

A Marriage of Rely/Guarantee and Separation Logic

Viktor Vafeiadis and Matthew Parkinson

University of Cambridge

Abstract. In the quest for tractable methods for reasoning about concurrent algorithms both rely/guarantee logic and separation logic have made great advances. They both seek to tame, or control, the complexity of concurrent interactions, but neither is the ultimate approach. Rely-guarantee copes naturally with interference, but its specifications are complex because they describe the entire state. Conversely separation logic has difficulty dealing with interference, but its specifications are simpler because they describe only the relevant state that the program accesses.

We propose a combined system which marries the two approaches. We can describe interference naturally (using a relation as in rely/guarantee), and where there is no interference, we can reason locally (as in separation logic). We demonstrate the advantages of the combined approach by verifying a lock-coupling list algorithm, which actually disposes/frees removed nodes.

1 Introduction

Reasoning about shared variable concurrent programs is difficult, because the interference between the simultaneously executing threads must be taken into account. Our aim is to find methods that allow this reasoning to be done in a modular and composable way.

On the one hand, we have rely/guarantee, a well-established method, introduced by Jones [11], that is popular in the derivation and the post-hoc verification of concurrent algorithms. RG provides a good way of describing interference by having two relations, the rely R and the guarantee G , which describe the state changes performed by the environment or by the program respectively. Its disadvantage is that the specification of interference is *global*: it must be checked against every state update, even if it is ‘obvious’ that the update cannot interfere with anything else. Even Jones [12] acknowledges this limitation and still considers the search for a satisfactory compositional approach to concurrency an ‘open problem.’

On the other hand, the recent development of separation logic [19,15] suggests that greater modularity is possible. There, the $*$ operator and the frame rule are used to carve all irrelevant state out of the specification and focus only on the state that matters for the execution of a certain component or thread. This makes specifications *local*; two components may interfere, only if they have

overlapping specifications. Its disadvantage is that, in dealing with concurrent programs, it took the simplest approach and uses invariants to specify thread interaction. This makes expressing the relational nature of interference often quite difficult and requires many auxiliary variables [17]. Even O’Hearn acknowledges the weaknesses of separation logic, and asks if “a marriage between separation logic and rely-guarantee is also possible” [15].

Here we present such a marriage of rely/guarantee and separation logic, which combines their advantages and eliminates some of their weaknesses. We split the state into two disjoint parts: (i) the shared state which is accessible by all threads, and (ii) the local state which is accessible by a single component. Then, we use rely/guarantee to deal with the shared state, and separation logic to deal with the local state. This is best illustrated by our parallel composition rule:

$$\frac{\vdash C_1 \text{ sat } (p_1, R \cup G_2, G_1, q_1) \quad \vdash C_2 \text{ sat } (p_2, R \cup G_1, G_2, q_2)}{\vdash C_1 \parallel C_2 \text{ sat } (p_1 * p_2, R, G_1 \cup G_2, q_1 * q_2)}$$

This rule is identical to the standard rely/guarantee rule except for the use of $*$ instead of \wedge in the pre- and post-conditions. In our specifications, the pre-conditions (e.g. p_1) and the postconditions (e.g. q_1) describe both the local and the shared state. The rely conditions (e.g. $R \cup G_2$) and the guarantee conditions (e.g. G_1) describe inter-thread interference: how the shared state gets modified.

The separating conjunction between assertions about both the local and the shared state splits local state (l) in two parts, but does not divide the shared state (s).

$$(p_1 * p_2)(l, s) \stackrel{\text{def}}{=} \exists l_1 l_2. l = l_1 \uplus l_2 \wedge p_1(l_1, s) \wedge p_2(l_2, s)$$

The parallel composition rules of rely/guarantee and separation logic are special cases of our parallel composition rule. (1) When the local state is empty, then $p_1 * p_2 = p_1 \wedge p_2$ and we get the standard rely/guarantee rule. (2) When the shared state is empty, we do not need to describe its evolution (R and G are the identity relation). Then $p_1 * p_2$ has the same meaning as separation logic $*$, and we get the parallel rule of concurrent separation logic without resource invariants (see §2.1).

An important aspect of our approach is that the boundaries between the local state and the shared state are not fixed, but may change as the program runs. This “ownership transfer” concept is fundamental to proofs in concurrent separation logic.

In addition, as we encompass separation logic, we can cleanly reason about dynamically allocated data structures and explicit memory management, avoiding the need to rely on a garbage-collector. In §4, we demonstrate this by verifying a lock-coupling list algorithm, which actually disposes/frees removed nodes.

2 Technical Background

In this paper, we reason about a parallel programming language with pointer operations. Let x, y and z range over logical variables, and x, y and z over

program variables. We assume `tid` is a special variable that identifies the current thread. Commands C and expressions e are given by the following grammar,

$$\begin{aligned} C ::= & \mathbf{x} := e \mid \mathbf{x} := [e] \mid [e_1] := e_2 \mid \mathbf{x} := \text{cons}(e_1, \dots, e_n) \mid \text{dispose}(e) \\ & \mid C_1; C_2 \mid C_1 \parallel C_2 \mid \mathbf{if}(b)\{C_1\} \mathbf{else} \{C_2\} \mid \mathbf{while}(b)\{C\} \mid \mathbf{atomic}(b)\{C\} \\ e ::= & x \mid \mathbf{x} \mid e + e \mid n \end{aligned}$$

where b ranges over boolean expressions. Note that expressions e are *pure*: they do not refer to the heap. In the grammar, each assignment contains at most one heap access; assignments with multiple heap accesses can be performed using multiple assignments and temporary variables to store the intermediate results.

The semantics of **atomic** are that C will be executed in one indivisible step. This could be implemented through locking, hardware atomicity, transactional memories, etc. Choosing **atomic** over a given synchronisation primitive (e.g. locks) enables our reasoning to be applied at multiple abstraction levels. In any case, any synchronisation primitive can be encoded using **atomic**.

