

CSimpl: A Rely-Guarantee-Based Framework for Verifying Concurrent Programs

David Sanán¹(✉), Yongwang Zhao^{1,2}, Zhe Hou¹, Fuyuan Zhang¹, Alwen Tiu¹,
and Yang Liu¹

¹ School of Computer Science and Engineering, Nanyang Technological University,
Singapore, Singapore

`sanan@ntu.edu.sg`

² School of Computer Science and Engineering, Beihang University, Beijing, China

Abstract. It is essential to deal with the interference of the environment between programs in concurrent program verification. This has led to the development of concurrent program reasoning techniques such as rely-guarantee. However, the source code of the programs to be verified often involves language features such as exceptions and procedures which are not supported by the existing mechanizations of those concurrent reasoning techniques. Schirmer et al. have solved a similar problem for sequential programs by developing a verification framework in the Isabelle/HOL theorem prover called Simpl, which provides a rich sequential language that can encode most of the features in real world programming languages. However Simpl only aims to verify sequential programs, and it does not support the specification nor the verification of concurrent programs. In this paper we introduce CSimpl, an extension of Simpl with concurrency-oriented language features and verification techniques. We prove the compositionality of the CSimpl semantics and we provide inference rules for the language constructors to reason about CSimpl programs using rely-guarantee, showing that the inference rules are sound w.r.t. the language semantics. Finally, we run a case study where we use CSimpl to specify and prove functional correctness of an abstract communication model of the XtratuM partitioning separation micro-kernel.

1 Introduction

In the past two decades, formal methods have been successfully applied in the verification of many critical systems. To improve confidence on the reliability of computer systems, verification of functional correctness and security properties is applied not only at the specification level [19], but also at the implementation [9] or even at the machine code level [5]. Verification of the implementation requires modelling languages that are able to capture the features in programming languages such as exceptions and procedure calls. Verification of sequential programs at implementation and machine code level has gained much attention both in academia and in industry [11], and now there is a reasonably strong tool support in this area [10, 15]. However, nowadays critical and high-assurance systems are often designed for multi-core architectures where multiple processes

run in parallel, but verification techniques and tools for concurrent programs are relatively less developed than those for sequential programs.

In order to tackle the verification of concurrent programs, first Owicki and Gries's work [14] introduced techniques for the verification of parallel programs. Later Jones [4] introduced the rely-guarantee method to improve Owicki and Gries's method by allowing compositional verification. The Owicki-Gries method has been mechanized in the Isabelle/HOL theorem prover in [13], and Jones's rely-guarantee method has been mechanized in Isabelle/HOL in [12] which follows the specification in [17]. Also, [1] models in Isabelle/HOL an algebraic specification of rely-guarantee. Although the languages used in previous mechanizations of the above mentioned methods are suitable for verifying system specifications, many implementations cannot be directly captured in those mechanizations. Therefore there is a need to develop a richer modelling language to accurately capture the behaviour of programs at the implementation level.

Simpl [15] is a while-language that supports most of the features of real world programming languages. The syntax and semantics of Simpl are modelled in Isabelle/HOL and Simpl has been used in the verification of seL4 source code [9]. However, its design aims only at reasoning about sequential programs, consequently, this language lacks constructors for parallel composition of programs. Moreover, its proof system is based on Hoare Logic, also for the verification of sequential languages, which cannot be used for reasoning about concurrent programs.

Building on the Simpl framework and the rely-guarantee method, we develop a formal verification framework in Isabelle/HOL, called CSimpl, for verifying partial correctness of high-assurance concurrent systems. The main contributions of this paper are as below:

(1) We extend Simpl using the notion of computation [12,17] to introduce parallelism in two layers: the bottom layer is the execution of sequential Simpl programs extended with a synchronization primitive `Await` over shared variables; and the top layer is the parallel execution of the bottom layer programs by means of a parallel composition operator. While existing rely-guarantee methods are mechanized for reasoning about abstract specification languages [12,13], our method goes one step further and covers most of the features of system programming languages such as exceptions, procedures, and pointers, among others.

(2) We define a compositional semantics of rely-guarantee for CSimpl. We also provide a set of inference rules for the rely-guarantee proof system and we prove that they are sound w.r.t. the semantics. The rich expressibility of CSimpl means that the number of inference rules of the rely-guarantee proof system is much higher than the work in [12] and their complexity is significantly increased. The CSimpl semantics, the rely-guarantee proof system specification and its soundness proof comprise more than 15k lines of proof and specification in Isabelle/HOL and Isar¹.

¹ Due to space reasons we only show some excerpts of the semantics and proofs, the whole model can be found at: <http://securify.scse.ntu.edu.sg/MicroVer/CSimpl>.

(3) As a case study, we specify in CSimpl two XtratuM [3] services for queuing inter-partition communication and we prove the correctness of an invariant on the queuing communication structure. Inter-partition communication is the mechanism used to implement information flow and is critical in proving event-based non-interference. XtratuM is a separation micro-kernel for space and time partitioning of applications. XtratuM supports multi-core architectures, being able to run several instances of the micro-kernel in parallel in multiple cores. Using our new rely-guarantee proof system, we prove that the specification of the inter-partition communication services correctly introduces and removes messages in the communication channel. The specification and the proofs comprise 3500 lines of formalization. To the best of our knowledge, this is the first attempt on the verification of separation micro-kernels targeting multi-core architectures. Other works such as [18–20] verify functional correctness and non-interference for sequential micro-kernels, and the work in [2] focuses on the verification of sequential applications using the ARINC standard.

2 CSimpl Language

2.1 Simpl Overview

Schirmer introduces in [15] a verification framework for imperative sequential programs developed in Isabelle/HOL. The verification framework includes a generic imperative language, called Simpl, which is composed of the necessary constructors to capture most of the features present in common sequential languages, such as conditional branching, loops, abrupt termination and exceptions, assertions, mutually recursive functions, expressions with side effects, and non-determinism. Additionally, Simpl can express memory related features like the memory heap, pointers, and pointers to functions. The Simpl verification framework also includes a Floyd/Hoare-like logic to reason about partial and total correctness, and on top of it, the framework implements a verification condition generator (VCG) to ease the verification process.

In order to capture all aspects of abrupt termination, assertions, and function calls, the program state `'s` in Simpl is modelled in Isabelle/HOL as a datatype `xstate` (shown in Fig. 1), which is composed of four different constructors: `Normal 's`, representing a regular execution; `Fault 'f`, representing a failed assertion; `Abrupt 's`, representing an exceptional state; and `Stuck`, representing a state where a call to a non-defined function is made. Additionally, the semantics requires an environment Γ containing procedure definitions, i.e., a partial function from the set `'p` of procedure names to the body of the procedures. Both features regarding the state and procedures definitions are used in CSimpl.

