Practical Abstractions for Automated Verification of Shared-Memory Concurrency

Conference Paper**Author(s):**

Oortwijn, Wytse; Gurov, Dilian; Huisman, Marieke

Publication date:

2020

Permanent link:

<https://doi.org/10.3929/ethz-b-000465582>

Rights / license:

In Copyright - Non-Commercial Use Permitted

Originally published in:

Lecture Notes in Computer Science 11990, https://doi.org/10.1007/978-3-030-39322-9_19

Practical Abstractions for Automated Verification of Shared-Memory Concurrency

Wytse Oortwijn^{1(✉)}, Dilian Gurov^{2(✉)}, and Marieke Huisman^{3(✉)}

¹ Department of Computer Science, ETH Zürich, Zürich, Switzerland
woortwijn@inf.ethz.ch

² KTH Royal Institute of Technology, Stockholm, Sweden
dilian@kth.se

³ University of Twente, Enschede, The Netherlands
m.huisman@utwente.nl

Abstract. Modern concurrent and distributed software is highly complex. Techniques to reason about the correct behaviour of such software are essential to ensure its reliability. To be able to reason about realistic programs, these techniques must be modular and compositional as well as practical by being supported by automated tools. However, many existing approaches for concurrency verification are theoretical and focus on expressivity and generality. This paper contributes a technique for verifying behavioural properties of concurrent and distributed programs that makes a trade-off between expressivity and usability. The key idea of the approach is that program behaviour is abstractly modelled using process algebra, and analysed separately. The main difficulty is presented by the typical abstraction gap between program implementations and their models. Our approach bridges this gap by providing a deductive technique for formally linking programs with their process-algebraic models. Our verification technique is modular and compositional, is proven sound with Coq, and has been implemented in the automated concurrency verifier VerCors. Moreover, our technique is demonstrated on multiple case studies, including the verification of a leader election protocol.

1 Introduction

Modern software is typically composed of multiple concurrent components that communicate via shared or distributed interfaces. The concurrent nature of the interactions between (sub)components makes such software highly complex as well as notoriously difficult to develop correctly. To ensure the reliability of modern software, verification techniques are much-needed to aid software developers to comprehend all possible concurrent system behaviours. To be able to reason about *realistic* programs, these techniques must be modular and compositional, but must also be practical by being supported by automated verifiers.

Even though verification of concurrent and distributed software is a very active research field [11, 13, 30, 41, 44, 50], most work is theoretical and focuses

primarily on expressivity and generality. This paper contributes a scalable and practical technique for verifying global behavioural properties of concurrent and distributed programs that makes a trade-off between expressivity and usability: rather than aiming for a unified approach to concurrency reasoning, we propose a powerful sound technique that is implemented in an automated verification tool, to reason automatically about realistic programs.

Reasoning about complex concurrent program behaviours is only practical if conducted at a suitable level of abstraction that hides irrelevant implementation details. This is because any real concurrent programming language with shared memory, threads and locks, has only very little algebraic behaviour. In contrast, *process algebra* offers an abstract, mathematically elegant way of expressing program behaviour. For this reason, many believe that process algebra provides a language for modelling and reasoning about the behaviour of concurrent programs at a suitable level of abstraction [1]. Our approach therefore uses process algebra as a language for *specifying* program behaviour. Such a specification can be seen as a model, the properties of which can additionally be checked (say, by model checking against temporal logic formulas). The main difficulty of this approach is dealing with the typical abstraction gap between program implementations and their models. The unique contribution of our approach is that it bridges this gap by providing a deductive technique for formally linking programs with their process-algebraic models. These formal links preserve *safety* properties; we leave the preservation of liveness properties for future work.

The key idea of the approach rests in the use of concurrent separation logic to reason not only about data races and memory safety, which is standard, but also about process-algebraic models (i.e., specified program behaviours), viewing the latter as *resources* that can be split and consumed. This results in a modular and compositional approach to establish that a program behaves as specified by its abstract model. Our approach is formally justified by correctness results that have mechanically been proven using Coq, including a machine-checked soundness proof of the proof system, stating that any verified program is a refinement of its abstract model. The verification technique has been implemented in the VerCors verifier for automated deductive verification of concurrent software [6]. Finally, the approach has been applied on various case studies [34], including a leader election protocol that is included in this paper.

We also recently successfully applied the techniques presented in this paper on an industrial case study, concerning the formal verification of a safety-critical traffic tunnel control system that is currently in use in Dutch traffic [36]. For this case study we made a process algebraic model of the control software that we analysed with mCRL2, and used the techniques presented in this paper to prove that this model is a sound abstraction of the program's behaviour.

An extended version of this paper is available as a technical report [46], which contains more details on the formalisation of the approach and the case study.

Contributions. This paper contributes a verification technique to reason about the behaviour of shared-memory concurrent programs that is modular, compositional, sound (proven with Coq), and implemented in an automated verifier.

First Sect. 2 illustrates the technique on a small Owicki–Gries example. Then Sect. 3 gives theoretical justification of the verification technique, as a concurrent separation logic with special constructs to handle process-algebraic models. Section 4 gives more details on the Coq embedding of the program logic and its soundness proof, and on its implementation in VerCors. Section 5 demonstrates the approach on a larger case study: the verification of a leader election protocol. Finally, Sect. 6 discusses related work and Sect. 7 concludes.

2 Approach

We first illustrate the approach on a simple example. In short, we abstractly specify concurrent program behaviour as process algebra terms. Process algebra terms are composed of atomic, indivisible *actions*. In our approach the actions are *logical descriptions of shared-memory modifications*: they describe what changes are allowed to a specified region of shared memory in the program. Actions are then *linked* to the concrete instructions in the program code that compute the memory updates. These links between the program and its abstract model are established deductively, using a concurrent separation logic that is presented later. Well-known techniques for process-algebraic reasoning can then be applied to guarantee safety properties over all possible state changes, as described by their compositions of actions. The novelty of the approach is that these safety properties can then be relied upon in the program logic due to the established formal connection between the program and its process-algebraic model.

Example Program. Consider the following program, which is a simple variant of the classical concurrent Owicki–Gries example [38].

$$\text{atomic } \{X := [E]; [E] := X + 4\} \parallel \text{atomic } \{Y := [E]; [E] := Y * 4\}$$

This program consists of two concurrent threads: one that atomically increments the value at heap location E by four, while the other atomically multiplies the value at E by four. The notation $[E]$ denotes *heap dereferencing*, where E is an expression whose evaluation determines the heap location to dereference.

The challenge is to modularly deduce the classical Owicki–Gries postcondition: after termination of both threads, the value at heap location E is either $4 * (\text{old}_E + 4)$ or $(4 * \text{old}_E) + 4$ (depending on the interleaving of threads), where old_E is the old value at E —the value of E at the pre-state of the computation.

Well-known classical approaches to deal with such concurrent programs [42] include auxiliary state [38] and interference abstraction via rely-guarantee reasoning [20]. Modern program logics employ more intricate constructs, like atomic Hoare triples [41] in the context of TaDa, or higher-order ghost state [23] in the

context of Iris. However, the mentioned classical approaches typically do not scale well, whereas such modern, theoretical approaches are hard to integrate into (semi-)automated verifiers like VeriFast or VerCors.

In contrast, our approach is a balanced trade-off between expressivity and usability: it is scalable as well as implemented in an automated deductive verifier. The approach consists of the following three steps: (1) defining a process-algebraic model $OG = \text{incr}(4) \parallel \text{mult}(4)$ that is composed out of two actions, incr and mult , that abstract the atomic sub-programs; (2) verifying the Owicki–Gries postcondition algorithmically on the OG process; and (3) deductively verifying that OG is a correct behavioural specification of the program’s execution flow (i.e., verifying that all atomic state changes in the program have a corresponding action in OG). The following paragraphs give more detail on these three steps.

Step 1: Specifying Program Behaviour. The first step is to construct a behavioural specification OG of the example program. This process is defined as the parallel composition of the actions $\text{incr}(4)$ and $\text{mult}(4)$, which specify the behaviour of the atomic increment and multiplication in the program, respectively. In our approach, program behaviour is specified logically, by associating a *contract* to every action. For our example program, incr and mult have the following contract:

requires true; ensures $x = \backslash\text{old}(x) + n$; action $\text{incr}(\text{int } n)$;	requires true; ensures $x = \backslash\text{old}(x) * n$; action $\text{mult}(\text{int } n)$;
---	---

The variable x is a free, *process-algebraic variable* that is later linked to a concrete heap location in the program (namely E). Moreover, the increment and multiplication of 4 has now been generalised to an arbitrary integer n .

These two actions may be composed into a full behavioural specification of the example program, by also assigning a top-level contract to OG :

requires true;
ensures $x = (\backslash\text{old}(x) + n) * n \vee x = (\backslash\text{old}(x) * n) + n$;
process $OG(\text{int } n) := \text{incr}(n) \parallel \text{mult}(n)$;

Step 2: Process-Algebraic Reasoning. The next step is to verify that OG satisfies its contract, which can be reduced to standard process-algebraic analysis. We say that OG satisfies its contract if all finite, action contract-complying traces of OG satisfy the ensures clause. The standard approach to analyse OG is to first linearise it to the bisimilar process term $\text{incr}(n) \cdot \text{mult}(n) + \text{mult}(n) \cdot \text{incr}(n)$, and then to prove its correctness by analysing all branches. VerCors currently does the analysis by encoding the linearised process as input to the Viper verifier [29]. VerCors can indeed automatically establish that OG satisfies its postcondition.

