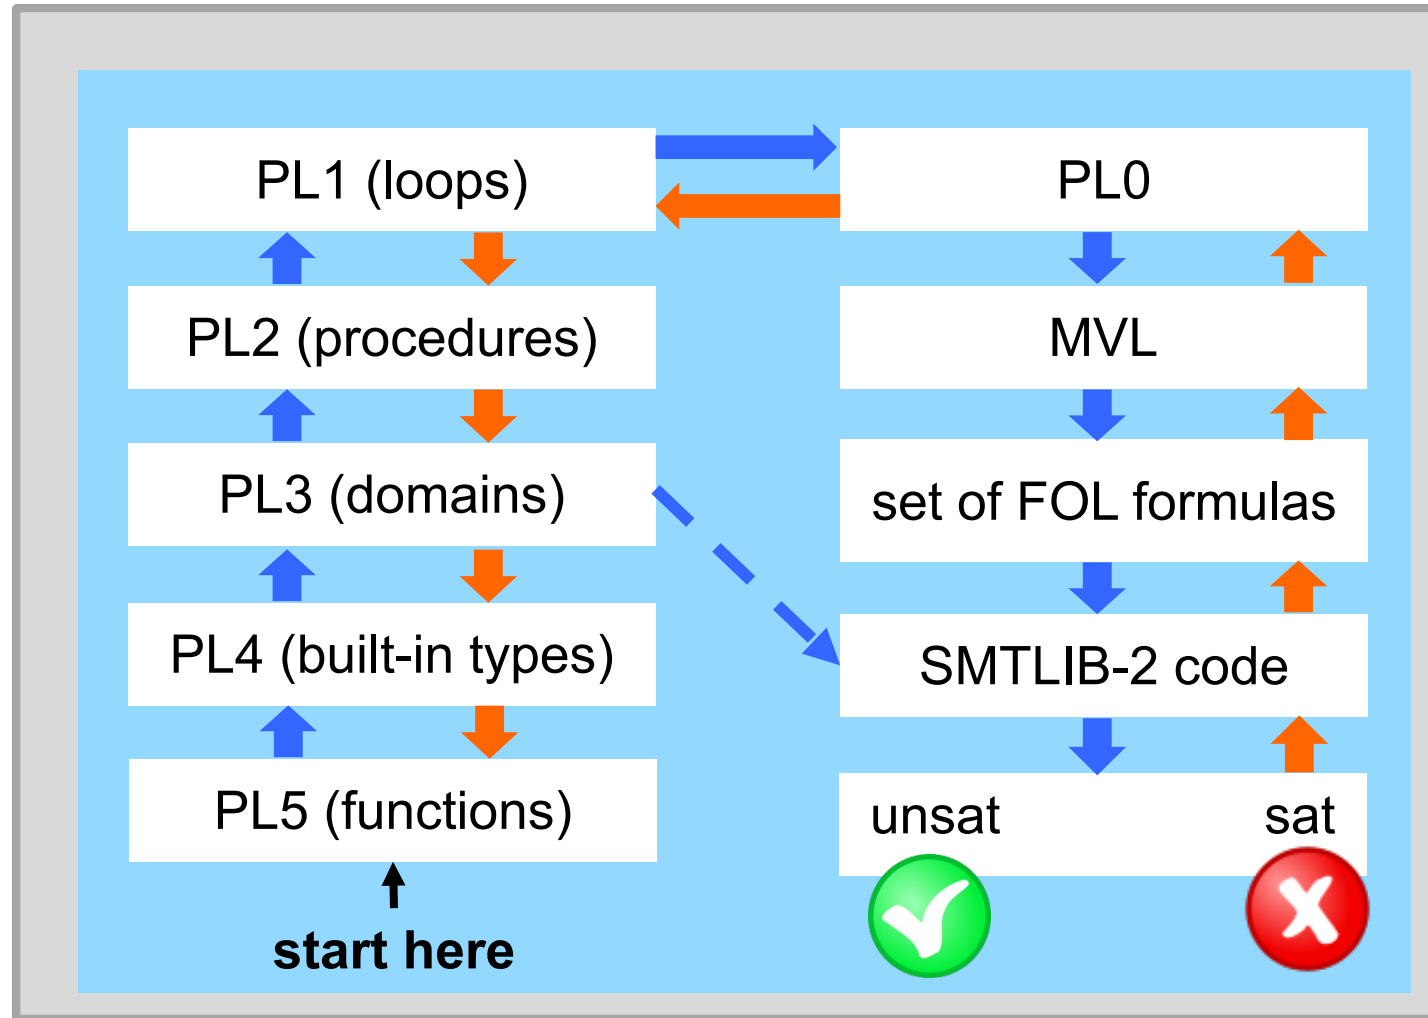


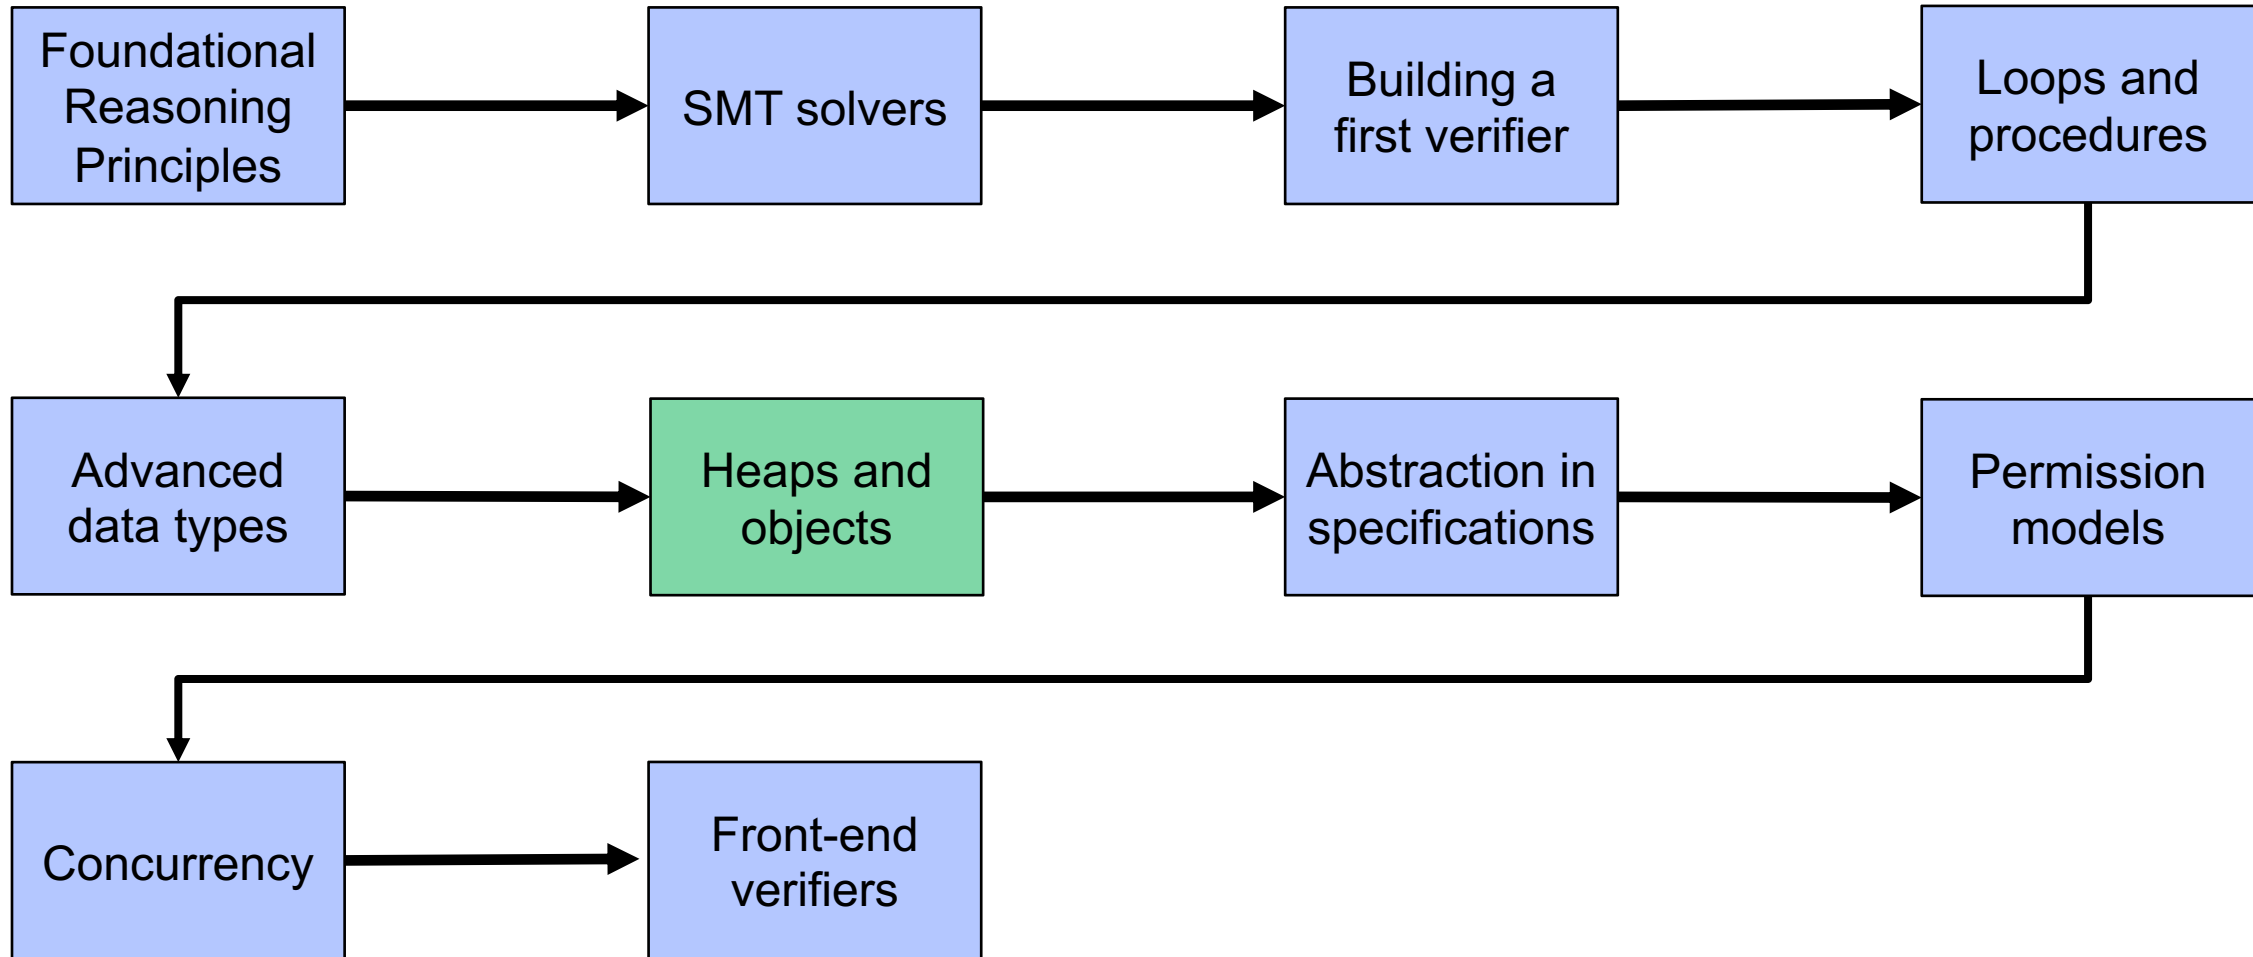
02245 – Module 7

# HEAPS AND OBJECTS

# Previously...



# Tentative course outline



# Why objects and heap-based data structures?

## ■ Static data structures

- Examples: arrays, all mathematical data structures from module 5
- Fixed size, stack-allocated
- Immutable, no memory reuse
- To update the data structure we create an updated copy

```
// static array A = [0,0,0]
A := cons(3, 0)

// create updated copy
B := set(A, 1, 17)

assert lookup(A, 1) == 0
```



## ■ Dynamic data structures

- Examples: resizable arrays, linked lists or trees, object graphs, ...
- Dynamic size, heap-allocated
- Mutable
- To update the data structure, we efficiently change it in-place

```
// dynamic array A = [0,0,0]
A := new Array(3, 0) // not Viper!

B := A // A, B reference same array
B[1] := 17 // in-place mutation


assert A[1] == 17
```



# Why verification of heap-manipulating programs?

- **Memory safety** is the absence of errors related to memory accesses

- dereferencing null-pointers
- accessing unallocated (heap) memory
- accessing dangling pointers
- double-free bugs
- use-after-free bugs



- **Heap-manipulating programs are a prime target for program verification**
  - Efficient algorithms need efficient data structures
  - Device drivers, embedded systems, ...
- Same concepts apply to concurrent programs

# Objects and the heap

1. Heap model
2. Reasoning about objects and references
3. Ownership and access permissions
4. Encoding

# Heap model: an object-based language

→ `00-heap.vpr`

```
field val: Int

method foo() returns (res: Int)
{
  var cell: Ref

  // create object with field val
  cell := new(val)

