

Reasoning about Risk and Trust in an Open World

Sophia Drossopoulou¹, James Noble², Mark S. Miller³, and Toby Murray⁴

¹ Imperial College London `scd@doc.ic.ac.uk`

² Victoria University of Wellington `kjx@ecs.vuw.ac.nz`

³ Google, Inc. `erights@google.com`

⁴ NICTA and UNSW `toby.murray@nicta.com.au`

Abstract. Contemporary open systems use components developed by different parties, linked together dynamically in unforeseen constellations. Code needs to live up to strict *security requirements*, and ensure the correct functioning of its objects even when they collaborate with external, potentially malicious, objects. In this paper we propose special specification predicates that model risk and trust in open systems. We specify Miller, Van Cutsem, and Tulloh’s escrow exchange example, and discuss the meaning of such a specification.

We propose a novel Hoare logic, based on four-tuples, including an invariant describing properties preserved by the execution of a statement as well as a post-condition describing the state after execution. We model specification and programming languages based on the Hoare logic, prove soundness, and prove the key steps of the Escrow protocol.

1 Introduction

Traditional systems designs are based on a closed world assumption: drawing a sharp border around a system where the system as a whole can be trusted because every component inside the border is known to be trustworthy, or is *confined* [25] by trustworthy mechanisms. Open systems, on the other hand, have an open world assumption: they must interact with a wide range of component objects with different levels of mutual trust (or distrust) — and whose configuration dynamically changes. Given a method request $x.m(y)$, what can we conclude about the behaviour of this request if we know nothing about the receiver x ?

In this paper, we lay the foundations for reasoning about the correctness of these kinds of open systems. Building on the object-capability security model [30] we introduce a first-class notion of *trust*, where we write “ \circ obeys S_{spec} ” to mean that object \circ can be trusted to obey specification S_{spec} . The **obeys** predicate is hypothetical: there is no central authority that can assign trustworthiness (or not) to objects; there is no trust bit that we can test. Rather, “ \circ obeys S_{spec} ” is an assumption that may or may not be true, and we will use that assumption to reason by cases. If we trust an object, we can use the object’s specification S_{spec} to determine the results of a method call on that object. If we don’t trust the object, we determine the maximum amount of damage the call could do: the *risk* of calling a method that turns out not to meet its specification.

Risks are effects against which we want to guard our objects: bounds on the potential damage caused by calls to untrusted objects. The key to delineating risks are two further hypothetical predicates: *MayAccess* and *MayAffect*. We write $MayAffect(\circ, p)$

to mean that it is possible that some method invocation on o would affect the object or property p , and $MayAccess(o, p)$ to mean that it is possible that the code in object o could potentially gain a capability to access to p . This first-class notion of risk complements our first-class notion of trust: $MayAccess$ and $MayAffect$ let us reason about the potential damage to a system when one or more objects are not trustworthy.

Our complementary notions of trust and risk are set within a very flexible specification language, and supported by an Hoare logic, enabling us to reason whether or not objects can be trusted to meet their specifications, providing sufficient security guarantees while mitigating any risks. Building on our earlier work [15, 17] we formalise and prove correctness, trust, and risk for the Escrow Exchange [31] a trusted third party that manages exchanges of different goods (e.g. money and shares) between untrusting counterparties [44]. We were surprised to find that the specification for the Escrow Exchange is weaker than originally anticipated in two significant aspects: the escrow *cannot guarantee* that a reported successful transaction implies a) that the participants were trustworthy, nor that b) the participants are exposed to no risk by an untrustworthy participant (but we were able to characterize the risk to which participants are exposed). We were even more surprised to realize that it is *impossible to write* an escrow which would give guarantees a) and b) — all the more striking given that a co-author is one of the original developers of the escrow example.

Common approaches to reasoning about programs cannot deal with the escrow exchange example. Most program specification and verification methods have an implicit underlying assumption that components are meant to be trustworthy (i.e. meet their specifications). Our approach first makes that assumption explicit (as **obeys**), lets us reason hypothetically and conditionally about those trust assumptions, even in cases where those assumptions fail (by quantifying risk via $MayAccess$ and $MayAffect$).

Paper Organization Section 2 introduces the Escrow Exchange example, shows why a traditional specification is not descriptive enough and why a naive implementation is not robust enough. Section 3 introduces our constructs and Hoare logic for modelling trust and risk, which we apply to a revised implementation of the Escrow to reason formally about its correctness. Section 4 discusses related work, and Section 5 concludes.

Disclaimers Throughout this paper, we make the simplifying assumptions that no two different arguments to methods are aliases, that the program is executed sequentially, that we can quantify over the entire heap, that objects do not breach their own encapsulation or throw exceptions, that machines on open networks are not mutually suspicious, and that any underlying network is error-free. This allows us to keep the specifications short, and to concentrate on the questions of risk and trust. Aliasing, concurrency, quantification, confinement, network errors, and exceptions can be dealt with using known techniques, but doing so would not shed further light on the questions we address.

Contribution This paper extends earlier informal work presented at the PLAS workshop [17]. Here we contribute the full formal foundations of the system, defining **obeys**, $MayAccess$, and $MayAffect$ in the context of the *Focal* and *Chainmail* languages (details in the full technical report [18]). We present a novel Hoare logic based on four-tuples to specify properties preserved during execution: this allows us to model trust and delineate risk even when a method’s receiver is unknown. We use our logic to prove formally that the key steps of the escrow example meet the specification.

2 Escrow Exchange

Figure 1 shows a first attempt to implement an escrow exchange, also shown in previous work [31, 36]. We model both money and goods by `Purses` (a resource model proposed in E [32]). The call `dst.deposit(amt, src)` will either transfer `amt` resources from the `src` purse to the `dst` purse and return `true`, or do nothing and return `false`. A new, empty purse can be created at any time by asking an existing purse to `sprout` — the new purse has a zero balance but can then be filled via `deposit`.

```
1 method deal_version1( ) {
2
3     // make temporary money Purse
4     escrowMoney = sellerMoney.sprout
5     // make temporary goods Purse
6     escrowGoods = buyerGoods.sprout
7
8     res = escrowMoney.deposit(price, buyerMoney)
9     if (!res) then
10        // insufficient money in buyerMoney
11        // or different money mints
12        { return false }
13
14    // sufficient money; same mints.
15    // price transferred to escrowMoney
16    res = escrowGoods.deposit(amt, sellerGoods)
17    if (!res) then
18        // insufficient goods in sellerGoods
19        // or different goods mints
20        { // undo the goods transaction
21            buyerMoney.deposit(price, escrowMoney)
22            return false }
23
24    // price in escrowMoney; amt in escrowGoods.
25    // now complete the transaction
26    sellerMoney.deposit(price, escrowMoney)
27    buyerGoods.deposit(amt, escrowGoods)
28    return true
29 }
```

Fig. 1. First attempt at Escrow Exchange deal method

The goal of the escrow is to exchange `amt` goods for `price` money, between the purses of a seller and buyer. To make the exchange transactional, we use two private *escrow* purses, one for on each side of the transaction (money and goods). Lines 3–6 of Figure 1 show how we first set up the escrow purses, by sprouting two new purses (`escrowMoney` and `escrowGoods`) from their respective input purses.

It is important that the escrow purses are newly created within the method, and cannot have been manipulated or retained by the buyer or seller, which is why the escrow asks `sellerMoney` to make one, and `buyerGoods` to make the other. The requirements of an open system means that the escrow method cannot have the escrow purses before the transaction, because the escrow cannot know the right kind of purses to create, and there is no central trusted authority that could provide them. Buyers and sellers cannot provide escrows purses directly, precisely because we must assume they don't trust each other: if they did, they wouldn't need to use an escrow.