Step 3: Deductively Linking Processes to Programs. The key idea of our approach is that, by analysing how contract-complying action sequences change the values of process-algebraic variables, we may indirectly reason about how the content

```

1   $old_E := [E];$ 
2   $M := \text{process } OG(4) \text{ over } \{x \mapsto E\};$ 
3  atomic {
4     $X := [E];$ 
5    action  $\text{incr}(4)$  do {
6       $[E] := X + 4;$ 
7    }
8  }
9  atomic {
10    $Y := [E];$ 
11   action  $\text{mult}(4)$  do {
12      $[E] := Y * 4;$ 
13   }
14 }
15 finish  $M;$ 
16 assert  $E \xrightarrow{1} (old_E + 4) * 4 \vee E \xrightarrow{1} (old_E * 4) + 4;$ 
    
```

Fig. 1. The annotated Owicki–Gries example (the annotations are coloured blue). (Color figure online)

at heap location E evolves over time. So the final step is to project this process-algebraic reasoning onto program behaviour, by annotating the program.

Figure 1 shows the required program annotations. First, x is connected to E by initialising a new model M on line 2 that executes according to $OG(4)$. The actions incr and mult are then linked to the corresponding subprograms on lines 5–7 and 11–13 by identifying *action blocks* in the code, using special program annotations. We use these **action** annotations to verify in a thread-modular way that the left thread performs the $\text{incr}(4)$ action (on lines 5–7) and that the right thread performs $\text{mult}(4)$ (lines 11–13). As a result, when the program reaches the **finish** annotation on line 15 all the actions of OG will have been performed. This indirectly means that the content at heap location E has evolved as described by OG , thus allowing the asserted postcondition on line 16 to be derived.

3 Formalisation

This section gives theoretical justification of the verification approach and explains the underlying logical machinery. First, Sects. 3.1 and 3.2 briefly discuss the syntax and semantics of process algebraic models and programs, respectively. Then Sect. 3.3 presents the program logic as a concurrent separation logic with assertions that allow to specify program behaviour as a process algebraic model. Section 3.4 discusses the proof rules. Finally, Sect. 3.6 discusses soundness of the approach. All these components have been fully formalised in Coq.

Due to space constraints, the technical presentation assumes a certain familiarity with process algebra and separation logic. For more details we refer to the accompanying technical report or to the Coq formalisation [46].

3.1 Process Algebraic Models

Process algebraic models are defined by the language *Proc* as follows, with $a, b, \dots \in Act$ the domain of *actions*, $x, y, z, \dots \in ProcVar$ the domain of *process algebraic variables*, and $m, n, \dots \in Lit$ the domain of *literals*.

Successful termination

$$\varepsilon \downarrow \quad \frac{P \downarrow \quad Q \downarrow}{P \cdot Q \downarrow} \quad \frac{P \downarrow}{P + Q \downarrow} \quad \frac{Q \downarrow}{P + Q \downarrow} \quad \frac{P \downarrow \quad Q \downarrow}{P \parallel Q \downarrow} \quad P^* \downarrow$$

Small-step reduction rules (excerpt)

$$\begin{array}{c} \text{PSTEP-SEQ-L} \\ \frac{(P, \sigma) \xrightarrow{a} (P', \sigma')}{(P \cdot Q, \sigma) \xrightarrow{a} (P' \cdot Q, \sigma')} \end{array} \quad \begin{array}{c} \text{PSTEP-SEQ-R} \\ \frac{P \downarrow \quad (Q, \sigma) \xrightarrow{a} (Q', \sigma')}{(P \cdot Q, \sigma) \xrightarrow{a} (Q', \sigma')} \end{array} \quad \begin{array}{c} \text{PSTEP-ACT} \\ \frac{\llbracket \text{pre}(a) \rrbracket(\sigma) \quad \llbracket \text{post}(a) \rrbracket(\sigma')}{(a, \sigma) \xrightarrow{a} (\varepsilon, \sigma')} \end{array}$$

Fig. 2. The small-step operational semantics of process algebraic models.**Definition 1 (Process expressions, Process conditions, Processes).**

$$\begin{aligned} e \in \text{ProcExpr} &::= m \mid x \mid e + e \mid e - e \mid \dots \\ b \in \text{ProcCond} &::= \text{true} \mid \text{false} \mid \neg b \mid b \wedge b \mid e = e \mid e < e \mid \dots \\ P, Q \in \text{Proc} &::= \varepsilon \mid \delta \mid a \mid P \cdot Q \mid P + Q \mid P \parallel Q \mid P^* \end{aligned}$$

As usual, ε is the empty process that has no behaviour, whereas δ is the deadlocked process that neither progresses nor terminates. The process $P \cdot Q$ is the sequential composition of P and Q , while $P + Q$ denotes their non-deterministic choice. The process $P \parallel Q$ is the parallel composition P and Q . Finally, P^* is the Kleene iteration of P and denotes a sequence of zero or more P 's.

The verification approach uses process algebraic models in the presence of data, implemented via *action contracts*. These action contracts make the process algebra language non-standard. Action contracts consist of pre- and postconditions that logically describe the state changes imposed by the action. Each action is assumed to have an associated contract that can be obtained via the functions $\text{pre}, \text{post} : \text{Act} \rightarrow \text{ProcCond}$. All pre- and postconditions are of type ProcCond , which is the domain of Boolean expressions over process algebraic variables.

Semantics. The operational semantics of processes is expressed as a binary reduction relation $\cdot \xrightarrow{\cdot} \cdot \subseteq \text{ProcConf} \times \text{Act} \times \text{ProcConf}$ over *process configurations* $\text{ProcConf} \triangleq \text{Proc} \times \text{ProcStore}$, labelled with actions from Act . The notion of data is implemented via *process stores* $\sigma \in \text{ProcStore} \triangleq \text{ProcVar} \rightarrow \text{Val}$ that map process algebraic variables to a semantic domain Val of *values*.

Most of the reduction rules are standard. Figure 2 gives an overview of the non-standard rules. All other transition rules are deferred to [46].

To define the transition rule PSTEP-SEQ-R for sequential composition, it is common in process algebra with ε to use an explicit notion of *successful termination* [4]. Successful termination $P \downarrow$ of any process P intuitively means that P has the choice to have no further behaviour and thus to behave as ε . Furthermore, the PSTEP-ACT transition rule for action handling permits state to change in any way that complies with the corresponding action contract.

$$\begin{array}{c}
 \text{STEP-PROC-INIT} \\
 (X := \text{process } P \text{ over } \Pi, h, s) \rightsquigarrow (\text{skip}, h, s) \qquad \text{STEP-PROC-FINISH} \\
 (\text{finish } X, h, s) \rightsquigarrow (\text{skip}, h, s) \\
 \\
 \text{STEP-ACT} \\
 \frac{(C, h, s) \rightsquigarrow (C', h', s')}{(\text{action } X.a \text{ do } C, h, s) \rightsquigarrow (\text{action } X.a \text{ do } C', h', s')} \\
 \\
 \text{STEP-ACT-FINISH} \\
 (\text{action } X.a \text{ do skip}, h, s) \rightsquigarrow (\text{skip}, h, s)
 \end{array}$$

Fig. 3. An excerpt of the small-step operational semantics of programs.

The program logic allows one to handle process algebraic models *up to bisimulation*. We write $P \cong Q$ to denote that P and Q are *bisimilar* (i.e., behaviourally equivalent). Bisimilarity is a congruence with respect to all process algebraic connectives. Moreover, we indeed have that $P \downarrow$ implies $P \cong P + \varepsilon$ for any P .

3.2 Programs

Our approach is formalised on the following simple concurrent pointer language, where $X, Y, Z, \dots \in \text{Var}$ are *(program) variables*.

Definition 2 (Expressions, Conditions, Programs).

$$\begin{aligned}
 E \in \text{Expr} &::= n \mid X \mid E + E \mid E - E \mid \dots \\
 B \in \text{Cond} &::= \text{true} \mid \text{false} \mid \neg B \mid B \wedge B \mid E = E \mid E < E \mid \dots \\
 \Pi \in \text{AbstrBinder} &::= \{x_0 \mapsto E_0, \dots, x_n \mapsto E_n\} \\
 C \in \text{Cmd} &::= \text{skip} \mid X := E \mid X := [E] \mid [E] := E \mid C; C \mid C \parallel C \\
 &\mid X := \text{alloc } E \mid \text{dispose } E \mid \text{atomic } C \\
 &\mid \text{if } B \text{ then } C \text{ else } C \mid \text{while } B \text{ do } C \\
 &\mid X := \text{process } P \text{ over } \Pi \mid \text{action } X.a \text{ do } C \mid \text{finish } X
 \end{aligned}$$

This language is a variation of the language proposed by [9, 32], extended with *specification-only* commands (displayed in blue) for handling process algebraic models in the logic. These commands are ignored during program execution.

Specification-wise, $X := \text{process } P \text{ over } \Pi$ initialises a new process algebraic model that is represented by the process term P , with Π a finite mapping from process algebraic variables to heap locations. Π is used to connect abstract state (i.e., the state of process algebraic models) to concrete program state (i.e., heap entries) and is therefore referred to as an *abstraction binder*.

