

Transaction Protocol Verification with Labeled Synchronization Logic

Mohsen Lesani

University of California, Riverside

Abstract. Synchronization algorithms that provide the transaction interface are intricate. We present an algorithm description language that explicitly captures the type of the used synchronization objects and associates labels to method calls to explicitly capture their intra-thread order. We use the language to capture architecture independent representations of transactional memory (TM) algorithms. We present a novel logic that enables reasoning about synchronization algorithms that are described in the language. The logic quantifies over program labels and provides specific predicates and intuitive inference rules to reason about the inter-thread execution and linearization orders of labeled method calls. In particular, the logic assertions can directly capture orders that are fundamental to the correctness of transactions. We present a denotational semantics for the language and prove the soundness of the logic. We have formalized the logic in the PVS proof assistant and mechanically constructed the challenging correctness proof of the TL2 TM algorithm.

1 Introduction

Synchronization algorithms such as mutual exclusion, concurrent data structures and transactional memory are subtle. Designing a synchronization algorithm involves choosing *synchronization objects* and programming the coordination logic using them. There is a trade-off in the choice between consistency and efficiency of a synchronization object. For example, although an atomic register maintains consistency in every concurrent execution, it is less efficient than a basic register that does not provide any guarantee in the presence of race. In addition, algorithm designers have to decide and properly program the *order* of method calls in each thread. These orders are crucial to the correctness of the algorithm in every possible interleaving. Intra-thread orders are usually specified using architecture-dependent fence instructions. As a result, a synchronization algorithm is complicated, low-level and prone to bugs. Engineering reliable software stacks built on top of these algorithms requires their precise description and rigorous verification. In this paper, we present a *description language* for synchronization algorithms, a novel *logic* to reason about synchronization algorithm descriptions and apply the logic to *mechanize* the verification of TM algorithms.

Dekker Example in the Description Language. The language explicitly captures the type of the base synchronization objects and the intra-thread order of method calls on them. As an example, consider the description of the Dekker

	$\mathcal{T}: f_1, f_2 : \mathbf{AtomicRegister}$			
	$\mathcal{D}:$			
	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 33%; border-right: 1px solid black; padding: 2px;"> $\mathbf{def\ } init()$ $W_{01} \triangleright f_1.write(0),$ $W_{02} \triangleright f_2.write(0),$ </td> <td style="width: 33%; border-right: 1px solid black; padding: 2px;"> $\mathbf{def\ } tryLock_1()$ $W_1 \triangleright f_1.write(1),$ $R_2 \triangleright x_2 = f_2.read(),$ $\mathbf{if\ } (x_2 = 0),$ $\quad C_{1t} \triangleright \mathbf{return\ } true$ \mathbf{else} $\quad C_{1f} \triangleright \mathbf{return\ } false,$ $\{W_1 \rightarrow R_2\},$ </td> <td style="width: 33%; padding: 2px;"> $\mathbf{def\ } tryLock_2()$ $W_2 \triangleright f_2.write(1),$ $R_1 \triangleright x_1 = f_1.read(),$ $\mathbf{if\ } (x_1 = 0)$ $\quad C_{2t} \triangleright \mathbf{return\ } true$ \mathbf{else} $\quad C_{2f} \triangleright \mathbf{return\ } false,$ $\{W_2 \rightarrow R_1\},$ </td> </tr> </table>	$\mathbf{def\ } init()$ $W_{01} \triangleright f_1.write(0),$ $W_{02} \triangleright f_2.write(0),$	$\mathbf{def\ } tryLock_1()$ $W_1 \triangleright f_1.write(1),$ $R_2 \triangleright x_2 = f_2.read(),$ $\mathbf{if\ } (x_2 = 0),$ $\quad C_{1t} \triangleright \mathbf{return\ } true$ \mathbf{else} $\quad C_{1f} \triangleright \mathbf{return\ } false,$ $\{W_1 \rightarrow R_2\},$	$\mathbf{def\ } tryLock_2()$ $W_2 \triangleright f_2.write(1),$ $R_1 \triangleright x_1 = f_1.read(),$ $\mathbf{if\ } (x_1 = 0)$ $\quad C_{2t} \triangleright \mathbf{return\ } true$ \mathbf{else} $\quad C_{2f} \triangleright \mathbf{return\ } false,$ $\{W_2 \rightarrow R_1\},$
$\mathbf{def\ } init()$ $W_{01} \triangleright f_1.write(0),$ $W_{02} \triangleright f_2.write(0),$	$\mathbf{def\ } tryLock_1()$ $W_1 \triangleright f_1.write(1),$ $R_2 \triangleright x_2 = f_2.read(),$ $\mathbf{if\ } (x_2 = 0),$ $\quad C_{1t} \triangleright \mathbf{return\ } true$ \mathbf{else} $\quad C_{1f} \triangleright \mathbf{return\ } false,$ $\{W_1 \rightarrow R_2\},$	$\mathbf{def\ } tryLock_2()$ $W_2 \triangleright f_2.write(1),$ $R_1 \triangleright x_1 = f_1.read(),$ $\mathbf{if\ } (x_1 = 0)$ $\quad C_{2t} \triangleright \mathbf{return\ } true$ \mathbf{else} $\quad C_{2f} \triangleright \mathbf{return\ } false,$ $\{W_2 \rightarrow R_1\},$		
	$\mathcal{P}:$			
	$L_0 \triangleright init(),$ $L_1 \triangleright l_1 = tryLock_1() \quad \quad L_2 \triangleright l_2 = tryLock_2()$			

(a)

$\frac{\text{X2L} \quad \mathcal{T}(o) \in LT \quad \pi, \Gamma \vdash obj(l) = obj(l') = o \quad \pi, \Gamma \vdash l \prec l'}{\pi, \Gamma \vdash l \prec_o l'}$	$\frac{\text{P2X} \quad \pi, \Gamma \vdash exec(\zeta'c_1) \quad \pi, \Gamma \vdash exec(\zeta'c_2) \quad c_1 \rightarrow_\pi c_2}{\pi, \Gamma \vdash \zeta'c_1 \prec \zeta'c_2}$
--	---

(b)

Fig. 1: (a) Dekker Description $\pi_D = (\mathcal{T}, \mathcal{D}, \mathcal{P})$. (b) Example inference rules.

mutual exclusion algorithm [10] in Figure 1.(a). The description comprises three sections. The first section, typing \mathcal{T} , describes the base synchronization objects that the algorithm uses. Dekker uses two atomic registers as flags. Using basic registers as flags can lead to a race and violation of mutual exclusion. The second section, definitions \mathcal{D} , describes the definition of methods. The initialization method initializes the two flags to zero. In the two try-lock methods, each thread first writes to its own flag and then reads from the flag of the other thread. Each try-lock method allows entry to the critical region only if it finds the flag of the other thread unset. Every method call in the description is uniquely marked with a *label*. The order of writing the flag of the current thread and then reading the flag of the other thread is crucial to the correctness. Reordering these two method calls can violate mutual exclusion. The required orders for the body of each method definition are declared after the body. For example, the order $W_1 \rightarrow R_2$ requires the methods call W_1 to be executed before the method call R_2 . The third section, program \mathcal{P} , represents the concurrent client program for the defined methods. First, the initialization method is called and then two concurrent threads are executed, each calling one of the try-lock methods.

