

Garbage Collector Verification for Proof-Carrying Code

Chun-Xiao Lin, Yi-Yun Chen, Long Li, and Bei Hua

Department of Computer Science and Technology, University of Science and Technology of China, Hefei 230027, P.R.China

E-mail: {cxlin3, liwis}@mail.ustc.edu.cn; {yiyun, bhua}@ustc.edu.cn

Received September 19, 2006; revised March 20, 2007.

Abstract We present the verification of the machine-level implementation of a conservative variant of the standard mark-sweep garbage collector in a Hoare-style program logic. The specification of the collector is given on a machine-level memory model using separation logic, and is strong enough to preserve the safety property of any common mutator program. Our verification is fully implemented in the Coq proof assistant and can be packed immediately as foundational proof-carrying code package. Our work makes important attempt toward building fully certified production-quality garbage collectors.

Keywords program verification, garbage collector, proof-carrying code, program safety

1 Introduction

The latest developments in theorem provers, certifying compilers and other program verification tools have enabled the building of more trustworthy software systems using the *Proof-Carrying Code* (PCC) paradigm^[1]. A PCC package contains the native code of the software, its safety specification, and a machine checkable proof saying that the code behaves as specified. The client can then use a tiny checker to check the safety proof before running the software.

Software verified in a high-level language is hardly trustworthy if the compiler fails to correctly translate the verified program into native code, which is often the case due to the various bugs in the compiler. However, in the PCC style of verification, the specification and proof are directly given on, or translated to, the machine level, and the compiler is thus removed from the *Trusted Computing Base* (TCB). Therefore, the PCC packages are endowed with a higher level of safety.

However, in existing PCC frameworks like the Touchstone-PCC^[1] and *Typed Assembly Language* (TAL)^[2], memory management is either regarded as trusted primitive instructions, or included in the trusted library. On the other hand, automatic memory management, especially *Garbage Collection* (GC)^[3], is acknowledged as complex and error-prone. Thus, the overall safety level of the PCC system will be undermined if an unverified garbage collector is included in TCB.

In this paper, we present the formal verification of a conservative variant^[4] of the standard mark-sweep garbage collector^[3] in the PCC style. The technical highlights and contributions of this paper include:

- Our verification is performed directly on the machine level with the Hoare-style program logic *Stack-based Certified Assembly Programming* (SCAP)^[5]. Thus, we are forced to specify and prove the concrete

machine-level behavior of the collector, which though often abstracted out in a high-level language implementation, is indispensable by the PCC-style mutators for the safe usage of the collector. On the other hand, as we will show later, the major effort we spend in proving the collector is on the manipulation of language independent properties like heap predicates and finite sets. Thus, verifying a low-level language, especially C, implementation of the collector would require almost the same effort. And a certifying compiler, if one could be constructed, is still needed for making the C-level verification useful to other PCC packages.

- We use *separation logic*^[6] to define the machine-level heap predicates for building the specification of the collector. With separation-logic operators, we define new heap predicates to assert reachability on a concrete memory model, which could evade troubles caused by cyclic links in heap objects. Our specification of the collector states that all the live objects are preserved, and all the unreachable objects (in a conservative definition) are collected, hence is strong enough to preserve the safety of common mutator programs.

- Our verification is fully mechanized within the Coq proof assistant^[7]. We follow the Foundational-PCC (FPCC)^[8] style to give the proof of the collector along side the soundness proof of the whole verification framework. The collector's proof in SCAP can be ported to an open FPCC framework^[9] without much difficulty as shown in [9], which enables our collector's possible interoperability with other mutator systems. Our verification can also be used as a model for other PCC systems where the verification of a garbage collector is necessary.

- During our verification of the collector, we make important improvements to the verification framework. We build a sound verification condition generator (VC-Gen) for the SCAP system, which alleviates the user from understanding the details of the SCAP rules, and introduces the potential for automatic verification of basic safety properties. We also build automatic tactics in Coq to deal with proof goals involving separation logic and finite set operations. These improvements greatly

Regular paper.

Supported by the National Natural Science Foundation of China under Grant No.60673126(Verification and Compilation of Software Safety); Intel China Research Center

```

inf_loop() {while (1);}
markbit(x) {
  return (ED+(x-ST)/2);
}
stack_push(ptr) {
  if (top>=buf) inf_loop();
  *(top++)=ptr;
}
stack_pop() {
  return *(--top);
}
stack_empty() {
  return (bot==top);
}

gc() {
  mark_field(root);
  while(!stack_empty()){
    ptr=stack_pop();
    mark_field(ptr->first);
    mark_field(ptr->second);
  }
  for(addr=ST; addr<ED; addr++){
    if (markbit(addr)==WHITE){
      addr->first=freeptr;
      freeptr=addr;
    }
    else markbit(addr)=WHITE;
  }
}

mark_field(val) {
  if (val<ST || val>=ED) return;
  if (val mod 8 != 0) return;
  if (markbit(val)==BLACK) return;
  markbit(val)=BLACK;
  stack_push(val);
}

alloc() {
  if (freeptr==NULL) gc();
  if (freeptr==NULL) loop();
  l=freeptr;
  freeptr=freeptr->first;
  return l;
}

```

Fig. 1. Pseudo code of our collector

simplify the proof construction.

The work presented in this paper is a part of our ongoing project^[10] to build an unified framework for certifying the mutator-collector interaction based on modern garbage collectors like the incremental ones and the generational ones. As an extension to the work in this paper, we have successfully linked a TAL with the verified collector^[11] following the ideas in [9], however, this part is beyond the scope of this paper. We also believe that our improvements to the verification framework will benefit future researches in this field.

The rest of this section is a brief introduction to the collector we verified. Then, we introduce in Section 2 the verification framework and our extension. In Section 3, we discuss the heap predicates that formalize the heap during a collection, and give the safety specification of the collector using these predicates. Section 4 outlines the methodology we used to certify the collector in the SCAP framework. In Section 5, we discuss and evaluate our Coq implementation. Finally, we talk about the related work in Section 6 and give the conclusion in Section 7.

Note that since all the lemmas and theorems in this paper are mechanically proved in Coq, we skip their detailed proofs, interested readers may refer to [12].

1.1 The collector

Fig. 1 shows the pseudo code of the garbage collector we verified, a conservative variant of the standard stop-the-world mark-sweep collector. To simplify the problem, we adopt a heap layout shown in Fig. 2: the size of each heap object is two words; all heap objects reside in a continuous subheap; the collector also keeps the mark bits, a mark stack and a record of global variables. This simplification implies that our allocator only works with mutators that always require two-word objects.

The conservative valid-pointer check in the `mark_field` procedure follows [4]. A value is considered as a valid object pointer only if it is inside the boundaries of the allocatable heap, and aligns to the size of two words.

A GC cycle begins with tracing and marking the ob-




Fig. 2. Collector's heap layout

jects reachable from root. A mark stack is used to temporarily store the marked objects whose fields are to be examined. Then, the collector examines the mark bits of all the objects, reclaims those unreachable objects to form a free list, and resets all the mark bits. The free list can then be used by the allocator when the GC cycle is finished.

2 The verification framework

We discuss in this section the three components of our verification framework: a MIPS-like abstract machine model; a Hoare-style program logic, the SCAP system; and a heap specification language with separation-logic operators.

The whole framework is formalized within a mechanized meta logic, the *Calculus of inductive Construction* (CiC)^[13]. CiC is a higher-order predicate logic extended with inductive definitions. The CiC terms in this paper are written with the standard logic notations. We let *Prop* be the universe of all logical propositions, and let *Set* be the universe of all computational terms.

2.1 Abstract machine

We show the syntax of the abstract machine in Fig. 3. A program \mathbb{P} is a triple of a code heap \mathbb{C} , a machine state \mathbb{S} and an instruction sequence \mathbb{I} . A code heap \mathbb{C} is a map from code label f to instruction sequence \mathbb{I} . A

(Prog)	\mathbb{P}	$::= (\mathbb{C}, \mathbb{S}, \mathbb{I})$
(CdHeap)	\mathbb{C}	$::= \{\mathbf{f} \sim \mathbb{I}\}^*$
(State)	\mathbb{S}	$::= (\mathbb{H}, \mathbb{R})$
(Heap)	\mathbb{H}	$::= \{\mathbf{l} \sim \mathbf{w}\}^*$
(RFile)	\mathbb{R}	$::= \{\mathbf{r} \sim \mathbf{w}\}^*$
(Reg)	\mathbf{r}	$::= \{\mathbf{rk}\}^{k \in \{0 \dots 31\}}$
(Wd, Lab)	\mathbf{w}, \mathbf{f}	$::= 0 \mid 1 \mid 2 \mid 3 \mid \dots$
(Address)	\mathbf{l}	$::= 0 \mid 4 \mid 8 \mid 12 \mid \dots$
(ISeq)	\mathbb{I}	$::= c; \mathbb{I} \mid \text{beq } \mathbf{r}_s, \mathbf{r}_t, \mathbf{f}; \mathbb{I}$ $\mid \text{bne } \mathbf{r}_s, \mathbf{r}_t, \mathbf{f}; \mathbb{I}$ $\mid \text{jf} \mid \text{jalf } \mathbf{f}, \mathbf{f}_{\text{ret}} \mid \text{jr } \mathbf{r}_s$
(Comm)	c	$::= \text{addu } \mathbf{r}_d, \mathbf{r}_s, \mathbf{r}_t \mid \text{addiu } \mathbf{r}_d, \mathbf{r}_s, \mathbf{w}$ $\mid \text{subu } \mathbf{r}_d, \mathbf{r}_s, \mathbf{r}_t \mid \text{srl } \mathbf{r}_d, \mathbf{r}_s, 1$ $\mid \text{sltu } \mathbf{r}_d, \mathbf{r}_s, \mathbf{r}_t \mid \text{andi } \mathbf{r}_d, \mathbf{r}_s, 7$ $\mid \text{lw } \mathbf{r}_d, \mathbf{w}(\mathbf{r}_s) \mid \text{sw } \mathbf{r}_s, \mathbf{w}(\mathbf{r}_d)$

Fig. 3. Abstract machine syntax

machine state \mathbb{S} contains a data heap \mathbb{H} and a register file \mathbb{R} . A data heap \mathbb{H} is a map from address \mathbf{l} (aligns to 4) to word value \mathbf{w} , while a register file \mathbb{R} is a map from register \mathbf{r} to word value, with $\mathbf{r0}$ always maps to 0. We use the standard MIPS register aliases^[14] in the rest of the paper. A command c is a non-controlflow instruction such as a register add or a heap load. An instruction sequence \mathbb{I} , or code block, is a series of commands followed by an unconditional jump instruction. For simplicity, we separate the code heap \mathbb{C} from the mutable data heap \mathbb{H} . Also, we use instruction sequence instead of the standard pc register, and this results in the additional return address \mathbf{f}_{ret} in the jump and link instruction $\text{jalf } \mathbf{f}, \mathbf{f}_{\text{ret}}$, which can be viewed as a macro for the MIPS instruction pair $\text{jalf } \mathbf{f}$ and $\text{j } \mathbf{f}_{\text{ret}}$.

