



Abstraction and Subsumption in Modular Verification of C Programs

Lennart Beringer^(✉) and Andrew W. Appel

Princeton University, Princeton, NJ 08544, USA
{eberinge,appel}@cs.princeton.edu

Abstract. Representation predicates enable *data abstraction* in separation logic, but when the same concrete implementation may need to be abstracted in different ways, one needs a notion of *subsumption*. We demonstrate *function-specification subtyping*, analogous to subtyping, with a *subsumption* rule: if ϕ is a **funspec_sub** of ψ , that is $\phi <: \psi$, then $x : \phi$ implies $x : \psi$, meaning that any function satisfying specification ϕ can be used wherever a function satisfying ψ is demanded. We extend previous notions of Hoare-logic sub-specification, which already included parameter adaption, to include *framing* (necessary for separation logic) and impredicative bifunctors (necessary for higher-order functions, i.e. function pointers). We show intersection specifications, with the expected relation to subtyping. We show how this enables compositional modular verification of the functional correctness of C programs, in Coq, with foundational machine-checked proofs of soundness.

Keywords: Foundational program verification · Separation logics · Specification subsumption

1 Introduction

Even in the 21st century, the world still runs on C: operating systems, run-time systems, network stacks, cryptographic libraries, controllers for embedded systems, and large swaths of critical infrastructure code are either directly hand-coded in C or employ C as intermediate target of compilation or code synthesis. Analysis methods and verification tools that apply to C thus remain a vital area of research. The Verified Software Toolchain (VST) [4] is a semi-automated proof system for functional-correctness verification of C programs that integrates two long-standing lines of research: (i) program logics with machine-checked proofs of soundness; (ii) practical verification tools for industry-strength programming languages. VST consists of three main components:

Verifiable C [3] is a higher-order impredicative concurrent separation logic covering almost all the control-flow and data-structuring features of C (we currently omit goto and by-copy whole-struct assignment);

VST-Floyd [7] is a library of lemmas, definitions, and automation tactics that assist the user in applying the program logic to a program, using forward symbolic execution, with separation logic assertions as symbolic states;

The semantic model justifies the proof rules, exploiting the theories of step-indexing, impredicative quantification, separation algebras, and concurrent ghost state. The semantic model is the basis of a machine-checked proof [4], in Coq, that the Verifiable C program logic is sound w.r.t. the operational semantics of CompCert Clight. Thus the user’s Coq proof *in* Verifiable C composes with our soundness proof *of* Verifiable C and with Leroy’s CompCert compiler correctness proof [15] to yield an end-to-end proof of the functional correctness of the assembly-language program.

VST’s key feature—distinguishing it from tools such as VCC [8], Frama-C [11], or VeriFast [9]—is that it is *entirely* implemented in the Coq proof assistant. A user imports C code into the Coq development environment and applies VST-Floyd’s automation—computational decision procedures from Coq’s standard library, plus custom-built tactics for forward symbolic execution and entailment checking—to construct formal derivations in the Verifiable C program logic. The full power of Coq and its libraries are available to manipulate application-specific mathematics. The semantic validity of the proof rules—machine-checked by Coq’s kernel—connects these derivations to Clight, i.e. CompCert’s representation of parsed and determinized C code.

Recent applications of VST include the verification of cryptographic primitives from OpenSSL [2,6] and mbedTLS [24], an asynchronous communication mechanism [17], and an internet-facing server component [13]. Ongoing efforts elsewhere include a generational garbage collector and a malloc-free library.

Motivated by these applications, we now add support for data abstraction, a key enabler of scalability. As shown in previous work [21], separation logic can easily express data abstraction, using abstract predicates: just as the client program of an abstract data type (ADT) can be written without knowing the representation, verification of the client can proceed without knowing the representation. In type theory, this is the principle of *existential types* [18].

But in real-life modular programming, the same function may want more than one specification. For example, a function may expose a concrete specification to “friend” functions that know the representation of internal data and a more abstract specification for clients that do not. In this case, one should not have to verify the function-body twice, once for each specification; instead, one should verify the function-body with respect to the concrete specification, then prove the concrete implies the abstract. Again, type theory provides an appropriate notion: *subtyping* [22]. In other cases, it may be desirable to specify different use cases of a function—applying, for example, to different input configurations, or to different control flow paths—using different specifications, perhaps using different abstract predicates. Yet again, type theory provides a useful analogue: *intersection types*, a form of ad-hoc polymorphism.

These observations motivate the use of type-theoretic principles as guidelines for developing specification mechanisms and automation features for abstrac-

tion. We now take a step in this direction, focusing primarily on the notion of subtyping. The observation that Hoare’s original rule of consequence is insufficiently powerful in languages with (recursive) procedures motivated research into *parameter adaptation*, by (among others) Kleymann, Nipkow, and Naumann [12, 19, 20]. Indeed, Kleymann observed that ([12], p. 9).

- *in proving that the postcondition has been weakened, one may also assume the precondition of the conclusion holds...*
- *one may adjust the auxiliary variables in the premise. Their value may depend on the value of auxiliary variables in the conclusion and the value of all program variables in the initial state.*


But these developments were carried out for small languages and predate the emergence of separation logic. The present article hence revisits these ideas in the context of VST, by developing a powerful notion of function-specification subtyping for higher-order impredicative separation logic. Our treatment improves on previous work in several regards:

- We support function-specifications of function pointers, as part of our support for almost the entire C language. Kleymann only considers a single (anonymous, parameterless, but possibly recursive) procedure, while Nipkow supports mutual recursion between named procedures.
- Our notion of subtyping avoids direct quantification over states, thus permitting a higher-order impredicative separation logic in the style of VST and Iris [10], where “assertion” must be an abstract type with a step-indexed model rather than simply $\text{state} \rightarrow \text{Prop}$. This is necessary to fully support function pointers and higher-order resource invariants (for concurrent programming). In contrast, Kleymann’s and Nipkow’s assertions are predicates over states, and the side conditions of their adaptation rules explicitly quantify over states. Naumann’s formulation using predicate transformers captures the same relationship in a slightly more abstract manner.
- VST associates function specifications to globally named functions in its proof context Δ and includes a separation logic assertion `func.at` that attaches specifications to function-pointer values. Our treatment integrates subsumption coherently into proof contexts, `func.at`, and the soundness judgment. We support subsumption at function call sites but also incorporate subsumption in a notion of (proof) context subtyping that is reminiscent of record subtyping [22]. This will allow bundling function specifications into specifications of objects or modules that can be abstractly presented to client programs and are compatible with behavioral subtyping [14, 16, 23].
- We introduce intersection specifications and show that their interaction with subsumption precisely matches that of intersection types.

