Zha JP, Feng XY, Qiao L. Modular verification of SPARCv8 code. JOURNAL OF COMPUTER SCIENCE AND TECHNO-LOGY 35(6): 1382–1405 Nov. 2020. DOI 10.1007/s11390-020-0536-9

Modular Verification of SPARCv8 Code

Jun-Peng Zha^{1,2}, Xin-Yu Feng^{1,2,*}, Member, CCF, ACM, and Lei Qiao³, Member, CCF

¹Department of Computer Science and Technology, Nanjing University, Nanjing 210023, China ²State Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China ³Beijing Institute of Control Engineering, Beijing 100080, China

E-mail: jpzha@smail.nju.edu.cn; xyfeng@nju.edu.cn; fly2moon@aliyun.com

Received April 11, 2020; revised October 24, 2020.

Abstract Inline assembly code is common in system software to interact with the underlying hardware platforms. The safety and correctness of the assembly code is crucial to guarantee the safety of the whole system. In this paper, we propose a practical Hoare-style program logic for verifying SPARC (Scalable Processor Architecture) assembly code. The logic supports modular reasoning about the main features of SPARCv8 ISA (instruction set architecture), including delayed control transfers, delayed writes to special registers, and register windows. It also supports relational reasoning for refinement verification. We have applied it to verify that there is a contextual refinement between a context switch routine in SPARCv8 and a switch primitive. The program logic and its soundness proof have been mechanized in Coq.

Keywords Scalable Processor Architecture Version 8 (SPARCv8), assembly code verification, context switch, Coq, refinement verification

1 Introduction

Operating system kernels are at the most foundational layer of computer software systems. To interact directly with hardware, many important components in OS kernels are implemented in assembly, such as the context switch code or the code that manages interrupts. In addition, there is also code written directly in assembly to achieve high performance (e.g., memcpy in linux v2.6.17.10⁽¹⁾). Correctness of such assembly code is crucial to ensure the safety and security of the whole system. However, assembly code verification remains a challenging task in existing work on OS kernel verification (e.g., [1-3]), where the assembly code is either unverified or verified based on operational semantics without a general program logic.

SPARC (Scalable Processor Architecture) is a CPU instruction set architecture (ISA) with highperformance and great flexibility⁽²⁾. It has been widely used in various processors for workstations and embedded systems. SPARCv8 ISA has some interesting features, which make it a non-trivial task to design a Hoare-style program logic for assembly code.

• Delayed Control Transfers. SPARCv8 has two program counters pc and npc. The npc register points to the next instruction to run. Control-transfer instructions in SPARCv8 change npc instead of pc to the target program point, while pc takes the original value of npc. This makes the control transfer happen one cycle later than the execution of the control transfer instructions.

• Delayed Writes. The wr instruction that writes a special class of registers such as the window invalid mask register wim does not take effect immediately. That is, "they may take until completion of the third instruction following the write instruction to consummate

Regular Paper

Special Section on Software Systems 2020

A preliminary version of the paper was published in the Proceedings of APLAS 2018.

This work was supported by the National Natural Science Foundation of China under Grant No. 61632005. *Corresponding Author

⁽¹⁾Linux v2.6.17.10. https://elixir.bootlin.com/linux/v2.6.17.10/source/arch, Sept. 2020.

⁽²⁾The SPARC Architecture Manual Version 8. https://gaisler.com/doc/sparcv8.pdf, Sept. 2020.

[©]Institute of Computing Technology, Chinese Academy of Sciences 2020

their write operation. The number of delay instructions (0 to 3) is implementation-dependent."

• *Register Windows*. SPARCv8 uses register windows and a window rotation mechanism to avoid saving contexts in the stack directly and achieves high performance in context management.

We use a simple example in Fig.1 to show these three features.

CALLER :	ChangeY :						
: 1 mov 1, %o ₀ 2 call ChangeY 3 save %sp, -64, %sp 4 mov %o ₀ , %l ₀ :	5 rd Y, %l ₀ 6 wr %i ₀ , 0, Y 7 nop 8 nop 9 nop 10 ret 11 restore %l ₀ , 0, %o ₀						
(a)	(b)						

Fig.1. Example for SPARC code $^{[4]}.$ (a) Function CALLER. (b) Function ChangeY.

SPARCv8 has 32 general registers, which are split into four logic groups as global $(\mathbf{r}_0-\mathbf{r}_7)$, out $(\mathbf{r}_8-\mathbf{r}_{15})$, local $(\mathbf{r}_{16}-\mathbf{r}_{23})$ and in $(\mathbf{r}_{24}-\mathbf{r}_{31})$ registers. Correspondingly, we give aliases " $\% \mathbf{g}_0 - \% \mathbf{g}_7$ ", " $\% \mathbf{o}_0 - \% \mathbf{o}_7$ ", " $\% \mathbf{l}_0 - \% \mathbf{l}_7$ " and " $\% \mathbf{i}_0 - \% \mathbf{i}_7$ " for these groups respectively.

CALLER in Fig.1(a) calls ChangeY, which updates the "special register" Y and returns its original value.

ChangeY in Fig.1(b) requires an input parameter as the new value for the special register Y. CALLER calls ChangeY at line 2, and pc and npc point to lines 2 and 3 respectively at this moment. The call instruction changes the value of pc to npc and lets npc point to the entry of ChangeY at line 5, which means the control-flow will not transfer to ChangeY in the next cycle, but in the cycle after the execution of the save instruction following the call. Similarly, when ChangeY returns (at line 10), the control is transferred back to the caller after executing the restore instruction at line 11. We call this feature "delayed control transfers".

SPARCv8 uses the save instruction (at line 3 in the example) to save the current context and the restore instruction (at line 10) to restore it. As explained above, among the 32 general registers, the out, local and in registers form the "current" register window. The local registers are for private use in the current context. The in and out registers are shared with adjacent register windows for parameters passing. The save instruction rotates the register window from the current one to the next. Then the local and in registers in the original window are no longer accessible, and the original out registers become the in registers in the current window.

In Fig.1, CALLER uses the save instruction (at line 3) to save its context (local and in registers) and rotate the register window (so that the out registers become the in registers). Thus, the $\%_{i_0}$ register assessed in ChangeY at line 6 is the same register as the $\%_{o_0}$ modified in CALLER at line 1. The restore instruction does the inverse. The arguments taken by the save and restore instructions are irrelevant here and can be ignored.

At line 6, the wr instruction tries to update the special register Y with the value of $\%i_0 \oplus 0$ (bitwise exclusive OR). Note that the $\%i_0$ register here is the same register as $\%o_0$ at line 1. However, the write is delayed for X cycles, where X is some predefined system parameter that ranges from 0 to 3. For portability, programmers usually do not rely on the exact value of X and assume it takes the maximum value of 3. Therefore three nop instructions are inserted. Reading of Y earlier than line 9 may give us the old value. This feature is called "delayed writes".

These features make the semantics of the SPARCv8 code context-dependent. For instance, a read of a special register (e.g., the register Y in the above example) needs to make sure that there are enough instructions executed since the most recent "delayed write". As another example, the instruction following the call instruction can be any instruction in general, but it is not supposed to update register r_{15} , which contains the return address saved by the call instruction. In addition, the delayed control transfer and the register windows also allow highly flexible calling conventions. Together, they make it a challenging task to have a Hoare-style program logic for local and modular reasoning of the SPARCv8 assembly code.

Working towards a certified OS kernel for aerospace crafts whose inline assembly is written in SPARCv8, we try to address these challenges and propose a practical program logic for realistically modelled SPARCv8 code. We have applied our logic to verify the main body of the task context switch routine in the kernel^[4].

However, the OS kernel is implemented as C language mainly and SPARCv8 as inline assembly. Just having a traditional Hoare-style program logic for SPARCv8, which can only make sure the safe execution of the SPARCv8 program if the initial state satisfies the precondition, is insufficient. Xu *et al.*^[1] proposed a program logic for verifying the correctness of OS kernel implemented in C language with inline assembly, but they used "abstract assembly primitives" to substitute the inline assembly in their verification work. As a supplement to their work, we extend our program logic so that it can ensure the contextual refinement relation, shown in (1), between the implementations and their corresponding abstract assembly primitives. Here, we use \mathbb{C} , A, and C_{as} to represent the C language program, the set of abstract assembly primitives and the implementations of abstract assembly primitives respectively. It means C_{as} refines A under "any context" \mathbb{C} :

$$\forall \mathbb{C}. \mathbb{C}[C_{\rm as}] \subseteq \mathbb{C}[A]. \tag{1}$$

Here we use " \subseteq " to represent the refinement relation.

However, if we use the C program as a client code to call inline assembly code, we need to define the semantics of C-assembly interaction.

Since the goal of this paper is to verify the correctness of the assembly code with respect to the abstract assembly primitives, we want to avoid the C-assembly interaction (and leave it as future work). We decompose the OS verification tasks into two steps, as shown in Fig.2, and focus on S2 only in this paper.

$$\begin{bmatrix} \mathbb{C}[A] \end{bmatrix}^{\mathsf{C}}$$
(S1) $\cup \mathsf{I} \iff \overline{C = Comp(\mathbb{C})}$

$$\begin{bmatrix} C[\Omega] \end{bmatrix}^{\mathsf{P}-\mathsf{SPARCv8}}$$
(S2) $\cup \mathsf{I} \iff \overline{\mathsf{P} \cdot C_{\mathrm{as}} : \Omega}$

$$\begin{bmatrix} C[C_{\mathrm{as}}] \end{bmatrix}^{\mathsf{SPARCv8}}$$

Fig.2. Idea to establish contextual refinement.

The source program of OS kernel $\mathbb{C}[A]$, which executes under the C language semantics (shown as $[-]^{C}$), is implemented as C language with a set of abstract assembly primitives A. The compiler (Comp) translates the C program \mathbb{C} to the SPARCv8 code C. As S1 shows, we assume the compilation ensures the refinement between $\mathbb{C}[A]$ and $C[\Omega]$ that executes under Pseudo-SPARCv8 semantics shown as []_]^{P-SPARCv8}. Here, Ω represents the set of abstract assembly primitives in the middle layer. We use distinguished notations to represent the set of abstract assembly primitives in the source and intermediate level, since they execute on different program states and have different semantics. The Pseudo-SPARCv8 language $C[\Omega]$, which uses SPARCv8 as client code, is able to call abstract assembly primitives in Ω . In S2, we verify using our refinement logic that the whole SPARCv8 program $C[C_{as}]$ executing under the realistically modelled SPARCv8 semantics, represented as $[\![_]\!]^{\mathsf{SPARCv8}},$ refines the program $C[\Omega]$ executing under the Pseudo-SPARCv8 semantics.

Finally, we can get $\llbracket C[C_{as}] \rrbracket^{\mathsf{SPARCv8}} \subseteq \llbracket \mathbb{C}[A] \rrbracket^{\mathsf{C}}$. In this work, we focus on S2, and leave S1 as future work.

Our work is based on earlier work on assembly code verification but makes the following contributions.

• We propose a new program logic which supports relational reasoning for refinement verification. It ensures that a verified SPARCv8 function contextually refines its corresponding abstract assembly primitive. Our logic supports all the above features of SPARCv8. We redefine basic blocks to include the instruction following the jump or return as the tail of a block, which models the delayed control transfer. To reason about delayed writes, we introduce a modal assertion $\triangleright_t \mathbf{sr} \mapsto w$, saying that the special register \mathbf{sr} will hold the value w in up to t cycles. We also give logic rules for \mathbf{save} and $\mathbf{restore}$ instructions that do register window rotation.

• In order to support refinement verification, we define a Pseudo-SPARCv8 language as the language to implement the high-level specification. It also hides the details of sophisticated register window mechanism in SPARCv8 by abstraction, and makes its language model simpler than realistic SPARCv8. Therefore, it can provide some convenience to write the abstract assembly primitive and reason in the Pseudo-SPARCv8 level.

• Following SCAP^[5], our logic supports modular reasoning of function calls in a direct style. However, we use the traditional pre- and post-conditions as function specifications, instead of the assertion g used in SCAP. This allows us to reuse existing techniques (e.g., Coq tactics) to simplify the verification process. The logic rules for function call and return are general and independent of any specific calling convention.

• We give direct-style semantic interpretation for the logic judgments, based on which we establish the soundness. This is different from previous work, which either does syntactic-based soundness proof (e.g., SCAP^[5]) or treats return code pointers as firstclass code pointers and gives CPS-style (continuationpassing style) semantics. Those approaches for soundness make it difficult to verify the interaction between the inline assembly and the C code in the kernel, the latter being verified following a direct-style program logic.

• Context switch of concurrent tasks is an important component in OS kernels. It is usually implemented as inline assembly because of the need to access registers and the stack. We apply our logic and verify that there is a contextual refinement between a context switch routine implemented in SPARCv8 and an abstract switch

primitive.

The program logic and its soundness proof have been mechanized in Coq. Coq proofs and a comparison technical report are available⁽³⁾.

This paper extends our conference paper in APLAS 2018^[4]. The program logic there can only verify the partial correctness of SPARCv8 code and we make the following expansions.

• We propose a new program logic to do relational reasoning for refinement verification (Subsection 4.4 for details).

• In order to support refinement verification, this paper presents a new Pseudo-SPARCv8 language as the high-level specification language (Subsection 4.1 for details).

• We also use the new logic to verify the implementation of a context switch routine, by showing that the implementation contextually refines an abstract switch primitive (Section 5 for details).

In the remainder of the paper, we present the program model and operational semantics of SPARCv8 in Section 2. For the clear presentation, we use a simplified version of our program logic that does not support refinement verification to demonstrate how our logic supports the three main features of SPARCv8 in Section 3, which is the main point of our work and irrelevant to refinement verification. We present the Pseudo-SPARCv8 program and the relational program logic for refinement reasoning in Section 4. We show the verification of a context switch routine in SPARCv8 in Section 5. Finally, we discuss more on related work in Section 6 and conclude the paper in Section 7.

2 SPARCv8 Assembly Language

We introduce the key SPARCv8 instructions, the model of machine states, and the operational semantics

in this section.

2.1 Language Syntax and States

The machine model and the syntax of SPARCv8 assembly language are defined in Fig.3. Here, we follow the block-based memory^[6] introduced in CompCert to define the memory in our work. The memory address lis defined as a pair of its block ID and the offset. Block IDs (Block) are integers in mathematics presented as \mathbb{Z} , and offsets (Word) are 32-bit integers (called "words"). The value (Val) is either word w or address l. The whole program configuration (Prog) P includes the code heap (CodeHeap) C, the machine state (State) S, and the program counters pc and npc. The code heap C is a partial function from labels f to commands c. Labels are also 32-bit integers, which can be viewed as locations where the commands are saved in the code heap. The operand expression (OpExp) o, which is either a general register \mathbf{r} or a word w, and the address expression (AddrExp) **a**, which is either an operand expression or a sum of the values of register r and an operation expression, are auxiliary definitions used as parameters of commands. Commands (Comm) in SPARCv8 are classified into two categories: 1) the simple instructions (SimpIns) i, which do sequential operation, e.g., arithmetic operation "add", or memory operations "ld" (load) and "st" (store), or register window operations "save" and "restore", or special register operations "rd" (read) and "wr" (write), or "nop", whose execution changes no program state (the program counters pc and npc); 2) the control-transfer instructions, e.g., call and retl for function call and return, or jmp and be for unconditional and conditional branch.

