

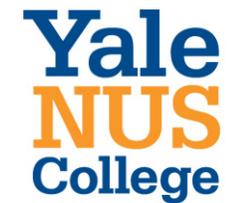
# CoSplit

Practical Smart Contract Sharding with  
Static Program Analysis

George Pîrlea

Amrit Kumar

*Ilya Sergey*



problem

# Smart Contract Sharding

*with* Static Program Analysis

technique

# Scaling blockchains *that* support smart contracts

*with* Static Program Analysis

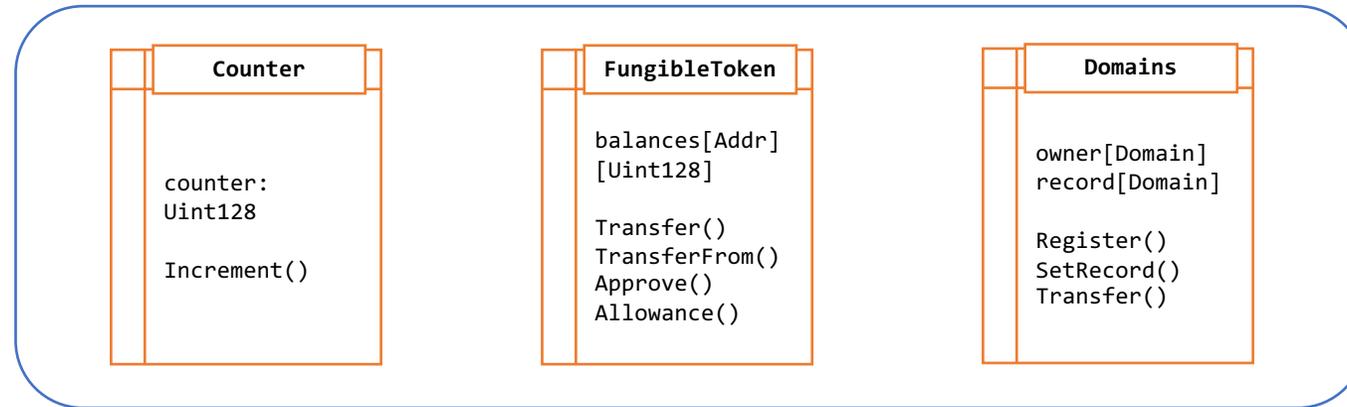
**Scaling blockchains**  
*that* support smart contracts

