



Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Improving User-Defined Permission Models in Viper

Practical Work

Anqi Li

March 15, 2023

Advisors: Thibault Dardinier, Prof. Dr. Peter Müller

Department of Computer Science, ETH Zürich

Abstract

Viper is a powerful toolchain and infrastructure for program verification. It automates and provides native support for permission-based reasoning by using separation logic. Currently, the fractional permission model is the only permission model available in Viper. A permission model is a partial commutative monoid useful for reasoning about shared data structures. Despite the merits of fractional permissions such as infinite splittability (i.e. any non-empty share can be split into two non-empty shares), alternative permission models can be useful depending on the use cases. This project aims at assessing and improving a prototype that introduces user-defined permission models to Viper. The goal is achieved by first conducting several case studies where multiple permission models are encoded using the prototype, and then adding to the prototype a new feature that allows users to specify a permission model for a field so that multiple permission models can exist in the same program.

Contents

Contents	iii
1 Introduction	1
1.1 Motivation	1
1.2 Related Work	2
1.3 Report Outline	2
2 Case Studies	5
2.1 Background	5
2.1.1 Architecture of Viper	5
2.1.2 Separation Algebra	5
2.2 Counting Permissions	6
2.2.1 Permission Shares	6
2.2.2 Joining Permission Shares	7
2.2.3 Result and Application	7
2.3 Tree Shares	10
2.3.1 Permission Shares	10
2.3.2 Joining Permission Shares	11
2.3.3 Result	11
2.4 Duplicable Permissions	12
2.4.1 Permission Shares	12
2.4.2 Joining Permission Shares	13
2.4.3 Result and Application	13
2.5 Summary	15
3 Permission-Model-per-Field Feature	17
3.1 Syntax	17
3.2 Implementation	18
3.2.1 Parser	18
3.2.2 Typechecker	18

CONTENTS

3.2.3	Permission Module	18
4	Conclusion and Future Work	21
4.1	Conclusion	21
4.2	Future Work	21
A	Appendix	23
A.1	Viper encoding of tree shares	23
	Bibliography	31

Introduction

1.1 Motivation

Program verification is the process of formally proving that a program satisfies its specification. It is useful for proving the absence of errors in various software systems. Nowadays, many software tools have been developed to automate program verification, such as Viper [10]. Dissimilar to other verification tools like Boogie [9] and Why3 [8], Viper uses separation logic [11] to natively support permission-based reasoning via permission models.

A permission model is, typically, a partial commutative monoid [4, 7]. Each permission model has its own set of permission shares that are used to express heap location ownership. Permission shares can be split and recombined. By using a permission model, one can reason about method calls with the frame rule, and also efficiently verify concurrent programs that manipulate shared data structures.

For the time being, Viper supports only one permission model – fractional permissions [2, 3]. In the fractional permission model, permission amounts are rational values in the range of $[0, 1]$. A permission amount of 1 allows one to write to a heap location, while any non-zero permission amount allows one to read from a heap location. The fractional permission model can be used to efficiently reason about programs involving concurrent reads, divide-and-conquer algorithms and so on, but it is not suitable for every situation. For instance, consider the code snippet in Listing 1.1 that attempts to formulate an inductive predicate for a binary tree data structure. When $p > \frac{1}{2}$, the predicate always describes a binary tree, while when $p \leq \frac{1}{2}$, it can describe a directed acyclic graph (DAG) rather than a binary tree. To understand why this is the case, consider a DAG with 4 distinct nodes T, L, R, B as shown in Figure 1.1.

```
1 field leftChild: Ref
2 field rightChild: Ref
3 field elem: Int
4 predicate Tree (this: Ref, p: Perm) {
5     this != null && p > 0/1 ==>
6         acc(this.elem, p) & &
7         acc(this.leftChild, p) & &
8         acc(this.rightChild, p) &&
9         Tree(this.leftChild, P) &&
10        Tree(this.rightChild, p)
11 }
```

Listing 1.1: Tree predicate written in Viper

Assume that $\text{Tree}(T, p)$ holds. The definition of the predicate implies that $\text{Tree}(L, p)$, $\text{Tree}(R, p)$ and $\text{Tree}(B, p+p)$ must also hold. When $p > \frac{1}{2}$, $p + p > 1$ is an undefined permission amount, which invalidates such DAG. In contrast, when $p \leq \frac{1}{2}$, $p + p \leq 1$ is a valid permission amount, so this DAG can be “mistakenly” considered as a tree. As a result, if a method has some positive permission amount to a binary tree and would like to read it, it needs to know that the left and right subtrees of the root node are disjoint to be sure that it is indeed reading a binary tree instead of a DAG. This example demonstrates the limitations of fractional permissions in some situations. Consequently, it is desirable to allow users to define alternative permission models and switch between them depending on the use cases.




Figure 1.1: A DAG described by the tree predicate in Listing 1.1

1.2 Related Work

User-defined permission models were introduced to Viper for the first time as an experimental feature in a master’s thesis project. The thesis [12] documents the theoretical foundations, methodology and implementation details for this feature. As an extension of the thesis, this project builds upon the existing prototype that implements such feature.

1.3 Report Outline

The content of the report is organized as follows. In Chapter 2, multiple case studies are presented to explore the limits of the existing prototype.

Next, in Chapter 3, we describe how a new feature, which enables different permission models in the same Viper program depending on the field, has been implemented based on the prototype. Finally, we draw a conclusion and discuss future work in Chapter 4.

Chapter 2

Case Studies

In this chapter, we present the case studies done with multiple permission models, including counting permissions [2], tree shares [7] and duplicable permissions. The main purpose of the case studies is to test the limits of the existing prototype that enables users to define alternative permission models in Viper programs. To achieve this goal, for each permission model, we study its semantics, encode it using the prototype and test the encoded model. The details of each case study are given in the following sections.