Second, we attempt to escrow the buyer's money by transferring it from the `buyerMoney` purse into the new `escrowMoney` purse — line 8. If this `deposit` request returns true, then the money will have been transferred. If the deposit fails we abort the transaction. Third, we attempt to escrow the seller's goods — line 16, again by depositing them into the other escrow purse. If we are unsuccessful, we again abort the transaction, after we have returned the escrowed money to the buyer — lines 21 and 22. At this point (line 26) the `deal` method should have sole access to sufficient money and goods in the escrow purses. The method completes the transaction by transferring the escrowed money and goods into the respective destination purses — lines 26 and 27. Thanks to the escrow purses, these transfers should not fail, and indeed, if `deal_version1` is called in good faith it will carry out the transaction correctly. Unfortunately, we cannot assume good faith in a mutually untrusting open system.

2.1 The failure of `deal_version1`

The method in Figure 1 does not behave correctly in an open system. The critical problems are assumptions about trust: both the code and the specification implicitly trust the purse objects with which they interact.

Imagine if `sellerMoney` was a malicious, untrustworthy object. At line 4, the `sprout` call could itself return a malicious object, which would then be stored in `escrowMoney`. At line 8, `escrowMoney.deposit(price, buyerMoney)` would let the malicious `escrowMoney` purse steal all the money out of `buyerMoney` purse, and still return `false`. As a result, the seller would lose all their money, and receive no goods! Even if the seller was more cautious, and themselves sprouted a special temporary purse with a balance of exactly `price` to pass in as `sellerMoney`, they would still lose all this money without any recompense.

Perhaps there is something else we could do — a `trusted` method on every object, say, that returns `true` if the object is trusted, and `false` otherwise? The problem, of course, is that an object that is untrustworthy is, well, untrustworthy: we cannot expect a `trusted` method ever to return `false`. This leads to our definition of trust: trust is *hypothetical*, and in relation to some specification of expected behaviour.

2.2 Modelling Trust and Risk: `obeys`, *MayAccess* and *MayAffect*

The key claim of this paper is that, to reason about the behaviour of systems in an open world, we need specifications that let us talk about trust and risk explicitly. In the rest of this section, we informally introduce three novel specification language constructs: **obeys** to model trust, and *MayAccess* and *MayAffect* to model risk, show

how they can be used to specify the purse and escrow examples, and argue a revised `deal_version2` method can meet that specification. Section 3 formalises these ideas.

To model trust, we introduce a special predicate, **obeys**, of the form o **obeys** $Spec$ which we interpret to mean that the current object trusts o to adhere to the specification $Spec$. Because we generally can't be sure that an object — especially one supplied from elsewhere in an open system — can actually be trusted to obey a particular specification, our reasoning and specifications are hypothetical: analysing the same piece of code under different trust hypotheses — i.e. assuming that particular objects may or may not be trusted to obey particular specifications.

Thus, if object o can be trusted to obey specification $Spec$, and $Spec$ had a policy describing the behaviour of some method m , then we may expect the method call $o.m(\dots)$ to behave according to that policy — otherwise, all bets are off.

To model risk, we introduce predicates *MayAccess* and *MayAffect*, which express whether an object may read or may affect another object or property. We will write *MayAffect* (o, p) to mean that it is possible that some method invocation on o would affect the object or property p . Similarly, we will write *MayAccess* (o, p) to mean that it is possible that the code in object o could potentially gain a capability to access to p — that is, a reference to p . In practice, *MayAccess* (o, p) means that p is in the transitive closure of the points-to relation on the heap starting from o including both public and private references.

2.3 Valid Purse: Specifying Purse

Using **obeys**, *MayAccess*, and *MayAffect*, we write the `ValidPurse` specification in Figure 2 that makes trust and risk explicit.

`ValidPurse` consists of five policies. `Pol_deposit_1` and `Pol_deposit_2` taken together distinguish between a successful and an unsuccessful deposit, signalled by returning `true` or `false` respectively. In the first case, i.e. `Pol_deposit_1` where the result is `true`, argument `src` must have been a valid purse (`src obeys ValidPurse`) which can trade with the receiver, and `src` must have sufficient balance. In the second case, i.e. `Pol_deposit_2` where the result is `false`, either `src` was not a valid purse, or would not trade with the receiver, or had insufficient funds. To quote Miller et al. [32]: “A reported successful deposit can be trusted as much as one trusts the purse one is depositing into”.

The last two lines in the postcondition of `Pol_deposit_1` and `Pol_deposit_2` provide framing conditions. In the first case, the transaction will happen, but all other purses will be unmodified (line 14 in figure 2), whereas in the second case no purses will be modified (line 24 in figure 2). Another framing condition, appears on lines 15, 25 and 36 of figure 2, and requires that the methods do not leak access to any `ValidPurse` object. In other words, if *after* the method call, a pre-existing o has access to a `ValidPurse` object p , then o already had access to a p *before* the call.

`Pol_sprout` promises that the result is a trusted purse that can trade with the receiver, no other valid purse's balance is affected, and references have not been leaked.

`Pol_can_trade_constant` guarantees that whether or not two purses can trade with each other can *never* change, no matter what code is run. This is another key

```

1 specification ValidPurse {
2   field balance // Number
3
4   policy Pol_deposit_1      // 1st case:
5     amt ∈ ℕ
6     { res = this.deposit(amt, src) }
7     res → (
8       // TRUST
9       src obeyspreValidPurse ∧ CanTrade(this, src)pre
10      // FUNCTIONAL SPECIFICATION
11      ∧ 0 ≤ amt ≤ src.balancepre ∧
12      this.balance = this.balancepre + amt ∧
13         src.balance = src.balancepre - amt ∧
14      //RISK
15      ∀ p. (p obeyspreValidPurse ∧ p ∉ {this, src} →
16         p.balance = p.balancepre) ∧
17      ∀ o:preObject. ∀ p obeyspreValidPurse.
18         MayAccess(o, p) → MayAccesspre(o, p) )
19
20  policy Pol_deposit_2      // 2nd case:
21    amt ∈ ℕ
22    { res = this.deposit(amt, src) }
23    ¬res → (
24      // TRUST and FUNCTIONAL SPECIFICATION
25      ¬( src obeyspreValidPurse ∧ CanTrade(this, src)pre ∧
26         0 ≤ amt ≤ src.balancepre) ∧
27      // RISK
28      ∀ p. (p obeyspreValidPurse → p.balance = p.balancepre) ∧
29      ∀ o:preObject. ∀ p obeyspreValidPurse.
30         MayAccess(o, p) → MayAccesspre(o, p) )

```

Fig. 2. ValidPurse specification

ingredient of our approach: we can require that our code must preserve properties in the face of unknown code.

`Pol_protect_balance` guarantees that a valid purse `p`'s balance can only be changed: — $May.Affect(o, p.balance)$ — by an object `o` that may access that purse: $May.Access(o, p)$.

Finally, the abstract predicate `CanTrade` holds when two Purses can trade with each other. `CanTrade` must be reflexive, but does not require that its arguments have the same class. It guarantees that `deposit` can transfer resources from one purse to another. This could involve a clearing house, interbank exchange, or other mechanisms abstract predicates can be implemented in different ways.

