

**Known Errata in the Dissertation**  
*Ownership-Based Order Reduction and Simulation in*  
*Shared-Memory Concurrent Computer Systems*  
 by Christoph Baumann, Saarland University, 2014

Maintained by Christoph Baumann  
 (baumann@wjpservers.cs.uni-saarland.de)  
 May 9, 2014

1. reported by Geng Chen (Saarland University)

**Definition 14:** The second and fourth occurrence of  $C'$  in the definition of  $safe(C, \tau\alpha)$  should be  $C''$ .

$$safe(C, \tau\alpha) \stackrel{def}{=} safe(C, \tau) \wedge \exists C', C''. C \xrightarrow{\tau} C' \xrightarrow{\alpha} C'' \wedge safe_{step}(C', \alpha)$$

2. **Figure 18:** The stack frame is missing its *rds* component.
3. **Definition** of  $info_{IL}$ , p.128: The first parameter of  $info_{IL}.crso$  is redundant, also it should be called *crso*.

$$info_{IL}.crso : \mathbb{F}_{name} \times \mathbb{N} \times \mathbb{B}^5 \rightarrow \mathbb{N}_0$$

4. reported by Artem Alekhin (Saarland University)

**Definition 64:** In the fifth case of the definition of  $vol_f^{\pi, \theta}(e)$ , when expression  $e$  identifies the field of a **struct**-type variable, we should not give this field the same name as the function  $f$  in whose context we are evaluating  $e$ , thus we write  $(e').f'$  instead of  $(e').f$ .

$$vol_f^{\pi, \theta}(e) \equiv \begin{cases} \vdots \\ \mathbf{volatile} \in q \vee vol_f^{\pi, \theta}(e') & : \quad \dots \vee e = (e').f' \wedge \dots \\ \vdots \end{cases}$$

5. reported by Artem Alekhin (Saarland University)

**Definition 68:** Since the C-IL stack is indexed in ascending order from bottom to top, the top frame has index  $top$ , not 1.

$$\begin{aligned} consis_{C-IL}^{regs}(c_{IL}, \pi, \theta, info_{IL}, h) &\stackrel{def}{=} \begin{aligned} &(i) \quad h.gpr(bp) = bin_{32}(base(top)) \\ &(ii) \quad h.gpr(sp) = bin_{32}(base(top) - dist(top)) \end{aligned} \end{aligned}$$

6. reported by Artem Alekhin (Saarland University)

**Definition 70:** The definition of  $crsa_{i,j}$  is missing the term “.crso” after  $info_{IL}$  and it contains the typo “csrbase”. The correct definition should read:

$$crsa_{i,j} \equiv bin_{32}(crsbase_i - info_{IL}.crso(f_i, loc_i, r_{i,j}))$$

Also in the second case of the definition of the C-IL local variable consistency, the callee-save registers of a caller in frame  $i$  must be saved in the callee-save area of the next frame  $i + 1$ .

$$\mathcal{M}_{\mathcal{E}_i}(v_{i,j}) = \begin{cases} \vdots \\ h.m_4(csa_{i+1,j}) & : \quad r_{i,j} \in CS \wedge i < top \\ \vdots \end{cases}$$

7. reported by Geng Chen (Saarland University)

**Definition of  $stackovf$** , p.135: Since the C-IL stack is indexed in ascending order from bottom to top, we need index  $top$  instead of 1.

$$stackovf(c_{IL}, \pi, \theta, info_{IL}) \stackrel{def}{=} (base(top) - dist(top)) < msp_{IL}$$

8. **Definition** of the  $R_{S_{MIPS}^n, S_{C-IL}^n}$ , p.190: Since every C-IL unit operates its own stack, every sequential simulation needs a different compiler information parameter  $info_{IL}$  with an individual stack base address. However in the given generalized simulation theory only a single fixed simulation parameter is provided. One would need to extend the theory to provide for individual simulation parameters for each computation unit.

9. reported by Artem Alekhin (Saarland University)

**Theorems 2 and 3:** In the generalized sequential simulation theorem (Theorem 5), we demand that the resulting simulation configurations are in consistency points. This condition should also be present in the special cases for MASM and C-IL. We extend the claim of Theorem 2 by:

$$(iv) \quad h_{n+1}.c.pc \in A_{cp}^{MASM}$$

Note here that all MASM configurations are consistency points wrt.  $consis_{MASM}$ . The claim of Theorem 3 is extended as follows:

$$(v) \quad h_{n+1}.c.pc \in A_{cp}^{C-IL} \wedge cp(c'_{IL}, info_{IL})$$

10. **Proofs of Assumption 2 for MASM and C-IL:** The shared invariant guarantees memory consistency only for the shared and read-only memory. For unshared memory owned by other processes we cannot prove memory consistency from the given hypotheses of Assumption 2. Thus the assumption cannot be discharged for the current formulations of  $consis_{MASM}$  and  $consis_{C-IL}$  which demand consistency for all memory besides the stack and code regions.<sup>1</sup> In order to solve the issue, one needs to introduce a set of addresses  $BM \subseteq \mathbb{B}^{32}$  (bad memory) in the assembler and compiler information and exclude these addresses from the memory consistency. Naturally the code and stack region should be contained in  $BM$  and we redefine software condition  $badmemop$  to forbid all accesses to addresses in  $BM$ .

