

Modular Program Verification

Peter Müller

ETH zürich

OpenJDK's `java.util.Collection.sort()` is broken: The good, the bad and the worst case*

Stijn de Gouw^{1,2}, Jurriaan Rot^{3,1}, Frank S. de Boer^{1,3}, Richard Bubel⁴, and
Reiner Hähnle⁴

- TimSort is the default sorting algorithm for Collections in Sun's JDK, OpenJDK, and Android SDK
- Certain large arrays ($\geq 67\text{M}$) lead to index-out-of-bounds errors
- Bug was detected during a verification attempt
- Previous attempts to fix related errors were ineffective

Testing and code reviews are not sufficient to detect certain bugs

Priority Inheritance Protocols: An Approach to Real-Time Synchronization

LUI SHA, MEMBER, IEEE, RAGUNATHAN RAJKUMAR, MEMBER, IEEE, AND JOHN P. LEHOCZKY, MEMBER, IEEE

- Real-time operating systems use priority inheritance protocol to ensure that low-priority processes do not block high-priority processes
- Several operating systems implement the protocol incorrectly, leading to deadlocks and priority inversion

Code-level verification should complement reasoning on the algorithm/design level

Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications

Ion Stoica; Robert Morris, David Karger, M. Frans Kaashoek, Hari Balakrishnan†
MIT Laboratory for Computer Science
chord@lcs.mit.edu
<http://pdos.lcs.mit.edu/chord/>

- Chord is a distributed hash table developed at MIT

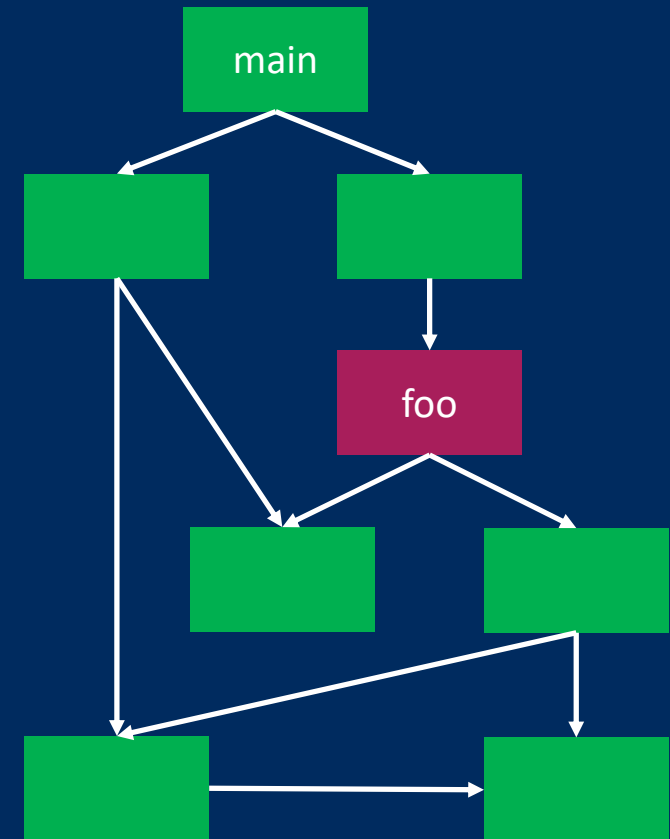
Three features that distinguish Chord from many other peer-to-peer lookup protocols are its simplicity, **provable correctness**, and provable performance.

- None of the seven properties claimed invariant of the original version is actually an invariant

Reasoning must be supported by tools

Modular Verification

- Verify each method separately
 - Scalability
- Do not use the implementation of callees
 - Software evolution
 - Dynamic method binding
- Do not use the implementation of callers and other methods
 - Correctness guarantees for libraries
 - Software evolution




Outline

- Permission-based Verification
- The Viper Intermediate Language
- Building Verifiers
- Encoding of Advanced Verification Techniques

Contracts

- Contracts specify the intended behavior of parts of the program
- For the verification of a method, use the contracts of the rest of the program, not the implementation
- Verify calls in terms of method pre- and postconditions



Example: Contracts

```
class Account {  
  var bal: int  
  
  method deposit(amount: int)  
    requires 0 < amount  
    ensures bal == old(bal) + amount  
  { ... }  
  
  method withdraw(amount: int)  
    requires 0 < amount && amount <= bal  
    ensures bal == old(bal) - amount  
  { ... }  
}
```

```
method demo(a: Account)  
  requires 0 <= a.bal  
{  
  a.deposit(200);  
  a.withdraw(100);  
}
```


Example: Side Effects

```
class Account {  
  var bal: int  
  
  method deposit(amount: int)  
    requires 0 < amount  
    ensures bal == old(bal) + amount  
  { ... }  
}
```

```
method demo(a: Account, l: List)  
  ensures l.len == old(l.len)  
{  
  a.deposit(200)  
}
```

Example: Side Effects

```
class Account {  
  var bal: int  
  var transactions: List  
  
  method deposit(amount: int)  
    requires 0 < amount  
    ensures bal == old(bal) + amount  
    { transactions.add(amount) ... }  
  
  method getTransactions() returns (t: List)  
    { t := transactions }  
}
```

```
method demo(a: Account, l: List)  
  ensures l.len == old(l.len)  
{  
  a.deposit(200)  
}
```

```
demo(a, a.getTransactions())
```

The Frame Problem



$$\begin{array}{c}
 \{P\} \ S \ \{Q\} \\
 FV(\{P\}) \cap FV(\{Q\}) = \{\} \\
 \hline
 \{P, R\} \ S \ \{Q, R\}
 \end{array}$$

Footprints



$$\frac{\{P\} S \{Q\}}{\{P \wedge R\} S \{Q \wedge R\}}$$

$\text{footprint}(S) \cap \text{footprint}(R) = \{\}$

Separation Logic


- Heap properties are specified via points-to assertions

$x.f \rightarrow v$


- Holds in a partial heap that maps the memory location $x.f$ to value v

- Each heap access to $x.f$ requires the current partial heap to contain $x.f$

$\{ x.f \rightarrow _ \} x.f := v \{ x.f \rightarrow v \}$



Footprints in Separation Logic



$$\frac{\{P\} \ S \ \{Q\}}{\{P \wedge R\} \ S \ \{Q \wedge R\}}$$

Separation and Framing

- Composition of partial heaps is described using separating conjunction

$P * R$

- Holds in a partial heap if it can be split into two disjoint partial heaps, in which P and Q hold
- $x.f \rightarrow _ * x.f \rightarrow _$ is equivalent to false
- $x.f \rightarrow _ * y.f \rightarrow _$ implies $x \neq y$

- Frame rule

$$\frac{\{P\} S \{Q\}}{\{P * R\} S \{Q * R\}}$$

$P \mid R$



Example: Framing

```
class Account {  
  var bal: int  
  
  method deposit(amount: int)  
    requires this.bal → B * 0 < amount  
    ensures this.bal → B + amount  
  { ... }  
}
```


```
method demo(a: Account, l: List)  
  requires a.bal → B * l.len → L  
  ensures l.len → L  
  { a.deposit(200) }
```



Example: Framing

```
class Account {  
  var bal: int  
  var ta: List  
  
  method deposit(amount: int)  
    requires this.bal → B * 0 < amount *  
             this.ta → T * this.ta.len → TL  
    ensures ...  
    { ta.add(amount) ... }  
}
```


```
method demo(a: Account, l: List)  
  requires a.bal → B * l.len → L *  
           a.ta → T * a.ta.len → TL  
  ensures l.len → L  
  { a.deposit(200) }
```



Recall $x.f \rightarrow _ * y.f \rightarrow _ \text{ implies } x \neq y$

```
demo(a, a.getTransactions())
```