2.1 Background

2.1.1 Architecture of Viper

The prototype and all implementation done in this project are integrated into Carbon – one of the two backends of Viper. It is essentially a compiler that takes a program written in the Viper intermediate language as input and produces a Boogie program as output. Figure 2.1 gives an overview of the architecture of the Viper verification toolchain.

2.1.2 Separation Algebra

A separation algebra (SA) [4] is defined as a tuple $\langle A, \oplus, u \rangle$, where A is a set of elements, \oplus is a partial operation of joining two elements, and u is the unit element in A . Every SA satisfies the following axioms [7]:

- Commutative: $a \oplus b = b \oplus a$
- Associative: $a \oplus (b \oplus c) = (a \oplus b) \oplus c$
- Identity: $a \oplus u = a$
- Cancellative: $a_1 \oplus b = a_2 \oplus b \implies a_1 = a_2$

A permission model is a SA used for permission accounting.

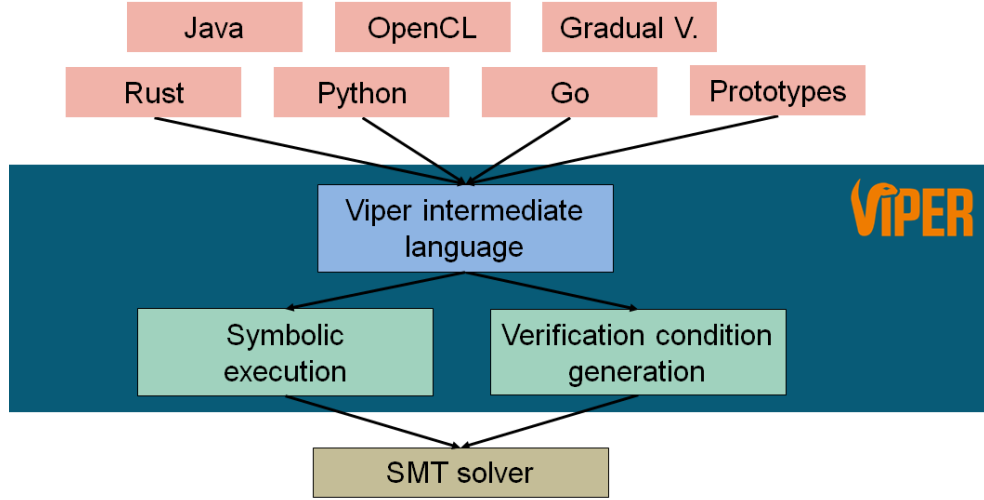


Figure 2.1: Viper architecture overview

2.2 Counting Permissions

The counting permission model allows one to count the number of permissions given out. An intuitive explanation of this permission model can be found in Section 2.1.7 of Roshardt [12]’s thesis, while a more formal definition and extensive details are available in the paper authored by Bornat et al. [2].

2.2.1 Permission Shares

In the model of counting permissions, permission shares are represented using integers. More specifically, 0 means a full permission, -1 means a read permission, and $+k$ means a source permission from which k read permissions have been taken. The unit share is represented by a symbol $u \notin \mathbb{Z}$.

As an example, consider three threads A , B and C . Initially, thread A has a full permission to some resource r , and the other threads have no permissions to access r . In other words, thread A has permission share 0, and threads B and C both have unit permissions represented by u . Thread A later gives out read permissions to threads B and C respectively. As a consequence, thread A now has permission share $+2$, while each of threads B and C has permission share -1 .

2.2.2 Joining Permission Shares

To join two shares a and b , the following rules apply:

$$a \oplus b = \begin{cases} \perp & a \geq 0 \wedge b \geq 0 \\ \perp & (a \geq 0 \vee b \geq 0) \wedge (a + b < 0) \\ a + b & \text{Otherwise} \end{cases} \quad (2.1)$$

2.2.3 Result and Application

By using the prototype, we encoded counting permissions in Viper as shown in Listing 2.1. The function `share(x:Int)` returns the permission share represented by integer x , so the full share, represented by the return value of `full()`, is `share(0)`. Based on Rules 2.1, we can easily formulate the axiom for `plus(x:Perm, y:Perm)` by computing the arithmetic sum of the two input shares, and we may also derive the joinable axiom `axJoinable` by negating the disjunction of all conditions under which two shares $a \oplus b = \perp$. Lastly, to understand the encoding of `minus(x:Perm, y:Perm)` and `geq(x:Perm, y:Perm)`, one should refer to Section 5.3.3 of Roshardt [12]’s thesis.

```

1  domain _Perm {
2      function unit(): Perm
3      function full(): Perm
4      function plus(x: Perm, y:Perm): Perm
5      function minus(x: Perm, y: Perm): Perm
6      function geq(x: Perm, y: Perm): Bool
7      function joinable(x: Perm, y: Perm): Bool
8
9      function share(x: Int): Perm
10
11     axiom axFullShare { full() == share(0) }
12
13     axiom axJoinable {
14         forall a: Int, b: Int ::
15             joinable(share(a), share(b)) <==>
16                 ((a < 0 && b < 0) ||
17                 (a < 0 && b > 0 && a + b >= 0) ||
18                 (a > 0 && b < 0 && a + b >= 0))
19     }
20
21     axiom axJoinableUnit { forall t: Perm ::
22         joinable(unit(), t) && joinable(t, unit())
23     }
24

