralocks. In Audun Jøsang, Pierangela Samarati, and Marinella Petrocchi, editors, *Proceedings of the 8th International Workshop on Security and Trust Management (STM)*, pages 305–320. Springer, 2012.
30. J. Whaley and M. Rinard. Compositional pointer and escape analysis for Java programs. In *Proceedings of the 14th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications, OOPSLA '99*, pages 187–206. ACM, 1999.
31. J. Yang, K. Yessenov, and A. Solar-Lezama. A language for automatically enforcing privacy policies. In *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, 2012.
32. L. Zheng and A. C. Myers. Dynamic security labels and static information flow control. *International Journal of Information Security*, 6, 2007.