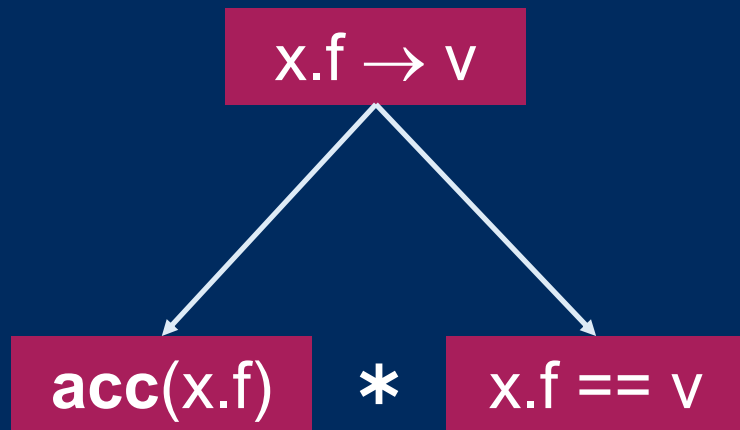
Parallel Composition



Disjointness of footprints ensures data-race freedom

$$\frac{\{P_1\} S_1 \{Q_1\} \quad \{P_2\} S_2 \{Q_2\}}{\{P_1 * P_2\} S_1 \parallel S_2 \{Q_1 * Q_2\}}$$

Implicit Dynamic Frames



```
method demo(a: Account, l: List)
  requires a.bal  $\rightarrow$  B * l.len  $\rightarrow$  L
  ensures l.len  $\rightarrow$  L
  { a.deposit(200) }
```

```
method demo(a: Account, l: List)
  requires acc(a.bal) * acc(l.len)
  ensures acc(l.len) * l.len == old(l.len)
  { a.deposit(200) }
```




- Assertions must be self-framing

Total-Heap Semantics

- Partial-heap semantics is not suitable

$$\mathbf{acc}(x.f) * x.f == v$$

- Define semantics relative to a total heap and permission mask
 - $x.f == v$ holds if $\text{heap}(x.f)$ yields v
 - $\mathbf{acc}(x.f)$ holds if $\text{permission mask}(x.f)$ yields true
 - $P * Q$ holds if the mask can be split into two compatible masks, in which P and Q hold




Ownership Transfer


```
class Account {  
  var bal: int  
  
  method deposit(amount: int)  
    requires acc(this.bal) * ...  
    ensures acc(this.bal) * ...  
  { ... }  
}
```

V'

```
method demo(a: Account, l: List)  
  requires acc(a.bal) * acc(l.len)  
  ensures acc(l.len) * l.len == old(l.len)
```



a.deposit(200)



Fractional Permissions

- Permissions can be split and recombined

$\text{acc}(x.f, 1/2)$


- Read access requires a non-zero permission
- Write access requires full permission

- Separating conjunction adds permissions

$\text{acc}(x.f, 1/2) * \text{acc}(x.f, 1/2) \equiv \text{acc}(x.f)$

Summary

- Modularity is important for scalability, components, and evolution
- Contracts enable modular verification
- Permissions
 - provide a solution to the frame problem



- enable the verification of concurrent programs

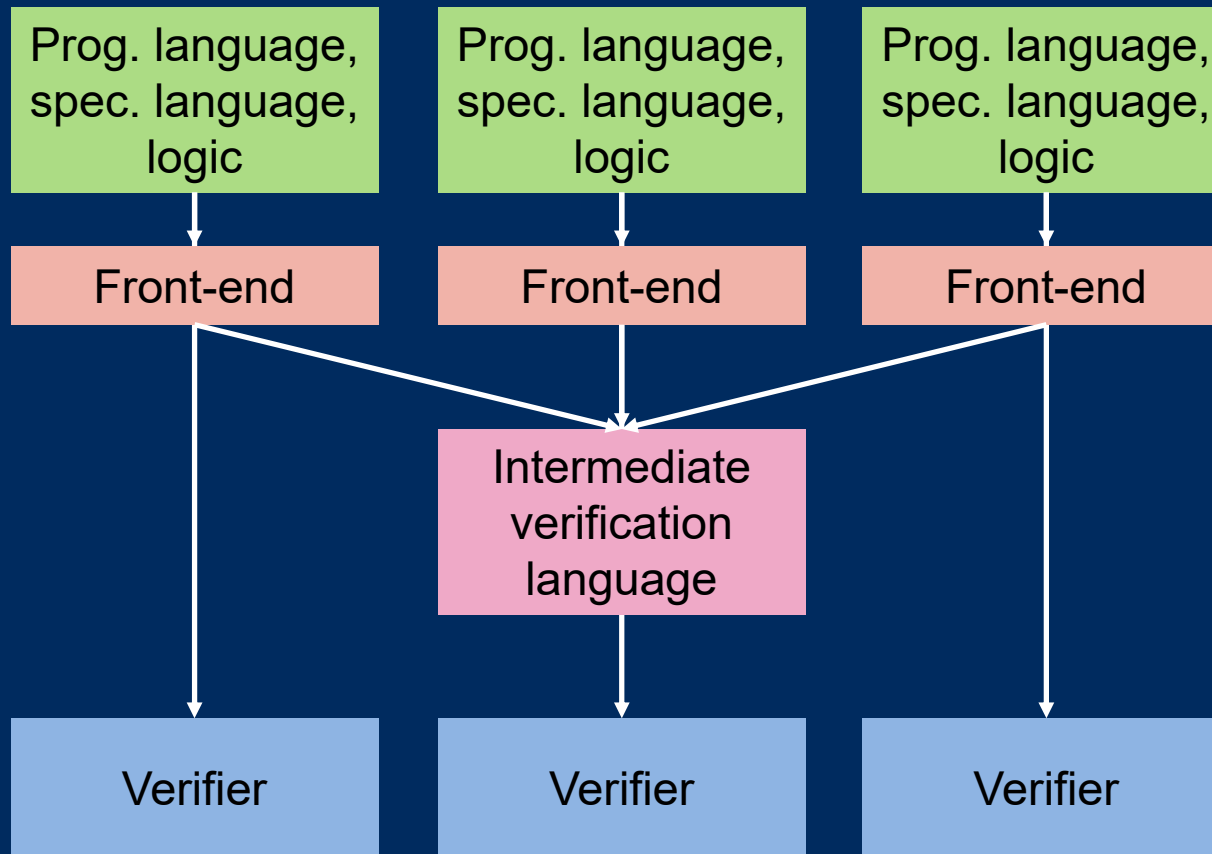
$$\frac{\{P\} S \{Q\}}{\{P * R\} S \{Q * R\}}$$

$$\frac{\{P_1\} S_1 \{Q_1\} \quad \{P_2\} S_2 \{Q_2\}}{\{P_1 * P_2\} S_1 \parallel S_2 \{Q_1 * Q_2\}}$$

Chalice



- Home page: www.pm.inf.ethz.ch/research/chalice.html
- Try online: www.rise4fun.com/Chalice
- Download: chalice.codeplex.com



Outline


- Permission-based Verification
- The Viper Intermediate Language
- Building Verifiers
- Encoding of Advanced Verification Techniques

Ownership Transfer


```
class Account {  
  var bal: int  
  
  method deposit(amount: int)  
    requires acc(this.bal) * ...  
    ensures acc(this.bal) * ...  
  { ... }  
}
```

V'

```
method demo(a: Account, l: List)  
  requires acc(a.bal) * acc(l.len)  
  ensures acc(l.len) * l.len == old(l.len)
```



a.deposit(200)



Ownership Transfer

$\{P\} \text{ method } m \{Q\}$
 $\{P\} \text{ e.m() } \{Q\}$
 $\{P * R\} \text{ e.m() } \{Q * R\}$

$\{P_1\} S_1 \{Q_1\} \quad \{P_2\} S_2 \{Q_2\}$
 $\{P_1 * P_2\} S_1 \parallel S_2 \{Q_1 * Q_2\}$
 $\{P_1 * P_2 * R\} S_1 \parallel S_2 \{Q_1 * Q_2 * R\}$