```

```
25     axiom axPlus { forall a: Int, b: Int ::
26         joinable(share(a), share(b)) ==>
27         plus(share(a), share(b)) == share(a+b)
28     }
29
30     axiom axUnitPlus { forall t: Perm ::
31         plus(unit(), t) == t
32         && plus(t, unit()) == t
33     }
34
35     axiom axMinus { forall a: Int, b: Int ::
36         geq(share(a), share(b)) && a != b ==>
37         minus(share(a), share(b)) == share(a-b)
38     }
39
40     axiom axUnitMinus { forall p: Perm ::
41         minus(p, p) == unit()
42     }
43
44     axiom axGeq {
45         forall a: Int, b: Int ::
46         geq(share(a), share(b)) <==>
47         ((a < 0 && b < 0 && a < b) ||
48          (a >= 0 && (b < 0 || (b >= 0 && a < b)))) ||
49         (a == b))
50     }
51
52     axiom axUnitGeq { forall p: Perm ::
53         geq(p, unit())
54     }
55 }
```

Listing 2.1: Counting permissions encoded using the prototype

The encoded counting permission model has been used to verify a program with a readers-and-writers pattern where multiple threads concurrently read from and write to some shared resource in the critical section. It is useful in this scenario due to its ability to count the permissions given out. Listing 2.2 shows how we modeled this scenario in Viper by referring to Bornat et al. [2]’s work.

In the following code, the variable `count` records the number of readers that are currently accessing the shared resource. To enter the critical section, a reader first needs to gain a read permission. If it is the first reader that attempts to read the resource, it sets the variable `m` to 0 to prevent any writers

from writing to the shared resource. Before the reader leaves the critical section, if it is the last reader to leave, i.e. if `count == 0`, it sets `m` to 1 to unblock the writers. When `m == 1`, one writer will gain a full permission to the shared resource and therefore become able to write.

```

1  field f: Int
2
3  define I_write(m)
4      m != 0 ==> acc(res.f, full())
5
6  define I_read(num)
7      num != 0 ==> acc(res.f, share(num))
8
9  define P() {
10     var m: Int
11     m := havoc_int_binary()
12     assume m != 0
13     inhale I_write(m)
14     m := 0
15     exhale I_write(m)
16 }
17
18 define V() {
19     var m: Int
20     m := havoc_int_binary()
21     assume true
22     inhale I_write(m)
23     m := 1
24     exhale I_write(m)
25 }
26
27 method havoc_int() returns (res: Int)
28 ensures res >= 0
29
30 method havoc_int_binary() returns (res: Int)
31 ensures res == 0 || res == 1
32
33 method reader(res: Ref) {
34     var count: Int
35     count := havoc_int()
36     assume true
37     inhale I_read(count)
38     if (count == 0) {
39         P()

```

```

40     }
41     count := count + 1
42     exhale I_read(count)
43
44     var content: Int
45     content := res.f    // Read
46
47     count := havoc_int()
48     assume count > 0
49     inhale I_read(count)
50     count := count - 1
51     if (count == 0) {
52         V()
53     }
54     exhale I_read(count)
55 }
56
57 method writer(res: Ref) {
58     P()
59     res.f := 2    // Write
60     V()
61 }

```

Listing 2.2: Viper program modeling the readers-and-writers problem

2.3 Tree Shares

The tree share permission model uses binary trees to represent permission shares. We are interested in this model because of its desirable properties: infinite splittability and disjointness. The former property means that any non-empty shares can be split into two non-empty shares, whilst the latter one means that no non-empty share can join with itself. A more extensive description of this model can be found in Section 2.1.7 of Roshardt [12]’s thesis.

2.3.1 Permission Shares

Permission shares are represented using binary trees defined as $\tau = \circ|\bullet|(\tau, \tau)$ where

- \circ is a leaf with boolean value *False*, representing an empty permission. It is also the unit share.
- \bullet is a leaf with boolean value *True*, representing a full permission.

- $\langle \tau, \tau \rangle$ is a binary tree with unlabeled internal nodes and boolean-valued leaves, which can represent any permission amount.

Note that, under such definition, trees of different shapes can represent the same permission amount. For instance, both \bullet and $\langle \bullet, \bullet \rangle$ represent a full permission. We say that such trees are equivalent to one another.

2.3.2 Joining Permission Shares

To join two shares represented by trees of the same shape, one should apply the following rules leaf-wise:

$$\left\{ \begin{array}{l} \circ \oplus \circ = \circ \\ \circ \oplus \bullet = \bullet \\ \bullet \oplus \circ = \bullet \\ \langle \tau_1, \tau_2 \rangle \oplus \langle \tau_3, \tau_4 \rangle = \langle \tau_1 \oplus \tau_3, \tau_2 \oplus \tau_4 \rangle \end{array} \right. \quad (2.2)$$

In the case where two shares are represented by trees τ_1, τ_2 of different shapes, one needs to perform the following steps to convert them to equivalent trees of the same shape before joining:

1. Let h_1, h_2 be the height of τ_1, τ_2 respectively. Compute $h = \max(h_1, h_2)$.
2. For each leaf of τ_1 (resp. τ_2), keep replacing it with two children of the same value as itself until τ_1 (resp. τ_2) becomes a full binary tree¹ of height h .

As an example, Figure 2.2 shows how a binary tree of height 1 is unfolded to an equivalent full binary tree of height 2.




Figure 2.2: Tree unfolding example

2.3.3 Result

In the initial attempt to encode tree shares, permission shares were encoded directly as trees, and equality of two shares was expressed using a function $\text{eq}(p1:\text{Perm}, p2:\text{Perm})$ whose semantics were specified in an axiom. This

¹A full binary tree is a binary tree in which every node except for the leaves has exactly 2 children.

is because two different trees can represent the same permission amount, which cannot be directly expressed with the equality operator `==`. Upon detecting the declaration of `eq(p1:Perm, p2:Perm)` in the permission model, the prototype syntactically replaced all occurrences of the equality operator with calls to this function when translating the Viper program.