2.1 Local Reasoning – Separation Logic

In Hoare logic [9], assertions describe properties of the *whole* memory, and hence specifications, e.g. $\{P\} C \{Q\}$, describe a change of the whole memory. This is inherently *global reasoning*. Anything that is not explicitly preserved in the specification could be changed, for example $\{\mathbf{x} = 4\} \mathbf{y} := 5 \{\mathbf{x} = 4\}$. Here \mathbf{y} is allowed to change, even though it is not mentioned in the specification.¹

The situation is different in *separation logic* [19]. Assertions describe properties of *part* of the memory, and hence specifications describe changes to *part* of the memory. The rest of the memory is guaranteed to be unchanged. This is the essence of *local reasoning*, specifications describe only the memory used by a command, its footprint.

The strength of separation logic comes from a new logical connective: the separating conjunction, $*$. $P * Q$ asserts the state can be split into two parts, one described by P and the other by Q . The separating conjunction allows us to formally capture the essence of *local reasoning* with the following rules:

$$\frac{\{P\} C \{Q\}}{\{P * R\} C \{Q * R\}} \text{ (Frame)} \qquad \frac{\{P_1\} C_1 \{Q_1\} \quad \{P_2\} C_2 \{Q_2\}}{\{P_1 * P_2\} C_1 \parallel C_2 \{Q_1 * Q_2\}} \text{ (Par)}$$

The first rule says, if P is separate from R , and C transforms P into Q then if C finishes we have Q and separately still have R . The second rule says that if two threads have disjoint memory requirements, they can execute safely in parallel, and the postcondition is simply the composition of the two threads' postconditions.²

¹ 'Modifies clauses' solve this problem, but they are neither pretty nor general.

² Originally, separation logic did not consider global variables as resource; hence the proof rules had nasty side-conditions. Later, this problem was solved by Bornat et al. [2]. By disallowing direct assignments to global variables, we avoid the problem.

Separation logic has the following assertions for describing the heap, h :

$$P, Q, S ::= \text{false} \mid \mathbf{emp} \mid e = e' \mid e \mapsto e' \mid \exists x. P \mid P \Rightarrow Q \mid P * Q \mid P -\circledast Q$$

We encode $\neg, \wedge, \vee, \forall$, and **true** in the classical way. **emp** stands for the empty heap; $e \mapsto e'$ for the heap consisting of a single cell with address e and contents e' . Separating conjunction, $P * Q$, is the most important operator of separation logic. A heap h satisfies $P * Q$, if it can be split in two parts, one of which satisfies P and the other satisfies Q . There remains one new connective to describe: *septraction*, $P -\circledast Q$.³ Intuitively, $P -\circledast Q$ represents removing P from Q . Formally, it means the heap can be extended with a state satisfying P , and the extended state satisfies Q .

$$\begin{aligned} h, i \models_{\text{SL}} (P * Q) &\stackrel{\text{def}}{=} \exists h_1, h_2. (h_1 \uplus h_2 = h) \wedge h_1, i \models_{\text{SL}} P \wedge h_2, i \models_{\text{SL}} Q \\ h, i \models_{\text{SL}} (P -\circledast Q) &\stackrel{\text{def}}{=} \exists h_1, h_2. (h_1 \uplus h = h_2) \wedge h_1, i \models_{\text{SL}} P \wedge h_2, i \models_{\text{SL}} Q \end{aligned}$$

Finally, $e \mapsto e_1, \dots, e_n$ is a shorthand for $(e \mapsto e_1) * \dots * (e + n - 1 \mapsto e_n)$.

3 The Combined Logic

3.1 Describing Interference

The strength of rely/guarantee is the careful description of interference between parallel processes. We describe interference in terms of actions $P \rightsquigarrow Q$ which describe the changes performed to the shared state. These resemble Morgan's *specification statements* [13], and P and Q will typically 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 $P \rightsquigarrow Q$ is that it replaces the part of the state that satisfies P before the action with a part satisfying Q . Its semantics is the following relation:

$$[[P \rightsquigarrow Q]] = \{(h_1 \uplus h_0, h_2 \uplus h_0) \mid h_1, i \models_{\text{SL}} P \wedge h_2, i \models_{\text{SL}} Q\}$$

It relates some initial state h_1 satisfying the precondition P to a final state h_2 satisfying the postcondition. In addition, there may be some disjoint state h_0 which is not affected by the action. In the spirit of separation logic, we want action specifications 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 rely and guarantee conditions are simply sets of actions. Their semantics as a relation is the reflexive and transitive closure of the union of the semantics of each action in the set. We shall write R for a syntactic rely condition (i.e. a set of actions) and \mathcal{R} for a semantic rely condition (i.e. a binary relation).

³ Sometimes called "existential magic wand", as it is the dual to "magic wand": $P -\circledast Q \stackrel{\text{def}}{=} \neg(P * \neg Q)$. It has been used in the connection with modal logic in [4].

3.2 Stability

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

Definition 1 (Stability). $S; \mathcal{R} \Longrightarrow S$ iff for all s, s' and i such that $s, i \models_{\text{SL}} S$ and $(s, s') \in \mathcal{R}$, then $s', i \models_{\text{SL}} S$

By representing the interference \mathcal{R} as a set of actions, we reduce stability to a simple syntactic check. For a single action $\llbracket P \rightsquigarrow Q \rrbracket$, the following separation logic implication is necessary and sufficient:

Lemma 1. $S; \llbracket P \rightsquigarrow Q \rrbracket \Longrightarrow S$ iff $\models_{\text{SL}} (P \multimap S) * Q \Longrightarrow S$.

Informally, it says that if from a state that satisfies S , we subtract the part of the state satisfying P , and replace it with some state satisfying Q , then the result should still satisfy S . When the action cannot fire because there is no substate of S satisfying P , then $P \multimap S$ is false and the implication holds trivially.

An assertion S is stable under interference of a set of actions R when it is stable under interference of every action in R .

Lemma 2. $S; (\mathcal{R}_1 \cup \mathcal{R}_2)^* \Longrightarrow S$ iff $S; \mathcal{R}_1 \Longrightarrow S$ and $S; \mathcal{R}_2 \Longrightarrow S$.

Finally, we define $\text{wssa}_{\mathcal{R}}(Q)$ to be the weakest assertion that is stronger than Q and stable under \mathcal{R} .

