

# RGSep Lecture Notes

Viktor Vafeiadis

May 13, 2014

## 1 Introduction

In this document, we shall introduce RGSep, a logic that combines rely/guarantee reasoning and separation logic. The presentation here is largely based on [1, Chapter 3] (correcting a number of errors and adapting it to the setting of the programming language introduced in the CSL lecture notes).

## 2 RGSep Assertions

RGSep divides the heap into two logical components: the local heap,  $h_L$ , and the shared heap,  $h_S$ , such that their combination,  $h_L \uplus h_S$ , is defined.

Its assertion language describes both the local and the shared part of the heap. Here is the syntax of RGSep assertions:

$p, q, r ::=$	$P$	Local assertion
	$\boxed{P}$	Shared assertion
	$p * q$	Separating conjunction
	$p \wedge q$	Normal conjunction
	$p \vee q$	Disjunction
	$\forall x. p$	Universal quantification
	$\exists x. p$	Existential quantification

where  $P$  stands for any separation logic assertion. We will often call  $\boxed{P}$  assertions as boxed assertions. Note, however, that boxes are not modalities in the usual sense as they cannot be nested.

Formally, the semantics of assertions are given as a function of type:

$$\llbracket - \rrbracket : \text{RGSepAssn} \rightarrow \text{Stack} \times \text{Heap} \times \text{Heap} \rightarrow \{true, false\}$$

This function is defined recursively over the syntax of RGSep assertions:

$$\begin{aligned}
\llbracket P \rrbracket(s, h_L, h_S) &\stackrel{\text{def}}{\iff} \llbracket P \rrbracket_{\text{SL}}(s, h_L) \wedge \text{def}(h_L \uplus h_S) \\
\llbracket \overline{P} \rrbracket(s, h_L, h_S) &\stackrel{\text{def}}{\iff} \llbracket P \rrbracket_{\text{SL}}(s, h_S) \wedge (h_L = \mathbf{emp}) \\
\llbracket p * q \rrbracket(s, h_L, h_S) &\stackrel{\text{def}}{\iff} \exists h_1, h_2. (h_L = h_1 \uplus h_2) \wedge \llbracket p \rrbracket(s, h_1, h_S) \wedge \llbracket q \rrbracket(s, h_2, h_S) \\
\llbracket p \wedge q \rrbracket(s, h_L, h_S) &\stackrel{\text{def}}{\iff} \llbracket p \rrbracket(s, h_L, h_S) \wedge \llbracket q \rrbracket(s, h_L, h_S) \\
\llbracket p \vee q \rrbracket(s, h_L, h_S) &\stackrel{\text{def}}{\iff} \llbracket p \rrbracket(s, h_L, h_S) \vee \llbracket q \rrbracket(s, h_L, h_S) \\
\llbracket \forall x. p \rrbracket(s, h_L, h_S) &\stackrel{\text{def}}{\iff} \forall v. \llbracket p \rrbracket(s[x := v], h_L, h_S) \\
\llbracket \exists x. p \rrbracket(s, h_L, h_S) &\stackrel{\text{def}}{\iff} \exists v. \llbracket p \rrbracket(s[x := v], h_L, h_S)
\end{aligned}$$

Note that the definition of  $*$  splits the local state, but not the shared state. We say that  $*$  is multiplicative over the local state, but additive over the shared state. In particular,  $\llbracket \overline{P} * \overline{Q} \rrbracket \iff \llbracket \overline{P \wedge Q} \rrbracket$ . The semantics of shared assertions,  $\llbracket \overline{P} \rrbracket$ , could alternatively be presented without  $l=u$ . This results in an equally expressive logic, but the definition above leads to shorter assertions in practice.

RGSep formulas include the separation logic formulas and overload the definition of some separation logic operators ( $*$ ,  $\wedge$ ,  $\vee$ ,  $\exists$  and  $\forall$ ) to act on RGSep assertions. This overloading is intentional and justified by the following Lemma (writing  $\text{Local}(P)$  for the first RGSep assertion kind):

**Lemma 1** (Properties of local assertions).

$$\begin{aligned}
\text{Local}(B) &\iff B \\
(\text{Local}(P) * \text{Local}(Q)) &\iff \text{Local}(P * Q) \\
(\text{Local}(P) \wedge \text{Local}(Q)) &\iff \text{Local}(P \wedge Q) \\
(\text{Local}(P) \vee \text{Local}(Q)) &\iff \text{Local}(P \vee Q) \\
(\exists x. \text{Local}(P)) &\iff \text{Local}(\exists x. P) \\
(\forall x. \text{Local}(P)) &\iff \text{Local}(\forall x. P)
\end{aligned}$$

These follow directly from the semantic definitions. Because of this lemma, we can reduce the notational overhead by making the  $\text{Local}$  implicit. This should not cause any confusion, because according to Lemma 1, the RGSep operators and the separation logic operators coincide for local assertions.