*with* **Static Program Analysis**

Blockchains don't scale

## Blockchain state

## block 0

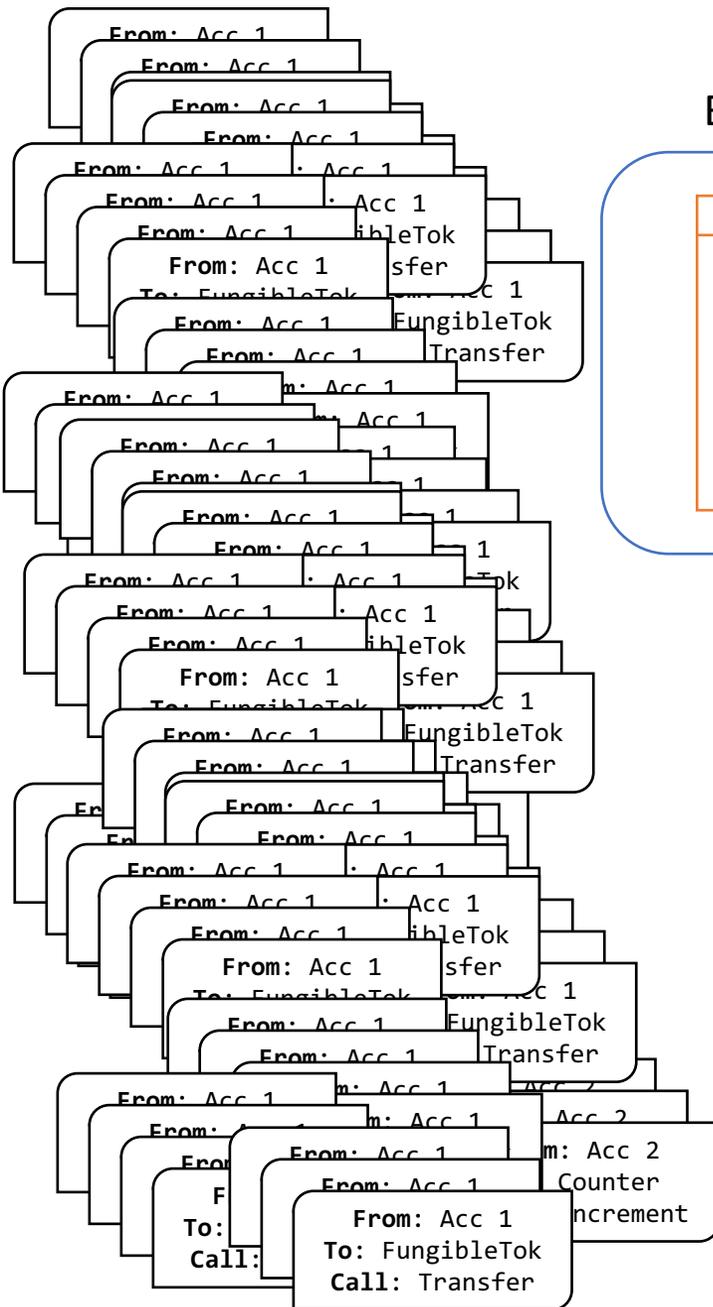


From: Acc 1  
To: FungibleTok  
Call: Transfer

From: Acc 2  
To: Counter  
Call: Increment

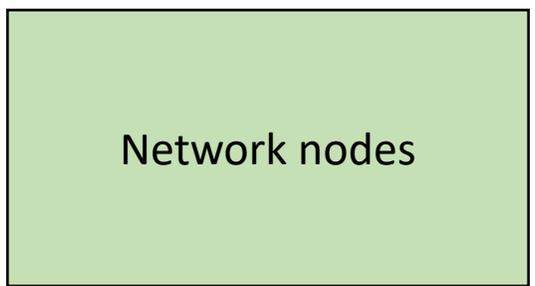

From: Acc 1  
To: Domains  
Call: Register

Network nodes




### Blockchain state

### block 2

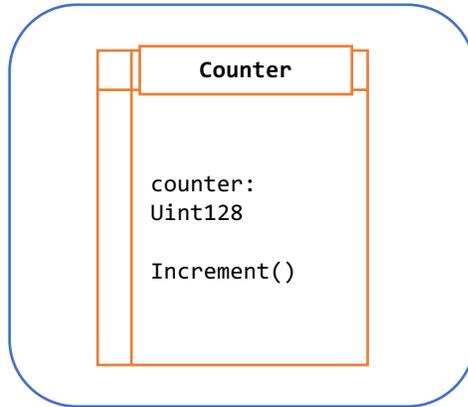


nodes have limited bandwidth  
and limited processing power




Network nodes


### Shard 1 state



### Shard 2 state



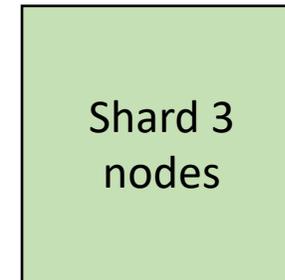


### Shard 3 state



From: Acc 1  
To: FungibleTok  
Call: Transfer

From: Acc 2  
To: Counter  
Call: Increment


From: Acc 1  
To: Domains  
Call: Register




# Sharding has limitations

Especially for smart contracts!