Reasoning and Orders. The mutual exclusion property states that at most one of the two threads can enter the critical section. More precisely, if either of the two methods calls labeled L_1 and L_2 returns *true*, the other one returns *false*. An intuitive classical proof for the mutual exclusion property of the Dekker algorithm is as follows. We directly reason about execution and linearization orders across threads and use the properties of linearizable registers such as the totality of the linearization order. In particular, we use the real-time-preservation property [19] of linearizability that states that if a method call is executed before another on a linearizable object, then the former is linearized before the latter as well. Assume that L_1 returns true, we prove that L_2 does not return true.

If L_1 has returned true, it should have been by C_{1t} . Therefore, the condition of the if statement is satisfied (that is $x_2 = 0$). Therefore, the read operation R_2 from the flag f_2 returns 0. There are two write method calls W_{02} and W_2 on f_2 . The initialization method call W_{02} is executed before both R_2 and W_2 ; therefore, by the real-time preservation property, is linearized before them. Now, W_2 can either linearize before or after R_2 . The first case is not possible because otherwise, W_2 would be the last write to f_2 before R_2 . Therefore, R_2 would return the value that W_2 writes. However, W_2 writes 1, and R_2 has returned 0.

In the second case, the method call R_2 linearizes before W_2 . Therefore, (1) R_2 is executed before or concurrent to W_2 . (This holds because otherwise, R_2 would be executed after W_2 . Thus, by the real-time-preservation property, R_2 would be linearized before W_2 as well that contradicts the assumption of this case.) According to the explicit program orders, (2) W_1 is executed before R_2 and (3) W_2 is executed before R_1 . From the transitivity of the execution order on the three orders 2, 1, and 3 above, we have that W_1 is executed before R_1 . Therefore, W_1 linearizes before R_1 as well. The initialization method call W_{01} is executed before R_1 and W_1 and is therefore linearized before the two. Therefore, W_1 is the last write to f_1 before R_1 . Therefore, R_1 returns the value that W_1 writes that is 1. Therefore, $x_1 = 1$. Therefore, as the condition of the if statement is not satisfied, C_{2f} is executed. Therefore, L_2 returns false.

Labeled Synchronization Logic. Can we build a rigorous foundation for this intuitive style of reasoning? Is it possible to construct formal proofs in this style? We present a novel first-order logic called *labeled synchronization logic (LSL)* that enables reasoning about synchronization algorithms based on the *execution and linearization orders* of method calls on the base synchronization objects. It quantifies over program labels and provides specific predicates for execution order, execution overlap and linearization orders of labeled method calls across threads. These assertions capture critical orders between concurrent operations. In addition, LSL provides simple-to-use *inference rules* to reason about these orders and deduce algorithm correctness. For example, we applied LSL to stated and prove the mutual exclusion property of the Dekker algorithm. The following theorem states that for the Dekker description (π_D) with no prior assumptions (\cdot), it can be inferred that if L_1 returns true, then L_2 returns false. If the first thread can enter the critical section, then the second cannot. The symmetric property can be state and proved similarly.

Theorem 1. $\pi_D, \cdot \vdash (retv(L_1) = true) \Rightarrow (retv(L_2) = false)$.

The full proof is available in the appendix [27] § 10 and the mechanised proof is available at [26]. We show two inference rules of LSL as examples in Figure 1.(b). LSL uses dynamic labels l to uniquely identify method call instances. For example, the label $L_2'W_2$ is a call string that refers to the instance of the method call labeled W_2 (in the definitions section \mathcal{D}) that is executed in the body of the caller method labeled L_2 (in the program section \mathcal{P}).

The rule X2L states the real-time-preservation property of linearizability [19]: the execution order \prec of method calls on an object o of a linearizable type LT is preserved in the linearization order \prec_o of the object o . If a method call

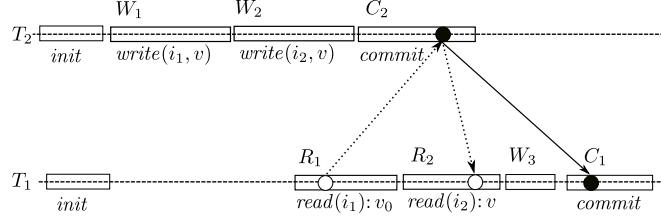


Fig. 2: Illustration of Read-Preservation

is executed before another on a linearizable object, then the former is linearized before the latter as well. Using the rule X2L, a step that we saw in the informal proof of the Dekker algorithm can be formalized as follows. From the fact that method call $L_0'W_{02}$ is executed before \prec the method call $L_2'W_2$ and that both are on the atomic register f_2 ,

$$\pi_D, \Gamma \vdash \text{obj}(L_0'W_{02}) = \text{obj}(L_2'W_2) = f_2 \wedge L_0'W_{02} \prec L_2'W_2$$

we can deduce that $L_0'W_{02}$ is before $L_2'W_2$ in the linearization order \prec_{f_2} of the atomic register f_2

$$\pi_D, \Gamma \vdash L_0'W_{02} \prec_{f_2} L_2'W_2$$

In the sequent, Γ is any set of assumption assertions.

The rule P2X states the *program-order-preservation* property: the program order is preserved in the execution order. (The prefix label variable is denoted by ς .) For example, the Dekker description π_D declares the method call W_1 to be ordered before the method call R_2 i.e. $W_1 \rightarrow_\pi R_2$. Using the rule P2X, from the declared order and that both $L_1'W_1$ and $L_1'R_2$ are executed

$$\pi_D, \Gamma \vdash \text{exec}(L_1'W_1) \wedge \text{exec}(L_1'R_2)$$

we can deduce that $L_1'W_1$ is executed before $L_1'R_2$

$$\pi_D, \Gamma \vdash L_1'W_1 \prec L_1'R_2$$

Correctness of Transactional Memory. Execution and linearization orders of method calls on the base synchronization objects play a critical role in the reasoning about the correctness of transactional memory algorithms. In a previous work [30], we represented a decomposition called *markability* of the TM correctness condition opacity [14]. Markability decomposes opacity to separate intuitive invariants that can be separately verified. An execution history is markable if there is a specific ordering relation on the set of transactions and their read operations called marking such that three invariants are satisfied.

A marking of a transaction history is a relation on the union of the *transactions* and the *read operations* in the history. We can think of the marking as the union of a collection of orders: The *effect order*: The effect order is a total order of the transactions. It represents the order in which the transactions appear to

take effect. The *access orders*: Let us refer to the committed transactions that have write operation(s) to location i as *writers of i* . Consider an unaborting read operation R on a location i . For each such R , the access order is an antisymmetric relation that orders R and every writer of i . The access order of R represents where R has read i between the writers of i .