Inhale and Exhale

▪ inhale A means:


- obtain all permissions required by A
- assume all logical constraints

▪ exhale A means:


- assert all logical constraints
- check and remove all permissions required by A
- havoc any locations to which all permission is lost

▪ Analogues of **assume** and **assert**

inhale $\text{acc}(x.f) * x.f == v$



exhale $\text{acc}(x.f) * x.f == v$



Encoding Ownership Transfer

$\frac{\frac{\{P\} \text{ method } m \{Q\}}{\{P\} \text{ e.m() } \{Q\}}}{\{P * R\} \text{ e.m() } \{Q * R\}}$

exhale P
inhale Q

$\frac{\frac{\{P_1\} S_1 \{Q_1\} \quad \{P_2\} S_2 \{Q_2\}}{\{P_1 * P_2\} S_1 \parallel S_2 \{Q_1 * Q_2\}}}{\{P_1 * P_2 * R\} S_1 \parallel S_2 \{Q_1 * Q_2 * R\}}$

exhale P₁
exhale P₂
inhale Q₁
inhale Q₂


```
method demo(a: Account, l: List)
  requires acc(a.bal) * acc(l.len)
  ensures acc(l.len) * l.len == old(l.len)
{
  a.deposit(200)
}
```

```
method deposit(amount: int)
  requires acc(this.bal)
  ensures acc(this.bal)
```

```
inhale acc(a.bal) * acc(l.len)
exhale acc(a.bal)
inhale acc(a.bal)
exhale acc(l.len) * l.len == old(l.len)
```

Encoding Monitors

```
class Account {  
  var bal: int  
  
  invariant acc(this.bal)  
  
  method deposit(amount: int)  
  {  
    acquire this  
    this.bal := this.bal + amount  
    release this  
  }  
}
```



```
inhale acc(this.bal)  
this.bal := this.bal + amount  
exhale acc(this.bal)
```



Abstraction

```
method demo(a: Account, l: List)
  requires acc(a.bal) * acc(l.len)
  ensures acc(l.len) * l.len == old(l.len)
{
  a.deposit(200)
}
```

Mentioning field names in contracts:

- Violates information hiding
- Cannot express access to a statically-unknown set of locations

Recursive Predicates




```
predicate list(this: Ref) {  
  acc(this.next) * acc(this.data) *  
  (this.next != null ==> list(this.next))  
}
```

- Predicate instances are manipulated similarly to access permissions
 - Ownership is transferred via inhale and exhale
 - But $P(x) * P(x)$ is **not** equivalent to false

```
inhale list(a)  
exhale list(a)
```

```
predicate P(this: Ref)  
{ acc(this.f, 1/3) }
```

- Folding and unfolding is done manually



```
unfold list(l)  
l.data := 7  
fold list(l)
```


Examples: Recursive Predicates

```
predicate lseg(this: Ref, last: Ref) {  
  this != last ==>  
    acc(this.next) * acc(this.data) *  
    lseg(this.next, last)  
}
```

```
predicate list(this: Ref) {  
  acc(this.next) * acc(this.data) *  
  (this.next != null ==> list(this.next)) *  
  0 <= this.data  
}
```


```
predicate list(this: Ref) {  
  acc(this.next) * acc(this.data) *  
  (this.next != null ==> list(this.next) *  
    this.data <= unfolding list(this.next) in this.next.data)  
}
```

Limitations




- Extending footprints

Limitations



- Extending footprints
- Sharing


Limitations



- Extending footprints
- Sharing
- Traversal order

Quantified Permissions

```
predicate list( nodes: Set[ Ref ] ) {  
  forall n: Ref :: n in nodes ==>  
    acc(n.next) *  
    n.next in nodes  
}
```



```
list(nodes) *  
v in nodes * w in nodes
```


Abstraction

```
method demo(a: Account, l: List)
  requires acc(a.bal) * acc(l.len)
  ensures acc(l.len) * l.len == old(l.len)
{
  a.deposit(200)
}
```

Mentioning field names in contracts:

- Violates information hiding
- Cannot express access to a statically-unknown set of locations

Heap-Dependent Functions

- Predicates abstract over permissions
- Functions abstract over expressions

```
predicate list(this: Ref) {  
  acc(this.next) * acc(this.data) *  
  (this.next != null ==> list(this.next))  
}
```

```
function length(this: Ref): Int  
  requires list(this)  
{  
  unfolding list(this) in  
  this.next == null ? 1 : 1 + length(this.next)  
}
```

Abstraction

```
method demo(a: Account, l: List)
  requires account(a) * list(l)
  ensures list(l) * length(l) == old(length(l))
{
  a.deposit(200)
}
```

- Predicates and functions need not have definitions

Function Framing



- Heap-dependent functions are mathematical functions of their arguments and their footprint

```
function length(this: Ref): Int  
  requires list(this)  
{  
  unfolding list(this) in  
  this.next == null ? 1 : 1 + length(this.next)  
}
```

Summary

- Intermediate verification languages facilitate the development of program verifiers
- Inhale and exhale primitives express ownership transfer
- Predicates and functions abstract over permissions and values