Our presentation is example-driven: we illustrate several use cases of subsumption on concrete code fragments in Verifiable C. Technical adaptations of the model that support these verifications have been machine-checked for soundness, but in the paper we only sketch them. The full Coq proofs of our example are in the VST repo, github.com/PrincetonUniversity/VST in directory `proofs/pile`.

2 Function Specifications in Verifiable C

Our main example is an abstract data type (ADT) for *piles*, simple collections of integers. Figure 1 (on the next page) shows a modular C program that throws numbers onto a pile, then adds them up.



The diagram at left shows that `pile.c` is imported by `onepile.c` (which manages a single pile), `apile.c` (which manages a single pile in a different way), and `triang.c` (which computes the n th triangular number). The latter three modules are imported by `main.c`. `Onepile.c` and `triang.c` import the abstract interface `pile.h`; `apile.c` imports also the low-level concrete interface `pile_private.h` that exposes the representation—a typical use case for this organization might be when `apile.c` implements representation-dependent debugging or performance monitoring.

When—as shown on the right—`pile.c` is replaced by a faster implementation `fastpile.c` (code in Fig. 3) using a different data structure, `apile.c` must be replaced with `fastapile.c`, but the other modules need not be altered, *and neither should their specification or verification*.

Figure 2 presents the specification of the pile module, in the Verifiable C separation logic. Each C-language function identifier (such as `_Pile.add`) is bound to a `funspec`, a function specification in separation logic.

Before specifying the functions (with preconditions and postconditions), we must first specify the data structures they receive as arguments and return as results. Linked lists are specified as usual in separation logic: `listrep` is a recursive definition over the abstract (“mathematical”) list value σ , specifying how it is laid out in a memory footprint rooted at address p . Then `pilerep` describes a memory location containing a pointer to a `listrep`.

A `funspec` takes the form, `WITH  $\vec{x} : \vec{\tau}$  PRE ... POST ...`. For example, take `Pile.add_spec` from Fig. 2: the $\vec{x}$ are bound Coq variables visible in both the precondition and postcondition, in this case, $p:val$, $n:\mathbb{Z}$, $\sigma: \text{list } \mathbb{Z}$, $gv: \text{globals}$, where p is the address of a pile data structure, n is the number to be added to the pile, σ is the sequence currently represented by the pile, and gv is a way to access all named global variables. The PREcondition is parameterized by the C-language formal parameter names `_p` and `_n`. An assertion in Verifiable C takes the form, `PROP(propositions) LOCAL(variable bindings) SEP(spatial conjuncts)`. In

```

/* pile.h */
typedef struct pile *Pile;
Pile Pile_new(void);
void Pile_add(Pile p, int n);
int Pile_count(Pile p);
void Pile_free(Pile p);

/* onepile.h */
void Onepile_init(void);
void Onepile_add(int n);
int Onepile_count(void);

/* apile.h */
void Apile_add(int n);
int Apile_count(void);

/* triang.h */
int Triang_nth(int n);

/* triang.c */
#include "pile.h"
int Triang_nth(int n) {
    int i, c;
    Pile p = Pile_new();
    for (i=0; i<n; i++)
        Pile_add(p, i+1);
    c = Pile_count(p);
    Pile_free(p);
    return c;
}

/* onepile.c */
#include "pile.h"
Pile the_pile;
void Onepile_init(void)
{ the_pile = Pile_new(); }
void Onepile_add(int n)
{ Pile_add(the_pile, n); }
int Onepile_count(void)
{ return Pile_count(the_pile); }

/* pile_private.h */
struct list {int n; struct list *next;};
struct pile {struct list *head;};

/* pile.c */
#include <stddef.h>
#include "stdlib.h"
#include "pile.h"
#include "pile_private.h"
Pile Pile_new(void) {
    Pile p = (Pile)surely_malloc(sizeof *p);
    p->head=NULL;
    return p;
}
void Pile_add(Pile p, int n) {
    struct list *head = (struct list *)
        surely_malloc(sizeof *head);
    head->n=n;
    head->next=p->head;
    p->head=head;
}
int Pile_count(Pile p) {
    struct list *q;
    int c=0;
    for(q=p->head; q; q=q->next)
        c += q->n;
    return c;
}
void Pile_free(Pile p) { . . . }

/* apile.c */
#include "pile.h"
#include "pile_private.h"
#include "apile.h"
struct pile a_pile = {NULL};
void Apile_add(int n)
{ Pile_add(&a_pile, n); }
int Apile_count(void)
{ return Pile_count(&a_pile); }