However, the tree shares encoded in this way did not have the expected behavior because the equality relation did not propagate to functions. In other words, we do not automatically have $\text{eq}(a, b) \not\Rightarrow \text{eq}(f(a), f(b))$ for some function f . This can happen for many reasons such as missing axioms and bad triggers. Without such implication, one has to perform more complex steps to arrive at the same conclusion as when such implication is available. Therefore, in the final encoding of tree shares, we decided to use a domain to define the binary tree data structure and another one to define the permission shares represented by the trees. Equality of two tree shares is expressed with `==` and the axiom in Listing 2.3. The complete Viper encoding of tree shares is available in Appendix A.1.

```

1 function eqTree(Tree t1, Tree t2): Bool
2 function share(Tree t): _Perm_TreeShare
3 axiom axShareEq { forall t1: Tree, t2: Tree ::
4     eqTree(t1, t2) ==> share(t1) == share(t2)
5 }
```

Listing 2.3: Equality axiom for tree shares

2.4 Duplicable Permissions

Duplicable permissions model the behavior of `wildcard` in Viper, which is an unspecified positive permission amount useful for implementing duplicable read-only resources [1].

2.4.1 Permission Shares

Duplicable permissions have 3 kinds of permission shares:

- \circ : empty permission, also the unit share
- \bullet : full permission
- \odot : some positive permission amount that allows read-only access

2.4.2 Joining Permission Shares

To join two permission shares, the following rules apply:

$$\left\{ \begin{array}{l} \circ \oplus \circ = \circ \\ \circ \oplus \bullet = \bullet \oplus \circ = \bullet \\ \circ \oplus \bullet = \bullet \oplus \circ = \bullet \\ \bullet \oplus \bullet = \bullet \end{array} \right. \quad (2.3)$$

Among these rules, the last one is the most important one. It allows any read-only permissions to be duplicated.

2.4.3 Result and Application

The encoding of duplicable permissions is shown in Listing 2.4.

```

1  domain _Perm {
2      function unit(): Perm
3      function full(): Perm
4      function plus(x: Perm, y: Perm): Perm
5      function minus(x: Perm, y: Perm): Perm
6      function geq(x: Perm, y: Perm): Bool
7      function joinable(x: Perm, y: Perm): Bool
8
9      function share(): Perm
10
11     axiom axUnitAndFull {
12         full() != share() && unit() != share()
13     }
14
15     axiom axPlus { forall p1: Perm, p2: Perm ::
16         joinable(p1, p2) &&
17         p1 != unit() && p2 != unit() ==>
18         plus(p1, p2) == share()
19     }
20
21     axiom axPlusUnit { forall p: Perm ::
22         plus(unit(), p) == plus(p, unit()) &&
23         plus(unit(), p) == p
24     }
25
26     axiom axMinus { forall p1: Perm, p2: Perm ::
27         geq(p1, p2) && p1 != unit() &&
28         p2 != unit() && p1 != full() &&

```

```
29     p2 != full() ==>
30         minus(p1, p2) == share()
31 }
32
33 axiom axMinusFullFull {
34     minus(full(), full()) == unit()
35 }
36
37 axiom axMinusUnit { forall p: Perm ::
38     minus(p, unit()) == p
39 }
40
41 axiom axMinusFull { forall p: Perm ::
42     p != unit() && p != full()
43     ==> minus(full(), p) == share()
44 }
45
46 axiom axJoinable { forall p1: Perm, p2: Perm ::
47     joinable(p1, p2) <==>
48         (p1 != full() && p2 != full()) ||
49         (p1 == full() && p2 == unit()) ||
50         (p1 == unit() && p2 == full())
51 }
52
53 axiom axGeq { forall p1: Perm, p2: Perm ::
54     geq(p1, p2) <==>
55         (p1 == full() && p2 != full()) ||
56         (p1 != unit() && p2 == unit()) ||
57         (p1 != unit() && p2 != unit()
58         && p1 != full() && p2 != full())
59 }
60
61 axiom axGeUnit { forall p: Perm ::
62     geq(p, unit())
63 }
64 }
```

Listing 2.4: Duplicable permissions encoded using the prototype

The encoded duplicable permission model was later used in a quick runtime comparison experiment with wildcard. In the experiment, we created a Viper program with references `ref0`, `ref1`, `ref2`, `ref3`, `ref4`, `ref5` and a block of code, shown in Listing 2.5, where a wildcard permission amount was inhaled and exhaled once for each of these references.

```

1 inhale acc(ref0.f, wildcard)
2     ...
3 inhale acc(ref5.f, wildcard)
4 exhale acc(ref0.f, wildcard)
5     ...
6 exhale acc(ref5.f, wildcard)

```

Listing 2.5: Test code for runtime comparison between wildcard and duplicable permissions

Such code block was then repeated for a different number of times to form a set of test programs. We ran the programs with `wildcard` and recorded the runtime of Viper, and then replaced `wildcard` with `share()` from duplicable permissions, which is `wildcard` in practice, in all programs and measured the runtime again. Table 2.1 shows that duplicable permissions always yield a shorter runtime than `wildcard`. The reduction in runtime is more and more significant as the number of code blocks increases.