  cell.val := 5
  res := cell.val
}
```

- A heap is a set of objects
- No classes: each object can have all fields declared in the entire program
  - Type rules of a source language can be encoded
  - Memory consumption is not a concern since programs are not executed
- Objects are accessed via **references**
  - Field read and update operations
  - No information hiding
- No explicit de-allocation (garbage collector)
  - Conceptually, objects could remain allocated

# Extended programming language

(PL6)

## Declarations

$D ::= \dots \mid \text{field } f: T$

Fields are declared globally

## Types

$T ::= \dots \mid \text{Ref}$

Only one type of references

## Expressions

$E ::= \dots \mid \text{null} \mid E.f$

Pre-defined null-reference

Field read expression

## Statements

$S ::= \dots$   
|  $x := \text{new}(\bar{f})$   
|  $x := \text{new}(*)$   
|  $x.f := E$

Allocation with given fields  
or with all fields

Field update of Ref-typed var.



# Objects and the heap

1. Heap model
2. Reasoning about objects and references
3. Ownership and access permissions
4. Encoding

# Proof rule for field read

→ 01-field-read.vpr

- Idea: treat field accesses like variable assignment

## Field read

$$\frac{}{\{ E \neq \text{null} \ \&\& \ Q[x / E.f] \} \ x := E.f \ \{ Q \}}$$

- Additional well-definedness condition prevents null-dereferencing

```
{ true }  
assume r != null && r.val == 5  
{ r != null && r.val == 5 }  
x := p.val  
{ x == 5 }  
assert x == 5  
{ true }
```

# Exercise: Naïve proof rule for field update

- Idea: treat field accesses like variable assignment

Field update

---

$$\{ s \neq \text{null} \ \&\& \ Q[x.f / E] \} \ x := E.f \ \{ Q \}$$