### Shard 1 state




### Shard 2 state




### Shard 3 state




### Shard 1 state



### Shard 2 state



### Shard 3 state



Shard 1  
nodes

Shard 2  
nodes

Shard 3  
nodes

# Sharded contracts: intuition

- A *blockchain* is a state transition system, consisting of:
  - State
  - Rules that say which state updates are legal
- We have strategies for sharding blockchains (for certain kinds of rules\*)
- A **smart contract** is a state transition system, consisting of:
  - State
  - Code that says which state updates are legal
- ...

\* - some rules shard better than others  
and some don't shard at all

Does the code of the contract  
define a shardable state  
machine?

Does the code have  
property X?

We can shard contracts the  
same way we shard  
blockchains.

Static analysis uncovers the opportunities.

# CoSplit


Practical Smart Contract Sharding with  
Static Program Analysis

# Our contributions


- Identify enabling mechanisms for sharding Ethereum-style contracts
  - and show their adequacy for some realistic contracts
- CoSplit, a static analysis tool that infers sharding strategies for smart contracts written in Scilla, an ML-style smart contract language
- End-to-end integration of CoSplit with a production-grade sharded blockchain (Zilliqa)
- Evaluation of the inferred sharding strategies
  - almost linear throughput increase as number of shards goes up

# Mechanism (1): disjoint state ownership


Shard 1 state



Shard 2 state




Shard 3 state




```
transition Donate ()  
  blk <- & BLOCKNUMBER;  
  in_time = blk_leq blk max_block;  
  match in_time with  
  | True =>  
    c <- exists backers[_sender];  
    match c with  
    | False =>  
      accept;  
      backers[_sender] := _amount;
```

# Mechanism (1): disjoint state ownership


Shard 1 state



Shard 2 state



Shard 3 state




# Mechanism (2): commutative effects

```
field counter : Uint128 = Uint128 0
```

```
transition Increment ()  
  c <- counter;  
  inc = Uint128 1;  
  new_c = builtin add c inc;  
  counter := new_c  
end
```

Cumulative result can be obtained by *joining* the contributions from each shard.

# Mechanism (2): commutative effects



# Commutative operations do not imply commutative effects!

```
field counter : Uint128 = Uint128 0
```

```
transition Increment ()  
  c <- counter;  
  inc = Uint128 1;  
  new_c = builtin add c inc;  
  counter := new_c  
end
```

```
transition Double ()  
  c <- counter;  
  new_c = builtin add c c;  
  counter := new_c  
end
```

$$2 * (2x + 1) \neq 2 * 2x + 1$$

# Static analysis for transition effects

- Produce an **effect summary** for every transition in the contract
  - Effects include: reads, writes, accepting funds, sending messages, conditioning on values derived from mutable fields
  - The effect summary *over-approximates* the behaviour of the transition

# Static analysis for transition effects

- Produce an **effect summary** for every transition in the contract
  - Effects include: reads, writes, accepting funds, sending messages, conditioning on values derived from mutable fields
  - The effect summary *over-approximates* the behaviour of the transition
- Determine which effects are commutative using a **linearity-tracking flows-to analysis**
  - The analysis is expressed as a type system for “contribution types” and is compositional (but gives uninformative types in some cases)

Constant	$x, y$	constant contract field or transition parameter
Mutable	$f$	mutable field or map-field access via parameter
<hr/>		
Contrib. src.	$cs ::= x \mid f$	
Cardinality	$card ::= \text{None} \mid \text{Linear} \mid \text{NonLinear}$	
Operation	$op ::= + \mid - \mid * \mid \dots$	
Abstr. expr.	$e ::= \top \mid \overline{(cs, card, op)}$	
<hr/>		
Effect	$\varepsilon ::= \text{Read}(f) \mid \text{Write}(f, e) \mid \text{AcceptFunds} \mid$ $\text{Condition}(e) \mid \text{Event}(e) \mid \text{SendMsg}(e) \mid \top$	

```

1 transition Transfer(to: ByStr20, amount: Uint)
2   from_bal <- balances[_sender];           Read(balances[_sender])
3   match from_bal with                     Condition(balances[_sender])
4   | Some bal =>                           (balances[_sender], Linear, ∅)
5     match amount ≤ bal with              Condition(balances[_sender], amount)
6     | True =>                             {(balances[_sender], Linear, sub),
7     new_from_bal = builtin sub bal amount; (amount, Linear, sub)}
8     balances[_sender] := new_from_bal;    Write(balances[_sender],
9     to_bal <- balances[to];              {(balances[_sender], Linear, sub),
10    new_to_bal = match to_bal with        (amount, Linear, sub)}}
11    | Some bal => builtin add bal amount  Read(balances[to])
12    | None => amount
13  end;
14  balances[to] := new_to_bal             Write(balances[to],
                                           {(balances[to], Linear, add),
                                           (amount, Linear, add)}})