The machine state (State) S consists of three parts: the memory (Mem) M, the register state (Rstate) Qwhich is a pair of register file (RegFile) R and frame

(Word) Int32 (Block) $b \in \mathbb{Z}$ (Addr) $l \in \text{Block} \times \text{Word}$ (Val) $v ::= w \mid l$ $w, f \in$ (Prog) P::= (C, S, pc, npc)(CodeHeap) C $\in \quad \mathrm{Word} \rightharpoonup \mathrm{Comm}$ (State) S(M, Q, D)(RState) ::= (R, F)..= 0 (ProgCount) $pc, npc \in Word$ (Mem) M \in $\operatorname{Addr} \rightharpoonup \operatorname{Val}$ (OpExp) $::= \mathbf{r} \mid w$ (AddrExp) ::= o | r + o 0 a $::= \ \texttt{i} \mid \texttt{call f} \mid \texttt{jmp a} \mid \texttt{retl} \mid \texttt{be f}$ (Comm) c(SimpIns) i ::= ld a r_d | st r_s a | nop | add r_s o r_d | save r_s o r_d | restore r_s o r_d $| \mathsf{rd} \mathsf{srr}_d | \mathsf{wrr}_s \mathsf{o} \mathsf{sr} |$. (InstrSeq) \mathbb{I} ::= i; $\mathbb{I} \mid jmp a; i \mid call f; i; \mathbb{I} \mid retl i \mid be f; i; \mathbb{I}$

Fig.3. Machine states and language for SPARCv8 $Code^{[4]}$.

⁽³⁾https://github.com/jpzha/VeriSparc, Sept. 2020.

1386

list (FrameList) F, and the delay buffer (DelayBuff) D. As defined in Fig.4, R is a partial mapping from register names (RegName) to values. Registers include the general registers \mathbf{r} , the processor state registers (PsrReg) \mathbf{psr} and the special registers (SpeReg) \mathbf{sr} . \mathbf{psr} contains the integer condition code fields $\mathbf{n}, \mathbf{z}, \mathbf{v}$ and \mathbf{c} , which can be modified by the arithmetic and logical instructions and used for conditional control-transfer, and \mathbf{cwp} recording the ID of the current register window. We explain the frame list F and the delay buffer D below.

Register Windows and Frame List. SPARCv8 provides 32 general registers that are split into four groups as global (r_0-r_7) , out (r_8-r_{15}) , $local(r_{16}-r_{23})$ and in

 $(\mathbf{r}_{24}-\mathbf{r}_{31})$ registers. The latter three groups (out, local and in) form the current register window.

At the entry and exit of functions and traps, one may need to save and restore some of the general registers as execution contexts. Instead of saving them into stacks in memory, SPARCv8 uses multiple register windows to form a circular stack, and does window rotation for efficient context save and restore. As shown in Fig.5 (the figure taken from "The SPARC Architecture Manual Version 8"⁽⁴⁾), there are N register windows (N = 8 here) consisting of $2 \times N$ groups of registers (each group containing eight registers). The cwp register (part of psr) records the ID number of the current window (cwp = 0 in this example).

(RegFile)	R	\in	$\operatorname{RegName} ightarrow \operatorname{Val}$	(RegName)	rn	::=	$r_0 \mid \ldots \mid r_{31} \mid psr \mid sr$
(PsrReg)	psr	::=	$\texttt{n} \mid \texttt{z} \mid \texttt{v} \mid \texttt{c} \mid \texttt{cwp}$	(SpeReg)	sr	::=	$\texttt{wim} \mid \texttt{Y} \mid \texttt{asr}_0 \mid \ldots \mid \texttt{asr}_3$
(FrameList)	F	::=	nil fm :: F	(Frame)	${\rm fm}$::=	$[v_0,\ldots,v_7]$
(DelayBuff)	D	::=	$\operatorname{nil} \mid (t, \mathtt{sr}, w) :: D$	$({\rm DelayCycle})$	t	\in	$\{0, 1, \dots, X\}$

Fig.4. Register file, frame list and delaybuffer^[4].



⁽⁴⁾The SPARC Architecture Manual Version 8. https://gaisler.com/doc/sparcv8.pdf, Sept. 2020.

The in and out registers of each window are shared with its adjacent windows for parameter passing. For example, the in registers of w_0 are the out registers of w_1 , and the out registers of w_0 are the in registers of w_7 . This explains why we need only $2 \times N$ groups of registers for N windows, while each window consists of three groups (out, local and in).

To save the context, the **save** instruction rotates the window by decrementing the **cwp** pointer (modulo N) and w_7 becomes the current window. The **out** registers of w_0 become the in registers of w_7 . The in and **local** registers of w_0 become inaccessible. This is like pushing them onto the circular stack. The **restore** instruction does the inverse, which is like a stack pop.

The wim register is used as a bit vector to record the end of the stack. Each bit in wim corresponds to a register window. The bit corresponding to the last available window is set to 1, which means "invalid". All the other bits are 0 (i.e., "valid"). When executing save (and restore), we need to ensure the next window is valid, in order to avoid the overflow of register window because of the limitation of the number of windows. We use the assertion win_valid(w_{id} , R) defined in Fig.6 to say the window pointed to by w_{id} is valid, given the value of wim in R.

We use the frame list F to model the circular stack consisting of register windows. As defined in Fig.4, a frame is an array of eight words, modeling a group of eight registers. F consists of a sequence of frames corresponding to all the register windows except the out, local and in registers in the current window. Then save saves the local and in registers onto the head of F and loads the two groups of registers at the "tail" of F to the local and out registers (and the original out registers become the in group). The restore instruction does the inverse. The operations are defined formally in Fig.6. Here, we use "::" for adding an element at the head of a list, and use " \cdot " for appending an element at the tail of a list.

Delay Buffer. The delay buffer D is a sequence of delayed writes. Because the wr instruction does not update the target register immediately, we put the write operation onto the delay buffer. A delayed write is recorded as a triple consisting of the remaining cycles t to be delayed, the target special register sr and the value w to be written. Note that the value of a special register is restricted to only words, since the special registers are used to record the state of processors, and it is impossible to store memory addresses in them.

Instruction Sequences. We use an instruction sequence \mathbb{I} to model a basic block, i.e., a sequence of commands ending with a control transfer. As defined in Fig.3, we require that a delayed control-transfer instruction must be followed by a simple instruction i, because the actual control-transfer occurs after the execution of i. The end of each instruction sequence can only be jmp or retl followed by a simple instruction i. Note that we do not view the call instruction as

$$\begin{aligned} & \text{out} \ ::= [\mathbf{r}_8, \dots, \mathbf{r}_{15}] \quad \text{local} \ ::= [\mathbf{r}_{16}, \dots, \mathbf{r}_{23}] \quad \text{in} \ ::= [\mathbf{r}_{24}, \dots, \mathbf{r}_{31}] \\ & R([\mathbf{r}_i, \dots, \mathbf{r}_{i+k}]) \ ::= [R(\mathbf{r}_i), \dots, R(\mathbf{r}_{i+k})] \\ & R\{[\mathbf{r}_i, \dots, \mathbf{r}_{i+7}] \rightsquigarrow \text{fm}\} \ ::= R\{\mathbf{r}_i \rightsquigarrow v_0\} \dots \{\mathbf{r}_{i+7} \rightsquigarrow v_7\} \quad \text{where} \quad \text{fm} = [v_0, \dots, v_7] \\ & \text{win_valid}(w_{id}, R) \ ::= 2^{w_{id}} \& R(\texttt{wim}) = 0 \\ & \text{where} \& \text{is the bitwise AND operation.} \\ & \text{next_cwp}(w_{id}) \quad ::= (w_{id} + N - 1)\% N \qquad \text{prev_cwp}(w_{id}) \ ::= (w_{id} + 1)\% N \\ & \text{save}(R, F) \quad ::= \begin{cases} (R', F'), \quad \text{if} \ w'_{id} = \texttt{next_cwp}(R(\texttt{cwp})), \texttt{win_valid}(w'_{id}, R), \\ & F = F'' \cdot \texttt{fm}_1 \cdot \texttt{fm}_2, \ F' = R(\texttt{local}) :: R(\texttt{in}) :: F'', \\ & R'' = R\{\texttt{in} \rightsquigarrow R(\texttt{out}), \texttt{local} \rightsquigarrow \texttt{fm}_2, \texttt{out} \rightsquigarrow \texttt{fm}_1\}, \\ & R' = R''\{\texttt{cwp} \rightsquigarrow w'_{id}\}, \\ \bot, \qquad \text{if} \neg \texttt{win_valid}(\texttt{next_cwp}(R(\texttt{cwp})), \texttt{win_valid}(w'_{id}, R), \\ & F = \texttt{fm}_1 :: \texttt{fm}_2 :: F'', \ F' = F'' \cdot R(\texttt{out}) \cdot R(\texttt{local}), \\ & R'' = R\{\texttt{in} \rightsquigarrow \texttt{fm}_2, \texttt{local} \rightsquigarrow \texttt{fm}_1, \texttt{out} \rightsquigarrow R(\texttt{in})\}, \\ & R'' = R\{\texttt{in} \rightsquigarrow \texttt{fm}_2, \texttt{local} \rightsquigarrow \texttt{fm}_1, \texttt{out} \rightsquigarrow R(\texttt{in})\}, \\ & R'' = R\{\texttt{in} \rightsquigarrow \texttt{fm}_2, \texttt{local} \rightsquigarrow \texttt{fm}_1, \texttt{out} \rightsquigarrow R(\texttt{in})\}, \\ & R'' = R\{\texttt{in} \rightsquigarrow \texttt{fm}_2, \texttt{local} \rightsquigarrow \texttt{fm}_1, \texttt{out} \rightsquigarrow R(\texttt{in})\}, \\ & R'' = R'' \{\texttt{cwp} \rightsquigarrow w'_{id}\}, \\ & \bot, \qquad \texttt{if} \neg \texttt{win_valid}(\texttt{prev_cwp}(R(\texttt{cwp})), R). \end{aligned}$$

Fig.6. Auxiliary definitions for instruction save and restore^[4].

the end of a basic block, since the callee is expected to return, following our direct-style semantics for function calls. We define $C[\mathbf{f}]$ to extract an instruction sequence starting from \mathbf{f} in C below.

$$C[\mathbf{f}] = \begin{cases} \mathbf{i}; \mathbb{I}, & \text{if } C(\mathbf{f}) = \mathbf{i} \text{ and } C[\mathbf{f}+4] = \mathbb{I}, \\ c; \mathbf{i}, & \text{if } c = C(\mathbf{f}) \text{ and } c = \mathtt{jmp} \mathbf{a} \text{ or retl} \\ & \text{and } C(\mathbf{f}+4) = \mathbf{i}, \\ c; \mathbf{i}; \mathbb{I}, & \text{if } c = C(\mathbf{f}) \text{ and } c = \mathtt{call} \mathbf{f} \text{ or } \mathtt{be} \mathbf{f} \\ & \text{and } C(\mathbf{f}+4) = \mathbf{i} \text{ and } C[\mathbf{f}+8] = \mathbb{I} \\ & \text{undefined, otherwise.} \end{cases}$$

2.2 Operational Semantics

The operational semantics is taken from Wang *et al.*^[7], but we use a block-based memory model and omit features like interrupts and traps. We show the selected rules in Fig.7. The program transition relation $(C, S'_{\text{max}}, \mathsf{npc}) ::\Rightarrow (C, S'_{\text{max}}, \mathsf{c}', \mathsf{npc}')$ is defined in Fig.7(a). Before the execution of the instruction pointed by pc, the delayed writes in D with 0 delay cycles are executed first. The execution of the delayed writes is defined in the form of $(R, D) \rightrightarrows (R', D')$ below.

$$\begin{split} & (R,D) \rightrightarrows (R',D') \\ & C \vdash ((M,(R',F),D'), \texttt{pc},\texttt{npc}) \longrightarrow ((M',(R'',F'),D''),\texttt{pc'},\texttt{npc'}) \\ & (C,(M,(R,F),D),\texttt{pc},\texttt{npc}) ::\Longrightarrow (C,(M',(R'',F'),D''),\texttt{pc'},\texttt{npc'}) \\ & (a) \end{split}$$

$$\begin{array}{l} C(\mathrm{pc}) = \mathrm{i} & (M, (R, F), D) \bullet^{-1} \to (M', (R', F'), D') \\ \hline C \vdash ((M, (R, F), D), \mathrm{pc}, \mathrm{npc}) \circ \longrightarrow ((M', (R', F'), D'), \mathrm{npc}, \mathrm{npc} + 4) \\ \hline & \frac{C(\mathrm{pc}) = \mathrm{jmp} \ \mathrm{a} \quad \llbracket \mathrm{a} \rrbracket_R = \mathrm{f}}{C \vdash ((M, (R, F), D), \mathrm{pc}, \mathrm{npc}) \circ \longrightarrow ((M, (R, F), D), \mathrm{npc}, \mathrm{f})} \\ \hline & \frac{C(\mathrm{pc}) = \mathrm{call} \ \mathrm{f} \quad \mathrm{r}_{15} \in \mathrm{dom}(R)}{C \vdash ((M, (R, F), D), \mathrm{pc}, \mathrm{npc}) \circ \longrightarrow ((M, (R_{\{\mathrm{r}_{15} \leftrightarrow \mathrm{pc}\}}, F), D), \mathrm{npc}, \mathrm{f})} \\ \hline & \frac{C(\mathrm{pc}) = \mathrm{call} \ \mathrm{f} \quad \mathrm{r}_{15} \in \mathrm{dom}(R)}{C \vdash ((M, (R, F), D), \mathrm{pc}, \mathrm{npc}) \circ \longrightarrow ((M, (R_{\{\mathrm{r}_{15}\}} \to \mathrm{pc}\}, F), D), \mathrm{npc}, \mathrm{f})} \\ \hline & \frac{C(\mathrm{pc}) = \mathrm{retl} \quad R(\mathrm{r}_{15}) = \mathrm{f}}{C \vdash ((M, (R, F), D), \mathrm{pc}, \mathrm{npc}) \circ \longrightarrow ((M, (R, F), D), \mathrm{npc}, \mathrm{f} + 8)} \\ \hline & \frac{C(\mathrm{pc}) = \mathrm{be} \ \mathrm{f} \quad R(\mathrm{z}) \neq 0}{C \vdash ((M, (R, F), D), \mathrm{pc}, \mathrm{npc}, \mathrm{f})} \quad \frac{C(\mathrm{pc}) = \mathrm{be} \ \mathrm{f} \quad R(\mathrm{z}) = 0}{C \vdash ((M, (R, F), D), \mathrm{pc}, \mathrm{npc}, \mathrm{o} \longrightarrow ((M, (R, F), D), \mathrm{pc}, \mathrm{npc} + 4))} \\ (\mathrm{b}) \end{array}$$

Fig.7. Selected operational semantics rules [4]. (a) Program transition. (b) Control transfer instruction transitions. (c) save, restore and wr instruction transitions. (d) Simple instruction transitions. (e) Expression semantics.