exhale `acc(x.f) * x.f == v`

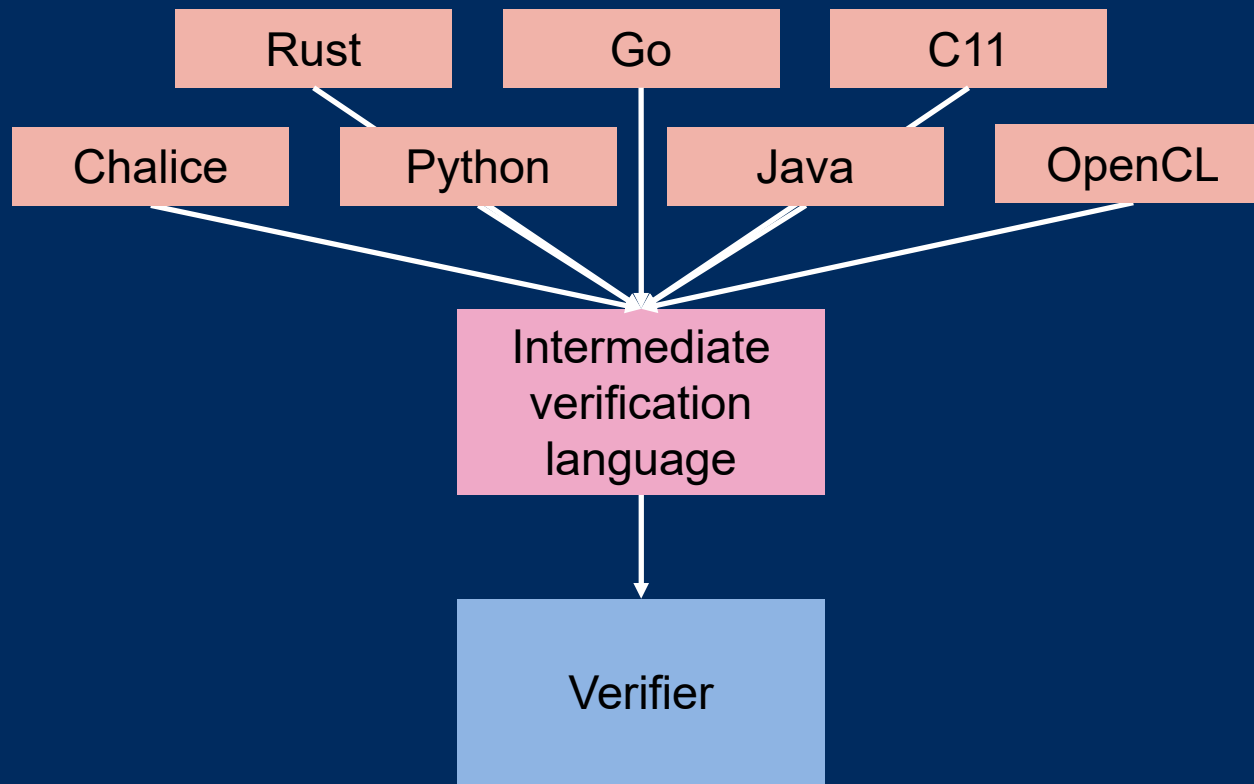


function `length(this: Ref): Int`
requires `list(this)`

Viper

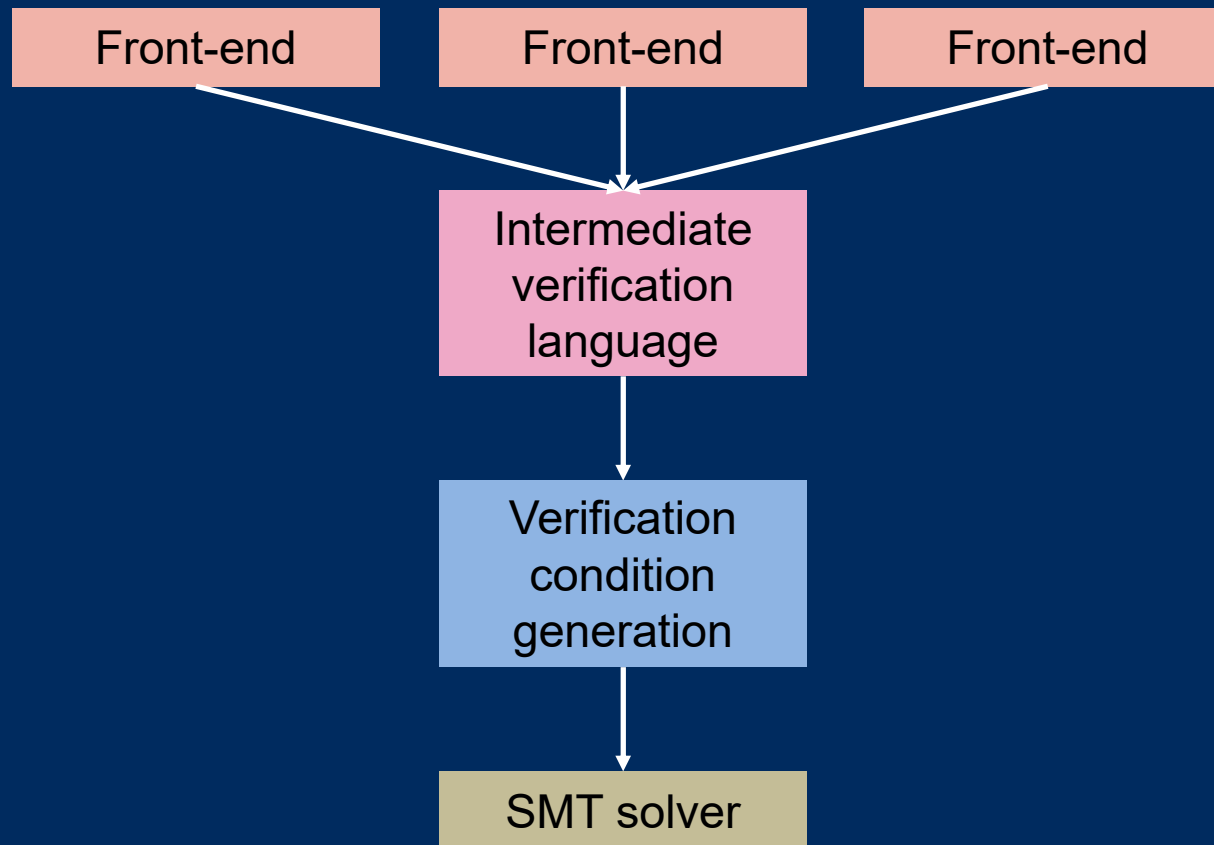
The logo for Viper features a stylized snake silhouette in a dark blue color, positioned behind the word "VIPER" which is written in large, bold, light blue capital letters. The snake's head is positioned over the letter 'V', and its body curves around the letters 'I', 'P', and 'E'. The letter 'R' is partially obscured by the snake's tail.

- Home page: viper.ethz.ch
- Try online: viper.ethz.ch/examples
- Download: bitbucket.org/viperproject
- Download a version for the tutorial on Friday



Outline

- Permission-based Verification
- The Viper Intermediate Language
- Building Verifiers
- Encoding of Advanced Verification Techniques



Verification Condition Generation

- For a given program, verification condition generation computes a logical formula whose validity implies the correctness of the program
- Verification condition reflects semantics of the program and its specification
- We will compute verification conditions in two steps:
 - Encode the Viper program into guarded commands
 - Compute weakest preconditions of guarded-commands programs

Guarded Commands

```
S ::= x := E
    | S1; S2
    | S1 □ S2
    | assert P
    | assume P
    | havoc x
```

- Assertions P are first-order logic formulas
 - Including quantifiers, uninterpreted functions, arithmetic, etc.

Encoding into Guarded Commands

```
method abs(a: Int) returns (res: Int)
  ensures 0 <= res
{
  if(0 <= a) { res := a }
  else      { res := -a }
}
```

```
(
  assume 0 <= a
  res := a
□
  assume a < 0
  res := -a
)
assert 0 <= res
```


Weakest Preconditions

$$\begin{aligned} \text{wp}(x := E, Q) &\equiv Q[E/x] \\ \text{wp}(S_1; S_2, Q) &\equiv \text{wp}(S_1, \text{wp}(S_2, Q)) \\ \text{wp}(S_1 \square S_2, Q) &\equiv \text{wp}(S_1, Q) \wedge \text{wp}(S_2, Q) \\ \text{wp}(\text{assert } P, Q) &\equiv P \wedge Q \\ \text{wp}(\text{assume } P, Q) &\equiv P \Rightarrow Q \\ \text{wp}(\text{havoc } x, Q) &\equiv \forall x \bullet Q \end{aligned}$$

To verify statement S , prove that $\text{wp}(S, \text{true})$ holds

Heap Model

- Define semantics relative to a total heap and permission mask
 - $x.f == v$ holds if $\text{heap}(x,f)$ yields v
 - $\text{acc}(x.f)$ holds if $\text{permission mask}(x,f)$ yields true



- Model heap as total map

$$\text{Ref} \times \text{Field}\langle T \rangle \rightarrow T$$

- Model permission mask as total map

$$\text{Ref} \times \text{Field}\langle T \rangle \rightarrow \text{Bool}$$

Encoding inhale

```
 $\langle \text{inhale } e \rangle \equiv \text{assume } \langle e \rangle$   
 $\langle \text{inhale acc}(e.f) \rangle \equiv \text{assume } \langle e \rangle \neq \text{null}$   
 $\text{assume } \neg \text{mask}(\langle e \rangle, f)$   
 $\text{mask}(\langle e \rangle, f) := \text{true}$   
 $\langle \text{inhale } a_1 * a_2 \rangle \equiv \langle \text{inhale } a_1 \rangle; \langle \text{inhale } a_2 \rangle$ 
```

- $\langle _ \rangle$ is the encoding function
- a is an assertion
- e is an expression (not containing permissions)

inhale $\text{acc}(x.f) * x.f == v$

inhale $\text{acc}(x.f) * \text{acc}(y.f)$

assume $x \neq \text{null}$
assume $\neg \text{mask}(x,f)$
 $\text{mask}(x,f) := \text{true}$
assume $\text{heap}(x,f) == v$

assume $x \neq \text{null}$
assume $\neg \text{mask}(x,f)$
 $\text{mask}(x,f) := \text{true}$
assume $y \neq \text{null}$
assume $\neg \text{mask}(y,f)$
 $\text{mask}(y,f) := \text{true}$