For example, Figure 2 presents the sketch of a transaction history with two transactions T_1 and T_2 . The horizontal lines from left to right show the time for the two threads executing T_1 and T_2 , and the boxes show execution of method calls on the transactional interface. These method calls may call multiple methods on the base synchronization objects. The dark circles show the effect points of the two transactions and the solid arrow shows the effect order. The transaction T_2 takes effect before T_1 . The transaction T_2 is a writer of both i_1 and i_2 . The white circles show the access points for the two reads R_1 and R_2 . The read R_1 reads i_1 before T_2 writes to it. Similarly, the read R_2 reads i_2 after T_2 writes to it. The relation is called marking as these points can be usually marked as particular *method calls on the base synchronization objects* in the algorithm and the orders are defined as execution and linearization orders on these calls.

As an example, the second invariant of markability called *read-preservation* requires that the location read by a read operation is not overwritten between the two points that the read takes place and the transaction takes effect. Consider an unaborting read operation R from a location i by a transaction T . Intuitively, read-preservation requires that no writer of i comes between R and T in the marking relation. The read-preservation property is violated in Figure 2. The read R_1 is an unaborting read from i_1 in T_1 . The transaction T_2 is a writer of i_1 . The read R_1 is before T_2 in the access order and T_2 is before T_1 in the effect order. The value that R_1 reads is overwritten by T_2 before T_1 is committed. The transaction T_2 writes a new value to both i_1 and i_2 . The read R_1 reads the old value of i_1 and R_2 reads the new value of i_2 that can be inconsistent. Read-preservation is usually simply verified by the validations checks in the commit and read operations.

In this project, we have used LSL to construct a new mechanized proof of TL2 [9], in the PVS proof assistant. Specifically, we have *expressed markability* in the assertion language and then applied LSL inference rules to deduce the markability assertion. This result shows that LSL is scalable to complicated transactional memory algorithms. We have proved the *soundness* of LSL: If an assertion is deduced using valid assumptions, then the deduced assertion is valid as well. An assertion is valid if it evaluates to true in every execution.

The Structure of the Paper. In § 2 and 3, we present the description language and LSL, and In § 4, we state a marking for TL2 and prove the markability of TL2 in LSL. We present the related works in § 5 before conclusion.

2 Description Language

We now present the language that we describe concurrent algorithms in. An algorithm description π is a triple $(\mathcal{T}, \mathcal{D}, \mathcal{P})$ where \mathcal{T} is the *type declarations*

for the used synchronization objects, \mathcal{D} is *the method definitions*, and \mathcal{P} is a concurrent *client program* that calls the defined methods. The set of descriptions Π is defined as follows:

$$\begin{aligned}
\pi \in \Pi &::= (\mathcal{T}, \mathcal{D}, \mathcal{P}) \\
\mathcal{T} &::= (\phi : ot)^* \\
\mathcal{D} &::= d^* \\
d &::= \mathbf{def} \ n_t(x^*) \ s, r \\
s &::= s, s \mid \mathbf{if} \ (b) \ s \ \mathbf{else} \ s \mid q \mid c \triangleright \mathbf{foreach} \ (i \in \mathit{set}) \ s \\
q &::= c \triangleright x = o.n_\tau(u^*) \mid c \triangleright \mathbf{return} \ u \\
b &::= \neg b \mid b \wedge b \mid u = u \mid u = u + u \mid u < u \\
r &::= \{ \} \ (c \rightarrow c)^* \{ \} \\
\mathcal{P} &::= p_0, (p_1 \parallel p_2 \parallel \dots \parallel p_n) \\
p &::= p; p \mid \mathbf{if} \ (b) \ p \ \mathbf{else} \ p \mid c \triangleright x = n_\tau(u^*)
\end{aligned}$$

The description of the Dekker algorithm is presented in Figure 1.(a) as an example. We look at each section in turn.

A typing \mathcal{T} is a mapping from object names ϕ to object types ot . We use x^* to denote a finite sequence of x 's. An object type ot is either a scalar or an array type. A scalar type is either a basic type BT such as `BasicRegister`, `BasicSet`, or `BasicMap`, or a linearizable type LT such as `AtomicRegister`, `AtomicCASRegister`, `Lock`, `TryLock`, or strong counter `SCounter`. As an example, in the Dekker description of Figure 1.(a), both flags f_1 and f_2 are declared to be atomic registers. Using basic registers can lead to a race and violation of mutual exclusion. We will revisit synchronization object types when we present their specific inference rules. An array type of a scalar type st is of the form $st[]$. A thread-local type is an array type and the well-formedness conditions enforce that a thread-local object is only indexed by the identifier of the calling thread.

The definitions \mathcal{D} is a sequence of method definitions d . We denote a method name by n , a value by v , a variable by x , a value or variable by u , a thread value by T , a thread variable by t , and a thread value or variable by τ . The method definition $\mathbf{def} \ n_t(x^*) \ s, r$ defines a method named n with thread parameter t and data parameters x^* with the body s and the declared order r . The Dekker description of Figure 1.(a) defines three methods: $init$, $tryLock_1$ and $tryLock_2$. A statement s is either a sequence, a conditional, a method call or a return statement. A condition b is a boolean expression on variables and values. In a method call $c \triangleright x = o.n_\tau(u^*)$, c is the label, x is the return variable, o is the receiving object, n is the method name, τ is the current thread argument, and u^* are the data arguments. The labels of statements are unique in π . Every variable is uniquely bound. An object o is either a single object ϕ or an element of an array $\phi[u]$. In a return statement $c \triangleright \mathbf{return} \ u$, c is the label and u is the returned value or variable. In the appendix [27] § 7, we define the **foreach** iteration statement on sets and maps as a syntactic sugar. As we will see, the semantics of the language supports out-of-order or *relaxed* execution. Any two labels that are left unordered by the description may be reordered in the execution. Data and control dependencies in the method body s impose order between statements. However, the programmer can explicitly require additional orders. The declared program order r of a method definition is a binary relation on the set of labels in the body s . For example, the Dekker

description of Figure 1.(a) declares the orders $W_1 \rightarrow R_2$ and $W_2 \rightarrow R_1$ that are crucial to the correctness. Programming fences is complicated and error-prone. This platform-independent description of the required orders can be used by compilers to optimize fence insertion [4] for different target architectures. The declared order facilitates architecture-independent verification. Further, if the order of two statements that is unnecessary for correctness is changed, the proof stays unchanged.

The client program section \mathcal{P} is of the form $p_0, (p_1 || p_2 || \dots || p_n)$ where p_0 is the initialization program, and $p_1, p_2, \dots,$ and p_n are the parallel programs. For example, the program section of the Dekker description in Figure 1.(a) has two parallel programs that each call one of the two defined try-lock methods. A sequential program p is either a sequence, a conditional or a method call. In a method call $c \triangleright x = n_\tau(u^*)$, n is name of a method that is defined in the method definitions section \mathcal{D} . The object **this** is the default receiver object and is therefore elided in the client calls. We use θ to denote a synchronization object o or the **this** object.

Let \rightarrow_n denote the irreflexive transitive closure of the data and control dependencies and the declared order of method n . Let the program order \rightarrow_π be the irreflexive partial order on $Labels(\pi)$ defined as the union of the following: (1) the initialization order (that orders the labels of p_0 before the labels of parallel programs), (2) the sequential order of the sequential programs p_i , and (3) For each method definition n , the order \rightarrow_n .