Jun-Peng Zha et al.: Modular Verification of SPARCv8 Code

$$\overline{(R, \operatorname{nil}) \rightrightarrows (R, \operatorname{nil})}$$
$$(R, D) \rightrightarrows (R', D')$$
$$\overline{(R, (t+1, \operatorname{sr}, w) :: D) \rightrightarrows (R', (t, \operatorname{sr}, w) :: D')}$$
$$\frac{(R, D) \rightrightarrows (R', D') \quad \operatorname{sr} \in \operatorname{dom}(R)}{(R, (0, \operatorname{sr}, w) :: D) \rightrightarrows (R' \{ \operatorname{sr} \rightsquigarrow w \}, D')}$$
$$\frac{(R, D) \rightrightarrows (R', D') \quad \operatorname{sr} \notin \operatorname{dom}(R)}{(R, (0, \operatorname{sr}, w) :: D) \rightrightarrows (R', D')}$$

Note that the write of sr has no effect if sr is not in the domain of R. Since R is defined as a partial map, we can prove the following lemma.

Lemma 1^[4]. $(R, D) \Rightarrow (R', D')$ and $R = R_1 \uplus R_2$, if and only if there exist R'_1 and R'_2 , such that $(R_1, D) \Rightarrow (R'_1, D'), (R_2, D) \Rightarrow (R'_2, D'), and R' = R'_1 \uplus R'_2.$

Here the disjoint union $R_1 \uplus R_2$ represents the union of R_1 and R_2 if they have disjoint domains, and undefined otherwise. This lemma is important to give sound semantics to delay buffer related assertions, as discussed in Section 3.

The transition steps for individual instructions are classified into three categories: the control transfer steps (_ \vdash _ $\circ \longrightarrow$ _), the steps for save, restore and wr instructions (_ $\bullet \implies$ _), and the steps for other simple instructions (_ \implies _). The corresponding step transition relations are defined inductively in Fig.7(b)– Fig.7(d) respectively.

Note that, after the control-transfer instructions, pc is set to npc and npc contains the target code pointer. This explains the one cycle delay for the control transfer. The call instruction saves pc into register \mathbf{r}_{15} , while retl uses $\mathbf{r}_{15} + 8$ as the return address (which is the address for the second instruction following the call). The conditional branch be f jumps to f (after one-cycle delay) if the value in the register z is not 0. Evaluation of expressions a and o are defined in Fig.7(e). Here, we define the sum of two values v_1 and v_2 below. The result of v_1+v_2 is legal, if both of v_1 and v_2 are words (Int32), or v_1 is an address and v_2 is an offset. The offset is a word, which acts as an immediate value in the calculation of address.

$$v_1 + v_2 ::= \begin{cases} w_1 + w_2, & \text{if } v_1 = w_1, \text{ and } v_2 = w_2, \\ (b, w_1 + w_2), & \text{if } v_1 = (b, w_1), \text{ and } v_2 = w_2, \\ \text{undefined}, & \text{otherwise.} \end{cases}$$

wr wants to save the bitwise exclusive OR of the operands into the special register sr, but it puts the

write into the delay buffer D instead of updating R immediately. The operation set_delay(sr, w, D) is defined below:

$$\mathsf{set_delay}(\mathtt{sr}, w, D) ::= (X, \mathtt{sr}, w) :: D,$$

where X ($0 \le X \le 3$) is a predefined system parameter for the delay cycle.

The save and restore instructions rotate the register windows and update the register file. Their operations over F and R are defined in Fig.6.

3 Program Logic

In this section, we use a simplified version of our program logic that does not support refinement verification to present how our logic handles the features of SPARCv8. The relational program logic for refinement reasoning will be introduced in Section 4.

3.1 Assertions

We define the syntax of assertions (Asrt) in Fig.8, and their semantics in Fig.9. We extend separation logic assertions with specifications of delay buffers and register windows. Registers are like variables in separation logic, but treated as resources. The assertion emp says that the memory and the register file are both empty. $l \mapsto v$ specifies a singleton memory cell with value v stored in address l. $\mathbf{rn} \mapsto v$ says that \mathbf{rn} is the only register in the register file and it contains the value v. Also \mathbf{rn} is not in the delay buffer. Separating conjunction p * q has the standard semantics as in separation logic^[8].

(Asrt)
$$p, q ::= emp \mid l \mapsto v \mid rn \mapsto v \mid \rhd_t sr \mapsto w \mid p \downarrow$$

 $\mid cwp \mapsto (\!(w_{id}, F)\!) \mid p \land q \mid p \lor q \mid p \ast q$
 $\mid a =_a v \mid o = v \mid \forall x. p \mid \exists x. q \mid \dots$



The assertion $\triangleright_t \mathbf{sr} \mapsto w$ describes a delayed write in the delay buffer D. It describes the uncertainty of $\mathbf{sr's}$ value in R, which is unknown for now but will become w in up to t+1 cycles. We use $_ \Rightarrow^k _$ to represent the k-step execution of the delayed writes in D. It also requires that there be at most one delayed write for a specific special register \mathbf{sr} in D (i.e., $\mathsf{noDup}(\mathbf{sr}, D)$). This prevents more than one delayed write to the register within four instruction cycles, which practically have no restrictions on programming. By the semantics

J. Comput. Sci. & Technol., Nov. 2020, Vol.35, No.6

$$\begin{split} S &\models \mathsf{emp} & ::= S.M = \emptyset \land S.Q.R = \emptyset \\ S &\models l \mapsto v & ::= S.M = \{l \rightsquigarrow v\} \land S.Q.R = \emptyset \\ S &\models \mathsf{rn} \mapsto v & ::= S.Q.R = \{\mathsf{rn} \rightsquigarrow v\} \land \mathsf{rn} \notin \mathsf{dom}(S.D) \land S.M = \emptyset \\ S &\models \mathsf{rs} \mapsto v & ::= \exists k, R', D'. 0 \leqslant k \leqslant t + 1 \land (R, D) \Rightarrow^k (R', D') \land \mathsf{noDup}(D, \mathsf{sr}) \\ & ((M, (R', F), D') \models \mathsf{sr} \mapsto w) \land \mathsf{noDup}(D, \mathsf{sr}) \\ & \mathsf{where} S = (M, (R, F), D) \\ S &\models \mathsf{ra} \Rightarrow v & ::= \exists R', D'. ((M, (R', F), D') \models p) \land (R', D') \Rightarrow (R, D) \\ & \mathsf{where} S = (M, (R, F), D) \\ S &\models \mathsf{cwp} \mapsto (w_{id}, F) ::= (S \models \mathsf{cwp} \mapsto w_{id}) \land \exists F'. F \cdot F' = S.Q.F \\ S &\models \mathsf{a} = a v & ::= \llbracket \mathsf{a} \rrbracket_{S,Q.R} = v \\ S &\models \mathsf{o} = v & ::= \llbracket \mathsf{a} \rrbracket_{S,Q.R} = v \\ S &\models \mathsf{p}_1 \ast p_2 & ::= \exists S_1, S_2. S_1 \models p_1 \land S_2 \models p_2 \land S = S_1 \uplus S_2 \\ M_1 \bot M_2 & ::= (\mathsf{dom}(M_1) \cap \mathsf{dom}(M_2)) = \emptyset & R_1 \bot R_2 & ::= (\mathsf{dom}(R_1) \cap \mathsf{dom}(R_2)) = \emptyset \\ S_1 &\Downarrow S_2 & ::= \begin{cases} (M_1 \cup M_2, (R_1 \cup R_2, F), D), & \text{if } M_1 \bot M_2 \land R_1 \bot R_2 \land \\ S_1 = (M_1, (R_1, F), D) \land S_2 = (M_2, (R_2, F), D), \\ \mathsf{undefined}, & \text{otherwise.} \end{cases} \\ \mathsf{dom}(D) & ::= \begin{cases} \{\mathsf{sr}\} \cup \mathsf{dom}(D'), & \text{if } D = (t, \mathsf{sr}, w) :: D', \\ \emptyset, & \text{if } D = \mathsf{nil}. \end{cases} \\ \mathsf{noDup}(D, \mathsf{sr}) & ::= \begin{cases} \mathsf{sr} \notin \mathsf{dom}(D'), & \text{if } D = (t, \mathsf{sr}, w) :: D', \\ \mathsf{True}, & \text{if } D = \mathsf{nil}. \end{cases} \end{cases}$$

Fig.9. Semantics of assertions $^{[4]}$.

we have $(p \Rightarrow q \text{ means for any } S$, if $S \models p$ holds, then $S \models q$ holds):

$$s\mathbf{r} \mapsto w \Rightarrow \rhd_t s\mathbf{r} \mapsto w,$$
$$\rhd_t s\mathbf{r} \mapsto w \Rightarrow \rhd_{t+k} s\mathbf{r} \mapsto w.$$

The assertion $p \downarrow$ allows us to reduce the uncertainty by executing one step of the delayed writes. It specifies states reachable after executing one step of delayed writes from those states satisfying p. Therefore we know:

$$(\rhd_0 \mathbf{sr} \mapsto w) \downarrow \Rightarrow \mathbf{sr} \mapsto w, \\ (\rhd_{t+1} \mathbf{sr} \mapsto w) \downarrow \Rightarrow \rhd_t \mathbf{sr} \mapsto w$$

Also it is easy to see that if p syntactically does not contain sub-terms in the form of $\triangleright_t \mathbf{sr} \mapsto w$, then $(p\downarrow) \Leftrightarrow p$.

The following lemma shows $(_)\downarrow$ is distributive over separating conjunction.

Lemma $2^{[4]}$. $(p * q) \downarrow \Leftrightarrow (p \downarrow) * (q \downarrow)$.

The lemma can be proved following Lemma 1. We present the proof of Lemma 2 below.

Proof. " \Rightarrow ": if $(M, (R, F), D) \models (p * q) \downarrow$, then $(M, (R, F), D) \models (p) \downarrow *(q) \downarrow$. We first get there exist R', D', M_1, M_2, R'_1 and R'_2 such that $(R_1 \uplus R_2)$ represents the union of R_1 and R_2 if they have disjoint domains, and $M_1 \uplus M_2$ has the same meaning): (A1) $(R', D') \Rightarrow (R, D)$; (A2) $R' = R'_1 \uplus R'_2$; (A3) $M = M_1 \uplus M_2$; (A4) $(M_1, (R'_1, F), D') \models p$; (A5) $(M_2, (R'_2, F), D') \models q$. By applying Lemma 1 on (A.1), we get that there exist R_1 and R_2 , where $R = R_1 \uplus R_2$, such that: $(R'_1, D') \Rightarrow (R_1, D), (R'_2, D') \Rightarrow (R_2, D)$. Thus, we get $(M_1, (R_1, F), D) \models p\downarrow$; $(M_2, (R_2, F), D) \models q\downarrow$. Finally, we prove that $(M, (R, F), D) \models (p\downarrow) * (q\downarrow)$ holds.

" \Leftarrow ": if $(M, (R, F), D) \models (p \downarrow) * (q \downarrow)$, then $(M, (R, F), D) \models (p * q) \downarrow$. We first get there exist $M_1, M_2, R'_1, R'_2, R_1, R_2, D'_1$ and D'_2 , where $R = R_1 \uplus R_2$, such that: (B1) $(R'_1, D'_1) \rightrightarrows (R_1, D)$, $(R'_2, D'_2) \rightrightarrows (R_2, D)$; (B2) $(M_1, (R'_1, F), D'_1) \models p$; (B3) $(M_2, (R'_2, F), D'_2) \models q$.

By the definition of the step of the delayed writes, we can prove that $D'_1 = D'_2$. Let $D' = D'_1 = D'_2$. By applying Lemma 1 on (B1), we get there exists R', where $R' = R'_1 \uplus R'_2$, such that $(R', D') \rightrightarrows (R, D)$. Thus, we prove that $(M, (R, F), D) \models (p * q) \downarrow$ holds. \Box

We use $\operatorname{cwp} \mapsto (|w_{id}, F|)$ to describe the pointer cwp of the current register window and the frame list as a circular stack. Note that F is just a prefix of the frame list, since usually we do not need to know contents of the full list. Here we use $F \cdot F'$ to represent the concatenation of lists F and F'. Therefore we have $\operatorname{cwp} \mapsto$ $(|w_{id}, F \cdot F'|) \Rightarrow \operatorname{cwp} \mapsto (|w_{id}, F|).$

The assertions $\mathbf{a} =_a v$ and $\mathbf{o} = v$ describe the value of \mathbf{a} and \mathbf{o} respectively. They are intuitionistic assertions. Since \mathbf{a} is used as an address, we also require it to be properly aligned on a 4-byte boundary. We define word_align to represent this restriction below:

word_align $(v) ::= \exists w. (v = w \lor v = (\underline{}, w)) \land w$ %4 = 0.

The result of the address expression **a** may be a word, if it is a pointer in code heap, or a memory address if

Jun-Peng Zha et al.: Modular Verification of SPARCv8 Code

it is a location of memory.

3.2 Inference Rules

The code specification θ and the code heap specification Ψ are defined below.

(valList)
$$\iota \in \text{list value (pAsrt) fp}, \text{fq} \in \text{valList} \to \text{Asrt}$$

(CdSpec) $\theta ::= (\text{fp}, \text{fq})$ (CdHpSpec) $\Psi ::= \{\mathbf{f} \rightsquigarrow \theta\}^*$

The code heap specification Ψ maps the code labels for basic blocks to their specifications θ , which is a pair of pre- and post-conditions. Instead of using normal assertions, the pre- and post-conditions are assertions parameterized over a list of values ι . They play the role of auxiliary variables—feeding the pre- and the post-conditions with the same ι allows us to establish the relationship of states specified in the pre- and post-conditions.

Although we assign a specification θ to each basic block, the post-condition does not specify the states reached at the end of the block. Instead, it specifies the condition that needs to be specified in the future when the "current function" returns. This follows the idea developed in SCAP^[5], but we use the standard unary state assertion instead of the binary state assertions used in SCAP so that existing proof techniques (such as Coq tactics) for standard Hoare-triples can be applied to simplify the verification process.

We give a simple example in Fig.10 to show a specification for a function, which simply sums the values of the registers $\%i_0$, $\%i_1$ and $\%i_2$ and writes the result into register $\%l_7$. The specification (fp, fq) says that, when provided with the same lv as argument, the function preserves the value of $\%i_0$, $\%i_1$ and $\%i_2$, $\%l_7$ in the beginning contains any value and at the end contains the sum of $\%i_0$, $\%i_1$ and $\%i_2$, and the function also preserves the value of \mathbf{r}_{15} , which it uses as the return address. To verify the function, we need to prove that it satisfies (fp lv, fq lv) for all lv. Here, lv[1] and lv[2] cannot be a memory address, because a value plus a memory address is illegal. lv[3] should also be a word, because it is a return code pointer whose type is word.