Following [14], we give the small step operational semantics of the abstract machine in Fig. 4. We write $X(z)$ for the value bound to z in the map X , and $X\{z \rightsquigarrow v\}$ for the map obtained by updating the binding of z to v in X . We also write $\mathbb{S}.\mathbb{R}$ for the register file in state \mathbb{S} . Note that for a lw or sw command, if the source address is not in the domain of the heap, the next state is undefined. The next step of a program is undefined if it jumps to a wrong label, or the next state of its initial command is undefined.

2.2 Program logic

The SCAP system^[5] is a Hoare-style program logic for modular verification of assembly code with procedure call/return. Procedures can be verified separately and then linked together to form a verified program. As shown in Fig. 5, a SCAP code specification is a partial correctness specification containing a pair of precondition p and guarantee g . p reassembles the precondition in Hoare logic, while g relates the entry state of the code block to the return state of the corresponding procedure. Thus the g at the entry label of a procedure asserts the guarantee of the procedure, as we will see in the later sections. A code heap specification Ψ maps the labels of the code blocks to their specifications. We also define the implication relation \Rightarrow on SCAP specifications.

if $\mathbb{I} =$	then $(\mathbb{C}, (\mathbb{H}, \mathbb{R}), \mathbb{I}) \mapsto$
$\text{jf } \mathbf{f}$	if $\mathbf{f} \in \text{dom}(\mathbb{C}), (\mathbb{C}, (\mathbb{H}, \mathbb{R}), \mathbb{C}(\mathbf{f}))$
$\text{jalf } \mathbf{f}, \mathbf{f}_{\text{ret}}$	if $\mathbf{f} \in \text{dom}(\mathbb{C}),$ $(\mathbb{C}, (\mathbb{H}, \mathbb{R})\{\mathbf{r}_{31} \rightsquigarrow \mathbf{f}_{\text{ret}}\}), \mathbb{C}(\mathbf{f})$
$\text{jr } \mathbf{r}_s$	if $\mathbb{R}(\mathbf{r}_s) \in \text{dom}(\mathbb{C}),$ $(\mathbb{C}, (\mathbb{H}, \mathbb{R}), \mathbb{C}(\mathbb{R}(\mathbf{r}_s)))$
$\text{beq } \mathbf{r}_s, \mathbf{r}_t, \mathbf{f}; \mathbb{I}'$	if $\mathbb{R}(\mathbf{r}_s) \neq \mathbb{R}(\mathbf{r}_t), (\mathbb{C}, (\mathbb{H}, \mathbb{R}), \mathbb{I}')$, else if $\mathbf{f} \in \text{dom}(\mathbb{C}), (\mathbb{C}, (\mathbb{H}, \mathbb{R}), \mathbb{C}(\mathbf{f}))$
$\text{bne } \mathbf{r}_s, \mathbf{r}_t, \mathbf{f}; \mathbb{I}'$	if $\mathbb{R}(\mathbf{r}_s) = \mathbb{R}(\mathbf{r}_t), (\mathbb{C}, (\mathbb{H}, \mathbb{R}), \mathbb{I}')$, else if $\mathbf{f} \in \text{dom}(\mathbb{C}), (\mathbb{C}, (\mathbb{H}, \mathbb{R}), \mathbb{C}(\mathbf{f}))$
$c; \mathbb{I}'$	if $\text{Next}_c((\mathbb{H}, \mathbb{R})) = \mathbb{S}', (\mathbb{C}, \mathbb{S}', \mathbb{I}')$

if $c =$	then $\text{Next}_c(\mathbb{H}, \mathbb{R}) =$
$\text{addu } \mathbf{r}_d, \mathbf{r}_s, \mathbf{r}_t$	$(\mathbb{H}, \mathbb{R}\{\mathbf{r}_d \rightsquigarrow \mathbb{R}(\mathbf{r}_s) + \mathbb{R}(\mathbf{r}_t)\})$
$\text{addiu } \mathbf{r}_d, \mathbf{r}_s, \mathbf{w}$	$(\mathbb{H}, \mathbb{R}\{\mathbf{r}_d \rightsquigarrow \mathbb{R}(\mathbf{r}_s) + \mathbf{w}\})$
$\text{subu } \mathbf{r}_d, \mathbf{r}_s, \mathbf{r}_t$	$(\mathbb{H}, \mathbb{R}\{\mathbf{r}_d \rightsquigarrow \mathbb{R}(\mathbf{r}_s) - \mathbb{R}(\mathbf{r}_t)\})$
$\text{srl } \mathbf{r}_d, \mathbf{r}_s, 1$	$(\mathbb{H}, \mathbb{R}\{\mathbf{r}_d \rightsquigarrow \mathbb{R}(\mathbf{r}_s)/2\})$
$\text{sltu } \mathbf{r}_d, \mathbf{r}_s, \mathbf{r}_t$	$(\mathbb{H}, \mathbb{R}\{\mathbf{r}_d \rightsquigarrow k\})$ if $\mathbb{R}(\mathbf{r}_s) < \mathbb{R}(\mathbf{r}_t), k = 1$, else $k = 0$
$\text{andi } \mathbf{r}_d, \mathbf{r}_s, 7$	$(\mathbb{H}, \mathbb{R}\{\mathbf{r}_d \rightsquigarrow \mathbb{R}(\mathbf{r}_s) \bmod 8\})$
$\text{lw } \mathbf{r}_d, \mathbf{w}(\mathbf{r}_s)$	if $(\mathbb{R}(\mathbf{r}_s) + \mathbf{w}) \in \text{dom}(\mathbb{H}),$ $(\mathbb{H}, \mathbb{R}\{\mathbf{r}_d \rightsquigarrow \mathbb{H}(\mathbb{R}(\mathbf{r}_s) + \mathbf{w})\})$
$\text{sw } \mathbf{r}_s, \mathbf{w}(\mathbf{r}_d)$	if $(\mathbb{R}(\mathbf{r}_d) + \mathbf{w}) \in \text{dom}(\mathbb{H}),$ $(\mathbb{H}\{\mathbb{R}(\mathbf{r}_d) + \mathbf{w} \rightsquigarrow \mathbb{R}(\mathbf{r}_s)\}, \mathbb{R})$

Fig. 4. Abstract machine semantics

(SPred)	$p, q \in \text{State} \rightarrow \text{Prop}$
(Guar)	$g \in \text{State} \rightarrow \text{State} \rightarrow \text{Prop}$
(BSpec)	$\sigma ::= (p, g)$
(CHSpec)	$\Psi ::= \{\mathbf{l} \rightsquigarrow \sigma\}^*$

$\text{WFST}(0, g, \mathbb{S}, \Psi)$	$\stackrel{\text{def}}{=} \neg \exists \mathbb{S}'. g \mathbb{S} \mathbb{S}'$
$\text{WFST}(n, g, \mathbb{S}, \Psi)$	$\stackrel{\text{def}}{=} \forall \mathbb{S}'. g \mathbb{S} \mathbb{S}' \rightarrow \mathbb{S}'. \mathbb{R}(\mathbf{ra}) \in \text{dom}(\Psi) \wedge$ $p' \mathbb{S}' \wedge \text{WFST}(n-1, g', \mathbb{S}', \Psi)$ where $(p', g') = \Psi(\mathbb{S}'. \mathbb{R}(\mathbf{ra}))$
$(p, g) \Rightarrow (p', g')$	$\stackrel{\text{def}}{=} \forall \mathbb{S}, p \mathbb{S} \rightarrow (p' \mathbb{S} \wedge (\forall \mathbb{S}'. g' \mathbb{S} \mathbb{S}' \rightarrow g \mathbb{S} \mathbb{S}'))$

Fig. 5. SCAP specification constructs

A set of operational-semantics-based inference rules, as partly shown in Fig.6, is provided to build a well-formed program from bottom-up. The abstract stack predicate $\text{WFST}(n, g, \mathbb{S}, \Psi)$ in Fig. 5 generally asserts that the current procedure can return to the label in \mathbf{ra} of its return state. n indicates the number of stack frames, once it becomes 0, the current code will never return. A detailed knowledge of these rules and WFST is not required for understanding the rest of the paper, interested readers may refer to [5].

Lemma 1 states that if a code block is well-formed with some σ' , it is also well-formed with a stronger assert σ , and the proof of a well-formed code block can be lifted from a local Ψ to a global Ψ' .

Lemma 1 (Weakening).

1. If $\sigma \Rightarrow \sigma'$ and $\Psi; \sigma' \vdash \mathbb{I}$, then: $\Psi; \sigma \vdash \mathbb{I}$;
2. If $\Psi \subseteq \Psi'$ and $\Psi; \sigma \vdash \mathbb{I}$, then: $\Psi'; \sigma \vdash \mathbb{I}$.

Theorem 1 ensures that a well-formed program will run safely without stopping at any program step undefined in Fig. 4.