LSL uses the following functions that are directly derived from the program description. The names of methods defined in a description are unique. Thus, we define the functions $par1_\pi$ and $par2_\pi$ that map method names to their first and second parameters. Similarly, $tpar_\pi$ maps method names to their thread parameter. As the labels in a description are unique, we define the function obj_π that maps the label of a method call to its receiver object. Similarly, the functions $index_\pi, name_\pi, thread_\pi, arg1_\pi, arg2_\pi$ and $retv_\pi$ map the label of a call to the array index of the receiver object, the name of the method, the thread identifier, first and second arguments and the return variable of the method call. For a return statement, we let $name_\pi$ and $arg1_\pi$ map to the name *return* and the argument of the return statement respectively.

We call the conjunction of all the enclosing if (and else) conditions of a statement, its *enclosing condition*. Let the function $cond_\pi$ map statement labels to their enclosing conditions. Let $Labels_\pi(n)$ denote the set of labels in the body of the method n . Let $Returns_\pi(n)$ denote the set of labels of return statements in the body of n . Let $PreReturns_\pi(c)$ denote the set of labels of the return statements before the statement labeled c in π .

3 Labeled Synchronization Logic

Now, we present our first-order logic to reason about synchronization algorithm descriptions.

$$\begin{array}{c}
\text{XASYM} \\
\frac{\pi, \Gamma \vdash l \prec l'}{\pi, \Gamma \vdash \neg(l' \prec l) \wedge \neg(l' \sim l) \wedge \neg(l' = l)} \\
\\
\text{XTRANS} \\
\frac{\pi, \Gamma \vdash l \prec l' \quad \pi, \Gamma \vdash l' \prec l''}{\pi, \Gamma \vdash l \prec l''} \\
\\
\text{X2TRANS} \\
\frac{\pi, \Gamma \vdash l_1 \prec l_2 \quad \pi, \Gamma \vdash l_2 \sim l_3 \quad \pi, \Gamma \vdash l_3 \prec l_4}{\pi, \Gamma \vdash l_1 \prec l_4} \\
\\
\text{XTOTAL} \\
\frac{\pi, \Gamma \vdash \text{exec}(l) \wedge \text{exec}(l')}{\pi, \Gamma \vdash (l \prec l') \vee (l' \prec l) \vee (l \sim l') \vee (l = l')} \quad \text{(a)} \\
\\
\text{X2X} \\
\frac{\pi, \Gamma \vdash l \prec l'}{\pi, \Gamma \vdash \text{exec}(l) \wedge \text{exec}(l')} \\
\\
\text{X2L} \\
\frac{\mathcal{T}(o) \in LT \quad \pi, \Gamma \vdash l \prec l'}{\pi, \Gamma \vdash \text{obj}(l) = \text{obj}(l') = o} \\
\\
\text{LASYM} \\
\frac{\pi, \Gamma \vdash l \prec_o l'}{\pi, \Gamma \vdash \neg(l' \prec_o l) \wedge \neg(l = l')} \\
\\
\text{LTRANS} \\
\frac{\pi, \Gamma \vdash l \prec_o l' \quad \pi, \Gamma \vdash l' \prec_o l''}{\pi, \Gamma \vdash l \prec_o l''} \\
\\
\text{LTOTAL} \\
\frac{\mathcal{T}(o) \in LT \quad \pi, \Gamma \vdash \text{exec}(l) \wedge \text{exec}(l')}{\pi, \Gamma \vdash \text{obj}(l) = \text{obj}(l') = o} \\
\\
\text{L2X} \\
\frac{\pi, \Gamma \vdash l \prec_o l'}{\pi, \Gamma \vdash \text{exec}(l) \wedge \text{exec}(l') \wedge \text{obj}(l) = \text{obj}(l') = o} \quad \text{(c)}
\end{array}$$

Fig. 3: The Basic Inference Rules. (a) properties of execution orders, (b) linearization orders, and (c) derived rules.

Assertions. We first define the set of dynamic labels that uniquely identify method call instances. As an example, in execution histories for the Dekker algorithm (Figure 1.(a)), we have the two labels L_1 and $L_1'W_1$. The label L_1 refers to a call site in the program section \mathcal{P} . On the other hand, the label $L_1'W_1$ is a call string that refers to the instance of the method call labeled W_1 (in the definitions section \mathcal{D}) that is executed in the body of the client method call labeled L_1 . Thus, a dynamic label l is of the form $\zeta'c$ where the pre-label ζ is either a static label c or no label ϵ . The symbol ϵ is the left identity element of the prefixing operator.

A method can be called several times in the program. To have unique local variable names, every variable (including the parameters) of the method is prefixed by the caller label. The labeled variable $c'x$ denotes the instance of the local variable x in the body of the method call labeled with c . Similarly, the prefixing operator is lifted to expressions over variables. The assertions language of LSL is described by the following grammar.

$$\begin{array}{ll}
e ::= \text{obj}(\ell) \mid \text{name}(\ell) \mid \text{thread}(\ell) \mid & \text{Expression} \\
\text{arg1}(\ell) \mid \text{arg2}(\ell) \mid \text{retv}(\ell) \mid & \\
\text{initOf}(\tau) \mid \text{commitOf}(\tau) \mid & \\
o \mid n \mid \zeta'x \mid v \mid \zeta't \mid T & \\
\mathcal{R} ::= e = e \mid e < e \mid \ell = \ell' \mid c = c \mid & \text{Proposition} \\
\text{exec}(\ell) \mid \ell \prec \ell' \mid \ell \sim \ell' \mid \ell \prec_o \ell' \mid \tau \ll \tau' & \\
\mathcal{A} ::= \mathcal{R} \mid \neg \mathcal{A} \mid \mathcal{A} \wedge \mathcal{A} \mid \forall \ell: \mathcal{A} \mid \forall t: \mathcal{A} & \text{Assertion}
\end{array}$$

Here, c is a program label (such as C_{1t} in the Dekker description of Figure 1.(a)), l is a constant label (such as the dynamic labels L_1 and $L_1'C_{1t}$ for the Dekker description), ℓ is a label variable, T is a thread (or transaction) identifier value, t is a thread identifier variable, τ is a thread value or variable, x is a variable, v is a value, o is an object name, and n is a method name. Expressions use six function symbols. The functions obj , $name$, $thread$, $arg1$, $arg2$ and $retv$ map a label of the program to its object, method name, thread name, first and second argument, and return value. The function $initOf$ maps each transaction to the label of its *init* method call. The function $commitOf$ maps each committed transaction to the label of its *commit* method call. Propositions use seven predicates. The first two are equality ($=$) and integer comparison ($<$). The proposition $exec(\ell)$ states that the method call labeled ℓ is *executed*. The proposition $\ell < \ell'$ asserts that ℓ is *executed before* ℓ' . the proposition $\ell \sim \ell'$ asserts that ℓ is *executed concurrent* to ℓ' . For a linearizable object o , the proposition $\ell_1 \prec_o \ell_2$ states that ℓ_1 is *linearized before* ℓ_2 in the linearization order of o . (As we will describe in the semantics section, any concurrent execution on a linearizable object is equivalent to a correct sequential execution. The total order of method calls in that sequential execution is called the linearization order.) The proposition $\tau \ll \tau'$ asserts that all the labels of thread τ are executed before all the labels of thread τ' . This assertion is used to state that a transaction is executed before another.

