

A Java Framework for Smart Contracts

Fausto Spoto

Università di Verona, Italy

WTSC 2019, Saint Kitts

February 22, 2019

Programming Smart Contracts

Transactions are atomic computation steps of a blockchain

Smart contracts are an object-oriented presentation of transactions over a shared heap

Which programming language?

- Bitcoin uses a low-level non-Turing complete bytecode for programming transactions
- Ethereum uses Turing-complete languages compiled into EVM bytecode (mainly Solidity)

Compared to Bitcoin bytecode, Solidity was a revolution, but...

- minimal toolbelt (IDE, builders, integrators, testing, analyzers...)
- low-level semantics (memory/storage distinction is explicit)
- no exception handling
- weak typing
- no inner classes, nor anonymous classes
- no lambda expressions, nor method references
- no generics
- very small support library
- limited production of libraries for blockchain
- another language to learn!

Compared to Bitcoin bytecode, Solidity was a revolution, but...

- minimal toolbelt (IDE, builders, integrators, testing, analyzers...)
- low-level semantics (memory/storage distinction is explicit)
- no exception handling
- weak typing
- no inner classes, nor anonymous classes
- no lambda expressions, nor method references
- no generics
- very small support library
- limited production of libraries for blockchain
- another language to learn!

What about using Java instead?

Java for blockchain programming

- there are Java implementations of blockchain nodes
- EthereumJ allows Java clients to query Ethereum blockchains
- NEO allows one to use the Java syntax for writing smart contracts:

```
1 public class ICOTemplate extends SmartContract {
2     public static Object deploy() { // static Object !!!
3         if (getTotalSupply().length != 0) {
4             Runtime.log("Insufficient token supply-No action");
5             return false;
6         }
7         Storage.put(Storage.currentContext(), TemplateToken.getOwner(), ...);
8         Storage.put(Storage.currentContext(), NEP5Template.TOTAL_SUPPLY, ...);
9         return true;
10    }
11 }
```

- Aion seems to go in a similar direction

Takamaka: Transparent Use of Java for Smart Contracts

```
1 import takamaka.lang.Contract; import takamaka.lang.Payable;
2 import takamaka.lang.Storage; import takamaka.util.StorageList;
3
4 public class CrowdFunding extends Contract {
5     private final StorageList<Campaign> campaigns = new StorageList<>(); // generics!
6
7     // callable from everywhere
8     public int newCampaign(Contract beneficiary, int goal) {
9         int campaignId = campaigns.size();
10        campaigns.add(new Campaign(beneficiary, goal));
11        return campaignId;
12    }
13
14    // only callable from another instance of Contract; requires payment
15    public @Payable @Entry void contribute(int amount, int campaignID) {
16        campaigns.elementAt(campaignID).addFunder(caller(), amount);
17    }
18
19    // callable from everywhere
20    public boolean checkGoalReached(int campaignID) {
21        return campaigns.elementAt(campaignID).payIfGoalReached();
22    }
}
```

Takamaka: Transparent Use of Java for Smart Contracts

```
24 private class Campaign extends Storage { // inner class
25     private final Contract beneficiary;
26     private final int fundingGoal;
27     private final StorageList<Funder> funders = new StorageList<>();
28     private int amount; // one could also use BigInteger for this field
29
30     private Campaign(Contract beneficiary, int fundingGoal) {
31         this.beneficiary = beneficiary; this.fundingGoal = fundingGoal;
32     }
33
34     private void addFunder(Contract who, int amount) {
35         funders.add(new Funder(who, amount)); this.amount += amount;
36     }
37
38     private boolean payIfGoalReached() {
39         if (amount >= fundingGoal) {
40             pay(beneficiary, amount); amount = 0; return true;
41         }
42         else
43             return false;
44     }
45 }
46 }
```

How Takamaka Nodes Execute Transactions

Execute `cf = new CrowdFunding()`:

```
1 | cf = new CrowdFunding(); // Java execution in RAM only
2 | updates = emptyset;
3 | cf.extractUpdates(updates); // Takamaka provides this method
4 | blockchain.store(updates); // expands the blockchain
5 | return cfRef = cf.storageReference to the wallet
```