Fig.11 shows selected inference rules in our logic. Our logic divides the proof work into three layers. We define the well-formed code heap in the form of $(\vdash C : \Psi)$ to verify the code heap C, the well-formed instruction sequence in the form of $(\Psi \vdash \{(p,q)\} \mathbf{f} : \mathbb{I})$ to verify the instruction sequence \mathbb{I} starting from \mathbf{f} in code heap, and well-formed instruction in the form of $(\vdash \{p\} \mathbf{i} \{q\})$ to verify the single simple instruction \mathbf{i} .

$$\begin{array}{l} - \ \left\{ (\mathrm{fp},\mathrm{fq}) \right\} \\ & \mathrm{add} \quad \%_{10}, \ \%_{11}, \ \%_{17} \\ & \mathrm{add} \quad \%_{17}, \ \%_{12}, \ \%_{17} \\ & \mathrm{retl} \\ & \mathrm{nop} \end{array}$$
 fp ::= $\lambda \, lv. \left((\%_{10} \mapsto lv[0]) * (\%_{11} \mapsto lv[1]) * (\%_{12} \mapsto lv[2]) \\ & *\%_{17} \mapsto _ * (\mathbf{r}_{15} \mapsto lv[3])) \\ & \wedge (lv[1], lv[2], lv[3] \in \mathrm{Word})$ fq ::= $\lambda \, lv. (\%_{10} \mapsto lv[0]) * (\%_{11} \mapsto lv[1]) * (\%_{12} \mapsto lv[2]) \\ & *(\%_{17} \mapsto lv[0] + lv[1] + lv[2]) * (\mathbf{r}_{15} \mapsto lv[3])$



The top rule CDHP verifies the code heap C. It requires that every basic block specified in Ψ can be verified with respect to the specification, with any argument ι used to instantiate the pre- and post-conditions.

The SEQ rule is applied when meeting an instruction sequence starting with a simple instruction \mathbf{i} . The instruction \mathbf{i} is verified by the corresponding wellformed instruction rules, with the precondition $p\downarrow$ and some post-condition p'. We use $p\downarrow$ because there is an implicit step executing delayed writes before executing every instruction. The post-condition p' for \mathbf{i} is then used as the precondition to verify the remaining part of the instruction sequence.

Delayed Control Transfers. We distinguish the jmp and call instructions — the former makes an "intrafunction" control transfer, while the latter makes function calls. The JMP rule requires that the target address is a valid one specified in Ψ . Starting from the precondition p, after executing the instruction i following JMP and the corresponding delayed writes, the post-condition p' of i should satisfy the precondition of the target instruction sequence, with some instantiation ι of the logical variables and a frame assertion p_r . Since the target instruction sequence of jmp is in the same function as the jmp instruction itself, the postcondition fq specified at the target address (with the same instantiation ι of the logical variables and the frame assertion p_r) should meet the post-condition qof the current function. As we explained before, the post-condition q does not specify the states reached at the end of the instruction sequence (which are specified by p' instead).

The CALL rule is similar to the JMP rule in that it also requires the post-condition p_2 of the instruction i following the call satisfy the precondition of the target instruction sequence, with some instantiation ι of the logical variables and a frame assertion p_r . Here we need to record that the code label **f** is saved in \mathbf{r}_{15} by the call instruction. When the callee returns, its postcondition fq (with the same instantiation of auxiliary

 $\vdash C: \Psi$

(Well-Formed Code Heap)

$$\frac{\text{for all } \mathbf{f} \in \text{dom}(\Psi), \ \iota : \Psi(\mathbf{f}) = (\text{fp}, \text{fq}) \quad \Psi \vdash \{(\text{fp } \iota, \text{fq } \iota)\} \ \mathbf{f} : C[\mathbf{f}]}{\vdash C : \Psi} \ (\mathsf{CDHP})$$

$$\Psi \vdash \{(p,q)\} \texttt{f} \,:\, \mathbb{I}$$

(Well-Formed Instruction Sequences)

$$\frac{\vdash \{p\downarrow\} \mathbf{i} \{p'\} \quad \Psi \vdash \{(p',q)\} \mathbf{f} + 4 : \mathbb{I}}{\Psi \vdash \{(p,q)\} \mathbf{f} : \mathbf{i} ; \mathbb{I}} \quad (SEQ)$$

$$p\downarrow \Rightarrow (\mathbf{a}=_a \mathbf{f}') \quad \mathbf{f}' \in \operatorname{dom}(\Psi) \quad \Psi(\mathbf{f}') = (\operatorname{fp}, \operatorname{fq})$$

$$\frac{\vdash \{p\downarrow\downarrow\} \mathbf{i} \{p'\} \quad \exists \iota, p_r. (p' \Rightarrow \operatorname{fp} \iota * p_r) \land (\operatorname{fq} \iota * p_r \Rightarrow q)}{\Psi \vdash \{(p,q)\} \mathbf{f} : \operatorname{jmp} \mathbf{a}; \mathbf{i}} \quad (JMP)$$

$$\frac{f' \in \operatorname{dom}(\Psi) \quad \Psi(\mathbf{f}') = (\operatorname{fp}, \operatorname{fq}) \quad \Psi \vdash \{(p',q)\} \mathbf{f} + 8 : \mathbb{I}}{p\downarrow \Rightarrow (\mathbf{r}_{15} \mapsto _) * p_1 \quad \vdash \{(\mathbf{r}_{15} \mapsto \mathbf{f} * p_1)\downarrow\} \mathbf{i} \{p_2\}}$$

$$\frac{\exists \iota, p_r. (p_2 \Rightarrow \operatorname{fp} \iota * p_r) \land (\operatorname{fq} \iota * p_r \Rightarrow p') \land (\operatorname{fq} \iota \Rightarrow \mathbf{r}_{15} = \mathbf{f})}{\Psi \vdash \{(p,q)\} \mathbf{f} : \operatorname{call} \mathbf{f}'; \mathbf{i}; \mathbb{I}} \quad (CALL)$$

$$\frac{p\downarrow\downarrow \Rightarrow (\mathbf{r}_{15} \mapsto \mathbf{f}') * p_1 \quad \vdash \{p_1\} \mathbf{i} \{p_2\} \quad (\mathbf{r}_{15} \mapsto \mathbf{f}') * p_2 \Rightarrow q}{\Psi \vdash \{(p,q)\} \mathbf{f} : \operatorname{retl}; \mathbf{i}} \quad (RETL)$$

$$\frac{p\downarrow \Rightarrow (\mathbf{z} \mapsto w) * \operatorname{true} \quad \mathbf{f}' \in \operatorname{dom}(\Psi) \quad \Psi(\mathbf{f}') = (\operatorname{fp}, \operatorname{fq})}{\vdash \{p\downarrow\downarrow\} \mathbf{i} \{p'\} \quad \Psi \vdash \{(p' \land w = 0, q)\} \mathbf{f} + 8 : \mathbb{I}} \quad \\ \frac{\exists \iota, p_r. ((p' \land w \neq 0) \Rightarrow \operatorname{fp} \iota * p_r) \land (\operatorname{fq} \iota * p_r \Rightarrow q)}{\Psi \vdash \{(p,q)\} \mathbf{f} : \operatorname{tb} \mathbf{f}'; \mathbf{i}; \mathbb{I}} \quad (BE)$$

 $\vdash \{p\}\,\mathtt{i}\,\{q\}$

(Well-Formed Instructions)

$$\frac{\operatorname{sr} \mapsto _ * p \Rightarrow (\mathbf{r}_s = w_1 \land \mathbf{o} = w_2)}{\vdash \{\operatorname{sr} \mapsto _ * p\} \operatorname{wr} \mathbf{r}_s \circ \operatorname{sr} \{(\triangleright_3 \operatorname{sr} \mapsto (w_1 \oplus w_2)) * p\}} (WR) \\ \xrightarrow{\vdash \{\operatorname{sr} \mapsto w * \mathbf{r}_d \mapsto _\}} \operatorname{rd} \operatorname{sr} \mathbf{r}_d \{\operatorname{sr} \mapsto w * \mathbf{r}_d \mapsto w\}} (RD) \\ p \Rightarrow (\mathbf{r}_s = v_1 \land \mathbf{o} = v_2) \qquad w'_{id} = \operatorname{next_cwp}(w_{id}) \qquad w \& 2^{w'_{id}} = 0 \qquad v = v_1 + v_2 \\ p \Rightarrow (\operatorname{cwp} \mapsto (w_{id}, F \cdot _ \cdot _)) * (\operatorname{out} \mapsto \operatorname{fm}_o) * (\operatorname{local} \mapsto \operatorname{fm}_l) * (\operatorname{in} \mapsto \operatorname{fm}_i) * p_1 \\ \frac{(\operatorname{cwp} \mapsto (w'_{id}, \operatorname{fm}_l :: \operatorname{fm}_i :: F)) * (\operatorname{out} \mapsto _) * (\operatorname{local} \mapsto _) * (\operatorname{in} \mapsto \operatorname{fm}_o) * p_1 \Rightarrow \mathbf{r}_d \mapsto _ * p_2}{\vdash \{(\operatorname{wim} \mapsto w) * p\} \operatorname{save} \mathbf{r}_s \circ \mathbf{r}_d \{(\operatorname{wim} \mapsto w) * (\mathbf{r}_d \mapsto v) * p_2\}} (SAVE) \\ \text{where } [\mathbf{r}_i, \dots, \mathbf{r}_{i+7}] \mapsto [w_0, \dots, w_7] ::= \mathbf{r}_i \mapsto w_0 * \cdots * \mathbf{r}_{i+7} \mapsto w_7 \\ \text{and out, local and in are defined in Fig.6.} \\ p \Rightarrow (\mathbf{r}_s = v_1 \land \mathbf{o} = v_2) \qquad w'_{id} = \operatorname{prev_cwp}(w_{id}) \qquad w \& 2^{w'_{id}} = 0 \qquad v = v_1 + v_2 \\ p \Rightarrow (\operatorname{cwp} \mapsto (w_{id}, \operatorname{fm}_1 :: \operatorname{fm}_2 :: F)) * (\operatorname{out} \mapsto _) * (\operatorname{local} \mapsto _) * (\operatorname{in} \mapsto \operatorname{fm}_i) * p_1 \\ (\operatorname{cwp} \mapsto (w'_{id}, \operatorname{fm}_1 :: \operatorname{fm}_2 :: F)) * (\operatorname{out} \mapsto _) * (\operatorname{local} \mapsto _) * (\operatorname{in} \mapsto \operatorname{fm}_i) * p_1 \\ (\operatorname{cwp} \mapsto (w'_{id}, \operatorname{fm}_1 :: \operatorname{fm}_i) * (\operatorname{local} \mapsto \operatorname{fm}_1) * (\operatorname{in} \mapsto \operatorname{fm}_i) * p_1 \\ \vdash \{(\operatorname{wim} \mapsto w) * p\} \operatorname{restore} \mathbf{r}_s \circ \mathbf{r}_d \{(\operatorname{wim} \mapsto w) * (\mathbf{r}_d \mapsto v) * p_2\}} (RESTORE)$$

Fig.11. Selected inference rules^[4].

variables ι) needs to ensure \mathbf{r}_{15} still contains \mathbf{f} , so that the callee returns to the correct address. Also the fq with the frame p_r needs to satisfy the precondition p'for the remaining instruction sequences of the caller.

The RETL rule simply requires that the postcondition q hold at the end of the instruction i following ret1. Also i cannot touch the register \mathbf{r}_{15} ; therefore \mathbf{r}_{15} specified in p must be the same as in q. Since at the calling point we have already required that the postcondition of the callee guarantees \mathbf{r}_{15} contain the correct return address, we know \mathbf{r}_{15} contains the correct value before ret1. The BE rule checks the "current" value of the register z and decides whether the branch will be taken after executing the following instruction i. If z is not zero, the branch is taken and it does the same things as the JMP rule; otherwise, the branch is not taken and the remaining instruction sequence I should be well-formed.

Delayed Writes and Register Windows. The bottom layer of our logic is for well-formed instructions. The WR rule requires the ownership of the target register \mathbf{sr} in the precondition ($\mathbf{sr} \mapsto \underline{}$). Also it implies there is no delayed write to \mathbf{sr} in the delay buffer (see the semantics defined in Fig.9). At the end of the delayed

write, we use $\triangleright_3 \mathbf{sr} \mapsto w_1 \oplus w_2$ to indicate the new value will be ready in up to three cycles. Since the maximum delay cycle X cannot be bigger than 3 and the value of X may vary in different systems, programmers usually take a conservative approach to assume X = 3 for the portability of code. Our rule reflects this conservative view. The RD rule says the special register can be read only if it is not in the delay buffer. The SAVE and RESTORE rules reflect the save and recovery of the execution contexts, respectively, which is consistent with the operational semantics of the **save** and **restore** instructions given in Figs.6 and 7 respectively.

4 Refinement Verification of SPARCv8

In this section, we present our relational program logic for refinement verification of SPARCv8 code⁽⁵⁾. As an extension of our conference paper, it consists of the following work.

• We define a new Pseudo-SPARCv8 language as the high-level specification language in Subsection 4.1.

• We make some modifications to the SPARCv8 language defined in Section 2 and let it act as the low-level language in our refinement verification. We present our low-level SPARCv8 language and the modifications in Subsection 4.2.

• We define the correctness of abstract assembly primitives in Subsection 4.3, which is formulated as "contextual refinement" between the implementations and their corresponding abstract assembly primitives.

• We present a new program logic to do refinement verification in Subsection 4.4.

• We show that our new program logic is sound in Subsection 4.5. The semantics of judgements, different from the previous work in our conference paper^[4] which only ensures the partial correctness of verified programs, is defined as simulation relations between the low- and high-level programs, which guarantees contextual refinement.

4.1 High-Level Pseudo-SPARCv8 Language

The Pseudo-SPARCv8 language contains two parts: the SPARCv8 code as the client and the set of abstract assembly primitives. Here, the execution of the client SPARCv8 code is required to preserve a restriction between register windows and stacks in memory, shown in Fig.12(a) (cwp points to the current window and wim marks the invalid window, and the details of overlapping of adjacent windows are omitted in the figure).



Fig.12. Abstraction of context management. (a) Reigster windows and stack in memory. (b) High-level frame list.

During the execution of the SPARCv8 program, parts of previous procedures' contexts (in Fig.12(a)) are saved in register windows, and the others (the dark gray part in Fig.12(a)) are stored in the stack in memory, since the number of windows is limited. The restriction is that the stack pointer (% sp) of each procedure, including the current and previous ones, whose context is saved in register windows currently, should point to the top of its stack frame (shown as the thick arrow in Fig. 12(a), so that the contexts in these windows can be stored correctly in memory when needed. For instance, the context switch routine will check whether the previous window is valid (in clockwise direction), and use the **restore** instruction to set it as the current one and save its contents into the stack (in memory) until the previous one (filled with east north lines in Fig. 12(a)) is invalid. The executions of client code are required to preserve such a restriction. Otherwise, some SPARCv8 functions, e.g., context switch routine, whose execution stores the contexts saved in register windows into stacks in memory, cannot be verified if it is unclear where to save the contents of some windows. We do the following when defining the Pseudo-SPARCv8 program to make the client code execution preserve such a restriction.