2.2 CSimpl Syntax

The syntax of CSimpl is shown in Fig. 1. CSimpl extends Simpl by adding two constructors for concurrency: `Await`, which takes two parameters `cond` and `body`,

```

type_synonym 's bexp = "'s set"
datatype ('s, 'p, 'f) com =
  Skip | Throw | Basic "'s  $\Rightarrow$  's" | Spec "('s  $\times$  's) set"
  | Seq "('s, 'p, 'f) com" "('s, 'p, 'f) com"
  | Cond "'s bexp" "('s, 'p, 'f) com" "('s, 'p, 'f) com"
  | While "'s bexp" "('s, 'p, 'f) com" | Call "'p"
  | DynCom "'s  $\Rightarrow$  ('s, 'p, 'f) com"
  | Guard "'f" "'s bexp" "('s, 'p, 'f) com"
  | Catch "('s, 'p, 'f) com" "('s, 'p, 'f) com"
  | Await "'s bexp" "('s, 'p, 'f) Simpl.com"
datatype ('s, 'f) xstate = Normal 's | Abrupt 's | Fault 'f | Stuck
type_synonym ('s, 'p, 'f) config = "('s, 'p, 'f)com  $\times$  ('s, 'f) xstate"
type_synonym ('s, 'p, 'f) body = "'p  $\Rightarrow$  ('s, 'p, 'f) com option"
type_synonym ('s, 'p, 'f) par_Simpl = "('s, 'p, 'f) com list"

```

Fig. 1. Syntax and state definition of the CSimpl language

and Parallel Composition. **Await** allows synchronization of processes under the boolean condition **cond** and then it atomically executes **body**, which is a pure sequential Simpl program. This allows us to use Hoare logic for sequential programs and the original Simpl VCG in the verification of the atomic blocks. Parallel composition happens at the top layer (root program), and it can not be nested with other constructors like in the approach followed in [7]. Therefore, a parallel program launches n sequential programs that are executed concurrently and that do not create new concurrent threads. A parallel program is defined as a list of sequential programs. Since we are aiming the verification of programs without dynamic creation of process, this approach is not a problem for our goal and simplify the mechanization.

CSimpl's syntax, following the syntax of Simpl, is defined in terms of states, of type 's; a set of fault types, of type 'f; and a set of procedure names of type 'p. The constructor **Skip** indicates program termination; **Seq** **s1** **s2**, **Cond** **b** **c1** **c2**, and **While** **b** **c** are respectively the standard constructors for sequential, conditional, and loop statements. **Throw** and **Throw** **c1** **c2** are the complements for abrupt termination of programs of **Skip** and **Seq** **c1** **c2**, and they allow to model exceptions. **Call** **p** invokes procedure **p**; **Guard** **f** **g** **c** represents assertions, where **c** is executed if the guard **g** holds in the current state, fault of type 'f is raised otherwise. Finally, **Spec** **r** and **DynCom** **cs** respectively introduce a nondeterministic behavior expressed by relation **r**, and a state dependent dynamic command transformation which is used to model blocks and functions with arguments.

2.3 CSimpl Semantics

The small-step operational semantics of CSimpl is a predicate inductively defined based on an environment for procedures I and a pair of component configurations $((P, s), (P', s'))$ where the program **P** in the state **s**, transits to the

program P' and the state \mathbf{s}' . It is represented as $\Gamma \vdash_c (P, \mathbf{s}) \rightarrow (P', \mathbf{s}')$, where c indicates it is a step transition in CSimpl. A CSimpl component configuration is defined as a tuple (P, \mathbf{s}) where P is a CSimpl program and \mathbf{s} is of type \mathbf{xstate} . A component configuration (\mathbf{p}, \mathbf{s}) is called **final** if $\mathbf{p} = \text{Skip}$ or $\mathbf{p} = \text{Throw}$ and there exists a state \mathbf{s}' such that $\mathbf{s} = \text{Normal } \mathbf{s}'$. A **final** configuration cannot progress to another configuration.

CSimpl extends Simpl with rules for synchronization on shared variables, **Await**, and the parallel computation shown below. For space reason we only provide the small-step semantics rules **Await** and **AwaitAb** for the **Await** command (Fig. 2). The rest are similar to those defined in [15].

$$\begin{array}{c}
 \frac{s \in b \quad \Gamma_{-a} \vdash \langle c, \text{Normal } s \rangle \Rightarrow t \quad \neg(\exists t'. t = \text{Abrupt } t')}{\Gamma \vdash_c (\text{Await } b \ \mathbf{p}, \text{Normal } s) \rightarrow (\text{Skip}, t)} \text{AWAIT} \quad \frac{}{\Gamma \vdash_c (P, \text{Normal } s) \rightarrow_e (P, t)} \text{ENV} \\
 \\
 \frac{s \in b \quad \Gamma_{-a} \vdash \langle c, \text{Normal } s \rangle \Rightarrow t \quad t = \text{Abrupt } t'}{\Gamma \vdash_c (\text{Await } b \ \mathbf{p}, \text{Normal } s) \rightarrow (\text{Throw}, \text{Normal } t')} \text{AWAITAB} \quad \frac{\forall s'. s \neq \text{Normal } s'}{\Gamma \vdash_c (P, s) \rightarrow_e (P, s)} \text{ENV_N} \\
 \\
 \frac{i < \text{length } Ps \quad \Gamma \vdash_c (Ps!i, s) \rightarrow (r, t)}{\Gamma \vdash_p (Ps, s) \rightarrow (Ps[i := r], t)} \text{PAR} \quad \frac{}{\Gamma \vdash_p (Ps, \text{Normal } s) \rightarrow_e (Ps, \text{Normal } t)} \text{P_ENV}
 \end{array}$$

Fig. 2. Small step and environment CSimpl semantic rules

The **Await** rules leverage Simpl's big step semantics to atomically transit from the initial configuration (\mathbf{p}, \mathbf{s}) to the next state \mathbf{t} resulting from the execution of \mathbf{p} from \mathbf{s} and it is expressed as $\Gamma \vdash \langle \mathbf{p}, \mathbf{s} \rangle \Rightarrow \mathbf{t}$. The two rules express the situation where from a current state \mathbf{s} satisfying the synchronization condition. The atomic program in Simpl ends in a state \mathbf{t} that can be an abrupt state as a result of an exception thrown in the sequential program for the rule **AwaitAb**, or any other possible state for the rule **Await**. This distinction is necessary since a Simpl program can finish in an **Abrupt** state, however the small-step semantics does not use the state **Abrupt**. Instead, a CSimpl program finishes in an exception state when the last configuration of a computation is a pair composed of the program **Throw**, together with a **Normal** state. Note that big step transitions use sequential Simpl programs, therefore the environment in the atomic step has to be a function from procedure names to Simpl programs, which do not contain **Await** instructions (for the same reason the body of **Await** cannot contain nested **Await** neither). Γ_{-a} translates bodies of procedures in Γ into Simpl programs if they do not contain any **Await** instruction, removing from Γ those procedures containing **Await** instructions.