The use of assertions about the pre-state in methods' postconditions increases the expressive power of our specifications. For example, consider:

(A) $amt \in \mathbb{N} \{res = this.deposit(amt, src)\} res \rightarrow src \text{ obeys}_{pre} ValidPurse$
 This allows us to deduce properties about the pre-state by observing the result of the

```

27 policy Pol_sprout
28   true
29   { res = this.sprout() }
30   // TRUST
31   res obeys ValidPurse  $\wedge$  CanTrade(this, res)pre  $\wedge$ 
32   // FUNCTIONAL SPECIFICATION
33   res.balance=0  $\wedge$ 
34   // RISK
35    $\forall p. (p \text{ obeys } \text{preValidPurse} \rightarrow$ 
36     p.balance=p.balancepre  $\wedge$  res  $\neq$  p)  $\wedge$ 
37    $\forall o:\text{preObject}. \forall p \text{ obeys } \text{preValidPurse}.$ 
38     MayAccess(o, p)  $\rightarrow$  MayAccesspre(o, p) )
39
40 policy Pol_can_trade_constant
41   true
42   { any_code }
43    $\forall \text{prsl}, \text{prs2} \text{ obeys } \text{preValidPurse}.$ 
44     CanTrade(prs1, prs2)  $\leftrightarrow$  CanTradepre(prs1, prs2)
45
46 policy Pol_protect_balance
47   // RISK
48    $\forall o, p:\text{Object}. p \text{ obeys } \text{ValidPurse} \wedge$ 
49     MayAffect(o, p.balance)  $\rightarrow$  MayAccess(o, p)
50 }
51
52 abstract predicate CanTrade(prs1, prs2) is reflexive

```

Fig. 2. ValidPurse specification (contd.)

method call. Such a specification cannot be easily translated into one which does not make use of this facility, as in:

(B) `scr obeys ValidPurse \wedge amt \in \mathbb{N} {res=this.deposit(amt, src)} res`
 (B) differs from (A) in that (B) requires us to establish that `scr obeys ValidPurse` before making the call, while (A) does not.

2.4 Establishing Mutual Trust

An escrow must build a two-way, trusted transfer by combining one-way transfers. From `Pol_deposit_1` we obtain that the call `res1=dest.deposit(amt, src)` lets us conclude `res1 \wedge dest obeys ValidPurse \rightarrow src obeys ValidPurse`. This trust is just one way: from the destination to the source purse. We can establish mutual trust between two purses by then attempting to perform a second deposit *in the reverse direction* from destination to source: `res2=src.deposit(amt, dest)` which in turn gives `res2 \wedge src obeys ValidPurse \rightarrow dest obeys ValidPurse`. Reasoning conditionally, on a path where `res1 \wedge res2` are true, we can then establish mutual trust:

$$\text{dest obeys ValidPurse} \leftrightarrow \text{src obeys ValidPurse}$$

We establish this formally in Section 3.4, having only argued informally earlier [17, 36].

As with much of our reasoning, this is both conditional and hypothetical: at a particular code point, when two `deposit` requests have succeeded (or rather, that they have both *reported* success) then we can conclude that either both are trust worthy, or both are untrustworthy: we have only *hypothetical* knowledge of the **obeys** predicate.

2.5 Escrow with Explicit Mutual Trust

```
1 method deal_version2( ) // returns Boolean
2 {
3   // setup and validate Money purses
4   escrowMoney = sellerMoney.sprout
5   res=escrowMoney.deposit(0, sellerMoney)
6   if (!res) then {return false}
7   res = buyerMoney.deposit(0, escrowMoney)
8   if (!res) then {return false}
9   res = escrowMoney.deposit(0, buyerMoney)
10  if (!res) then {return false}
11
12  // setup and validate Goods purses
13  // similar to lines 4–10 from above, but for Goods
14
15  // make the transaction
16  // as in lines 8–29 from Fig.1
17 }
```

Fig. 3. Revised `deal_version2` method

Two way deposit calls are sufficient to establish mutual trust, but come with risks. For example, as part of validating that a buyer's purse the seller's purse, we must pass the buyer's purse as an argument in a `deposit` call to the seller's purse, e.g.

```
sellerMoney.deposit(0, buyerMoney)
```

If the seller's purse is not in fact trustworthy, then it can take this opportunity to steal all the money in the buyer's purse before the transaction officially starts, even if the `amt` that is supposed to be deposited is 0.

We can minimise this risk by careful use of escrow purses. Rather than mutually validating buyers and sellers directly, we can create an escrow purse on the destination side of the transaction (the seller's money and the buyer's goods) and then mutually validate the buyer's and sellers actual purses against the escrow — resulting in a chain of mutual trust between the destination purse and the escrow purse, and the escrow purse and the source purse. This allows us to hypothesise that the source and destination purses are mutually trusting before we start on the transaction proper.

The resulting escrow method is in Figure 3. Line 4 creates a `escrowMoney` purse and then lines 5–10 hypothetically establish mutual trust between the `escrowMoney`, `sellerMoney`, and `buyerMoney` purses. The `sellerMoney` purse doesn't need to validate `escrowMoney` explicitly (`sellerMoney.deposit(0, escrowMoney)`) because

the `sprout` method specification says sprouted purses can be trusted as much as their parent purses. (Figure 4 illustrates the trust relationships.) If any of these `deposit` requests




Fig. 4. Establishing Mutual Trust. Dashed arrows show purse validation.

fail, we abort. Afterwards we do exactly the same, but for goods purses rather than money purses. Finally, we carry out the escrow exchange itself, in exactly the same manner as lines 8–29 of the first implementation in Figure 1.

2.6 Specifying the Mutual Trust Escrow

Figure 5 shows a specification for the revised escrow deal method from Figure 3. This specification uses conditional and hypothetical reasoning to distinguish four cases, based on the value of the result and the trustworthiness of the participants. We use these auxiliary definitions:

```

GoodPrs = { p | p obeyspre ValidPurse }
PPrs = { sellerMoney, sellerGoods, buyerMoney, buyerGoods }
OthrPrs = GoodPrs \ PPrs
BadPPrs = PPrs \ GoodPrs
  
```

The set `PPrs` contains the four *participant purses* passed as arguments. `BadPPrs` contains the untrustworthy participant purses. `GoodPrs` are all trustworthy purses in the system that do conform to the `ValidPurse` specification, and `OthrPrs` are the trustworthy purses that do *not* participate in this particular deal. We can now discuss the four cases of the policy:

1st case: The result is `true` and all participant purses are trustworthy. Then, the goods and money purses can trade with each other, and there was sufficient money in the buyer’s purse and sufficient goods in the seller’s purse. In this case, everything is fine, so the transfer can proceed: `price` will have been transferred from the buyer’s to the seller’s money purse, and `amt` will have been transferred from the seller’s to the buyer’s goods purse. No risk arises: no other purses’ balance will change (whether passed in to the method or not).

2nd case: The result is `false` and all participant purses are trustworthy. Then one or more of the functional correctness conditions are not satisfied: purses were unable to trade with each other, or input purses did not have sufficient balance. Again, no risk arises to any purses.