### 3 Describing Interference using RGSep Actions

The strength of rely/guarantee is the relational description of interference between parallel processes. Instead of using relations directly, RGSep describes interference in terms of actions  $\vec{x}. P \rightsquigarrow Q$  that describe the changes

performed to the shared state. Typically,  $P$  and  $Q$  will be linked with some existentially quantified logical variables. (We do not need to mention separately the set of modified shared locations, because these are all included in  $P$ .) The meaning of an action  $\vec{x}. P \rightsquigarrow Q$  is that it replaces the part of the shared state that satisfies  $P$  prior to the action with a part satisfying  $Q$  without changing the rest of the shared state. For example, consider the following action:

$$M, N. \quad \mathbf{x} \mapsto M \rightsquigarrow \mathbf{x} \mapsto N \wedge N \geq M \quad (\text{Increment})$$

It specifies that the value in the heap cell  $\mathbf{x}$  may be changed, but its value is never decremented. The logical variables  $M$  and  $N$  are existentially bound with scope ranging over both the precondition and the postcondition. In this action, the heap footprints of the precondition and of the postcondition both consist of the location  $\mathbf{x}$ . The footprints of the precondition and the postcondition, however, need not be the same. When they are different, this indicates a transfer of ownership between the shared state and the local state of a thread. For instance, consider a simple lock with two operations: **Acquire** which changes the lock bit from 0 to 1, and removes the protected object,  $list(y)$ , from the shared state; and **Release** which changes the lock bit from 1 to 0, and puts the protected object back into the shared state. We can represent these two operations formally as

$$(\mathbf{x} \mapsto 0) * list(y) \rightsquigarrow \mathbf{x} \mapsto 1 \quad (\text{Acquire})$$

$$\mathbf{x} \mapsto 1 \rightsquigarrow (\mathbf{x} \mapsto 0) * list(y) \quad (\text{Release})$$

An action  $P \rightsquigarrow Q$  represents the modification of some shared state satisfying  $P$  to some state satisfying  $Q$ . Its semantics is the following relation:

$$\llbracket \vec{x}. P \rightsquigarrow Q \rrbracket \stackrel{\text{def}}{=} \{(s, h_1 \uplus h_0, h_2 \uplus h_0) \mid \exists \vec{v}. \llbracket P[\vec{v}/\vec{x}] \rrbracket(s, h_1) \wedge \llbracket Q[\vec{v}/\vec{x}] \rrbracket(s, h_2)\}$$

It relates some initial shared heap  $h_1$  satisfying the precondition  $P$  to a final heap  $h_2$  satisfying the postcondition. In addition, there may be some disjoint shared heap  $h_0$  which is not affected by the action. In the spirit of separation logic, we want the action specification as ‘small’ as possible, describing  $h_1$  and  $h_2$  but not  $h_0$ , and use the frame rule to perform the same update on a larger state. The existential quantification over the interpretation,  $\exists \vec{v}$ , allows  $P$  and  $Q$  to have shared logical variables, such as  $M$  and  $N$  in **Increment**.

RGSep represents the rely and guarantee conditions as sets of actions. The relational semantics of a set of actions is the reflexive and transitive closure of the union of the semantics of each action in the set. This allows each action to run any number of times in any interleaved order with respect to the other actions.

$$\llbracket P_1 \rightsquigarrow Q_1, \dots, P_n \rightsquigarrow Q_n \rrbracket = \left( \bigcup_{i=1}^n \llbracket P_i \rightsquigarrow Q_i \rrbracket \right)^*$$

A specification,  $P_1 \rightsquigarrow Q_1$  is allowed by a guarantee  $G$  if its effect is contained in  $G$ .

**Definition 2.**  $(\vec{x}. P \rightsquigarrow Q) \subseteq G \stackrel{\text{def}}{\iff} \llbracket \vec{x}. P \rightsquigarrow Q \rrbracket \subseteq \llbracket G \rrbracket$ .

## 4 Stability of assertions

Rely/guarantee reasoning requires that every pre- and post-condition in a proof is stable under environment interference. A separation logic assertion  $S$  is stable under interference of a relation  $R$  if and only if whenever  $S$  holds initially and we perform an update satisfying  $R$ , then the resulting state still satisfies  $S$ .

**Definition 3** (Stability). *sem\_stable( $S, R$ ) if and only if for all  $s, h$ , and  $h'$  such that  $\llbracket S \rrbracket_{\text{SL}}(s, h)$  and  $(s, h, h') \in R$ , then  $\llbracket S \rrbracket_{\text{SL}}(s, h, h')$ .*

By representing the interference  $R$  as a set of actions, we can reduce stability to a simple syntactic check.

**Lemma 4** (Checking stability).

- $\text{sem\_stable}(S, \llbracket \vec{x}. P \rightsquigarrow Q \rrbracket)$  iff  $((P[\vec{v}/\vec{x}] -\otimes S) * Q[\vec{v}/\vec{x}]) \Rightarrow S$  for all  $\vec{v}$ .
- $\text{sem\_stable}(P, (R_1 \cup R_2)^*)$  iff  $\text{sem\_stable}(S, R_1)$  and  $\text{sem\_stable}(S, R_2)$ .

Informally, the first property says that if from a state that satisfies  $S$ , we remove the part of the state satisfying  $P$ , and replace it with some state satisfying  $Q$ , then this should imply that  $S$  holds again. In the case when the action cannot run, because there is no sub-state of  $S$  satisfying  $P$ , then  $P -\otimes S$  is *false* and the implication holds trivially. An assertion  $S$  is stable under interference of a set of actions  $R$  if and only if it is stable under interference by every action in  $R$ .

RGSep forbids interference on the local state, but permits interference on the shared state. Hence, given an RGSep assertion, only the parts of it that describe the shared state may be affected by interference. We shall say that an RGSep assertion is (syntactically) stable under  $R$ , if all the boxed assertions it contains are stable under  $R$ .