• In order to ensure that the stack pointer (%sp) always points to the top of its stack frame, we require that each instruction, e.g., add and ld, whose execution does not operate register windows, is prohibited to modify %sp. As for the save and restore, we restrict them to be used in specific forms. We define "Psave w" as a macro of "save %sp, -w, %sp", whose execution generates a new %sp pointing to the stack frame size w allocated newly for the next window. We also define

⁽⁵⁾We give more details about the language definition and the soundness proof in the Coq implementation and the technical report, which is available at https://github.com/jpzha/VeriSparc, Sept. 2020.

 $^{^{(6)}}$ In SPARCv8, $\% g_0$ is always 0, and usually used as a parameter when instructions do not require a specific parameter.

"Prestore" as a macro of "restore $\%g_0, \%g_0, \%g_0$ "⁽⁶⁾, whose execution just restores the previous window and does not modify the value of any register in the previous window. The original save and restore instructions have no semantics in the high-level client code.

• The special registers in SPARCv8 usually play specific roles and modifying them should be done carefully. For example, wim marks which window is invalid.

If we change its value shown in Fig.12(a) to mark another window invalid, as shown in Fig.13, and call context switch routine to save contents of the previous windows into memory until the invalid one at this moment, a problem will arise since we do not know where to save the contents of window marked invalid originally (filled with dots in Fig.13). Therefore, we forbid the client to modify special registers and give *no* semantics to instruction wr in high-level client code. Modifying them is hidden in the implementations of abstract assembly primitives in the low level. And the delay buffer can be omitted in high-level program state.



Fig.13. Problem of modifying wim arbitrary.

• As shown in Fig.12, we find that we can abstract the register windows and the stack in memory storing

contexts into a list (shown in Fig.12(b)). After this abstraction, we do not need to care about whether contexts are saved in register windows or memory, and do not need to describe the contents of the windows unused (the windows in white in Fig.12(a), but excluding the current one pointed by cwp) in Pseudo-SPARCv8. The cwp register is no longer needed in Pseudo-SPARCv8 since the register windows are abstracted away. The low-level program in our work does not use such abstraction, since the low-level program should be realistically modelled, and the implementations of some primitives need to know the existence of register windows, e.g., the context switch routine that saves the contents of register windows into stacks (in memory).

We define the syntax of the high-level Pseudo-SPARCv8 language in Fig.14. The code (HCode) Π includes the code heap C and the set of abstract primitives (PrimSet) Ω , which is a partial mapping from labels to abstract assembly primitives. The code heap Cin Π acts as the client to call abstract assembly primitives. The abstract assembly primitive (Prim) Υ is defined as a relation that takes a list of values as arguments and maps a high-level program state (defined in Fig.15) to another. We add three pseudo instructions in simple instruction (SimpIns). The Psave w and Prestore restrict the save and restore instructions to be used in the specific form as mentioned before. We also introduce print r, whose execution outputs the value v in general register r and generates a message out(v) as an observable event. The high-level message (HMsg) α can be either an empty message τ , or an output out(v), or a $call(f, \overline{v})$ meaning to call a primitive labelled **f** with arguments \overline{v} .

The machine states (HState) of Pseudo-SPARCv8 are defined in Fig.15. The high-level program \mathbb{P}

Fig.14. Syntax of Pseudo-SPARCv8 code.

(HProg) (Tid)	\mathbb{P} ::= t \in	$\mathbb{Z}^{(\Pi,\mathbb{S})}$	(HState) (ThrdLcSt)	$ \begin{split} \mathbb{S} & ::= (T, t, \mathcal{K}, \\ \mathcal{K} & ::= (\mathbb{Q}, pc, r) \end{split} $,M)		(ThrdPool) (HRstate)	$\begin{array}{c} T & ::= \\ \mathbb{Q} & ::= \end{array} \left(\begin{array}{c} \\ \end{array} \right)$	$\{t\rightsquigarrow\mathcal{K}\}^{*}$ (\mathbb{R},\mathbb{F})
(HRegFile) (HFrmList)	$\mathbb{R} \in \mathbb{F} ::=$	HRegNam nil $ (fm_1, fm_2) $	$e \rightarrow Val$ fm ₂):: \mathbb{F}	(HRegName) (HFrame)	rn fm	::= ::=	$\mathbf{r}_0 \mid \ldots \mid \mathbf{r}_3$ $[v_0, \ldots, v_7]$	n z	c v

Fig.15. Machine states for Pseudo-SPARCv8 code.

(HProg) is a pair of the high-level code Π and the highlevel state S. The high-level state S is a tuple including: the thread pool (ThrdPool) T, the current thread ID (Tid) t, the thread local state (ThrdLcSt) \mathcal{K} of the current thread, and the memory (Mem) M.

Thread Local State. The thread local state \mathcal{K} is a triple of high-level register state (HRstate) \mathbb{Q} , and program counters pc and npc. The high-level register state $\mathbb Q$ consists of the high-level register file (HRegFile) $\mathbb R,$ and the high-level frame list (HFrmList) \mathbb{F} , which is the abstraction of the register windows and the memory storing contexts in the low level. rn is the high-level register names (HRegName), where the cwp register is omitted as introduced before and we also omit special registers for simplicity, because we forbid the high-level client code to modify them⁽⁷⁾. The high-level frame list \mathbb{F} is a list of pairs (fm₁, fm₂), which is used to save the contexts (local and in registers) fm_1 and fm_2 of previous procedures. Then, we define the switch primitive as an instantiation of Υ below.

switch ::= $\lambda \overline{v}, \mathbb{S}, \mathbb{S}', \exists t'. M(\mathsf{TaskNew}) = (t', 0) \land T(t') = (\mathbb{Q}', \mathsf{pc}', \mathsf{npc}')$ $\wedge T' = T\{\mathsf{t} \rightsquigarrow (\mathbb{Q}, \mathsf{pc}, \mathsf{npc})\} \land \mathsf{t} \neq \mathsf{t}' \land \overline{v} = \mathrm{nil},$ where $\mathbb{S} = (T, \mathsf{t}, (\mathbb{Q}, \mathsf{pc}, \mathsf{npc}), M),$ $\mathbb{S}' = (T', \mathsf{t}', (\mathbb{Q}', \mathsf{f} + 8, \mathsf{f} + 12), M), \mathsf{f} = \mathbb{Q}'.\mathbb{R}(\mathsf{r}_{15}).$

The switch primitive takes no arguments ($\overline{v} = nil$), and changes the identifier of the current thread according to the value in the location TaskNew. We use parameters S and S' to represent the machine states before and after the execution of switch respectively.

Operational Semantics in High Level. The operational semantics for Pseudo-SPARCv8 is defined in Fig.16. The high-level program transition relation

 $(\mathbf{D} \mathbf{D})$

 $(\Pi, \mathbb{S}) : \stackrel{\alpha}{\Longrightarrow} (\Pi, \mathbb{S}')$ is defined in Fig.16(a). In each step, the program may either execute the instruction pointed by pc and generate empty message τ or an output out(v), or call an abstract assembly primitive in the primitive set. When calling an abstract assembly primitive, a message $call(f, \overline{v})$ will be generated. It means that we hope to call the abstract assembly primitive labelled **f** with the arguments \overline{v} (args($\mathbb{Q}, M, \overline{v}$), whose definition is omitted here, means getting the arguments \overline{v} from \mathbb{Q} and M).

The thread local step is defined in Fig.16(b). The step for simple instruction i is represented as "exec(i,_,_)". We show the state transition relation for pseudo instructions Psave w and Prestore in Fig.16(c). Supposing the current register state \mathbb{Q} is (\mathbb{R}, \mathbb{F}) , the execution of Psave *w* will save the local and in registers into the high-level frame list \mathbb{F} . It also allocates a new block b of size 64-byte to w bytes as a new stack frame in memory (shown as alloc(M, b, 64, w) =M'). The reason why it starts from 64-byte is that the 0-64 bytes (16 words) in a stack frame are usually reserved to save the context in window (local and in registers) by convention, and this part of memory is abstracted away in the Pseudo-SPARCv8 program as we have explained and shown in Fig.12. Prestore does the reverse, freeing the block of the current stack frame (shown as free(b, M) = M'), and restoring the context of the previous procedure saved in \mathbb{F} .

4.2 Low-Level SPARCv8 Program

The global program transition of the low-level SPARCv8 program is defined as the following form.

$$\begin{split} & (R,D) \rightrightarrows (R',D') \\ & C \vdash ((M,(R',F),D'),\texttt{pc},\texttt{npc}) \circ \xrightarrow{\tau/\texttt{out}(v)} ((M',(R'',F'),D''),\texttt{pc}',\texttt{npc}') \\ & (C,(M,(R,F),D),\texttt{pc},\texttt{npc}) :: \xrightarrow{\tau/\texttt{out}(v)} (C,(M',(R'',F'),D''),\texttt{pc}',\texttt{npc}') \end{split}$$

The low-level SPARCv8 program is slightly different from the SPARCv8 program defined in Section 2.

1) The low-level SPARCv8 program uses the instructions defined in Fig.14. It means that we need to give semantics to the pseudo instructions Psave, Prestore and print in the low-level SPARCv8 program. Since Psave and Prestore are simply special forms of save and restore as explained, and print is a primitive responsible for generating observable events, defining their semantics is not a challenge and the translation of programs in this modified language into ones in the standard SPARCv8 language is trivial.

2) The program transition defined in Section 2 does not generate observable events. Here, since we want to support refinement verification and use the event trace refinement^[9], each step of the program generates either an empty message τ , or an output out(v) produced by print.

⁽⁷⁾There is no problem to reserve special registers in the high-level register file and permit the high-level client code to read them.

Fig.16. Selected operational semantics rules for high-level program. (a) High-level program transitions. (b) High-level thread local transitions. (c) High-level instruction transitions.

Note that the client code and the implementations of primitives in the low level are both SPARCv8 code heap. There is no need to define their linking in semantics.

4.3 Primitive Correctness

We first establish a state relation between low- and high-level program states. We define this relation below (\uplus means disjoint union), shown as " $S \sim \mathbb{S}$ ".

$$\frac{M = M_c \uplus M_T \uplus \{\mathsf{TaskCur} \rightsquigarrow (\mathfrak{t}, 0)\} \uplus M'}{(M_c, Q) \Downarrow_{\mathsf{c}} (\mathfrak{t}, \mathcal{K}) \qquad M_T \Downarrow_{\mathsf{r}} T \setminus \{\mathfrak{t}\} \qquad D = \operatorname{nil}}{(M, Q, D) \sim (T, \mathfrak{t}, \mathcal{K}, M')}$$

The low-level memory M is split into four parts: M_c used saving the context of the current thread t; M_T saving the contexts of the ready threads, except the current thread t; a singleton memory cell located TaskCur saving the current thread ID; and the shared memory M' that is the same as the high-level memory. The delayed buffer D is nil, since the client is not permitted to modify any special register through the wr instructions. M_T is "abstracted" as a thread pool in the high-level program. Their relation is represented as " $M_T \Downarrow_r T \setminus \{t\}$ ". We use " $(M_c, Q) \Downarrow_c (t, \mathcal{K})$ " to represent the state relation of the current thread in low- and high-level programs.

The correctness of abstract assembly primitives is defined in terms of contextual refinement. We give its formal definition in Definition 1. And we use the "event trace refinement" proposed by Liang *et al.*^[9]

Definition 1 (Primitive Correctness). $C_{as} \sqsubseteq \Omega$ *iff* for any C, S, \mathbb{S} , pc and npc, if $S \sim \mathbb{S}$ and $ProgSafe(\mathbb{P})$, then $P \subseteq \mathbb{P}$, where $P = (C \uplus C_{as}, S, pc, npc)$, $\mathbb{P} = ((C, \Omega), \mathbb{S})$ and $\mathbb{S}.\mathcal{K} = (_, pc, npc)$.

We use the code heap $C_{\rm as}$ to represent the implementations and Ω to represent the set of corresponding

abstract assembly primitives. The contextual refinement, denoted as $C_{\rm as} \sqsubseteq \Omega$, says that if and only if for any client code (or context) C, low-level program state S, high-level program state \mathbb{S} , program counters pc and npc, if the low- and high-level program states satisfy the state relation $S \sim \mathbb{S}$ and the high-level program will never get stuck (shown as $\operatorname{ProgSafe}(\mathbb{P})$), then there is an event trace refinement ^[9], which means that P produces no more observable behaviors than \mathbb{P} and is denoted as $P \subseteq \mathbb{P}$, between low- and high-level programs. $\operatorname{ProgSafe}(\mathbb{P})$ is defined formally below:

$$\mathsf{ProgSafe}(\mathbb{P}) ::= \forall \mathbb{P}'.(\mathbb{P} :\Longrightarrow^* \mathbb{P}') \Longrightarrow (\exists \mathbb{P}''.\mathbb{P}' :\Longrightarrow \mathbb{P}'')$$

The client code C is the SPARCv8 code. Therefore, the high-level code is a pair of C and Ω , and the low-level code is just the union of C and C_{as} , shown as $(C \uplus C_{as})$, because both of C and C_{as} are SPARCv8 code heaps.

$$C \uplus C_{as} ::= C \cup C_{as}$$
 if $\operatorname{dom}(C) \cap \operatorname{dom}(C_{as}) = \emptyset$.

4.4 Relational Program Logic for Refinement Verification

Relational Assertion. Fig.17 gives the relational assertion language, and its semantics is given in Fig.18. The relational assertions are interpreted over the relational state (S, S, A, w), which contains the low-level state S, the high-level state S, the abstract assembly primitive command A defined in Fig.18, and the word w recording the number of the tokens. The high-level primitive command A is either an abstract assembly primitive Υ parameterized over arguments \overline{v} , or a \bot meaning the primitive has already been executed. The relational assertion p reserves the original assertion pdescribing the low-level state S.

$$\begin{array}{ll} \text{(RelAsrt)} & \texttt{p}, \texttt{q} & ::= p \mid \hat{\texttt{rn}} \rightarrowtail v \mid l \rightarrowtail v \mid \texttt{Emp} \\ & \mid \texttt{t} \leadsto_{\texttt{c}} \mathcal{K} \mid \texttt{t} \leadsto_{\texttt{r}} \mathcal{K} \mid (A) \mid \blacklozenge(w) \\ & \mid \texttt{p} \downarrow \mid \texttt{p} \land \texttt{q} \mid \texttt{p} \lor \texttt{q} \mid \texttt{p} \ast \texttt{q} \mid .. \end{array}$$



We define $\hat{\mathbf{rn}} \to v$ and $l \to v$ to describe the state of register files and memory in the high level. The assertion Emp says that the high-level memory and the thread pool are both empty, and the low-level state satisfies emp defined in Fig.9. The assertions $\mathbf{t} \rightsquigarrow_{\mathbf{c}} \mathcal{K}$ and $\mathbf{t} \rightsquigarrow_{\mathbf{r}} \mathcal{K}$ represent the thread local state of the current and ready thread respectively. Note that the threads in the thread pool are viewed as resources and can be separated by separation conjunction.

The assertion (A) means the current high-level primitive command is A. And the assertion $\blacklozenge(w)$ takes a word w, which can also be separated by separation conjunction (*), to state that the number of tokens in current state is "no less" than w. Tokens are used to avoid infinite loops and recursive function calls to make sure the termination preserving refinement.