```

Fig. 1. The `pile.h` abstract data type has operations *new*, *add*, *count*, *free*. The `triang.c` client adds the integers $1-n$ to the pile, then counts the pile. The `pile.c` implementation represents a pile as header node (`struct pile`) pointing to a linked list of integers. At bottom, there are two modules that each implement a single “implicit” pile in a module-local global variable: `onepile.c` maintains a pointer to a pile, while `apile.c` maintains a `struct pile` for which it needs knowledge of the representation through `pile_private.h`.

```

(* spec.pile.v *)
(* representation of linked lists in separation logic *)
Fixpoint listrep ( $\sigma$ : list Z) ( $x$ : val) : mpred :=
  match  $\sigma$  with
  |  $h::hs \Rightarrow$  EX  $y$ :val, !! ( $0 \leq h \leq \text{Int.max\_signed}$ ) &&
    data.at Ews tlist (Vint (Int.repr  $h$ ),  $y$ )  $x$ 
    * malloc_token Ews tlist  $x$  * listrep  $hs$   $y$ 
  | nil  $\Rightarrow$  !! ( $x = \text{nullval}$ ) && emp
  end.

(* representation predicate for piles *)
Definition pilerep ( $\sigma$ : list Z) ( $p$ : val) : mpred :=
  EX  $x$ :val, data.at Ews tpile  $x$   $p$  * listrep  $\sigma$   $x$ .

Definition pile.freeable ( $p$ : val) :=
  malloc_token Ews tpile  $p$ .

Definition Pile_new_spec :=
  DECLARE _Pile_new
  WITH  $gv$ : globals
  PRE [ ] PROP() LOCAL( $gvars$   $gv$ ) SEP(mem_mgr  $gv$ )
  POST[ tptr tpile ]
    EX  $p$ : val,
      PROP() LOCAL(temp ret_temp  $p$ )
      SEP(pilerep nil  $p$ ; pile.freeable  $p$ ; mem_mgr  $gv$ ).

Definition Pile_add_spec :=
  DECLARE _Pile_add
  WITH  $p$ : val,  $n$ : Z,  $\sigma$ : list Z,  $gv$ : globals
  PRE [ _p OF tptr tpile, _n OF tint ]
    PROP( $0 \leq n \leq \text{Int.max\_signed}$ )
    LOCAL(temp _p  $p$ ; temp _n (Vint (Int.repr  $n$ )));
     $gvars$   $gv$ )
    SEP(pilerep  $\sigma$   $p$ ; mem_mgr  $gv$ )
  POST[ tvoid ]
    PROP() LOCAL()
    SEP(pilerep ( $n::\sigma$ )  $p$ ; mem_mgr  $gv$ ).

Definition sumlist : list Z  $\rightarrow$  Z := List.fold_right Z.add 0.

Definition Pile_count_spec :=
  DECLARE _Pile_count
  WITH  $p$ : val,  $\sigma$ : list Z
  PRE [ _p OF tptr tpile ]
    PROP( $0 \leq \text{sumlist } \sigma \leq \text{Int.max\_signed}$ ) LOCAL(temp _p  $p$ )
    SEP(pilerep  $\sigma$   $p$ )
  POST[ tint ]
    PROP() LOCAL(temp ret_temp (Vint (Int.repr (sumlist  $\sigma$ ))))
    SEP(pilerep  $\sigma$   $p$ ).

```

Notation key

mpred predicate on memory

EX existential quantifier

!! injects Prop into mpred

&& nonseparating conjunction

data.at $\pi \tau v p$ is $p \mapsto v$,
separation-logic mapsto
at type τ , permission π

malloc_token $\pi \tau x$ represents
“capability to deallocate x ”

Ews the “extern write share”
gives write permission

_Pile_new is a C identifier

WITH quantifies variables
over PRE/POST of funspec

The C function’s return type,
tptr tpile, is “pointer
to **struct pile**”

PROP(...) are pure propositions
on the WITH-variables

LOCAL(... temp _p p ...)
associates C local var **_p**
with Coq value p

gvars gv establishes gv as
mapping from C global
vars to their addresses

SEP(R_1 ; R_2) are separating
conjuncts $R_1 * R_2$

mem_mgr gv represents
different states of the
malloc/free system in
PRE and POST of
any function that
allocates or frees

Fig. 2. Specification of the pile module (Pile_free_spec not shown).

this case the `PROP` asserts that n is between 0 and `max-int`; `LOCAL` asserts¹ that address p is the current value of C variable `_p`, integer n is the value of C variable `_n`, and gv is the global-variable access map. The precondition’s `SEP` clause has two conjuncts: the first one says that there’s a *pile* data structure at address p representing sequence σ ; the second one represents the memory-manager library. The spatial conjunct `(mem_mgr gv)` represents the private data structure of the memory-manager library, that is, the global variables in which the malloc-free system keeps its free lists.

The `SEP` clause of the `POST`condition says that the *pile* at address p now represents the list $n::\sigma$, and that the memory manager is still there.

Verifying that `pile.c`’s functions satisfy the specifications in Fig. 2 using VST-Floyd is done by proving Lemmas like this one (in file `verif-pile.v`):

Lemma `body.Pile_new: semax_body Vprog Gprog f.Pile_new Pile_new.spec.`

Proof. ... (*7 lines of Coq proof script*).... **Qed.**

This says, in the context `Vprog` of global-variable types, in the context `Gprog` of function-specs (for functions that `Pile_new` might call), the function-body `f.Pile_new` satisfies the function-specification `Pile_new.spec`.

Linking

A modular proof of a modular program is organized as follows: CompCert parses each module `M.c` into the AST file `M.v`. Then we write the specification file `spec.M.v` containing funspecs as in Fig. 2. We write `verif.M.v` which imports `spec` files of all the modules from which `M.c` calls functions, and contains `semax_body` proofs of correctness (such as `body.Pile_new` at the end of Sect. 2), for each of the functions in `M.c`.

What’s special about the `main()` function is that its separation-logic precondition has all the initial values of the global variables, merged from the global variables of each module. In `spec_main` we merge the ASTs (global variables and function definitions) of all the `M.v` by a simple, computational, syntactic function. This is illustrated in the Coq files in `VST/progs/pile`.

VST’s main soundness statement is that, when running `main()` in CompCert’s operational semantics, in the initial memory induced from all global-variable initializers, the program is safe and correct—with a notion of partial correctness in interacting with the world via effectful external function calls [13] and returning the “right” value from `main`.

3 Subsumption of Function Specifications

We now turn to the replacement of `pile.c` by a more performant implementation, `fastpile.c`, and its specification—see Fig. 3. As `fastpile.c` employs a differ-

¹ A `LOCAL` clause `temp _p p` asserts that the current value of C local variable `_p` is the Coq value p . If n is a mathematical integer, then `Int.repr n` is its projection into 32-bit machine integers, and `Vint` projects machine integers into the type of scalar C-language values.

```

/* fastpile_private.h */
struct pile { int sum; };

/* fastpile.c */
#include ...
#include "pile.h"
#include "fastpile_private.h"
Pile Pile_new(void)
{ Pile p = (Pile)surely_malloc(sizeof *p); p->sum=0; return p; }
void Pile_add(Pile p, int n)
{ int s = p->sum; if (0 ≤ n && n ≤ INT_MAX-s) p->sum = s+n; }
int Pile_count(Pile p) { return p->sum; }
void Pile_free(Pile p) { free(p); }

(* spec_fastpile.v *)
Definition pilerep ( $\sigma$ : list Z) ( $p$ : val) : mpred :=
  EX  $s$ :Z, !! (0 ≤  $s$  ≤ Int.max_signed ∧ Forall (Z.le 0)  $\sigma$  ∧
    (0 ≤ sumlist  $\sigma$  ≤ Int.max_signed →  $s$ =sumlist  $\sigma$ ))
    && data_at Ews tpile (Vint (Int.repr  $s$ ))  $p$ .

Definition pile_freeable := (* looks identical to the one in fig.2 *)
Definition Pile_new.spec := (* looks identical to the one in fig.2 *)
Definition Pile_add.spec := (* looks identical to the one in fig.2 *)
Definition Pile_count.spec := (* looks identical to the one in fig.2 *)