A Parallel CSimpl configuration is defined as a tuple (Ps, \mathbf{s}) where Ps is a list of CSimpl programs and \mathbf{s} is of type \mathbf{xstate} . Parallel CSimpl semantics is inductively defined by means of rule **PAR** in Fig. 2. A parallel configuration (Ps, \mathbf{s}) transits to another parallel configuration $(Ps[i := r], \mathbf{s}')$ when there

is a program i in \mathbf{Ps} such that $\Gamma \vdash_c (Ps!i, s) \rightarrow (r, s')$. It is represented with $\Gamma \vdash_p (Ps, s) \rightarrow (Ps[i:=r], s')$. Similarly to component configurations, a parallel CSimpl configuration $(\mathbf{xs}, \mathbf{s})$ is called final if \mathbf{xs} is not empty and every component configuration $(\mathbf{xs}!i, s)$, with i smaller than the length of \mathbf{xs} , is final. $P\mathbf{s}!i$ means accessing the i element in the list $P\mathbf{s}$, whilst $P\mathbf{s}[i := r]$ means substitute the i element in $P\mathbf{s}$ for r .

Together with the semantic representing component transitions, it is necessary to define semantics for environment transitions. They are inductively defined using rules **Env** and **Env_n** in Fig. 2, where e is to express that it is an environment transition. CSimpl semantics for components can transit from a **Normal** state to a different type. However it is not possible to transit from a non **Normal** state to a different type of state, i.e. $\Gamma \vdash_p (P, \mathbf{Stuck}) \rightarrow (P', \mathbf{Normal\ t})$. Moreover, the component semantics always transits from a configuration (\mathbf{p}, \mathbf{s}) with $\mathbf{p} = \mathbf{Skip}$ and $\exists s'. \mathbf{s} = \mathbf{Normal\ s'}$ to a final transition $(\mathbf{Skip}, \mathbf{s})$. Therefore, the environment at the sequential layer needs to model this behaviour in the rules **Env** and **Env_n** in order to make the semantics at the parallel layer compositional. Environment transitions at the parallel level are defined in such a way that they can transit from a **Normal** state to another **Normal** state as shown in rule **P_ENV** in Fig. 2.

3 Rely-Guarantee for CSimpl

The rely-guarantee [7] method extends the specification of a program with two relations R and G characterizing, respectively, how the environment interferes with the program (Rely) and how the program modifies the environment (Guarantee). Therefore a specification for the verification of parallel systems using rely-guarantee is composed of four elements: precondition, postcondition, rely, and guarantee.

In order to take into account CSimpl state specification \mathbf{xstate} (which can take multiple forms to express different execution issues), the semantic for procedure calls, and the dual postcondition for normal or exception termination, the rely-guarantee specification and proof rules need to be modified accordingly. In the proof system itself, a total of 8 new rules have been added to the work in [12] to deal with all the language constructors present in CSimpl. Finally, soundness of the axiomatic rules for the proof system w.r.t. the rely-guarantee specification of validity is proven. The multiple forms of states makes the proof considerably more complex and larger than the work in [12]. While the work in [12] consists of around 2300 lines of proofs and specification, the current work consists of more than 13000 lines of proofs and specification.

3.1 Definition of Computation for CSimpl

The formal validity of a rely-guarantee tuple in this work is based on the definition of computation, which is the set of all possible sequences of configurations resulting of transiting the component or the environment, starting from an initial configuration.

Definition 1 (Sequential Component Computation). A computation is a tuple (Γ, \mathbf{confs}) where Γ is an environment for procedures and \mathbf{confs} is a list of sequential configurations. The set of possible computations \mathbf{cptn} is inductively defined as follows:

- $(\Gamma, [(P, s)]) \in \mathbf{cptn}$
- if $\Gamma \vdash_c (P, s) \rightarrow_e (P, t)$ and $(\Gamma, (P, t) \# xs) \in \mathbf{cptn}$ then $(\Gamma, (P, s) \# (P, t) \# xs) \in \mathbf{cptn}$
- if $\Gamma \vdash_c (P, s) \rightarrow (Q, t)$ and $(\Gamma, (Q, t) \# xs) \in \mathbf{cptn}$ then $(\Gamma, (P, s) \# (Q, t) \# xs) \in \mathbf{cptn}$

Definition 2 (cp Γ P s). The set of possible computations of an environment for procedures Γ starting from an initial configuration (P, s) is the set of tuples (Γ, l) such that $l!0 = (P, s)$ and $(\Gamma, l) \in \mathbf{cptn}$.

The set of parallel computations $\mathbf{par_cp}$ is defined similarly to \mathbf{cp} using parallel configurations and the semantic rules for parallel and environment step transitions.

Definition 3 (α). Conjoin [17] represented by α , defines an equivalence relation between a parallel computation p of n CSimpl components and a list clist of n component computations, where for all $i < n$. $\mathit{clist}!i = (\Gamma_i, \mathit{cptn}_i)$. $(\Gamma, p) \alpha \mathit{clist}$ iff:

- for all $i < n$, $\mathit{length} \mathit{cptn}_i = \mathit{length} p$ and $\Gamma_i = \Gamma$.
- for all $i < n$ and $k < \mathit{length} p$, $\mathit{cptn}_i!k = (c_i^k, s_i^k)$ and $p!k = (cs, s)$ with $cs!i = c_i^k$ and $s = s_i^k$.
- for all k such that $k + 1 < \mathit{length} p$, if $\Gamma \vdash_p p!k \rightarrow_e p!(k + 1)$, then for all $i < n$, $\Gamma_i \vdash \mathit{cptn}_i!k \rightarrow_e \mathit{cptn}_i!(k + 1)$; if $\Gamma \vdash_p p!k \rightarrow p!(k + 1)$ then there exists an $i < n$ where $\Gamma_i \vdash \mathit{cptn}_i!k \rightarrow (\mathit{cptn}_i)!(k + 1)$ and $\forall j. j \neq i \rightarrow \Gamma_j \vdash \mathit{cptn}_j!k \rightarrow_e \mathit{cptn}_j!(k + 1)$.

The last condition of conjoin states that for any step k in p , if k is an environment step in p , then k is also an environment step in all cptn_i ; and if it is a component step, then there is some cptn_i where k is a component step and for any other cptn_j , with $j \neq i$, k is an environment step.

Lemma 1 (Parallel computation as component computation).

$$xs \neq [] \implies \mathit{par_cp} \Gamma xs s = \{(\Gamma_1, c). \Gamma_1 = \Gamma \wedge (\exists \mathit{clist}. (\mathit{length} \mathit{clist} = (\mathit{length} xs) \wedge (\forall i < \mathit{length} \mathit{clist}. (\mathit{clist}!i) \in \mathit{cp}\Gamma(xs!i)s) \wedge (\Gamma, c) \alpha \mathit{clist})\}$$

Lemma 1 states that given a parallel configuration $(\mathbf{xs}, \mathbf{s})$ such that \mathbf{xs} is not empty ($[]$), then for any parallel computation (Γ, c) starting from $(\mathbf{xs}, \mathbf{s})$ there is a list of component computations clist with the same length of \mathbf{xs} and $(\Gamma, c) \alpha \mathit{clist}$. That is, the execution of a parallel number of components $xs_0 \dots xs_n$ can be expressed as the execution of one single component xs_i , with i smaller than n , where the execution of any other component xs_j is simulated by a component environment transition, with j smaller than n and different than i . The right and left implications of the equality in Lemma 1 are proven first by induction on the parallel computation and then by cases on the type of parallel and component events using conjoin.