The assertion $p \downarrow$, which is similar to $p \downarrow$ defined in Fig.9, describes the state after executing one step of

$$\begin{split} & (S, \mathbb{S}, A, w) \models \operatorname{Emp} & :::= \mathbb{S}.M = \emptyset \land \mathbb{S}.T = \emptyset \land S \models \operatorname{emp} \\ & (S, \mathbb{S}, A, w) \models \widehat{\mathbf{n}} \mapsto v & :::= S \models p \land \mathbb{S}.M = \emptyset \land \mathbb{S}.T = \emptyset \\ & (S, \mathbb{S}, A, w) \models \widehat{\mathbf{n}} \mapsto v & :::= \mathbb{S}.K.\mathbb{Q}.\mathbb{R}(\widehat{\mathbf{n}}) = v \land (S, \mathbb{S}, A, w) \models \operatorname{Emp} \\ & (S, \mathbb{S}, A, w) \models i \mapsto v & :::= \mathbb{S}.M = \{l \rightsquigarrow v\} \land \mathbb{S}.T = \emptyset \land S \models \operatorname{emp} \\ & (S, \mathbb{S}, A, w) \models i \rightsquigarrow_{\mathbf{c}} \mathcal{K} & :::= \mathbb{S}.T \setminus \{t\} = \emptyset \land \mathbb{S}.t = t \land \mathbb{S}.\mathcal{K} = \mathcal{K} \land \mathbb{S}.M = \emptyset \land S \models \operatorname{emp} \\ & (S, \mathbb{S}, A, w) \models i \rightsquigarrow_{\mathbf{c}} \mathcal{K} & :::= \mathbb{S}.T = \{t \rightsquigarrow \mathcal{K}\} \land \mathbb{S}.M = \emptyset \land t \neq \mathbb{S}.t \land S \models \operatorname{emp} \\ & (S, \mathbb{S}, A, w) \models (\mathcal{A}') & :::= A = A' \land (S, \mathbb{S}, A, w) \models \operatorname{Emp} \\ & (S, \mathbb{S}, A, w) \models (\mathcal{A}') & :::= w' \leqslant w \land (S, \mathbb{S}, A, w) \models \operatorname{Emp} \\ & (S, \mathbb{S}, A, w) \models (\mathcal{A}') & :::= i \exists S'. ((S', \mathbb{S}, A, w) \models \mathbb{P}) \land (R', D') \Rightarrow (R, D) \\ & where \quad S = (M, (R, F), D), \ S' = (M, (R', F), D') \\ & (S, \mathbb{S}, A, w) \models \mathbb{P} \star q & :::= \exists S_1, S_2, \mathbb{S}_1, \mathbb{S}_2, w_1, w_2. S = S_1 \uplus S_2 \land \mathbb{S} = \mathbb{S}_1 \uplus \mathbb{S}_2 \land w = w_1 + w_2 \land (S_1, \mathbb{S}_1, A, w_1) \models \mathbb{P} \land (S_2, \mathbb{S}_2, A, w_2) \models \mathfrak{q} \\ & T_1 \perp T_2 & :::= (\operatorname{dom}(T_1) \cap \operatorname{dom}(T_2)) = \emptyset \\ & \mathbb{S}_1 \uplus \mathbb{S}_2 & :::= \begin{cases} (T_1 \cup T_2, t, \mathcal{K}, M_1 \cup M_2), & \operatorname{if} \ T_1 \perp T_2 \land M_1 \perp M_2 \land & \mathbb{S}_1 = (T_1, t, \mathcal{K}, M_1) \land \mathbb{S}_2 = (T_2, t, \mathcal{K}, M_2), \\ & \operatorname{undefined}, & \operatorname{otherwise.} \\ \\ & (\operatorname{HPrimCom}) \ A & :::= \Upsilon(\overline{v}) \mid \perp & \underbrace{\Upsilon(\overline{v})(\mathbb{S})(\mathbb{S}')}{(\Upsilon(\overline{v}, \mathbb{S}) = \cdots \to (\bot, \mathbb{S}')} \end{cases} \end{cases}$$

Fig.18. Semantics of relation assertion.

delayed writes.

Inference Rules in Relational Program Logic. The code specification $\hat{\theta}$ and the code heap specification Ψ in relational program logic are defined below.

$$\begin{split} &(\mathrm{valList})\iota \in \mathrm{list} \ \mathrm{value} \\ &(\mathrm{pAsrt}) \ \mathrm{fp}, \mathrm{fq} \in \mathrm{valList} \to \mathrm{RelAsrt} \\ &(\mathrm{CdSpec})\hat{\theta} ::= (\mathrm{fp}, \mathrm{fq}) \\ &(\mathrm{CdHpSpec}) \ \Psi \ ::= \{ \mathtt{f} \rightsquigarrow \hat{\theta} \}^* \end{split}$$

Here, fp and fq are relational assertions parameterized over a list of values ι . Fig.19 shows selected inference rules for refinement verification in our logic. The top rule WfPrim verifies the contextual refinement between the code heap $C_{\rm as}$ and the corresponding abstract assembly primitive set Ω . It requires that each code block specified in Ψ can be verified with respect to its specification, shown as $(\vdash C_{as} : \Psi)$, and the specifications of the implementations of abstract assembly primitives need to meet some restrictions, shown as wdSpec(fp, fq, Υ), which we will discuss in more details following. The inference rules for jmp and call in relational program logic will consume a token, shown as (1), in order to avoid infinite loops and recursive function calls and ensure termination-preserving. $wf(p_r)$, whose definition is omitted here, means there is no subterm in form of $(\mathbf{t} \rightsquigarrow_{\mathbf{c}} \mathcal{K})$, $(\hat{\mathbf{rn}} \rightarrowtail v)$ and (A) in the frame p_r , because the state they described is not separated by separation conjunction *. The ABSCSQ rule allows us to execute the high-level primitive command specified in precondition. The implication $p \Rightarrow p'$ is defined below formally:

$$\begin{aligned} (\mathfrak{p} \Rightarrow \mathfrak{p}') \lor \\ (\forall S, \mathbb{S}, A, w. ((S, \mathbb{S}, A, w) \models \mathfrak{p}) \Longrightarrow \\ (\exists \mathbb{S}', A', w'. ((A, \mathbb{S}) \dashrightarrow (A', \mathbb{S}')) \land ((S, \mathbb{S}', A', w') \models \mathfrak{p}')). \end{aligned}$$

The inference rules for instructions are not shown here, since they are not different from the rules in Fig.11.

Well-Defined Specification. wdSpec(fp, fq, Υ) is defined formally in Definition 2. It contains three properties that the specification needs to satisfy.

Definition 2 (Well-Defined Specification). wdSpec(fp, fq, Υ) *holds, iff*:

1) for any \overline{v} , \mathbb{S} , \mathbb{S}' , \mathbb{S}_r , if $\Upsilon(\overline{v})(\mathbb{S})(\mathbb{S}')$, and $\mathbb{S} \perp \mathbb{S}_r$, then the following holds (where $\mathbf{f} = \mathbb{S}'.\mathcal{K}.\mathbb{Q}.\mathbb{R}(\mathbf{r}_{15})$):

• $\mathbb{S}'.\mathcal{K}.pc = f + 8$, $\mathbb{S}'.\mathcal{K}.npc = f + 12$;

• there exist \mathbb{S}'' and \mathbb{S}'_r , such that $\Upsilon(\overline{v})(\mathbb{S} \uplus \mathbb{S}_r)(\mathbb{S}'')$, $\mathbb{S}'' = \mathbb{S}' \uplus \mathbb{S}'_r$, and $\mathbb{S}_r.T = \mathbb{S}'_r.T$, $\mathbb{S}_r.M = \mathbb{S}'_r.M$;

2) for any ι , there exists \overline{v} , such that

 $\texttt{fp } \iota \Rightarrow (\Upsilon(\overline{v})) * \texttt{true}, \text{ and } \texttt{fq } \iota \Rightarrow (\bot) * \texttt{true};$

3) for any $\overline{v}, S, \mathbb{S}$, if $(S, \mathbb{S}, _, _) \in \mathsf{INV}(\Upsilon(\overline{v}), \overline{v})$, there exist ι, \mathfrak{p}_r and w, such that



Fig.19. Selected inference rules for refinement verification.

 $(S, \mathbb{S}, \Upsilon(\overline{v}), w) \models (\mathfrak{fp} \ \iota * \mathfrak{p}_r), (\mathfrak{fq} \ \iota * \mathfrak{p}_r) \Rightarrow \mathsf{INV}(\bot, _),$ and $\mathsf{Sta}(\Upsilon(\overline{v}), \mathfrak{p}_r)$ hold.

First, the program counters should equal $\mathbf{f} + 8$ and $\mathbf{f} + 12$ respectively, where \mathbf{f} is contained in the \mathbf{r}_{15} register after the execution of the abstract assembly primitive Υ . It ensures that low-level implementations and corresponding high-level abstract assembly primitives return to the same code pointers. We also require that if an abstract assembly primitive can execute safely on a part of program state, it can also execute safely on the whole program state, and keeps the additional program state unchanged. $\mathbb{S} \perp \mathbb{S}_r$ is defined formally below:

$$\begin{split} \mathbb{S} \perp \mathbb{S}_r &::= T \perp T' \land M \perp M' \land \mathsf{t} = \mathsf{t}' \land \mathcal{K} = \mathcal{K}' \\ \text{where } & \mathbb{S} = (T, \mathsf{t}, \mathcal{K}, M), \, \mathbb{S}_r = (T', \mathsf{t}', \mathcal{K}', M') \end{split}$$

Second, the abstract assembly primitive should be specified in the precondition, and its execution should be done in the final state. Third, an "invariant" holds between low- and high-level programs at the entry of the function. Our logic needs to ensure that such invariant can be reestablished when the function returns. We define such invariant as INV formally below:

$$\begin{aligned} \mathsf{INV}(A,\overline{v}) &::= \{(S,\mathbb{S},A,w) \mid S \sim \mathbb{S} \land (\exists \, \mathbb{S}'(\fbox{\bullet}) \dashrightarrow (\bot,\mathbb{S}')) \\ \land \operatorname{args}(\mathbb{S}.\mathcal{K}.\mathbb{Q}\mathbb{S}.M,\overline{v}) \}. \end{aligned}$$

The invariant consists of the state relation between lowand high-level program states, shown as $S \sim S$ and defined in Subsection 4.3, and the safe execution of the primitive command A, which means that A can execute zero (if $A = \bot$) or one step (if $A = \Upsilon(\overline{v})$) from the current state, shown as $\exists S'. (A, S) \dashrightarrow^* (\bot, S')$. Including the safe execution of A in the invariant is essential, since we can get some knowledge of the high-level program state from it. For example, if INV(switch(nil), nil) holds, we know that the location TaskNew must save a pointer pointing to a ready thread in the thread pool; otherwise the switch primitive cannot execute.

$$\begin{array}{l} \mathsf{INV}(\mathsf{switch}(\mathrm{nil}),\mathrm{nil}) \implies \\ \exists \, \mathsf{t}, \mathcal{K}. \, (\mathsf{t} \rightsquigarrow_{\mathsf{r}} \mathcal{K}) * (\mathsf{TaskNew} \rightarrowtail (\mathsf{t}, 0)) * \mathrm{true} \end{array}$$

We use the frame p_r for local reasoning. $\mathsf{Sta}(\Upsilon(\overline{v}), p_r)$, whose definition is omitted here, says that p_r still holds after the execution of $\Upsilon(\overline{v})$.

4.5 Semantics and Soundness

We first define the simulation relation for instruction sequences. It says $C_{\rm as}$ can execute safely from S, pc and npc until reaching the end of the current instruction sequence $(C_{\rm as}[pc])$, and q holds if $C_{\rm as}[pc]$ ends with the return instruction ret1, and for each step of the low-level execution, the high-level program will execute zero or one step. It is formally defined in Definition 3. Here we use " $_ \mapsto^n _$ " to represent *n*-step execution. The *w* in simulation records the number of tokens. It will be consumed when meeting the jmp and call instructions to avoid infinite loops and recursive in low level for termination preserving, and reset when the high-level abstract assembly primitive executes.

Definition 3 (Simulation for Instruction Sequence). $q; \Psi \models (C_{as}, S, pc, npc) \preccurlyeq_w (A, \mathbb{S})$ holds if and only if the followings are true (we omit the case for be here, which is similar to jmp).

- 1) If $C_{as}(pc) = i$ then:
- there exist $S', pc', npc', such that C_{as} \vdash (S, pc, npc) \longmapsto (S', pc', npc'),$
- for any S', pc', npc',

if
$$C_{as} \vdash (S, pc, npc) \longmapsto (S', pc', npc')$$
, then there exist A', S' and w' , such that:

- a) either A' = A, S' = S and w' = w; or (A,S) --→ (A',S'),
 b) q; Ψ⊨(Cas, S', pc', npc') ≼_{w'} (A', S').
- 2) If $C_{as}(pc) = jmp a$ then:
- there exist S', pc', npc', such that
 C_{as} ⊢ (S, pc, npc) →² (S', pc', npc'),
- for any S', pc', npc', if C_{as} ⊢ (S, pc, npc) →² (S', pc', npc'), then there exist fp, fq, ι, A', S', w', w'' < w' and p_r such that the followings hold:
 - a) $\operatorname{npc}' = \operatorname{pc}' + 4$, $\Psi(\operatorname{pc}') = (\operatorname{fp}, \operatorname{fq})$,
 - b) either A' = A, $\mathbb{S}' = \mathbb{S}$ and w' = w; or $(A, \mathbb{S}) \dashrightarrow (A', \mathbb{S}')$,
 - c) $(S', \mathbb{S}', A', w'') \models (\mathfrak{fp} \iota) * \mathfrak{p}_r, (\mathfrak{fq} \iota) * \mathfrak{p}_r \Rightarrow \mathfrak{q},$ wf(\mathfrak{p}_r).

3) If $C_{\rm as}(pc) =$ bef then ...

4) If $C_{as}(pc) = call f$ then:

- there exist S', pc', npc', such that
 C_{as} ⊢ (S, pc, npc) →² (S', pc', npc'),
- for any S', pc' and npc', if C_{as} ⊢ (S,pc,npc) →² (S',pc',npc'), then there exist fp,fq,ι,A',S',w',w'' < w' and p_r, such that the followings hold:
 - a) $\operatorname{npc}' = \operatorname{pc}' + 4$, $\Psi(\operatorname{pc}') = (\operatorname{fp}, \operatorname{fq})$,
 - b) either A' = A, $\mathbb{S}' = \mathbb{S}$ and w' = w; or $(A, \mathbb{S}) \dashrightarrow (A', \mathbb{S}')$,
 - c) $(S', \mathbb{S}', A', w'') \models (\mathfrak{fp} \iota) * \mathfrak{p}_r, wf(\mathfrak{p}_r),$
 - d) for any S_0 , S_0 , A_0 , w_0 , if $(S_0, \mathbb{S}_0, A_0, w_0) \models (\mathfrak{fq} \ \iota) * \mathfrak{p}_r$, then $\mathfrak{q}; \Psi \models (C_{\mathrm{as}}, S_0, \mathfrak{pc} + 8, \mathfrak{pc} + 12) \preccurlyeq_{w_0} (A_0, \mathbb{S}_0)$,

e) (fg
$$_{L}$$
) \Rightarrow (r₁₅ = pc).
5) If $C_{as}(pc) = retl$ then:

- there exist S', pc', npc', such that
 C_{as} ⊢ (S, pc, npc) →² (S', pc', npc'),
- for any S', pc' and npc',
 if C_{as} ⊢ (S, pc, npc) →² (S', pc', npc'), then
 there exist A', S', and w', such that:
 - a) either A' = A, $\mathbb{S}' = \mathbb{S}$ and w' = w; or $(A, \mathbb{S}) \dashrightarrow (A', \mathbb{S}')$,
 - b) $(S', \mathbb{S}', A', w') \models q$, $pc' = S'.Q.R(r_{15}) + 8$, and $npc' = S'.Q.R(r_{15}) + 12$.