The $\text{finish } X$ command *concludes* the model that is identified by X , given that the associated process successfully terminates. By *concludes* we mean that the model's postcondition can be relied upon and used in the proof system.

Finally, $\text{action } X.a \text{ do } C$ executes the command C in the context of the abstract model X as the action a . In particular, this specification command

states that, by executing C (according to the operational semantics of programs), the action a is executed in the specified process algebraic model.

Semantics. The operational semantics of programs is expressed as a small-step reduction relation $\cdot \rightsquigarrow \cdot \subseteq \text{Conf} \times \text{Conf}$, between (*program*) *configurations* $\text{Conf} \triangleq \text{Cmd} \times \text{Heap} \times \text{Store}$. Program configurations $(C, h, s) \in \text{Conf}$ consist of a program C , as well as a heap $h \in \text{Heap} \triangleq \text{Val} \rightarrow_{\text{fin}} \text{Val}$ that models shared memory and a store $s \in \text{Store} \triangleq \text{Var} \rightarrow \text{Val}$ that models thread-local memory.

Figure 3 shows an excerpt of the new reduction rules for ghost commands. All other reduction rules are standard in spirit and are deferred to [46].

Most importantly, all ghost commands are specification constructs: they do not affect the program state and are essentially handled as if they were comments. However, observe that STEP-PROC-FINISH and STEP-ACT-FINISH are auxiliary transition steps that reduce a finished process or action to **skip**. These are not strictly needed, but make it more convenient to prove soundness of the logic.

3.3 Program Logic

Our program logic builds on *intuitionistic*¹ concurrent separation logic (CSL), where the assertion language is defined by the following grammar.

Definition 3 (Assertions).

$$\begin{aligned} t \in \text{PointsToType} &::= \text{std} \mid \text{proc} \mid \text{act} \\ \mathcal{P}, \mathcal{Q}, \mathcal{R}, \dots \in \text{Assn} &::= B \mid \forall X. \mathcal{P} \mid \exists X. \mathcal{P} \mid \mathcal{P} \vee \mathcal{Q} \mid \mathcal{P} * \mathcal{Q} \mid \mathcal{P} \multimap \mathcal{Q} \\ &\mid *_{i \in I} \mathcal{P}_i \mid E \xrightarrow{\pi}_t E \mid \text{Proc}_{\pi}(X, b, P, \Pi) \end{aligned}$$

The assertion $\mathcal{P} * \mathcal{Q}$ is the separating conjunction of separation logic and states that $\mathcal{P}$ and $\mathcal{Q}$ hold on *disjoint* parts of the heap. This for example means that $\mathcal{P}$ and $\mathcal{Q}$ cannot both express write access to the same heap entry. The assertion $*_{i \in I} \mathcal{P}_i$ is the *iteration of $*$* and is equivalent to $\mathcal{P}_0 * \dots * \mathcal{P}_n$ given that $I = \{0, \dots, n\}$. Furthermore, the $\multimap$ connective from separation logic is known as the *magic wand* and expresses hypothetical modifications of the current state.

Apart from these standard CSL connectives, the assertion language contains three different heap ownership predicates $\xrightarrow{\pi}_t$, with $\pi \in (0, 1]_{\mathbb{Q}}$ a *fractional permission* in the style of Boyland [8] and t the *heap ownership type*, where:

- $E \xrightarrow{\pi}_{\text{std}} E'$ is the *standard heap ownership predicate* from separation logic, that provides read-only access for $0 < \pi < 1$ and write access in case $\pi = 1$.
- $E \xrightarrow{\pi}_{\text{proc}} E'$ is the *process heap ownership predicate*, which indicates that the heap location E is bound to an active process algebraic model, but in a *read-only* manner: it only provides read-only access, even when $\pi = 1$.

¹ This intuitively means that the program logic is able to “forget” about resources, which fits naturally with garbage collecting languages like Java and C#.

- $E \xrightarrow{\pi}_{\text{act}} E'$ is the *action heap ownership predicate*, which indicates that the heap location E is bound by an active process algebraic model and is used in the context of an action block, in a *read/write* manner.

The distinction between different types of heap ownership is needed for the program logic to be sound, for example to disallow the deallocation of memory that is bound by a process algebraic model. Moreover, observe that $E \xrightarrow{\pi}_{\text{proc}} E'$ predicates never provide write access to E . However, we shall later see that the proof system allows one to upgrade $\xrightarrow{\pi}_{\text{proc}}$ predicates to $\xrightarrow{\pi}_{\text{act}}$ inside **action** blocks, and $\xrightarrow{\pi}_{\text{act}}$ again provides write access if $\pi = 1$. This system of upgrading enforces that all modifications to E happen in the context of **action** $X.a$ **do** C commands, and can therefore be recorded in the model X as the action a .

Finally, the $\text{Proc}_{\pi}(X, b, P, \Pi)$ assertion expresses ownership of the program model that is identified by X and is represented by the process P . The condition b is the postcondition of the abstract model. Furthermore, Π connects the abstract model to the concrete program, by mapping the models' process algebraic variables to heap locations in the program. And last, the fractional permission π is needed to implement the ownership system of program models. Fractional permissions are only used here to be able to reconstruct the full Proc_1 predicate.

Semantics of Assertions. The interpretation of assertions is defined as a modelling relation $ph, pm, s, g \models \mathcal{P}$, where the models (ph, pm, s, g) consist of the following four components:

- A *permission heap*, $ph \in \text{PermHeap} \triangleq \text{Var} \rightarrow \text{free} \mid \langle v \rangle_t^{\pi}$, that maps values (heap locations) to either **free** (unoccupied) or to occupied entries $\langle v \rangle_t^{\pi}$. Occupied heap cells store a value v , as well as a type t to associate heap cells to the three different kinds of heap ownership predicates used in the logic.
- A *process map*, $pm \in \text{ProcMap} \triangleq \text{Var} \rightarrow \text{free} \mid \langle b, P, \Lambda \rangle^{\pi}$, defined as a total mapping from values (process identifiers) to *process map entries*. Occupied entries have the form $\langle b, P, \Lambda \rangle^{\pi}$ and model ownership of process algebraic models in the program logic. The components $\Lambda \in \text{ProcVar} \rightarrow_{\text{fin}} \text{Val}$ in turn define the models of the abstraction binders (that were defined in Definition 2).
- Two stores, $s, g \in \text{Store}$, that gives an interpretation to all variables used in program and ghost code, respectively. Ghost variables do not interfere with regular program execution and are therefore separated from program variables and maintained in an extra store g , referred to as the *ghost store*.

The semantics of assertions is defined as a modelling relation $\cdot \models \cdot$ between models of the logic $\text{PermHeap} \times \text{ProcMap} \times \text{Store}^2$ and assertions Assn as follows:

Definition 4 (Semantics of assertions (excerpt)). *The interpretation of assertions $ph, pm, s, g \models \mathcal{P}$ is defined by structural recursion on $\mathcal{P}$ in the standard way, except for the following two cases:*

$$\begin{array}{c}
\hookrightarrow\text{-SPLITMERGE} \\
E_1 \xrightarrow{\pi_1 + \pi_2}_t E_2 \dashv\vdash E_1 \xrightarrow{\pi_1}_t E_2 * E_1 \xrightarrow{\pi_2}_t E_2 \\
\\
\text{PROC-SPLITMERGE} \\
\text{Proc}_{\pi_1 + \pi_2}(X, b, P_1 \parallel P_2, \Pi) \dashv\vdash \text{Proc}_{\pi_1}(X, b, P_1, \Pi) * \text{Proc}_{\pi_2}(X, b, P_2, \Pi) \\
\\
\text{PROC-}\cong \\
\frac{P \cong Q}{\text{Proc}_{\pi}(X, b, P, \Pi) \dashv\vdash \text{Proc}_{\pi}(X, b, Q, \Pi)}
\end{array}$$

Fig. 4. Selected entailment rules of the program logic.

$$\begin{array}{ll}
ph, pm, s, g \models E_1 \xrightarrow{\pi}_t E_2 & \text{iff } ph(\llbracket E_1 \rrbracket(s)) = \langle \llbracket E_2 \rrbracket(s) \rangle_t^{\pi'} \wedge \pi \leq \pi' \\
ph, pm, s, g \models \text{Proc}_{\pi}(X, b, P, \Pi) & \text{iff } \exists P'. pm(g(X)) = \langle b, P \parallel P', \llbracket \Pi \rrbracket(s) \rangle^{\pi'} \\
& \wedge \pi \leq \pi' \wedge (\pi = 1 \implies P' = \varepsilon)
\end{array}$$

The full definition of the semantics of assertions can be found in [46].

Clarifying the non-standard cases, $E \xrightarrow{\pi}_t E'$ is satisfied if ph holds an entry at location E that matches with the ownership type t , with an associated fractional permission that is at least π . Process ownership assertions $\text{Proc}_{\pi}(X, b, P, \Pi)$ are satisfied if pm holds a matching entry with a fractional permission at least π , as well as a process that has at least the behaviour of P . The denotation $\llbracket \Pi \rrbracket(s)$ gives the model of the abstraction binder Π , and is defined as follows:

Definition 5 (Semantics of abstraction binders).

$$\llbracket \{x_0 \mapsto E_0, \dots, x_n \mapsto E_n\} \rrbracket(s) \triangleq \{x_0 \mapsto \llbracket E_0 \rrbracket(s), \dots, x_n \mapsto \llbracket E_n \rrbracket(s)\}$$

3.4 Entailment Rules

Figure 4 shows the non-standard entailment rules of the program logic. All other, standard rules can be found in [46]. The notation $\mathcal{P} \dashv\vdash \mathcal{Q}$ is a shorthand notation for $\mathcal{P} \vdash \mathcal{Q}$ and $\mathcal{Q} \vdash \mathcal{P}$, and indicates that the rule can be used in both directions. All rules have shown to be sound in the standard sense, using Coq.