An assertion is either a proposition, negation of an assertion, conjunction of two assertions, or existential quantification over labels or transactions. As usual, we can define, disjunction \vee , universal quantification \forall , less than or equal \leq , executes before or equal \preceq , linearized before or equal \preceq_o , thread executed before or equal \preceq as syntactic sugar.

For an algorithm description π , a judgement is of the form $\pi, \Gamma \vdash \mathcal{A}$, where Γ is the context, that is, a list of assertions, and \mathcal{A} is a closed assertion. The judgement is read as for every execution of the program π , if the assertions in Γ hold, then the assertion \mathcal{A} holds. For example, $\pi, \Gamma \vdash l < l'$ says that in every execution of π where Γ holds, the statement labeled l is executed before the statement labeled l' .

Algorithms can be verified modularly. The client program section \mathcal{P} of an algorithm description π can specify general clients. For example a lock algorithm description π (such as Dekker) can be verified separately. Then, the lock object (with its abstracted implementation) can be used to implement a TM. Verification of the description π' of the TM is restricted to the labels of π' and does not involve π .

Inference Rules. We now present the inference rules of LSL. The inference rules can be conceptually divided into four groups. First, the first-order logic rules (which are standard and omitted here). Second, the basic rules that axiomatize the properties of execution and linearization orders and their interdependence (Figure 3). Third, the synchronization object rules that axiomatize the properties of common synchronization object types (Figure 4). Fourth, the inference rules that axiomatize the relation of the algorithm description and

$\frac{\text{AREG} \quad \mathcal{T}(r) = \text{AtomicRegister} \quad \pi, \Gamma \vdash \text{isRead}_r(l_R)}{\pi, \Gamma \vdash \exists \ell_W : \text{isWriter}_r(\ell_W, l_R) \wedge \text{retv}(l_R) = \text{arg1}(\ell_W)}$	$\begin{aligned} \text{isRead}_r(l_R) &\Leftrightarrow \text{exec}(l_R) \wedge \text{obj}(l_R) = r \wedge \text{name}(l_R) = \text{read} \\ \text{isWrite}_r(l_W) &\Leftrightarrow \text{exec}(l_W) \wedge \text{obj}(l_W) = r \wedge \text{name}(l_W) = \text{write} \\ \text{isWriter}_r(l_W, l_R) &\Leftrightarrow \text{isWrite}_r(l_W) \wedge l_W \prec_r l_R \wedge \\ &\quad \forall \ell'_W : \text{isWrite}_r(\ell'_W) \Rightarrow (\ell'_W \preceq_r l_W \vee l_R \prec_r \ell'_W) \end{aligned}$
$\frac{\text{BREG} \quad \mathcal{T}(r) = \text{BasicRegister} \quad \pi, \Gamma \vdash \text{isSingleWriter}(r) \quad \pi, \Gamma \vdash \text{isRead}_r(l_R) \quad \pi, \Gamma \vdash \text{isRaceFree}_r(l_R)}{\pi, \Gamma \vdash \exists \ell_W : \text{isEWriter}_r(\ell_W, l_R) \wedge \text{retv}(l_R) = \text{arg1}(\ell_W)}$	$\begin{aligned} \text{isEWriter}_r(l_W, l_R) &\Leftrightarrow \text{isWrite}_r(l_W) \wedge l_W \prec l_R \wedge \\ &\quad \forall \ell'_W : \text{isWrite}_r(\ell'_W) \Rightarrow (\ell'_W \preceq l_W \vee l_R \prec \ell'_W) \\ \text{isRaceFree}_r(l) &\Leftrightarrow \forall \ell_W : \text{isWrite}_r(\ell_W) \Rightarrow (l_W \prec l \vee l \prec l_W) \\ \text{isSingleWriter}(r) &\Leftrightarrow \forall \ell_w : \text{isWrite}_r(\ell_w) \Rightarrow \text{isRaceFree}_r(\ell_w) \end{aligned}$
$\frac{\text{LOCKUNLOCKPAIR} \quad \mathcal{T}(o) = \text{Lock} \quad \pi, \Gamma \vdash \text{isOwnerRespect}(o) \quad \pi, \Gamma \vdash \text{isLock}_o(l_{a_1}) \quad \pi, \Gamma \vdash \text{isUnlock}_o(l_{r_2}) \quad \pi, \Gamma \vdash l_{a_1} \prec_o l_{r_2}}{\pi, \Gamma \vdash \exists \ell_{r_1}, \ell_{a_2} : \text{isUnlock}_o(\ell_{r_1}) \wedge \text{thread}(\ell_{r_1}) = \text{thread}(l_{a_1}) \wedge \text{isLock}_o(\ell_{a_2}) \wedge \text{thread}(\ell_{a_2}) = \text{thread}(l_{r_2}) \wedge \ell_{r_1} \prec_o \ell_{a_2}}$	$\begin{aligned} \text{isLock}_o(l) &\Leftrightarrow \text{exec}(l) \wedge \text{obj}(l) = o \wedge \text{name}(l) = \text{lock} \\ \text{isUnlock}_o(l) &\Leftrightarrow \text{exec}(l) \wedge \text{obj}(l) = o \wedge \text{name}(l) = \text{unlock} \\ \text{isOwnerRespect}(o) &\Leftrightarrow \forall \ell : \text{isUnlock}_o(\ell) \Rightarrow \exists \ell' : \\ &\quad \text{isLock}_o(\ell') \wedge \text{thread}(\ell') = \text{thread}(\ell) \wedge \\ &\quad \ell' \prec \ell \wedge \\ &\quad \forall \ell'' : \\ &\quad (\text{isUnLock}_o(\ell'') \wedge \text{thread}(\ell'') = \text{thread}(\ell)) \Rightarrow \\ &\quad \ell'' \prec \ell' \vee \ell \preceq \ell'' \end{aligned}$
$\frac{\text{COUNTSEQ} \quad \mathcal{T}(o) = \text{SCounter} \quad \pi, \Gamma \vdash \text{exec}(l_1) \wedge \text{obj}(l_1) = o \wedge \text{name}(l_1) = \text{iaf} \quad \pi, \Gamma \vdash \text{exec}(l_2) \wedge \text{obj}(l_2) = o \quad \pi, \Gamma \vdash \text{retv}(l_1) < \text{retv}(l_2)}{\pi, \Gamma \vdash l_1 \prec_o l_2}$	

Fig. 4: Synchronization Object Inference Rules. Four of the rules for atomic and basic registers, lock and strong counter.

the execution. We showcase a few rules. The full set of rules for the common synchronization objects is available in the appendix [27] § 9.