Definition 2 (Weakest stable stronger assertion). (1) $\text{wssa}_{\mathcal{R}}(Q) \Rightarrow Q$,
 (2) $\text{wssa}_{\mathcal{R}}(Q); \mathcal{R} \Longrightarrow \text{wssa}_{\mathcal{R}}(Q)$, and
 (3) for all P , if $P; \mathcal{R} \Longrightarrow P$ and $P \Rightarrow Q$, then $P \Rightarrow \text{wssa}_{\mathcal{R}}(Q)$.

3.3 Local and Shared State Assertions

We can specify a state using two assertions, one describing the local state and the other the shared state. However, this approach has some drawbacks: specifications are longer, and extending the logic to a setting with multiple disjoint regions of shared state is clumsy.

Instead, we consider a unified assertion language that describes both the local and the shared state. This is done by extending the positive fragment of separation logic assertions with ‘boxed’ terms. We could use boxes for both local and shared assertions: for example, \boxed{P}_{local} and $\boxed{P}_{\text{shared}}$. However, since $\boxed{P}_{\text{local}} * \boxed{Q}_{\text{local}} \iff \boxed{P * Q}_{\text{local}}$ holds for $*$, and all the classical operators, we can omit the $\boxed{}_{\text{local}}$ and the “shared” subscript. Hence the syntax of assertions is

$$p, q, r ::= P \mid \boxed{P} \mid p * q \mid p \wedge q \mid p \vee q \mid \exists x. p \mid \forall x. p$$

Semantically, we split the state, σ , of the system into two components: the local state l , and the shared state s . Each component state may be thought to be

a partial finite function from locations to values. We require that the domains of the two states are disjoint, so that the total state is simply the (disjoint) union of the two states. Assertions without boxes describe purely the local state l , whereas a boxed assertion \boxed{P} describes the shared state s . Formally, we give the semantics with respect to a ‘rely’ condition R , a set of actions describing the environment interference:

$$\begin{aligned}
l, s, i \models_R P &\iff l, i \models_{\text{SL}} P \\
l, s, i \models_R \boxed{P} &\iff l = \emptyset \wedge s, i \models_{\text{SL}} \text{wssa}_{\llbracket R \rrbracket}(P) \\
l, s, i \models_R p_1 * p_2 &\iff \exists l_1, l_2. (l = l_1 \uplus l_2) \wedge (l_1, s, i \models_R p_1) \wedge (l_2, s, i \models_R p_2) \\
l, s, i \models_R p_1 \wedge p_2 &\iff (l, s, i \models_R p_1) \wedge (l, s, i \models_R p_2) \\
\dots &
\end{aligned}$$

Note that $*$ is multiplicative over the local state, but additive over the shared state. Hence, $\boxed{P} * \boxed{Q} \iff \boxed{P \wedge Q}$. The semantics of shared assertions, \boxed{P} , could alternatively be presented without $l = \emptyset$. This results in an equally expressive logic, but the definition above leads to shorter assertions in practice.

We use $\text{wssa}_{\llbracket R \rrbracket}(_)$ to make assertions semantically resistant to interference:

Lemma 3. *If $(l, s, i \models_R p)$, $s' \uplus l$ defined and $\llbracket R \rrbracket(s, s')$ then $(l, s', i \models_R p)$.*

We define an assertion to be syntactically stable if each of the assertions about the shared state is stable. By construction, any assertion about the local state of a component is unaffected by other components, because interference can happen only on the shared state. On the other hand, a boxed assertion \boxed{S} may be affected.

Definition 3 (Stable assertion). *P stable under R always; \boxed{P} stable under R iff $P; \llbracket R \rrbracket \implies P$; $(p \text{ op } q)$ stable under R iff p stable under R and q stable under R ; and $(qu \ x. p)$ stable under R iff p stable under R where $\text{op} ::= \wedge \mid \vee \mid *$ and $\text{qu} ::= \forall \mid \exists$.*

This syntactic condition allows us to change the interpretation of a formula to a more permissive rely.

Lemma 4. *If $(l, s, i \models_R p)$, $\llbracket R \rrbracket \subseteq \llbracket R' \rrbracket$ and p stable under R' then $(l, s, i \models_{R'} p)$.*

We present a few entailments for formulae involving shared states.

$$\begin{array}{ccccccc}
\frac{P \vdash_{\text{SL}} Q}{\boxed{P} \vdash \boxed{Q}} & \boxed{P} \wedge \boxed{Q} \vdash \boxed{P \wedge Q} & \boxed{P} \vee \boxed{Q} \vdash \boxed{P \vee Q} & \boxed{P} * \boxed{Q} \vdash \boxed{P \wedge Q} \\
\forall x. \boxed{P} \vdash \boxed{\forall x. P} & \exists x. \boxed{P} \vdash \boxed{\exists x. P} & \boxed{P} \vdash \boxed{P} * \boxed{P} & \boxed{P} \vdash \text{emp}
\end{array}$$

3.4 Ownership Transfer

Usually the precondition and postcondition of an action have the same heap footprint. For example, consider the action saying that \mathbf{x} can be incremented:

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

If they have a different footprints, this indicates a transfer of ownership between the shared state and the local state of a thread. Consider a simple lock with