```

# Sharding Constraints

A language for restricting a set of shards that can execute a certain transition of a contract.

Constraint       $oc ::= \text{Owns}(f) \mid \text{UserAddr}(x) \mid \text{NoAliases}(\langle x, y \rangle) \mid$   
SenderShard  $\mid$  ContractShard  $\mid \perp$

Constraint  $oc ::= \text{Owns}(f) \mid \text{UserAddr}(x) \mid \text{NoAliases}(\langle x, y \rangle) \mid$   
 $\text{SenderShard} \mid \text{ContractShard} \mid \perp$

Join  $\uplus_f ::= \text{OwnOverwrite} \mid \text{IntMerge}$

Read(balances[\_sender])

Condition(balances[\_sender])

**Weak reads**

Condition(balances[\_sender], amount)

**Owns(balances[\_sender])**

**OwnOverwrite** join for owned contributions

Write(balances[\_sender],  
{(balances[\_sender], Linear, sub),  
(amount, Linear, sub)})

**IntMerge** join for un-owned contributions

Read(balances[to])

Write(balances[to],  
{(balances[to], Linear, add),  
(amount, Linear, add)})

```

1 transition Transfer(to: ByStr20, amount: Uint)
2   from_bal <- balances[_sender];
3   match from_bal with
4   | Some bal =>
5     match amount ≤ bal with
6     | True =>
7       new_from_bal = builtin sub bal amount;
8       balances[_sender] := new_from_bal;
9       to_bal <- balances[to];
10      new_to_bal = match to_bal with
11      | Some bal => builtin add bal amount
12      | None => amount
13      end;
14      balances[to] := new_to_bal

```

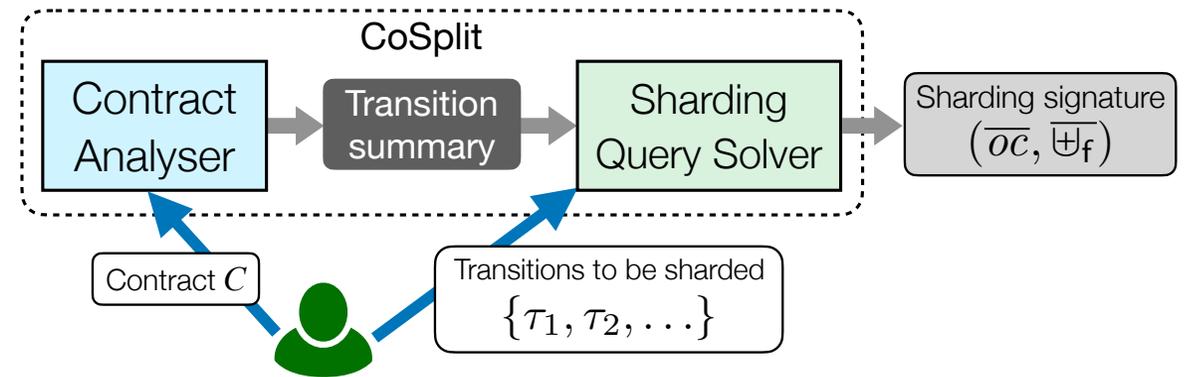
**Owns**(balances[\_sender])