- Additional well-definedness condition prevents null-dereferencing
- The above rule for field update is *unsound*. Give an example that illustrates that.

# Reminder: method framing with global variables

- Method specification declares which variables may get modified

```
var x, y: Int

method set(v: Int)
  modifies x
  ensures x == v
  { ... }
```

```
y := 7
set(5)

assert x > 0 && y == 7
```



- Frame rule (for any statement S)

Frame rule

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

where S does not assign to a variable that is free in R

- Encoding

```
y := 7
havoc x // havoc vars in mod-clause
assume x == 5
assert x > 0 && y == 7
```

# Method framing with heap locations: modifies clause

- Idea: method specification declares which locations may get modified

```
method set(x: Ref, v: Int)
  modifies x.f
  ensures x.f == v
  { ... }
```

## Frame rule

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

where S does not assign to a variable that is free in R

- Two ways to adapt the frame rule
  - «variable» means local or global variable, or «field»
  - «variable» means local or global variable, but not «field»

# Method framing with heap locations: naïve approach

```
method set(x: Ref, v: Int)
  modifies x.f
  ensures x.f == v
  { ... }
```

## Frame rule

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

where S does not assign to a variable that is free in R

«variable» may mean «field»

```
assume y != z
y.f := 7
set(z, 5)
assert y.f == 7
```



- Incomplete: framing is very weak, as information about all objects is lost

«variable» does not mean «field»

```
assume y == r
y.f := 7
set(z, 5)
assert y.f == 7
```



- Unsound: this interpretation of the frame rule ignores aliasing!