Clarifying the entailment rules, $\hookrightarrow\text{-SPLITMERGE}$ expresses that heap ownership predicates $\xrightarrow{\pi}_t$ of any type t may be *split* (in the left-to-right direction) and be *merged* (right-to-left) along π . This allows one to distribute heap ownership among the different threads in the program. Likewise, PROC-SPLITMERGE allows one to split and merge process ownership along parallel compositions inside abstract models, to distribute them over different threads. More specifically, by splitting a predicate $\text{Proc}_{\pi_1 + \pi_2}(X, b, P_1 \parallel P_2, \Pi)$ into two, both parts can be distributed over different concurrent threads, so that thread i can establish that it executes as prescribed by its part $\text{Proc}_{\pi_i}(X, b, P_i, \Pi)$ of the abstraction. Afterwards, when the threads join again, the remaining partial abstractions can be

$$\begin{array}{c}
 \text{HT-PROCINIT} \\
 \frac{I = \{0, \dots, n\} \quad \text{fv}(b_1) \subseteq \text{dom}(\Pi) = \{x_0, \dots, x_n\} \quad X \notin \text{fv}(\mathcal{R}, E_0, \dots, E_n) \quad B = b_1[x_i/E_i]_{\forall i \in I}}{\Gamma, \{b_1\} P \{b_2\}; \mathcal{R} \vdash \begin{array}{c} \left\{ *_{i \in I} \Pi(x_i) \xrightarrow{1}_{\text{std}} E_i * B \right\} \\ \text{X} := \text{process } P \text{ over } \Pi \\ \left\{ *_{i \in I} \Pi(x_i) \xrightarrow{1}_{\text{proc}} E_i * B * \right. \\ \quad \left. \text{Proc}_1(X, b_2, P, \Pi) \right\} \end{array}} \\
 \\
 \text{HT-PROCUPDATE} \\
 \frac{\begin{array}{c} \text{fv}(a) = \{x_0, \dots, x_n\} \subseteq \text{dom}(\Pi) \quad I = \{0, \dots, n\} \\ B_1 = \text{pre}(a)[x_i/E_i]_{\forall i \in I} \quad B_2 = \text{post}(a)[x_i/E_i]_{\forall i \in I} \\ \Gamma; \mathcal{R} \vdash \{ *_{i \in I} \Pi(x_i) \xrightarrow{\pi_i}_{\text{act}} E_i * B_1 * \mathcal{P} \} C \{ *_{i \in I} \Pi(x_i) \xrightarrow{\pi_i}_{\text{act}} E'_i * B_2 * \mathcal{Q} \} \end{array}}{\Gamma; \mathcal{R} \vdash \begin{array}{c} \left\{ *_{i \in I} \Pi(x_i) \xrightarrow{\pi_i}_{\text{proc}} E_i * B_1 * \right. \\ \quad \left. \text{Proc}_\pi(X, b, a \cdot P + Q, \Pi) * \mathcal{P} \right\} \\ \text{action } X.a \text{ do } C \\ \left\{ *_{i \in I} \Pi(x_i) \xrightarrow{\pi_i}_{\text{proc}} E'_i * B_2 * \right. \\ \quad \left. \text{Proc}_\pi(X, b, P, \Pi) * \mathcal{Q} \right\} \end{array}} \\
 \\
 \text{HT-PROCFINISH} \\
 \frac{\text{fv}(b) \subseteq \text{dom}(\Pi) = \{x_0, \dots, x_n\} \quad I = \{0, \dots, n\} \quad B = b[x_i/E_i]_{\forall i \in I} \quad P \downarrow}{\Gamma; \mathcal{R} \vdash \left\{ *_{i \in I} \Pi(x_i) \xrightarrow{1}_{\text{proc}} E_i * \right. \text{finish } X \left. \left\{ *_{i \in I} \Pi(x_i) \xrightarrow{1}_{\text{std}} E_i * B \right\} \right. \text{Proc}_1(X, b, P, \Pi) \left. \right\}}
 \end{array}$$

Fig. 5. The non-standard Hoare proof rules related to abstract models.

merged back into a single predicate. This system thus provides a compositional way of verifying that programs meet their abstract models.

Finally, $\text{PROC} \cong$ allows one to replace program abstractions by bisimilar ones. This rule is used to rewrite processes in a canonic form used by some other rules.

3.5 Program Judgments

Judgments of programs are defined as sequents of the form $\Gamma; \mathcal{R} \vdash \{\mathcal{P}\} C \{\mathcal{Q}\}$, where $\mathcal{R}$ is a *resource invariant* [9], and Γ is a *process environment*:

Definition 6 (Process environment).

$$\Gamma ::= \emptyset \mid \Gamma, \{b\} P \{b\}$$

Process environments are defined in the style of *interface specifications* [33], and are essentially a series of Hoare-triples $\{b_1\} P \{b_2\}$ for processes P , that constitute the top-level contracts of the programs' abstract models.

The intuitive meaning of a program judgment $\Gamma; \mathcal{R} \vdash \{\mathcal{P}\} C \{\mathcal{Q}\}$ is that, starting from any state satisfying $\mathcal{P} * \mathcal{R}$, the invariant $\mathcal{R}$ is maintained throughout execution of C , and any final state upon termination of C will satisfy $\mathcal{Q} * \mathcal{R}$. Moreover, the proof derivation of C may use any abstract model that is in Γ .

Figure 5 presents the proof rules that handle process algebraic abstractions. All other proof rules are deferred to [46] due to space constraints.

The HT-PROCINIT rule handles initialisation of an abstract model P over a set of heap locations as specified by Π . Standard points-to predicates with write-permission are required for any heap location that is to be bound by P , and these are converted to $\hookrightarrow_{\text{proc}}^1$. Moreover, HT-PROCINIT requires that the precondition of P holds, which is constructed from b_1 by replacing all process variables by the symbolic values at the corresponding heap locations. A Proc_1 predicate with full permission is ensured, containing the postcondition b_2 of the abstract model.

The HT-PROCUPDATE rule handles updates to program abstractions, by performing an action a in the context of an **action** $X.a$ **do** C program, provided that C respects the contract of a . As a precondition, a predicate of the form $\text{Proc}_\pi(X, b, a \cdot P + Q, \Pi)$ is required for some π . The process component of this predicate must be of the form $a \cdot P + Q$ to allow performing the a action. After performing a , this process component will be reduced to P , thereby discarding Q as the choice is made not to follow execution as prescribed by Q . In order to get process components into the required format $a \cdot P + Q$, the $\text{PROC-}\cong$ rule can be used to rewrite process components up to bisimilarity. Furthermore, $\hookrightarrow_{\text{proc}}^\pi$ predicates are required for any heap location that is bound by Π . These points-to predicates are needed to resolve the pre- and postcondition of a .

Finally, HT-PROCFINISH handles finalisation of program models that successfully terminate. A predicate $\text{Proc}_1(X, b, P, \Pi)$ with *full* permission is required, which means that no other thread can have any fragment of the model. This predicate is exchanged for the postcondition of the abstraction. This postcondition can be established, since (i) the contracts of processes in Γ are assumed as their validity is checked externally, and b is a postcondition of one of these contracts; (ii) the abstraction has been initialised in a state satisfying the precondition of that contract; and (iii) the leftover process P is able to successfully terminate. Lastly, all $\hookrightarrow_{\text{proc}}^1$ predicates are converted back to $\hookrightarrow_{\text{std}}^1$ to indicate that the associated heap locations are no longer bound by the abstraction.

3.6 Soundness

The soundness proof of the program logic has been fully mechanised using the Coq proof assistant, as a deep embedding that is inspired by [53]. The overall Coq implementation comprises roughly 15.000 lines of code. Proving soundness was non-trivial and required substantial auxiliary definitions. The Coq development and its documentation can be found at [46].

The soundness theorem relates program judgments to the operational semantics of programs, and amounts to the following: if a proof $\Gamma; \mathcal{R} \vdash \{P\} C \{Q\}$ can be derived for any program C , and if the contracts in Γ of all abstract models of C are satisfied, then C *executes safely for any number of computation steps*. To concretise this, we first define the semantics of program judgments.

Definition 7 (Semantics of program judgments).

$$\begin{aligned} \Gamma; \mathcal{R} \models \{\mathcal{P}\} C \{\mathcal{Q}\} &\triangleq \models \Gamma \implies \forall n, ph, pm, s, g. \\ &ph, pm, s, g \models \mathcal{P} \implies \text{safe}_\Gamma^n(C, ph, pm, s, g, \mathcal{R}, \mathcal{Q}) \end{aligned}$$

The entailment $\models \Gamma$ intuitively means that, for any Hoare triple $\{b_1\} P \{b_2\}$ in Γ and for any σ such that $\llbracket b_1 \rrbracket(\sigma)$, we have that any run $(P, \sigma) \longrightarrow^* (P', \sigma')$ that terminates (i.e., $P' \downarrow$) ends up with a store σ' for which $\llbracket b_2 \rrbracket(\sigma')$ holds.