Number of code blocks	Runtime of wildcard	Runtime of duplicable permissions
1	2.34s	2.55s
2	3.35s	2.80s
3	152s	3.04s
4	366s	9.91s

Table 2.1: Runtime comparison between wildcard and duplicable permissions

2.5 Summary

The case studies show that the prototype could encode various permission models and that different permission models have different properties. The properties that users might want in a permission model are as follows:

- Infinite splittability: each non-empty share can be split into two non-empty shares.
- Disjointness: no non-empty share joins with itself.
- Countability: the permission model can count the number of permissions given out.
- Recombinability: shares split from a full share can recombine to form a full share.
- Complexity: defined with respect to the workload for SMT solvers to work with the permission model.

Table 2.2 summarizes the properties of the aforementioned permission models.

2. CASE STUDIES

Properties	Fractional permissions	Counting permissions	Tree shares	Duplicable permissions
Infinite splittability	✓		✓	✓
Disjointness			✓	
Countability		✓		
Recombinability	✓	✓	✓	
Complexity	+++	++	++++	+

Table 2.2: Properties of permission models

Consequently, users might prefer one permission model to another based on the use cases, or they might need multiple permission models in one Viper program, which leads to the development of a new feature, namely permission-model-per-field, that we are going to discuss in the next chapter.

Permission-Model-per-Field Feature

This chapter describes a new feature, permission-model-per-field, that allows users to specify a permission model for a field in a Viper program. This feature enables multiple permission models to coexist in the same program. In the following sections, we first show the syntax of this feature, and then illustrate how it has been implemented.

3.1 Syntax

It is obvious that this new feature requires that multiple permission models could be defined in one Viper program. This leads to some changes in the syntax of permission model declaration. To define a permission model, users should declare a Viper domain with identifier `_Perm_<id>` and a list of mandatory function declarations as follows:

- `function unit_<id>():_Perm_<id>`
This function returns a unit share.
- `function full_<id>():_Perm_<id>`
This function returns a full share.
- `function plus_<id>(x:_Perm_<id>, y:_Perm_<id>):_Perm_<id>`
This function returns the result of joining two shares x and y .
- `function minus_<id>(x:_Perm_<id>, y:_Perm_<id>):_Perm_<id>`
This function returns the result of inversely joining two shares x and y .
It is used to compute z such that $x = y \oplus z$.
- `function joinable_<id>(x:_Perm_<id>, y:_Perm_<id>):Bool`
This function checks if two shares x and y can join with each other.
- `function geq_<id>(x:_Perm_<id>, y:_Perm_<id>):Bool`
This function checks if share x is greater than or equal to share y .

An error message will be thrown if any of the functions above is missing or does not have the expected signature. The semantics of these functions depend on the encoded permission model and should be specified using axioms.

To specify a permission model for a field, the following syntax should be used for field declarations:

```
field ([_Perm_<id>])? <field_id>: <field.type>
```

Note that when the permission model of a field is left unspecified, the default permission model, i.e. fractional permissions, is used.

3.2 Implementation

This section gives implementation details of the permission-model-per-field feature. We first show the changes made in the parser and typechecker of Silver, and then list the updates in the permission module of Carbon.

3.2.1 Parser

The grammar rule of field declarations has been updated to allow users to specify `_Perm_<id>` for the declared field. This leads to the addition of one argument, i.e. `permId`, to the constructor of `PField` with a default value `Fractions`.

3.2.2 Typechecker

The typechecker checks that `_Perm_<id>` in every field declaration is indeed the identifier of a domain declared in the program. It also ensures that strings starting with `unit`, `full`, `plus`, `minus`, `geq` and `joinable` are only used as the identifier of functions declared in permission models, i.e. any domain whose identifier has the prefix `_Perm_`. Moreover, the typechecking rules are updated such that domain types are allowed wherever only the built-in `Perm` type was expected in the prototype. This has caused changes in the typechecking of access predicates, `perm()` function calls, and wherever `wildcard` can appear. Furthermore, the typechecker makes sure that the permission model used to access a field matches with the permission model in the declaration of the field. For instance, if `x` has type `Ref` and field `f` is declared to use counting permissions, the typechecker rejects the statement `inhale acc(x.f, 1/2)`.

3.2.3 Permission Module

The permission module specializes in the handling of permission models. It checks that the declaration of permission models has the expected syntax, and

then creates and maintains a Map object that allows us to access everything in a permission model by providing its identifier.

To enable the use of multiple permission models, the module makes use of polymorphism. Recall that every permission model is required to have its own `unit`, `full`, `plus`, `minus`, `geq` and `joinable` functions. Instead of emitting Boogie functions corresponding to these functions for every different permission model encountered, the module emits the following polymorphic Boogie functions only once:

```

1 function full<P>(): P;
2 function unit<P>(): P;
3 function joinable<P>(x: P, y: P): bool;
4 function plus<P>(x: P, y: P): P;
5 function geq<P>(x: P, y: P): bool;
6 function minus<P>(x: P, y: P): P;

```

Listing 3.1: Polymorphic Boogie functions emitted by the permission module

The semantics of these functions from each different permission model are defined via axioms with the type parameter `P` specified. For example, Listing 3.2 shows an axiom that describes the semantics of the `plus` function for fractional permissions, while the axiom in Listing 3.3 formulates the semantics of the same function for duplicable permissions. The type parameter `P` is instantiated as `FracPerm` in the former axiom and `_Perm_wildcardDomainType` in the latter one.

```

1 axiom (forall x: FracPerm, y: FracPerm ::
2   joinable(x, y) ==> plus(x, y) == x + y
3 );

```

Listing 3.2: Axiom of `plus` function from fractional permissions

```

1 axiom (forall p1: _Perm_wildcardDomainType, p2:
2   _Perm_wildcardDomainType ::
3   joinable(p1, p2) && p1 != unit() && p2 != unit()
4   ==> (plus(p1, p2): _Perm_wildcardDomainType) == (
5     share(): _Perm_wildcardDomainType)
6 );