**OwnOverwrite** join for owned contributions

**IntMerge** join for un-owned contributions

```
type expr_type =  
  | ETop  
  | EVal of known_contrib  
  | ECompositeVal of expr_type * expr_type  
  
  | EOp of contrib_op * expr_type  
  
  | EComposeSequence of expr_type list  
  | EComposeParallel of expr_type * expr_type list  
  
  | EFun of efun_desc  
  | EApp of efun_desc * expr_type list
```

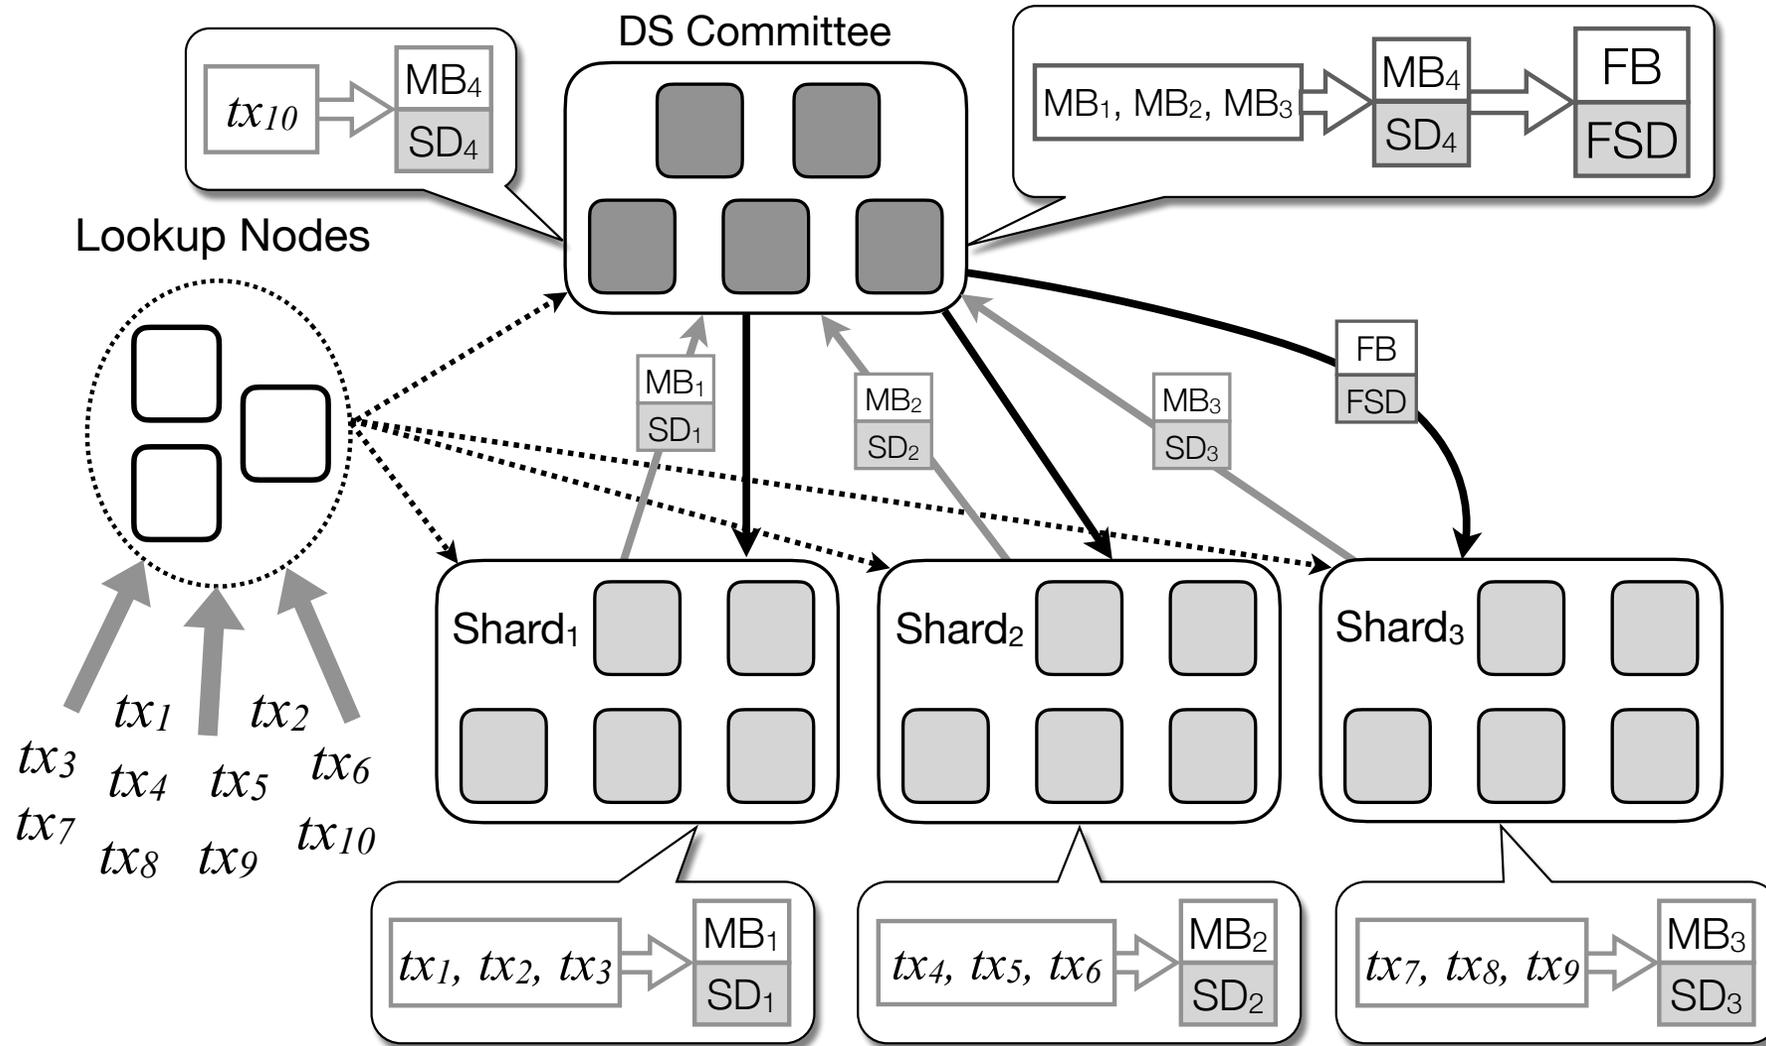
# Analysis Pipeline



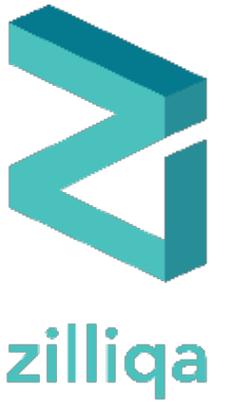
1. Derive summaries for all contract transitions (R/W, operations, lin-ty)
2. Take from a user a set of transitions she wants to shard + weak reads
3. Produce an optimal (the most permissive) set of sharding constraints. These constraints determine conditions a shard need to satisfy in order to run the transaction with this transition.
4. Sharding more transitions of a contract => stronger constraints