$$\begin{array}{c}
\frac{\frac{\frac{\vdash C \text{ sat } (p, R, G, q)}{\left(\begin{array}{l} (r \text{ stable under } R \cup G) \\ \vee (C \text{ has no atomics}) \end{array} \right)}}{\vdash C \text{ sat } (p * r, R, G, q * r)}}{\vdash C_1 \text{ sat } (p, R, G, q)} \quad \frac{Q = (P * X \mapsto Y) \quad x \notin \text{fv}(P)}{\vdash (x := [e]) \text{ sat } (\boxed{Q} \wedge e = X, R, G, \boxed{Q} \wedge x = Y)} \\
\frac{\frac{\frac{\vdash C_1 \text{ sat } (p, R, G, q)}{\vdash C_2 \text{ sat } (q, R, G, r)}}{\vdash C_1; C_2 \text{ sat } (p, R, G, r)}}{\vdash C_1 \parallel C_2 \text{ sat } (p_1 * p_2, R, G_1 \cup G_2, q_1 * q_2)} \quad \frac{\frac{\frac{\vdash C_1 \text{ sat } (p_1, R \cup G_2, G_1, q_1)}{\vdash C_2 \text{ sat } (p_2, R \cup G_1, G_2, q_2)} \quad p_1 \text{ stable under } R \cup G_1 \quad p_2 \text{ stable under } R \cup G_2}{\vdash C_1 \parallel C_2 \text{ sat } (p_1 * p_2, R, G_1 \cup G_2, q_1 * q_2)}}{\vdash C_1 \parallel C_2 \text{ sat } (p_1 * p_2, R, G_1 \cup G_2, q_1 * q_2)} \\
\frac{\frac{\frac{\frac{\vdash C \text{ sat } (P_1 * P_2, \{\}, \{\}, Q_1 * Q_2)}{\overline{y} \cap \text{FV}(P_2) = \emptyset} \quad P \Rightarrow P_1 * F \quad Q_1 * F \Rightarrow Q}{} \quad \frac{\boxed{Q} \text{ stable under } R}{(P_1 \rightsquigarrow Q_1) \subseteq G}}{\vdash (\text{atomic}\{C\}) \text{ sat } (\overline{\exists y}. \boxed{P}) * P_2, R, G, \exists \overline{y}. \boxed{Q} * Q_2)}
\end{array}$$

Fig. 1. Proof rules

two operations: (Acq) which changes the lock bit from 0 to 1, and removes the protected object, $\text{list}(y)$, from the shared state; and (Rel) which changes the lock bit from 1 to 0, and replaces the protected object into the shared state. We can represent these two operations formally as

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

3.5 Specifications and Proof Rules

The judgement $\vdash C \text{ sat } (p, R, G, q)$ semantically says that any execution of C from an initial state satisfying p and under interference at most R , (i) does not fault (e.g. access unallocated memory), (ii) causes interference at most G , and, (iii) if it terminates, its final state satisfies q .

The key proof rules are presented in Figure 1. The rest can be found in the technical report [22]. From separation logic, we inherit the frame rule. This rule says that a program safely running with initial state p can also be executed with additional state r . As the program runs safely without r , it cannot access r when it is present; hence, r is still true at the end. The additional premise is needed because r might mention the shared state and C might modify it in an **atomic**.

We adopt all of the small axioms for local state from separation logic (not presented) [14]. Additionally, we have a read axiom (Fig. 1 top right) for shared state, which allows a non-atomic read from a shared location if we can rely on its value not changing. Note that we do not need to check stability for this read.

The next rule is that of conditional critical regions **atomic**(b) $\{C\}$. For clarity, we present the rule where the guard b is just **true**. The general case, where b is non-trivial and may access the heap, just complicates the essential part of the rule. A simple rule for critical regions would be the following:

$$\frac{\vdash C \text{ sat } (P, \{\}, \{\}, Q) \quad (P \rightsquigarrow Q) \subseteq G \quad \boxed{Q} \text{ stable under } R}{\vdash (\text{atomic}\{C\}) \text{ sat } (\boxed{P}) R, G, \boxed{Q}}$$

$$\begin{array}{c}
 \frac{}{x \mapsto y \rightsquigarrow x \mapsto y \subseteq G} \text{G-EXACT} \\
 \frac{P_1 \rightsquigarrow S * Q_1 \subseteq G \quad P_2 * S \rightsquigarrow Q_2 \subseteq G}{P_1 * P_2 \rightsquigarrow Q_1 * Q_2 \subseteq G} \text{G-SEQ} \\
 \frac{\vDash_{\text{SL}} P' \Rightarrow P \quad P \rightsquigarrow Q \subseteq G \quad \vDash_{\text{SL}} Q' \Rightarrow Q}{P' \rightsquigarrow Q' \subseteq G} \text{G-CONS} \\
 \frac{P \rightsquigarrow Q \in G}{P \rightsquigarrow Q \subseteq G} \text{G-AX} \\
 \frac{P \rightsquigarrow Q \subseteq G}{P[e/x] \rightsquigarrow Q[e/x] \subseteq G} \text{G-SUB} \\
 \frac{(P * F) \rightsquigarrow (Q * F) \subseteq G}{P \rightsquigarrow Q \subseteq G} \text{G-COFRM}
 \end{array}$$

Fig. 2. Rules and axioms for guarantee allows an action

As in RG, we must check that the postcondition is stable under interference from the environment, and that changing the shared state from P to Q is allowed by the guarantee G .

This rule is sound, but too weak in two ways. First, it does not allow critical regions to access any local state, as the precondition \boxed{P} requires that the local state is empty. Second, it requires that the critical region changes the *entire* shared state from P to Q and that the guarantee condition allows such a change. Thus, we extend the rule by (i) adding a precondition P_2 and a postcondition Q_2 for the local state, and (ii) allowing the region to change a part P_1 of P into a part Q_1 of Q , ensuring that the rest F does not change. Additionally, we allow some existential quantifiers, \overline{y} in the shared state to be pulled out over both the shared and local state.

A specification, $P_1 \rightsquigarrow Q_1$ is allowed by a guarantee G if its effect is contained in G . Fig. 2 provides rules to approximate this definition in proofs. The rule G-SEQ allows actions to be sequenced and builds in a form of framing. Note that, if S is empty, then the rule is a parallel composition of two actions; if P_2 and Q_1 are empty, then the rule sequences the actions. It would be simpler, if we simply included the frame rule however this is unsound. In fact, the coframe rule G-COFRM is admissible. G-CONS is similar to the rule of consequence, but the second implication is reversed, $Q \Rightarrow Q'$. Semantically, the property is defined as follows:

Definition 4. $P \rightsquigarrow Q \subseteq G$ iff $\llbracket P \rightsquigarrow Q \rrbracket \subseteq \llbracket G \rrbracket$.

There is a side-condition to the atomic rule requiring that Q is a *precise* assertion. This is formally defined in §5 (Footnote. 7). This is a technical requirement inherited from concurrent separation logic. It ensures that the splitting of the resultant state into local and shared portions is unambiguous.