3.2 Validity of Formulas for Rely-Guarantee in CSimpl

Based on the rely-guarantee definitions, we define the validity of a rely-guarantee tuple from the set of all possible computations from an initial configuration. This uses the notions of *assumption* of preconditions and the environment, and *commitment* of the component and the postcondition.

Definition 4 (assum(pre, rely)). *The assumption of a predicate pre and an environment relation rely for an environment of procedures Γ is the set of component computations (Γ, cptn) such that $\text{cptn}!0 = \text{Normal } s$ and $s \in \text{pre}$, and for any step transition in the computation $\Gamma \vdash_c \text{cptn}!k \rightarrow_e \text{cptn}!(k+1)$, where $k+1 < \text{length } \text{cptn}$, $\text{cptn}!k = (p_k, s_k)$, and $\text{cptn}!(k+1) = (p_{k+1}, s_{k+1})$ then $(s_k, s_{k+1}) \in \text{rely}$.*

The predicate *assum* represents the set of component computations (Γ, cptn) , such that the state component of the initial configuration of the computation is a `Normal` state satisfying *pre*. Also, in any transition of the environment $\Gamma \vdash_c \text{cptn}!k \rightarrow_e \text{cptn}!(k+1)$, the states of the configurations $\text{cptn}!k$ and $\text{cptn}!(k+1)$ belong to the rely relation.

To take advantage of automatic methods such as model checking, and following the original notion of validity for Hoare triples in `Simpl`, the commitment assumes that the last configuration in a computation does not end in a `Fault` state belonging to the set `F`, which is a set of non-reachable states previously calculated using external tools. Then the commitment is the set of computations such that component transitions belong to the *guarantee* relation, and that their last configuration are final (therefore with the program state equal to `Skip` or `Throw`) with the state component belonging to `q` or `a`.

Definition 5 (comm(guar, (q, a) F). *The commitment of a relation guar, a pair of predicates (q, a) , and a set of `Fault` states F , for an environment of procedures Γ , is the set of component computations (Γ, cptn) such that if $\text{cptn}!(\text{length } l - 1) = (l_p, l_s)$ and there is not a fault f such that $l_s = \text{Fault } f$ and $f \in F$, then (1) if for any component transition in the computation $\Gamma \vdash_c \text{cptn}!k \rightarrow \text{cptn}!(k+1)$ where $k < \text{length } \text{cptn}$, $\text{cptn}!k = (p_k, s_k)$, and $\text{cptn}!(k+1) = (p_{k+1}, s_{k+1})$ then $(s_k, s_{k+1}) \in \text{guar}$, and (2) if l is final then $l_s = \text{Normal } l'_s$ and if $l_p = \text{Skip}$ then $l'_s \in q$ and if $l_p = \text{Throw}$ then $l'_s \in a$.*

Definition 6 (com_validity). *Validity of a specification of a component P w.r.t. a precondition p , postcondition (q, a) , a rely relation R , a guarantee relation G , an environment of procedures Γ , and a set F of `Faults`, is represented as $\Gamma \models_{/F} P \text{ sat}[p, R, G, q, a]$ iff for all s , $\text{cp } \Gamma \text{ } P \text{ } s \cap \text{assum}(p, R) \subseteq \text{comm}(G(q, a)) \text{ } F$.*

Following [15], we use a set of procedure specifications Θ that are used during the procedure verification. The set of procedure specifications Θ , is a tuple which elements represent a procedure name and its specification in terms of precondition, rely and guarantee relations, and postcondition. Note that procedures in specifications belonging to Θ do not need to match the procedures defined in the environment Γ .

Definition 7 (com_cvalidity). *CValidity of a specification of a component P w.r.t. a precondition p , postcondition (q, a) , a rely relation R , a guarantee relation G , an environment of procedures Γ , a specification of procedures Θ , and a set F of Faults, represented as $\Gamma, \Theta \models_{/F} P \text{ sat}[p, R, G, q, a]$ iff for all tuples $(c, p', R', G', q', a') \in \Theta$ such that $\Gamma \models_{/F} (\text{Call } c) \text{ sat}[p', R', G', q', a']$ then $\Gamma \models_{/F} P \text{ sat}[p, R, G, q, a]$.*

Validity and CValidity for parallel computations are respectively represented by $\Gamma \models_{/F} P \text{ SAT}[p, R, G, q, a]$ and $\Gamma, \Theta \models_{/F} P \text{ SAT}[p, R, G, q, a]$. They are defined similarly to the ones for computation of components, using the definitions of computation, assumption, and commitment for parallel programs. We omit these definitions due to space reasons. Theorem 1 shows compositionality of validity of parallel rely-guarantee specifications.

Theorem 1 (validity_compositionality).

$$\begin{aligned}
 & \forall i < \text{length } xs. \Gamma, \Theta \models_{/F} C(xs!i) \text{ sat } [P(xs!i), R(xs!i), G(xs!i), Q(xs!i), A(xs!i)] \longrightarrow \\
 (1) & \forall i < \text{length } xs. R_p \cup (\bigcup j \in \{j. j < \text{length } xs \wedge j \neq i\}. G(xs!j)) \subseteq R(xs!i) \longrightarrow \\
 (2) & \bigcup j < \text{length } xs. G(xs!j) \subseteq G_p \longrightarrow (3) p \subseteq (\bigcap i < \text{length } xs. P(xs!i)) \longrightarrow \\
 (4) & (\bigcap i < \text{length } xs. (Q(xs!i))) \subseteq q \longrightarrow (5) (\bigcap i < \text{length } xs. (A(xs!i))) \subseteq a \longrightarrow \\
 & \Gamma, \Theta \models_{/F} \text{ParCom } xs \text{ SAT } [p, R_p, G_p, q, a].
 \end{aligned}$$

Therefore, to show that a parallel rely-guarantee specification is true, it is only necessary to prove rely-guarantee validity for each one of the single components and that (1) the rely of individual components is implied by the parallel rely and the union of the guarantee relations of the other individual components of the parallel system, (2) the union of the guarantee relations of the component specifications implies the guarantee relation for the parallel specification, (3) the precondition of the parallel specification is included in all the component specifications, (4) the intersection of all the normal postconditions of the component specifications is included in the normal postcondition of the parallel specification, (5) the union of the abrupt postcondition of all the component specifications is included in the abrupt postcondition of the parallel specification. Theorem 1 is proven using Lemma 1 and the definition of parallel validity of a rely-guarantee specification.