**Definition 5.** *Let  $p$  stable under  $R$  be defined by induction on  $p$  as follows*

- $P$  stable under  $R$  always holds.
- $\boxed{P}$  stable under  $R$  if and only if  $(P; R) \Rightarrow P$ .
- For  $op ::= * \mid \wedge \mid \vee$ , let  $(p_1 \text{ op } p_2)$  stable under  $R$  if and only if  $p_1$  stable under  $R$  and  $p_2$  stable under  $R$ .
- For  $Q ::= \forall \mid \exists$ , let  $(Qx. p)$  stable under  $R$  if and only if  $p$  stable under  $R$ .

If an assertion is syntactically stable, then it is also semantically stable in the following sense:

**Lemma 6.** *If  $p$  stable under  $R$  and  $\llbracket p \rrbracket(s, h_L, h_S)$  and  $(s, h_S, h'_S) \in R$ , then  $\llbracket p \rrbracket(s, h_L, h'_S)$ .*

The converse is not true. For example, let

$$R \stackrel{\text{def}}{=} (M, N. x \mapsto M \rightsquigarrow x \mapsto N)$$

The relation,  $R$ , writes an arbitrary value to  $x$  without changing the rest of the heap. Then,  $\boxed{\exists n. x \mapsto n}$  is stable under  $R$ , whereas the assertions  $P_1 = \boxed{\exists n. x \mapsto n \wedge n \leq 0}$  and  $P_2 = \boxed{\exists n. x \mapsto n \wedge n > 0}$  are not. Therefore,  $P_1 \vee P_2$  is not syntactically stable although it is semantically equivalent to  $\boxed{\exists n. x \mapsto n}$ .

## 5 Specifications and proof rules

Specifications of a command  $C$  are quadruples  $(p, R, G, q)$ , where

- The precondition  $p$  describes the set of initial states in which  $C$  might be executed (both its local and shared parts).
- The rely  $R$  is a relation (i.e. a set of actions) describing the interference caused by the environment.
- The guarantee  $G$  is a relation describing the changes to the shared state, caused by the program.
- The postcondition  $q$  describes the possible resulting local and shared states, should the execution of  $C$  terminate.

The judgement  $C \text{ sat } (p, R, G, q)$  says that any execution of  $C$  from an initial state satisfying  $p$  and under environment interference  $R$  (*i*) does not fault (e.g., by accessing unallocated memory), (*ii*) causes interference at most  $G$ , and, (*iii*) if it terminates, its final state satisfies  $q$ . Its formal definition is given in the next section.

First, we have the familiar specification weakening rule:

$$\frac{C \text{ sat } (p', R', G', q') \quad \begin{array}{l} R \subseteq R' \quad p \Rightarrow p' \\ G' \subseteq G \quad q' \Rightarrow q \end{array}}{C \text{ sat } (p, R, G, q)} \quad (\text{WEAKEN})$$

Next, we have the disjunction and existential quantification rules:

$$\frac{C \text{ sat } (p_1, R, G, q) \quad C \text{ sat } (p_2, R, G, q)}{C \text{ sat } (p_1 \vee p_2, R, G, q)} \quad (\text{DISJ}) \qquad \frac{x \notin \text{fv}(q, C, R, G) \quad C \text{ sat } (p, R, G, q)}{C \text{ sat } (\exists x. p, R, G, q)} \quad (\text{EX})$$

From separation logic, RGSep inherits the frame rule: If a program runs safely with initial state  $p$ , it can also run with additional state  $r$  lying around. Since the program runs safely without  $r$ , it cannot access the additional state; hence,  $r$  is still true at the end. Since the frame,  $r$ , may also specify the shared state, the FRAME rule checks that  $r$  is stable under interference from both the program and its environment. Otherwise, they may invalidate  $r$  during their execution. In the simple case when  $r$  does not mention the shared state, the stability check is trivially satisfied. Similarly, as in CSL, we require that the frame not to mention any variables modified by the command.

$$\frac{C \text{ sat } (p, R, G, q) \quad r \text{ stable under } (R \cup G) \quad \text{writes}(C) \cap \text{fv}(r) = \emptyset}{C \text{ sat } (p * r, R, G, q * r)} \quad (\text{FRAME})$$

Then, there is a proof rule for each construct in the language. The rules for the empty program, sequential composition, non-deterministic choice, and loops are completely standard. Similarly to FRAME, SKIP checks that the precondition,  $p$ , is stable under the rely,  $R$ . Because the empty program does not change the state,  $p$  is trivially stable under interference from the program itself.

$$\frac{p \text{ stable under } R}{\text{skip sat } (p, R, G, p)} \quad (\text{SKIP}) \qquad \frac{p \text{ stable under } R \quad C \text{ sat } (p, R, G, p)}{C^* \text{ sat } (p, R, G, p)} \quad (\text{LOOP})$$

$$\frac{C_1 \text{ sat } (p, R, G, r) \quad C_2 \text{ sat } (r, R, G, q)}{(C_1; C_2) \text{ sat } (p, R, G, q)} \quad (\text{SEQ}) \qquad \frac{C_1 \text{ sat } (p, R, G, q) \quad C_2 \text{ sat } (p, R, G, q)}{(C_1 \oplus C_2) \text{ sat } (p, R, G, q)} \quad (\text{CHOICE})$$

For primitive commands,  $c$ , that do not access the shared state, we adopt the separation logic rules. We have the following rule scheme:

$$\frac{\vdash_{\text{SL}} \{P\} c \{Q\} \quad \text{writes}(c) \cap \text{fv}(R, G) = \emptyset}{c \text{ sat } (P, R, G, Q)} \quad (\text{PRIM})$$

The sidecondition that  $c$  does not modify any variables appearing in  $R$  and  $G$  is a general requirement that our rules enforce. (This is because for simplicity we assume that there are no shared updateable variables between threads.)