We reiterate the parallel composition rule from the introduction. As the interference experienced by thread C_1 can arise from C_2 or the environment of the parallel composition, we have to ensure that this interference $R \cup G_2$ is allowed. Similarly C_2 must be able to tolerate interference from C_1 and from the environment of the parallel composition. The precondition and postcondition of the composition are the separating conjunction, $*$, of the preconditions/postconditions of the individual threads. In essence, this is the conjunction of the shared

<pre>lock(p) { atomic(p.lock==0){ p.lock = tid; //p.oldn = p.next; } } unlock(p) { atomic(true) { p.lock = 0; } }</pre>	<pre>locate(e) { local p,c; p = Head; lock(p); c = p.next; while(c.value<e){ lock(c); unlock(p); p = c; c = p.next; } return (p,c); }</pre>	<pre>add(e) { local x,y,z; (x,z)=locate(e); if(z.value!=e){ y = cons(0,e,z); x.next = y; } unlock(x); }</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); // A dispose(y); } else { unlock(x); } }</pre>
---	--	---	--

Fig. 3. Source code for lock coupling list operations. For clarity, we use a field notation, hence we encode $p.lock$, $x.value$, $x.next$ and $p.oldn$ as $[p]$, $[x + 1]$, $[x + 2]$ and $[p + 3]$, respectively. Commented code is auxiliary, that is, required only for the proof.

state assertions, and the separating conjunction of the local state assertions (cf. the semantics of $*$ in §3.3).

The proof rules for conditional and iterative commands are completely standard (See [22].)

4 Example

This section uses the new logic to verify a fine-grained concurrent linked list implementation of a mutable set data structure (see Fig. 3). It has operations `add` which adds an element to the set, and `remove` which removes an element from the set.

The algorithm associates one lock per list node rather than have a single lock for the entire list. 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.

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 values $-\infty$ and $+\infty$ respectively. This allows us to avoid checking for the end of the list.

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

$$\begin{aligned}
N_s(x, v, y) &\stackrel{\text{def}}{=} x \mapsto s, v * \left((s = 0 \wedge x + 2 \mapsto y, _) \vee (s \neq 0 \wedge x + 3 \mapsto y) \right) \wedge x \bmod 4 = 0 \\
U(x, v, y) &\stackrel{\text{def}}{=} N_0(x, v, y) \quad L_t(x, v, y) \stackrel{\text{def}}{=} N_t(x, v, y) \wedge t > 0
\end{aligned}$$

We assume nodes are aligned, $x \bmod 4 = 0$, and **cons** returns aligned nodes.⁴ 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. The fourth field/cell is auxiliary. It is used to store the last value of the nodes tail before it was locked. Once a node is locked its tail field is released to the locking thread, allowing it to mutate the field outside of critical sections, the auxiliary field is used in the proof to track the list structure.

Actions. The algorithm does four kinds of actions: (1) **lock**, which locks a node, (2) **unlock**, which unlocks a node, (3) **add**, which inserts a new node to the list, and (4) **delete**, which removes a node from the list. All of these 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 of this set. Let $I(T)$ be the set of these four actions.

The first two actions are straightforward:

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

Now, consider adding a node to the list. We begin by describing 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 preserve the sorted order of the list, the actual **add** action must also mention the next node: the inserted value must be between the previous and the next values.

$$\begin{aligned}
(t \in T) \wedge (u < v < w) \wedge (L_t(x, u, n) * N_s(n, w, y)) \\
\rightsquigarrow L_t(x, u, m) * U(m, v, n) * N_s(n, w, y) &&& \text{(add)}
\end{aligned}$$

The final action we allow is removing an element from the list. We must lock the node we wish to delete, n , and its previous node, x . The tail of the previous node must be updated to the deleted node's tail, m .

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

⁴ Without this restriction a node could be formed by parts of two adjacent nodes. Instead of assuming alignment, this problem can also be solved by allowing contexts in actions, for example the node is reachable from the head.

List predicate. We use separation to describe the structure of the shared list. The 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 use \cdot as a list separator.

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

Note, as we use separation logic we do not need any reachability predicates, our predicate is simply a recursively defined predicate. The use of $*$ and the inequality $x \neq y$ ensures the list is acyclic. Removing a node from a list segment simply gives two list segments.

Proposition 1. $(N_s(x, v, y) \text{---} * ls(w, A, z))$ is equivalent to $\exists BC. (A = B \cdot v \cdot C) \wedge w \neq z \wedge (ls(w, B, x) \downarrow_z * ls(y, C, z) \downarrow_x)$ where $P \downarrow_x \stackrel{\text{def}}{=} P \wedge \neg(x \mapsto _ * \text{true})$

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

$$srt(+\infty \cdot \epsilon) \stackrel{\text{def}}{=} \mathbf{emp} \quad srt(a \cdot b \cdot A) \stackrel{\text{def}}{=} srt(b \cdot A) \wedge a < b \quad s(-\infty \cdot A) \stackrel{\text{def}}{=} srt(A)$$

Main proof. Appendix A contains the proof outline for the remove function. The outline presents the intermediate assertions in the proof. We present one step of the verification of remove function in detail: the unlock action labelled “A” in Fig. 3. For simplicity, we inline the unlock body.

$$\{ \boxed{\exists AB. ls(\text{Head}, A, x) * L_{\text{tid}}(x, u, y) * L_{\text{tid}}(y, e, z) * ls(z, B, \text{nil}) * s(A \cdot u \cdot B)} * (x+2 \mapsto z) \}$$

atomic $\{ \{ L_{\text{tid}}(x, u, y) * L_{\text{tid}}(y, e, z) * (x+2 \mapsto z) \} \cdot \text{x.lock} = 0; \{ U(x, u, z) * L_{\text{tid}}(y, e, z) \} \}$

$$\{ \boxed{\exists A. ls(\text{Head}, A, \text{nil}) * s(A)} * L_{\text{tid}}(y, e, z) \}$$

We must prove four things: (1) the body meets its specification; (2) the body’s specification is allowed by the guarantee; (3) the outer specification’s postcondition is stable; and (4) find a frame, F , that satisfies the two implications.