3.3 Inference Rules of the Proof System

The rely-guarantee proof system for CSimpl extends the previous mechanization of the logic in [12] with eight more inference rules. There are a total of fifteen rules, one for each language constructor, plus the consequence rule. Figure 3 shows those rules that are either new or substantially changed w.r.t. the work in [12]. Rules **Skip**, and **Throw** are added to handle program termination for normal and abrupt termination respectively. Since **Skip** deals with normal termination it requires the normal postcondition to be stable w.r.t. the rely relation, whilst in the case of **Throw** is the abrupt postcondition which has to be stable w.r.t. the rely relation. Similarly, **Catch** is the complement of the sequential rule for abrupt termination. In CSimpl, composition of programs can finish on an

$$\begin{array}{c}
\frac{Sta\ p\ R \quad Sta\ q\ R \quad p \subseteq \{s.f\ s \in q\} \\
\forall s\ t.s \in p \wedge (t = f\ s) \longrightarrow (Normal\ s, Normal\ s) \in G}{\Gamma, \Theta \vdash_F Basic\ f\ \mathbf{sat}\ [p, R, G, q, a]} \text{ BASIC} \\
\\
\frac{Sta\ p\ R \quad Sta\ q\ R \\
p \subseteq \{s.(\forall t.(s, t) \in r \longrightarrow t \in q) \wedge (\exists t.(s, t) \in r)\} \\
\forall s\ t.s \in p \wedge (s, t) \in r \longrightarrow (Normal\ s, Normal\ s) \in G}{\Gamma, \Theta \vdash_F Spec\ r\ \mathbf{sat}\ [p, R, G, q, a]} \text{ SPEC} \\
\\
\frac{Sta\ a\ R \\
\forall s.(Normal\ s, Normal\ s) \in G}{\Gamma, \Theta \vdash_F Throw\ \mathbf{sat}\ [a, R, G, q, a]} \text{ THROW} \qquad \frac{Sta\ q\ R \\
\forall s.(Normal\ s, Normal\ s) \in G}{\Gamma, \Theta \vdash_F Skip\ \mathbf{sat}\ [q, R, G, q, a]} \text{ SKIP} \\
\\
\frac{Sta\ p\ R \quad Sta\ q\ R \quad Sta\ a\ R \\
\forall V.\Gamma_{-a}, \{\} \vdash_F (p \cap b \cap \{V\})c \\
\{s.(Normal\ V, Normal\ s) \in G\} \cap q, \\
\{s.(Normal\ V, Normal\ s) \in G\} \cap a}{\Gamma, \Theta \vdash_F Await\ b\ c\ \mathbf{sat}\ [p, R, G, q, a]} \text{ AW} \qquad \frac{\Gamma, \Theta \vdash_F c\ \mathbf{sat}\ [p \cap g, R, G, q, a] \\
Sta\ (p \cap g)R \\
\forall s.(Normal\ s, Normal\ s) \in G}{\Gamma, \Theta \vdash_F Guard\ f\ g\ c\ \mathbf{sat}\ [p \cap g, R, G, q, a]} \text{ GD} \\
\\
\frac{\Gamma, \Theta \vdash_F c1\ \mathbf{sat}\ [p, R, G, q, r] \\
\Gamma, \Theta \vdash_F c2\ \mathbf{sat}\ [r, R, G, q, a] \\
Sta\ (p \cap g)R \quad Sta\ (a \cap g)R \\
\forall s.(Normal\ s, Normal\ s) \in G}{\Gamma, \Theta \vdash_F Catch\ c1\ c2\ \mathbf{sat}\ [p, R, G, q, a]} \text{ CATCH} \qquad \frac{\Gamma, \Theta \vdash_F c1\ \mathbf{sat}\ [p, R, G, r, a] \\
\Gamma, \Theta \vdash_F c2\ \mathbf{sat}\ [r, R, G, q, a] \\
Sta\ (p \cap g)R \quad Sta\ (a \cap g)R \\
\forall s.(Normal\ s, Normal\ s) \in G}{\Gamma, \Theta \vdash_F Seq\ c1\ c2\ \mathbf{sat}\ [p, R, G, q, a]} \text{ SEQ} \\
\\
\frac{\Gamma, \Theta \vdash_F c\ \mathbf{sat}\ [p \cap g, R, G, q, a] \\
Sta\ (p \cap g)R \quad f \in F \\
\forall s.(Normal\ s, Normal\ s) \in G}{\Gamma, \Theta \vdash_F Guard\ f\ g\ c\ \mathbf{sat}\ [p, R, G, q, a]} \text{ G} \qquad \frac{\Gamma, \Theta \vdash_F the(\Gamma\ c)\ \mathbf{sat}\ [p \cap g, R, G, q, a] \\
Sta\ (p \cap g)R \quad c \in dom\ \Gamma \\
\forall s.(Normal\ s, Normal\ s) \in G}{\Gamma, \Theta \vdash_F Call\ c\ \mathbf{sat}\ [p, R, G, q, a]} \text{ C} \\
\\
\frac{\forall s \in p.\Gamma, \Theta \vdash_F c\ s\ \mathbf{sat}\ [p \cap g, R, G, q, a] \\
Sta\ p\ R \quad \forall s.(Normal\ s, Normal\ s) \in G}{\Gamma, \Theta \vdash_F DynCom\ c\ \mathbf{sat}\ [p, R, G, q, a]} \text{ DYNCOM} \\
\\
\frac{\forall i < xs.R \cup (\bigcup j \in \{j. j < xs \wedge j \neq i\}.(Guar(xs!j))) \subseteq Rely(xs!i) \\
(\bigcup j < length\ xs.(Guard(xs!j))) \subseteq G \quad p \subseteq (\bigcap i < length\ xs.Pre(xs!i)) \\
(\bigcap j < length\ xs.(Post(xs!j))) \subseteq q \quad (\bigcup j < length\ xs.(Abr(xs!j))) \subseteq a \\
\forall i < xs.\Gamma, \Theta \vdash_F Com(xs!i)\ \mathbf{sat}\ [Pre(xs!i), Rely(xs!i), Guar(xs!i), Post(xs!i), Abr(xs!i)]}{\Gamma, \Theta \vdash_F xs\ \mathbf{SAT}\ [p, R, G, q, a]} \text{ COMP}
\end{array}$$

Fig. 3. Rely-guarantee proof rules for CSimpl

abrupt state without executing the second program. Hence it is necessary stability of the abrupt postcondition w.r.t. the *rely* relation. Similarly, the **Catch** rule requires stability of the normal postcondition with *rely*. We say that a predicate p is stable w.r.t. a relation R , $\text{Sta } p \ R$, if given two states s, s' , such that $p \ s$ is true and $(s, s') \in R$, then s' is also true in p .

The **Await** rule requires Hoare satisfiability of the sequential program representing its body, which is represented following traditional Hoare triplet notation $\{p\}c\{q\}$. In this case the postcondition is given as a pair to capture normal and abrupt termination. See [15] for more details on sequential **Simpl** program verification. Since the Hoare program can finish in either a normal state or an abrupt state, it is necessary that both postconditions are stable. The precondition should also be stable to remain unchanged under environment transitions. Since every component transition has to belong to the guarantee relation, we add this constraint into the Hoare triple, binding the initial states from the precondition to the final states of both the normal and abrupt postconditions.