The predicate safe_Γ^n defines execution safety for n computation steps, meaning that the program is: data-race free, memory safe, complies with its pre- and postconditions, and refines its process algebraic models, for n computation steps. This definition extends the well-known inductive definition of configuration safety of Vafeiadis [53] by adding machinery to handle process algebraic models. The most important extension is a *simulation argument* between program execution (with respect to $\rightsquigarrow$) and the execution of all active models (with respect to $\xrightarrow{a}$). However, as the reduction steps of these two semantics do not directly correspond one-to-one, this simulation is established via an intermediate, instrumented semantics. This intermediate semantics is defined in terms of $\rightsquigarrow_{\text{ghost}}$ transitions that define the *lock-step execution* of program transitions $\rightsquigarrow$ and the transitions $\xrightarrow{a}$ of their abstractions. Our definition of “*executing safely for n execution steps*” includes that all $\rightsquigarrow$ steps can be simulated by $\rightsquigarrow_{\text{ghost}}$ steps and vice versa, for n execution steps. Thus, the end-result is a refinement between programs and their abstract models.

Theorem 1 (Soundness). $\Gamma; \mathcal{R} \vdash \{\mathcal{P}\} C \{\mathcal{Q}\} \implies \Gamma; \mathcal{R} \models \{\mathcal{P}\} C \{\mathcal{Q}\}$

The underlying idea of the above definition, i.e., having a continuation-passing style definition for program judgments, has first been applied in [2] and has further been generalised in [16] and [17]. Moreover, the idea of defining (program) execution safety in terms of an inductive predicate originates from [3]. These two concepts have been reconciled in [53] into a formalisation for the classical CSL of Brookes [9], that has been encoded and mechanically been proven in both Isabelle and Coq. Our definition builds on the latter, by having a refinement between programs and abstractions encoded in *safe*.

4 Implementation

The verification approach has been implemented in the VerCors verifier, which specialises in automated verification of parallel and concurrent programs written in high-level languages, like (subsets of) Java and C [6]. VerCors applies a correctness-preserving translation of the input program into a sequential imperative language, and delegates the generation of verification conditions to the Viper verifier [29] and their verification ultimately to Z3.

Tool support for our technique has been implemented in VerCors for languages with fork/join concurrency and statically-scoped parallel constructs [34]. This is done defining an axiomatic domain for processes in Viper, consisting of

constructors for all process-algebraic connectives, supported by standard process-algebraic axioms. The Proc_π assertions are encoded as predicates over these process types. The three different ownership types $\xrightarrow{\pi}_t$ are encoded by defining extra fields that maintain the ownership status t for each global reference.

To analyse process-algebraic models, VerCors first linearises all processes and then encodes the linear processes and their contracts into Viper. The linearisation algorithm is based on a rewrite system that uses a subset of the standard process-algebraic axioms as rewrite rules [51] to eliminate parallel connectives.

The VerCors implementation of the abstraction approach is much richer than the simple language of Sect. 3 that is used to formalise the approach on. Notably, the abstraction language in VerCors supports general recursion instead of Kleene iteration, and allows parameterising process and action declarations by data. VerCors also has support for several axiomatic data types that enrich the expressivity of reasoning with abstractions, like (multi)sets and sequences.

5 Case Study

Finally, we demonstrate our verification approach on a well-known version of the leader election protocol [35] that is based on shared memory. Most importantly, this case study shows how our approach bridges the typical abstraction gap between process algebraic models and program implementations. In particular, it shows how a high-level process algebraic model of a leader election protocol, together with a contract for this model (checked with mCRL2 for various inputs), is formally connected to an actual program implementation of the protocol.

The protocol is performed by N concurrent workers that are organised in a ring, so that worker i only sends to worker $i + 1$ and only receives from worker $i - 1$, modulo N . The goal is to determine a leader among these workers. To find a leader, the election procedure assumes that each worker i receives a unique integer value to start with, and then operates in N rounds. In every round (i) each worker sends the highest value it encountered so far to its right neighbour, (ii) receives a value from its left neighbour, and (iii) remembers the highest of the two. The result after N rounds is that all workers know the highest unique value in the network, allowing its original owner to announce itself as leader.

The case study has been verified with VerCors using the presented approach. All workers communicate via two standard non-blocking operations for message passing: $\text{mp_send}(r, \text{msg})$ for sending a message msg to the worker with rank r^2 , and $\text{msg} := \text{mp_recv}(r)$ for receiving a message from worker r . The election protocol is implemented on top of this message passing system.

The main challenge of this case study is to define a message passing system on the process algebra level that matches this implementation. To design such a system we follow the ideas of [35]; by defining two actions, $\text{send}(r, \text{msg})$ and $\text{recv}(r, \text{msg})$, that abstractly describe the behaviour of the concrete implementations in mp_send and mp_recv , respectively. Moreover, *process algebraic summation* $\Sigma_{x \in D} P$ is used to quantify over the possible messages that mp_recv might

² The identifiers of workers are typically called *ranks* in message passing terminology.

```

1  seq⟨seq⟨Msg⟩⟩ chan; // communication channels between workers
2  int lead; // rank of the worker that is announced as leader
3
4  /* Action for sending messages. */
5  requires 0 ≤ rank < |chan|;
6  ensures chan[rank] = \old(chan[rank]) + {msg};
7  ensures ∀r' : int. (0 ≤ r' < |chan| ∧ r' ≠ rank) ⇒ chan[r'] = \old(chan[r']);
8  action send(int rank, Msg msg);
9
10 /* Action for receiving messages. */
11 requires 0 ≤ rank < |chan|;
12 ensures {msg} + chan[rank] = \old(chan[rank]);
13 ensures ∀r' : int. (0 ≤ r' < |chan| ∧ r' ≠ rank) ⇒ chan[r'] = \old(chan[r']);
14 action rcv(int rank, Msg msg);
15
16 /* Action for announcing a leader. */
17 requires 0 ≤ rank < |chan|;
18 ensures lead = rank;
19 action announce(int rank);
20
21 /* Local behavioural specification for each worker. */
22 requires 0 ≤ n ≤ |chan| ∧ 0 ≤ rank < |chan|;
23 process Elect(int rank, Msg v0, Msg v, int n) △
24   if 0 < n then send((rank + 1) % |chan|, v) ·
25     Σv' ∈ Msg rcv(rank, v') · Elect(rank, v0, max(v, v'), n - 1)
26   else (if v = v0 then announce(rank) else ε);
27
28 /* Global behavioural specification of the election protocol. */
29 requires |vs| = |chan|;
30 requires ∀i, j : int. (0 ≤ i < |vs| ∧ 0 ≤ j < |vs| ∧ vs[i] = vs[j]) ⇒ i = j;
31 ensures |vs| = |chan| ∧ 0 ≤ lead < |vs|;
32 ensures ∀i : int. (0 ≤ i < |vs|) ⇒ vs[i] ≤ vs[lead];
33 process ParElect(seq⟨Msg⟩ vs) △
34   Elect(0, vs[0], vs[0], |vs|) || ... || Elect(|vs| - 1, vs[|vs| - 1], vs[|vs| - 1], |vs|);

```

Fig. 6. Behavioural specification of the leader election protocol.

receive. The summation operator $\sum_{x \in D} P$ quantifies over a set $D = \{d_0, \dots, d_n\}$ of data and is defined as the (finite) sequence $P[x/d_0] + \dots + P[x/d_n]$ of non-deterministic choices. The following two rules illustrate how the abstract **send** and **rcv** actions are connected to **mp_send** and **mp_rcv** (observe that both these actions are parameterised by data³).

³ Recall that the VerCors implementation of our abstraction technique is much richer than the simple language of Sect. 3 that is used to formalise the approach on.

$$\{\text{send}(r, \text{msg}) \cdot P\} \text{mp_send}(r, \text{msg}) \{P\}$$

$$\{\sum_{x \in \text{Msg}} \text{recv}(r, x) \cdot P\} \text{msg} := \text{mp_recv}(r) \{P[x/\text{msg}]\}$$

Finally, we construct a process-algebraic model of the election protocol using `send` and `recv`, and verify that the implementation adheres to this model. This model has been analysed with mCRL2 for various inputs (since mCRL2 is essentially finite-state) to establish the global property of announcing the correct leader. The deductive proof of the program can then rely on this property.

5.1 Behavioural Specification

Our main goal is proving that the implementation determines the correct leader upon termination. To prove this, we first define a *behavioural specification* of the election protocol that hides all irrelevant implementation details, and prove the correctness property on this specification. Process algebra provides a proper abstraction level that suits our needs well, as the behaviour of leader election can concisely be specified in terms of sequences of sends and receives.

Figure 6 presents the process algebraic specification. In particular, `ParElect` specifies the *global* behaviour whereas `Elect` specifies the *thread-local* behaviour. The `ParElect` process encodes the parallel composition of all eligible participants. `ParElect` takes a sequence `vs` of initial values as argument, whose length equals the total number of workers by its precondition. `ParElect`'s postcondition states that `lead` must be a valid rank after termination and that `vs[lead]` be the highest initial worker value. It follows that worker `lead` is the correctly chosen leader.