The first is a simple proof in separation logic. The second follows as:

$$\frac{\begin{array}{l} L_{\text{tid}}(x, u, y) * L_{\text{tid}}(y, e, z) \rightsquigarrow L_{\text{tid}}(x, u, z) \subseteq I(\{\text{tid}\}) \\ L_{\text{tid}}(x, u, z) \rightsquigarrow U(x, u, Vz) \subseteq I(\{\text{tid}\}) \end{array}}{L_{\text{tid}}(x, u, y) * L_{\text{tid}}(y, e, z) \rightsquigarrow U(x, u, z) \subseteq I(\{\text{tid}\})} \text{G-SEQ}$$

Third, to show $\boxed{\exists A. ls(\text{Head}, A, \text{nil}) * s(A)}$ is stable, we use Lemma 1 for the four actions in the rely: `lock`, `unlock`, `add` and `delete`. The proof of stability is long (hence omitted), but the proof steps are largely automatic. We can automate these checks [6].

Finally, we define F as $ls(\text{Head}, A, x) * ls(z, B, \text{nil}) * s(A \cdot u \cdot B)$

Theorem 1. *The algorithm in Fig. 3 is safe and keeps the list always sorted.*

5 Semantics and Soundness

Our semantics follows the abstract semantics for separation logic of Calcagno, O’Hearn and Yang [5]. Rather than presenting the semantics with respect to a

$$\begin{array}{c}
\frac{l \uplus s = l_1 \quad b(l_1) \quad l' \uplus s' = l_2 \quad Q(s')}{(C, (l_1, \emptyset, o)) \xrightarrow{\text{Emp}^*} (\mathbf{skip}, (l_2, \emptyset, o'))} \quad \frac{(C_1, \sigma) \xrightarrow{\mathcal{R}} (C'_1, \sigma')}{\mathbf{p}} \\
\frac{(\mathbf{atomic}_Q(b)\{C\}, (l, s, o)) \xrightarrow{\mathcal{R}} (\mathbf{skip}, (l', s', o'))}{\mathbf{p}} \quad \frac{(C_1 \parallel C_2, \sigma) \xrightarrow{\mathcal{R}} (C'_1 \parallel C_2, \sigma')}{\mathbf{p}} \\
\\
\frac{A(l, l') \quad (l', s, o) \in \text{Heaps}}{(A, (l, s, o)) \xrightarrow{\mathcal{R}} (\mathbf{skip}, (l', s, o))}{\mathbf{p}} \quad \frac{(\neg \exists l'. A(l, l'))}{(A, (l, s, o)) \xrightarrow{\mathcal{R}} \mathbf{fault}}{\mathbf{p}} \\
\\
\frac{\mathcal{R}(s, s') \quad (l, s', o') \in \text{Heaps}}{(C, (l, s, o)) \xrightarrow{\mathcal{R}} (C, (l, s', o'))}{\mathbf{e}}
\end{array}$$

Fig. 4. Abridged operational semantics

particular model of the heap, we use a partial commutative cancellative⁵ monoid (M, \uplus, \emptyset) as an abstract notion of a heap. We use m, l, s and o to range over elements of M .

Our logic explicitly deals with the separation between a thread's own local state (l) and the shared state (s), and hence implicitly the environment's own state (o). Our semantics are given with respect to a structured heap, which separates these three components.⁶ This splitting is only used to prove the soundness of the logic. There is an obvious erasure to a semantics without a splitting.

Definition 5 (Structured heaps). *Heaps* $\stackrel{\text{def}}{=} \{(l, s, o) \mid \{l, s, o\} \subseteq M \wedge l \uplus s \uplus o \text{ is defined}\}$ and $(l_1, s_1, o_1) \uplus (l_2, s_2, o_2)$ defined as (l, s, o) , iff $s_1 = s_2 = s$, $l_1 \uplus l_2 = l$, $o_1 = l_2 \uplus o$, and $o_2 = l_1 \uplus o$; otherwise it is undefined.

We use σ to range over these structured heaps. Again following [5], we use abstract commands, A , and abstract boolean tests, b , for our abstract heap model. Note that by encoding each primitive command onto a pair of abstract commands, we can give our language a grainless semantics [20].

Definition 6. (i) *Primitive commands* A are represented by a subset of $M \times M$, satisfying: (1) If $A(l_1 \uplus l, l_2)$, then either there exists l'_2 such that $A(l_1, l'_2)$ and $l_2 = l \uplus l'_2$, or $\neg \exists l. A(l_1, l)$; and (2) If $\neg \exists l_2. A(l_1 \uplus l, l_2)$, then $\neg \exists l_2. A(l_1, l_2)$. (ii) *Boolean expressions* b are represented by $M \rightarrow \{\mathbf{true}, \mathbf{false}, \mathbf{fault}\}$, satisfying: if $b(l_1 \uplus l) = v$, then either $b(l_1) = v$ or $b(l_1) = \mathbf{fault}$.

We present the key rules of the semantics of the abstract programming language in Figure 4. The rest can be found in the extended version [22]. We define a reduction step $\text{Config}_1 \xrightarrow[\lambda]{\mathcal{R}} \text{Config}_2$, as configuration Config_1 makes a reduction step to Config_2 with possible interference \mathcal{R} and label λ . The label indicates whether this is a program action, \mathbf{p} , or an environment action, \mathbf{e} . Configurations are either **fault** or a pair of a command and a structured heap, (C, σ) . We use $\xrightarrow{\mathcal{R}}^*$ as the transitive and reflex closure of the reduction relation.

⁵ If $m_1 \uplus m = m_2 \uplus m$, then $m_1 = m_2$.

⁶ The assertions simply ignore the environment.