The rest of rules can be deduced intuitively from their semantics adding stability of the precondition for non-terminal commands, e.g., **if** for branching and **call** for non-recursive procedure calls; and adding also stability of the normal postcondition for commands modifying the state.

Finally **COMP** is the rule for parallel composition. To apply compositionality, the rule is applied over a tuple xs composed of a sequential component Com and rely-guarantee specification, i.e. $Pre, Rely, Guarantee, Post$ for Com . The rule follows [12] taking abrupt termination into consideration, since this is an exception state, and not all the individual computations may be in an exception state. Therefore, whilst we require that the intersection of all component postconditions is included in the postcondition of the parallel program, for abrupt termination we only require that the union of abrupt postconditions is in the parallel program.

3.4 Soundness of the Proof System

We prove that the set of inference rules in the proof system is sound w.r.t. the definition of validity for parallel systems. The proof is carried out in two steps, first we prove that the inference rules for single components are sound.

Theorem 2 (comp_rgsound).

$$\Gamma, \Theta \vdash_{/F} c \text{ sat } [p, R, G, q, a] \longrightarrow \Gamma, \Theta \models_{/F} c \text{ sat } [p, R, G, q, a]$$

This is proved by induction on the inference rules. Axioms, i.e., those rules without assumptions on the proof system induction, are proven based on the notion of stability and the fact that any computation starting from them only has one component step. Therefore we prove that the stability rule preserves the precondition under any environment step. We then show that the component step preserves the commitment.

The semantics for computation makes it cumbersome to prove the soundness for those CSimpl constructors whose semantic is recursively defined, such as **Seq**,

Catch, and **While**. Soundness for these constructors are proven using a modular notion of computation [17] and the equivalence of both types of computation. The modular computation serializes the recursive specification of computation for the CSimpl constructors. This alternative semantics for computation unfolds the computation of CSimpl constructors. Soundness for these constructors is proven based on the different cases these rules provide. The modular computation for CSimpl extends the one provided in [12] with rules for the new language constructors, and new rules for **seq** and **while**, considering that the program in a final configurations can be **Skip** or **Throw**. The constructors **If** and **Call**, for non-recursive function calls, are proven similarly to the axioms based on the existence of a first component step for non-final configurations. After applying the component transition, we prove the correctness by the inductive step.

Recursive procedure calls require to consider the maximum number of nested function calls invoked by an execution and we do not currently provide a rule for them. **cptn** serializes the small step semantics removing scopes, which does not allow to prove soundness of recursive procedure calls. Nevertheless, it is possible to provide such a rule for recursive procedure calls, by extending the modular computation, with a parameter n representing the limit of nested procedures for which the computation is valid. Also, the semantics for validity must be extended to express that a formula is valid when it invokes at least n nested function calls by intersecting the assumptions in **com.validity** with the set of modular computations with limit n . Soundness of recursive procedure calls can be proven similarly to [15], by monotonicity of the extended computation in n and equality of the semantics.

Finally we show soundness of the proof system for the parallel composition of programs using Theorems 1 and 2.

Theorem 3 (par_rgsound).

$$\Gamma, \Theta \vdash_{/F} Ps \text{ SAT } [p, R, G, q, a] \longrightarrow \Gamma, \Theta \models_{/F} (\text{ParCom } Ps) \text{ SAT } [p, R, G, q, a]$$

4 Case Study

We apply the proof system for the specification and the verification of two XtratuM services for inter-partition communication. The XtratuM separation micro-kernel [3] provides spatial and temporal partitioning of applications. In a separation micro-kernel, different partitions are executed in separated memory domains, and the only allowed communication among partitions is by means of static dedicated channels explicitly defined between two or more partitions. XtratuM provides, among others services to communicate partitions through channels, partitions health-monitoring, and a static cyclic scheduler. In this case study we provide a very abstract CSimpl specification of the services to send and receive messages using queuing channels in a parallel architecture, where the XtratuM micro-kernel is executed in several cores of a multi-core processor. Using the rely-guarantee proof system introduced in Sect. 3 we prove: (1) that the services correctly introduce and remove elements in the queues associated with each communication channel and (2) that the number of elements in

the queues do not exceed the channel maximum capacity. The specification and proofs are comprised of more than 3500 lines of specification.

4.1 Queuing Inter-Partition Communication Description

Queuing inter-partition communication services allow partitions to escape from the isolated environment that XtratuM provides, allowing them to send and receive messages to/from other partitions using communication channels by means of dedicated ports assigned to partitions. A communication channel is an entity storing the communication data and the source and destination ports involved in the communication.

XtratuM implements two types of communication: sampling and queuing communication. While the former only allows to store one message, and it is multicasting, i.e., a channel has one source port and a list of destination ports. The latter allows to store many messages in a bounded buffer implemented as a queue, and only allows peer to peer communication, i.e., the channel has one source port and one destination port. Channels and ports are classified according to the type of communication. Therefore, a channel and a port can be of type sampling or queuing, and a sampling channel can only allocate sampling ports, and vice-verse. The services have as input a port to/from which the message is sent/received, and the message to be sent in the case of the sending service. Prior to modifying the queue, the services check whether the input values are consistent, e.g., the port which receives the message belongs to the partition, or it is a source or destination port depending on the invoked service.

4.2 State and Specification Definition

The state definition provides global and local variables. Global variables represent those variables shared by multiple instances of the micro-kernel, they hold the data for inter-partition communication, partitions, and the partition scheduler. Since we are targeting only queuing inter-partition communication, the components for the scheduler and partitions contain the necessary information for those services. The scheduler is highly abstracted and only contains information about the partition that is currently being executed, and therefore invoking the service; partitions only contain the list of assigned ports to the partition. The communication datatype includes the specification of channels and ports. A channel is defined as a datatype with the two possible types of channels, having as parameters the source and destination ports, and the message shared between the partitions, for which the queue is abstracted in the model as a multiset. The queuing channel also has a parameter indicating the maximum size of the channel buffer. Messages are modelled just as an abstract entity.

```
record com = ports :: "port_id  $\rightarrow$  port" channels :: "chan_id  $\rightarrow$  channel"
record vars = p_ ' :: "part_id  $\rightarrow$  partition"   c_ ' :: "com"
           s_ ' :: "scheduler list"   l_ ' :: "locals list"
```

In the model, \acute{l} , \acute{c} , \acute{s} , and \acute{p} access the locals, communication, scheduler, and partition component of the state, respectively. Local variables for each

process are a structure with the necessary variables for the input and output parameters of the services. One of the limitations of rely-guarantee is that the relations lose track of the sequence of executed operations. To solve this, verification of the concurrent increment of a variable, or adding/removing elements from a set like in this example, requires using additional variables to help tracking the changes [17]. In our model, we include a variable of type `Message option` that is initialized to `None`, and when a message is correctly sent or received the model assigns it to the variable. Our state abstracts and maps `XtratuM` global structures `xmcCommChannelTab`, and `xmcCommChannelPorts`, storing channel and port data, into the components of `´c` and `xmcPartition`, storing partition data, into `´p` respectively.