The parallel composition rule of RGSep is very similar to the parallel composition rule of rely/guarantee. Its crucial difference is that the precondition and postcondition of the composition are the separating conjunction of the preconditions and postconditions of the individual threads. In essence,

this is the normal conjunction of the shared state assertions, and the separating conjunction of the local state assertions.

$$\frac{C_1 \text{ sat } (p_1, R \cup G_2, G_1, q_1) \quad \text{writes}(C_1) \cap \text{fv}(C_2, p_2, q_2) = \emptyset \\ C_2 \text{ sat } (p_2, R \cup G_1, G_2, q_2) \quad \text{writes}(C_2) \cap \text{fv}(C_1, p_1, q_1) = \emptyset}{(C_1 \parallel C_2) \text{ sat } (p_1 * p_2, R, G_1 \cup G_2, q_1 * q_2)} \quad (\text{PAR})$$

As the interference experienced by thread  $C_1$  can come from  $C_2$  or from the environment of the parallel composition, we have to ensure that both interferences ( $R \cup G_2$ ) are allowed. Similarly  $C_2$  must be able to tolerate interference from  $C_1$  and from the environment,  $R$ .

The most complex rule is that of atomic commands, **atomic**  $C$ . Instead of tackling the general case directly, it is easier if we have two rules. The first rule checks that the atomic block meets its specification in an empty environment, and then checks that the precondition and the postcondition are stable with respect to the actual environment,  $R$ . This reduces the problem from an arbitrary rely condition to an empty rely condition.

$$\frac{(\text{atomic } C) \text{ sat } (p, \emptyset, G, q) \quad p \text{ stable under } R \\ \text{writes}(c) \cap \text{fv}(R) = \emptyset \quad q \text{ stable under } R}{(\text{atomic } C) \text{ sat } (p, R, G, q)} \quad (\text{ATOMR})$$

The second rule is somewhat trickier. Here is a first attempt:

$$\frac{C \text{ sat } (P * P', \emptyset, \emptyset, Q * Q') \\ (P \rightsquigarrow Q) \subseteq G \quad \text{writes}(C) \cap \text{fv}(G) = \emptyset}{(\text{atomic } C) \text{ sat } (\overline{P} * P', \emptyset, G, \overline{Q} * Q')}$$

Within an atomic block, we can access the shared state  $\overline{P}$ , but we must check that changing the shared state from  $P$  to  $Q$  is allowed by the guarantee  $G$ . This rule is sound, but too weak in practice. It requires that the critical region changes the *entire* shared state from  $P$  to  $Q$  and that the guarantee condition allows such a change. We can extend the rule by allowing the region to change only *part* of the shared state  $P$  into  $Q$ , leaving the rest of the shared state ( $F$ ) unchanged, and checking that the guarantee permits the small change  $P \rightsquigarrow Q$ .

$$\frac{C \text{ sat } (P * P', \emptyset, \emptyset, Q * Q') \\ (P \rightsquigarrow Q) \subseteq G \quad \text{writes}(C) \cap \text{fv}(G) = \emptyset}{(\text{atomic } C) \text{ sat } (\overline{P * F} * P', \emptyset, G, \overline{Q * F} * Q')} \quad (\text{ATOM})$$

From these two rules for atomic blocks, we can derive the following more elaborate rule:

$$\frac{\begin{array}{l} (P' \rightsquigarrow Q') \subseteq G \\ \text{fv}(P'') \cap \{\bar{y}\} = \emptyset \\ \text{writes}(C) \cap \text{fv}(R, G) = \emptyset \end{array} \quad \begin{array}{l} C \text{ sat } (P' * P'', \emptyset, \emptyset, Q' * Q'') \\ P \Rightarrow P' * F \quad Q' * F \Rightarrow Q \\ \text{sem\_stable}(\exists \bar{y}. P, R) \quad \text{sem\_stable}(Q, R) \end{array}}{(\text{atomic } C) \text{ sat } (\exists \bar{y}. \boxed{P} * P'', R, G, \exists \bar{y}. \boxed{Q} * Q'')} \quad (\text{ATOMIC})$$

This rule is derivable by applying the rules ATOMR, CONSEQ, EX, and ATOM. In the other direction, ATOM is derivable from this complex rule, but ATOMR is not.