Figure 3 represents the set of basic inference rules that intuitively capture the properties of execution and linearization orders and their relation. The rule XASYM states the asymmetry property of the execution order. If a method call is executed before another method call, then the latter is not executed before the former and they are not executed concurrently. The rule XTRANS states the transitivity property of the execution order. The rule X2TRANS states the transitivity of the sequence of precedence, concurrency and precedence execution relations. If l_1 is executed before l_2 , l_2 is executed concurrent to l_3 and l_3 is executed before l_4 , then l_1 is executed before l_4 . The rule XTOTAL states the totality property of the precedence and concurrency execution relations. Every pair of executed method calls either execute in order or concurrently. The rule X2L states the *real-time-preservation* property of linearization orders: The execution order of two method calls on a linearizable object (specified by $\mathcal{T}(o) \in LT$) is preserved in the linearization order. The rule LASYM states the asymmetry

property of linearization orders. If a method call is linearized before another one, then the latter is not linearized before the former. The rule LTRANS states the transitivity property of linearization orders. The rule LTOTAL states the totality property of linearization orders. Every two executed method calls on a linearizable object are ordered in its linearization order. The two derived rules X2X and L2X can be established by an inductive reasoning on the length of the proof of $l \prec_o l'$. The rule X2X states that if a method call is executed before another one, then clearly both are executed. The rule L2X states that if a method call is linearized before another one, then clearly both are executed.

Now let us look at a few synchronization object rules in Figure 4. First, the rule AREG states that for every read method call l_R on an atomic register, there is a write method call ℓ_W on it that writes the same value that l_R reads and ℓ_W is the last write method call that is linearized before l_R . Second, the rule BREG states that if a basic register r is single-writer, for every race-free read method call l_R on r , there is a write method call ℓ_W on r that writes the same value that l_R reads and ℓ_W is the last write method call that is executed before l_R . A register r is single-writer if and only if every pair of write method calls on it are ordered in the execution order or in other words, every write method call on it is race-free. A method call r is race-free if and only if there is no *write* method call on r that executes concurrent to it. The rules AREG and BREG model Lamport's notion of atomic and safe registers [25]. Third, the rule LOCKUNLOCKPAIR states the *lock-unlock-pair* property: if ownership of a lock object o is respected and a *lock* method call on o by a thread τ_1 is linearized before an *unlock* method call on o by a thread τ_2 , then an *unlock* method call on o by τ_1 is linearized before a *lock* method call on o by τ_2 . The rule is derived from the fact that if the ownership of a lock is respected, its linearization order is a sequence of matching pairs of *lock* and *unlock* method calls. Intuitively, ownership for a lock o is respected, if and only if every thread unlocks o only if it has already locked o and has not unlocked o since then. Fourth, the rule COUNTSEQ states the *count-sequence* property: for a strong counter object co , if the return value of an *iaf* (inc-and-fetch) method call on co is less than the return value of another method call on co , then the former is linearized before the latter. The rule is derived from the fact that the return values of method calls in the linearization order of a strong counter is non-decreasing.

$$\begin{array}{c}
\text{P2X} \\
\frac{c_1 \rightarrow_{\pi} c_2 \quad \pi, \Gamma \vdash \text{exec}(\zeta'c_1) \quad \pi, \Gamma \vdash \text{exec}(\zeta'c_2)}{\pi, \Gamma \vdash \zeta'c_1 \prec \zeta'c_2} \\
\text{CALLEE} \\
\frac{c' \in \text{Labels}_{\pi}(n) \quad \text{tpar}_{\pi}(n) = t \quad \text{par}1_{\pi}(n) = x \quad \pi, \Gamma \vdash \text{exec}(\zeta'c')}{\pi, \Gamma \vdash \neg(\zeta = \epsilon) \wedge \text{exec}(\zeta) \wedge \text{obj}(\zeta) = \mathbf{this} \wedge \text{name}(\zeta) = n \wedge \text{thread}(\zeta) = \zeta't \wedge \text{arg}(\zeta) = \zeta'x}
\end{array}$$

The rules presented above refer to the algorithm description. The rule P2X states the *program-order-preservation* property: the program order is preserved in the execution order. If the algorithm description requires a method call l_1 to be ordered before another method call l_2 , the order is preserved in the execution

of them from any call site ζ . That is if $\zeta'l_1$ and $\zeta'l_2$ are executed, then $\zeta'l_1$ is executed before $\zeta'l_2$. The rule CALLEE states that if a method call c' in the body of the caller method call ζ is executed, then the later is also executed, is a *this* method call and its parameters and arguments are equal.

We define the semantics $\llbracket \pi \rrbracket$ of a description π as a set of execution histories \mathcal{X} . A specification π models an assertion \mathcal{A} , written as $\pi \models \mathcal{A}$, iff every execution \mathcal{X} of π models \mathcal{A} written as $\mathcal{X} \models \mathcal{A}$. The soundness theorem states that if LSL deduces an assertion \mathcal{A} for a description π using valid assumptions for π , then the deduced assertion \mathcal{A} is valid for π as well. Formally, $\forall \pi, \mathcal{A}: (\pi, \Gamma \vdash \mathcal{A} \wedge \pi \models \Gamma) \Rightarrow (\pi \models \mathcal{A})$. The semantics and soundness of LSL is available in the appendix [27] § 8 & 11.

4 TM Verification

We now state the correctness of TM algorithms as an LSL assertion and apply LSL to prove the correctness of the TL2 [9] algorithm. The challenge is to verify that any concurrent execution of any set of well-formed transactions on TL2 is opaque. Markability factors out a large part of the proof, allows specification of the critical points of the algorithm and reduces verification to separate proof obligations about the order of these points. LSL inference rules can be easily used to prove the obligations based on the validation checks in the algorithm.

Transactional Memory. A TM object encapsulates a set of locations and provides four methods $init_t()$, $read_t(i)$, $write_t(i, v)$, and $commit_t()$. A well-formed transaction first calls $init_t()$ and then calls a sequence of $read_t(i)$ and $write_t(i, v)$ methods, and finally calls $commit_t()$. The method $commit_t()$ tries to commit transaction t and returns \mathbb{C} (if it is successful). A TM object should detect if an inconsistency is about to happen between two concurrent transactions and should at least abort one of them. All methods may return abort \mathbb{A} and terminate the transaction.

Correctness Assertion. We presented a decomposition called *markability* [30] of the correctness condition *opacity* [14]. Markability restates opacity in terms of three intuitive invariants. We *state Markability as an assertion in LSL*. The markability assertion $isMarking(\sqsubseteq)$ is parametric with the marking relation \sqsubseteq . The marking assertion is available in the appendix [27] § 13. We briefly explain markability. A TM algorithm is markable iff there exists a *marking* relation for it that is *write-observant*, *read-preserving*, and *real-time-preserving*.