# Shortcomings of naïve method framing approach

- Sound encoding needs to consider aliasing

- Inherits shortcomings of candidate rule for field updates
- Explosion of cases
- Treatment of assertions that depend on heap locations implicitly

```
y.f := 7
// encoding of set(z, 5)
var tmp: Int
z.f := tmp // considers aliasing
assume z.f == 5
assert y.f == 7
```

- Many methods modify a statically-unknown set of heap locations

- Locations cannot be listed explicitly in a modifies clause

```
method sort(list: Ref)
  modifies list.val, list.next.val, list.next.next.val, ...
  { ... }
```

- Listing modified heap locations violates information hiding

# Summary of challenges

Heap data structures pose three major challenges for sequential verification

- Reasoning about [aliasing](#)
- [Framing](#), especially for dynamic data structures
- Writing specifications that preserve [information hiding](#)

Additional challenges for concurrent programs, e.g., data races

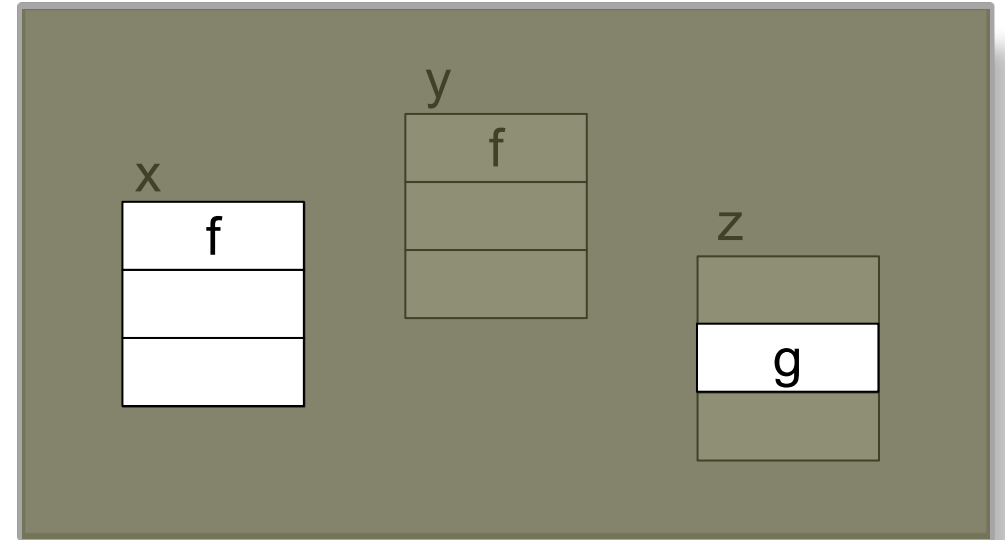


# Objects and the heap

1. Heap model
2. Reasoning about objects and references
3. Ownership and access permissions
4. Encoding

# Access permissions

- Associate each heap location with *at most one* permission
- Read or write access to a memory location requires permission
- Permissions are created when the heap location is allocated
- Permissions can be transferred, but not duplicated or forged



x.f := 5



y.f := 5



z.g := x.f



x.f := y.f



# Permission assertions

- Permissions are denoted by **access predicates**
  - Access predicates are *not* permitted under negations, disjunctions, and on the left of implications
- Predicates may contain both permissions and value constraints
- Predicates must be **self-framing**, that is, include all permissions to evaluate their heap accesses
- An assertion that does not contain access predicates is called **pure** or heap independent

## Predicates

$P ::= \dots \mid \text{acc}(E.f)$

$\text{acc}(p.f) \ \&\& \ p.f > \emptyset$

$\text{requires } p.f > \emptyset$



# Exercise: swapping the fields of two objects

→ 04-swap.vpr

- Implement a swap method that exchanges the field values of two objects.
- Specify its functional behavior.
- Write a client method that creates two objects and calls swap on them. Include an assertion to check that swap's specification is strong enough.
- Change your client method such that it calls swap, passing the same reference twice.

```
field f: Int  
  
method swap(a: Ref, b: Ref)  
{ ... }
```

# Permission assertions and aliasing

→ 05-alias.vpr

## Reminder:

- There is *at most one* permission for every heap location
- Permissions can be transferred, but not duplicated or forged

If we have two permissions `acc(a.f)` and `acc(b.f)`, can `a` and `b` be aliases?

```
field f: Int

method alias(a: Ref, b: Ref)
  requires acc(a.f) && acc(b.f)
{
  a.f := 5
  b.f := 7
  assert a.f == 5
}
```



```
field f: Int

method alias2(a: Ref, b: Ref)
  requires acc(a.f) && acc(b.f)
{
  assert a == b
}
```



→ How do we justify this?