Encoding exhale

$\langle\langle \text{exhale } e \rangle\rangle \equiv \text{assert } \langle e \rangle$

$\langle\langle \text{exhale acc}(e.f) \rangle\rangle \equiv \text{assert mask}(\langle e \rangle, f)$
 $\text{mask}(\langle e \rangle, f) := \text{false}$

$\langle\langle \text{exhale } a_1 * a_2 \rangle\rangle \equiv \langle\langle \text{exhale } a_1 \rangle\rangle; \langle\langle \text{exhale } a_2 \rangle\rangle$

$\langle \text{exhale } a \rangle \equiv \langle\langle \text{exhale } a \rangle\rangle$

$h := \text{heap}$

havoc heap

assume $\forall x, f \bullet \text{mask}(x, f) \Rightarrow \text{heap}(x, f) == h(x, f)$

exhale $\text{acc}(x.f) * x.f == v$

```
assert mask(x,f)
mask(x,f) := false
assert heap(x,f) == v
// havoc heap(x,f)
```

exhale $\text{acc}(x.f) * \text{acc}(y.f)$

```
assert mask(x,f)
mask(x,f) := false
assert mask(y,f)
mask(y,f) := false
// havoc heap(x,f), heap(y,f)
```

Predicates

- For each predicate, introduce a function that maps a predicate instance to a field name

```
predicate lseg(this: Ref, last: Ref)
```

```
lsegField: Ref × Ref → Field<Int>
```

- Use this field to index mask

```
inhale lseg(a, b)
```

```
mask(null, lsegField(a, b)) := true
```

Heap-Dependent Functions

- Encode Viper function as mathematical function

```
function balance(this: Ref): Int  
requires acc(this.bal)
```

```
balance: Ref × Int → Int
```

- Function applications evaluate footprint

```
balance(x)
```

```
balance(x, heap(x.bal))
```

Heap-Dependent Functions

- Encode Viper function as mathematical function

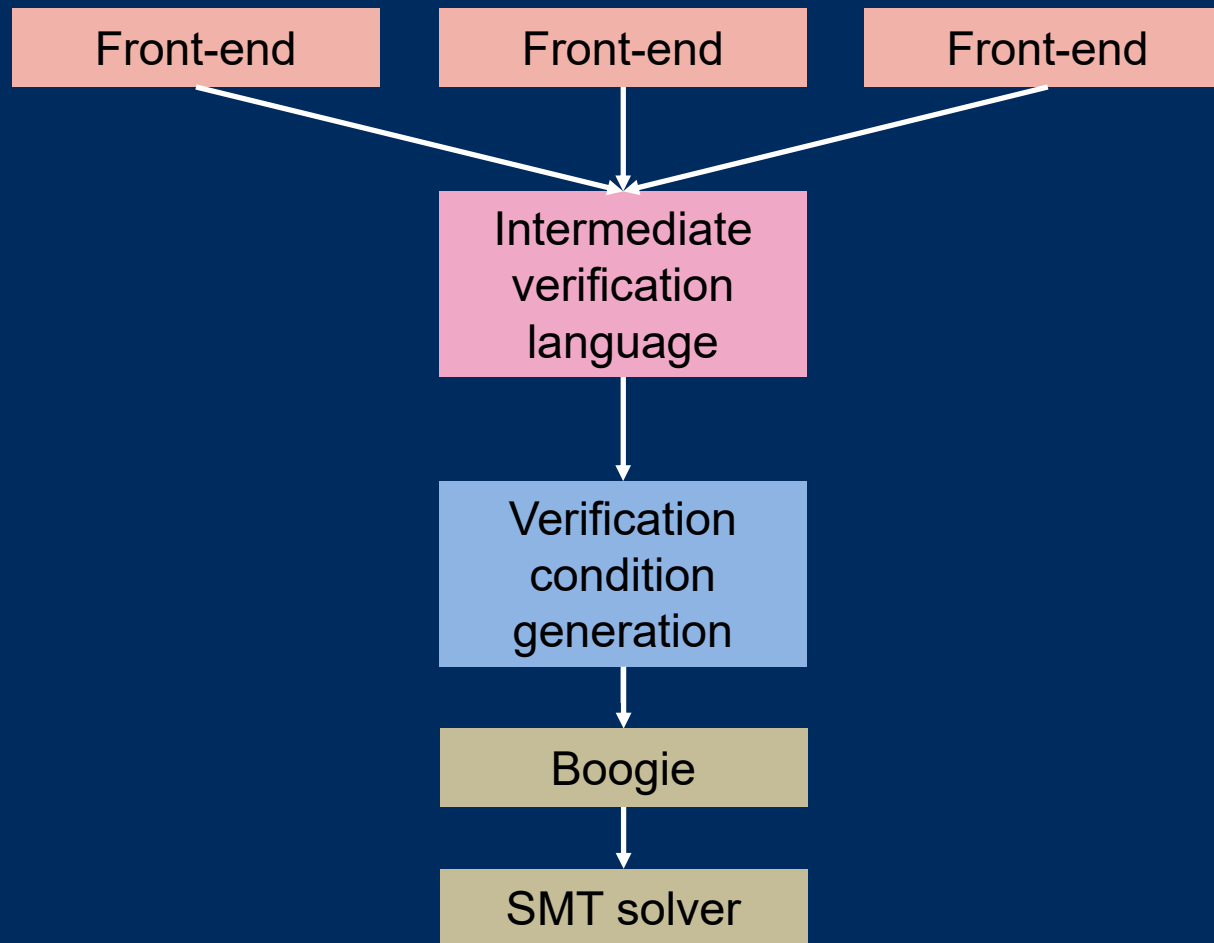
```
function length(this: Ref): Int  
requires list(this)
```