```

Fig. 3. `fastpile.c`, a more efficient implementation of the pile ADT. Since the only query function is `count`, there’s no need to represent the entire list, just the sum will suffice. In the verification of a client program, the `pilerep` separation-logic predicate has the same signature: `list Z → val → mpred`, even though the representation is a single number rather than a linked list.

ent data representation than `pile.c`, its specification employs a different representation predicate `pilerep`. As `pilerep`’s type remains unchanged, the function specifications look virtually identical²; however, the VST-Floyd proof scripts (in file `verif_fastpile.v`) necessarily differ. Clients importing only the `pile.h` interface, like `onepile.c` or `triang.c`, cannot tell the difference (except that things run faster and take less memory), and are specified and verified only once (files `spec_onepile.v/verif_onepile.v` and `spec_triang.v/verif_triang.v`).

But we may also equip `fastpile.c` with a more low-level specification (see Fig. 4) in which the function specifications refer to a different representation predicate, `countrep`. In reasoning about clients of this low-level interface, we do not need a notion of “sequence”—in contrast to `pilerep` in Fig. 3. The new specification is less abstract than the one in Fig. 3, and closer to the implementation. The subsumption rule (to be introduced shortly) allows us to exploit this relationship:

² Existentially abstracting over the internal representation predicates would further emphasize the uniformity between `fastpile.c` and `pile.c`—a detailed treatment of this is beyond the scope of the present article.

(* spec_fastpile_concrete.v *)

Definition countrep ($s: \mathbb{Z}$) ($p: \text{val}$) : mpred := EX $s': \mathbb{Z}$,
 !! ($0 \leq s \wedge 0 \leq s' \leq \text{Int.max_signed} \wedge (s \leq \text{Int.max_signed} \rightarrow s' = s)$) &&
 data_at Ews tpile (Vint (Int.repr s')) p .

Definition count_freeable ($p: \text{val}$) := malloc_token Ews tpile p .

Definition Pile_new_spec := ...

Definition Pile_add_spec :=
 DECLARE _Pile_add
 WITH $p: \text{val}, n: \mathbb{Z}, s: \mathbb{Z}, gv: \text{globals}$
 PRE [$_p$ OF tpitr tpile, $_n$ OF tint]
 PROP($0 \leq n \leq \text{Int.max_signed}$)
 LOCAL(temp $_p$ p ; temp $_n$ (Vint (Int.repr n)); gvars gv)
 SEP(countrep s p ; mem_mgr gv)
 POST[tvoid]
 PROP() LOCAL() SEP(countrep ($n + s$) p ; mem_mgr gv).

Definition Pile_count_spec := ...

Fig. 4. The fastpile.c implementation could be used in applications that simply need to keep a running total. That is, a *concrete* specification can use a predicate countrep: $\mathbb{Z} \rightarrow \text{val} \rightarrow \text{mpred}$ that makes no assumption about a sequence (list $\mathbb{Z}$). In countrep, the variable s' and the inequalities are needed to account for the possibility of integer overflow.

we only need to explicitly verify the code against the low-level specification and can establish satisfaction of the high-level specification by recourse to subsumption. This separation of concerns extends from VST specifications to model-level reasoning: for example, in our verification of cryptographic primitives we found it convenient to verify that the C program implements a *low-level functional model* and then separately prove that the low-level functional model implements a high-level specification (e.g. cryptographic security).³ In our running example, fastpile.c's low-level functional model is *integer* (the Coq $\mathbb{Z}$ type), and its high level specification is list $\mathbb{Z}$.

To formally state the desired subsumption lemma, observe that notation like DECLARE _Pile_add WITH ... PRE ... POST ... is merely VST's syntactic sugar

³ For example: in our proof of HMAC-DRBG [24], before VST had function-spec subsumption, we had two different proofs of the function f_mbedtls_hmac_drbg_seed, one with respect to a more concrete specification drbg_seed_inst256_spec and one with respect to a more abstract specification drbg_seed_inst256_spec_abs. The latter proof was 202 lines of Coq, at line 37 of VST/hmacdrbg/drbg_protocol_proofs.v in commit 3e61d2991e3d70f5935ae69c88d7172cf639b9bc of <https://github.com/PrincetonUniversity/VST>. Now, instead of reproving the function-body a second time, we have a funspec_sub proof that is only 60 lines of Coq (at line 42 of the same file in commit c2fc3d830e15f4c70bc45376632c2323743858ef).

for a pair that ties the identifier `_Pile.add` to the funspec `WITH...PRE...POST`. For `_Pile.add` we have two such specifications,

`spec_fastpile.Pile.add_spec`: `ident*funspec` (* in Figure 3 *)
`spec_fastpile_concrete.Pile.add_spec`: `ident*funspec` (* in Figure 4 *).

and our notion of *funspec subtyping* will satisfy the following lemma.

Lemma `sub_Pile.add`: `funspec_sub (snd spec_fastpile_concrete.Pile.add_spec)`
`(snd spec_fastpile.Pile.add_spec)`.

and similarly for `Pile.new` and `Pile.count`. Specifically, we permit related specifications to have different `WITH`-lists, in line with Kleymann’s adaptation-complete rule of consequence

$$\frac{\vdash \{P'\}c\{Q'\}}{\vdash \{P\}c\{Q\}} \forall Z. \forall \sigma. PZ\sigma \rightarrow \forall \tau. \exists Z'. (P'Z'\sigma \wedge (Q'Z'\tau \rightarrow QZ\tau))$$

where assertions are binary predicates over auxiliary and ordinary states, and Z, Z' are the `WITH` values.⁴

Our subsumption applies to function specifications, not arbitrary statements c . In the rule for function calls, it ensures that a concretely specified function can be invoked where callers expect an abstractly specified one, just like the subsumption rule of type theory: $\frac{\Gamma \vdash e : \sigma \quad \sigma <: \tau}{\Gamma \vdash e : \tau}$. It is also reflexive and transitive.