# Permission assertions, more formally

- We extend states to stack-heap pairs  $\sigma = (s, h)$
- The **stack**  $s: \mathbf{Var} \rightarrow \mathbf{Value}$  assigns values to variables
  - We used this as the full state state used in all previous classes
- The **heap**  $h$  assigns values to object-field pairs

$$h: \mathbf{Objects} \times \mathbf{Fields} \xrightarrow{\text{finite partial}} \mathbf{Value}$$

- $dom(h)$  is the set of all object-field pairs for which  $h$  is defined
- $(obj, f) \in dom(h)$  means we have permission to field  $f$  of object  $obj$

# Predicates over extended states

Predicate P	$\mathfrak{I} = (\mathfrak{A}, s, h) \models P$ if and only if
$acc(t.f)$	$(\mathfrak{I}(t), f) \in dom(h)$
$t_1 = t_2$	$\mathfrak{I}(t_1) = \mathfrak{I}(t_2)$
$R(t_1, \dots, t_n)$	$(\mathfrak{I}(t_1), \dots, \mathfrak{I}(t_n)) \in R^{\mathfrak{A}}$
$Q \wedge R$	$\mathfrak{I} \models Q$ and $\mathfrak{I} \models R$
$Q \Rightarrow R$	If $\mathfrak{I} \models Q$ , then $\mathfrak{I} \models R$
$\exists x: \mathbf{T}(Q)$	For some $v \in \mathbf{T}^{\mathfrak{A}}$ , $\mathfrak{I}[x := v] \models Q$
$\forall x: \mathbf{T}(Q)$	For all $v \in \mathbf{T}^{\mathfrak{A}}$ , $\mathfrak{I}[x := v] \models Q$

- Self-framing predicates are always well-defined

Assume  $s(a) == s(b)$  and  $h(a.f) == s(c)$

Does  $\mathfrak{I} = (\mathfrak{A}, s, h) \models acc(a.f) \wedge acc(b.f) \rightarrow b.f == c$  hold?

$\mathfrak{I}(t)$  is the value obtained from evaluating term  $t$  in interpretation  $\mathfrak{I}$

## Examples:

$$\mathfrak{I}(x) = s(x)$$

$$\mathfrak{I}(x + 17) = s(x) +^{\mathfrak{A}} 17^{\mathfrak{A}}$$


$$\mathfrak{I}(x.f) = h(s(x), f)$$

$$\mathfrak{I}(x.f.g) = h(h(s(x), f), g)$$

# Handling aliasing

- Problem: having permissions  $a.f$  and  $b.f$  should mean  $a$  and  $b$  are no aliases
- We introduce a new connective: **the separating conjunction  $P * Q$** 
  - $P * Q$  partitions the heap  $h$  into two chunks
  - Every permission assertion  $\text{acc}(E.f)$  is evaluated in its own heap chunk
  - All other predicates are evaluated in the full heap

$$(\mathcal{A}, s, h) \models_h \text{acc}(x.f) * \text{acc}(y.g) ?$$



$$(\mathcal{A}, s, h) \models_{h_1} \text{acc}(x.f)$$


$$(\mathcal{A}, s, h) \models_{h_2} \text{acc}(y.g)$$



# Handling aliasing

- Problem: having permissions  $a.f$  and  $b.f$  should mean  $a$  and  $b$  are no aliases
- We introduce a new connective: the separating conjunction  $P * Q$ 
  - $P * Q$  partitions the heap  $h$  into two chunks
  - Every permission assertion  $\text{acc}(E.f)$  is evaluated in its own heap chunk
  - All other predicates are evaluated in the full heap

$$(\mathcal{A}, s, h) \models_h \text{acc}(x.f) * \text{acc}(x.f) ?$$



We cannot partition heap  $h$  into  $h1$  and  $h2$  such that both give permission to  $x.f$

$$(\mathcal{A}, s, h) \models_{h1} \text{acc}(x.f)$$

$$(\mathcal{A}, s, h) \models_{h2} \text{acc}(x.f)$$

# Predicates with separating conjunction

Predicate P	$\mathfrak{S} = (\mathfrak{A}, s, h) \models_{h'} P$ if and only if
$acc(t.f)$	$(\mathfrak{S}(t), f) \in dom(h')$
$t_1 = t_2$	$\mathfrak{S}(t_1) = \mathfrak{S}(t_2)$
$R(t_1, \dots, t_n)$	$(\mathfrak{S}(t_1), \dots, \mathfrak{S}(t_n)) \in R^{\mathfrak{A}}$
$Q \wedge R$	$\mathfrak{S} \models_{h'} Q$ and $\mathfrak{S} \models_{h'} R$
$Q * R$	exists partition of $h'$ into $h_1, h_2$ such that $\mathfrak{S} \models_{h_1} Q$ and $\mathfrak{S} \models_{h_2} R$
...	...

evaluate access permissions in current heap chunk  $h'$  (initially  $h$ )

split current heap chunk into two

- $Q * R$  and  $Q \wedge R$  are equivalent if  $Q$  and  $R$  are pure
- Holding permission to  $x.f$  and  $y.f$  implies that  $x$  and  $y$  are no aliases

$$acc(x.f) * acc(y.f) \implies x \neq y$$

# Separating Conjunction in Viper

→ 04-swap.vpr  
→ 05-alias.vpr

- Viper's  $\&\&$  is the separating conjunction  $*$
- Viper has no ordinary conjunction  $\wedge$
- $Q * R$  and  $Q \wedge R$  are equivalent if  $Q$  and  $R$  are pure (heap independent)
- For the call `swap(x, x)`, the precondition is equivalent to false

```
method swap(a: Ref, b: Ref)  
  requires acc(a.f) && acc(b.f)
```