3rd case: The result is `false` and some participant purse is untrustworthy. In this case, no trustworthy purses’ balances have been changed — unless they were already accessible by an untrustworthy purse passed in to the method.

```

1 specification ValidEscrow {
2   fields sellerMoney, sellerGoods, buyerMoney, buyerGoods
3   fields price, amt // N
4
5   policy Pol_deal
6     price, amt ∈ N ∧ price, amt > 0
7     { this.deal( ) }
8     res ∧ BadPPrs=∅ → ( // 1st case:
9       CanTrade(buyerMoney, sellerMoney) ∧
10      CanTrade(buyerGoods, sellerGoods) ∧
11      buyerMoney.balance=buyerMoney.balancepre-price ∧
12      sellerMoney.balance=sellerMoney.balancepre+price ∧
13      buyerGoods.balance=buyerGoods.balancepre+amt ∧
14      sellerGoods.balance=sellerGoods.balancepre-amt ∧
15      ∀p:preOthrPrs. p.balance=p.balance.pre ∧
16      ∀o:preObject, p:preGoodPrs.
17        (MayAccess(o, p) → MayAccess(o, p)pre) )
18    ∧
19    ¬res ∧ BadPPrs=∅ → ( // 2nd case:
20      ¬( CanTrade(buyerMoney, sellerMoney) ∧
21        CanTrade(buyerGoods, sellerGoods) ∧
22        buyerMoney.balancepre ≥ price ∧
23        sellerGoods.balancepre ≥ amt ) ∧
24      ∀p:preGoodPrs. p.balance=p.balance.pre ∧
25      ∀o:preObject, p:preGoodPrs.
26        (MayAccess(o, p) → MayAccess(o, p)pre) )
27    ∧
28    ¬res ∧ BadPPrs≠∅ → ( // 3rd case:
29      ∀p:preGoodPrs. (p.balance=p.balance.pre ∨
30        ∃ bp∈BadPPrspre. MayAccess(bp, p)pre) ∧
31      ∀o:preObject, p:preGoodPrs. ( MayAccess(o, p) →
32        (MayAccess(o, p)pre ∨ ∃ b∈BadPPrspre. MayAccess(b, p)pre) ) )
33    ∧
34    res ∧ BadPPrs≠∅ → ( // 4th case:
35      buyerMoney obeys ValidPurse ↔ sellerMoney obeys ValidPurse ∧
36      buyerGoods obeys ValidPurse ↔ sellerGoods obeys ValidPurse ∧
37      ∀p:preOthrPrs. (p.balance=p.balance.pre ∨
38        ∃ bp∈BadPPrspre. MayAccess(bp, p)pre) ∧
39      ∀o:preObject, p:preGoodPrs. (MayAccess(o, p) →
40        (MayAccess(o, p)pre ∨ ∃ b∈BadPPrspre. MayAccess(b, p)pre) ) )
41  }

```

Fig. 5. ValidEscrow specification

4th case: The result is true and some participant purse is untrustworthy — actually at least two matching participant purses are untrustworthy. That is, a pair of matching purses have coöperated to suborn the escrow *and we cannot tell*. Therefore, either both money purses are untrustworthy, (as per line 35), or both goods purses are untrustwor-

thy, (as per line 36), or all four are bad. The risk is that an uninvolved trustworthy purse's balance can be changed if it was previously accessible from a bad purse. The first and second cases correspond to a traditional specification, because traditional specifications assume all objects are trustworthy. The third and fourth cases arise precisely because we are explicitly modelling the trust and risk involved in an open system.

Discussion The 3rd and 4th case represent more of a risk than we would like: ideally (as in the 2nd case) we'd hope nothing should have changed. But an escrow method cannot undo a system that is already suborned — if one of the participant purses is already benefiting from a security breach, passing that purse in to this method gives it an opportunity to exercise that breach. On the other hand, the risk is contained: this method cannot make things worse.

The 4th case does not prevent trustworthy participant purses from being modified, to cater e.g., for the possibility that the two money purses are trustworthy, while the two goods purses are not, in which case the money transaction will take place as expected, while all bets are off about the goods transaction. We can give the stronger guarantee for the 3rd case, because by the time the escrow starts making non-0 transactions it has established that the purses in each pair are both either trustworthy or both not.

Most importantly (perhaps surprisingly) the return value of the method, `res`, does *not* indicate whether the participants were trustworthy or not. A `true` result may be returned in the 1st case (all purses trustworthy) as well as the 4th (some purses are untrustworthy). The return value indicates *only* whether the escrow attempted to complete the transaction (returning `true`) or abort (returning `false`). This came as a surprise to us (and to the escrow's designers [31].) As with much of our reasoning around trust, this leads to yet more conditional reasoning, which must be interpreted hypothetically.

Nevertheless, the return value does communicate a valuable guarantee to an honest participant, whose money and goods purses are both trustworthy: If `deal` returns `true`, then the exchange has taken place. Furthermore if it returns `false`, the exchange has not taken place and with *no more* risk to the honest purses than existed before the call. The `ValidEscrow` specification also gives a guarantee to other purse objects even if they did not participate in the deal: dishonest purses can only change other purses' balances if they had prior access to those other purses.

3 A Formal Model of Trust and Risk

In this section we provide an overview of our core programming language, *Focal*, our specification language, *Chainmail*, and our Hoare logic. The Hoare logic uses *four-tuples* because it includes an invariant that must be preserved during the execution of a statement as well as a postcondition established afterwards. We also outline a key step required to prove that `deal_version2` meets the `ValidEscrow` specification: we prove that two purses can establish mutual trust, and formally delineate the risk. Many details are relegated to our technical report [18]; here we adopt its numbering for definitions.

3.1 *Focal*

We define a small object oriented language, *Focal* (Featherweight Object Capability Language, not to be confused with FOCAL [28]). *Focal* supports classes, fields and methods. (Figures 1 and 3 are effectively examples of *Focal*.) *Focal* is memory-safe: it does not allow addresses to be forged, or non-existent methods or fields to be called, read or written. *Focal* is dynamically typed: it does not check that the arguments to a method call or a field write are of the appropriate type either statically or dynamically: similar to JavaScript, Grace, E, and Dart’s unchecked mode.

Modules, M , are mappings from class identifiers to *Focal* class definitions, and from predicate identifiers to *Chainmail* assertions as described in section 3.2. The linking operator $*$ combines these definitions, provided that the modules’ mappings have separate domains, and performs no other checks. This reflects the open world setting, where objects of different provenance interoperate without a central authority. For example, taking M_p as a module implementing purses, and M_e as another module implementing the escrow, $M_p * M_e$ is defined but $M_e * M_p$ is not.

Focal enforces a weak form of privacy for fields; only the receiver may modify these fields, and anybody may read them.

The operational semantics of *Focal* takes a module M and a runtime state $\sigma = \text{frame} \times \text{heap}$ and maps statements onto a new state σ' .

Definition 6 (Shape of Execution).

$$\rightsquigarrow : \text{Module} \times \text{state} \times \text{Stmts} \longrightarrow \text{state}$$

Arising and Reachable Configurations Policies need to be satisfied in all *configurations* (pairs of states and statements) which may arise during execution of the program. For example, if a program contains a class which has field which is not exported, and where this field is initialized to 0 by the constructor, and incremented by 3 in all method calls, then in the arising configurations the value of this field is guaranteed to be a multiple of 3. Thus, through the concept of arising configurations we can ignore configurations which are guarantee not to arise.

To define arising configurations we need the concept of initial configuration, and reachability. A configuration is *reachable* from some starting configuration if it is reached during the evaluation of the starting configuration after any number of steps. We define the function $\text{Reach} : \text{Module} \times \text{state} \times \text{Stmts} \longrightarrow \mathcal{P}(\text{state} \times \text{Stmts})$ by cases on the structure of the statements. Note that $\text{Reach}(M, \sigma, \text{stmts})$ is defined, even when the execution should diverge. This is important, because it allows us to give meaning to capability policies without requiring termination.

We then define $\text{Arising}(M)$ as the set of runtime configurations which may be reached during execution of some initial context $(\sigma_0, \text{stmts}_0)$.

Definition 7 (Arising and Initial configurations).

$$\begin{aligned} \text{Init} & : \text{Module} \longrightarrow \mathcal{P}(\text{state} \times \text{Stmt}) \\ \text{Arising} & : \text{Module} \longrightarrow \mathcal{P}(\text{state} \times \text{Stmts}) \\ \text{Init}(M) & = \{ (\sigma_0, \text{new } c.m(\text{new } c')) \mid c, c' \in \text{dom}(M), \\ & \quad \text{where } \sigma_0 = ((\iota_0, \text{null}), \chi_0), \text{ and } \chi_0(\iota) = (\text{Object}, \emptyset) \} \\ \text{Arising}(M) & = \bigcup_{(\sigma, \text{stmts}) \in \text{Init}(M)} \text{Reach}(M, \sigma, \text{stmts}) \end{aligned}$$

3.2 Chainmail

Chainmail is a specification language where a specification is a conjunction of a set of named policies. (Figures 2 and 5 are examples of *Chainmail* specifications.)