```

Listing 3.3: Axiom of `plus` function from duplicable permissions

By using polymorphism, code duplication is avoided since functions shared by permission models are emitted only once in the Boogie program. Moreover, existing code can be reused without adaptation. For example, the code emitting the axiom that initializes the zero masks to unit shares, shown in Listing 3.4, does not need to be adapted, because there is no need to specify the permission model associated with each unit share thanks to

3. PERMISSION-MODEL-PER-FIELD FEATURE

polymorphism. More importantly, the modularity of Carbon code base is maintained: other modules in Carbon do not need to know the existence of multiple permission models in the permission module.

```
1 axiom (forall <A, B, P> o_2: Ref, f_5: (Field A B P)
2     :: { ZeroMask[o_2, f_5] }
3     !IsPredicateField(f_5) ==>
4     ZeroMask[o_2, f_5] == (unit(): P)
5 );
```

Listing 3.4: Zero mask initialization

Conclusion and Future Work

4.1 Conclusion

In this project, we have carried out case studies involving multiple permission models to assess the functionality of the prototype that allows users to define alternative permission models in Viper. Several bug fixes have also been made to the prototype at the same time. Additionally, we have added a new feature to the prototype to enable users to specify a permission model for a field so that multiple permission models can exist in one Viper program.

4.2 Future Work

There are several aspects that can be improved in the future to extend our work. Firstly, more extensive testing of the permission-model-per-field feature is needed. One possible way to do so is to encode real-world scenarios that use multiple permission models. Secondly, the automation of magic wands [6] needs to be adapted to the changes made in the permission module of Carbon. Last but not least, due to the introduction of the permission-model-per-field feature, there can exist multiple permission models in one predicate. Thereby, the multiplication of predicates [5] needs to be generalized.

Appendix A

Appendix

A.1 Viper encoding of tree shares

```
1
2 domain Tree {
3   function getRight(this: Tree): Tree
4   function getLeft(this: Tree): Tree
5   function node(t1: Tree, t2: Tree): Tree
6
7   function eqTree(t1: Tree, t2: Tree): Bool
8   function _eqTree(t1: Tree, t2: Tree): Bool
9   function geqTree(t1: Tree, t2: Tree): Bool
10  function _geqTree(t1: Tree, t2: Tree): Bool
11  function joinableTree(t1: Tree, t2: Tree): Bool
12  function _joinableTree(t1: Tree, t2: Tree):
    Bool
13  function plusTree(t1: Tree, t2: Tree): Tree
14  function _plusTree(t1: Tree, t2: Tree): Tree
15  function minusTree(t1: Tree, t2: Tree): Tree
16  function _minusTree(t1: Tree, t2: Tree): Tree
17  function isLeaf(t1: Tree): Bool
18
19
20  function toFullTree(t: Tree, h: Int): Tree
21  function sameShape(t1: Tree, t2: Tree): Bool
22  function getHeight(t: Tree): Int
23  function max(a: Int, b: Int): Int
24
25  function dummy(t: Tree): Bool
26
27  // Leaves
```

```
28     unique function zero(): Tree    // White circle
29     unique function one(): Tree    // Black circle
30
31     axiom axConst {
32         zero() != one()
33     }
34
35     axiom axMax { forall a: Int, b: Int ::
36         {max(a, b)}
37         (a >= b ==> max(a, b) == a) &&
38         (a < b ==> max(a, b) == b)
39     }
40
41     axiom axIsLeaf { forall t: Tree :: {isLeaf(t)}
42         isLeaf(t) <==> (t == zero() || t == one())
43     }
44
45     axiom axLeaf { forall t: Tree ::
46         (t == zero() || t == one()) ==>
47         (getRight(t) == t && getLeft(t) == t)
48     }
49
50     axiom axNonLeaf { forall t1: Tree, t2: Tree ::
51         getLeft(node(t1, t2)) == t1 &&
52         getRight(node(t1, t2)) == t2 &&
53         !isLeaf(node(t1, t2))
54     }
55
56     axiom axSameShape { forall t1: Tree, t2: Tree
57         :: {dummy(t1), dummy(t2)}
58         sameShape(t1, t2) <==>
59         (isLeaf(t1) && isLeaf(t2)) ||
60         (!isLeaf(t1) && !isLeaf(t2) &&
61         sameShape(getLeft(t1), getLeft(t2)) &&
62         sameShape(getRight(t1), getRight(t2))
63         )
64     }
65
66     axiom axGetHeight { forall t: Tree ::
67         {dummy(t), getHeight(t)}
68         (isLeaf(t) ==> getHeight(t) == 0) &&
69         (!isLeaf(t) ==> getHeight(t) == 1 + max(
70             getHeight(getLeft(t)), getHeight(
71                 getRight(t))))
```

```

69   }
70
71   axiom axToFullTree { forall t: Tree, h: Int ::
72     {dummy(t), toFullTree(t, h)}
73     (h == 0 ==> toFullTree(t, h) == t) &&
74     (h != 0 ==> toFullTree(t, h) == node(
75       toFullTree(getLeft(t), h-1), toFullTree(
76         getRight(t), h-1)))
77   }
78
79   axiom axEqTree { forall t1: Tree, t2: Tree :: {
80     eqTree(t1, t2)}
81     (sameShape(t1, t2) ==> eqTree(t1, t2) ==
82       _eqTree(t1, t2)) &&
83     (!sameShape(t1, t2) ==> eqTree(t1, t2) ==
84       _eqTree(toFullTree(t1, max(getHeight(t1),
85         getHeight(t2))), toFullTree(t2, max(
86         getHeight(t1), getHeight(t2)))))
87   }
88
89   // Assume that t1 & t2 have the same shape
90   axiom axEqTreeHelper { forall t1: Tree, t2:
91     Tree :: {dummy(t1), dummy(t2)}
92     _eqTree(t1, t2) <==>
93     ((isLeaf(t1) && isLeaf(t2) && t1 == t2)
94       ||
95       (!isLeaf(t1) && !isLeaf(t2) && _eqTree(
96         getRight(t1), getRight(t2)) &&
97         _eqTree(getLeft(t1), getLeft(t2))))
98   }
99
100  axiom axJoinableTree { forall t1: Tree, t2:
101    Tree :: {joinableTree(t1, t2)}
102    (sameShape(t1, t2) ==> joinableTree(t1, t2)
103      == _joinableTree(t1, t2)) &&
104    (!sameShape(t1, t2) ==> joinableTree(t1, t2)
105      == _joinableTree(toFullTree(t1, max(
106        getHeight(t1), getHeight(t2))),
107        toFullTree(t2, max(getHeight(t1),
108        getHeight(t2)))))
109  }
110
111  // Assume that t1 & t2 have the same shape