The `Elect` process takes four arguments, which are: the rank of the worker, the initial unique value `v0` of that worker, the current highest value `v` encountered by that worker, and finally the number `n` of remaining rounds. The rounds are implemented via general recursion. In each round all workers send their current highest value `v` to their right neighbour (on line 24), receive a value `v'` in return from their left neighbour (line 25), and continue with the highest of the two. The extra `announce` action is declared and used to announce the leader after `n` rounds. The postcondition of `announce` is that `lead` stores the leader's rank.

The contracts of `send` and `recv` describe the behaviour of standard non-blocking message passing. Communication on the specification level is implemented via *message queues*. Message queues are defined as sequences of messages that are taken from a finite domain `Msg`. Since workers are organised in a ring in this case, every worker can do with only a single queue and the global communication channel architecture can be defined as a sequence of message queues: `chan` in the figure. The action contract of `send(r, msg)` expresses enqueueing the message `msg` onto the message queue `chan[r]` of the worker with rank `r`. The postcondition of `send` is that `msg` has been enqueueing onto `chan[r]` and that the queues `chan[r']` for any `r' ≠ r` have not been altered. Likewise, the contract of `recv(r, msg)` expresses dequeuing `msg` from `chan[r]`. The expression `\old(e)` indicates that `e` is to be evaluated with respect to the pre-state of computation.

```

1  global seq⟨seq⟨Msg⟩⟩ C; // implementation of communication channels
2  global int N; // total number of workers
3  global int L; // rank of the leader to be announced
4
5  lock_invariant  $L \xrightarrow{1}_{\text{proc}} - * \exists c : \text{seq}\langle \text{seq}\langle \text{Msg}\rangle \rangle . C \xrightarrow{1}_{\text{proc}} c * N \xrightarrow{\frac{1}{2}}_{\text{proc}} |c|$ ;
6
7  given  $p, P, Q, \Pi, \pi, \pi'$ ;
8  context  $\{ \text{chan} \mapsto C \} \in \Pi * \exists n . N \xrightarrow{\pi}_{\text{proc}} n * 0 \leq \text{rank} < n$ ;
9  requires  $\text{Proc}_{\pi'}(X, p, \text{send}(\text{rank}, \text{msg}) \cdot P + Q, \Pi)$ ;
10 ensures  $\text{Proc}_{\pi'}(X, p, P, \Pi)$ ;
11 void mp_send(ref X, int rank, Msg msg) { /* omitted */ }
12
13 given  $p, P, Q, \Pi, \pi, \pi'$ ;
14 context  $\{ \text{chan} \mapsto C \} \in \Pi * \exists n . N \xrightarrow{\pi}_{\text{proc}} n * 0 \leq \text{rank} < n$ ;
15 requires  $\text{Proc}_{\pi'}(X, p, \Sigma_{m \in \text{Msg}} \text{recv}(\text{rank}, m) \cdot P + Q, \Pi)$ ;
16 ensures  $\text{Proc}_{\pi'}(X, p, P[m/\text{result}], \Pi)$ ;
17 Msg mp_recv(ref X, int rank) { /* omitted */ }
18
19 given  $n, p, \Pi, \pi, \pi'$ ;
20 context  $\{ \text{lead} \mapsto L, \text{chan} \mapsto C \} \in \Pi * N \xrightarrow{\pi}_{\text{proc}} n * 0 \leq \text{rank} < n$ ;
21 requires  $\text{Proc}_{\pi'}(X, p, \text{Elect}(\text{rank}, v_0, v, n), \Pi)$ ;
22 ensures  $\text{Proc}_{\pi'}(X, p, \varepsilon, \Pi)$ ;
23 void elect(ref X, int rank, Msg v0, Msg v) {
24   loop_invariant  $0 \leq i \leq n$ ;
25   loop_invariant  $\text{Proc}_{\pi'}(X, p, \text{Elect}(\text{rank}, v_0, v, n - i), \Pi)$ ;
26   for (int i := 0 to N) {
27     mp_send(X, (rank + 1) % N, v) with {
28        $P := \Sigma_{x \in \text{Msg}} \text{recv}(\text{rank}, x) \cdot \text{Elect}(\text{rank}, v_0, \max(v, x), n - i - 1)$ ,
29        $Q := \varepsilon, p := p, \Pi := \Pi, \pi := \pi, \pi' := \pi'$ 
30     };
31      $\text{Msg } v' := \text{mp\_recv}(X, \text{rank})$  with {
32        $P := \text{Elect}(\text{rank}, v_0, \max(v, v'), n - i - 1)$ ,
33        $Q := \varepsilon, p := p, \Pi := \Pi, \pi := \pi, \pi' := \pi'$ 
34     };
35      $v := \max(v, v')$ ;
36   }
37   if ( $v = v_0$ ) {
38     atomic { action X.announce(rank) do  $L := \text{rank}$ ; }
39   }
40 }

```

Fig. 7. The annotated implementation of the leader election protocol. Annotations of the form **context** $\mathcal{P}$ are shorthand for **requires** $\mathcal{P}$; **ensures** $\mathcal{P}$.

```

41 given  $p, \Pi$ ;
42 context  $N \xrightarrow{\frac{1}{2}}_{\text{proc}} |vs| * 0 < |vs|$ ;
43 requires  $\text{Proc}_1(X, p, \text{ParElect}(vs), \Pi)$ ;
44 ensures  $\text{Proc}_1(X, p, \varepsilon, \Pi)$ ;
45 void parelect(ref  $X$ , seq $\langle \text{Msg} \rangle$   $vs$ ) {
46   context  $0 \leq \text{rank} < |vs|$ ;
47   requires  $\text{Proc}_{1/|vs|}(X, p, \text{Elect}(\text{rank}, vs[\text{rank}], vs[\text{rank}], |vs|), \Pi')$ ;
48   ensures  $\text{Proc}_{1/|vs|}(X, p, \varepsilon, \Pi)$ ;
49   par (int  $\text{rank} := 0$  to  $N$ ) {
50     elect( $X$ ,  $vs[\text{rank}]$ ,  $vs[\text{rank}]$ ) with {
51        $n := N$ ,  $p := p$ ,  $\Pi := \Pi$ ,  $\pi := 1/(4|vs|)$ ,  $\pi' := 1/|vs|$ 
52     };
53   }
54 }
55
56 context  $N \xrightarrow{1}_{\text{std}} - * C \xrightarrow{1}_{\text{std}} - * L \xrightarrow{1}_{\text{std}} -$ ;
57 requires  $\forall i, j : \text{int}. (0 \leq i < |vs| \wedge 0 \leq j < |vs| \wedge vs[i] = vs[j]) \Rightarrow i = j$ ;
58 ensures  $0 \leq \text{result} < |vs|$ ;
59 ensures  $\forall i : \text{int}. (0 \leq i < |vs|) \Rightarrow vs[i] \leq vs[\text{result}]$ ;
60 int main(seq $\langle \text{Msg} \rangle$   $vs$ ) {
61    $N := |vs|$ ,  $C := \text{initialiseChannels}(N)$ ;
62    $X := \text{process } \text{ParElect}(vs)$  over  $\{ \text{chan} \mapsto C, \text{lead} \mapsto L \}$ ;
63   commitLock(); // initialise the lock invariant
64   parelect( $X$ ,  $vs$ ) with  $\{ p := \text{ParElect}(vs), \Pi := \{ \text{chan} \mapsto C, \text{lead} \mapsto L \} \}$ ;
65   uncommitLock(); // reclaim the lock invariant
66   finish  $X$ ; // obtain the global correctness property from the abstraction
67   return  $L$ ; // return rank of leader
68 }

```

Fig. 8. Bootstrap procedures of the leader election protocol.

5.2 Protocol Implementation

Figure 7 presents the annotated implementation of the election protocol⁴. The **elect** method contains the code that is executed by every worker. The contract of **elect**(X , rank , v_0 , v) states that the method body adheres to the behavioural description **Elect**(rank , v_0 , v , N) of the election protocol. Each worker performing **elect** enters a **for**-loop that iterates N times, whose loop invariant states that, at iteration i , the remaining program behaves as prescribed by the process **Elect**(rank , v_0 , v , $i - 1$). The invocations to **mp_send** and **mp_rcv** on lines 27 and 31 are annotated with **with** clauses that resolve the assignments required by the

⁴ It should be noted that the presentation is slightly different from the version that is verified by VerCors, to better connect to the theory discussed in the earlier sections to the case study. Notably, VerCors uses Implicit Dynamic Frames [27] as the underlying logical framework, which is equivalent to separation logic [39] but handles ownership slightly differently. The details of this are deferred to [6, 21].

given clauses in the contracts of `mp_send` and `mp_recv`. The **given** $\bar{\eta}$ annotation expresses that the parameter list $\bar{\eta}$ are extra ghost arguments for the sake of specification. After N rounds all workers with $v = v_0$ announce themselves as leader. However, since the initial values are chosen to be unique there can only be one such worker. Finally, we can verify that at the post-state of `elect` the abstract model has been fully executed and thus reduced to ε .