We alter the syntax of **atomic** to have a postcondition annotation Q , to specify how the state is split between shared and local on exit from the block. In CSL the resource invariant does this job, but we do not have a single resource invariant in this logic. Each of these postconditions must be precise, so there is a unique splitting.⁷ Consider the semantics of **atomic** (Figure 4). The non-faulting rule (1) combines the thread’s local state with the shared state to create a new local state, $l \uplus s = l_1$, (2) checks the guard holds of this new state, $b(l_1)$, (3) executes the command with no interference on the shared state (**Emp**), (4) splits the resulting local state into a new shared and local state, $l' \uplus s' = l_2$, and (5) finally checks the postcondition Q holds of the shared state s' . As Q is precise, it uniquely specifies the splitting in step (4). There are three more rules for **atomic** (not presented) where the program faults on the evaluation of the body, the evaluation of the guard, or fails to find a splitting to satisfy the postcondition.

Parallel composition is modelled by interleaving, we just present one of the rules. The three remaining rules concern abstract commands and environment transitions. The abstract command A executes correctly, if it runs correctly by accessing only the local state. Otherwise, A faults. Its execution does not affect the shared and environment states. An environment transition can happen anytime and affects only the shared state and the environment state, provided that the shared-state change describes the rely relation, \mathcal{R} ; the local state is unchanged.

We extend the standard separation logic notion of safety with a guarantee observed by each program action.

Definition 7 (Guarantee). (1) $(C, \sigma, \mathcal{R}) \text{ guar}_0 \mathcal{G}$ always holds; and (2) $(C, \sigma, \mathcal{R}) \text{ guar}_{n+1} \mathcal{G}$ iff if $(C, \sigma) \xrightarrow[\lambda]{\mathcal{R}} \text{Config}$ then there exist $C' \sigma'$ such that $\text{Config} = (C', \sigma')$; $(C', \sigma', \mathcal{R}) \text{ guar}_n \mathcal{G}$; and if $\lambda = \mathbf{p}$ then $(\sigma, \sigma') \in \mathcal{G}$.

Definition 8. $\models C \text{ sat } (p, R, G, q)$ iff for all $R' \subseteq R$ and $\sigma \models_{R'} (p)$, then (1) $\forall n. (C, \sigma, \llbracket R' \rrbracket) \text{ guar}_n \llbracket G \rrbracket$; and (2) if $(C, \sigma) \xrightarrow{\llbracket R' \rrbracket}^* (\text{skip}, \sigma')$ then $\sigma' \models_{R'} (q)$.

Theorem 2 (Soundness). If $\vdash C \text{ sat } (p, R, G, q)$, then $\models C \text{ sat } (p, R, G, q)$

6 Related Work

Owicki & Gries [16] introduced the concept of non-interference between the proofs of parallel threads. Their method is not compositional and does not permit top-down development of a proof because the final check of interference-freedom may fail rendering the whole development useless.

To address this problem, Jones [11] introduced the compositional rely/guarantee method. In the VDM-style, Jones opted for ‘two-state’ postconditions; other authors [23,18] have chosen single-state postconditions. Several authors

⁷ P is precise iff for every $l \in M$, there exists at most one l_P such that $l_P \models_{\text{SL}} P$ and $\exists l'. l_P \uplus l' = l$.

have proved the soundness and relative completeness of rely/guarantee [23,18,7]; Prensa's proof [18] is machine checked by the Isabelle theorem prover. The completeness results are all modulo the introduction of auxiliary variables. Abadi and Lamport [1] have adapted RG to temporal logic and have shown its soundness for safety specifications.

Separation logic [19,15] takes a different approach to interference by forbidding it except in critical regions [10]. An invariant, I , is used to describe the shared state. This is a simple case of our system where the interference specifications (i.e. R and G) are restricted to a very simple relation, $I \rightsquigarrow I$. Brookes has shown concurrent separation logic to be sound [3].

There have been attempts to verify fine-grained concurrent algorithms using both separation logic and rely/guarantee. Vafeiadis *et al.* [21] verify several list algorithms using rely/guarantee. Their proofs require reachability predicates to describe lists and they cannot deal with the disposal of nodes. Parkinson *et al.* [17] verify a non-blocking stack algorithm using concurrent separation logic. Their proof requires a lot of auxiliary state to encode the possible interference. With the logic presented in this paper much of the auxiliary state can be removed, and hence the proof becomes clearer.

Concurrently with our work, Feng, Ferreira and Shao [8] proposed a different combination of rely/guarantee and separation logic, SAGL. Both our approach and SAGL partition memory into shared and private parts. However, in SAGL, every primitive command is assumed to be atomic. Our approach is more flexible and allows one to specify what is atomic; everything else is considered non-atomic. By default, non-atomic commands cannot update shared state, so we only need stability checks when there is an atomic command: in the lock coupling list only at the lock and unlock operations. On the other hand, SAGL must check stability after every single command. Moreover, in SAGL, the rely and guarantee conditions are relations and stability checks are semantic implications. We instead provide convenient syntax for writing down these relations, and reduces the semantic implication into a simple logic implication. This allowed us to automate our logic [6], and hence automatically verify the safety of a collection of fine-grained list algorithms.

SAGL is presented as a logic for assembly code, and is thus hard to apply at different abstraction levels. It does not contain separation logic as a proper subsystem, as it lacks the standard version of the frame rule [19]. This means that it cannot prove the usual separation logic specification of procedures such as `copy_tree` [14]. In contrast, our system subsumes SL [19], as well as the single-resource variant of CSL [15]: hence, the same proofs there (for a single resource) go through directly in our system (for procedures see [22]). Of course, the real interest is the treatment of additional examples, such as lock coupling, that neither separation logic nor rely/guarantee can prove tractably. Our system also includes a rely-guarantee system, which is why we claim to have produced a marriage of the two approaches. It may be possible to extend SAGL to include the frame rule for procedures, but we understand that such extension is by no means obvious.

With this all being said, there are remarkable similarities between our work and SAGL; that they were arrived at independently is perhaps encouraging as to the naturalness of the basic ideas.

7 Conclusion

We have presented a marriage of rely/guarantee with separation logic. We proved soundness with respect to an abstract operational semantics in the style of abstract separation logic [5]. Hence, our proof can be reused with different languages and with different separation logics, e.g. permissions and variables as resource [2]. Our logic allows us to give a clear and simple proof of the lock-coupling list algorithm, which includes memory disposal. Moreover, our logic can be efficiently automated [6].