```

```

95     axiom axJoinableTreeHelper { forall t1: Tree,
96         t2: Tree :: {dummy(t1), dummy(t2)}
97         _joinableTree(t1, t2) <==>
98         (
99             (isLeaf(t1) && isLeaf(t2) &&
100                ((t1 == zero() && t2 == one())
101                 ||
102                 (t1 == one() && t2 == zero()))
103             ||
104             (t1 == zero() && t2 == zero()))
105         ) ||
106         (!isLeaf(t1) && !isLeaf(t2) &&
107            _joinableTree(getRight(t1),
108                          getRight(t2)) &&
109            _joinableTree(getLeft(t1),
110                          getLeft(t2)))
111     )
112 }
113
114 axiom axGeqTree { forall t1: Tree, t2: Tree ::
115     {geqTree(t1, t2)}
116     (sameShape(t1, t2) ==> geqTree(t1, t2) ==
117        _geqTree(t1, t2)) &&
118     (!sameShape(t1, t2) ==> geqTree(t1, t2) ==
119        _geqTree(toFullTree(t1, max(getHeight(t1),
120                                   getHeight(t2))), toFullTree(t2, max(
121            getHeight(t1), getHeight(t2)))))
122 }
123
124 // Assume that t1 & t2 have the same shape
125 axiom axGeqTreeHelper { forall t1: Tree, t2:
126     Tree :: {dummy(t1), dummy(t2)}
127     _geqTree(t1, t2) <==>
128     ((isLeaf(t1) && isLeaf(t2) &&
129        (t1 == t2) || (t1 == one() && t2 ==
130         zero())) ||
131        (!isLeaf(t1) && !isLeaf(t2) &&
132         _geqTree(getLeft(t1), getLeft(t2))
133         && _geqTree(getRight(t1),
134                     getRight(t2))))
135 }

```



```

124 axiom axPlusTree { forall t1: Tree, t2: Tree ::
    {plusTree(t1, t2)}
125 (sameShape(t1, t2) ==> plusTree(t1, t2) ==
    _plusTree(t1, t2)) &&
126 (!sameShape(t1, t2) ==> plusTree(t1, t2) ==
    _plusTree(toFullTree(t1, max(getHeight(t1),
    getHeight(t2))), toFullTree(t2, max(
    getHeight(t1), getHeight(t2)))))
127 }
128
129 // Assume that t1 & t2 have the same shape
130 axiom axPlusTreeHelper { forall t1: Tree, t2:
    Tree :: {dummy(t1), dummy(t2)}
131 _joinableTree(t1, t2) ==>
132 (
133     (isLeaf(t1) && isLeaf(t2) &&
134         (((t1 == zero() && t2 == one()
135             ) || (t1 == one() && t2 ==
136                 zero()))
137             ==> _plusTree(t1, t2)
138                 == one()) &&
139         ((t1 == zero() && t2 == zero()
140             ) ==> _plusTree(t1, t2) ==
141                 zero()))
142     ) ||
143     (!isLeaf(t1) && !isLeaf(t2) &&
144         _plusTree(t1, t2) == node(
145             _plusTree(getLeft(t1),
146                 getLeft(t2)), _plusTree(
147                 getRight(t1), getRight(t2)))
148         )
149 )
150 }
151
152 axiom axMinusTree { forall t1: Tree, t2: Tree
    :: {minusTree(t1, t2)}
153 (sameShape(t1, t2) ==> minusTree(t1, t2) ==
    _minusTree(t1, t2)) &&
154 (!sameShape(t1, t2) ==> minusTree(t1, t2)
    == _minusTree(toFullTree(t1, max(
    getHeight(t1), getHeight(t2))),
    toFullTree(t2, max(getHeight(t1),
    getHeight(t2)))))
155 }