$$\frac{\begin{array}{l} C \text{ sat } (P' * P'', \emptyset, G, Q' * Q'') \quad (\text{G}) \\ (\text{atomic } C) \text{ sat } (\boxed{P' * F} * P'', \emptyset, G, \boxed{Q' * F} * Q'') \quad \text{ATOM} \\ (\text{I}) \quad (\text{atomic } C) \text{ sat } (\exists \bar{y}. \boxed{P} * P'', \emptyset, G, \boxed{Q} * Q'') \quad \text{EX} \\ (\text{S}) \quad (\text{atomic } C) \text{ sat } (\exists \bar{y}. \boxed{P} * P'', \emptyset, G, \exists \bar{y}. \boxed{Q} * Q'') \quad \text{CONSEQ} \end{array}}{(\text{atomic } C) \text{ sat } (\exists \bar{y}. \boxed{P} * P'', R, G, \exists \bar{y}. \boxed{Q} * Q'')} \quad \text{ATOMR}$$

In the derivation, (S) stands for  $\text{sem\_stable}(\exists \bar{y}. P, R)$ ,  $\text{sem\_stable}(Q, R)$ , and  $\text{writes}(C) \cap \text{fv}(R) = \emptyset$ ; (I) stands for  $\text{fv}(P'') \cap \{\bar{y}\} = \emptyset$ ,  $P \Rightarrow P' * F$ , and  $Q' * F \Rightarrow Q$ ; and (G) stands for  $(P' \rightsquigarrow Q') \subseteq G$  and  $\text{writes}(C) \cap \text{fv}(G) = \emptyset$ .

## 6 Encodings of SL and RG

Separation logic and rely/guarantee are trivial special cases of RGSep. This is best illustrated by the parallel composition rule.

**Plain rely/guarantee:** When all the state is shared, we get the standard rely/guarantee rule. In this case, as the local state is empty, we get  $p_1 * p_2 \iff p_1 \wedge p_2$  and  $q_1 * q_2 \iff q_1 \wedge q_2$ .

Formally, we encode  $\vdash C \text{ sat}_{\text{RG}} (P, R, G, Q)$  as  $C \text{ sat } (\boxed{P}, R, G, \boxed{Q})$ .

**Plain separation logic:** When all the state is local, we get the separation logic rules. (Since there is no shared state, we do not need to describe its evolution:  $R$  and  $G$  are simply the identity relation.)

Formally, we encode  $\vdash_{\text{SL}} \{P\} C \{Q\}$  as  $C \text{ sat } (P, \emptyset, \emptyset, Q)$ .

**Concurrent separation logic:** We simply thread the resource invariant  $J$  through the assertions, and make the rely and guarantee conditions assert that it remains unaffected by interference.

Formally, we encode  $J \vdash_{\text{SL}} \{P\} C \{Q\}$  as  $C \text{ sat } (P * \boxed{J}, R, R, Q * \boxed{J})$ , where  $R = \{J \rightsquigarrow J\}$ .



## 7 Soundness

As with CSL, we define the meaning of RGSep judgments in terms of a **safe** predicate. Here, the **safe** predicate records not only the local heap,  $h_L$ , but also the shared heap,  $h_S$ , as this is needed for  $R$  and  $G$ :

**Definition 7.**  $\text{safe}_0(C, s, h_L, h_S, R, G, q)$  holds always.

$\text{safe}_{n+1}(C, s, h_L, h_S, R, G, q)$  if and only if

- (i) if  $C = \mathbf{skip}$ , then  $\llbracket q \rrbracket(s, h_L, h_S)$ ; and
- (ii) for all  $h_F$ ,  $\langle C, s, h_L \uplus h_S \uplus h_F \rangle \not\rightarrow \mathbf{abort}$ ; and
- (iii) whenever  $\langle C, s, h_L \uplus h_S \uplus h_F \rangle \rightarrow \langle C', s', h' \rangle$ , then there exist  $h'_L$  and  $h'_S$  such that  $h' = h'_L \uplus h'_S \uplus h_F$  and  $(s, h_S, h'_S) \in \llbracket G \rrbracket$  and  $\text{safe}_n(C', s', h'_L, h'_S, R, G, q)$ ;
- (iv) whenever  $(s, h_S, h'_S) \in \llbracket R \rrbracket$  and  $\text{def}(h_L \uplus h'_S)$ , then  $\text{safe}_n(C, s, h_L, h'_S, R, G, q)$ .

A configuration is safe for  $n + 1$  steps if (i) whenever it is a terminal configuration, it satisfies the postcondition; and (ii) it does not abort; and (iii) whenever it performs a transition, its change to the shared state satisfies the guarantee and the new configuration remains safe for  $n$  steps; and finally (iv) whenever the environment changes the shared state according to the rely, the resulting configuration remains safe for another  $n$  steps.

The semantics of RGSep judgments is then defined as follows:

**Definition 8.**  $C \text{ sat } (p, R, G, q)$  holds iff  $\text{writes}(C) \cap \text{fv}(R, G) = \emptyset$  and  $\neg \text{locked}(C)$  and  $\forall s, h_L, h_S, n. \llbracket P \rrbracket(s, h_L, h_S) \implies \text{safe}_n(C, s, h_L, h_S, R, G, q)$ .

The definition checks the basic syntactic properties of programs that are RGSep judgments ensure (i.e., that  $C$  does not modify the values of variables appearing in  $R$  and  $G$ , and that the command is a valid user command—not containing the intermediate **inatom** - marker). It then also requires that any initial configuration of  $C$  satisfying the precondition is safe for an arbitrary number of steps.

## 8 Example: Lock-Coupling List

This section demonstrates a fine-grained concurrent linked list implementation of a mutable set data structure. Instead of having a single lock for the entire list, there is one lock per list node.