# Exercise

- Reconsider the method on the right.
- Change the precondition such that we can call the method by passing both aliasing references and non-aliasing references to it as arguments without violating the precondition.
- Does the assertion still hold?  
Why (not)?

```
method alias(a: Ref, b: Ref)
  requires acc(a.f) && acc(b.f)
{
  a.f := 5
  b.f := 7
  assert a.f == 5
}
```

# Challenges revisited

Heap data structures pose three major challenges for sequential verification

- Reasoning about aliasing
  - [Permissions and separating conjunction](#)
- Framing, especially for dynamic data structures
- Writing specifications that preserve information hiding



And additional challenges for concurrent programs, e.g., data races

# Field access: proof rules with permissions

## Field read

$$\frac{}{\{ \text{acc}(E.f) * P[x / E.f] \} x := E.f \{ \text{acc}(E.f) * P \}}$$

## Field update

$$\frac{}{\{ \text{acc}(x.f) * x.f == N \} x.f := E \{ \text{acc}(x.f) * x.f == E[x.f / N] \}}$$

- Each field access **requires** (and **preserves**) the corresponding **permission**
- Permission to a location implies that the receiver is non-null
- Substitution with **logical variable N** in the field-update rule is needed to handle occurrences of  $x.f$  inside  $E$  (e.g.,  $x.f := x.f + 1$ )

# Framing

Frame rule

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

where  $S$  does not assign to  
a variable that is free in  $R$

Unsound if  $S$  assigns to  
heap locations constrained by  $R$

# Framing

## Frame rule

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

where  $S$  does not assign to a variable that is free in  $R$

- The frame  $R$  must be self-framing
  - If heap locations constrained by  $R$  are disjoint from those modified by  $S$ ,  $R$  is preserved
  - Otherwise, the precondition is equivalent to false (the triple holds trivially)

## Example

$$\frac{\frac{\{ \text{acc}(x.f) * x.f = N \} \quad x.f := 5 \quad \{ \text{acc}(x.f) * x.f = 5 \}}{\{ \text{acc}(x.f) * x.f = N * \text{acc}(y.f) * y.f = 7 \} \quad x.f := 5 \quad \{ \text{acc}(x.f) * x.f = 5 * \text{acc}(y.f) * y.f = 7 \}}}{\{ \text{acc}(x.f) * x.f = N * \text{acc}(y.f) * y.f = 7 \} \quad x.f := 5 \quad \{ \text{acc}(x.f) * x.f = 5 * \text{acc}(y.f) * y.f = 7 \}}$$



# Framing (cont'd)

- The following proof derives an incorrect triple. Why is it not a valid proof?

$$\frac{\frac{}{\{ \mathbf{acc}(x.f) * x.f = N \} \ x.f := 5 \ \{ \mathbf{acc}(x.f) * x.f = 5 \}}}{\{ \mathbf{acc}(x.f) * x.f = N * x.f = 1 \} \ x.f := 5 \ \{ \mathbf{acc}(x.f) * x.f = 5 * x.f = 1 \}}}$$

- Recall that the frame must be self-framing, which is not the case here
- Making the frame self-framing yields a valid (but vacuous) proof

$$\frac{\frac{}{\{ \mathbf{acc}(x.f) * x.f = N \} \ x.f := 5 \ \{ \mathbf{acc}(x.f) * x.f = 5 \}}}{\{ \mathbf{acc}(x.f) * x.f = N * \mathbf{acc}(x.f) * x.f = 1 \} \ x.f := 5 \ \{ \mathbf{acc}(x.f) * x.f = 5 * \mathbf{acc}(x.f) * x.f = 1 \}}}$$

# Framing for method calls

```
method set(p: Ref, v: Int)
  requires acc(p.f)
  ensures  acc(p.f) && p.f == v
{
  p.f := v
}
```

```
// assume we have acc(x.f) && acc(y.f)
assume y.f == 7
set(x, 5)
assert x.f == 5 && y.f == 7
```


$$\frac{\frac{\{ \mathbf{acc}(p.f) \} \text{ method set}(p, v) \{ \mathbf{acc}(p.f) * p.f = v \}}{\{ \mathbf{acc}(x.f) \} \text{ set}(x, 5) \{ \mathbf{acc}(x.f) * x.f = 5 \}}}{\{ \mathbf{acc}(x.f) * \mathbf{acc}(y.f) * y.f = 7 \} \text{ set}(x, 5) \{ \mathbf{acc}(x.f) * x.f = 5 * \mathbf{acc}(y.f) * y.f = 7 \}}}$$

- Frame rule enables framing without modifies clauses
- A method may modify only heap locations to which it has permission


# Permission transfer

```
method set(p: Ref, v: Int)  
  requires acc(p.f)  
  ensures  acc(p.f) && p.f == v
```

```
{
```




```
p.f := v
```



```
}
```

```
// assume we have acc(x.f) && acc(y.f)  
assume x.f == 2 && y.f == 7
```