```

```

147
148 // Assume that t1 & t2 have the same shape
149 axiom axMinusTreeHelper { forall t1: Tree, t2:
    Tree :: {dummy(t1), dummy(t2)}
150     _geqTree(t1, t2) ==>
151     (
152         (isLeaf(t1) && isLeaf(t2) &&
153             ((t1 == t2 ==> _minusTree(t1, t2)
154                 == zero()) &&
155                 ((t1 == one() && t2 == zero()) ==>
156                     _minusTree(t1, t2) == one()))
157         ) ||
158         (!isLeaf(t1) && !isLeaf(t2) &&
159             _minusTree(t1, t2) == node(
160                 _minusTree(getLeft(t1), getLeft(
161                     t2)), _minusTree(getRight(t1),
162                         getRight(t2))))
163     )
164 }
165
166 domain _Perm {
167     function unit(): Perm
168     function full(): Perm
169     function plus(x: Perm, y: Perm): Perm
170     function minus(x: Perm, y: Perm): Perm
171     function geq(x: Perm, y: Perm): Bool
172     function joinable(x: Perm, y: Perm): Bool
173     function share(tree: Tree): Perm
174
175     axiom axUnit { unit() == share(zero()) }
176     axiom axFull { full() == share(one()) }
177
178     axiom axJoinable { forall t1: Tree, t2: Tree ::
179         joinable(share(t1), share(t2)) ==
180         joinableTree(t1, t2)
181     }
182
183     axiom axPlus { forall t1: Tree, t2: Tree ::
184         joinable(share(t1), share(t2)) ==>
185         plus(share(t1), share(t2)) ==
186         share(plusTree(t1, t2))
187     }
188 }

```

```
184   axiom axGeq { forall t1: Tree, t2: Tree ::
185       geq(share(t1), share(t2)) ==
186       geqTree(t1, t2)
187   }
188
189   axiom axMinus { forall t1: Tree, t2: Tree ::
190       geq(share(t1), share(t2)) ==>
191       minus(share(t1), share(t2)) ==
192       share(minusTree(t1, t2))
193   }
194
195   axiom axEq { forall a: Tree, b: Tree ::
196       eqTree(a, b) ==> share(a) == share(b)
197   }
198
199   axiom axUnitPlus { forall p: Perm ::
200       plus(unit(), p) == p &&
201       plus(p, unit()) == p
202   }
203 }
```

Listing A.1: Equality axiom for tree shares

Bibliography

- [1] Viper tutorial. <https://viper.ethz.ch/tutorial/>.
- [2] Richard Bornat, Cristiano Calcagno, Peter O’Hearn, and Matthew Parkinson. Permission accounting in separation logic. *SIGPLAN Not.*, 40(1): 259–270, jan 2005. ISSN 0362-1340. doi: 10.1145/1047659.1040327. URL <https://doi.org/10.1145/1047659.1040327>.
- [3] John Boyland. Checking interference with fractional permissions. In Radhia Cousot, editor, *Static Analysis*, pages 55–72, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg. ISBN 978-3-540-44898-3.
- [4] Cristiano Calcagno, Peter W. O’Hearn, and Hongseok Yang. Local action and abstract separation logic. In *22nd Annual IEEE Symposium on Logic in Computer Science (LICS 2007)*, pages 366–378, 2007. doi: 10.1109/LICS.2007.30.
- [5] Thibault Dardinier, Peter Müller, and Alexander J. Summers. Fractional resources in unbounded separation logic. *Proc. ACM Program. Lang.*, 6 (OOPSLA2), oct 2022. doi: 10.1145/3563326. URL <https://doi.org/10.1145/3563326>.
- [6] Thibault Dardinier, Gaurav Parthasarathy, Noé Weeks, Peter Müller, and Alexander J. Summers. Sound automation of magic wands. In Sharon Shoham and Yakir Vizel, editors, *Computer Aided Verification*, pages 130–151, Cham, 2022. Springer International Publishing. ISBN 978-3-031-13188-2.
- [7] Robert Dockins, Aquinas Hobor, and Andrew W. Appel. A fresh look at separation algebras and share accounting. In Zhenjiang Hu, editor, *Programming Languages and Systems*, pages 161–177, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg. ISBN 978-3-642-10672-9.

- [8] Jean-Christophe Filliâtre and Andrei Paskevich. Why3 — where programs meet provers. In *Programming Languages and Systems*, pages 125–128, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. ISBN 978-3-642-37036-6.
- [9] K. Rustan M. Leino. This is boogie 2. June 2008. URL <https://www.microsoft.com/en-us/research/publication/this-is-boogie-2-2/>.
- [10] Peter Müller, Malte Schwerhoff, and Alexander J. Summers. Viper: A verification infrastructure for permission-based reasoning. In *Verification, Model Checking, and Abstract Interpretation*, pages 41–62, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg. ISBN 978-3-662-49122-5.
- [11] J.C. Reynolds. Separation logic: a logic for shared mutable data structures. In *Proceedings 17th Annual IEEE Symposium on Logic in Computer Science*, pages 55–74, 2002. doi: 10.1109/LICS.2002.1029817.
- [12] Matthias Roshardt. Extending the viper verification language with user-defined permission models. Master’s thesis, 2021.



Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Declaration of originality

The signed declaration of originality is a component of every semester paper, Bachelor's thesis, Master's thesis and any other degree paper undertaken during the course of studies, including the respective electronic versions.

Lecturers may also require a declaration of originality for other written papers compiled for their courses.

I hereby confirm that I am the sole author of the written work here enclosed and that I have compiled it in my own words. Parts excepted are corrections of form and content by the supervisor.

Title of work (in block letters):

Improving User-Defined Permission Models in Viper

Authored by (in block letters):

For papers written by groups the names of all authors are required.

Name(s):

Li

First name(s):

Anqi

With my signature I confirm that

- I have committed none of the forms of plagiarism described in the '[Citation etiquette](#)' information sheet.
- I have documented all methods, data and processes truthfully.
- I have not manipulated any data.
- I have mentioned all persons who were significant facilitators of the work.

I am aware that the work may be screened electronically for plagiarism.

Place, date

Zürich, Switzerland March 15, 2023

Signature(s)

李安琪

For papers written by groups the names of all authors are required. Their signatures collectively guarantee the entire content of the written paper.