Chainmail policies are based on one-state assertions (A) and two-state assertions (B). To express the state in which an expression is evaluated, we annotate it with a subscript. For example, $x > 1$ is a one-state, and $x_{pre} - x_{post} = 1$ is a two-state assertion. Validity of an assertion is defined in the usual manner, e.g. in a state σ with $\sigma(x) = 4$ we have $M, \sigma \models x > 1$. If we also have $\sigma'(x) = 3$, then we obtain $M, \sigma, \sigma' \models x_{pre} - x_{post} = 1$. *Chainmail* specifications may also express ghost information, which is not stored explicitly in the state σ but can be deduced from it — e.g. the length of a null-terminated string.

Policies can have one of the three following forms: 1) invariants of the form A , which require that A holds at all visible states of a program; or 2) $A \{ \text{code} \} B$, which require that execution of `code` in any state satisfying A will lead to a state satisfying B wrt the original state or 3) $A \{ \text{any_code} \} B$ which requires that execution of *any* code in a state satisfying A will lead to a state satisfying B wrt the original state.

Definition 12 (Policies).

$\text{Policy} ::= A \mid A \{ \text{code} \} B \mid A \{ \text{any_code} \} B$
 $\text{PolSpec} ::= \text{specification } S \{ \text{Policy}^* \}$

One-state assertions include assertions about expressions (such as $\leq, >$ *e.t.c.*) and four additional assertions: *Expr obeys SpecId* to model trust, i.e. that an object confirms to a specification; and *MayAccess* and *MayAffect* to model risk, i.e. whether one object may access another, or alter a property. These are *hypothetical*, in that they talk about the potential effects or behaviour of code: we cannot somehow evaluate their truth-value when executing the program. The fourth assertion *Expr:ClassId* simply tests class membership.

Validity of one-state assertions is expressed through the judgment $M, \sigma \models A$. The key case is that some expression obeys a specification if it satisfies that specification's policies in all reachable configurations arising from the module.

(from Definition 13):

- $M, \sigma \models e : c$ iff $\sigma(\lfloor e \rfloor_{M, \sigma}) \downarrow_1 = c$.
- $M, \sigma \models \text{MayAffect}(e, e')$ iff there exist method m , arguments \bar{a} , state σ' , identifier z , such that $M, \sigma[z \mapsto \lfloor e \rfloor_{M, \sigma}], z.m(\bar{a}) \rightsquigarrow \chi'$, and $\lfloor e' \rfloor_{M, \sigma} \neq \lfloor e' \rfloor_{M, \sigma \downarrow_1, \chi'}$.
- $M, \sigma \models \text{MayAccess}(e, e')$ iff there exist fields \bar{f} , such that $\lfloor z.\bar{f} \rfloor_{M, \sigma[z \mapsto \lfloor e \rfloor_{M, \sigma}]} = \lfloor e' \rfloor_{M, \sigma}$
- $M, \sigma \models e \text{ obeys } \text{PolSpecId}$ iff
 - $\forall (\sigma, \text{stmts}) \in \text{Arising}(M). \forall i \in \{1..n\}. \forall \sigma', \text{stmts}'.$
 - $(\sigma', \text{stmts}') \in \text{Reach}(M, \sigma, \text{stmts}) \longrightarrow M, \sigma'[z \mapsto \lfloor e \rfloor_{\sigma}] \models \text{Policy}_i[z/\text{this}]$
 where z is a fresh variable in σ' , and where we assume that *PolSpecId* was defined as $\text{specification } \text{PolSpecId} \{ \text{Policy}_1, \dots, \text{Policy}_n \}$.

Two-state assertions allow us to compare properties of two different states. Validity of two-state assertions $M, \sigma, \sigma' \models B$ is defined similarly to one-state assertions, using cases. We can now define adherence to policy, $M, \sigma \models_{\text{pol}} \text{Policy}$:

Definition 15 (Adherence to Policies).

- $M, \sigma \models_{pol} A$ iff $M, \sigma \models A$
- $M, \sigma \models_{pol} A \{code\} B$ iff
 $(M, \sigma \models A \wedge M, \sigma, code \rightsquigarrow \sigma' \longrightarrow M, \sigma, \sigma' \models B)$
- $M, \sigma \models_{pol} A \{any_code\} B$ iff
 $\forall code. ((\sigma, code) \in \mathcal{A}rising(M) \wedge M, \sigma \models A \wedge M, \sigma, code \rightsquigarrow \sigma' \longrightarrow M, \sigma, \sigma' \models B)$

3.3 Hoare Logic

The Hoare logic allows us to prove adherence to policies. In order to reflect that the code to be verified is executed in an open system, and that it calls code whose specification and trustworthiness is unknown to the code being verified, we use Hoare four-tuples rather than Hoare triples, so that not only do they guarantee a postcondition holds *after* execution of the code, but also guarantee that an invariant is preserved *during* execution of the code. These invariants are critical to modelling risk, as they let us talk about the absence of temporary but unwanted effects caused on objects during execution.

A Hoare four-tuple is either $M \vdash A \{stms\} A' \bowtie B$ (executing *stms* in any state satisfying *A* will lead to a state which satisfies *A'*) or $M \vdash A \{stms\} B' \bowtie B$ (executing *stms* in any state satisfying *A* will lead to a state where the relation of the old and new state is described by *B'*). Critically, both promise that the relation between the initial state, and *any* of the intermediate states reached by execution of *stms*, will maintain the invariant *B*. The execution of *stms* may call methods defined in *M*, and the predicates appearing in *A*, *A'*, *B'*, and *B*, may use predicates defined in *M*. When *M* is implicit from the context, we use the shorthand $\vdash A \{stms\} A' \bowtie B$.

In order to model open systems, we require that after linking *any* module with the module at hand, the policy will be satisfied. As stated in [34], “*A programmer should be able to prove that his programs have various properties and do not malfunction, solely on the basis of what he can see from his private bailiwick.*”

Definition 16 (Validity of Hoare Four-Tuples).

$$M \vdash A \{stms\} B' \bowtie B \text{ iff } \forall M', \sigma. \\ (\sigma, _) \in \mathcal{A}rising(M * M') \wedge M * M', \sigma \models A \wedge M * M', \sigma, stms \rightsquigarrow res, \sigma' \\ \longrightarrow \\ M * M', \sigma, \sigma' \models B' \wedge \forall \sigma'' \in \mathcal{R}each(M, \sigma, stms). M * M', \sigma, \sigma'' \models B$$

Figure 6 shows a selection of our Hoare rules. It starts with two familiar Hoare Logic rules: In (VARASG) and (FIELDASG) the postconditions use the previous value of the right-hand-side, and thus allow us to deduce, *e.g.* :

$$\vdash \mathbf{this}.f = 21 \{ \mathbf{this}.f = 2 * \mathbf{this}.f \} \mathbf{this}.f = 42 \bowtie true.$$

(METH-CALL-1) is also familiar, as it reasons over method calls under the assumption that the receiver **obeys** a specification *S*, and that the current state satisfies the precondition of *m* as defined in *S*.

The remaining rules are more salient.