$$\begin{array}{c}
\boxed{\Psi \vdash \mathbb{P}} \text{ (Well-formed Program)} \\
\frac{\Psi \vdash \mathbb{C} : \Psi \quad \text{p } \mathbb{S} \quad \Psi; (\text{p}, \text{g}) \vdash \mathbb{I} \quad \exists n. \text{WFST}(n, \text{g}, \mathbb{S}, \Psi)}{\Psi \vdash (\mathbb{C}, \mathbb{S}, \mathbb{I})} \\
\text{(PROG)} \\
\boxed{\Psi \vdash \mathbb{C} : \Psi'} \text{ (Well-formed Code Heap)} \\
\frac{\Psi'; \Psi(\mathbf{f}) \vdash \mathbb{C}(\mathbf{f}) \quad \forall \mathbf{f} \in \text{dom}(\Psi)}{\Psi' \vdash \mathbb{C} : \Psi} \\
\text{(CDHP)} \\
\boxed{\Psi; (\text{p}, \text{g}) \vdash \mathbb{I}} \text{ (Well-formed Instruction Sequence)} \\
\frac{\Psi; (\text{p}, \text{g}) \vdash \mathbb{I} \\
\forall \mathbb{S}. \text{p } \mathbb{S} \rightarrow \exists \mathbb{S}'. \text{Next}_c(\mathbb{S}) = \mathbb{S}' \wedge \text{p}' \mathbb{S}' \wedge \\
\forall \mathbb{S}''. \text{g}' \mathbb{S}' \mathbb{S}'' \rightarrow \text{g } \mathbb{S} \mathbb{S}''}{\Psi; (\text{p}, \text{g}) \vdash \mathbb{I}} \\
\text{(SEQ)} \\
\frac{\Psi(\mathbf{f}) = (\text{p}', \text{g}') \\
\forall \mathbb{S}. \text{p } \mathbb{S} \rightarrow \text{p}' \mathbb{S} \wedge \forall \mathbb{S}'. \text{g}' \mathbb{S} \mathbb{S}' \rightarrow \text{g } \mathbb{S} \mathbb{S}'}{\Psi; (\text{p}, \text{g}) \vdash \text{j } \mathbf{f}} \\
\text{(J)} \\
\frac{\Psi(\mathbf{f}) = (\text{p}', \text{g}') \quad \Psi(\mathbf{f}_{\text{ret}}) = (\text{p}'', \text{g}'') \\
\forall (\mathbb{H}, \mathbb{R}). \text{p } (\mathbb{H}, \mathbb{R}) \rightarrow \text{p}' (\mathbb{H}, \mathbb{R}\{\mathbf{ra} \rightsquigarrow \mathbf{f}_{\text{ret}}\}) \wedge \\
\forall \mathbb{S}'. \text{g}' (\mathbb{H}, \mathbb{R}\{\mathbf{ra} \rightsquigarrow \mathbf{f}_{\text{ret}}\}) \mathbb{S}' \rightarrow \\
\text{p}'' \mathbb{S}' \wedge \forall \mathbb{S}''. \text{g}'' \mathbb{S}' \mathbb{S}'' \rightarrow \text{g } (\mathbb{H}, \mathbb{R}) \mathbb{S}' \\
\forall (\mathbb{H}, \mathbb{R}), (\mathbb{H}', \mathbb{R}'). \\
\text{g}' (\mathbb{H}, \mathbb{R}) (\mathbb{H}', \mathbb{R}') \rightarrow \mathbb{R}(\mathbf{ra}) = \mathbb{R}'(\mathbf{ra})}{\Psi; (\text{p}, \text{g}) \vdash \text{j al } \mathbf{f}, \mathbf{f}_{\text{ret}}} \\
\text{(CALL)} \\
\frac{\forall \mathbb{S}. \text{p } \mathbb{S} \rightarrow \text{g } \mathbb{S} \mathbb{S}}{\Psi; (\text{p}, \text{g}) \vdash \text{j r } \mathbf{ra}} \\
\text{(RETURN)}
\end{array}$$

Fig. 6. SCAP inference rules (excerpt)

Theorem 1 (Soundness).

If $\Psi \vdash (\mathbb{C}, \mathbb{S}, \mathbb{I})$, for all natural number n there exists a $(\mathbb{C}, \mathbb{S}', \mathbb{I}')$, such that $(\mathbb{C}, \mathbb{S}, \mathbb{I}) \mapsto_n (\mathbb{C}, \mathbb{S}', \mathbb{I}')$.

We extend SCAP by building for it a VCGen, which is partly shown in Fig. 7. $\text{wp}(\Psi, \mathbb{I})$ is a code block specification formed from the code heap specification Ψ and the code block \mathbb{I} . The predicate $\text{rapred}(\Psi, \mathbb{I})$ makes sure that if \mathbb{I} calls some procedure \mathbf{f} , \mathbf{f} will keep \mathbf{ra} unchanged, as required by the CALL rule in Fig. 6. Since it is hard to directly embed rapred into wp , we just make it a standalone predicate. With these definitions, we obtain Theorem 2 from Lemma 1 and the CDHP rule in Fig. 6.

Theorem 2 (VCGen Correctness).

1. If $\sigma \Rightarrow \text{wp}(\Psi, \mathbb{I}) \wedge \text{rapred}(\Psi, \mathbb{I})$, then $\Psi; \sigma \vdash \mathbb{I}$;
2. If $\text{vc}(\Psi, \mathbb{C})$, then $\Psi \vdash \mathbb{C} : \Psi$.

With the help of VCGen, the SCAP proof construction is lifted onto the code block level. For each block, only the proof of a σ implication is required, instead of the proofs of σ implications for each instructions following the rules in Fig. 6. This alleviates the user from understanding the complex details of the SCAP rules, simplifies the proof construction, and makes the reasoning more natural. Besides, by eliminating the need for instantiating the specifications for each instruction inside a code block, the VCGen introduces the potential of automatic proof construction for programs verified against simple properties.

$$\begin{array}{l}
A, B \in \text{Heap} \rightarrow \text{Prop} \\
T \in \text{Set} \\
\mathbf{1} \mapsto \mathbf{w} \stackrel{\text{def}}{=} \lambda \mathbb{H}. \mathbb{H} = \{\mathbf{1} \rightsquigarrow \mathbf{w}\} \\
\text{emp} \stackrel{\text{def}}{=} \lambda \mathbb{H}. \text{dom}(\mathbb{H}) = \emptyset \\
\text{true} \stackrel{\text{def}}{=} \lambda \mathbb{H}. \text{True} \\
A * B \stackrel{\text{def}}{=} \lambda \mathbb{H}. \exists \mathbb{H}_1, \mathbb{H}_2. \\
\quad \mathbb{H}_1 \uplus \mathbb{H}_2 = \mathbb{H} \wedge (A \ \mathbb{H}_1) \wedge (B \ \mathbb{H}_2) \\
\exists x : T. A \stackrel{\text{def}}{=} \lambda \mathbb{H}. \exists x : T. (A \ \mathbb{H}) \\
!(P) \stackrel{\text{def}}{=} \lambda \mathbb{H}. P \wedge (\text{emp } \mathbb{H}) \\
\forall_* x \in \emptyset. A \stackrel{\text{def}}{=} \text{emp} \\
\forall_* x \in \{n\} \cup S. A \stackrel{\text{def}}{=} A[n/x] * \forall_* x \in S - n. A
\end{array}$$

Fig. 8. Separation logic operators

2.3 Heap predicate

To specify the behavior of a collector, we introduce the separation-logic operators. Separation logic was originally proposed as an extension to Hoare logic for reasoning mutable data structures^[6]. In order to employ it in the SCAP system, we directly embed a set of separation-logic operators as heap predicate constructors using CiC, as shown in Fig.8. Thus, the SCAP specifications p and g , which are state predicates, may contain heap specifications build with these predicates.

We write $\mathbb{H} \Vdash A$ if $(A \ \mathbb{H})$ is a valid proposition in CiC. The definitions are consistent with the semantics described in [6]. For simplicity of presentation, we use the same notation to denote the separation-logic connectives and the logic connectives in CiC. We also use the standard separation-logic abbreviations: $\mathbf{1} \mapsto \mathbf{w}_1, \mathbf{w}_2$ for $\mathbf{1} \mapsto \mathbf{w}_1 * \mathbf{1} + \mathbf{4} \mapsto \mathbf{w}_2$, $\mathbf{1} \mapsto -$ for $\exists x : \text{Nat}. \mathbf{1} \mapsto x$ and so on.

The inductively defined iterated separating conjunction $\forall_* x \in S. A$, taken from [15], asserts that the heap can be split into a finite set of subheaps according to the finite set S of natural numbers, while for each member n of S , there is a unique subheap on which $A[n/x]$ holds.

We prove as lemmas the properties of these separation-logic operators (axioms in [6]), and provide Lemma 2 for linking the heap predicates with the operational semantics of the lw and sw instructions.

Lemma 2 (Heap Operations).

1. If $\mathbb{H} \Vdash \mathbf{1} \mapsto \mathbf{w} * \text{true}$, then: $\mathbb{H}(\mathbf{1}) = \mathbf{w}$;
2. If $\mathbb{H} \Vdash \mathbf{1} \mapsto - * A$, then:
$$\mathbb{H}\{\mathbf{1} \rightsquigarrow \mathbf{w}\} \Vdash \mathbf{1} \mapsto \mathbf{w} * A.$$

Most of our separation logic lemmas follow the frame-rule^[6] style to support local reasoning. For example, the universal quantification of the heap predicate A in Lemma 2 ensures that a local heap update never affects the rest of the heap. We also use this kind of heap predicate quantification in our specifications of the collector to achieve local reasoning in sub-procedures.

$$\text{vc}(\Psi, \mathbb{C}) \stackrel{\text{def}}{=} \forall \mathbf{f} \in \text{dom}(\Psi). (\Psi(\mathbf{f}) \Rightarrow \text{wp}(\Psi, \mathbb{C}(\mathbf{f}))) \wedge \text{rapred}(\Psi, \mathbb{C}(\mathbf{f}))$$

if $\mathbb{I} =$	then $\text{wp}(\Psi, \mathbb{I}) =$	in case that
$c; \mathbb{I}'$	$(\lambda \mathbb{S}. \text{p Next}_c(\mathbb{S}), \lambda \mathbb{S}, \mathbb{S}'. \text{g Next}_c(\mathbb{S}) \mathbb{S}')$	$\text{wp}(\Psi, \mathbb{I}') = (\text{p}, \text{g})$
$\text{j } \mathbf{f}$	(p, g)	$\Psi(\mathbf{f}) = (\text{p}, \text{g})$
$\text{jal } \mathbf{f}, \mathbf{f}_{\text{ret}}$	$(\lambda(\mathbb{H}, \mathbb{R}). \text{p}'(\mathbb{H}, \mathbb{R}\{\mathbf{ra} \rightsquigarrow \mathbf{f}_{\text{ret}}\}) \wedge (\forall \mathbb{S}'. \text{g}'(\mathbb{H}, \mathbb{R}\{\mathbf{ra} \rightsquigarrow \mathbf{f}_{\text{ret}}\}) \mathbb{S}' \rightarrow \text{p}'' \mathbb{S}', \lambda(\mathbb{H}, \mathbb{R}), \mathbb{S}''. \exists \mathbb{S}'. \text{g}'(\mathbb{H}, \mathbb{R}\{\mathbf{ra} \rightsquigarrow \mathbf{f}_{\text{ret}}\}) \mathbb{S}' \wedge \text{g}'' \mathbb{S}' \mathbb{S}'')$	$\Psi(\mathbf{f}) = (\text{p}', \text{g}')$ $\Psi(\mathbf{f}_{\text{ret}}) = (\text{p}'', \text{g}'')$
$\text{jr } \mathbf{ra}$	$(\lambda \mathbb{S}. \text{True}, \lambda \mathbb{S}, \mathbb{S}'. \mathbb{S} = \mathbb{S}')$	

if $\mathbb{I} =$	then $\text{rapred}(\Psi, \mathbb{I}) =$	in case that
$c; \mathbb{I}'$	$\text{rapred}(\Psi, \mathbb{I}')$	
$\text{j } \mathbf{f}$	True	
$\text{jal } \mathbf{f}, \mathbf{f}_{\text{ret}}$	$\forall(\mathbb{H}, \mathbb{R}), (\mathbb{H}', \mathbb{R}'). \text{g}'(\mathbb{H}, \mathbb{R}) (\mathbb{H}', \mathbb{R}') \rightarrow \mathbb{R}(\mathbf{ra}) = \mathbb{R}'(\mathbf{ra})$	$\Psi(\mathbf{f}) = (\text{p}', \text{g}')$
$\text{jr } \mathbf{ra}$	True	