A marking is a relation on the union of the transactions and the read method calls. We can think of the marking relation as the union of a collection of orders: (1) The *effect order*: The effect order is a total order of the transactions. The effect order represents the order in which the transactions appear to take effect, that is, the order that justifies the correctness of the execution. (2) The *access orders*: Let writers of location i be the committed transactions that have write method call(s) to i . Consider a read method call l_R that reads from a location i and doesn't abort. For each such l_R , the access order is an antisymmetric

\mathcal{T}_{TL2} :	
reg : BasicReg [],	$rver$: ThreadLocal BasicReg ,
ver : AtomicReg [],	$rset$: ThreadLocal BasicSet ,
$lock$: TryLock [],	$wset$: ThreadLocal BasicMap ,
$clock$: SCounter ,	$lset$: ThreadLocal BasicSet
\mathcal{D}_{TL2} :	
def $init_t()$	def $commit_t()$
$I01 \triangleright snap = clock.read(),$	$C01 \triangleright$ foreach ($i \in wset[t]$)
$I02 \triangleright rver[t].write(snap),$	$C02_i \triangleright l' = lock[i].trylock(),$
$I03 \triangleright$ return	if ($\neg l'$)
def $read_t(i)$	$C03_i \triangleright lset[t].add(i)$
$R01 \triangleright pv = wset[t].get(i),$	else
if ($pv \neq \perp$)	$C04_j \triangleright$ foreach ($j \in lset[t]$)
$R02 \triangleright$ return $pv,$	$C05_{i,j} \triangleright lock[j].unlock(),$
	$C06_i \triangleright$ return $\mathbb{A},$
$R03 \triangleright t_1 = ver[i].read(),$	$C07 \triangleright wver = clock.iaf(),$
$R04 \triangleright v = reg[i].read(),$	$C08 \triangleright s = rver[t].read(),$
$R05 \triangleright l = lock[i].read(),$	if ($wver \neq s + 1$)
$R06 \triangleright t_2 = ver[i].read(),$	$C09 \triangleright$ foreach ($i \in rset[t]$)
$R07 \triangleright s = rver[t].read(),$	$C10_i \triangleright l = lock[i].read(),$
if ($\neg(\neg l \wedge t_1 = t_2$	$C11_i \triangleright cver = ver[i].read(),$
$\wedge t_2 \leq s)$)	if ($\neg(\neg l \wedge cver \leq s)$)
$R08 \triangleright$ return $\mathbb{A},$	foreach ($j \in lset[t]$)
$R09 \triangleright rset[t].add(i),$	$C12_i \triangleright lock[j].unlock(),$
$R10 \triangleright$ return $v,$	$C13_{i,j} \triangleright$
$\{R03 \rightarrow R04, R04 \rightarrow R05,$	$C14_i \triangleright$ return $\mathbb{A},$
$R05 \rightarrow R06\},$	
def $write_t(i, v)$	$C15 \triangleright$ foreach ($(i, v) \in wset[t]$)
$W01 \triangleright wset[t].put(i, v),$	$C16_i \triangleright reg[i].write(v),$
$W02 \triangleright$ return $ok,$	$C17_i \triangleright ver[i].write(wver),$
	$C18_i \triangleright lock[i].unlock(),$
	$C19 \triangleright$ return $\mathbb{C},$
	$\{C01 \rightarrow C07, C10 \rightarrow C11,$
	$C09 \rightarrow C15, C16 \rightarrow C17,$
	$C17 \rightarrow C18\},$
\mathcal{P} : $tran_0, (tran_1 \parallel tran_2 \parallel \dots \parallel tran_n)$	

Fig. 5: TL2 Algorithm Description $\pi_{TL2} = (\mathcal{T}_{TL2}, \mathcal{D}_{TL2}, \mathcal{P})$

relation that orders l_R and every writer of i . The access order represents where l_R 's access to location i has happened between the accesses by the writers of i .

Write-observation requires that each read method call should read the most current value. Read-preservation requires that the location read by a read method call is not overwritten between the read accesses the location and the transaction takes effect. The real-time-preservation condition requires that the marking relation preserves the real-time order of transactions.

Algorithm Description. We have represented the TL2 algorithm [9] in our description language in Figure 5. The description π_{TL2} provides implementations of the four TM methods $init_t()$, $read_t(i)$, $write_t(i, v)$, and $commit_t()$ in the

definitions section \mathcal{D} . The program in section \mathcal{P} represents well-formed general client transactions. A client program first runs $tran_0$ to initialize the shared variables and then concurrently runs n well-formed transactions $tran_1, \dots, tran_n$. A well-formed transaction t executes $init_t()$, then a sequence of $read_t(i)$ and $write_t(i, v)$ calls and finally a $commit_t()$; it finishes if any call returns \mathbb{A} .

TL2 is a subtle algorithm. We briefly review how it works. TL2 uses the basic register $reg[i]$ to store the value of a location i . The algorithm reads $reg[i]$ at $R04$ and writes to $reg[i]$ at $C16$. Additionally, TL2 uses synchronization objects to help abort executions that would violate consistency. The idea is to give the value written in $reg[i]$ a *version number* that is stored in the $ver[i]$ register. TL2 uses a strong counter $clock$, whose value increases monotonically, to create such version numbers. Specifically, TL2 takes snapshots of $clock$ both at $I01$ when a transaction starts and at $C07$ (with an increment-and-fetch operation, abbreviated *iaf*) during commit. TL2 validates the versions of read values before completing both the read and commit methods.

Verification. We state the marking relation for the TL2 algorithm as an assertion in LSL as follows:

The marking \sqsubseteq is the reflexive closure of \sqsubset . The relation \sqsubset is defined as follows:

$$\begin{aligned} \forall t, t': t \sqsubset t' &\Leftrightarrow Eff(t) \prec_{clock} Eff(t') \\ \forall \ell_R, t: isTRead(\ell_R) \wedge isTWriter_i(t) &\Rightarrow \\ & \text{Let } i = arg1(\ell_R): \\ & t \sqsubset \ell_R \Leftrightarrow writeAcc_i(t) \prec readAcc(\ell_R) \\ & \ell_R \sqsubset t \Leftrightarrow readAcc(\ell_R) \prec writeAcc_i(t) \end{aligned}$$

where

$$\begin{aligned} Eff(\tau) &= \begin{cases} initOf(\tau)'I01 & \text{if } isAborted(\tau) \\ commitOf(\tau)'C07 & \text{if } isCommitted(\tau) \end{cases} \\ readAcc(\ell_R) &= \ell_R'R04 \\ writeAcc_i(\tau) &= commitOf(\tau)'C16_i \end{aligned}$$

Intuitively, the effect order of transactions is the linearization order of their calls to $clock$ at $I01$ and $C07$. The access order of read operations and writer transactions to location i is the execution order of their access to the $reg[i]$ register at $R04$ and $C16$. The following theorem states that the relation \sqsubseteq defined above is a marking relation for TL2. The assertions I_0 are the properties of well-formed client transactions. We have mechanically checked the proof in PVS. The PVS theories for TL2 and Dekker are available [26].

Theorem 2 (TL2 Correctness). $\pi_{TL2}, I_0 \vdash isMarking(\sqsubseteq)$.

5 Related Works and Conclusion

Manovit et al. [35] applied random testing to the TCC TM system. Lourenco et al. [32] reported several bugs during the porting of the TL2 algorithm. Given a TM algorithm and a bug pattern, our previous work [29] constructs a bug trace if the algorithm is prone to the bug pattern. It showed the incorrectness of algorithms that were deemed verified.

Although testing can find bugs, it does not prove their absence. To verify the correctness of TM algorithms, researchers have employed model checking

and theorem proving. Model checkers from Cohen et al. [5,6], and Guerraoui et al. [15,17,16] are the pioneering approach to verification of TM. Subsequently, the same approach was taken by O’Leary et al. [39] and Baek et al. [3]. Model checking can automate the verification process but it has been dependent on assuming properties about the TM algorithm and only scalable to a finite number of threads and locations or simplified algorithms. Later, Emmi et al. [12] tried to infer algorithm invariants from small number of threads and memory locations. However, it worked on simplified algorithms due to scalability issues.