The parallel execution of services is modelled parametrically on the number of processes, which is defined as a fixed natural number within a Isabelle/HOL locale [8]. Each service is modelled as a procedure that is also parametrized by the process being executed; this allows that each specific procedure only accesses its local variables. The function F is generated by assigning a unique name for each service using a fold higher order function and assigning to each parametrized service name the corresponding parametrized body of the service. The parametrized event `receive` is shown below.

```

definition receive_q_message_i where "receive_q_message_i i ≡
(IF (¬ (ex_port_id ´c ((pt ((´l!)i)))))) ∨
  (¬ (port_q (the ((ports ´c) (pt (´l!)i)))))) ∨
  (¬ (port_dest (the ((ports ´c) (pt (´l!)i)))))) ∨
  (¬ port_in_part ´p ((´s!)i) (the ((ports ´c) (pt (´l!)i)))) THEN
  ´l ::= ´l[i:=((´l!)i)(ret_msg := None)])
ELSE AWAIT True
IFs port_empty (pt ((´l!)i)) ´c THEN
  ´l ::=s ´l[i:=((´l!)i)(ret_msg := None)])
ELSE
  ´l ::=s ´l[i:=((´l!)i) (ret_msg := port_get_msg (pt (´l!)i) ´c )];;s
  ´c ::=s port_rem_msg (pt (´l!)i) (the (ret_msg (´l!)i)) ´c ;;s
  ´l ::=s ´l[i:=((´l!)i) (aux_msg := (ret_msg (´l!)i) )];; FI FI"

```

The services abstract the low level behaviour of the `XtratuM` functions. They first check parameters validity, and then carries out the insertion/extraction of the message to/from the queue. Atomic blocks abstract `XtratuM`'s mutexes for mutual exclusion. Validation of correctness of the model w.r.t. the implementation is carried out at this stage by inspecting the code. For the `ReceiveQueuingPort` model, the event first checks that the accessed port exists in the current communication state, that it is a queuing and destination port, and that the partition that it belongs to the partition being executed. If any parameter is not valid then the service finishes returning `None`, otherwise it performs the operations over the channel queue after checking whether the queue is not empty for event `receive`, or not full for event `send`. The statements in the body of the `Await` statement are `IFs` and `::=s`. This is because the body of the `Await` is a sequential `Simpl` program and when embedded into a `CSimpl` program needs to modify the syntax. Event `SendQueuingPort` is modeled similarly.

4.3 Verification

For the parallel verification, we specify the rely and guarantee relations for the receive and send services. This relations are parameterized by a variable i , which refers to the i th process. We show the rely relation, the guarantee relation is similar to this, only differing in that local variables for any process j different than i will not be modified.

definition `Rely where "Rely B i`
 $\{(x,y). (\exists x1 y1. x=\text{Normal } x1 \wedge y=\text{Normal } y1 \wedge s_ ' x1 = s_ ' y1 \wedge$
 $(l_ ' x1)!i = (l_ ' y1)!i \wedge \text{ports } (c_ ' x1) = \text{ports } (c_ ' y1) \wedge$
 $p_ ' x1 = p_ ' y1 \wedge (x1 \in \text{Invariant } B \longrightarrow y1 \in \text{Invariant } B)) \}$ "

Since parallel programs do not change others programs' local variables, the i element in the list of local variables is not changed by the rely. Also, ports, partitions, and the scheduler are not changed by any service, therefore the rely relation does not change them. Finally, if the initial state of the relation preserves the invariant, it also preserves the shape of the channel's queues, then so does the final state of the relation.

The invariant establishes consistency of the port and channel structures that must be preserved by the services. Its most important specification is `channel_spec` which preserves the specification of the queue for every defined channel in the state.

definition `channel_spec where "channel_spec B`
 $\{\{ \forall c_id\ c. (\text{channels } \acute{c})\ c_id = \text{Some } c \longrightarrow$
 $\text{chan_get_msgs } c = (B\ c_id + \text{chan_sent_msgs } c_id\ \acute{c}\ \acute{l}) -$
 $\text{chan_rec_msgs } c_id\ \acute{c}\ \acute{l} \wedge$
 $(\text{size } (\text{chan_get_msgs } c) \leq \text{chan_get_max_bufs } c) \wedge$
 $\text{chan_rec_mes } c_id\ \acute{c}\ \acute{l} \sqsubseteq\# B\ c_id + \text{chan_sent_msgs } c_id\ \acute{c}\ \acute{l} \}$ "

`channel_spec` checks that the multiset modelling the queue for each defined c_id is equal to its initial value, which is given by $B\ c_id$; those messages correctly sent are pushed into the queue by the service; and that the received messages are popped out of the queue. `chan_sent\rec_msgs` gets for each c_id the multiset with the auxiliary variables different than `None`, meaning that the service has modified the queue for that channel. Also, for consistency of the multiset, the invariant needs to ensure that removed messages are a subset of the added messages.

Lemma 2 (Send_Rec_Correct).

$n > 0 \implies \Gamma, \{ \} \vdash / \{ \}$ (COBEGIN SCHEME $[0 \leq i < n]$
 $(\text{ex_service } i, \text{pre}_i\ B\ i, \text{Rely } B\ i, \text{Guar } B\ i, \text{Post_Arinc } B, \{ \{\text{True}\} \})$
 COEND) SAT $[\text{Pre_Arinc } B, \{ (x,y). x = y \}, \{ \{\text{True}\} \}, \text{Post_Arinc } B, \{ \{\text{True}\} \}]$

Lemma 2 proves the property on the parallel execution of the services. `ex_service` is a sequence of nested *ifs* controlling the call to the services, each *if* guarded by a local variable that indicates which service is invoked in each parallel process. In the parallel program, the identity relation indicates that the parallel environment does not change the state, being therefore a closed system, i.e., there is not any environment at the parallel level. The guarantee relation

is the universal set in which everything can be modified. The precondition `Pre_Arinc B` defines the invariant and auxiliary variables initialization to `None`. The precondition for each process `pre_i B i` sets the initial value for the auxiliary variable, the initial values of the channel queues, and it defines the invariant that is preserved by the postcondition for the normal termination `Post_Arinc B`. The abrupt postcondition is the universal set since we do not have any abrupt termination in this specification.

The proof obligations for the parallel rule are proven immediately after unfolding the definitions of the precondition, postcondition, and rely and guarantee relations. After applying the parallel rule on the parallel execution of the n components, it is necessary to prove that the parametrized `execute_service` satisfies the postcondition using the rely-guarantee rules for components. Once the `conditional` and `call` rules have been applied on `execute_service`, only the proof of the verification of each service body is left. Both send and receive services are similarly proven.