(METH-CALL-2) expresses the basic axiom of object-capability systems that “only connectivity begets connectivity” [30]. It promises in the postcondition that the result of the method call *v* cannot expose access to any object *z* that wasn’t accessible initially

$$\begin{array}{c}
\text{(VARASG)} \qquad \qquad \qquad \text{(FIELDASG)} \\
\hline
\vdash \text{true} \{ v := a \} v = a_{pre} \bowtie \text{true} \qquad \vdash \text{true} \{ \text{this.f} := a \} \text{this.f} = a_{pre} \bowtie \text{true} \\
\\
\text{(METH-CALL-1)} \\
\frac{M(S) = \text{spec } S \{ \overline{\text{Policy}}, A \{ \text{this.m(par)} \} B, \overline{\text{Policy}} \}}{\vdash x \text{obeys } S \wedge A[x/\text{this}, y/\text{par}] \{ v := x.m(y) \} B[x/\text{this}, y/\text{par}, v/\text{res}] \bowtie \text{true}} \\
\\
\text{(METH-CALL-2)} \\
\begin{array}{l}
B \equiv \forall z :_{pre} \text{Object}. \text{MayAccess}(v, z) \rightarrow (\text{MayAccess}_{pre}(x, z) \vee \text{MayAccess}_{pre}(y, z)) \\
B' \equiv \forall z, u :_{pre} \text{Object}. (\text{MayAccess}(u, z) \rightarrow \\
\qquad (\text{MayAccess}_{pre}(u, z) \vee \\
\qquad \qquad (\text{MayAccess}_{pre}(x, z) \vee \text{MayAccess}_{pre}(y, z)) \wedge \\
\qquad \qquad (\text{MayAccess}_{pre}(x, u) \vee \text{MayAccess}_{pre}(y, u)))) \\
\hline
\vdash \text{true} \{ v := x.m(y) \} B \bowtie B'
\end{array} \\
\\
\text{(FRAME-METHCALL)} \\
\begin{array}{l}
\vdash A \{ v := x.m(y) \} \text{true} \bowtie B \\
B \equiv \forall z. (\text{MayAffect}(z, A') \rightarrow B'(z)) \wedge \\
\qquad \forall z. ((\text{MayAccess}(x, z) \vee \text{MayAccess}(y, z) \vee \text{New}(z)) \rightarrow \neg B'(z)) \\
\hline
\vdash A \wedge A' \{ v := x.m(y) \} A' \bowtie \text{true}
\end{array} \\
\\
\text{(CODE-INVAR-1)} \qquad \qquad \qquad \text{(CODE-INVAR-2)} \\
\frac{M(S) \equiv \text{spec } S \{ \overline{\text{Policy}}, P, \overline{\text{Policy}}' \} \quad B \equiv \forall x. (x \text{obeys } S \rightarrow P[x/\text{this}])}{\vdash \text{true} \{ \text{stmts} \} \text{true} \bowtie B \qquad \vdash e \text{obeys } S \{ \text{stmts} \} \text{true} \bowtie e_{pre} \text{obeys } S} \\
\\
\text{(CONS-2)} \qquad \qquad \qquad \text{(CONS-3)} \qquad \qquad \qquad \text{(CONS-4)} \\
\begin{array}{l}
\vdash A \{ \text{stmts} \} B \bowtie B'' \qquad \vdash A \{ \text{stmts} \} B \bowtie B' \qquad \vdash A \{ \text{stmts} \} A' \bowtie B' \\
A', B' \rightarrow_M A, _ \qquad A, B \rightarrow_M _, A' \qquad A, A' \rightarrow_M B \\
\hline
\vdash A' \{ \text{stmts} \} B' \rightarrow B \bowtie B'' \qquad \vdash A \{ \text{stmts} \} A' \bowtie B' \qquad \vdash A \{ \text{stmts} \} B \bowtie B'
\end{array} \\
\\
\text{(SEQUENCE)} \\
\frac{\vdash A \{ s_1 \} B_1 \bowtie B' \quad \vdash A_2 \{ s_2 \} B_2 \bowtie B' \quad A, B_1 \rightarrow_M _, A_2 \quad B_1, B_2 \rightarrow_M B}{\vdash A \{ s_1; s_2 \} B \bowtie B'}
\end{array}$$

Fig. 6. A selection of Hoare Logic rules; we assume that the module M is globally given

to the method call's receiver x or argument y . Additionally, it also promises that, *during* execution of the method, accessibility does not change, unless the participants (here z and u) were accessible to the receiver and/or the argument *before* the call. Note that this latter promise is made via the invariant (fourth) rather than the postcondition (third) part of the Hoare-tuple. Note also that this rule is applicable *even if we know nothing* about the receiver of the call: this rule and the invariants are critical to reasoning about risk.

(CODE-INVAR-1) allows reasoning under the hypothesis that an object o **obeys** its specification S : in this case o can be trusted to act in accordance with S always.

(FRAME-METHCALL) also expresses an axiom of object-capability languages, namely that in order to cause some visible effect, one must have access to an object able to per-

form the effect. Coupled with “only connectivity begets connectivity”, this implies that a method can cause some effect only if the caller has (transitive) access to some object able to cause the effect (including perhaps the callee).

The remaining rules each make use of the entailment judgement \rightarrow_M , which allows converting back and forth between one-state and two-state assertions and comes in number of flavours; the relevant ones are defined as follows.

Definition 19 (Entailment).

1. $A, B \rightarrow_M A', A'$ iff
 $\forall \sigma, \sigma'. \sigma \models A \wedge \sigma, \sigma' \models B \longrightarrow \sigma \models A' \wedge \sigma' \models A'$
2. $A, A' \rightarrow_M B$ iff
 $\forall \sigma, \sigma'. \sigma \models A \wedge \sigma' \models A' \longrightarrow \sigma, \sigma' \models B$
3. $B, B' \rightarrow_M B''$ iff
 $\forall \sigma, \sigma', \sigma''. \sigma, \sigma' \models B \wedge \sigma', \sigma'' \models B' \longrightarrow \sigma, \sigma'' \models B''$

The rules (CONS-3) and (CONS-4) make use of the entailment judgement to allow converting between one- and two-state postconditions during Hoare logic reasoning. To reason across sequenced computations $s_1; s_2$, the (SEQUENCE) rule requires finding a one-state assertion A_2 that holds after s_1 and is the precondition of s_2 . It uses the entailment $A, B_1 \rightarrow_M _, A_2$ to require that s_1 's execution guarantees A_2 , and the entailment $B_1, B_2 \rightarrow_M B$ to require that the combined execution of s_1 and s_2 guarantees the top-level postcondition B .

Theorem 3 (Soundness of the Hoare Logic).

*For all modules M , statements $stms$ and assertions A, B and B' ,
 If $M \vdash A \{ stms \} B' \bowtie B$, then $M \models A \{ stms \} B' \bowtie B$.*

The theorem is proven in [18].

In summary, we have four “code agnostic” rules — rules which are applicable regardless of the underlying code. Rules (FRAME-METHCALL) and (METH-CALL-2) express restrictions on the effects of a method call. Normally such restrictions stem from the specification of the method being called, but here we can argue in the absence of any such specifications, allowing us to reason about risk even in open systems. Rules (CODE-INVAR-1) and (CODE-INVAR-2) are applicable on *any* code, and allow us to assume that an object which **obeys** a specification S , satisfies all policies from S , and that the object, once trusted, will always be **obeying** S .

3.4 Proving Mutual Trust

We now use our Hoare Logic to prove the key steps of the escrow protocol, establishing mutual trust and delineating the risk. Here we have space to show just one-way trust between the escrow and seller in full: the remaining reasoning to establish mutual trust is outlined in the technical report [18]. Figure 7 shows the Hoare tuple for the first statement in method `deal` (line 4 from Figure 3). Lines 3-8 of Figure 7 describe the postcondition in case `escrowMoney` indeed **obeys** `ValidPurse`, while lines 9-17 make absolutely no assumption about the trustworthiness, or provenance, of `escrowMoney`.

postcondition as in 4-8.