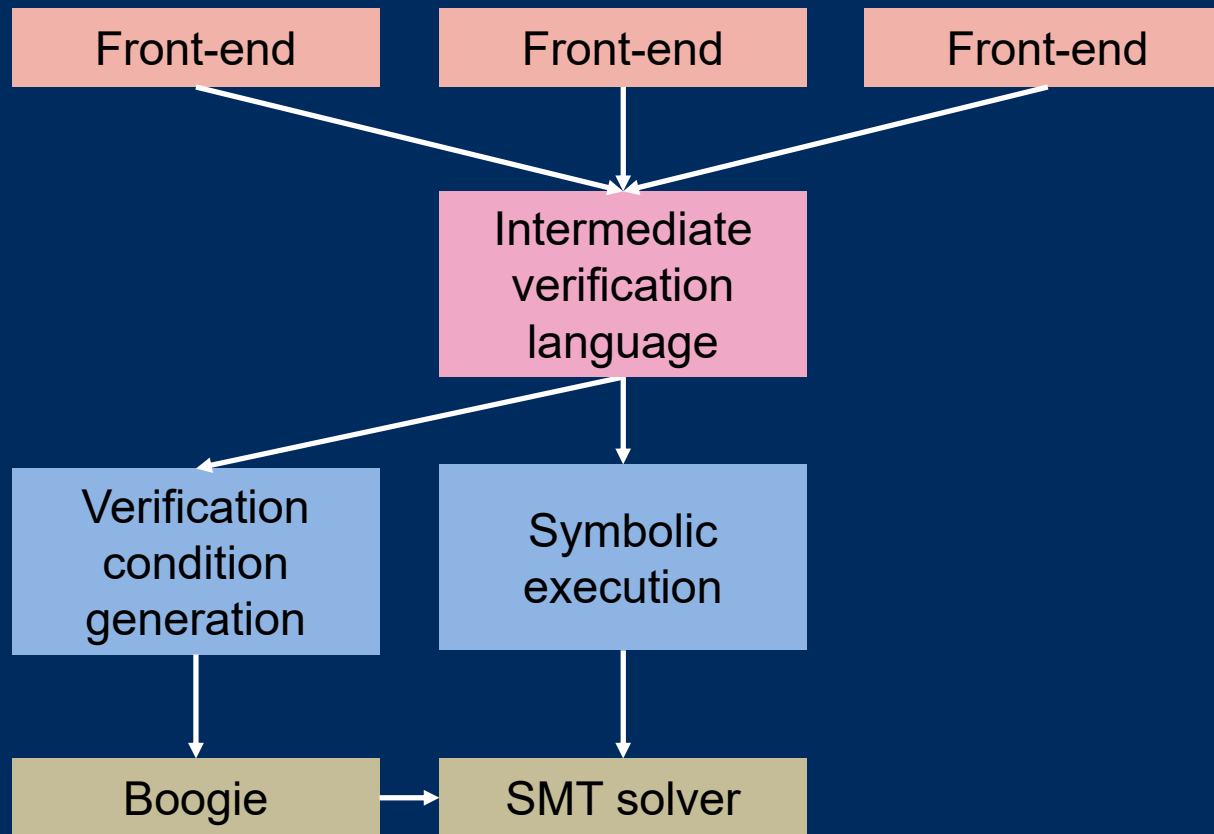
```
balance: Ref × Int → Int
```

- Give predicate instances a version number
 - Changes each time the predicate's footprint may change

```
length(x)
```

```
length(x, heap(null, listField(x)))
```





Symbolic Execution

- Symbolic execution simulates the execution of a program statically, using symbolic rather than concrete values
 - Introduce symbolic variables to represent inputs
- A configuration consists of
 - The statement to be executed
 - An environment, mapping program variables to expressions over the symbolic variables
 - A path condition, a logic formula representing information about the symbolic variables


```
method abs(a: Int) returns (res: Int)
```

```
  ensures 0 <= res
```

```
{  
  if(0 <= a) { res := a }  
  else      { res := -a }  
}
```

Statement

Environment

Path Condition

a res

```
if(0 <= a) { res := a }  
else      { res := -a }  
assert 0 <= res
```

A R true

```
res := a;  
assert 0 <= res
```

A R $0 \leq A$

```
assert 0 <= res
```

A A $0 \leq A$

check $0 \leq A \Rightarrow 0 \leq A$



```
res := -a;  
assert 0 <= res
```

A R $A < 0$

```
assert 0 <= res
```

A -A $A < 0$

check $A < 0 \Rightarrow 0 \leq -A$



Algorithm

```
 $\sigma := \{ (s; \text{stop}, \text{env}_0, \text{true}) \}$   
while  $\sigma \neq \{ \}$  do  
   $\gamma := \text{take}(\sigma)$   
   $r := \text{step}(\gamma)$   
  if  $r = \perp$  then return failure  
   $\sigma := \sigma \cup r$   
end  
return success
```

- s is the program to be verified
- stop is an artificial stop-marker
- env_0 maps each program variable to a fresh symbolic variable

Algorithm


$$\begin{aligned} \text{step}(x := e; s, \text{env}, \pi) &= \{ (s, \text{env}[x := \langle e \rangle], \pi) \} \\ \text{step}(\text{if}(e) \{ s_1 \} \text{ else } \{ s_2 \}; s, \text{env}, \pi) &= \{ (s_1; s, \text{env}, \pi \wedge \langle e \rangle), (s_2; s, \text{env}, \pi \wedge \neg \langle e \rangle) \} \\ \text{step}(\text{assert } e; s, \text{env}, \pi) &= \text{if check}(\pi \Rightarrow \langle e \rangle) \text{ then } \{ (s, \text{env}, \pi) \} \text{ else } \perp \\ \text{step}(\text{stop}, \text{env}, \pi) &= \{ \} \end{aligned}$$

- $\langle _ \rangle$ symbolically evaluates an expression to a symbolic expression
- check is a query to the SMT solver

Heap Model

- The symbolic heap models the partial-heap semantics

$\text{acc}(x.f)$



- Model heap as set of heap chunks (E_r and E_v are symbolic expressions)

$E_r.f \rightarrow E_v$

- Extend configurations by a symbolic heap

Executing inhale

$$\text{step}(\text{inhale } e; s, \text{env}, \pi, h) \equiv \\ \{ (s, \text{env}, \pi \wedge \langle e \rangle, h) \}$$
$$\text{step}(\text{inhale } \text{acc}(e.f); s, \text{env}, \pi, h) \equiv \\ \{ (s, \text{env}, \pi \wedge \langle e \rangle \neq \text{null} \wedge \bigwedge_{E.f \rightarrow _ \in h} E \neq \langle e \rangle, h \cup \{ \langle e \rangle.f \rightarrow V \}) \}$$
$$\text{step}(\text{inhale } a_1 * a_2; s, \text{env}, \pi, h) \equiv \\ \text{step}(\text{inhale } a_1; \text{inhale } a_2; s, \text{env}, \pi, h)$$

Statement	Environment	Path Condition	Heap
-----------	-------------	----------------	------

	X V		
--	--------	--	--

inhale acc(x.f) inhale x.f == v	X V	true	{ }
--	--------	------	-----

inhale x.f == v	X V	true	{ X.f → W }
------------------------	--------	------	-------------

	X V	W = V	{ X.f → W }
--	--------	-------	-------------

Statement	Environment	Path Condition	Heap
-----------	-------------	----------------	------

	X Y		
--	--------	--	--


inhale acc(x.f) inhale acc(y.f)	X Y	true	{ }
--	--------	------	-----

inhale acc(y.f)	X Y	true	{ X.f → W }
------------------------	--------	------	-------------

	X Y	X ≠ Y	{ X.f → W, Y.f → Z }
--	--------	-------	----------------------

Executing exhale

```
inhale acc(x.f) * x.f == v
exhale acc(x.f) * x.f == v
```

Environment		Path Condition	Heap
x	v		
X	V	true	{ }
X	V	true	{ X.f → W }
X	V	W = V	{ X.f → W }
X	V	W = V	{ } 

- Solution: evaluate expressions in the heap before the exhale

Executing exhale

$\text{step}(\text{exhale } a;s, \text{env}, \pi, h) \equiv \text{step}'(\text{exhale } a;s, \text{env}, \pi, h, h)$

$\text{step}'(\text{exhale } e;s, \text{env}, \pi, h, h_0) \equiv$
if $\text{check}(\pi \Rightarrow \langle e \rangle_0)$ then $\{ (s, \text{env}, \pi, h) \}$ else \perp

$\text{step}'(\text{exhale } \text{acc}(e.f);s, \text{env}, \pi, h, h_0) \equiv$
if $\langle e \rangle_0.f \rightarrow _ \in h$ then $\{ (s, \text{env}, \pi, h \setminus \{ \langle e \rangle_0.f \rightarrow _ \}) \}$ else \perp

$\text{step}'(\text{exhale } a_1 * a_2;s, \text{env}, \pi, h, h_0) \equiv$
let $\gamma = \text{step}'(\text{exhale } a_1;\text{exhale } a_2;s, \text{env}, \pi, h, h_0)$ in
if $\gamma = \perp$ then \perp else $\text{step}'(\gamma, h_0)$

Predicates

- Predicate instances are represented by heap chunks of the form

$$p(E_1, \dots, E_n) \rightarrow E_v$$

where E_v is a symbolic expression denoting the instance's footprint

```
predicate list(this: Ref) {  
  acc(this.next) * acc(this.data) *  
  (this.next != null ==> list(this.next))  
}
```

$$FP_{list} ::= (\text{Ref}, \text{Int}, FP_{list}) \mid \varepsilon$$

Heap-Dependent Functions

- Symbolic expressions may contain function applications

```
function balance(this: Ref): Int  
  requires acc(this.bal)
```

```
balance:  $\text{Ref} \times \text{Int} \rightarrow \text{Int}$ 
```