Fig. 7. The VCGen (excerpt)

null	$::=$	0	
st, ed	$::=$	$8 \mid 16 \mid 24 \mid \dots$	
ptrs	$\stackrel{\text{def}}{=}$	$\{1 \mid (1 \bmod 8 = 0) \wedge (\text{st} \leq 1 < \text{ed})\}$	
$\text{vptr}(1)$	$\stackrel{\text{def}}{=}$	$1 \in \text{ptrs}$	
$\text{rchrt}(\mathbb{H}, \mathbb{R}, 1)$	$\stackrel{\text{def}}{=}$	$\text{reach}(\mathbb{H}, \mathbb{R}(\text{t0}), 1)$	
		$\frac{\text{vptr}(1)}{\text{reach}(\mathbb{H}, 1, 1)}$	(REFL)
		$\frac{\text{vptr}(1) \text{ vptr}(1') \text{ reach}(\mathbb{H}, 1', 1')}{\mathbb{H}(1) = 1'' \vee \mathbb{H}(1+4) = 1''}$	(NEXT)
$\text{ok_val}(S, \mathbf{w})$	$\stackrel{\text{def}}{=}$	$\mathbf{w} \in \text{ptrs} \rightarrow \mathbf{w} \in S$	
$\text{ok_fld}(S, 1)$	$\stackrel{\text{def}}{=}$	$\exists \mathbf{w}. !(\text{ok_val}(S, \mathbf{w})) * 1 \mapsto \mathbf{w}$	
$\text{ok_obj}(S, 1)$	$\stackrel{\text{def}}{=}$	$\text{ok_fld}(S, 1) * \text{ok_fld}(S, 1+4)$	
$\text{obj_hp}(S', S)$	$\stackrel{\text{def}}{=}$	$\forall *x \in S. \text{ok_obj}(S', x)$	

Fig. 9. Specification interface

3 Specification

We present in this section the SCAP specification of the collector we proved. First, we give the reachability predicate on which our specification is build. Then, we model the collector's heap with separation-logic-based heap predicates. Finally, we give the specification of the collector's interface, the `alloc` procedure.

3.1 The specification interface

We define in Fig. 9 the heap conceptions that should be consistent between the collector and the mutator. The constant address `null` is set to be 0, while the lower and upper boundaries of the collector's allocatable heap, `st` and `ed`, should both align to 8, which is the size of a heap object. A value `1` is a valid pointer (`vptr(1)`) only if it is the address of an allocatable heap object.

The reachability predicate `reach`($\mathbb{H}, 1, 1'$) is inductively defined. In the base case, a valid pointer is self-reachable. And in the inductive case, $1'$ is reachable from 1 if it is reachable from the pointers in the heap object at 1 . The predicate `rchrt`($\mathbb{S}, 1$) asserts that 1 points

to a live heap object (reachable from the root) in the state \mathbb{S} . For the sake of simplicity, we consider the case of a single root register `t0`, which can be extended to support more root registers without much difficulty.

The heap predicate `obj_hp`(S', S) in Fig. 9 asserts an object heap with reachability information. A value \mathbf{w} is valid with a pointer set S (`ok_val`(S, \mathbf{w})) if it is either a member of S or a non-pointer value. `ok_fld`($S, 1$) asserts that the field value at 1 is valid with S . A valid object (`ok_obj`($S, 1$)) is composed of two valid fields. Thus, the proposition $\mathbb{H} \Vdash \text{obj_hp}(S', S)$ is true only if \mathbb{H} is formed exactly with the heap objects in S , and all its field values are valid with S' . It is noteworthy that once $\mathbb{H} \Vdash \text{obj_hp}(S, S)$ holds, \mathbb{H} will be a *closed* heap with no outgoing pointers. The relationship between `obj_hp` and `reach` is stated in Lemma 3.

Lemma 3 (Object Heap Reachability).

If $\mathbb{H} \Vdash \text{obj_hp}(S, S)$, $1 \in S$, and `reach`($\mathbb{H}, 1, 1'$), then $1' \in S$.

By virtue of the `obj_hp` predicate, we can avoid the problem caused by the possible cyclic links in the object heap when the `reach` predicate is directly used. This benefit greatly simplifies our verification of the mark phase of the garbage collection.

3.2 The collector's heap

The object heap in the collector's heap layout (described in Section 1.1) is formalized with the `obj_hp` predicate introduced in the previous subsection.

The rest of the formalization of the collector's heap is listed in Fig. 10. The free list is inductively defined by the set of free objects. The heap that storing the mark bits n for a objects set S is specified with `mbits`(S, n), which asserts that there is a mark bit n for each member x of S at the address $(\text{ed} + (x - \text{st})/2)$. The mark stack `mstk`(S, x, y, z) is represented by a continuous memory area, with the lower part containing pointers to a set of objects. The collector's record `gcinfo` stores three pointers to the collector's internal data structures. `flist`(\mathbb{R}, S) and `mstack`(\mathbb{R}, S) are used when the pointers are loaded to the corresponding registers. The free list header is

```

eq( $\mathbb{H}$ )  $\stackrel{\text{def}}{=} \lambda \mathbb{H}'. \mathbb{H}' = \mathbb{H}$ 
flst( $1, \emptyset$ )  $\stackrel{\text{def}}{=} !(1 = \text{null})$ 
flst( $1, \{1\} \cup S$ )  $\stackrel{\text{def}}{=} !(1 \neq \text{null}) * \exists 1'. 1 \mapsto -, 1' * \text{flst}(1', S-1)$ 
hdr( $1$ )  $\stackrel{\text{def}}{=} (\text{ed} + (1 - \text{st})/2)$ 
mbits( $S, n$ )  $\stackrel{\text{def}}{=} \forall_* x \in S. \text{hdr}(x) \mapsto n$ 
array_set( $1, \emptyset$ )  $\stackrel{\text{def}}{=} \text{emp}$ 
array_set( $1, \{w\} \cup S$ )  $\stackrel{\text{def}}{=} 1 \mapsto w * \text{array\_set}(1+4, S-w)$ 
buffer( $1, 1'$ )  $\stackrel{\text{def}}{=} \forall_* x \in \{x \mid x \bmod 4 = 0 \wedge 1 \leq x < 1'\}. x \mapsto -$ 
mstk( $S, x, y, z$ )  $\stackrel{\text{def}}{=}$ 
   $!(y - x = \text{size}(S)) * \text{array\_set}(x, S) * \text{buffer}(y, z)$ 
gcinfo( $1, l_1, l_2, l_3$ )  $\stackrel{\text{def}}{=} 1 \mapsto l_1 * 1 + 4 \mapsto l_2 * 1 + 8 \mapsto l_3$ 
flist( $\mathbb{R}, S$ )  $\stackrel{\text{def}}{=} \text{flst}(\mathbb{R}(\mathbf{t5}), S)$ 
mstack( $\mathbb{R}, S$ )  $\stackrel{\text{def}}{=} \text{mstk}(S, \mathbb{R}(\mathbf{t3}), \mathbb{R}(\mathbf{t2}), \mathbb{R}(\mathbf{t4}))$ 

```

Fig. 10. Auxiliary heap definitions

stored in $\mathbf{t5}$, and the stack pointers (top , bot and buf) are stored in $\mathbf{t2}$, $\mathbf{t3}$ and $\mathbf{t4}$.

3.3 The specification of alloc

We formalize the SCAP specification of the collector's main procedure `alloc` in Fig.11. The precondition p_{ALLOC} is defined in terms of the predicate gc_inv . The first conjunct of $\text{gc_inv}(\mathbb{S}, \mathbb{H}, \mathbf{l}_f)$ asserts that the allocatable object set ptrs is composed of the allocated object set B and the free object set F . The following two heap propositions ensure that the heap contains an allocated subheap \mathbb{H} , a free list with head at \mathbf{l}_f , an empty mark stack and the record of pointers. This faithfully follows the heap layout in Fig. 2.

The guarantee g_{ALLOC} comprises two parts. In the first part, $\text{reg_ok}(\mathbb{S}, \mathbb{S}')$ asserts that the MIPS callee-saved registers ($\mathbf{s0}$ to $\mathbf{s7}$, and \mathbf{ra}), together with the root register $\mathbf{t0}$ and the record pointer $\mathbf{t1}$, are identical in the two states.

The second part asserts that gc_inv is preserved in \mathbb{S} and \mathbb{S}' . If the free list is empty ($\mathbf{l}_f = 0$) at the entry state of `alloc`, a garbage collection will remove all unreachable objects in \mathbb{H} . The allocated subheap \mathbb{H}' at the return state thus contains exactly the untouched live object heap \mathbb{H}_T plus a new object at $\mathbb{S}'.\mathbb{R}(\mathbf{v0})$. Otherwise, \mathbb{H}' is simply formed by extending \mathbb{H} with a new object if the free list is not empty in \mathbb{S} .