In the generalized simulation framework we require  $BM$  as a mandatory component of simulation parameters in  $\mathcal{P}$  and we demand the absence of bad memory operations explicitly next to all other software conditions represented by the  $sc$  predicate. In the concurrent simulation, for a process  $p$  we also put all addresses that are currently unshared and owned by other processes into  $p$ 's  $BM$  set. Thus we do not need to prove memory consistency for resources that are not visible for  $p$ . We reformulate the claim of Assumption 2 allowing for the choice of a new  $BM$  parameter:

$$\dots \implies \exists par'. sim_p(D'.M, par', E'.M)$$

The proof of the concurrent simulation theorem is adapted accordingly. Here one has to use ownership-safety together with the shared invariant to show that the abstract system does not access addresses in  $BM$ .

11. **Definition 7:** In the definition of  $(u', m')$ ,  $\delta_p$  should be just  $\delta$ .
12. **Definition 71:** Confusion between  $rds_i$  (value in the C-IL configuration) and  $rds(i)$  (memory content in the stack implementation). The second line should read:

$$(ii) \quad rds_i \neq \perp \implies rds(i) = \begin{cases} a & : \quad rds_i = \mathbf{val}(a, t) \in val_{\mathbf{ptr}} \\ alv & : \quad rds_i = \mathbf{lref}((v, o), j, t) \in val_{\mathbf{lref}} \end{cases}$$

---

<sup>1</sup>In fact in MASM global memory consistency always holds because macros do not access the global memory. Thus global memory is updated by complete consistency blocks containing only a single write instruction on the ISA and MASM level. For C-IL this is not the case and when a process executes an incomplete consistency block, its owned global memory may differ on the ISA and C-IL level.

13. reported by Artem Alekhin (Saarland University)

**Definition** of  $one\mathcal{IO}(\sigma, \tau)$ , p. 158: The case where  $\sigma|_{io} \neq \varepsilon$  and  $\tau|_{io} = \varepsilon$  is not excluded by the case split. We need the following statement:

$$one\mathcal{IO}(\sigma, \tau) \stackrel{def}{=} (\sigma|_{io} = \varepsilon \Leftrightarrow \tau|_{io} = \varepsilon) \wedge one\mathcal{IO}(\sigma) \wedge one\mathcal{IO}(\tau)$$

14. **Definition** of  $S_{C-IL}^n.\mathcal{IO}$ , p. 143: In an assignment of the form  $v = *(p)$  with variable  $v \in \mathbb{V}$  and volatile pointer expression  $p$ , i.e.,  $\tau_{Q_{c_{IL}}}^{\pi, \theta}(p) = (\{\mathbf{volatile}\}, \mathbf{ptr}(q, t))$  holds,  $p$  could contain references to more than one volatile pointer variable, e.g., by being the sum of two volatile pointer variables. We must forbid this case in order to ensure that C-IL statements do not reference more than one volatile variable. We could implement this restriction by extending the  $no2vol$  predicate, such that it recurses over an expression, searching for multiple accesses to volatile variables. At the moment  $no2vol_{c_{IL}}^{\pi, \theta}(p)$  only asserts that  $p$  is not a volatile pointer expression pointing to a volatile variable. A new definition is given below.

$$n2v_{c_{IL}}^{\pi, \theta}(e) \equiv \begin{cases} n2v_{c_{IL}}^{\pi, \theta}(e') & : e \in \{\ominus e', (t)e', \&(*e')\} \\ /vol_{c_{IL}}^{\pi, \theta}(e') & : e = *(e') \wedge \tau_{Q_{c_{IL}}}^{\pi, \theta}(e') = (q', \mathbf{ptr}(q, t)) \\ & \wedge \mathbf{volatile} \in q' \\ n2v_{c_{IL}}^{\pi, \theta}(e') & : e = *(e') \wedge \tau_{Q_{c_{IL}}}^{\pi, \theta}(e') = (q', \mathbf{ptr}(q, t)) \\ & \wedge \mathbf{volatile} \notin q' \\ n2v_{c_{IL}}^{\pi, \theta}(e') \wedge n2v_{c_{IL}}^{\pi, \theta}(e'') & : e = e' \oplus e'' \vee \\ \wedge / (vol_{c_{IL}}^{\pi, \theta}(e') \wedge vol_{c_{IL}}^{\pi, \theta}(e'')) & e = (e' ? e'' : \bar{e}) \wedge zero(\theta, \llbracket e' \rrbracket_c^{\pi, \theta}) \vee \\ & e = (e' ? \bar{e} : e'') \wedge /zero(\theta, \llbracket e' \rrbracket_c^{\pi, \theta}) \\ 1 & : \text{otherwise} \end{cases}$$

Then for assignments  $e = e'$  or  $e' = e$  to be  $\mathcal{IO}$  steps we would demand:

$$/vol_{c_{IL}}^{\pi, \theta}(e) \wedge vol_{c_{IL}}^{\pi, \theta}(e') \wedge n2v_{c_{IL}}^{\pi, \theta}(e')$$

An alternative, quick and dirty solution would be to simply forbid pointer arithmetics using volatile variables. To this end we only need to change the last line of the definition slightly, adding ' $e'' \in \mathbb{V}$ '.

$$\dots \wedge /vol_{c_{IL}}^{\pi, \theta}(e) \wedge vol_{c_{IL}}^{\pi, \theta}(e') \wedge (e' \in \mathbb{V} \vee (e' = *(e'')) \wedge e'' \in \mathbb{V} \wedge no2vol_{c_{IL}}^{\pi, \theta}(e''))$$