**Locks** Here is the source code for locking and unlocking a node:

```
lock(p)  { atomic { t := p.lock; assume(t = 0); p.lock := TID; } }
unlock(p) { atomic { p.lock := 0; } }
```

These locks store the identifier of the thread that acquired the lock: **TID** represents the thread identifier of the current thread. Storing thread identifiers is not necessary for the algorithm's correctness, but it facilitates the

<pre> locate(e) {   local p, c;   p := Head;   lock(p);   c := p.next;   while (c.value &lt; e) {     lock(c);     unlock(p);     p := c;     c := p.next;   }   return(p, c); } </pre>	<pre> remove(e) {   local x, y, z;   (x, y) := locate(e);   if (y.value = e) {     lock(y);     z := y.next;     x.next := z;     unlock(x);     dispose(y);   } else {     unlock(x);   } } </pre>	<pre> add(e) {   local x, y, z;   (x, z) := locate(e);   if (z.value ≠ e) {     y := new Node();     y.lock := 0;     y.value := e;     y.next := z;     x.next := y;   }   unlock(x); } </pre>
---	---	---

Figure 1: Source code for lock coupling list operations.

proof. Similarly, `unlock` uses an `atomic` block to indicate that the write to `p.lock` must be done indivisibly.

The following three predicates represent a node in the list: (1)  $N_s(x, v, y)$  represents a node at location  $x$  with contents  $v$  and tail pointer  $y$  and with the lock status set to  $s$ ; (2)  $U(x, v, y)$  represents an unlocked node at location  $x$  with contents  $v$  and tail pointer  $y$ ; and (3)  $L_t(x, v, y)$  represents a node locked with thread identifier  $t$ . We will write  $N_-(x, v, y)$  for a node that may or may not be locked.

$$\begin{aligned}
N_s(x, v, y) &\stackrel{\text{def}}{=} x \mapsto \{\text{.lock}=s; \text{.value}=v; \text{.next}=y\} \\
U(x, v, y) &\stackrel{\text{def}}{=} N_0(x, v, y) \\
L_t(x, v, y) &\stackrel{\text{def}}{=} N_t(x, v, y) \wedge t > 0
\end{aligned}$$

The thread identifier parameter in the locked node is required to specify that a node can only be unlocked by the thread that locked it.

$$\begin{aligned}
t \in T \wedge U(x, v, n) &\rightsquigarrow L_t(x, v, n) && \text{(Lock)} \\
t \in T \wedge L_t(x, v, n) &\rightsquigarrow U(x, v, n) && \text{(Unlock)}
\end{aligned}$$

The `Lock` and `Unlock` actions are parameterised with a set of thread identifiers,  $T$ . This allows us to use the actions to represent both relies and guarantees. In particular, we take a thread with identifier `TID` to have the guarantee with  $T = \{\text{TID}\}$ , and the rely to use the complement on this set.

From the `ATOMIC` rule, we can derive the following rules for the lock primitives. (Arguably, these specifications are not as simple as one might hope. More advanced logics can give better specifications to `lock` and `unlock`.)

$$\frac{P \text{ stable under } R \quad Q \text{ stable under } R \quad P \Rightarrow N_{\perp}(p, v, n) * F \quad L_{\text{TID}}(p, v, n) * F \Rightarrow Q}{\vdash \text{lock}(p) \text{ sat } (\overline{P}, R, G, \overline{Q})} \quad (\text{LOCK})$$

$$\frac{P \text{ stable under } R \quad Q \text{ stable under } R \quad P \Rightarrow L_{\text{TID}}(p, v, n) * F \quad U(p, v, n) * F \Rightarrow Q}{\vdash \text{unlock}(p) \text{ sat } (\overline{P}, R, G, \overline{Q})} \quad (\text{UNLOCK})$$

**The Algorithm** Now we build a fine-grained concurrent linked list implementation of a set using the lock mechanism we have defined. The list has operations **add**, which adds an element to the set, and **remove**, which removes an element from the set. Traversing the list uses *lock coupling*: the lock on one node is not released until the next node is locked. Somewhat like a person climbing a rope “hand-over-hand,” you always have at least one hand on the rope.

Figure 1 contains the source code. An element is added to the set by inserting it in the appropriate position while holding the lock of its previous node. It is removed by redirecting the previous node’s pointer while both the previous and the current node are locked. This ensures that deletions and insertions can happen concurrently in the same list. The algorithm makes two assumptions about the list: (1) it is sorted; and (2) the first and last elements have sentinel values  $-\infty$  and  $+\infty$  respectively. This allows us to avoid checking for the end of the list.

First, consider the action of adding a node to the list. Here is an action that ignores the sorted nature of the list:

$$t \in T \wedge L_t(x, u, n) \rightsquigarrow L_t(x, u, m) * U(m, v, n)$$

To add an element to the list, we must have locked the previous node, and then we can swing the tail pointer to the added node. The added node must have the same tail as previous node before the update.