A mutator program accesses only the live parts of the allocated object heap. A proper specification of a mutator only asserts the values or data structures on this subheap. Since the guarantee g_{ALLOC} ensures that the live object heap is preserved between the entry/return states of the collector, the validity of any mutator side specification is preserved. In this sense, our GC specification is strong enough to preserve the safety of any common mutator programs. Note that the specification given here is based on the fact that a mark-sweep collector never moves objects, thus it is not suitable for a copying collector^[3]. The reader may find a specification

```

gc_inv( $(\mathbb{H}, \mathbb{R}), \mathbb{H}_T, \mathbf{l}_f$ )  $\stackrel{\text{def}}{=} \exists B, F, \mathbf{l}_b, \mathbf{l}_t.
  B \cup F = \text{ptrs} \wedge
  \mathbb{H} \Vdash \text{eq}(\mathbb{H}_T) * \text{flst}(\mathbf{l}_f, F) * \text{mbits}(\text{ptrs}, 0) *
  \text{mstk}(\emptyset, \mathbf{l}_b, \mathbf{l}_b, \mathbf{l}_t) * \text{gcinfo}(\mathbb{R}(\mathbf{t1}), \mathbf{l}_b, \mathbf{l}_t, \mathbf{l}_f) \wedge
  \mathbb{H}_T \Vdash \text{obj\_hp}(\text{ptrs}, B)$ 
saved_regs  $\stackrel{\text{def}}{=} \{\mathbf{s0} \dots \mathbf{s7}, \mathbf{ra}, \mathbf{t0}, \mathbf{t1}\}$ 
reg_ok( $(\mathbb{H}, \mathbb{R}), (\mathbb{H}', \mathbb{R}')$ )  $\stackrel{\text{def}}{=} \forall \mathbf{r} \in \text{saved\_regs}. \mathbb{R}(\mathbf{r}) = \mathbb{R}'(\mathbf{r})$ 

p_{\text{ALLOC}}  $\stackrel{\text{def}}{=} \lambda \mathbb{S}. \exists \mathbb{H}, \mathbf{l}_f. \text{gc\_inv}(\mathbb{S}, \mathbb{H}, \mathbf{l}_f)$ 
g_{\text{ALLOC}}  $\stackrel{\text{def}}{=} \lambda \mathbb{S}, \mathbb{S}'. \text{reg\_ok}(\mathbb{S}, \mathbb{S}') \wedge
  \forall \mathbb{H}, \mathbf{l}_f. \text{gc\_inv}(\mathbb{S}, \mathbb{H}, \mathbf{l}_f) \rightarrow \exists \mathbb{H}', \mathbf{l}'_f. \text{gc\_inv}(\mathbb{S}', \mathbb{H}', \mathbf{l}'_f) \wedge
  (\mathbf{l}_f = 0 \rightarrow \exists \mathbb{H}_T. \text{dom}(\mathbb{H}_T) = \{1, 1+4 \mid \text{rchrt}(\mathbb{S}, 1)\}) \wedge
  \mathbb{H} \Vdash \text{eq}(\mathbb{H}_T) * \text{true} \wedge
  \mathbb{H}' \Vdash \text{eq}(\mathbb{H}_T) * \mathbb{S}'.\mathbb{R}(\mathbf{v0}) \mapsto -, -) \wedge
  (\mathbf{l}_f \neq 0 \rightarrow \mathbb{H}' \Vdash \text{eq}(\mathbb{H}) * \mathbb{S}'.\mathbb{R}(\mathbf{v0}) \mapsto -, -)$ 

```

Fig. 11. Collector specification

for a copying collector in [15].

4 Proof construction

The SCAP well-formedness proof of the collector is constructed as follows. For each procedure, we firstly form a local code heap specification Ψ_L by composing the specifications for all its code blocks and the procedures it calls. Then, we generate for each code block \mathbb{I} the $\text{wp}(\Psi_L, \mathbb{I})$, and build their well-formedness proof using Theorem 2.

After verifying each procedure with its local Ψ_L , we sum up the specifications to form a global Ψ_G . With Lemma 1 and the CDHP rule in Fig. 6, we prove the well-formedness of the collector's code heap with Ψ_G . A mutator verified in SCAP can then link to this code heap through Lemma 1 and the CDHP rule.

The tricky part of the proof construction is to figure out the correct specification of the collector's internal code blocks. We present these specifications in the following subsection. After that, we briefly discuss the theorem proving issue.

4.1 Code heap specification

The SCAP specifications of the collector's code blocks follow the pattern of $(\text{p}_{\text{ALLOC}}, \text{g}_{\text{ALLOC}})$ in Fig. 11. Unfamiliar readers may just get the general idea instead of going through the complex specifications in Fig. 12, 13, 14, and 15.

We use the following shorthand throughout this section. $\text{nin_rid}(\mathbb{R}, \mathbb{R}', R)$ asserts that any register which is not a member of the register set R is identical in \mathbb{R} and \mathbb{R}' . $\text{sted_ok}(\mathbb{R})$ asserts that the allocatable heap boundaries are loaded into the corresponding registers in \mathbb{R} . We store st in $\mathbf{k0}$ and ed in $\mathbf{k1}$.

```

nin_rid( $\mathbb{R}, \mathbb{R}', R$ )  $\stackrel{\text{def}}{=} \forall \mathbf{r} \notin R. \mathbb{R}(\mathbf{r}) = \mathbb{R}'(\mathbf{r})$ 
sted_ok( $\mathbb{R}$ )  $\stackrel{\text{def}}{=} \mathbb{R}(\mathbf{k0}) = \text{st} \wedge \mathbb{R}(\mathbf{k1}) = \text{ed}$ 

```

Mark stack operations. We show the assembly code with specifications of the stack-related procedures

<pre> is_empty:($\mathbb{P}_{EMP}, \mathbb{G}_{EMP}$) beq \$t3 \$t2 empty2 addiu \$v0 \$0 0 jr \$ra empty2 : addiu \$v0 \$0 1 jr \$ra </pre>	<pre> push:($\mathbb{P}_{PUSH}, \mathbb{G}_{PUSH}$) addiu \$at \$t2 4 sltu \$at \$t4 \$1 bne \$at \$0 stk_loop sw \$a0 0(\$t2) jr \$ra </pre>	<pre> pop:($\mathbb{P}_{POP}, \mathbb{G}_{POP}$) addiu \$v0 \$0 4 subu \$t2 \$t2 \$v0 lw \$v0 0(\$t2) jr \$ra </pre>	<pre> mark_field:($\mathbb{P}_{MFLD}, \mathbb{G}_{MFLD}$) sltu \$t9 \$a0 \$k0 bne \$t9 \$0 return sltu \$t9 \$a0 \$k1 beq \$t9 \$0 return subu \$at \$a0 \$k0 andi7 \$t9 \$at bne \$t9 \$0 return srl1 \$at \$at addu \$at \$at \$k1 lw \$t9 0 \$at addiu \$v1 \$0 1 beq \$t9 \$v1 return sw \$v1 0 \$at j push </pre>
$\mathbb{P}_{EMP} \stackrel{\text{def}}{=} \lambda(\mathbb{H}, \mathbb{R}). \exists S. \mathbb{H} \Vdash \text{mstack}(\mathbb{R}, S) * \text{true}$ $\mathbb{G}_{EMP} \stackrel{\text{def}}{=} \lambda(\mathbb{H}, \mathbb{R}), (\mathbb{H}', \mathbb{R}'). \exists n. (\mathbb{H}', \mathbb{R}') = (\mathbb{H}, \mathbb{R}\{v0 \rightsquigarrow n\}) \wedge$ $\forall S. \mathbb{H} \Vdash \text{mstack}(\mathbb{R}, S) * \text{true} \rightarrow$ $(n = 0 \rightarrow S = \emptyset) \wedge (n \neq 0 \rightarrow S \neq \emptyset)$	$\mathbb{P}_{MFLD} ::= \lambda(\mathbb{H}, \mathbb{R}). \text{sted_ok}(\mathbb{R}) \wedge$ $\neg \text{vptr}(\mathbb{R}(\mathbf{a0})) \vee$ $\text{vptr}(\mathbb{R}(\mathbf{a0})) \wedge \exists G.$ $(\mathbb{R}(\mathbf{a0}) \in G \wedge \mathbb{H} \Vdash \text{mbits}(G, 1) * \text{mstack}(\mathbb{R}, G) * \text{true}) \vee$ $(\mathbb{R}(\mathbf{a0}) \notin G \wedge$ $\mathbb{H} \Vdash \text{mbits}(G, 1) * \text{mstack}(\mathbb{R}, G) * \text{hdr}(\mathbb{R}(\mathbf{a0})) \mapsto - * \text{true})$	$\mathbb{P}_{PUSH} \stackrel{\text{def}}{=} \lambda(\mathbb{H}, \mathbb{R}). \exists S. \mathbb{R}(\mathbf{a0}) \notin S \wedge \mathbb{H} \Vdash \text{mstack}(\mathbb{R}, S) * \text{true}$ $\mathbb{G}_{PUSH} \stackrel{\text{def}}{=} \lambda(\mathbb{H}, \mathbb{R}), (\mathbb{H}', \mathbb{R}'). \text{nin_rid}(\mathbb{R}, \mathbb{R}', \{\mathbf{t2}, \mathbf{at}\}) \wedge$ $\forall S, A. \mathbb{H} \Vdash \text{mstack}(\mathbb{R}, S) * A \rightarrow$ $\mathbb{H}' \Vdash \text{mstack}(\mathbb{R}', S + \mathbb{R}(\mathbf{a0})) * A$	$\mathbb{P}_{POP} \stackrel{\text{def}}{=} \lambda(\mathbb{H}, \mathbb{R}). \exists S \neq \emptyset. \mathbb{H} \Vdash \text{mstack}(\mathbb{R}, S) * \text{true}$ $\mathbb{G}_{POP} \stackrel{\text{def}}{=} \lambda(\mathbb{H}, \mathbb{R}), (\mathbb{H}', \mathbb{R}'). \text{nin_rid}(\mathbb{R}, \mathbb{R}', \{\mathbf{t2}, \mathbf{v0}\}) \wedge$ $\forall S \neq \emptyset, A. \mathbb{H} \Vdash \text{mstack}(\mathbb{R}, S) * A \rightarrow$ $\mathbb{R}'(\mathbf{v0}) \in S \wedge \mathbb{H}' \Vdash \text{mstack}(\mathbb{R}', S - \mathbb{R}'(\mathbf{v0})) * A$

Fig. 12. Assembly code with specifications: stack and mark_field

(`is_empty`, `push` and `pop`) in Fig. 12. Each precondition generally asserts that there is a mark stack in the entry state. The guarantees ensure that if the preconditions are satisfied, the intended operations on the stack are performed. Note that the universally quantified heap predicate A in the guarantees ensures the frame-rule style local reasoning.

Mark an object. The `mark_field` procedure is also shown in Fig. 12. The precondition \mathbb{P}_{MFLD} requires that either the argument $\mathbb{R}(\mathbf{a0})$ is a non-pointer value, or the heap predicate contains enough information for checking the mark bit of $\mathbb{R}(\mathbf{a0})$ and performing stack push. The guarantee \mathbb{G}_{MFLD} ensures that the procedure only modifies a few registers if the argument $\mathbb{R}(\mathbf{a0})$ is a non-pointer value or a pointer to a marked object. Otherwise, \mathbb{G}_{MFLD} guarantees that $\mathbb{R}(\mathbf{a0})$ is marked and pushed to the mark stack.

The mark phase. We show in Fig. 13 the code blocks implementing the mark phase of the collection. The specifications of these code blocks follow the same pattern. Each precondition asserts that the corresponding state predicates (`mpre`, `minv`, `mfst`, or `msnd`) is satisfied. On the other hand, each guarantee ensures that if the corresponding state predicate is satisfied, the `mark` procedure will return on a state satisfying `mpost`.

On the entry state of `mark`, as asserted by `mpre`, all objects are allocated and the heap boundaries are loaded into the corresponding registers. After marking the root register `t0`, we reach the state where the mark loop invariant `minv` holds. Following the tri-color abstraction^[3], `minv` asserts that the object set `ptrs` is divided into the sets of *black*, *gray* and *white* objects, which are separated by their mark bits and the mark stack. The *black* and *gray* objects are reachable from root, and the reachability between the three sets are

asserted by the `obj_hp` predicates on \mathbb{H}_T .