Execute `id = cf.newCampaign(beneficiary, 42)`:

```
1 | cf = blockchain.deserialize(cfRef);
2 | beneficiary = blockchain.deserialize(beneficiaryRef);
3 | id = cf.newCampaign(beneficiary, 42); // Java execution in RAM only
4 | updates = emptyset;
5 | cf.extractUpdates(updates); // Takamaka provides this method
6 | beneficiary.extractUpdates(updates);
7 | blockchain.store(updates); // expands the blockchain
8 | return id to the wallet
```


Key Features of Takamaka

- ① constructors and methods are completely normal Java code, that operates without explicit primitives for storage manipulation
- ② only white-listed methods of the standard Java library can be used: the deterministic ones
- ③ the wallet uses storage references to refer to objects in the blockchain, since actual memory addresses are node-dependent
- ④ deserialization of storage references is lazy
- ⑤ only updates are serialized at the end (this is not standard Java serialization!)

Instrumentation of Contract Classes

```
public class MyContract extends takamaka.lang.Contract {  
    public @Entry T m1(args  
        body  
    }  
  
    public @Payable @Entry T m2(int amount, args  
        body  
    }  
}
```

Instrumentation of Contract Classes

```
public class MyContract extends takamaka.lang.Contract {  
    public @Entry T m1(args, Contract caller) {  
        entry(caller);  
        body  
    }  
  
    public @Payable @Entry T m2(int amount, args                ) {  
  
        body  
    }  
}
```

Instrumentation of Contract Classes

```
public class MyContract extends takamaka.lang.Contract {  
    public @Entry T m1(args, Contract caller) {  
        entry(caller);  
        body  
    }  
  
    public @Payable @Entry T m2(int amount, args, Contract caller) {  
        payableEntry(caller, amount);  
        body  
    }  
}
```

Instrumentation for Gas Metering

Before each bytecode instruction, Takamaka adds a call to `takamaka.lang.Gas.tick(int amount)`

- the amount depends on the bytecode instruction
- tick can throw an `OutOfGasException`
- that exception cannot be caught in code

White-listed Methods

They must be deterministic

- methods of `java.lang.String*`
- methods of wrapper classes in `java.lang`.
- `java.util.Arrays`
- `java.util.ArrayList`, `java.util.LinkedList`, `java.util.PriorityQueue`
- most `java.util.Date`
- ...

Black-listed

- `java.lang.System.currentTimeMillis()`
- `java.lang.Object.hashCode()`
- ...

Should we white-list `java.util.HashSet` and `java.util.HashMap`?

- 1 no, since iteration on them is not deterministic! (behavioral non-determinism)
- 2 no, since (for instance) `set.add(element)` might call a variable number of times the `hashCode` and `equals` methods on the elements of the set (cost non-determinism)

A solution is to require that `hashCode` must be redefined on objects put inside these collections (enforced by static/dynamic verification)

The JVM verifies that some basic Java constraints are met:

- types are strong
- visibility modifiers are honored
- `final` methods and classes are not redefined
- called methods do exist (no fall-back methods!) or an exception is thrown
- accessed fields do exist or an exception is thrown

Takamaka Further Verification

Static

- storage classes have only fields of type primitive, storage, `java.lang.String`, `java.math.BigInteger` or `java.lang.Object` (for generics)
- `@Entry` is only applied to methods of contracts
- `@Payable` is only applied to `@Entry` methods with an `int` amount first parameter
- `@Entry` methods are only called from the code of a contract
- `caller()` can only be called on `this` and from an `@Entry` method

Dynamic

- accessed storage fields of type `java.lang.Object` actually hold a storage object, a `java.lang.String` or a `java.math.BigInteger`
- `@Entry` methods are only called from a distinct contract instance
- `@Payable` methods receive a non-negative amount and the caller contract has enough funds

Status of the Project

Done

- September 2018: project started
- December 2018: first draft of the working principles (this paper!)
- February 2019: Java bytecode instrumenter for storage objects and contracts (the transformation described in the previous slides)

To do

- March 2019: execution of instrumented contracts on an in-memory simulation of a blockchain
- April 2019: verification of jars before instrumentation
- fall 2019: integration into a real blockchain currently being developed by an independent company

Open Questions and Future Work

- is Takamaka actually keeping track of all storage updates?
- does the updates-only approach actually support scalability?
- which security guarantees can be proved for Takamaka?
- can we add a layer of non-trivial static analysis?
 - overflow/underflow checks
 - complexity analysis
 - inference of parametric gas costs
- how much can we enlarge the set of white-listed methods?
- which other types can we allow in storage objects? (wrapper types, arrays, enums. . .)

THANKS!