To obtain 9-11, we will apply several of the code-agnostic rules. After all, here we cannot appeal to the specification of `sprout`, as we do not know whether `sellerMoney` adheres to `ValidPurse`. We start by application of (METH-CALL-2), and a consequence rule ((CONS-1) in [18]):

$$\begin{array}{l}
 \text{(D)} \quad \text{true} \\
 \quad \quad \{ \text{escrowMoney} := \text{sellerMoney.sprout} \} \\
 \quad \quad \times \\
 \quad \quad \forall z. \text{MayAccess}(\text{sellerMoney}, z) \rightarrow \text{MayAccess}_{pre}(\text{sellerMoney}, z)
 \end{array}$$

By applying the fact that $\forall u, v, w, \text{MayAccess}(u, v) \wedge \text{MayAccess}(v, w) \rightarrow \text{MayAccess}(u, w)$, and conjunction and inference rules on (D), we get:

$$\begin{array}{l}
 \text{(E)} \quad \neg \text{MayAccess}(\text{sellerMoney}, p) \\
 \quad \quad \{ \text{escrowMoney} := \text{sellerMoney.sprout} \} \\
 \quad \quad \times \\
 \quad \quad \forall z. \text{MayAccess}(\text{sellerMoney}, z) \rightarrow \neg \text{MayAccess}(z, p)
 \end{array}$$

By application of rule (CODE-INV-1), we obtain:

$$\begin{array}{l}
 \text{(F)} \quad \text{true} \\
 \quad \quad \{ \text{escrowMoney} := \text{sellerMoney.sprout} \} \\
 \quad \quad \times \\
 \quad \quad \forall p. (p \text{obeys ValidPurse} \rightarrow (\forall z. \text{MayAffect}(z, p.\text{balance}) \rightarrow \text{MayAccess}(z, p)))
 \end{array}$$

Through a combination of (E) and (F) and application of conjunction, and application of (FRAME-METH-CALL), we obtain that

$$\begin{array}{l}
 \text{(G)} \quad \neg \text{MayAccess}(\text{sellerMoney}, p) \\
 \quad \quad \{ \text{escrowMoney} := \text{sellerMoney.sprout} \} \\
 \quad \quad \times \\
 \quad \quad p \text{obeys}_{pre} \text{ValidPurse} \rightarrow (p.\text{balance} = p.\text{balance}_{pre})
 \end{array}$$

Now by applying (CONS-2) on (F), we obtain

$$\begin{array}{l}
 \text{(H)} \quad \text{true} \\
 \quad \quad \{ \text{escrowMoney} := \text{sellerMoney.sprout} \} \\
 \quad \quad \times \\
 \quad \quad \forall p. p \text{obeys}_{pre} \text{ValidPurse} \rightarrow \\
 \quad \quad \quad (p.\text{balance} = p.\text{balance}_{pre} \vee \text{MayAccess}(\text{sellerMoney}, p))
 \end{array}$$

We now apply (CONS-1) from [18] to conjoin the invariant and postcondition, obtaining

$$\begin{array}{l}
 \text{(I)} \quad \text{true} \\
 \quad \quad \{ \text{escrowMoney} := \text{sellerMoney.sprout} \} \\
 \quad \quad \forall p. p \text{obeys}_{pre} \text{ValidPurse} \rightarrow \\
 \quad \quad \quad (p.\text{balance} = p.\text{balance}_{pre} \vee \text{MayAccess}(\text{sellerMoney}, p)) \\
 \quad \quad \times \\
 \quad \quad \text{true}
 \end{array}$$

Last, we obtain lines 11-12 from (METH-CALL-2). We also obtain lines 13-17 from (METH-CALL-2), and (CONS-1) from [18].

4 Related Work

Object Capabilities and Sandboxes. *Capabilities* were developed in the 60’s by Dennis and Van Horn [10] within operating systems, and were adapted to the programming languages setting in the 70’s [34]. *Object capabilities* were first introduced [30] in the early 2000s, and much recent work investigates the safety or correctness of object capability programs. Google’s Caja [33] applies sandboxes, proxies, and wrappers to limit components’ access to *ambient* authority. Sandboxing has been validated formally: Maffeis et al. [27] develop a model of JavaScript, demonstrate that it obeys two principles of object capability systems and show how untrusted applications can be prevented from interfering with the rest of the system.

JavaScript analyses. More practically, there are a range of recent analyses of JavaScript [23, 5, 38, 26, 43] based on static analyses or type checking. Lerner et al. extend these approaches to ensure browser extensions observe “*private mode*” [26], while Dimoulas et al. [11] enforce explicit access policies. The problem posed by the Escrow example is that it establishes a two-way dependency between trusted and untrusted systems — precisely the kind of dependencies these techniques prevent.

Concurrent Reasoning Our Hoare logic invariants are similar to the guarantees in Rely-Guarantee reasoning [22]. Deny-Guarantee [12] distinguishes between assertions guaranteed by a thread, and actions denied to all other threads. Deny properties correspond to our requirements that certain properties be preserved by all code linked to the current module. Compared with our work, rely-guarantee and deny-guarantee assumes coöperation: composition is legal only if threads adhere to their rely or deny properties and guarantees. Our modules have to be robust and ensure that their invariants cannot be affected by *any* arbitrary, uncertified, untrusted code.

Relational models of trust. Artz and Gil [4] survey various types of trust in computer science generally, although trust has also been studied in specific settings, ranging from peer-to-peer systems [2] and cloud computing [20] to mobile ad-hoc networks [9], the internet of things [19], online dating [37], and as a component of a wider socio-technical system [8, 45]. Considering trust (and risk) in systems design, Cahill et al.’s overview of the SECURE project [6] gives a good introduction to both theoretical and practical issues of risk and trust, including a qualitative analysis of an e-purse example. This project builds on Carbone’s trust model [7] which offers a core semantic model of trust based on intervals to capture both trust and uncertainty in that trust. Earlier Abdul-Rahman proposed using separate relations for trust and recommendation in distributed systems [1], more recently Huang and Nicol preset a first-order formalisation that makes the same distinction [21]. Solhaug and Stølen [42] consider how risk and trust are related to uncertainties over actual outcomes versus knowledge of outcomes. Compared with our work, these approaches produce models of trust relationships between high-level system components (typically treating risk as uncertainty in trust) but do not link those relations to the system’s code.

Logical models of trust. Various different logics have been used to measure trust in different kinds of systems. Some of the earliest work is Lampson et al.’s “speaks for” and “says” constructs [24], clear precursors to our “**obeys**” but for authentication rather than specifications. Murray [35] models object capability patterns in CSP, and applies automatic refinement checking to analyse various properties in the presence of untrusted

components. Ries et al. [40] evaluate trust under uncertainty by evaluating Boolean expressions in terms of real values. Carbone et al. [41] and Aldini [3] model trust using temporal logic. Primiero and Taddeo [39] have developed a modal type theory that treats trust as a second-order relation over relations between counterparties. Merro and Sibilio [29] developed a trust model for a process calculus based on labelled transition systems. Compared with ours, these approaches use process calculi or other abstract logical models of systems, rather than engaging directly with the system’s code.

Verification of Object Capability Programs. Drossopoulou and Noble [13, 36] have analysed Miller’s Mint and Purse example [30] by expressing it in Joe-E and in Grace [36], and discussed the six capability policies as proposed in [30]. In [16], they proposed a complex specification language, and used it to fully specify the six policies from [30]; uncovering the need for another four policies. More recently, [14] they have shown how different implementations of the underlying Mint and Purse systems coexist with different policies. In contrast, this work formalises the informal ideas from [17], proposes *Focal*, which is untyped and modelled on Grace and JavaScript rather than Java; a much simpler specification language *Chainmail*; the **obeys** predicate to model trust; *MayAccess* and *MayAffect* to model risk; a full specification of the `Escrow`; and a Hoare logic for reasoning about risk and trust, applied to the Escrow specification.