The `mp_send( $X$ ,  $rank$ ,  $msg$ )` method implements the operation of enqueueing msg onto the message queue of worker $rank$. Its implementation has been omitted for brevity. The contract of `mp_send` expresses that the enqueueing operation is encapsulated as a `send( $rank$ ,  $msg$ )` action that is prescribed by an abstract model identified by X . The `mp_recv( $X$ ,  $rank$ )` function implements the operation of dequeuing and returns the first message of the message queue of worker $rank$. The receive is prescribed as the `recv` action on the abstraction level, where the potential received message is ranged over by the summation on lines 15.

Figure 8 presents bootstrapping code for the implementation of message passing. The `main` function initialises the communication channels whereas `parelect` spawns all worker threads. `main( $vs$ )` additionally initialises and finalises the abstraction `ParElect( $vs$ )` on the specification level (on line 62 and 66, respectively), whose analysis allows one to establish the postconditions of `main`. The function `parelect( $X$ ,  $vs$ )` implements the abstract model `ParElect( $vs$ )` by spawning N workers that all execute the `elect` program. The contract associated to the parallel block (lines 46–48) is called an *iteration contract* and assigns pre- and postconditions to every parallel instance. For more details on iteration contracts we refer to [5]. Most importantly, the iteration contract of each parallel worker states (on line 47) that the worker behaves as specified by `Elect`. Thus, we deductively verify in a *thread-modular way* that the program implements its behavioural specification. Lastly, all the required ownership for the global fields and the `Proc1` predicate is split and distributed among the individual workers via the iteration contract and the **with** clause on lines 50–52.

6 Related Work

Significant progress has been made on the theory of concurrent software verification over the last years [11–13, 30, 41, 48–50]. This line of research proposes advanced program logics that all provide some notion of expressing and restricting thread interference of various complexity, via *protocols* [24]: formal descriptions of how shared-memory is allowed to evolve over time. In our approach protocols have the form of process algebraic abstractions.

The original work on CSL [32] allows specifying simple thread interference in shared-memory programs via resource invariants and critical regions. Later, RGSep [54] merges CSL with rely-guarantee reasoning to enable describing more fine-grained inter-thread interference by identifying atomic concurrent actions. Many modern program logics build on these principles and propose even more advanced ways of verifying shared-memory concurrency. For example, TaDa [41] and CaReSL [50] express thread interference protocols through state-transition

systems. iCAP [48] and Iris [25] propose a more unified approach by accepting user-defined monoids to express protocols on shared state, together with invariants restricting these protocols. Iris provides reasoning support for proving language properties in Coq, where our focus is on proving programs correct.

In the distributed setting, Disel [44] allows specifying protocols for distributed systems. Disel builds on dependent type theory and is implemented as a shallow embedding in Coq. Even though their approach is more expressive than ours, it can only semi-automatically be applied in the context of Coq. Villard et al. [55] present a program logic for message passing concurrency, where threads may communicate over channels using native send/receive primitives. This program logic allows specifying protocols via *contracts*, which are state-machines in the style of Session Types [18], to describe channel behaviour. Our technique is more general, as the approach of Villard et al. is tailored specifically to basic shared-memory message passing. Actor Services [47] is a program logic with assertions to express the consequences of asynchronous message transfers between actors. However, the meta-theory of Actor Services has not been proven sound.

Most of the related work given so far is essentially theoretical and mainly focuses on expressiveness and generality. Our approach is a trade-off between expressivity and usability. It allows specifying process algebraic protocols over a general class of concurrent systems, while also allowing the approach to be implemented in automated verifiers for concurrency like VerCors. Related concurrency verifiers are SmallfootRG [10], VeriFast [19], CIVL [45], THREADER [14] and Viper [22, 29]; the latter tool is used as the main back-end of VerCors. SmallfootRG is a memory-safety verifier based on RGSep. VeriFast is a rich toolset for verifying (multi-threaded) Java and C programs using separation logic. Notably, Penninckx et al. [40] extend VeriFast with a Petri-net extension to reason about the I/O behaviour of programs. This Petri-net approach is similar to ours, however our technique supports reasoning about abstract models and allows reasoning about more than just I/O behaviour. The CIVL framework can reason about race-freedom and functional correctness of MPI programs written in C [28, 57]. The reasoning is done via bounded model checking combined with symbolic execution. THREADER is an automated verifier for multi-threaded C, based on model checking and counterexample-guided abstraction refinement.

Apart from the proposed technique, VerCors also allows using process algebraic abstractions as *histories* [7, 56]. Also related in this respect are the time-stamped histories of [43], which records atomic state changes in concurrent programs as a history, which are, likewise to our approach, handled as resources in the logic. However, history recording is only suitable for terminating programs.

Finally, there is a lot of more general work on proving linearisability [15, 26, 52], which essentially allows reasoning about fine-grained concurrency by using sequential verification techniques. Our technique, as well as the history-based technique of [7], uses process algebraic linearisation to do so.

7 Conclusion

To reason effectively about realistic concurrent and distributed software, we have presented a verification technique that performs the reasoning at a suitable level of abstraction that hides irrelevant implementation details, is scalable to realistic programs by being modular and compositional, and is practical by being supported by automated tools. The approach is expressive enough to allow reasoning about realistic software as is demonstrated by the case study as well as by [36], and can be implemented as part of an automated deductive program verifier (viz. VerCors). The proof system underlying our technique has mechanically been proven sound using Coq. Our technique is therefore supported by a strong combination of theoretical justification and practical usability.

We consider the presented technique as just the beginning of a comprehensive verification framework that aims to capture many different concurrent and distributed programming paradigms. To illustrate, we recently adapted the presented approach to the distributed case, by allowing process algebraic models to describe message passing behaviour of distributed programs [37].

We are currently further investigating the use of mCRL2 and Ivy to reason algorithmically about program abstractions, e.g., [31]. Moreover, we are planning to investigate the preservation of liveness properties in addition to safety.

Acknowledgements. This work is partially supported by the NWO VICI 639.023.710 Mercedes project and by the NWO TOP 612.001.403 VerDi project.

References

1. Aldini, A., Bernardo, M., Corradini, F.: A Process Algebraic Approach to Software Architecture Design. Springer, London (2010). <https://doi.org/10.1007/978-1-84800-223-4>
2. Appel, A.W., Blazy, S.: Separation logic for small-step CMINOR. In: Schneider, K., Brandt, J. (eds.) TPHOLs 2007. LNCS, vol. 4732, pp. 5–21. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74591-4_3
3. Appel, A., Melliès, P., Richards, C., Vouillon, J.: A very modal model of a modern, major, general type system. In: POPL, vol. 42, pp. 109–122. ACM (2007). <https://doi.org/10.1145/1190216.1190235>
4. Baeten, J.: Process Algebra with Explicit Termination. Eindhoven University of Technology, Department of Mathematics and Computing Science (2000)
5. Blom, S., Darabi, S., Huisman, M.: Verification of loop parallelisations. In: Egyed, A., Schaefer, I. (eds.) FASE 2015. LNCS, vol. 9033, pp. 202–217. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46675-9_14
6. Blom, S., Darabi, S., Huisman, M., Oortwijn, W.: The VerCors tool set: verification of parallel and concurrent software. In: Polikarpova, N., Schneider, S. (eds.) IFM 2017. LNCS, vol. 10510, pp. 102–110. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66845-1_7
7. Blom, S., Huisman, M., Zaharieva-Stojanovski, M.: History-based verification of functional behaviour of concurrent programs. In: Calinescu, R., Rumpe, B. (eds.) SEFM 2015. LNCS, vol. 9276, pp. 84–98. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-22969-0_6