Support for Framing. An important principle of separation logic is the frame rule:

$$\frac{\{P\}c\{Q\}}{\{P * R\}c\{P * R\}} (\text{modifiedvars}(c) \cap \text{freevars}(R) = \emptyset)$$

We have found it useful to explicitly incorporate framing in `funspec_sub`, because abstract specifications may have useless data. Consider a function that performs some action (e.g., increment a variable) using some auxiliary data (e.g., an array of 10 integers):

int `incr1(int i, unsigned int *auxdata)` {`auxdata[i%10] += 1; return i+1;`}

The function specification makes clear that the `private` contents of the `auxdata` is, from the client’s point of view, unconstrained; the implementation is free to store anything in this array:

Definition `incr1_spec` := `DECLARE _incr1`
`WITH i: Z, a: val,  $\pi$ : share, private: list val`
`PRE [ $i$  OF tint, _auxdata OF tptr tuint]`
`PROP ( $0 \leq i < \text{Int.max\_signed}$ ; writable_share  $\pi$ )`

⁴ We give Kleymann’s rule for total correctness here. VST is a logic for partial correctness, but its preconditions also guarantee safety; Kleymann’s partial-correctness adaptation rule cannot guarantee safety.

```

    LOCAL(temp _i (Vint (Int.repr i)); temp _auxdata a)
    SEP(data_at sh (tarray tuint 10) private a)
  POST [tint]
    EX private': list val, PROP() LOCAL(temp ret_temp (Vint (Int.repr (i+1))))
      SEP(data_at  $\pi$  (tarray tuint 10) private' a).

```

You might think the `auxdata` is useless, but (i) real-life interfaces often have useless or vestigial fields; and (ii) this might be where the implementation keeps profiling statistics, memoization, or other algorithmically useful information.

Here is a different implementation that should serve any client just as well:

```
int incr2(int i, unsigned int *auxdata) {return i+1;}
```

Its natural specification has an empty SEP clause:

```

Definition incr2_spec := DECLARE _incr2
  WITH i: Z
  PRE [_i OF tint, _auxdata OF tptr tuint]
    PROP (0 ≤ i < Int.max_signed) LOCAL(temp _i (Vint (Int.repr i))) SEP()
  POST [tint]
    PROP() LOCAL(temp ret_temp (Vint (Int.repr (i+1)))) SEP().

```

The *formal* statement that `incr2` serves any client just as well as `incr1` is another case of subsumption:

Lemma `sub_incr12`: `funspec.sub (snd incr2_spec) (snd incr1_spec)`.

In the proof, we use `(data_at  $\pi$  (tarray tuint 10) private a)` as the *frame*.

If the `auxdata` is a global variable instead of a function parameter, all the same principles apply:

```

int global_auxdata[10];
int incr3(int i) {global_auxdata[i%10] += 1; return i+1;}
int incr4(int i) {return i+1;}

```

We define a `funspec` for `incr3` whose SEP clause mentions the `auxdata`, we define a `funspec` for `incr4` whose SEP clause is empty, and we can prove,

Lemma `sub_incr34`: `funspec.sub (snd incr4_spec) (snd incr3_spec)`.

For another example of framing, consider again Fig. 2, the specification of `pile_rep`, `pile_freeable`, `Pile_new_spec`, etc. One might think to combine `pile_freeable` (the memory-deallocation capability) with `pile_rep` (capability to modify the contents) yielding a single combined predicate `pile_rep'`. That way, proofs of client programs would not have to manage two separate conjuncts.

That would work for clients such as `triang.c` and `onepile.c`, but not for `apile.c` which has an initialized global variable (`a_pile`) that satisfies `pile_rep` but *not* `pile_freeable` (since it was not obtained from the `malloc-free` system). Furthermore, the specifications of `pile_add` and `pile_count` do not mention `pile_freeable` in their pre- or postconditions, since they have no need for this capability.

By using `funspec.sub` (with its framing feature), we can have it both ways. One can easily make a more abstract spec in which the `funspecs` of `pile_new`,

`pile_add`, `pile_count`, `pile_free` all take `pilerep'` in their pre- and postconditions; `onepile` and `triang` will still be verifiable using these specs. But in proving `funspec_sub`, therefore, specifications for `pile_add` and `pile_count` now *do* implicitly take `pile_freeable` in their pre- and postconditions, even though they have no use for it; this is the essence of the frame rule.

4 Definitions of Funspec Subtyping

Except in certain higher-order cases, we use this notion of function specification:

`NDmk_funspec` (`f`: `funsig`) (`cc`: `calling_convention`)
 (`A`: `Type`) (`Pre Post`: $A \rightarrow \text{environ} \rightarrow \text{mpred}$): `funspec`.

To construct a *nondependent* (ND) function spec, one gives the function's C-language type signature (`funsig`), the calling convention (usually `cc=cc_default`), the precondition, and the postcondition. A gives the type of variable (or tuple of variables) “shared” between the precondition and postcondition. `Pre` and `Post` are each applied to the shared value of type A , then to a local-variable environment (of type `environ`) containing the formal parameters or result-value (respectively), finally yielding an `mpred`, a spatial predicate on memory.

For example, to specify an increment function with formal parameter `_p` pointing to an integer in memory, we let $A = \text{int}$, so that

$$\text{Pre} = \lambda i : A. \lambda \rho. \rho(_p) \mapsto i \quad \text{and} \quad \text{Post} = \lambda i : A. \lambda \rho. \rho(_p) \mapsto (i + 1).$$

This form suffices for most C programming. But sometimes in the presence of higher-order functions, one wants impredicativity: A may be a tuple of types that includes the type `mpred`. If this is done naively, it cannot typecheck in CiC (there will be universe inconsistencies); see the Appendix.