To ensure that the sorted order of the list is preserved, the actual action must be specified with respect to the next node as well. We ensure the value we add is between the previous and next values.

$$(t \in T) \wedge (u < v < w) \wedge (L_t(x, u, n) * N_s(n, w, y)) \rightsquigarrow L_t(x, u, m) * U(m, v, n) * N_s(n, w, y) \quad (\text{Insert})$$

The final permitted action is to remove an element from the list. The action specifies that to remove node  $n$  from the list, both  $n$  and the previous node ( $x$ ) must be locked. The tail of the previous node is then updated to the removed node’s tail,  $m$ .

$$(t \in T) \wedge (v < \infty) \wedge (L_t(x, u, n) * L_t(n, v, m)) \rightsquigarrow L_t(x, u, m) \quad (\text{Remove})$$

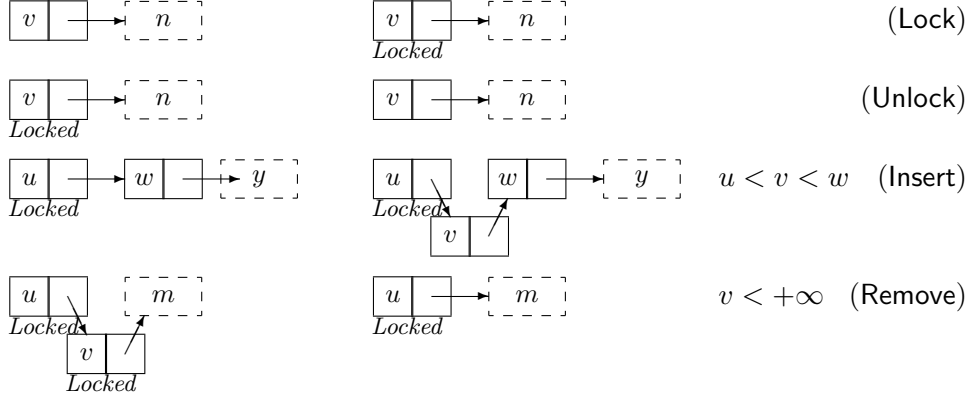


Figure 2: Pictorial representation of the actions

We define  $\mathcal{I}(T)$  as the four actions given above: **Lock**, **Unlock**, **Insert** and **Remove**. These are depicted in Figure 2.  $\mathcal{I}(\{t\})$  allows the thread  $t$ : (i) to lock an unlocked node, (ii) to unlock a node that it had locked, (iii) to insert a node in the list immediately after a node that it had locked, and (iv) if two adjacent nodes in the list are locked by  $t$ , to remove the second node from the list by swinging a pointer past it. For a thread with thread identifier  $t$ , take  $R = \mathcal{I}(\{t\})$  and  $G = \mathcal{I}(\{t\})$ .

We can use separation to describe the structure of the shared list. The following predicate,  $ls(x, A, y)$ , describes a list segment starting at location  $x$  with the final tail value of  $y$ , and with contents  $A$ . We write  $\epsilon$  for the empty sequence and  $\cdot$  for the concatenation of two sequences.

$$ls(x, A, y) \stackrel{\text{def}}{=} (x = y \wedge A = \epsilon \wedge \mathbf{emp}) \vee (\exists v z B. x \neq y \wedge A = v \cdot B * N_-(x, v, z) * ls(z, B, y))$$

Because of separation logic, we do not need any reachability predicates. Instead, the ‘list segment’ predicate is simply a recursively defined predicate. The definition above ensures that the list does not contain any cycles.

The algorithm works on sorted lists with the first and last values being  $-\infty$  and  $+\infty$  respectively. We define  $s(A)$  to represent this restriction on a logical list  $A$ .

$$sorted(A) \stackrel{\text{def}}{=} \begin{cases} true & \text{if } (A = \epsilon) \vee (A = a \cdot \epsilon) \\ (a < b) \wedge sorted(b \cdot B) & \text{if } (A = a \cdot b \cdot B) \end{cases}$$

$$s(A) \stackrel{\text{def}}{=} (\exists B. A = -\infty \cdot B \cdot +\infty) \wedge sorted(A) \wedge \mathbf{emp}$$

Figures 3 and 4 contain the proof outlines of **locate**, **add**, and **remove**. The outline presents the intermediate assertions in the proof. Further, we must prove that every shared state assertion is stable under the rely. These

```

locate(e) {
  local p, c, t;
  { $\exists A. ls(\text{Head}, A, \text{nil}) * s(A) \wedge -\infty < e$ }
  p := Head;
  { $\exists ZB. ls(\text{Head}, \epsilon, p) * N(p, -\infty, Z) * ls(Z, B, \text{nil}) * s(-\infty \cdot B)$ }
  lock(p);
  { $\exists Z. \exists B. ls(\text{Head}, \epsilon, p) * L(p, -\infty, Z) * ls(Z, B, \text{nil}) * s(-\infty \cdot B)$ }
  (c := p.next; )
  { $\exists B. ls(\text{Head}, \epsilon, p) * L(p, -\infty, c) * ls(c, B, \text{nil}) * s(-\infty \cdot B)$ }
  (t := c.value; )
  { $\exists u. \exists ABZ. ls(\text{Head}, A, p) * L(p, u, c) * N(c, t, Z) * ls(c, B, \text{nil}) * s(A \cdot u \cdot t \cdot B)$ }
  while (t < e) {
    { $\exists u. \exists ABZ. ls(\text{Head}, A, p) * L(p, u, c) * N(c, t, Z) * ls(c, B, \text{nil}) * s(A \cdot u \cdot t \cdot B)$ }
    lock(c);
    { $\exists uZ. \exists AB. ls(\text{Head}, A, p) * L(p, u, c) * L(c, t, Z) * ls(Z, B, \text{nil}) * s(A \cdot u \cdot t \cdot B)$ }
    unlock(p);
    { $\exists Z. \exists AB. ls(\text{Head}, A, c) * L(c, t, Z) * ls(Z, B, \text{nil}) * s(A \cdot t \cdot B)$ }
    p := c;
    { $\exists uZ. \exists AB. ls(\text{Head}, A, p) * L(p, u, Z) * ls(Z, B, \text{nil}) * s(A \cdot u \cdot B)$ }
    (c := p.next; )
    { $\exists u. \exists AB. ls(\text{Head}, A, p) * L(p, u, c) * ls(c, B, \text{nil}) * s(A \cdot u \cdot B)$ }
    (t := c.value; )
    { $\exists u. \exists ABZ. ls(\text{Head}, A, p) * L(p, u, c) * N(c, t, Z) * ls(Z, B, \text{nil}) * s(A \cdot u \cdot t \cdot B)$ }
  }
  { $\exists uv. \exists ABZ. ls(\text{Head}, A, p) * L(p, u, c) * N(c, v, Z) * ls(Z, B, \text{nil}) * s(A \cdot u \cdot v \cdot B)$ }
  return (p, c);
}