In the mark loop, a *gray* object is popped into `v0`, and the code blocks `mfirst` and `msecond` examine the two fields of this object. The predicate `mfst` makes no different than `minv` except for splitting out the object in `v0` from the *gray* object set G . The predicate `msnd` goes one step further by asserting that if the first field of the object in `v0` is a valid pointer, it will point to a marked object. The call to `mark_field` in `msecond` returns to `mloop`, where we can safely move the object in `v0` into the *black* object set B and get `minv` again.

There is no more *gray* object on the return state of `mark`. Thus we get a closed subheap in \mathbb{H}_T formed with the live objects in B .

The sweep phase. In Fig. 14, we specify the first two code blocks of `sweep` in same way we specify the `mark` code blocks. The predicate `spre` on the entry state of `sweep` is different from `mpost` only in that we move the unreachable objects out of the allocated subheap \mathbb{H}_T .

In the sweep loop invariant `sinv`, G stands for the objects before the sweep pointer in `a0`. The unreachable objects in G are collected into the free list with set F , and the mark bits of the reachable objects in G are set to 0. Finally, G becomes a superset of `ptrs` when the sweep loop finishes. Thus all the unreachable objects will be collected into the free list, as asserted by `spost`. The other two code block specifications are simply generated by the `wp` function in Section 2.2.

The collector. The rest of the codes, including the `gc` procedure and the `alloc` procedure, are shown in Fig. 15. The specification of `gc` is a trivial composition of the `mark` and `sweep` specifications. The guarantee \mathbb{G}_{GC} asserts that the reachable subheap \mathbb{H}'_T is preserved during the collection. The specifications of `alloc` is listed in Fig. 11, and the other specifications are generated by

$\text{mark}:(\mathcal{P}_{\text{MARK}}, \mathcal{G}_{\text{MARK}})$ <code>addiu \$t6 \$ra 0</code> <code>addiu \$a0 \$t0 0</code> <code>jal mark_field mloop</code>	$\text{mloop}:(\mathcal{P}_{\text{MLOOP}}, \mathcal{G}_{\text{MLOOP}})$ <code>jal is_empty mloop2</code> $\text{mloop2}:$ <code>bne \$v0 \$0 return1</code> <code>jal pop mfirst</code>	$\text{mfirst}:(\mathcal{P}_{\text{MFST}}, \mathcal{G}_{\text{MFST}})$ <code>lw \$a0 0(\$v0)</code> <code>jal mark_field msecond</code> <code>jal pop mfirst</code>	$\text{msecond}:(\mathcal{P}_{\text{MSND}}, \mathcal{G}_{\text{MSND}})$ <code>lw \$a0 4(\$v0)</code> <code>jal mark_field mloop</code>
$\mathcal{P}_{\text{MARK}} \stackrel{\text{def}}{=} \lambda S. \exists \mathbb{H}. A. \text{mpre}(S, \mathbb{H}, A)$ $\mathcal{G}_{\text{MARK}} \stackrel{\text{def}}{=} \lambda S, S'. (\forall \mathbb{H}. A. \text{mpre}(S, \mathbb{H}, A) \rightarrow \text{mpost}(S', \mathbb{H}, A)) \wedge$ $\text{nin_rid}(\{\mathbf{t6}, \mathbf{a0}, \mathbf{at}, \mathbf{v0}, \mathbf{t9}, \mathbf{t2}, \mathbf{v1}\}, S.\mathbb{R}, S'.\mathbb{R})$	$\mathcal{P}_{\text{MLOOP}} \stackrel{\text{def}}{=} \lambda S. \exists \mathbb{H}. A. \text{minv}(S, \mathbb{H}, A)$ $\mathcal{G}_{\text{MLOOP}} \stackrel{\text{def}}{=} \lambda S, S'. (\forall \mathbb{H}. A. \text{minv}(S, \mathbb{H}, A) \rightarrow \text{mpost}(S', \mathbb{H}, A)) \wedge$ $\text{nin_rid}(\{\mathbf{ra}, \mathbf{t6}, \mathbf{a0}, \mathbf{at}, \mathbf{v0}, \mathbf{t9}, \mathbf{t2}, \mathbf{v1}\}, S.\mathbb{R}, S'.\mathbb{R}) \wedge$ $S'.\mathbb{R}(\mathbf{ra}) = S.\mathbb{R}(\mathbf{t6})$	$\mathcal{P}_{\text{MFST}} \stackrel{\text{def}}{=} \lambda S. \exists \mathbb{H}. A. \text{mfst}(S, \mathbb{H}, A)$ $\mathcal{G}_{\text{MFST}} \stackrel{\text{def}}{=} \lambda S, S'. (\forall \mathbb{H}. A. \text{mfst}(S, \mathbb{H}, A) \rightarrow \text{mpost}(S', \mathbb{H}, A)) \wedge$ $\text{nin_rid}(\{\mathbf{ra}, \mathbf{t6}, \mathbf{a0}, \mathbf{at}, \mathbf{v0}, \mathbf{t9}, \mathbf{t2}, \mathbf{v1}\}, S.\mathbb{R}, S'.\mathbb{R}) \wedge$ $S'.\mathbb{R}(\mathbf{ra}) = S.\mathbb{R}(\mathbf{t6})$	$\text{mpre}((\mathbb{H}, \mathbb{R}), \mathbb{H}_T, A) \stackrel{\text{def}}{=} \text{sted_ok}(\mathbb{R}) \wedge$ $\mathbb{H} \Vdash \text{eq}(\mathbb{H}_T) * \text{mbits}(\text{ptrs}, 0) * \text{mstack}(\mathbb{R}, \emptyset) * A \wedge$ $\mathbb{H}_T \Vdash \text{obj_hp}(\text{ptrs}, \text{ptrs})$
$\mathcal{P}_{\text{MSND}} \stackrel{\text{def}}{=} \lambda S. \exists \mathbb{H}. A. \text{msnd}(S, \mathbb{H}, A)$ $\mathcal{G}_{\text{MSND}} \stackrel{\text{def}}{=} \lambda S, S'. (\forall \mathbb{H}. A. \text{msnd}(S, \mathbb{H}, A) \rightarrow \text{mpost}(S', \mathbb{H}, A)) \wedge$ $\text{nin_rid}(\{\mathbf{ra}, \mathbf{t6}, \mathbf{a0}, \mathbf{at}, \mathbf{v0}, \mathbf{t9}, \mathbf{t2}, \mathbf{v1}\}, S.\mathbb{R}, S'.\mathbb{R}) \wedge$ $S'.\mathbb{R}(\mathbf{ra}) = S.\mathbb{R}(\mathbf{t6})$	$\text{minv}((\mathbb{H}, \mathbb{R}), \mathbb{H}_T, A) \stackrel{\text{def}}{=} \exists B, G, W.$ $\text{sted_ok}(\mathbb{R}) \wedge B \cup G \cup W = \text{ptrs} \wedge$ $\text{ok_val}(B \cup G, \mathbb{R}(\mathbf{t0})) \wedge (\forall x \in (B \cup G). \text{reach}(\mathbb{H}_T, \mathbb{R}(\mathbf{t0}), x)) \wedge$ $\mathbb{H} \Vdash \text{eq}(\mathbb{H}_T) * \text{mbits}(B \cup G, 1) * \text{mbits}(W, 0) * \text{mstack}(\mathbb{R}, G) * A \wedge$ $\mathbb{H}_T \Vdash \text{obj_hp}(B \cup G, B) * \text{obj_hp}(\text{ptrs}, G \cup W)$	$\text{minv2}((\mathbb{H}, \mathbb{R}), \mathbb{H}_T, A, S_1, S_2, B, G, W) \stackrel{\text{def}}{=} \text{sted_ok}(\mathbb{R}) \wedge B \cup G \cup W \cup \{\mathbb{R}(\mathbf{v0})\} = \text{ptrs} \wedge$ $\text{ok_val}(B \cup G \cup \{\mathbb{R}(\mathbf{v0})\}, \mathbb{R}(\mathbf{t0})) \wedge$ $(\forall x \in B \cup G \cup \{\mathbb{R}(\mathbf{v0})\}. \text{reach}(\mathbb{H}_T, \mathbb{R}(\mathbf{t0}), x)) \wedge$ $\mathbb{H} \Vdash \text{eq}(\mathbb{H}_T) * \text{mbits}(B \cup G \cup \{\mathbb{R}(\mathbf{v0})\}, 1) * \text{mbits}(W, 0) * \text{mstack}(\mathbb{R}, G) * A \wedge$ $\mathbb{H}_T \Vdash \text{obj_hp}(B \cup G \cup \{\mathbb{R}(\mathbf{v0})\}, B) * \text{obj_hp}(\text{ptrs}, G \cup W) * \text{ok_fld}(S_1, \mathbb{R}(\mathbf{v0})) * \text{ok_fld}(S_2, \mathbb{R}(\mathbf{v0}) + 4)$	$\text{mfst}(S, \mathbb{H}_T, A) \stackrel{\text{def}}{=} \exists B, G, W. \text{minv2}(S, \mathbb{H}_T, A, \text{ptrs}, \text{ptrs}, B, G, W)$
$\text{mpost}((\mathbb{H}, \mathbb{R}), \mathbb{H}_T, A) \stackrel{\text{def}}{=} \exists B, W. \text{sted_ok}(\mathbb{R}) \wedge B \cup W = \text{ptrs} \wedge$ $\text{ok_val}(B, \mathbb{R}(\mathbf{t0})) \wedge (\forall x \in B. \text{reach}(\mathbb{H}_T, \mathbb{R}(\mathbf{t0}), x)) \wedge$ $\mathbb{H} \Vdash \text{eq}(\mathbb{H}_T) * \text{mbits}(B, 1) * \text{mbits}(W, 0) * \text{mstack}(\mathbb{R}, \emptyset) * A \wedge$ $\mathbb{H}_T \Vdash \text{obj_hp}(B, B) * \text{obj_hp}(\text{ptrs}, W)$	$\text{msnd}(S, \mathbb{H}_T, A) \stackrel{\text{def}}{=} \exists B, G, W.$ $\text{minv2}(S, \mathbb{H}_T, A, B \cup G \cup \{S.\mathbb{R}(\mathbf{v0})\}, \text{ptrs}, B, G, W)$	$\text{mfirst}(\mathbb{H}, \mathbb{R}, \mathbb{H}_T, A) \stackrel{\text{def}}{=} \exists B, G, W. \text{minv2}(S, \mathbb{H}_T, A, \text{ptrs}, \text{ptrs}, B, G, W)$	$\text{msecond}(\mathbb{H}, \mathbb{R}, \mathbb{H}_T, A) \stackrel{\text{def}}{=} \exists B, G, W.$ $\text{minv2}(S, \mathbb{H}_T, A, B \cup G \cup \{S.\mathbb{R}(\mathbf{v0})\}, \text{ptrs}, B, G, W)$

Fig. 13. Assembly code with specifications: mark

the wp function.