Then we define the semantics for well-formed instruction sequences and well-formed code heap below.

Definition 4 (Judgment Semantics).

• $\Psi \models \{(\mathbf{p}, \mathbf{q})\} \mathbf{f} : \mathbb{I} \text{ if and only if, for all } C_{\mathrm{as}}, S, \mathbb{S}, A \text{ and } w \text{ such that } C_{\mathrm{as}}[\mathbf{f}] = \mathbb{I} \text{ and } (S, \mathbb{S}, A, w) \models \mathbf{p}, we have \mathbf{q}; \Psi \models (C_{\mathrm{as}}, S, \mathbf{f}, \mathbf{f} + 4) \preccurlyeq_w (A, \mathbb{S}).$

• $\models C_{as}$: Ψ if and only if, for all f, fp and fq such that $\Psi(f) = (fp, fq)$, we have $\Psi \models \{(fp \iota, fq \iota)\} f$: $C_{as}[f]$ for all ι .

Next, we define the simulation for functions in Definition 5. It says that if there exists a relational state (S, \mathbb{S}, A, w) satisfying the precondition \mathfrak{p} , then we have the simulation $\mathfrak{q} \models (C_{\mathrm{as}}, S, \mathfrak{f}, \mathfrak{f} + 4) \preccurlyeq_i^0 (A, \mathbb{S})$ defined in Definition 6. The simulation, which ensures the safe execution of the low-level SPARCv8 function and its corresponding high-level abstract assembly primitive, in Definition 6 carries an index *i*, which is used to ensure the termination preserving, and the depth *k* of function calls, which increases by the call instruction and decreases by ret1 (unless k = 0).

Definition 5 (Simulation for Function).

$$(C_{\mathrm{as}}, \mathbf{f}) \preccurlyeq^{(\mathbf{p}, \mathbf{q})} A ::= \forall S, \mathbb{S}, w. (S, \mathbb{S}, A, w) \models \mathbf{p} \Longrightarrow$$
$$\exists i \in Index. \ \mathbf{q} \models (C_{\mathrm{as}}, S, \mathbf{f}, \mathbf{f} + 4) \preccurlyeq^{0}_{i} (A, \mathbb{S}),$$

where $q \models (C_{as}, S, pc, npc) \preccurlyeq_i^k (A, \mathbb{S})$ is defined in Definition 6.

Definition 6. $q \models (C_{as}, S, pc, npc) \preccurlyeq_i^k (A, \mathbb{S})$ holds if and only if the followings are true:

- 1) if $C_{as}(pc) \in \{i, jmp a, be f\}, then:$
- there exist S', pc', npc', such that $(C_{as}, S, pc, npc) :: \stackrel{\tau}{\Longrightarrow} (C_{as}, S', pc', npc');$
- for any S', pc', npc', if (C_{as}, S, pc, npc) :: → (C_{as}, S', pc', npc'), then one of the followings hold:

a)
$$\exists j < i. q \models (C_{as}, S', pc', npc') \preccurlyeq_{i}^{k} (A, \mathbb{S});$$

J. Comput. Sci. & Technol., Nov. 2020, Vol.35, No.6

- b) there exists $\mathbb{S}', j \in Index$, such that $(A, \mathbb{S}) \dashrightarrow (\bot, \mathbb{S}')$ and $\mathfrak{q} \models (C_{\mathrm{as}}, S', \mathfrak{pc}', \mathfrak{npc}') \preccurlyeq_{j}^{k} (\bot, \mathbb{S}')$ holds;
- 2) if $C_{as}(pc) = call f$, then:
- there exist S', pc', npc', such that $(C_{as}, S, pc, npc) :: \stackrel{\tau}{\Longrightarrow}^2 (C_{as}, S', pc', npc');$
- for any S', pc', npc', if (C_{as}, S, pc, npc) ::=^τ→² (C_{as}, S', pc', npc'), then one of the followings holds:
 a) ∃ j < i. q ⊨ (C_{as}, S', pc', npc') ≼^{k+1}_i (A,
 - a) ∃j < i. q ⊨ (C_{as}, S', pc', npc') ≼_j^{k+1} (A, S);
 b) there exist S', j ∈ Index, such that

 (A, S) --→ (⊥, S') and
 q ⊨ (C_{as}, S', pc', npc') ≼_j^{k+1} (⊥, S') holds;

3) if $C_{as}(pc) = retl$, then:

- there exist S', pc', npc', such that $(C_{as}, S, pc, npc) :: \stackrel{\tau}{\Longrightarrow}^2 (C_{as}, S', pc', npc');$
- for any S', pc', npc', if (C_{as}, S, pc, npc) :: →^τ (C_{as}, S', pc', npc'), then there exist j ∈ Index, S' and A', such that the followings hold:
 - a) either $j < i, \mathbb{S}' = \mathbb{S}$ and A' = A; or $(A, \mathbb{S}) \dashrightarrow (A', \mathbb{S}')$;
 - b) if k = 0, then there exists w', such that (where $f = S'.Q.R(r_{15})$):

$$(S', \mathbb{S}', A', w') \models q, A' = \bot,$$

$$pc' = f + 8, and npc' = f + 12;$$

$$else$$

$$q \models (C_{as}, S', pc', npc') \preccurlyeq_{j}^{k-1} (A', \mathbb{S}').$$

Then we give the semantics for well-formed primitive in Definition 7.

Definition 7 (Well-Defined Primitive Set Semantics).

$$\begin{split} \Psi &\models C_{\mathrm{as}} : \Omega ::= \forall \, \mathbf{f} \in \mathrm{dom}(\Omega), \, \iota. \, \exists \, \Upsilon, \, \overline{v}, \, \mathrm{fp}, \mathrm{fq}, \\ & \mathsf{wdSpec}(\mathrm{fp}, \mathrm{fq}, \Upsilon) \, \land \, (\mathrm{fp} \, \iota \Rightarrow (\!\! \uparrow \Upsilon(\overline{v})) \!\!) * \mathrm{true}) \\ & \land (C_{\mathrm{as}}, \mathbf{f}) \preccurlyeq^{(\mathrm{fp} \, \iota, \, \mathrm{fq} \, \iota)} \Upsilon(\overline{v}), \end{split}$$

where $\Omega(f) = \Upsilon, \Psi(f) = (fp, fq).$

It says that for any high-level abstract assembly primitive in the primitive set Ω , we can establish a simulation relation defined in Definition 5 between its low-level implementation in code heap $C_{\rm as}$ and itself.

Theorem 1 shows the soundness of our logic.

Soundness Proof. We show the soundness proof of our logic. We first give Lemma 3 to show that the simulation relation for functions can be decomposed into the simulation relations for individual instruction sequences.

Lemma 3. If $\Psi \models \{(\mathbf{p}, \mathbf{q})\}\ \mathbf{pc} : C_{\mathrm{as}}[\mathbf{pc}]\ and \models C_{\mathrm{as}} : \Psi, \ then \ (C_{\mathrm{as}}, \mathbf{f}) \preccurlyeq^{(\mathbf{p}, \mathbf{q})} A.$

Then, in Lemma 4, we show that our logic ensures the simulations.

- Lemma 4 (Logic Ensures Simulation).
- $\Psi \vdash \{(p,q)\} f : \mathbb{I} \implies \Psi \models \{(p,q)\} f : \mathbb{I};$
- $\vdash C_{\mathrm{as}} : \Psi \Longrightarrow \models C_{\mathrm{as}} : \Psi;$
- $\Psi \vdash C_{\mathrm{as}} : \Omega \Longrightarrow \Psi \models C_{\mathrm{as}} : \Omega$.

Next, we give Lemma 7, which says that the simulation for functions implies the primitive correctness. We define a whole program simulation in Definition 8 and divide the proof of Lemma 7 into two steps. First, we prove that the simulation for functions implies the whole program simulation in Lemma 5. Second, we prove that the whole program simulation implies the refinement relation between low- and high-level programs in Lemma 6.

Definition 8 (Whole Program Simulation). Whenever $P \leq^{i} \mathbb{P}$ holds, the followings are true:

- 1) if $P ::= \stackrel{\tau}{\Longrightarrow} P'$, then:
- $\exists j < i. P' \leq^{j} \mathbb{P}; or$
- $\exists j, \mathbb{P}' . \mathbb{P} :\Longrightarrow^{\tau} \mathbb{P}', and P' \leq j \mathbb{P}';$
- 2) if $P ::\stackrel{e}{\Longrightarrow} P'$, then $\exists j, \mathbb{P}' . \mathbb{P} :\stackrel{e}{\Longrightarrow} + \mathbb{P}'$, and $P \leq j \mathbb{P}$; 3) if $P ::\stackrel{\tau}{\Longrightarrow} abort$, then $\mathbb{P} :\stackrel{\tau}{\Longrightarrow} + abort$.

Lemma 5. If $\Psi \models C_{as} : \Omega, S \sim \mathbb{S}$, $\operatorname{ProgSafe}((C, \Omega), \mathbb{S})$ and $\mathbb{S}.\mathcal{K} = (_, \operatorname{pc}, \operatorname{npc})$, then there exists $i \in Index$, such that $(C \uplus C_{as}, S, \operatorname{pc}, \operatorname{npc}) \leq^i ((C, \Omega), \mathbb{S})$.

Lemma 6. $P \leq i \mathbb{P} \implies P \subseteq \mathbb{P}$.

Lemma 7 (Simulation Implies Primitive Correctness).

$$\Psi \models C_{\rm as} : \Omega \Longrightarrow C_{\rm as} \sqsubseteq \Omega.$$

Proof. We unfold $C_{as} \subseteq \Omega$ according to its definition in Definition 1, and finish the proof by applying Lemma 5 and Lemma 6.

Theorem 1 (Logic Soundness).

$$\Psi \vdash C_{\mathrm{as}} : \Omega \Longrightarrow C_{\mathrm{as}} \sqsubseteq \Omega.$$

Proof. The soundness proof can be done by applying Lemma 4 and Lemma 7. \Box

5 Verifying Context Switch Routine

We apply our program logic to verify that a context switch routine implemented in SPARCv8, which saves the current task's context and restores the new task's context, contextually refines the switch primitive defined in Subsection 4.1. Fig.20 shows the structure of the context switch routine that we proved. • SwitchEntry is the entry of the context switch routine. It saves the local and in registers of the current window into the current task's stack in memory, and calls **reg_save** to save the other registers into the current task's TCB.

Fig.20. Structure of context switch routine.

• Save_UsedWindows saves the register windows (except the current one) into the current task's stack in memory.

• Switch_NewContext restores the general registers from the new task's TCB (by calling reg_restore) and its stack in memory respectively. Then it sets the new task as the current one.

The main complexity of the verification lies in the code that manages the register window. To save all the used register windows, Save_UsedWindows repetitively restores the next window into general registers (as the current window) and then saves them into memory, until all the windows are saved.

Specification. Below we give the pre- and postconditions (a_{pre} and a_{post}) of the verified module respectively.

 $\begin{array}{l} a_{\mathrm{pre}}(\mathbf{t}_{c},\mathbf{t}_{n},\mathit{env},\mathit{nst},\mathcal{K}_{c},\mathcal{K}_{n}) & :::=\\ & \mathsf{Env}(\mathit{env})*(\mathsf{TaskNew} \Leftrightarrow (\mathbf{t}_{n},0))* \blacklozenge (10)*\\ & \mathsf{CurT}(\mathbf{t}_{c},_,\mathit{env},\mathcal{K}_{c})*\mathsf{RdyT}(\mathbf{t}_{n},\mathit{nst},\mathcal{K}_{n})*(\mathsf{switch}(\mathrm{nil}))\\ & a_{\mathrm{post}}(\mathbf{t}_{c},\mathbf{t}_{n},\mathit{env},\mathit{nst},\mathcal{K}_{c},\mathcal{K}_{n}) & ::=\\ & \exists \mathit{env}',\mathcal{K}'.\,\mathsf{Env}(\mathit{env}')*(\mathsf{TaskNew} \Leftrightarrow (\mathbf{t}_{n},0))*\\ & (\mathsf{CurT}(\mathbf{t}_{n},\mathit{nst},\mathit{env}',\mathcal{K}') \land \mathsf{p_env}(\mathit{env}') = \mathit{nst})*\\ & \mathsf{RdyT}(\mathbf{t}_{c},\mathtt{p_env}(\mathit{env}),\mathcal{K}_{c})*(\Downarrow \bot) \end{array}$

Each of them takes six arguments, the ID of the current task t_c , the ID of the new task t_n , the values *env* of general registers and other register windows saving contexts, the new task's context *nst* to be restored, the current task's thread local state \mathcal{K}_c and the new task's thread local state \mathcal{K}_n . In the specification, we use $\mathsf{Env}(env)$ to specify the values of general registers and the register windows. We describe the state of the current task using $\mathsf{CurT}(\mathsf{t}_c,_,env,\mathcal{K}_c)$. It describes the states of the current task t_c in the low and high level, and their state relation. Similarly, $\mathsf{RdyT}(\mathsf{t}_n, nst, \mathcal{K}_n)$ describes the states of the new task t_n in the low and high level, and their state relation. The memory location TaskNew records the identifier of the new task. TaskNew \Rightarrow (t_n, 0), where $l \Rightarrow v$ is defined as $(l \mapsto v) * (l \mapsto v)$, denotes that TaskNew saves (t_n, 0) in both the low- and high-level memory.

The precondition takes 10 tokens (\blacklozenge (10)). As we have explained, verifying call and jmp instructions will consume a token. Therefore, verifying calling functions reg_save and reg_restore will both consume a token. And Save_Usedwindows, which saves the context of each previous window into memory repetitively until the invalid one, will execute at most eight times, because the upper bound of the number of windows is eight. Therefore, 10 tokens are sufficient (two for reg_save and reg_restore, and eight for Save_Usedwindows).