```

Figure 3: Outline verification of locate.

```

add(e) { local x, y, z, t;
  { $\exists A. ls(\text{Head}, A, \text{nil}) * s(A)$ }  $\wedge -\infty < e$ 
  (x, z) := locate(e);
  { $\exists uv. \exists ZAB. ls(\text{Head}, A, x) * L(x, u, z) * N(z, v, Z)$ 
   *  $ls(Z, B, \text{nil}) * s(A \cdot u \cdot v \cdot B)$ }  $\wedge u < e \wedge e \leq v$ 
  (t = z.value;) if(t  $\neq$  e) {
    { $\exists uv. \exists ZAB. ls(\text{Head}, A, x) * L(x, u, z) * N(z, v, Z)$ 
     *  $ls(Z, B, \text{nil}) * s(A \cdot u \cdot v \cdot B)$ }  $\wedge u < e \wedge e < v$ 
    y = cons(0, e, z);
    { $\exists uv. \exists ZAB. ls(\text{Head}, A, x) * L(x, u, z) * N(z, v, Z)$ 
     *  $ls(Z, B, \text{nil}) * s(A \cdot u \cdot v \cdot B)$ } *  $U(y, e, z) \wedge u < e \wedge e < v$ 
    (x.next = y;)
    { $\exists uv. \exists ZAB. ls(\text{Head}, A, x) * L(x, u, y) * N(y, e, Z) * ls(Z, B, \text{nil}) * s(A \cdot u \cdot e \cdot B)$ }}
  }
  unlock(x);
  { $\exists v. \exists A. ls(\text{Head}, A, \text{nil}) * s(A)$ }
}

remove(e) { local x, y, z, t;
  { $\exists A. ls(\text{Head}, A, \text{nil}) * s(A)$ }  $\wedge -\infty < e \wedge e < +\infty$ 
  (x, y) = locate(e);
  { $\exists uv. \exists ZAB. ls(\text{Head}, A, x) * L(x, u, y) * N(y, v, Z)$ 
   *  $ls(Z, B, \text{nil}) * s(A \cdot u \cdot v \cdot B)$ }  $\wedge u < e \wedge e \leq v \wedge e < +\infty$ 
  (t = y.value;) if(t = e) {
    { $\exists u. \exists ZAB. ls(\text{Head}, A, x) * L(x, u, y) * N(y, e, Z)$ 
     *  $ls(Z, B, \text{nil}) * s(A \cdot u \cdot e \cdot B)$ }  $\wedge e < +\infty$ 
    lock(y);
    { $\exists u. \exists ZAB. ls(\text{Head}, A, x) * L(x, u, y) * L(y, e, Z)$ 
     *  $ls(Z, B, \text{nil}) * s(A \cdot u \cdot e \cdot B)$ }  $\wedge e < +\infty$ 
    (z := y.next;)
    { $\exists u. \exists AB. ls(\text{Head}, A, x) * L(x, u, y) * L(y, e, z)$ 
     *  $ls(z, B, \text{nil}) * s(A \cdot u \cdot B)$ }  $\wedge e < +\infty$ 
    (x.next := z;)
    { $\exists u. \exists AB. ls(\text{Head}, A, x) * L(x, u, z) * ls(z, B, \text{nil}) * s(A \cdot u \cdot B)$ } *  $L(y, e, z)$ 
    unlock(x);
    { $\exists A. ls(\text{Head}, A, \text{nil}) * s(A)$ } *  $L(y, e, z)$ 
    dispose(y);
  } else { { $\exists u. \exists ZAB. ls(\text{Head}, A, x) * L(x, u, y) * ls(y, B, \text{nil}) * s(A \cdot u \cdot B)$ }}
    unlock(x);
  }
  { $\exists A. ls(\text{Head}, A, \text{nil}) * s(A)$ }
}

```

Figure 4: Outline verification of add and remove.

proofs involve reasoning about the septraction operator ( $-\otimes$ ); they are long, but straightforward.

**Theorem 9.** *The lock coupling algorithm is safe, and maintains the sorted nature of the list.*

## References

- [1] Vafeiadis, V., *Modular fine-grained concurrency verification*. PhD thesis. University of Cambridge (2007) See chapter 3.