4.2 Theorem proving

With the VCGen in Section 2.2, the well-formedness proofs of the code blocks are constructed without the direct use of SCAP rules. Instead, the main difficulty in the proof construction is to solve the domain specific problem. For our collector verification, the problem is to prove those properties about the finite set and the heap predicate, as the specifications are mainly constructed with assertions about them.

We tackle this problem by breaking the proof construction into two steps. First, we prove as lemmas the common properties of the finite set and the heap predicate, such as the introduction/elimination rules of the set operators, and the associative/communicative properties of the separating conjunction $*$. These are collected in two lemma libraries for the finite set and the heap predicate. Both the finite set and the heap predicate contain non-trivial inductive definitions hence it is infeasible to implement fully automatic decision procedures for either of them. However, there are still many patterns in using these lemmas during our proof construction. Following these patterns, we design automatic tactics to simplify the proof construction, an example of such tactics is shown in the next section.

With the well-designed domain specific lemma libraries, the well-formedness proofs of the code blocks are constructed with ease.

5 Coq implementation

Our verification is fully mechanized within the Coq proof assistant^[7]. Coq is an interactive theorem prover based on the formulae-as-type notion^[16], where theorems and proofs are constructed in CiC as types and terms, respectively. Proof checking in Coq is thus type checking of terms in CiC, which is simple to implement and more trustworthy than the methodologies used by the other systems like PVS.

Coq provides a rich language which is able to define both logical and computational terms using inductive constructors and pattern matching. With this language, we build the abstract machine, the SCAP rule set, and other components for specifying the collector.

Theorem proving in Coq is goal-directed and tactic-based. A set of predefined tactics is provided to solve the proof goals. This includes the basic tactics like `apply`, `intro` and `elim`, which transform the goal according to the corresponding CiC typing rules; an `auto` tactic, which tries to apply previous lemmas declared as hints; and several predefined tactics representing some well-established decision procedures, such as the `omega` tactic, which solves a goal in Presburger arithmetic. A tactic language is also provided to build user-defined tactics using pattern matching on proof goals, recursion and other programming facilities. Our tactics for the finite set and the heap predicate are defined using this language.

In the following part of this section, we illustrate our

<pre> sweep:(P_{SWEEP}, G_{SWEEP}) addiu \$a0 \$k0 0 addiu \$t5 \$0 0 j sloop </pre>	<pre> sloop:(P_{SLOOP}, G_{SLOOP}) sltu \$at \$a0 \$k1 beq \$at \$0 return subu \$at \$a0 \$k0 srl1 \$at \$at addu \$at \$at \$k1 </pre>	<pre> l w \$v0 0(\$at) beq \$v0 \$0 sadd sw \$0 0(\$at) j snext </pre>	<pre> sadd:(P_{SADD}, G_{SADD}) sw \$t5 4(\$a0) addiu \$t5 \$a0 0 j snext return: jr \$ra </pre>	<pre> snext:(P_{SNXT}, G_{SNXT}) addiu \$a0 \$a0 8 j sloop return1: addiu \$ra \$t6 0 jr \$ra </pre>
$P_{SWEEP} \stackrel{\text{def}}{=} \lambda S. \exists H, A. \text{spre}(S, H, A)$ $G_{SWEEP} \stackrel{\text{def}}{=} \lambda S, S'. \text{nin_rid}(\{a0, t5, at, v0\}, S.R, S'.R) \wedge \forall H, A. \text{spre}(S, H, A) \rightarrow \text{spost}(S', H, A)$ $P_{SLOOP} \stackrel{\text{def}}{=} \lambda S. \exists H, A. \text{sinv}(S, H, A)$ $G_{SLOOP} \stackrel{\text{def}}{=} \lambda S, S'. \text{nin_rid}(\{a0, t5, at, v0\}, S.R, S'.R) \wedge \forall H, A. \text{sinv}(S, H, A) \rightarrow \text{spost}(S', H, A)$ $(P_{SNXT}, G_{SNXT}) \stackrel{\text{def}}{=} \text{wp}(\{\text{sloop} \rightsquigarrow (P_{SLOOP}, G_{SLOOP})\}, \mathbb{I}_{SNXT})$ $(P_{SADD}, G_{SADD}) \stackrel{\text{def}}{=} \text{wp}(\{\text{snxt} \rightsquigarrow (P_{SNXT}, G_{SNXT})\}, \mathbb{I}_{SADD})$ $\text{spost}((H, R), H_T, A) \stackrel{\text{def}}{=} \exists B, W. \text{sted_ok}(R) \wedge B \cup W = \text{ptrs} \wedge \text{ok_val}(B, R(t0)) \wedge (\forall x \in B. \text{reach}(H_T, R(t0), x)) \wedge H \Vdash \text{eq}(H_T) * \text{flist}(R, W) * \text{mbits}(\text{ptrs}, 0) * \text{mstack}(R, \emptyset) * A \wedge H_T \Vdash \text{obj_hp}(B, B)$	$\text{spre}((H, R), H_T, A) \stackrel{\text{def}}{=} \exists B, W. \text{sted_ok}(R) \wedge B \cup W = \text{ptrs} \wedge \text{ok_val}(B, R(t0)) \wedge (\forall x \in B. \text{reach}(H_T, R(t0), x)) \wedge H \Vdash \text{eq}(H_T) * \text{obj_hp}(\text{ptrs}, W) * \text{mstack}(R, \emptyset) * \text{mbits}(B, 1) * \text{mbits}(W, 0) * A \wedge H_T \Vdash \text{obj_hp}(B, B)$ $\text{sinv}((H, R), H_T, A) \stackrel{\text{def}}{=} \exists B, W, F, G. \text{sted_ok}(R) \wedge \text{ok_val}(B, R(t0)) \wedge G = \{x \mid x \in \text{ptrs} \wedge x < R(a0)\} \wedge W \cap G = \emptyset \wedge F \subseteq G \wedge (R(a0) - \text{st}) \bmod 8 = 0 \wedge B \cup W \cup F = \text{ptrs} \wedge (\forall x \in B. \text{reach}(H_T, R(t0), x)) \wedge H \Vdash \text{eq}(H_T) * \text{obj_hp}(\text{ptrs}, W) * \text{flist}(R, F) * \text{mstack}(R, \emptyset) * \text{mbits}(W \cup F \cup (B \cap G), 0) * \text{mbits}(B - G, 1) * A \wedge H_T \Vdash \text{obj_hp}(B, B)$			

Fig. 14. Assembly code with specifications: `sweep`

<pre> gc:(P_{GC}, G_{GC}) addiu \$t7 \$ra 0 jal mark gc2 gc2: addiu \$ra \$t7 0 j sweep </pre>	<pre> alloc:(P_{ALLOC}, G_{ALLOC}) addiu \$t8 \$ra 0 addiu \$k0 \$0 st addiu \$k1 \$0 ed lw \$t3 0(\$t1) lw \$t4 4(\$t1) </pre>	<pre> addu \$t2 \$0 \$t3 lw \$t5 8(\$t1) bne \$t5 \$0 retp jal gc aloop </pre>	<pre> retp:(P_{RETP}, G_{RETP}) addiu \$ra \$t8 0 addiu \$v0 \$t5 0 sw \$t5 8(\$t1) jr \$ra </pre>	<pre> aloop:(P_{ALOOP}, G_{ALOOP}) beq \$t5 \$0 loop j retp loop: j loop </pre>
$P_{GC} \stackrel{\text{def}}{=} \lambda S. \exists H, A. \text{gcpre}(S, H, A)$ $G_{GC} \stackrel{\text{def}}{=} \lambda S, S'. (\forall H, A. \text{gcpre}(S, H, A) \rightarrow \text{gcpost}(S', H, A)) \wedge \text{nin_rid}(\{t6, t7, t5, a0, at, v0, t9, v1\}, S.R, S'.R)$ $\text{gcpre}((H, R), H_T, A) \stackrel{\text{def}}{=} \text{mpre}((H, R\{t7 \rightsquigarrow R(ra)\}), H_T, A)$	$(P_{RETP}, G_{RETP}) \stackrel{\text{def}}{=} \text{wp}(\{\}, \mathbb{I}_{RETP})$ $(P_{ALOOP}, G_{ALOOP}) \stackrel{\text{def}}{=} \text{wp}(\{\text{loop} \rightsquigarrow (\lambda S. \text{True}, \lambda S, S'. \text{False}), \text{retp} \rightsquigarrow (P_{RETP}, G_{RETP})\}, \mathbb{I}_{ALOOP})$ $\text{gcpost}(S, H_T, A) \stackrel{\text{def}}{=} \exists H'_T. \text{spost}(S, H'_T, A) \wedge H_T \Vdash \text{eq}(H'_T) * \text{true}$			

Fig. 15. Assembly code with specifications: `gc` and `alloc`

Coq implementation methodology with several examples: the state transition function for commands in the abstract machine, the SCAP rule set for well-formed instruction sequence, the `wp` function in VCGen, and a user-defined tactic for solving goals with heap predicates. Interested readers may get more details of our Coq implementation from [12].

5.1 Verification framework

The `nextc` function in Fig.16 is defined in Section 3.1 as the small step state transition relation for commands. It is implemented as a function with pattern matching. The tags such as `Some` and `None` on the return value of such functions are the constructors of an option type, like `option state` in this function. A value with an option type is either a meaningful value with the tag `Some`, or a bad value represented by `None`. For example, a `None` returned by the heap read function `hget` stands for a bad heap value, while a `None` returned by the program step function stands for a stuck program.

We implement the SCAP well-formed instruction se-

quence rule set as a recursion function which performs pattern matching on the instruction sequence and returns different proof obligations based on the rules in Fig.6. The `wp` function in the VCGen follows the same pattern to generate the corresponding specifications.

5.2 Tactics for automation

The most complicated part of the collector's proof lies in heap manipulation with separation logic, that is to say, to prove that one proposition with heap predicates implies another. For complex heap predicates with many separating conjunctions, like the ones mentioned in Section 4.1, this will be enormously difficult using only the communicative, associative and monotonic properties described in [6]. Usually hundreds of lines are needed to prove one such goal, and it is not uncommon to encounter several such goals in proving one code block.