General Funspec. Higher-order function specs are (mostly) beyond the scope of this paper. When precondition and postcondition must predicate over predicates, we must ensure that each is a *bifunctor*, that is, we must keep track of *covariant* and *contravariant* occurrences, and so on. This approach was outlined by America and Rutten [1] and has been implemented both in Iris [10] and VST.⁵

VST's most general form of function spec is,

Inductive `funspec` :=

`mk_funspec`: `forall` (`f`: `funsig`) (`cc`: `calling_convention`) (`A`: `TypeTree`)
 (`P Q`: `forall ts, dependent_type_functor_rec ts (AssertTT A) mpred`)
 (`P_ne`: `super_non_expansive P`) (`Q_ne`: `super_non_expansive Q`), `funspec`.

Here, `super_non_expansive` is a proof that the precondition (or postcondition) is a nonexpansive (in the step-indexing sense) bifunctor; see the Appendix. The *nondependent* (ND) form of `mk_funspec` shown above is simply a derived form of dependent `mk_funspec`.

⁵ Bifunctor function-specs in VST were the work of Qinxiang Cao, Robert Dockins, and Aquinas Hobor.

Too-Special Funspec Subtyping. Let's consider the obvious notion of funspec subtyping: ϕ_1 is a subtype of ϕ_2 if the precondition of ϕ_2 entails the precondition of ϕ_1 , and the postcondition of ϕ_1 entails the postcondition of ϕ_2 .

Definition `far.too.special.NDfunspec_sub` ($f_1\ f_2 : \text{funspec}$) :=

let $\Delta := \text{funsig_tycontext}(\text{funsig_of_funspec } f_1)$ **in**

match f_1, f_2 **with**

`NDmk.funspec` $f_{sig_1}\ cc_1\ A_1\ P_1\ Q_1$, `NDmk.funspec` $f_{sig_2}\ cc_2\ A_2\ P_2\ Q_2 \Rightarrow$

$f_{sig_1} = f_{sig_2} \wedge cc_1 = cc_2 \wedge A_1 = A_2 \wedge (\forall x : A_1, \Delta, P_2\ \text{nil } x \vdash P_1\ \text{nil } x) \wedge$
 $(\forall x : A_1, (\text{ret0_tycon } \Delta), Q_1\ \text{nil } x \vdash Q_2\ \text{nil } x)$

end.

We write $\Delta, P_2\ \text{nil } x \vdash P_1\ \text{nil } x$, where P_1 and P_2 are the preconditions of f_1 and f_2 , nil expresses that these are nondependent funspecs (no bifunctor structure), and x is the value shared between precondition and postcondition. The type-context Δ provides the additional guarantee that the formal parameters are well typed, and `ret0_tycon` Δ guarantees that the return-value is well typed.

This notion of funspec-sub is sound (w.r.t. subsumption), but barely useful: (1) it requires that the witness types of the two funspecs be the same ($A_1 = A_2$), (2) it doesn't support framing, and (3) it requires $Q_1 \vdash Q_2$ even when P_2 is not satisfied. *Each* of these omissions prevents the practical use of funspec-sub in real verifications, but only (1) and (3) were addressed in previous work [12,20].

Useful, Ordinary Funspec Subtyping. If `NDmk.funspec` were a constructor, we could define,

Definition `NDfunspec_sub` ($f_1\ f_2 : \text{funspec}$) :=

let $\Delta := \text{funsig_tycontext}(\text{funsig_of_funspec } f_1)$ **in**

match f_1, f_2 **with**

`NDmk.funspec` $f_{sig_1}\ cc_1\ A_1\ P_1\ Q_1$, `NDmk.funspec` $f_{sig_2}\ cc_2\ A_2\ P_2\ Q_2 \Rightarrow$

$f_{sig_1} = f_{sig_2} \wedge cc_1 = cc_2 \wedge$

$\forall x_2 : A_2,$

$\Delta, P_2\ \text{nil } x_2 \vdash$

$\text{EX } x_1 : A_1, \text{EX } F : \text{mpred}, (((\lambda \rho. F) * P_1\ \text{nil } x_1) \&\&$

$!! ((\text{ret0_tycon } \Delta), (\lambda \rho. F) * Q_1\ \text{nil } x_1 \vdash Q_2\ \text{nil } x_2))$

end.

Here, each of the three deficiencies is remedied: the witness value $x_1 : A_1$ is existentially derived from $x_2 : A_2$, the frame F is existentially quantified, and the entailment $Q_1 \vdash Q_2$ is conditioned on the precondition P_2 being satisfied.

This version of funspec-sub is, we believe, fully general for `NDmk.funspec`, that is, for function specifications whose witness types A do not contain (covariant or contravariant) occurrences of `mpred`. We present the general, dependent funspec-sub in the Appendix, with its constructor `mk.funspec`, and show the construction of `NDmk.funspec` as a derived form. And actually, since `NDmk.funspec` is not really a constructor (it is a function that applies the constructor `mk.funspec`), we must define `NDfunspec_sub` as a pattern-match on `mk.funspec`; see the Appendix.

5 The Subsumption Rules

The purpose of `funspec.sub` is to support subsumption rules.

Our Hoare-logic judgment takes the form $\Delta \vdash \{P\}c\{Q\}$ where the context Δ describes the types of local and global variables and the funspecs of global functions. We say $\Delta <: \Delta'$ if Δ is at least as strong as Δ' ; in Verifiable C this is written `tycontext.sub  $\Delta$   $\Delta'$` . Again, this relation is reflexive and transitive.

Definition (*glob_specs*): If i is a global identifier, write $(\text{glob_specs } \Delta)!i$ to be the `option(funspec)` that is either `None` or `Some  $\phi$` .

Lemma *funspec.sub.tycontext.sub*: **Suppose** Δ agrees with Δ' on types attributed to global variables, types attributed to local variables, current function return type (if any), and differs only in *specifications* attributed to global functions, in particular: For every global identifier i , if $(\text{glob_specs } \Delta)!i = \text{Some } \phi$ then $(\text{glob_specs } \Delta')!i = \text{Some } \phi'$ and `funspec.sub  $\phi$   $\phi'$` . **Then** $\Delta <: \Delta'$.

Proof. Trivial from the definition of $\Delta <: \Delta'$.