If we compare a_{pre} and a_{post} , we can see that t_n becomes the current task ($\text{CurT}(t_n, nst, env', \mathcal{K}')$), and its general registers and stack, specified by Env(env'), are loaded from the saved context nst (i.e., $p_env(env') =$ nst). Here $p_env(env')$ refers to the part of the environment that we want to save or restore as context. Correspondingly, t_c becomes a non-current-thread, and part of its environment env at the entry of the context switch is saved, as specified by $\text{RdyT}(t_c, p_env(env), \mathcal{K}_c)$. The execution of switch should be done in the final state. We use \mathcal{K}' to represent the thread local state of t_n instead of \mathcal{K}_n in the final state, since the execution of switch will modify the program counters in \mathcal{K}_n .

Proof Outline. We show how to use our relational program logic defined in Fig.19 to verify the correctness of the context switch routine. We first instantiate the set of abstract assembly primitives (2) and the code heap specification (3) below:

$$\Omega ::= \{\texttt{SwitchEntry} \rightsquigarrow \texttt{switch}\}.$$
(2)

$$\Psi ::= \{ SwitchEntry \rightsquigarrow (a_{pre}, a_{post}), \\ reg_save \rightsquigarrow (fp_{rs}, fq_{rs}), \\ reg_restore \rightsquigarrow (fp_{rr}, fq_{rr}), \\ Save_Usedwindows \rightsquigarrow (fp_{su}, fq_{su}), \\ Switch_NewContext \rightsquigarrow (fp_{sn}, fq_{sn}) \}.$$
(3)

The set of abstract assembly primitives Ω contains only one abstract assembly primitive switch. The code heap specification Ψ contains the specifications of each code block. We use $(\mathbf{fp}_{rs}, \mathbf{fq}_{rs})$, $(\mathbf{fp}_{rr}, \mathbf{fq}_{rr})$, $(\mathbf{fp}_{su}, \mathbf{fq}_{su})$ and $(\mathbf{fp}_{sn}, \mathbf{fq}_{sn})$ to represent the specifications of reg_save, reg_restore, Save_Usedwindows and Switch_NewContext respectively. Since the postcondition in our logic specifies the state when the current function returns, the specification of SwitchEntry is $(a_{\rm pre}, a_{\rm post})$.

First, we prove that the specification of the context switch routine is well-defined in Lemma 8. J. Comput. Sci. & Technol., Nov. 2020, Vol.35, No.6

Lemma 8. wdSpec(a_{pre}, a_{post} , switch).

Proof. It is proved by the definition of wdSpec (in Definition 2). \Box

We use $C_{\texttt{switch}}$ to represent the code heap storing the code of the context switch routine, which includes the code blocks SwitchEntry, $\texttt{reg_save}$, $\texttt{reg_restore}$, $\texttt{Save_Usedwindows}$ and $\texttt{Switch_NewContext}$. We prove that $C_{\texttt{switch}}$ is well-defined in Lemma 9.

Lemma 9. $\vdash C_{\texttt{switch}} : \Psi$.

Proof. From the WfInt rule in Fig.19, we need to prove that for any ι_1 , ι_2 , ι_3 , ι_4 and ι_5 , the followings hold (\mathbf{f}_{se} , \mathbf{f}_{rs} , \mathbf{f}_{rr} , \mathbf{f}_{su} and \mathbf{f}_{sn} represent the starting labels of SwitchEntry, reg_save, reg_restore, Save_Usedwindows and Switch_NewContext respectively):

Ψ ⊢ {(a_{pre} ι₁, a_{post} ι₁)} f_{se} : C_{switch}[f_{se}];
Ψ ⊢ {(fp_{rs} ι₂, fq_{rs} ι₂)} f_{rs} : C_{switch}[f_{rs}];
Ψ ⊢ {(fp_{rr} ι₃, fq_{rr} ι₃)} f_{rr} : C_{switch}[f_{rr}];
Ψ ⊢ {(fp_{su} ι₄, fq_{su} ι₄)} f_{su} : C_{switch}[f_{su}];
Ψ ⊢ {(fp_{su} ι₅, fq_{su} ι₅)} f_{sn} : C_{switch}[f_{sn}].

The correctness proof of each code block can be done by applying the inference rules for instruction sequences shown in Fig.19. We can choose any place to apply the ABSCSQ rule to execute switch. Here, we apply the AB-SCSQ rule in verifying Switch_NewContext, when the context switch routine returns, as shown in Fig.21. In Fig.21, we use the solid circle to represent the point applying the ABSCSQ rule, and after the execution of switch, the state relation, defined in Subsection 4.3 and represented as the solid lines in Fig.21, between the low level and the high level can be reestablished.



Fig.21. Point doing refinement reasoning.

Theorem 2 (Context Switch Routine Correctness). $\Psi \vdash C_{\text{switch}} : \Omega.$

Proof. The proof follows the WfPrim rule in Fig.19, and Lemma 8 and Lemma 9. \Box

This part of work has not been mechanized in Coq. In our conference paper^[4], we show that we apply our logic for partial correctness to verify the main body of the context switch routine in a realistic embedded OS kernel for aerospace crafts, which consists of around 250 lines of SPARCv8 code, by 6 690 lines of Coq proof scripts. Here, the context switch routine verified by applying our relational program logic is a simplified version of such context switch routine, which omits some details like judging whether the current thread is a valid one. Verifying that each code block is well-defined using the inference rules in our new logic is not different from the previous work. The additional proof efforts include: 1) proving that the specification of context switch routine is well-defined (shown in Lemma 8); 2) applying the ABSCSQ rule to execute the switch primitive and proving that the state relation between lowand high-level programs can be reestablished when the context switch routine returns (as noted in the proof of Lemma 9).

6 Related Work

There has been much work on assembly or machine code verification. Most of them do not support function calls or simply treat function calls in the continuation-passing style where return addresses are viewed as first-class code pointers $^{[10-16]}$. SCAP $^{[5]}$ supports assembly code verification with various stackbased control abstractions, including function calls and returns. We follow the same idea here. However, SCAP gives a syntactic-based soundness proof by establishing the preservation of the syntactic judgment, which makes it difficult to interact with other modules verified in different logic. Since our goal is to verify inline assembly and link the verified code with the verified C programs, we give a direct-style semantic model of the logic judgments. And it allows us to extend our program logic in conference version^[4] to support verifying contextual refinement without many challenges. Also SCAP is based on a simplified subset of assembly instructions, while our work is focused on a realistically modeled subset of SPARCv8 instructions.

In terms of the support of realistic instruction sets, previous work on proof-carrying code (PCC) and typed assembly language (TAL) mostly supports subsets of x86. Myreen and Gordon's work^[17] presents a framework for ARM verification based on a realistic model (but it does not support function calls and returns).

As part of the Foundational Proof-Carrying Code (FPCC) project^[11], Tan and Appel presented a program logic \mathcal{L}_c for reasoning about control flow in assembly code^[16]. Although \mathcal{L}_c is implemented on top of the SPARC machine language, the underlying logic is a type system instead of a full-blown program logic for functional correctness. It reasons about functions in the continuation-passing style. Also handling SPARC features such as delayed writes or delayed control transfers is not the focus of \mathcal{L}_c . There has been work on mechanized semantics of the SPARCv8 ISA. Hou et al.^[18] modelled the SPARCv8 ISA in Isabelle/HOL, and tested their formal model against LENON3 simulation board, which is a synthesisable VHDL model of a 32-bit processor compliant with the SPARCv8 architecture, through more than 100000 instruction instances. Wang et al.^[7] formalized its semantics in Coq. Our operational semantics of SPARCv8 follows Wang et al.^[7] However neither Wang et al.^[7] nor we validate the formalization against actual hardware, and this remains as future work.

Ni *et al.*^[19] verified a context switch module of 19 lines in x86 code to showcase the support of embedded code pointers (ECP) in XCAP^[15]. We use our program logic to verify the contextual refinement between a context switch routine in SPARCv8 and switch primitive. The context switch routine implemented in SPARCv8 that we verified is more complicated than that implemented in x86, because of the requirement to save the contexts stored in register window in memory.

Yang and Hawblitzel^[20] verified Verve, an x86 implementation of an experimental operating system. Verve has two levels, the high-level TAL code and the low-level "Nucleus" that provides primitive access to hardware and memory. The Nucleus code is verified automatically using the Z3 SMT solver, while the goal of our work is to generate machine checkable proofs. Another key difference is the use of different ISAs. Here we give details to verify specific features of SPARCv8 programs.

There have been many techniques and tools proposed for automated program verification (e.g., [21,22]). It is possible to adapt them to verify SPARCv8 code. We propose a new program logic and do the verification in Coq mainly because the work is part of a big project for a fully certified OS kernel for aerospace crafts whose inline assembly is written in SPARCv8. We already have a program logic implemented in Coq for C programs, which allows us to verify C code with Coq proofs. Therefore we want to have a program logic for SPARCv8 so that it can be linked with the correctness proof of C and can generate machine-checkable Coq proofs too. That is, many of the automated verification techniques can be applied to reduce the manual efforts to write Coq proofs, which we would like to study in the future work.

7 Conclusions

We defined a relational program logic for SPARCv8. Our logic is based on a realistic semantics model and supports the main features of SPARCv8, including delayed control transfer, delayed writes, and register windows. It also supports modular reasoning of function calls in a direct-style and refinement verification. We applied our logic to verify that a context switch routine implemented in SPARCv8 contextually refines the switch primitive for task switching.

Our current work does not consider interrupts in the machine model. We would like to extend it for concurrency verification and finish S1 shown in Fig.2 in the future. S1 shown in Fig.2 says that the compilation ensures that the Pseudo-SPARCv8 code refines the C program with abstract assembly primitives.

References

- Xu F, Fu M, Feng X, Zhang X, Zhang H, Li Z. A practical verification framework for preemptive OS kernels. In *Proc.* the 28th International Conference, July 2016, pp.59-79.
- [2] Klein G, Elphinstone K, Heiser G, Andronick J, Cock D, Derrin P, Elkaduwe D, Engelhardt K, Kolanski R, Norrish M, Sewell T, Tuch H, Winwood S. seL4: Formal verification of an OS Kernel. In Proc. the 22nd ACM Symposium on Operating Systems Principles, Oct. 2009, pp.207-220.
- [3] Gu R, Koenig J, Ramananandro T, Shao Z, Wu X N, Weng S C, Zhang H, Guo Y. Deep specifications and certified abstraction layers. In Proc. the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Jan. 2015, pp.59-5608.
- [4] Zha J, Feng X, Qiao L. Modular verification of SPARCv8 code. In Proc. the 16th Asian Symposium on Programming Languages and Systems, December 2018, pp.245-263.
- [5] Feng X, Shao Z, Vaynberg A, Xiang S, Ni Z. Modular verification of assembly code with stack-based control abstractions. In Proc. the ACM SIGPLAN 2006 Conference on Programming Language Design and Implementation, June 2006, pp.401-414.
- [6] Leroy X, Blazy S. Formal verification of a C-like memory model and its uses for verifying program transformations. *Journal of Automated Reasoning*, 2008, 41(1): 1-31.
- [7] Wang J, Fu M, Qiao L, Feng X. Formalizing SPARCv8 instruction set architecture in Coq. In Proc. the 3rd International Symposium on Dependable Software Engineering: Theories, Tools, and Applications, Oct. 2017, pp.300-316.
- [8] Reynolds J. Separation logic: A logic for shared mutable data structures. In Proc. the 17th IEEE Symposium on Logic in Computer Science, July 2002, pp.55-74.

J. Comput. Sci. & Technol., Nov. 2020, Vol.35, No.6

- [9] Liang H, Feng X, Shao Z. Compositional verification of termination-preserving refinement of concurrent programs. In Proc. the Joint Meeting of the 23rd EACSL Annual Conference on Computer Science Logic and the 29th Annual ACM/IEEE Symposium on Logic in Computer Science, July 2014, Article No. 65.
- [10] Necula G C, Lee P. Safe kernel extensions without run-time checking. In Proc. the 2nd USENIX Symp. Operating System Design and Implementation, October 1996, pp.229-243.
- [11] Appel A W. Foundational proof-carrying code. In Proc. the 16th Annual IEEE Symposium on Logic in Computer Science, June 1998, pp.247-256.
- [12] Morrisett G, Crary K, Glew N, Grossman D, Samuels R, Smith F, Walker D, Weirich S, Zdancewic S. TALx86: A realistic typed assembly language. In Proc. the 1999 ACM SIGPLAN Workshop on Compiler Support for System Software, May 1996, pp.25-35.
- [13] Morrisett G, Walker D, Crary K, Glew N. From system F to typed assembly language. In Proc. the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Jan. 1998, pp.85-97.
- [14] Yu D, Nadeem A H, Shao Z. Building certified libraries for PCC: Dynamic storage allocation. *Science of Computer Programming*, 2004, 50(1/2/3): 101-127.
- [15] Ni Z, Shao Z. Certified assembly programming with embedded code pointers. In Proc. the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, January 2006, pp.320-333.
- [16] Tan G, Appel A W. A compositional logic for control flow. In Proc. the 7th International Conference on Verification, Model Checking, and Abstract Interpretation, Jan. 2006, pp.80-94.
- [17] Myreen M O, Gordon M J. Hoare logic for realistically modelled machine code. In Proc. the 13th International Conference on Tools and Algorithms for Construction and Analysis of Systems, March 2007, pp.568-582.
- [18] Hou Z, Sanán D, Tiu A, Liu Y, Hoa K C. An executable formalisation of the SPARCv8 instruction set architecture: A case study for the LEON3 processor. In *Proc. the 21st International Symposium on Formal Methods*, November 2016, pp.388-405.
- [19] Ni Z, Yu D, Shao Z. Using XCAP to certify realistic systems code: Machine context management. In Proc. the 20th International Conference on Theorem Proving in Higher Order Logics, Sept. 2007, pp.189-206.
- [20] Yang J, Hawblitzel C. Safe to the last instruction: Automated verification of a type-safe operating system. In Proc. the 2010 ACM SIGPLAN Conference on Programming Language Design and Implementation, June 2010, pp.99-110.
- [21] Berdine J, Calcagno C, O'Hearn P W. Symbolic execution with separation logic. In Proc. the 3rd Asian Symposium on Programming Languages and Systems, November 2005, pp.52-68.
- [22] Berdine J, Calcagno C, O'Hearn P W. Smallfoot: Modular automatic assertion checking with separation logic. In Proc. the 4th International Symposium on Formal Methods for Components and Objects, November 2005, pp.115-137.



Jun-Peng Zha is a Ph.D. candidate in the Department of Computer Science and Technology at Nanjing University, Nanjing. He received his M.S. degree in computer science from the University of Science and Technology of China, Hefei, in 2019, and his B.S. degree in software engineering from Shandong

University at Weihai in 2016. His research interests are in programming languages and formal methods, with a focus on the verification of compiler and low-level assembly code.



Xin-Yu Feng is a professor in Department of Computer Science and Technology at Nanjing University, Nanjing. He obtained his Ph.D. degree in computer science from Yale University, New Haven, in 2007, and his M.E. and B.S. degrees in computer science from Nanjing University, Nanjing, in

2002 and 1999, respectively. His research interests are on theories of programming languages and formal program verification.



Lei Qiao is a professor in the Center of On-Board Computer and Electronics at Beijing Institute of Control Engineering, Beijing. He received his Ph.D. degree in computer science from the University of Science and Technology of China, Hefei, in 2007. His research interests are in operating system design

and formal verification.