We build a tactic `simplesep` that automatically matches and eliminates the identical or trivial parts of the heap predicates in the hypothesis and proof goal, leaving the different ones for manually proving. For

```

Definition nextc(c:comm)(s:state): option state :=
  match c with
  | addu rd rs rt =>
    Some (rset s rd (rget s rs + rget s rt))
  ...
  | lw rt offset rs =>
    match (hget s (rget s rs + offset)) with
    | Some v => Some (rset s rt v)
    | None => None
    end
  end.

end.

Fixpoint scap_iseq_ok(P:chspect)(p:pre)
  (g:guar)(is:iseq){struct is}: Prop :=
  match is with
  | jr r31 => forall s, p s -> g s s
  ...
  | seq c is' => exists p', exists g',
    (forall s, p s ->
      match (nextc c s) with
      | Some s' => p' s' /\
        forall s'', g' s' s'' -> g s s''
      | None => False
      end) /\ scap_iseq_ok P p' g' is'
  end.

end.

Fixpoint wp(P:chspect)(is:iseq){struct is}
  : option (pre, guar) :=
  match is with
  | jr r31 => (fun s=>True, fun s s' =>s=s')
  ...
  | seq c is' =>
    match (wp P is') with
    | Some (p', g') => Some (fun s =>
      match (nextc c s) with
      | Some s'=>p' s' | None=>False
      end, fun s s'' =>
      match (nextc c s) with
      | Some s'=>g' s' s'' | None=>False
      end)
    | None => None
    end
  end
end.

```

Fig. 16. Coq implementation

example, the following goal is trivially provable with `simplsep`.

$$\forall \mathbb{H}, A, B, C. \mathbb{H} \Vdash A * B * C * \text{emp} \rightarrow \mathbb{H} \Vdash C * \text{true} * B$$

In most of the cases, a code block changes only a small part of the heap. This tactic alleviates us from complicated reasoning, and greatly reduces the proof script size. The principle of `simplsep` is to regard the heap predicates in the goal and hypotheses as two lists, and prove that one is a permutation of another using the classic algorithm. During this procedure, the identical parts are eliminated using the monotonic lemma of separating conjunction. The running time of the tactic is bounded by the size of the heap predicates.

We also implement other automatic tactics for reasoning about heap lookup and update, the manipulation of finite sets, auto rewriting on complicate states values generated by VCGen and so on. The detailed examples

Lines	Component
816	Basic properties and tactics
1387	Ordered heap library
451	Abstract machine encoding and lemmas
1263	Ordered finite set library
844	Separation logic library
547	SCAP, VCGen and related tactics
529	Collector's heap definitions and lemmas
1930	Code, specification and proof of the collector

Fig. 17. Proof script size

of these tactics can be found in [12].

5.3 Evaluation

Fig. 17 lists the proof script size of our Coq implementation, in terms of the number of non-empty lines. The work takes several man months for programmers familiar with the Coq system. The intensive use of automatic tactics and the VCGen results in a 3/4 drop of the proof script size and makes the proof much easier to follow. The benefit of these facilities is also demonstrated by our evaluation on the CDSA example in [5]. The `malloc/free` functions are verified within two days by a single person, while each proof script has a length of only 1/6 when comparing to their original ones.

Our implementation relies heavily on the Coq standard library. We also extend these libraries with various lemmas and tactics. The heap model and the finite set is designed with the ideas in [17], where a heap/set is a list of data (as in the Coq `ListSet` library) together with a proof that the list is well-ordered and with no redundancy. Thus we obtain the heap/set extensionally from the Law of Exclude Middle, and the `Fixpoint` functions on sets, such as the iterated separating conjunction and the mark stack predicate, can be built with ease.

6 Related work

Intensive efforts have been dedicated to the building of trustworthy PCC systems [1, 2, 8, 5, 9]. Our work complements theirs by introducing verified garbage collector into PCC, which will reduce the overall TCB.

The work on the formal correctness proof of a garbage collector dates back to [18, 19], both with informal proofs. Researches on mechanized verification of GC algorithms^[20, 21, 22] focus mostly on algorithms with abstract memory model, and some are implemented with model checking. Unlike these work, our collector is verified as a machine-level implementation, which forces us to employ a more concrete specification and thus our verified program is more trustworthy to run directly on a real machine.

The recent work by Birkedal *et al.*^[15] use the axiom based separation logic to reason a copying collector. They give a formal paper proof of the Cheney collector^[3] against a safety specification based on heap isomorphism. The iterated separating conjunction on

finite set, which is intensively used in our work, comes from this paper.

Prior work on type-safe garbage collection^[23, 24, 25] mainly focus on including garbage collection into the mutator's type system. As a result, for reasoning each of the non-trivial safety properties, complex type systems must be constructed and their soundness be proved, which are unlikely easy tasks. On the other hand, as analyzed in Section 3.3, our collector specification is strong enough to preserve any common safety properties of the mutator program.

7 Conclusion

We demonstrate in this paper the mechanical verification of a conservative variant of the standard mark-sweep garbage collector in the PCC style. The specification of our collector is given on a machine-level memory model using separation logic. Our verification is fully implemented in Coq, and can be packed immediately as an FPCC package. The verification can also be used as a model for other PCC systems where the verification of a garbage collector is required. We also make non-trivial improvement to the verification framework, which is helpful for the future researches.

The work presented in this paper is a part of our ongoing project for building an SCAP-based framework to verify the mutator-collector interaction, where we adopt an unified basic safety specification interface for any of the copying, non-copying or incremental collectors. We also utilize the OCAP^[9] framework as the platform for verifying the safe interaction between TAL and the collector discussed in this paper. The natural choice of future work includes verifying more efficient collectors, such as incremental and generational ones, and their interaction with various mutator systems.

References

- [1] G. Necula. Proof-carrying code. In *Proc. 24th ACM Symp. on Principles of Prog. Lang.*, pages 106–119, New York, January 1997. ACM Press.
- [2] G. Morrisett, D. Walker, K. Crary, and N. Glew. From system F to typed assembly language. *ACM Trans. on Prog. Lang. and Sys.*, 21(3):527–568, 1999.
- [3] R. E. Jones. *Garbage Collection: Algorithms for Automatic Dynamic Memory Management*. Wiley, Chichester, July 1996. With a chapter on Distributed Garbage Collection by R. Lins.
- [4] H. Boehm and M. Weiser. Garbage collection in an uncooperative environment. *Software Practice and Exp.*, 18(9):807–820, 1988.
- [5] X. Feng, Z. Shao, A. Vaynberg, S. Xiang, and Z. Ni. Modular verification of assembly code with stack-based control abstractions. In *Proc. 2006 ACM conf. on Prog. Lang. Design and Impl.*, Ottawa, Canada, pages 401–414, June 2006. ACM Press.
- [6] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proc. 17th IEEE Symp. on Logic in Comp. Sci.*, pages 55–74, Washington, DC, USA, 2002. IEEE Comp. Soc..
- [7] Coq Development Team. The Coq proof assistant reference manual. Coq release v8.0, October 2005.
- [8] A. W. Appel. Foundational proof-carrying code. In *Proc. 16th IEEE Symp. on Logic in Comp. Sci.*, pages 247–258. IEEE Comp. Soc., June 2001.
- [9] X. Feng, Z. Ni, Z. Shao, and Y. Guo. An open framework for foundational proof-carrying code. In *The third ACM SIGPLAN Workshop on Types in Lang. Design and Impl.*, pages 67–78, Nice, France, January 2007. ACM Press.
- [10] A. McCreight, Z. Shao, C. Lin, and L. Li. A General Framework for Certifying Garbage Collectors and Their Mutators. In *2007 ACM SIGPLAN conf. on Prog. Lang. Design and Impl.*, paper to appear, San Diego, CA, USA, June 2007. ACM Press.
- [11] C. Lin, A. McCreight, Z. Shao, Y. Chen, and Y. Guo. Foundational Typed Assembly Language with Certified Garbage Collection. In *1st IEEE & IFIP International Symp. on Theoretical Aspects of Soft. Eng.*, pages (to appear), Shanghai, China, June 2007. IEEE Comp. Soc..
- [12] C. Lin, Y. Chen, L. Li, and B. Hua. Garbage collector verification for proof-carrying code (documents and coq implementation). <http://sug.ustcsz.edu.cn/~cxlin/gcpaper/>, 2006.
- [13] C. Paulin-Mohring. Inductive definitions in the system Coq—rules and properties. In *In Proc. 1st Int. Conf. on Typed Lambda Calculi and Applications*, Utrecht, The Netherlands, volume 664 of *Lecture Notes in Computer Science*, pages 328–345, 1993. Springer-Verlag.
- [14] MIPS Technologies, Inc. MIPS32™ Architecture For Programmers Volume II: The MIPS32™ Instruction Set, v2.50.
- [15] L. Birkedal, N. Torp-Smith, and J. C. Reynolds. Local reasoning about a copying garbage collector. In *Proc. 31st ACM Symp. on Principles of Prog. Lang.*, pages 220–231, New York, NY, USA, 2004. ACM Press.
- [16] W. A. Howard. The formulas-as-types notion of construction. In *To H. B. Curry: Essays on Combinatory Logic, Lambda Calculus, and Formalism*, pages 479–490. Academic Press, 1980.
- [17] N. Marti, R. Affeldt, and A. Yonezawa. Formal verification of the heap manager of an operating system using separation logic. In *Proc. ICFEM2006*, volume 4260 of *Lecture Notes in Computer Science*, pages 225–244, Canberra, September 1998. Springer-Verlag.
- [18] E. W. Dijkstra, L. Lamport, A. J. Martin, C. S. Scholten, and E. F. M. Steffens. On-the-fly garbage collection: an exercise in cooperation. *Commun. ACM*, 21(11):966–975, 1978.
- [19] M. Ben-Ari. Algorithms for on-the-fly garbage collection. *ACM Trans. on Prog. Lang. and Sys.*, 6(3):333–344, 1984.
- [20] D. M. Russinoff. A mechanically verified incremental garbage collector. *Formal Aspects of Computing*, 6:359–390, 1994.
- [21] P. Jackson. Verifying a garbage collection algorithm. In *Proc. of 11th Int'l Conference on Theorem Proving in Higher Order Logics*, volume 1479 of *Lecture Notes in Computer Science*, Canberra, Australia, pages 225–244, 2006. Springer-Verlag.
- [22] L. Burdy. B vs. Coq to prove a garbage collector. In R. J. Boulton and P. B. Jackson, editors, *14th Int'l*

Conference on Theorem Proving in Higher Order Logics, Edinburgh, UK, pages 85–97, September 2001.

- [23] D. C. Wang and A. W. Appel. Type-preserving garbage collectors. In *Proc. 28th ACM Symp. on Principles of Prog. Lang.*, pages 166–178, New York, NY, USA, 2001. ACM Press.
- [24] S. Monnier, B. Saha, and Z. Shao. Principled scavenging. In *Proc. 2001 ACM conf. on Prog. Lang. Design and Impl.*, pages 81–91, New York, 2001. ACM Press.
- [25] C. Hawblitzel, H. Huang, L. Wittie, and J. Chen. A garbage-collecting typed assembly language. In *The third ACM SIGPLAN Workshop on Types in Language Design and Implementation*, pages 41–52, Nice, France, January 2007. ACM Press.