5 Conclusions and Further Work

In this paper we addressed the questions of specification of risk, trust, and reasoning about such specifications. To answer these questions, we contributed:

- *Hypothetical* predicates **obeys** to model trust, *MayAccess* and *MayAffect* to model risk, and their formal semantics.
- *Open Assertions* and *Open Policies* whose validity must be guaranteed, even when linked with *any* other code.
- *Formal models* of *Focal* and *Chainmail*.
- *Hoare four-tuples* that make invariants explicit.
- A *Hoare logic* incorporating code agnostic inference rules.
- *Formal reasoning* to prove key steps of the Escrow Exchange.

In further work we will extend our approach to deal with concurrency, distribution, exceptions, networking, aliasing, and encapsulation. Finally, we hope to develop automated reasoning techniques to make these kinds of specifications practically useful.

Bibliography

- [1] A. Abdul-Rahman and S. Halles. A distributed trust model. In *New Security Paradigms Wkshp.*, 1988. Langdale, Cumbria.
- [2] K. Aberer and Z. Despotovic. Managing trust in a peer-2-peer information system. In *CKIM*, 2001.
- [3] A. Aldini. A calculus for trust and reputation systems. In *IFIPTM*, 2014.
- [4] D. Artz and Y. Gil. A survey of trust in computer science and the semantic web. *Journal of Web Semantics*, 2007.
- [5] K. Bhargavan, A. Delignat-Lavaud, and S. Maffei. Language-based defenses against untrusted browser origins. In *USENIX Security*, 2013.
- [6] Cahill et al. Using trust for secure collaboration in uncertain environments. *Pervasive Computing*, July 2003.
- [7] M. Carbone, M. Nielsen, and V. Sassone. A formal model for trust in dynamic networks. In *SEFM*, 2003.
- [8] J.-H. Cho and K. S. Shan. Building trust-based sustainable networks. *IEEE Tech. and Soc.*, Summer, 2013.
- [9] J.-H. Cho, A. Swami, and I.-R. Chen. A survey on trust management for mobile ad hoc networks. *IEEE Comms. Surv. & Tuts.*, 13(4), 2011.
- [10] J. B. Dennis and E. C. V. Horn. Programming Semantics for Multiprogrammed Computations. *Comm. ACM*, 9(3), 1966.
- [11] C. Dimoulas, S. Moore, A. Askarov, and S. Chong. Declarative policies for capability control. In *Computer Security Foundations Symposium*, 2014.
- [12] M. Dodds, X. Feng, M. Parkinson, and V. Vafeiadis. Deny-guarantee reasoning. In *ESOP*. Springer, 2009.
- [13] S. Drossopoulou and J. Noble. The need for capability policies. In *FTfJP*, 2013.
- [14] S. Drossopoulou and J. Noble. How to break the bank: Semantics of capability policies. In *iFM*, 2014.
- [15] S. Drossopoulou and J. Noble. Invited Talk: Towards Reasoning about Risk and Trust in the Open World, 2014. slides from "<http://www/doc.ic.ac.uk/~scd>".
- [16] S. Drossopoulou and J. Noble. Towards capability policy specification and verification, May 2014. ecs.victoria.ac.nz/Main/TechnicalReportSeries.
- [17] S. Drossopoulou, J. Noble, and M. S. Miller. Swapsies on the Internet. In *PLAS*, 2015.
- [18] S. Drossopoulou, J. Noble, M. S. Miller, and T. Murray. More Reasoning about Risk and Trust in an Open World. Technical Report ECSTR-15-08, VUW, 2015.
- [19] L. Gu, J. Wang, and B. Sun. Trust management mechanism for Internet of Things. *China Communications*, Feb. 2014.
- [20] S. M. Habib and M. M. Sebastian Ries and. Towards a trust management system for cloud computing. In *TrustCom*, 2011.
- [21] J. Huang and D. M. Nicol. A formal-semantics-based calculus of trust. *IEEE INTERNET COMPUTING*, 2010.
- [22] C. Jones. Specification and design of (parallel) programs. In *IFIP Congress*, 1983.

- [23] R. Karim, M. Dhawan, V. Ganapathy, and C.-C. Shan. An Analysis of the Mozilla Jetpack Extension Framework. In *ECOOP*, Springer, 2012.
- [24] B. Lampson, M. Abadi, M. Burrows, and E. Wobblers. Authentication in Distributed Systems: Theory and Practice. *ACM TOCS*, 10(4):265–310, 1992.
- [25] B. W. Lampson. A note on the confinement problem. *Communications of the ACM*, 16:613–615, 1973.
- [26] B. S. Lerner, L. Elberty, N. Poole, and S. Krishnamurthi. Verifying web browser extensions’ compliance with private-browsing mode. In *ESORICS*, Sept. 2013.
- [27] S. Maffei, J. Mitchell, and A. Taly. Object capabilities and isolation of untrusted web applications. In *Proc of IEEE Security and Privacy*, 2010.
- [28] R. Merrill. focal: new conversational language. DEC, 1969. homepage.cs.uiowa.edu/~jones/pdp8/focal/focal69.html.
- [29] M. Merro and E. Sibilio. A calculus of trustworthy ad hoc networks. *Formal Aspects of Computing*, page 25, 2013.
- [30] M. S. Miller. *Robust Composition: Towards a Unified Approach to Access Control and Concurrency Control*. PhD thesis, Baltimore, Maryland, 2006.
- [31] M. S. Miller, T. V. Cutsem, and B. Tulloh. Distributed electronic rights in JavaScript. In *ESOP*, 2013.
- [32] M. S. Miller, C. Morningstar, and B. Frantz. Capability-based financial instruments: From object to capabilities. In *Financial Cryptography*. Springer, 2000.
- [33] M. S. Miller, M. Samuel, B. Laurie, I. Awad, and M. Stay. Safe active content in sanitized JavaScript. code.google.com/p/google-caja/.
- [34] J. H. Morris Jr. Protection in programming languages. *CACM*, 16(1), 1973.
- [35] T. Murray. *Analysing the Security Properties of Object-Capability Patterns*. D.Phil. thesis, University of Oxford, 2010.
- [36] J. Noble and S. Drossopoulou. Rationally reconstructing the escrow example. In *FTJP*, 2014.
- [37] G. Norcie, E. D. Cristofaro, and V. Bellotti. Bootstrapping trust in online dating: Social verification of online dating profiles. In *Fin. Crypt. & Data Sec.*, 2013.
- [38] J. G. Politz, S. A. Eliopoulos, A. Guha, and S. Krishnamurthi. Adsafety: Type-based verification of JavaScript sandboxing. In *USENIX Security*, 2011.
- [39] G. Primiero and M. Taddeo. A modal type theory for formalizing trusted communications. *J. Applied Logic*, 10, 2012.
- [40] S. Ries, S. M. Habib, M. M. Sebastian Ries and, and V. Varadharajan. Certain-logic: A logic for modeling trust and uncertainty. In *TRUST*, 2011. LNCS 6740.
- [41] Roberto Carbone et al. Towards formal validation of trust and security in the Internet of services. In *Future Internet Assembly*, 2001. LNCS 6656.
- [42] Solhaug and Stølen. Uncertainty, subjectivity, trust and risk: How it all fits together. In *STM*, 2011.
- [43] A. Taly, U. Erlingsson, J. C. Mitchell, M. S. Miller, and J. Nagra. Automated Analysis of Security-Critical JavaScript APIs. In *SOSP*, 2011.
- [44] The Swapsies. Got Got Need. In *5: A February Records Anniversary Compilation*. February Records, 2015.
- [45] M. Walterbusch, B. Martens, and F. Teuteberg. Exploring trust in cloud computing: A multi- method approach. In *ECIS*, page 145, 2013.