To prove the body of the services, it is necessary to apply the conditional rule to generate the proof obligations for the execution of two branches of the `if`. The first corresponds to the case in which the service is not invoked with the appropriate parameters and is immediately proven after apply the `Basic` rule since it does not modify any channel or auxiliary variable. For the second branch, after invoking `Await`, the sequential `Simpl` program representing its body is automatically unfolded using `Simpl`'s VCG. The resulting goal, now without any embedded `Simpl` specification, is solved by proving that the state after removing or inserting a message from/to the channel associated to the input port, and after assigning the removed/inserted message to the auxiliary variable, satisfies `channel_spec`. We use some auxiliary lemmas to prove it: first, that the modification of the auxiliary variable in a component does not modify the sets `chan_sent_msgs` and `chan_rec_msgs` for any channel other than the one associated to the port the service access; second, that the modification of a variable only modifies one of these sets. Using these auxiliary lemmas the postcondition is proven immediately by applying the properties over multisets.

5 Conclusions and Future Work

In this work we have presented `CSimpl`, a framework for specifying concurrent programs and verifying their partial correctness using rely-guarantee. This framework allows us to specify programs written in a large subset of the C language. Currently we are working also on axiomatic separation rules for the proof system following works on separation logic and rely-guarantee [6, 16]. This will help to cope with local variables and to hide global variables, thus improving scalability of the approach. There are, however, some aspects where this framework can be improved. First, we can introduce deadlock freedom and weak total correctness, which enable us to reason about termination of programs. Second, we can provide VCG tactics to achieve a higher level of automation. Currently, the language supports annotation to provide loop invariant, but the soundness

of annotated rules is yet to be proven. Third, it is also desirable to have completeness of the proof system to introduce properties proven at the language and semantics levels. The complexity of proving completeness make us to consider this as future work. Finally, the current proof system do not include a rule for recursive procedure calls, but our framework can be easily extended to support it, with minimal modifications on the rules already proven.

Acknowledgement. This research is supported (in part) by the National Research Foundation, Prime Ministers Office, Singapore under its National Cybersecurity R&D Program (Award No. NRF2014NCR-NCR001-30) and administered by the National Cybersecurity R & D Directorate.

References

1. Armstrong, A., Gomes, V.B.F., Struth, G.: Algebraic principles for rely-guarantee style concurrency verification tools. In: Jones, C., Pihlajasaari, P., Sun, J. (eds.) FM 2014. LNCS, vol. 8442, pp. 78–93. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-06410-9_6](https://doi.org/10.1007/978-3-319-06410-9_6)
2. de la Cámara, P., Mar Gallardo, M., Merino, P.: Model extraction for ARINC 653 based avionics software. In: Bošnački, D., Edelkamp, S. (eds.) SPIN 2007. LNCS, vol. 4595, pp. 243–262. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-73370-6_16](https://doi.org/10.1007/978-3-540-73370-6_16)
3. Carrascosa, E., Coronel, J., Masmano, M., Balbastre, P., Crespo, A.: XtratuM hypervisor redesign for LEON4 multicore processor. *SIGBED Rev.* **11**(2), 27–31 (2014)
4. Coleman, J.W., Jones, C.B.: A structural proof of the soundness of rely/guarantee rules. *J. Logic Comput.* **17**(4), 807–841 (2007)
5. Dam, M., Guanciale, R., Khakpour, N., Nemati, H., Schwarz, O.: Formal verification of information flow security for a simple arm-based separation kernel. In: Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security, CCS 1913, pp. 223–234. ACM, New York (2013)
6. Feng, X.: Local rely-guarantee reasoning. *SIGPLAN Not.* **44**(1), 315–327 (2009)
7. Jones, C.B.: Development methods for computer programs including a notion of interference. Ph.D. thesis. Oxford University, June 1981
8. Kammüller, F., Wenzel, M., Paulson, L.C.: Locales a sectioning concept for Isabelle. In: Bertot, Y., Dowek, G., Théry, L., Hirschowitz, A., Paulin, C. (eds.) TPHOLs 1999. LNCS, vol. 1690, pp. 149–165. Springer, Heidelberg (1999). doi:[10.1007/3-540-48256-3_11](https://doi.org/10.1007/3-540-48256-3_11)
9. 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: Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles (SOSP), pp. 207–220. ACM, New York (2009)
10. Myreen, M.O., Gordon, M.J.C., Slind, K.: Machine-code verification for multiple architectures: an application of decompilation into logic. In: Proceedings of the 2008 International Conference on Formal Methods in Computer-Aided Design, FMCAD 2008, pp. 20:1–20:8. IEEE Press, Piscataway (2008)

11. Newcombe, C., Rath, T., Zhang, F., Munteanu, B., Brooker, M., Deardeuff, M.: How Amazon web services uses formal methods. *Commun. ACM* **58**(4), 66–73 (2015)
12. Nieto, L.P.: The rely-guarantee method in Isabelle/HOL. In: Degano, P. (ed.) *ESOP 2003*. LNCS, vol. 2618, pp. 348–362. Springer, Heidelberg (2003). doi:[10.1007/3-540-36575-3_24](https://doi.org/10.1007/3-540-36575-3_24)
13. Nipkow, T., Nieto, L.P.: Owicki/Gries in Isabelle/HOL. In: Finance, J.-P. (ed.) *FASE 1999*. LNCS, vol. 1577, pp. 188–203. Springer, Heidelberg (1999). doi:[10.1007/978-3-540-49020-3_13](https://doi.org/10.1007/978-3-540-49020-3_13)
14. Owicki, S., Gries, D.: An axiomatic proof technique for parallel programs I. *Acta Informatica* **6**(4), 319–340 (1976)
15. Schirmer, N.: Verification of sequential imperative programs in Isabelle/HOL. Ph.D. thesis, Technischen Universität Munchen (2006)
16. Vafeiadis, V., Parkinson, M.: A marriage of rely/guarantee and separation logic. In: Caires, L., Vasconcelos, V.T. (eds.) *CONCUR 2007*. LNCS, vol. 4703, pp. 256–271. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-74407-8_18](https://doi.org/10.1007/978-3-540-74407-8_18)
17. Xu, Q., de Roever, W.P., He, J.: The rely-guarantee method for verifying shared variable concurrent programs. *Formal Aspects Comput.* **9**(2), 149–174 (1997)
18. Zhao, Y., Yang, Z., Sanán, D., Liu, Y.: Event-based formalization of safety-critical operating system standards: an experience report on ARINC 653 using event-B. In: *Proceedings of IEEE 26th International Symposium on Software Reliability Engineering (ISSRE)*, pp. 281–292 November 2015
19. Zhao, Y., Sanán, D., Zhang, F., Liu, Y.: Formal specification and analysis of partitioning operating systems by integrating ontology and refinement. *IEEE Trans. Industr. Inf.* **12**(4), 1321–1331 (2016)
20. Zhao, Y., Sanán, D., Zhang, F., Liu, Y.: Reasoning about information flow security of separation kernels with channel-based communication. In: Chechik, M., Raskin, J.-F. (eds.) *TACAS 2016*. LNCS, vol. 9636, pp. 791–810. Springer, Heidelberg (2016). doi:[10.1007/978-3-662-49674-9_50](https://doi.org/10.1007/978-3-662-49674-9_50)