- Predicate footprints allow framing

```
function length(this: Ref): Int  
  requires list(this)
```

```
length:  $\text{Ref} \times \text{FP}_{\text{list}} \rightarrow \text{Int}$ 
```

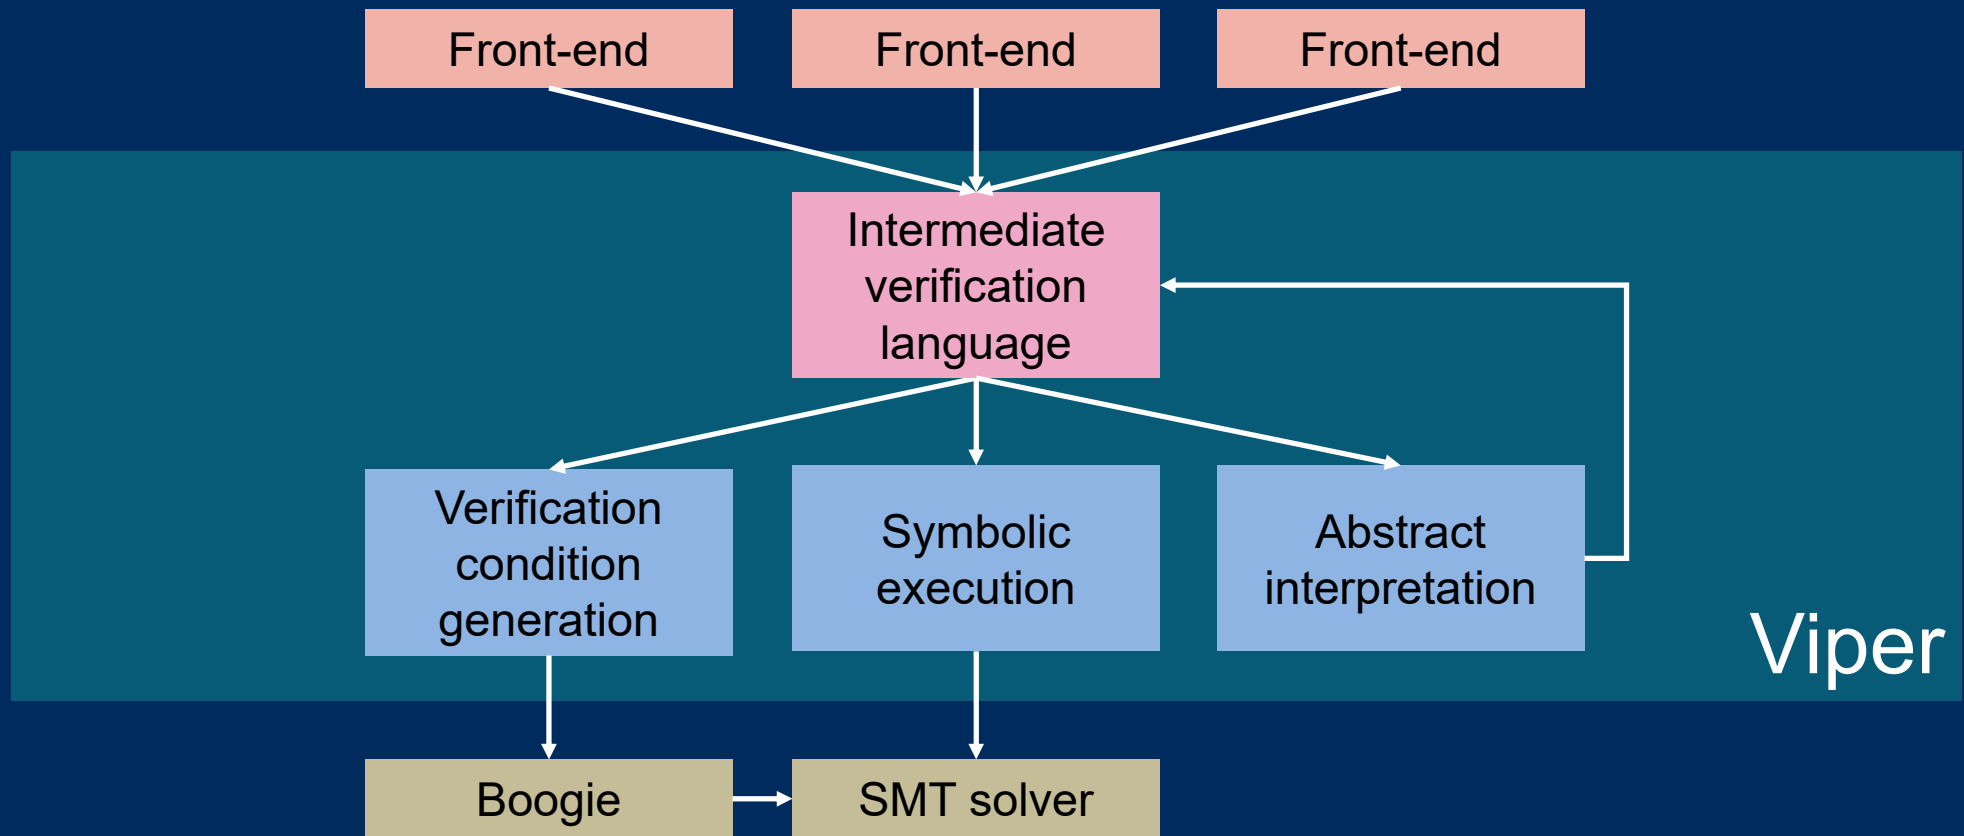
Pros and Cons of Symbolic Execution

Pros

- Small prover queries lead to better, more predictable performance
- Higher control over proof search enables dedicated algorithms

Cons

- Separation of path and heap information can lead to incompleteness
- Internal treatment of heap reasoning requires dedicated algorithms



Summary

- SMT solvers can be used to automate permission-based verification
- Verification condition generation encodes total-heap semantics
- VCG can be implemented as an encoding into existing languages and tools
- Symbolic execution implements partial-heap semantics
- Performs heap-reasoning internally, and uses SMT solver to reason about value information

Outline

- Permission-based Verification
- The Viper Intermediate Language
- Building Verifiers
- Encoding of Advanced Verification Techniques

Reminder: Encoding Monitors

```
class Account {  
  var bal: int  
  
  invariant acc(this.bal)  
  
  method deposit(amount: int)  
  {  
    acquire this  
    this.bal := this.bal + amount  
    release this  
  }  
}
```

```
inhale acc(this.bal)  
this.bal := this.bal + amount  
exhale acc(this.bal)
```

Non-Blocking Data Structures

- Permissions ensures data race freedom
 - Monitors and other synchronization are used to transfer ownership between threads
- Non-blocking data structures can increase performance by allowing extra concurrency
 - Synchronization is done through atomic operations such as compare-and-swap (CAS)
 - Data races are permitted

```
typedef int SpinLock;  
  
void Lock(SpinLock* sl) {  
    while(CAS(sl, 0, 1))  
        ;  
}  
  
void UnLock(SpinLock* sl) {  
    *sl = 0;  
}
```


Weak Memory

- Modern hardware often does not provide sequentially consistent shared memory
- Weak memory permits behaviors that are not possible under sequential consistency
- However, data-race free programs have only sequentially consistent behaviors

```
        a = 0;  
        b = 0;  
  
a = 1;   ||   b = 1;  
print(b); ||   print(a);
```

Possible results:

under SC: 10, 01, 11

under WM: also 00

C11

- The C11 memory model provides several kinds of variables


- Non-atomic variables

- Data races are errors

- Atomic variables with release-write and acquire-read

- Writes and reads are synchronized

- Relaxed separation logic (RSL) supports some features of the C11 memory model




Non-Atomic Variables

- Permissions prevent data races on non-atomic variables

$$\{ \text{true} \} \ x = \text{alloc}_{\text{NA}}(v) \ \{ x \rightarrow v \}$$
$$\{ x \rightarrow _ \} \ *x = v \ \{ x \rightarrow v \}$$
$$\{ x \rightarrow V \} \ t = *x \ \{ x \rightarrow V \ * t = V \}$$