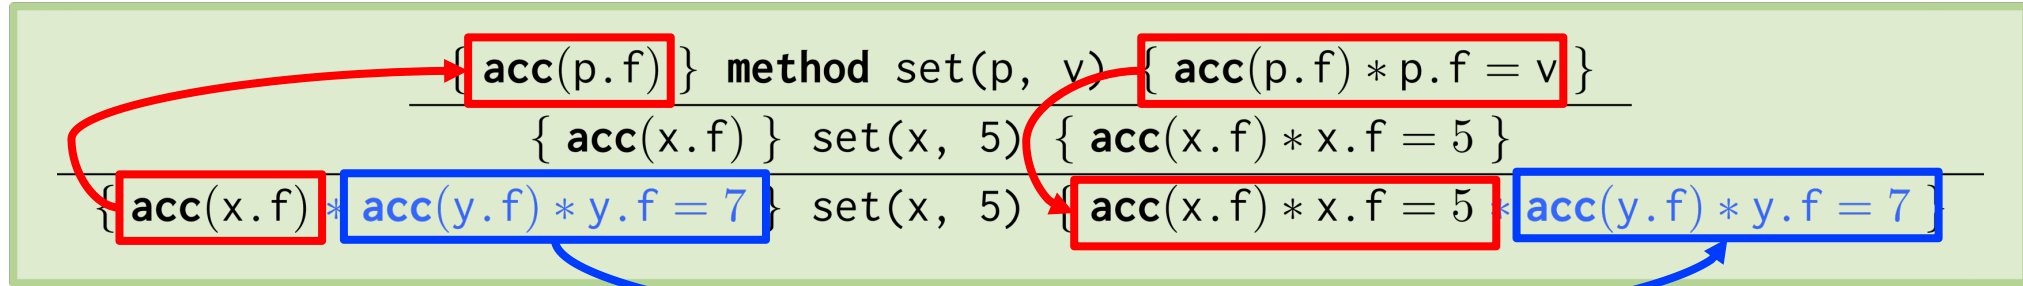


```
set(x, 5)
```

Framing!

```
assert x.f == 5 && y.f == 7
```

# Permission transfer for method calls




- Permissions are held by **method executions** or loop iterations
- Calling a method **transfers permissions from the caller to the callee** (according to the method precondition)
- Returning from a method **transfers permissions from the callee to the caller** (according to the method postcondition)
- **Residual permissions are framed around the call**

# Framing for loops

```
// assume we have acc(x.f) && acc(y.f)  
x.f := 0  
y.f := 7  
while (x.f < 10)  
  invariant acc(x.f)  
  {  
    x.f := x.f + 1  
  }  
assert y.f == 7
```

$$\frac{\frac{\{ \mathbf{acc}(x.f) * x.f < 10 \} \ x.f := x.f + 1 \ \{ \mathbf{acc}(x.f) \}}{\{ \mathbf{acc}(x.f) \} \ \mathbf{while}(x.f < 10) \ \{ \dots \} \ \{ \mathbf{acc}(x.f) * \neg x.f < 10 \}}}{\{ \mathbf{acc}(x.f) * \mathbf{acc}(y.f) * y.f = 7 \} \ \mathbf{while}(x.f < 10) \ \{ \dots \} \ \{ \mathbf{acc}(x.f) * \neg x.f < 10 * \mathbf{acc}(y.f) * y.f = 7 \}}}$$

# Permission transfer for loops



- Permissions are held by method executions or **loop iterations**
- Entering a loop **transfers permissions from the enclosing context to the loop** (according to the loop invariant)
- Leaving a loop **transfers permissions from the loop to the enclosing context** (according to the loop invariant)
- **Residual permissions are framed around the loop**


# Permission transfer: inhale and exhale operations

- **inhale**  $P$  means:

- obtain all permissions required by assertion  $P$
- assume all logical constraints

**inhale**  $\text{acc}(x.f) \ \&\& \ x.f == 2$


2



- **exhale**  $P$  means:

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

**exhale**  $\text{acc}(x.f) \ \&\& \ x.f == 2$



# Encoding of method bodies and calls

```
method foo() returns (...)  
  requires P  
  ensures Q  
  { S }
```

```
x := foo()
```

## ▪ Encoding *without heap and globals*

- Body

```
assume P  
// encoding of S  
assert Q
```

- Call

```
assert P[...]  
havoc x  
assume Q[...]
```

## ▪ Encoding *with heap*

- Body

```
inhale P  
// encoding of S  
exhale Q
```

- Call

```
exhale P[...]  
havoc x  
inhale Q[...]
```

## ▪ **inhale** and **exhale** are permission-aware analogues of **assume** and **assert**



# Exercise: definition of exhale

- **exhale**  $P$  means:
  - assert all logical constraints
  - check and remove all permissions required by  $P$
  - havoc (reset) any locations to which all permission is lost
- Write an example that demonstrates that omitting the havoc from the exhale encoding would be unsound

# Encoding of loops

```
while(b)
  invariant I
{ S }
```

- Reminder: encoding **without heap**

```
assert I
havoc targets
assume I
if(*) {
  assume b
  // encoding of S
  assert I
  assume false
} else {
  assume !b
}
```

- Encoding **with heap**

```
exhale I
havoc targets
inhale I
if(*) {
  assume b
  // encoding of S
  exhale I
  assume false
} else {
  assume !b
}
```

# Encoding of allocation

- new-expression specifies the relevant fields

```
x := new(f, g)
```

- Encoding chooses an arbitrary reference and inhales permissions to relevant fields

```
var x: Ref  
inhale acc(x.f) && acc(x.g)
```

- Incomplete information about freshness of new object

```
x := new(f)  
y := new(f)  
assert x != y
```



```
method foo(y: Ref)  
{  
  var x: Ref  
  x := new(f)  
  assert x != y  
}
```