Acknowledgements. We should like to thank Josh Berdine, Richard Bornat, Cristiano Calcagno, Joey Coleman, Tony Hoare, Cliff Jones, Xinyu Feng, Alan Mycroft, Peter O'Hearn, Uday Reddy, John Reynolds and Hongseok Yang for discussions and feedback on this work. We acknowledge funding from an RAEng/EPSRC fellowship (Parkinson) and a Gates scholarship (Vafeiadis).

References

1. Abadi, M., Lamport, L.: Conjoining specifications. *ACM Trans. Prog. Lang. Syst.* 17(3), 507–534 (1995)
2. Bornat, R., Calcagno, C., Yang, H.: Variables as resource in separation logic. *ENTCS* 155, 247–276 (2006)
3. Brookes, S.D.: A semantics for concurrent separation logic. In: Gardner, P., Yoshida, N. (eds.) *CONCUR 2004*. LNCS, vol. 3170, pp. 16–34. Springer, Heidelberg (2004)
4. Calcagno, C., Gardner, P., Zarfaty, U.: Context logic as modal logic: completeness and parametric inexpressivity. In: *POPL*, pp. 123–134. ACM Press, New York (2007)
5. Calcagno, C., O'Hearn, P., Yang, H.: Local action and abstract separation logic. *LICS* (to appear, 2007)
6. Calcagno, C., Parkinson, M., Vafeiadis, V.: Modular safety checking for fine-grained concurrency. In: *SAS*. LNCS, Springer, Heidelberg (2007)
7. Coleman, J.W., Jones, C.B.: A structural proof of the soundness of rely/guarantee rules. Technical Report CS-TR-987, Newcastle University (October 2006)
8. Feng, X., Ferreira, R., Shao, Z.: On the relationship between concurrent separation logic and assume-guarantee reasoning. In: *Proceedings of ESOP* (2007)
9. Hoare, C.A.R.: An axiomatic basis for computer programming. *Commun. ACM* 12(10), 576–580 (1969)
10. Hoare, C.A.R.: Towards a theory of parallel programming. *Operating Systems Techniques* (1971)
11. Jones, C.B.: Specification and design of (parallel) programs. In: *IFIP Congress*, pp. 321–332 (1983)
12. Jones, C.B.: *Wanted: a compositional approach to concurrency*, pp. 5–15. Springer, New York (2003)

13. Morgan, C.: The specification statement. *ACM Trans. Program. Lang. Syst.* 10(3), 403–419 (1988)
14. O’Hearn, P., Reynolds, J., Yang, H.: Local reasoning about programs that alter data structures. In: *Proceedings of CSL*, pp. 1–19 (2001)
15. O’Hearn, P.W.: Resources, concurrency and local reasoning. In: Gardner, P., Yoshida, N. (eds.) *CONCUR 2004*. LNCS, vol. 3170, pp. 49–67. Springer, Heidelberg (2004)
16. Owicki, S.S., Gries, D.: An axiomatic proof technique for parallel programs. *Acta Informatica* 6, 319–340 (1976)
17. Parkinson, M.J., Bornat, R., O’Hearn, P.W.: Modular verification of a non-blocking stack. In: *POPL* (2007)
18. Prensa Nieto, L.: The rely-guarantee method in Isabelle/HOL. In: Degano, P. (ed.) *ESOP 2003*. LNCS, vol. 2618, pp. 348–362. Springer, Heidelberg (2003)
19. Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: *LICS*, Washington, DC, USA, pp. 55–74. IEEE Computer Society Press, Los Alamitos (2002)
20. Reynolds, J.C.: Toward a grainless semantics for shared-variable concurrency. In: Lodaya, K., Mahajan, M. (eds.) *FSTTCS 2004*. LNCS, vol. 3328, pp. 35–48. Springer, Heidelberg (2004)
21. Vafeiadis, V., Herlihy, M., Hoare, T., Shapiro, M.: Proving correctness of highly-concurrent linearisable objects. In: *PPoPP*, ACM Press, New York (2006)
22. Vafeiadis, V., Parkinson, M.: A marriage of rely/guarantee and separation logic. Technical Report UCAM-CL-TR-687, University of Cambridge (June 2007), <http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-687.html>
23. Xu, Q., de Roever, W.P., He, J.: The rely-guarantee method for verifying shared variable concurrent programs. *Formal Aspects of Computing* 9(2), 149–174 (1997)

A Proof Outline: remove

```

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);
  { $\exists uv. \exists ZAB. ls(\text{Head}, A, x) * L_{\text{tid}}(x, u, y) * N(y, v, Z) * ls(Z, B, \text{nil}) * s(A \cdot u \cdot v \cdot B)$ }
  { $* (x+2 \mapsto y) \wedge u < e \wedge e \leq v \wedge e < +\infty$ }
  t = y.value; if (t == e) {
  { $\exists u. \exists ZAB. ls(\text{Head}, A, x) * L_{\text{tid}}(x, u, y) * N(y, e, Z) * ls(Z, B, \text{nil}) * s(A \cdot u \cdot e \cdot B)$ }
  { $* (x+2 \mapsto y) \wedge e < +\infty$ }
  lock(y);
  { $\exists uZ. \exists AB. ls(\text{Head}, A, x) * L_{\text{tid}}(x, u, y) * L_{\text{tid}}(y, e, Z) * ls(Z, B, \text{nil}) * s(A \cdot u \cdot e \cdot B)$ }
  { $* (x+2 \mapsto y) * (y+2 \mapsto Z) \wedge e < +\infty$ }
  z = y.next; x.next = z;
  { $\exists u. \exists AB. ls(\text{Head}, A, x) * L_{\text{tid}}(x, u, y) * L_{\text{tid}}(y, e, z) * ls(z, B, \text{nil}) * s(A \cdot u \cdot B)$ }
  { $* (x+2 \mapsto z) * (y+2 \mapsto z)$ }
  unlock(x);
  { $\exists A. ls(\text{Head}, A, \text{nil}) * s(A) * L_{\text{tid}}(y, e, z) * (y+2 \mapsto z)$ }
  dispose(y);
  } else { unlock(x); }
  { $\exists A. ls(\text{Head}, A, \text{nil}) * s(A)$ }
}

```