8. Boyland, J.: Checking interference with fractional permissions. In: Cousot, R. (ed.) SAS 2003. LNCS, vol. 2694, pp. 55–72. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-44898-5_4
9. Brookes, S.: A semantics for concurrent separation logic. *Theor. Comput. Sci.* **375**(1–3), 227–270 (2007). <https://doi.org/10.1016/j.tcs.2006.12.034>
10. Calcagno, C., Parkinson, M., Vafeiadis, V.: Modular safety checking for fine-grained concurrency. In: Nielson, H.R., Filé, G. (eds.) SAS 2007. LNCS, vol. 4634, pp. 233–248. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74061-2_15
11. Dinsdale-Young, T., Dodds, M., Gardner, P., Parkinson, M.J., Vafeiadis, V.: Concurrent abstract predicates. In: D’Hondt, T. (ed.) ECOOP 2010. LNCS, vol. 6183, pp. 504–528. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14107-2_24
12. Feng, X.: Local rely-guarantee reasoning. In: POPL, vol. 44, pp. 315–327. ACM (2009). <https://doi.org/10.1145/1480881.1480922>
13. Feng, X., Ferreira, R., Shao, Z.: On the relationship between concurrent separation logic and assume-guarantee reasoning. In: De Nicola, R. (ed.) ESOP 2007. LNCS, vol. 4421, pp. 173–188. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-71316-6_13
14. Gupta, A., Popeea, C., Rybalchenko, A.: Threader: a constraint-based verifier for multi-threaded programs. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 412–417. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22110-1_32
15. Herlihy, M., Wing, J.: Linearizability: a correctness condition for concurrent objects. *TOPLAS* **12**(3), 463–492 (1990). <https://doi.org/10.1145/78969.78972>
16. Hobor, A.: Oracle semantics. Ph.D. thesis, Princeton University (2008)
17. Hobor, A., Appel, A.W., Nardelli, F.Z.: Oracle semantics for concurrent separation logic. In: Drossopoulou, S. (ed.) ESOP 2008. LNCS, vol. 4960, pp. 353–367. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78739-6_27
18. Honda, K., Vasconcelos, V.T., Kubo, M.: Language primitives and type discipline for structured communication-based programming. In: Hankin, C. (ed.) ESOP 1998. LNCS, vol. 1381, pp. 122–138. Springer, Heidelberg (1998). <https://doi.org/10.1007/BFb0053567>
19. Jacobs, B., Smans, J., Philippaerts, P., Vogels, F., Penninckx, W., Piessens, F.: VeriFast: a powerful, sound, predictable, fast verifier for C and Java. In: Bobaru, M., Havelund, K., Holzmann, G.J., Joshi, R. (eds.) NFM 2011. LNCS, vol. 6617, pp. 41–55. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20398-5_4
20. Jones, C.: Tentative steps toward a development method for interfering programs. *TOPLAS* **5**(4), 596–619 (1983). <https://doi.org/10.1145/69575.69577>
21. Joosten, S., Oortwijn, W., Safari, M., Huisman, M.: An exercise in verifying sequential programs with VerCors. In: Summers, A. (ed.) FTFJP, pp. 40–45. ACM (2018). <https://doi.org/10.1145/3236454.3236479>
22. Juhasz, U., Kassios, I., Müller, P., Novacek, M., Schwerhoff, M., Summers, A.: Viper: a verification infrastructure for permission-based reasoning. Technical report, ETH Zürich (2014)
23. Jung, R., Krebbers, R., Birkedal, L., Dreyer, D.: Higher-order ghost state. In: ICFP, vol. 51, pp. 256–269. ACM (2016). <https://doi.org/10.1145/2951913.2951943>
24. Jung, R., et al.: Iris: monoids and invariants as an orthogonal basis for concurrent reasoning. In: POPL, vol. 50, pp. 637–650. ACM (2015). <https://doi.org/10.1145/2676726.2676980>

25. Krebbers, R., Jung, R., Bizjak, A., Jourdan, J.-H., Dreyer, D., Birkedal, L.: The essence of higher-order concurrent separation logic. In: Yang, H. (ed.) ESOP 2017. LNCS, vol. 10201, pp. 696–723. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54434-1_26
26. Krishna, S., Shasha, D., Wies, T.: Go with the flow: compositional abstractions for concurrent data structures. POPL **2**, 1–31 (2017). <https://doi.org/10.1145/3158125>
27. Leino, K.R.M., Müller, P., Smans, J.: Verification of concurrent programs with Chalice. In: Aldini, A., Barthe, G., Gorrieri, R. (eds.) FOSAD 2007-2009. LNCS, vol. 5705, pp. 195–222. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-03829-7_7
28. Luo, Z., Zheng, M., Siegel, S.: Verification of MPI programs using CIVL. In: EuroMPI. ACM (2017). <https://doi.org/10.1145/3127024.3127032>
29. Müller, P., Schwerhoff, M., Summers, A.J.: Viper: a verification infrastructure for permission-based reasoning. In: Jobstmann, B., Leino, K.R.M. (eds.) VMCAI 2016. LNCS, vol. 9583, pp. 41–62. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49122-5_2
30. Nanevski, A., Ley-Wild, R., Sergey, I., Delbianco, G.A.: Communicating state transition systems for fine-grained concurrent resources. In: Shao, Z. (ed.) ESOP 2014. LNCS, vol. 8410, pp. 290–310. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54833-8_16
31. Neele, T., Willemse, T.A.C., Groote, J.F.: Solving parameterised Boolean equation systems with infinite data through quotienting. In: Bae, K., Ölveczky, P.C. (eds.) FACS 2018. LNCS, vol. 11222, pp. 216–236. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-02146-7_11
32. O’Hearn, P.: Resources, concurrency and local reasoning. Theor. Comput. Sci. **375**(1–3), 271–307 (2007). <https://doi.org/10.1016/j.tcs.2006.12.035>
33. O’Hearn, P., Yang, H., Reynolds, J.: Separation and information hiding. In: POPL, vol. 39, pp. 268–280. ACM (2004). <https://doi.org/10.1145/964001.964024>
34. Oortwijn, W., Blom, S., Gurov, D., Huisman, M., Zaharieva-Stojanovski, M.: An abstraction technique for describing concurrent program behaviour. In: Paskevich, A., Wies, T. (eds.) VSTTE 2017. LNCS, vol. 10712, pp. 191–209. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-72308-2_12
35. Oortwijn, W., Blom, S., Huisman, M.: Future-based static analysis of message passing programs. In: Programming Language Approaches to Concurrency- & Communication-centric Software (PLACES), pp. 65–72. Open Publishing Association (2016). <https://doi.org/10.4204/EPTCS.211.7>
36. Oortwijn, W., Huisman, M.: Formal verification of an industrial safety-critical traffic tunnel control system. In: Ahrendt, W., Tapia Tarifa, S.L. (eds.) IFM 2019. LNCS, vol. 11918, pp. 418–436. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-34968-4_23
37. Oortwijn, W., Huisman, M.: Practical abstractions for automated verification of message passing concurrency. In: Ahrendt, W., Tapia Tarifa, S.L. (eds.) IFM 2019. LNCS, vol. 11918, pp. 399–417. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-34968-4_22
38. Owicki, S., Gries, D.: An axiomatic proof technique for parallel programs. Acta Informatica **6**, 319–340 (1975)
39. Parkinson, M.J., Summers, A.J.: The relationship between separation logic and implicit dynamic frames. In: Barthe, G. (ed.) ESOP 2011. LNCS, vol. 6602, pp. 439–458. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-19718-5_23

40. Penninckx, W., Jacobs, B., Piessens, F.: Sound, modular and compositional verification of the input/output behavior of programs. In: Vitek, J. (ed.) ESOP 2015. LNCS, vol. 9032, pp. 158–182. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46669-8_7
41. da Rocha Pinto, P., Dinsdale-Young, T., Gardner, P.: TaDA: a logic for time and data abstraction. In: Jones, R. (ed.) ECOOP 2014. LNCS, vol. 8586, pp. 207–231. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44202-9_9
42. da Rocha Pinto, P., Dinsdale-Young, T., Gardner, P.: Steps in modular specifications for concurrent modules. In: MFPS, pp. 3–18 (2015). <https://doi.org/10.1016/j.entcs.2015.12.002>
43. Sergey, I., Nanevski, A., Banerjee, A.: Specifying and verifying concurrent algorithms with histories and subjectivity. In: Vitek, J. (ed.) ESOP 2015. LNCS, vol. 9032, pp. 333–358. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46669-8_14
44. Sergey, I., Wilcox, J., Tatlock, Z.: Programming and proving with distributed protocols. POPL **2**, 1–30 (2017). <https://doi.org/10.1145/3158116>
45. Siegel, S., et al.: CIVL: the concurrency intermediate verification language. In: International Conference for High Performance Computing, Networking, Storage and Analysis (SC), p. 61. ACM (2015). <https://doi.org/10.1145/2807591.2807635>
46. Supplementary Material. The supplementary material for this paper, consisting of a technical report, the Coq formalisation and the case study, can be found online at <https://github.com/wytseoortwijn/VMCAI20-SharedMemAbstr>
47. Summers, A.J., Müller, P.: Actor services – modular verification of message passing programs. In: Thiemann, P. (ed.) ESOP 2016. LNCS, vol. 9632, pp. 699–726. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49498-1_27
48. Svendsen, K., Birkedal, L.: Impredicative concurrent abstract predicates. In: Shao, Z. (ed.) ESOP 2014. LNCS, vol. 8410, pp. 149–168. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54833-8_9
49. Svendsen, K., Birkedal, L., Parkinson, M.: Modular reasoning about separation of concurrent data structures. In: Felleisen, M., Gardner, P. (eds.) ESOP 2013. LNCS, vol. 7792, pp. 169–188. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-37036-6_11
50. Turon, A., Dreyer, D., Birkedal, L.: Unifying refinement and Hoare-style reasoning in a logic for higher-order concurrency. In: ICFP, pp. 377–390. ACM (2013). <https://doi.org/10.1145/2500365.2500600>
51. Usenko, Y.: Linearization in μCRL . Technische Universiteit Eindhoven (2002)
52. Vafeiadis, V.: Automatically proving linearizability. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 450–464. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14295-6_40
53. Vafeiadis, V.: Concurrent separation logic and operational semantics. In: MFPS, volume 276 of ENTCS, pp. 335–351 (2011). <https://doi.org/10.1016/j.entcs.2011.09.029>
54. 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). https://doi.org/10.1007/978-3-540-74407-8_18

55. Villard, J., Lozes, É., Calcagno, C.: Proving copyless message passing. In: Hu, Z. (ed.) APLAS 2009. LNCS, vol. 5904, pp. 194–209. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-10672-9_15
56. Zaharieva-Stojanovski, M.: Closer to reliable software: verifying functional behaviour of concurrent programs. Ph.D. thesis, University of Twente (2015). <https://doi.org/10.3990/1.9789036539241>
57. Zheng, M., Rogers, M., Luo, Z., Dwyer, M., Siegel, S.: CIVL: formal verification of parallel programs. In: ASE, pp. 830–835. IEEE (2015). <https://doi.org/10.1109/ASE.2015.99>