# Exercise: working with permissions

→ 07-account.vpr

- Implement, specify, and verify a class for bank accounts with the following methods:
  - `create` returns a fresh account with initial balance 0
  - `deposit` deposits a non-negative amount to an account
  - `transfer` transfers a non-negative amount between two accounts
  - Account balances are integers.
- Verify the client program on the right.

```
method client()
{
  var x: Ref
  var y: Ref
  var z: Ref
  x := create()
  y := create()
  z := create()
  deposit(x, 100)
  deposit(y, 200)
  deposit(z, 300)
  transfer(x, y, 100)
  assert x.bal == 0
  assert y.bal == 300
  assert z.bal == 300
}
```

# Verifying memory safety

→ 08-memory-safety.vpr

- Memory safety is the absence of errors related to memory accesses, such as, null-pointer dereferencing, access to un-allocated memory, dangling pointers, out-of-bounds accesses, double free, etc.
- Using permissions, Viper verifies **memory safety by default**

```
var x: Ref  
x.f := 5
```



```
var x: Ref  
x := null  
x.f := 5
```



```
method free(p: Ref)  
  requires acc(p.f)
```

model de-allocation  
via method call

```
free(x)  
x.f := 5
```



```
free(x)  
free(x)
```




See module 8 for arrays

# Challenges revisited

Heap data structures pose three major challenges for sequential verification

- Reasoning about aliasing
  - Permissions and separating conjunction
- Framing, especially for dynamic data structures
  - Sound frame rule, but no support yet for unbounded data structures
- Writing specifications that preserve information hiding



And additional challenges for concurrent programs, e.g., data races

# Objects and the heap

1. Heap model
2. Reasoning about objects and references
3. Ownership and access permissions
4. Encoding

# Heaps

- Encode references and fields

```
type Ref           // type for references
const null: Ref   // null references

type Field T      // polymorphic type for field names
```

```
field f: Int
field g: Ref
```

```
const f: Field int
const g: Field Ref
```

- Heaps map references and field names to values

```
type HeapType = Map<T>[(Ref, Field T), T] // polymorphic map
```

- Represent the program heap as one global variable

```
var Heap: HeapType
```



# Permissions and field access

- Permissions are tracked in a global permission mask

```
type MaskType = Map<T>[(Ref, Field T), bool]  
var Mask: MaskType
```

- Convention:  $\neg$ Mask[null, f] for all fields f

- Field access

```
v := x.f
```

```
assert Mask[x, f]  
v := Heap[x, f]
```

```
x.f := E
```

```
assert Mask[x, f]  
Heap[x, f] := E
```

- Field access requires permission!

- **inhale** P means:
  - obtain all permissions required by assertion P
  - assume all logical constraints
- Encoding is defined recursively over the structure of P

**inhale** E

**assume** [[E]]

[[.]] encoding

**inhale** acc(E.f)

**assume**  $\neg$ Mask[ [[E]], f ]  
Mask[ [[E]], f ] := true

Reaching more than full permission goes to magic

**inhale** E => P

**if**([[E]]) { **inhale** P }

**inhale** P && Q

[[**inhale** P]]; [[**inhale** Q]]

Separating conjunction:  
add sum of permissions

- The encoding also asserts that E is well-defined (omitted here)

# Exhale (1<sup>st</sup> attempt)

- **exhale** P means:
  - assert all logical constraints
  - check and remove all permissions required by assertion P
  - havoc any locations to which all permission is lost
- Encoding is defined recursively over the structure of P

```
exhale E
```

```
assert [[E]]
```

```
exhale acc(E.f)
```

```
assert Mask[ [[E]], f ]  
Mask[ [[E]], f ] := false  
havoc Heap[ [[E]], f ]
```

havoc e.g. by assigning to a fresh variable

```
exhale E => P
```

```
if([[E]]) { [[exhale P]] }
```

```
exhale P && Q
```

```
[[exhale P]]; [[exhale Q]]
```

Separating conjunction:  
remove sum of permissions


- The encoding also asserts that E is well-defined (omitted here)

# Example

```
inhale acc(x.f) && x.f == 5
```

```
assume  $\neg$ Mask[x,f]  
Mask[x,f] := true
```


```
assert Mask[x,f] // well-definedness check  
assume Heap[x,f] == 5
```



```
exhale acc(x.f) && x.f == 5
```

```
assert Mask[x,f]  
Mask[x,f] := false  
havoc Heap[x,f]
```

```
assert Mask[x,f] // well-definedness check  
assert Heap[x,f] == 5
```



# Exhale (fixed)

- Conceptually, permissions should be removed **after** checking logical constraints
- Adapt encoding
  - Check well-definedness against mask at the beginning of the exhale
  - Delay havoc until the end of the exhale

exhale P

```
var oldMask: MaskType
var newHeap: HeapType
oldMask := Mask
[[exhale P]]      // Like before, but no havoc and with
                  // well-definedness check on oldMask
assume forall y,g :: Mask[y,g] ==> newHeap[y,g] == Heap[y,g]
Heap := newHeap // effectively havocs all locations to which
                // permission was lost
```

# Exercise: encoding of exhale

- Encode the operation


```
exhale acc(x.f) && x.f == 5
```

with the fixed encoding.

# Challenges revisited

Heap data structures pose three major challenges for sequential verification

- Reasoning about aliasing
  - Permissions and separating conjunction
- Framing, especially for dynamic data structures
  - Sound frame rule, but no support yet for unbounded data structures
- Writing specifications that preserve information hiding
  - Not solved, but see next module



And additional challenges for concurrent programs, e.g., data races

- Permissions are an excellent basis, but see later