Release-Acquire



- Races on atomic variables are permitted
- Release-acquire can be seen as message passing
- Messages may transfer ownership to non-atomic variables

Reminder: Monitor Invariants


- Monitor invariant Q specifies an assertion that holds when monitor is not currently held
- Acquire transfers ownership of Q from monitor to thread
- Release transfers ownership of Q from thread to monitor

```
class Account {  
  var bal: int  
  
  invariant acc(this.bal)  
  
  method deposit(amount: int)  
  {  
    acquire this  
    this.bal := this.bal + amount  
    release this  
  }  
}
```

Location Invariants

- Location invariant $Q(v)$ specifies an assertion that holds when the location has value v
- Acquire-read of value v transfers ownership of $Q(v)$ from atomic variable to thread
- Release-write of value v transfers ownership of $Q(v)$ from thread to atomic variable

$$Q(v) \equiv \begin{cases} n \rightarrow 5 & \text{if } v = 1 \\ \text{true} & \text{otherwise} \end{cases}$$



Proof Rules

- Choose location invariant when allocating an atomic location

$$\{ Q(v) \} \ x = \text{alloc}_{RA}(v) \ \{ \text{Rel}_Q(x) * \text{Acq}_Q(x) \}$$

- Release-write gives up ownership

$$\{ \text{Rel}_Q(x) * Q(v) \} \ x.\text{store}(v) \ \{ \text{Rel}_Q(x) \}$$

- Acquire-read gains ownership

$$\{ \text{Acq}_Q(x) \} \ t = x.\text{load}() \ \{ Q(t) * \text{Acq}_Q(x) \}$$

Proof Outline

```
    { true }  
    n = allocNA(0);  
    { n → 0 }  
    a = allocRA(0);  
    { n → 0 * RelQ(a) * AcqQ(a) }
```

```
{ n → 0 * RelQ(a) }
```

```
  *n = 5;
```

```
{ n → 5 * RelQ(a) }
```

```
  a.store(1);
```

```
{ RelQ(a) }
```

||

```
{ AcqQ(a) }
```

```
if(a.load() == 1)
```

```
{ Q(1) * AcqQ(a) }
```

```
  assert *n == 5;
```

$$Q(v) \equiv \begin{cases} n \rightarrow 5 & \text{if } v = 1 \\ \text{true} & \text{otherwise} \end{cases}$$


Proof Rules

$$\{ \text{Acq}_Q(x) \} \ t = x.\text{load}() \ \{ Q(t) * \text{Acq}_{Q[t := \text{true}]}(x) \}$$

- Reading the same value more than once would duplicate permissions

$$Q(v) \equiv \begin{cases} n \rightarrow 5 & \text{if } v = 1 \\ \text{true} & \text{otherwise} \end{cases}$$

```
x = a.load();  
y = a.load();  
if(x == 1 && y == 1)  
    assert false;
```



Intermediate verification language

Verification condition generation


Symbolic execution

Abstract interpretation

Boogie

SMT solver

Viper

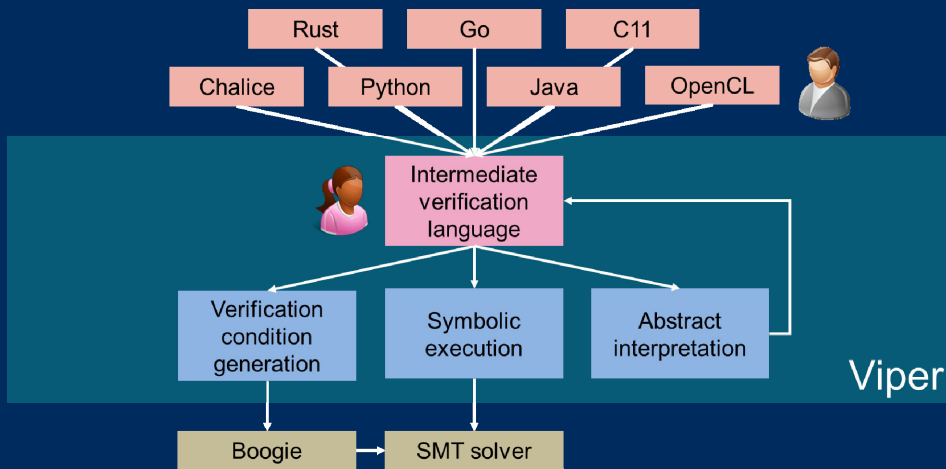


Modularity is important for scalability, components, and evolution

$$\frac{\{P\} S \{Q\}}{\{P * R\} S \{Q * R\}}$$

$$\frac{\{P_1\} S_1 \{Q_1\} \quad \{P_2\} S_2 \{Q_2\}}{\{P_1 * P_2\} S_1 \parallel S_2 \{Q_1 * Q_2\}}$$

Permissions enable framing and reasoning about concurrency



Intermediate languages enable reuse of infrastructure



Viper lets you encode a wide variety of reasoning techniques

References: Lecture 1

- J. C. Reynolds: Separation Logic: A Logic for Shared Mutable Data Structures. LICS 2002
- P. W. O'Hearn: Resources, Concurrency and Local Reasoning. CONCUR 2004
- J. Smans, B. Jacobs, F. Piessens: Implicit Dynamic Frames: Combining Dynamic Frames and Separation Logic. ECOOP 2009
- M. J. Parkinson, A. J. Summers: The Relationship between Separation Logic and Implicit Dynamic Frames. Logical Methods in Computer Science 8(3), 2012
- K. R. M. Leino, P. Müller: A Basis for Verifying Multi-threaded Programs. ESOP 2009

References: Lecture 2

- P. Müller, M. Schwerhoff, A. J. Summers: Viper: A Verification Infrastructure for Permission-Based Reasoning. VMCAI 2016
- K. R. M. Leino, P. Müller: A Basis for Verifying Multi-threaded Programs. ESOP 2009
- P. Müller, M. Schwerhoff, A. J. Summers: Automatic Verification of Iterated Separating Conjunctions using Symbolic Execution. CAV 2016
- M. J. Parkinson, G. M. Bierman: Separation logic and abstraction. POPL 2005

References: Lecture 3

- K. R. M. Leino, P. Müller: A Basis for Verifying Multi-threaded Programs. ESOP 2009
- S. Heule, I. T. Kassios, P. Müller, A. J. Summers: Verification Condition Generation for Permission Logics with Abstract Predicates and Abstraction Functions. ECOOP 2013
- M. Barnett, B.-Y. E. Chang, R. DeLine, B. Jacobs, K. R. M. Leino: Boogie: A Modular Reusable Verifier for Object-Oriented Programs. FMCO 2005
- J. Berdine, C. Calcagno, P. W. O'Hearn: Smallfoot: Modular Automatic Assertion Checking with Separation Logic. FMCO 2005
- J. Smans, B. Jacobs, F. Piessens: Heap-Dependent Expressions in Separation Logic. FMOODS/FORTE 2010
- I. T. Kassios, P. Müller, M. Schwerhoff: Comparing Verification Condition Generation with Symbolic Execution: An Experience Report. VSTTE 2012

References: Lecture 4

- A. J. Summers, P. Müller: Automating Deductive Verification for Weak-Memory Programs, TACAS 2018
- P. Müller, M. Schwerhoff, A. J. Summers: Viper: A Verification Infrastructure for Permission-Based Reasoning. VMCAI 2016
- V. Vafeiadis, C. Narayan: Relaxed separation logic: a program logic for C11 concurrency. OOPSLA 2013