# Integration

# A sharded blockchain design

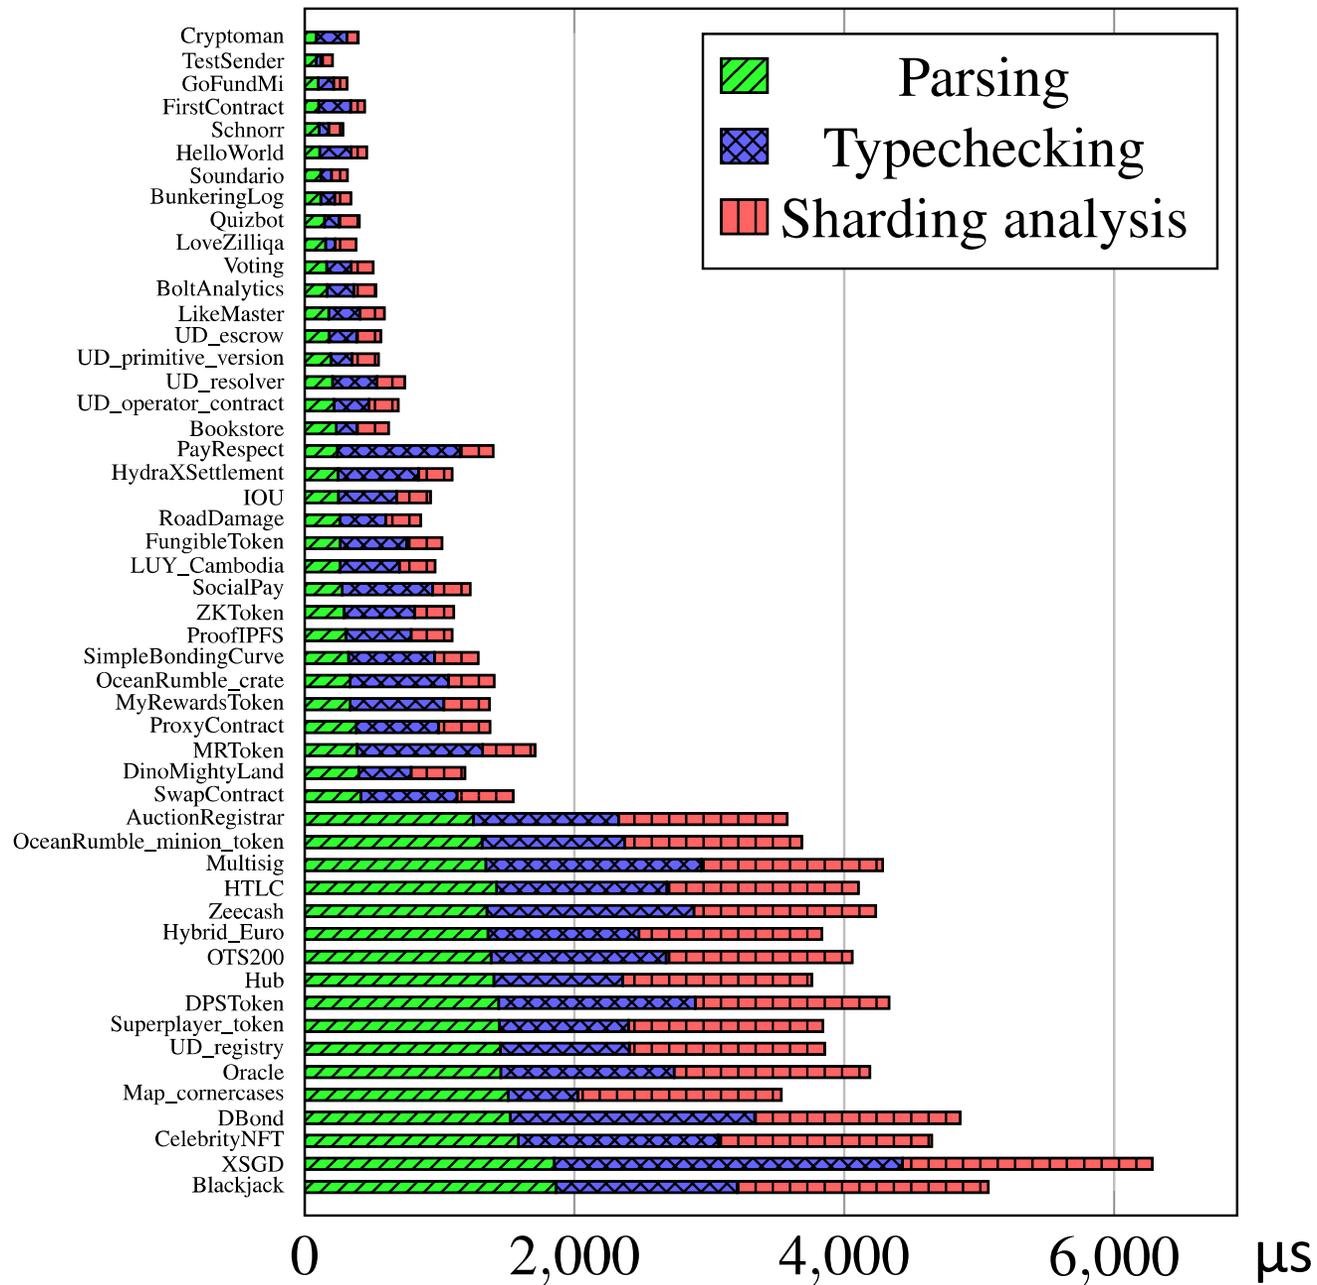



# Integrating CoSplit with Zilliqa



1. Run the static analysis upon contract deployment
2. Store the resulting sharding signature  
= set of transition constraints + join instructions for each field in the contract
3. Dispatch transactions to shards using the constraints to determine which shard the transaction needs to go to
4. Merge (join) contributions from shards before sequential/cross-shard transactions are processed

# Evaluation





*Average TPS for different contract transitions as a function of number of shards, over 10 epochs.*

# Limitations

- Currently no support for sharding multi-contract transactions
  - would need to somehow combine the signatures from multiple contracts
- At the moment, the contract is *always* sharded
  - non-parallelizable operations become cross-shard (synchronized)
  - inefficient if the contract has either low transaction volume (overhead dominates) or many transactions to non-parallelizable transitions (forced into expensive synchronization)
  - would want dynamic sharding – only shard the contract when transactions are of a profile that is known to shard well; otherwise keep it to a single shard

# Future work

- Implement dynamic sharding to alleviate fixed-strategy limitation
- Produce sharding signatures that allow sharding of multi-contract calls with commutative effects
  - proxy contracts are a common example of this