Theorem (*semax_Delta_subsumption*):

$$\frac{\Delta <: \Delta' \quad \Delta' \vdash \{P\}c\{Q\}}{\Delta \vdash \{P\}c\{Q\}}$$

Proof. Nontrivial. Because this is a logic of higher-order recursive function pointers, our Coq proof⁶ in the modal step-indexed model uses the Löb rule to handle recursion, and unfolds our rather complicated semantic definition of the Hoare triple [4].

But this is not the only subsumption rule we desire. Because C has function-pointers, the general function-call rule is for $\Delta \vdash \{P\}e_f(e_1, \dots, e_n)\{Q\}$ where e_f is an expression that evaluates to a function-pointer. Therefore, we cannot simply look up e_f as a global identifier in Δ . Instead, the precondition P must associate the value of e_f with a funspec. Without subsumption, the rules are:

$$\frac{\begin{array}{l} (\text{glob_specs } \Delta)!f = \text{Some } \phi \\ \Delta \vdash f \Downarrow v \\ \Delta \vdash \{\text{func_ptr } v \phi \wedge P\}c\{Q\} \end{array}}{\Delta \vdash \{P\}c\{Q\}} \quad \frac{\begin{array}{l} \Delta \vdash e_f \Downarrow v \\ \Delta \vdash e_1 \Downarrow v_1 \dots \Delta \vdash e_n \Downarrow v_n \\ P * F \vdash \text{func_ptr } v \phi \\ \phi(w) = \{P\}\{Q\} \end{array}}{\Delta \vdash \{P * F\}e_f(e_1, e_2, \dots, e_n)\{Q * F\}}$$

The rule `semax_fun_id` at left says, if the global context Δ associates identifier f with funspec ϕ , and if f evaluates to the address v , then for the purposes of proving $\{P\}c\{Q\}$ we can assume the stronger precondition in which address v has the funspec ϕ .

The `semax_call` rule says, if e_f evaluates to address v , and the precondition factors into conjuncts $P * F$ that imply address v has the funspec ϕ , then choose a

⁶ See `file veric/semax.lemmas.v` in the VST repo.

witness w (for the **WITH** clause), instantiate the witness of ϕ with w , and match the precondition and postcondition of $\phi(w)$ with P and Q ; then the function-call is proved. (Functions can return results, but we don't show that here.)

To turn **semax_call** into a rule that supports subsumption, we simply replace the hypothesis $\phi(w) = \{P\}\{Q\}$ with $\phi <: \phi' \wedge \phi'(w) = \{P\}\{Q\}$.

To reconcile **semax_Delta_subsumption** and **semax_fun_id**, we build $<:$ into the definition of the predicate **func_ptr** $v \phi$, i.e. permit ϕ to be more abstract than the specification associated with address v in VST's semantic model ("rmap").

6 Intersection Specifications

In some of our verification examples, we found it useful to separate different use cases of a function into separate function specifications. One can easily do this using a pattern that discriminates on a boolean value from the **WITH** list jointly in the pre- and postcondition:

WITH $b : \text{bool}, \vec{x} : \vec{\tau}$ **PRE** if b then P_1 else P_2 **POST** if b then Q_1 else Q_2 .

To attach different **WITH**-lists to different cases, we may use Coq's sum type to define a type such as **Variant** $T := \text{case1: int} \mid \text{case2: string}$. and use it in a specification

WITH $\vec{x} : \vec{\tau}, t : T, \vec{y} : \vec{\sigma}$
PRE [...] **match** t **with** **case1** $i \Rightarrow P_1(\vec{x}, i, \vec{y}) \mid \text{case2 } s \Rightarrow P_2(\vec{x}, s, \vec{y})$ **end**
POST [...] **match** t **with** **case1** $i \Rightarrow Q_1(\vec{x}, i, \vec{y}) \mid \text{case2 } s \Rightarrow Q_2(\vec{x}, s, \vec{y})$ **end**.

which amounts to the *intersection* of

WITH $\vec{x} : \vec{\tau}, i : \text{int}, \vec{y} : \vec{\sigma}$ **PRE** [...] $P_1(\vec{x}, i, \vec{y})$ **POST** [...] $Q_1(\vec{x}, i, \vec{y})$ and
WITH $\vec{x} : \vec{\tau}, s : \text{string}, \vec{y} : \vec{\sigma}$ **PRE** [...] $P_2(\vec{x}, s, \vec{y})$ **POST** [...] $Q_2(\vec{x}, s, \vec{y})$.

Generalizing to arbitrary index sets, we may—for a given function signature and calling convention—combine specifications into specification *families*. (We show the nondependent (ND) case; the Coq proofs cover the general case.)

Definition **funspec_Pi_ND** sig cc (I:Type) (A : I $\rightarrow$ Type)
 (Pre Post: forall i, A i $\rightarrow$ environ $\rightarrow$ mpred): funspec := ...

In previous work [5] we showed how relational (2-execution) specifications can be encoded as unary VDM-style specifications. Intersection specifications internalize VDM's "sets of specifications" feature.

The interaction between this construction and subtyping follows precisely that of intersection types in type theory: the lemmas

Lemma **funspec_Pi_ND_sub**: forall fsig cc I A Pre Post i,
 funspec_sub (funspec_Pi_ND fsig cc I A Pre Post)
 (NDmk.funspec fsig cc (A i) (Pre i) (Post i)).

Lemma **funspec_Pi_ND_sub3**: forall fsig cc I A Pre Post g (i:I)
 (H1: forall i, funspec_sub g (NDmk.funspec fsig cc (A i) (Pre i) (Post i))),
 funspec_sub g (funspec_Pi_ND fsig cc I A Pre Post).

are counterparts of the typing rules $\bigwedge_{j \in I} \tau_j <: \tau_i$ (for all $i \in I$) and $\frac{\forall i, \sigma <: \tau_i}{\sigma <: \bigwedge_{i \in I} \tau_i}$, the specializations of which to the binary case appear on page 206 of TAPL [22]. We expect these rules to be helpful for formalizing Leavens and Naumann’s treatment of specification inheritance in object-oriented programs [14].

7 Conclusion

Even without funspec subtyping, separation logic easily expresses data abstraction [21]. But real-world code is modular (as in our running example) and reconfigurable (as in the substitution of `fastpile.c` for `pile.c`). Therefore a notion of specification re-abstraction is needed. We have demonstrated how to extend Kleymann’s notion from commands to functions, and from first-order Hoare logic to higher-order separation logic with framing. We have a full soundness proof for the extended program logic, in Coq. Our `funspec_sub` integrates nicely with our existing proof automation tools and our existing methods of verifying individual modules. As a bonus, one’s intuition that function-specs are like the “types” of functions is borne out by our theorems relating `funspec_sub` to intersection types.