- Automatic contract repair to make contracts shardable, e.g.:
  - split records into a separate map for each component
  - translate to compare-and-swap transitions

```

transition transfer(to: ByStr20, tokenId: Uint256)
  getTokenOwner <- tokenOwners[tokenId];
  match getTokenOwner with
  | None => throw
  | Some tokenOwner =>
    isOwner = builtin eq _sender tokenOwner;
    (* ... *)
    getOperatorStatus <-
      operatorApprovals[tokenOwner][_sender];
    (* ... *)
    tokenOwners[tokenId] := to;

```

```

transition transfer(tokenOwner: ByStr20,
                    to: ByStr20, tokenId: Uint256)
  getTokenOwner <- tokenOwners[tokenId];
  match getTokenOwner with
  | None => throw
  | Some actual =>
    isCorrectOwner = builtin eq tokenOwner actual;
    match isCorrectOwner with
    | False => throw
    | True =>
      isOwner = builtin eq _sender tokenOwner;
      (* ... *)
      getOperatorStatus <-
        operatorApprovals[tokenOwner][_sender];
      (* ... *)
      tokenOwners[tokenId] := to;

```

# To Take Away

- Sharding is a solution to the blockchain scalability problem
- Some smart contract logic can be sharded (“pessimistically parallelized”) in the same way as simple blockchain transactions.
- We use static analysis to soundly determine conditions under which (parts of) smart contracts can be executed in parallel.
- The technique has been integrated into real-world blockchain and shown to give observable increase in the throughput.

Thanks!