Future Work: When a client module respects data abstraction, such as `onepile.c` and `triang.c` in our example, its Coq proof script does not vary if the implementation of the abstraction changes (such as changing `pile.c` to `fastpile.c`). But our current proofs need to rerun the proof scripts on the modified definition of `pilesep`. As footnote 2 suggests, this could be avoided by the use of existential quantification, in Coq, to describe data abstraction at the C module level.

Acknowledgements. This work was funded by the National Science Foundation under the awards 1005849 (*Verified High Performance Data Structure Implementations*, Beringer) and 1521602 *Expedition in Computing: The Science of Deep Specification*, Appel). We are grateful to the members of both projects for their feedback, and we greatly appreciate the reviewers’ comments and suggestions.

Appendix: Fully General funspec_sub

`NDfunspec_sub` as introduced in Sect. 4 specializes the “real” subtype relation $\phi <: \psi$ in two regards: first, it only applies if ϕ and ψ are of the `NDfunspec` form, i.e. the types of their `WITH`-lists (“witnesses”) are trivial bifunctors as they do not contain co- or contravariant occurrences of `mpred`. Second, it fails to exploit step-indexing and is hence unnecessarily strong.

The technical report (www.cs.princeton.edu/~eberinge/funspec_sub.pdf) contains a brief appendix presenting the fully general `funspec_sub_si`.

References

1. America, P., Rutten, J.: Solving reflexive domain equations in a category of complete metric spaces. *J. Comput. Syst. Sci.* **39**(3), 343–375 (1989)
2. Appel, A.W.: Verification of a cryptographic primitive: SHA-256. *ACM Trans. on Program. Lang. Syst.* **37**(2), 7:1–7:31 (2015)
3. Appel, A.W., Beringer, L., Cao, Q., Dodds, J.: Verifiable C: applying the verified software toolchain to C programs (2019). <https://vst.cs.princeton.edu/download/VC.pdf>
4. Appel, A.W., et al.: *Program Logics for Certified Compilers*. Cambridge University Press, Cambridge (2014)
5. Beringer, L.: Relational decomposition. In: van Eekelen, M., Geuvers, H., Schmaltz, J., Wiedijk, F. (eds.) *ITP 2011. LNCS*, vol. 6898, pp. 39–54. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22863-6_6
6. Beringer, L., Petcher, A., Ye, K.Q., Appel, A.W.: Verified correctness and security of OpenSSL HMAC. In: 24th USENIX Security Symposium, pp. 207–221. USENIX Association, August 2015
7. Cao, Q., Beringer, L., Gruetter, S., Dodds, J., Appel, A.W.: VST-Floyd: a separation logic tool to verify correctness of C programs. *J. Autom. Reason.* **61**(1–4), 367–422 (2018)
8. Cohen, E., et al.: VCC: a practical system for verifying concurrent C. In: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) *TPHOLs 2009. LNCS*, vol. 5674, pp. 23–42. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-03359-9_2
9. Jacobs, B., Smans, J., Philippaerts, P., Vogels, F., Penninckx, W., Piessens, F.: VeriFast: a powerful, sound, predictable, fast verifier for C and Java. In: Bobaru, M., Havelund, K., Holzmann, G.J., Joshi, R. (eds.) *NFM 2011. LNCS*, vol. 6617, pp. 41–55. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20398-5_4
10. Jung, R., Krebbers, R., Jourdan, J.-H., Bizjak, A., Birkedal, L., Dreyer, D.: Iris from the ground up: a modular foundation for higher-order concurrent separation logic. *J. Funct. Program.* **28** (2018)
11. Kirchner, F., Kosmatov, N., Prevosto, V., Signoles, J., Yakobowski, B.: Frama-C: a software analysis perspective. *Formal Aspects Comput.* **27**(3), 573–609 (2015)
12. Kleymann, T.: Hoare logic and auxiliary variables. *Formal Aspects Comput.* **11**(5), 541–566 (1999)
13. Koh, N., et al.: From C to interaction trees: specifying, verifying, and testing a networked server. In: *Proceedings of the 8th ACM SIGPLAN International Conference on Certified Programs and Proofs*, pp. 234–248. ACM (2019)
14. Leavens, G.T., Naumann, D.A.: Behavioral subtyping, specification inheritance, and modular reasoning. *ACM Trans. Program. Lang. Syst.* **37**(4), 13:1–13:88 (2015)
15. Leroy, X.: Formal verification of a realistic compiler. *Commun. ACM* **52**(7), 107–115 (2009)
16. Liskov, B., Wing, J.M.: A behavioral notion of subtyping. *ACM Trans. Program. Lang. Syst.* **16**(6), 1811–1841 (1994)
17. Mansky, W., Appel, A.W., Nogin, A.: A verified messaging system. In: *Proceedings of the 2017 ACM International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA 2017*. ACM (2017)
18. Mitchell, J.C., Plotkin, G.D.: Abstract types have existential type. *ACM Trans. Program. Lang. Syst.* **10**(3), 470–502 (1988)

19. Naumann, D.A.: Deriving sharp rules of adaptation for Hoare logics. Technical report 9906, Department of Computer Science, Stevens Institute of Technology (1999)
20. Nipkow, T.: Hoare logics for recursive procedures and unbounded nondeterminism. In: Bradfield, J. (ed.) CSL 2002. LNCS, vol. 2471, pp. 103–119. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45793-3_8
21. Parkinson, M.J., Bierman, G.M.: Separation logic and abstraction. In: 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2005), pp. 247–258 (2005)
22. Pierce, B.C.: Types and Programming Languages. MIT Press, Cambridge (2002)
23. Pierik, C., de Boer, F.S.: A proof outline logic for object-oriented programming. *Theor. Comput. Sci.* **343**(3), 413–442 (2005)
24. Ye, K.Q., Green, M., Sanguansin, N., Beringer, L., Petcher, A., Appel, A.W.: Verified correctness and security of mbedTLS HMAC-DRBG. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS 2017). ACM (2017)