Attiya et al. [2] proved that opacity is sufficient for observational refinement of high-level atomic block semantics. Our previous work [30] showed the equivalence of opacity and markability and an informal proof of correctness for TL2. Koskinen and Parkinson presented a semantic model of serializability based on pulls from and pushes to an abstract shared log. Khyzha et al. [24] extended opacity to account for non-transactional accesses. In contrast to the current work, these works consider the correctness criteria, include only informal or non-mechanized proofs and do not include a logic and its soundness.

Singh [40] developed a runtime verification tool for TM algorithms. Although the tool is optimized with sound approximation techniques, the runtime overhead is still not negligible. Our previous work [28] presented a machine-checked theorem proving framework based on simulation between specifications and implementations [33,18] represented as IOA [34] and verified the NORec algorithm [7]. Doherty et al. [11] adopted the same approach and proved the correctness of a pessimistic TM algorithm [36]. In follow-up works, Derrick et. al and Armstrong et al. [8,1] simplified their simulation proofs by first model checking or proving the linearizability of the TM algorithm. In contrast to LSL that can reason about the algorithm description, these works require the algorithm to be translated to a transition system. In addition, they do not feature a logic.

To the best of our knowledge, LSL is the first logic that is applied to verification of transaction algorithms. In particular, it provides assertions for inter-thread execution and linearization orders that can directly capture the marking relation and the markability condition. Leveraging the proof of sufficiency of markability for opacity, verification of opacity is reduced to separate markability conditions that can be proved by the logic based on the validation checks in the algorithm. Logics based on concurrent separation logic [38,20] and rely-guarantee reasoning [21] such as RGSep [42], LRG [13,31], FCSL [37], GPS [41] and Iris [22,23] require the specification of inter-thread relations as complicated global rely and guarantee conditions. Further, they need auxiliary variables even for the simple Dekker algorithm which may obscure the underlying design intuitions of the algorithms.

Conclusion. We presented a logic that supports syntactic reasoning about synchronization algorithm descriptions and features novel assertions and inference rules for execution and linearization orders. These assertions enable capturing critical orders between concurrent operations and in particular markability orders between transactions. We proved the soundness of the logic and used it to machine-check a significant proof of TL2.

References

1. Alasdair Armstrong, Brijesh Dongol, and Simon Doherty. Proving opacity via linearizability: a sound and complete method. In *International Conference on Formal Techniques for Distributed Objects, Components, and Systems*, pages 50–66. Springer, 2017.
2. Hagit Attiya, Alexey Gotsman, Sandeep Hans, and Noam Rinetzky. A programming language perspective on transactional memory consistency. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing*, PODC '13, pages 309–318, New York, NY, USA, 2013. ACM.
3. Woongki Baek, N. Bronson, C. Kozyrakis, and K. Olukotun. Implementing and evaluating a model checker for transactional memory systems. In *Engineering of Complex Computer Systems (ICECCS), 2010 15th IEEE International Conference on*, pages 117–126, 2010.
4. John Bender, Mohsen Lesani, and Jens Palsberg. Declarative fence insertion. In *OOPSLA '15*, pages 367–385, 2015.
5. Ariel Cohen, John W. O’Leary, Amir Pnueli, Mark R. Tuttle, and Lenore D. Zuck. Verifying correctness of transactional memories. In *FMCAD*, 2007.
6. Ariel Cohen, Amir Pnueli, and Lenore D. Zuck. Mechanical verification of transactional memories with non-transactional memory accesses. In *CAV*, 2008.
7. Luke Dalessandro, Michael F. Spear, and Michael L. Scott. Norec: streamlining stm by abolishing ownership records. In *PPoPP*, 2010.
8. John Derrick, Brijesh Dongol, Gerhard Schellhorn, Oleg Travkin, and Heike Wehrheim. Verifying opacity of a transactional mutex lock. In Nikolaj Bjørner and Frank de Boer, editors, *FM 2015: Formal Methods*, pages 161–177, Cham, 2015. Springer International Publishing.
9. D. Dice, O. Shalev, and N. Shavit. Transactional locking II. In *DISC (LNCS 4167)*, 2006.
10. Edsger Wybe Dijkstra. Cooperating sequential processes, technical report ewd-123. 1965.
11. Simon Doherty, Brijesh Dongol, John Derrick, Gerhard Schellhorn, and Heike Wehrheim. Proving opacity of a pessimistic stm. In *Leibniz International Proceedings in Informatics*, volume 70, pages 35–1. Dagstuhl Publishing, 2017.
12. Michael Emmi, Rupak Majumdar, and Roman Manevich. Parameterized verification of transactional memories. In *PLDI*, 2010.
13. Xinyu Feng. Local rely-guarantee reasoning. In *POPL '09*, 2009.
14. R. Guerraoui and M. Kapalka. On the correctness of transactional memory. In *PPOPP*, 2008.
15. Rachid Guerraoui, Thomas A. Henzinger, Barbara Jobstmann, and Vasu Singh. Model checking transactional memories. In *ACM SIGPLAN Conference on Programming Languages Design and Implementation*, pages 372–382, 2008.
16. Rachid Guerraoui, Thomas A. Henzinger, and Vasu Singh. Software transactional memory on relaxed memory models. In *CAV*, 2009.
17. Rachid Guerraoui, Thomas A. Henzinger, and Vasu Singh. Model checking transactional memories. *Distributed Computing*, 2010.
18. Chris Hawblitzel, Erez Petrank, Shaz Qadeer, and Serdar Tasiran. Automated and modular refinement reasoning for concurrent programs. In *International Conference on Computer Aided Verification*, pages 449–465. Springer, 2015.
19. M. P. Herlihy and J. M. Wing. Linearizability: a correctness condition for concurrent objects. *TOPLAS*, 12(3):463–492, July 1990.

20. Aquinas Hobor, Andrew W. Appel, and Francesco Zappa Nardelli. Oracle semantics for concurrent separation logic. In *ESOP*, 2008.
21. Cliff B. Jones. Specification and design of (parallel) programs. In *Information Processing 83*, volume 9, pages 321–332, 1983.
22. Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. Higher-order ghost state. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming*, ICFP 2016, pages 256–269, New York, NY, USA, 2016. ACM.
23. Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. Iris: Monoids and invariants as an orthogonal basis for concurrent reasoning. In *POPL '15*, 2015.
24. Artem Khyzha, Hagit Attiya, Alexey Gotsman, and Noam Rinetzkky. Safe privatization in transactional memory. In *Proceedings of the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 233–245. ACM, 2018.
25. Leslie Lamport. On interprocess communication. part i: Basic formalism. *Distributed Computing*, 1(2):77–85, 1986.
26. Mohsen Lesani. PVS Proof Theories. <http://www.cs.ucr.edu/%7Elesani/companion/nfm19/PVSTheories.tar.gz>, 2018.
27. Mohsen Lesani. Submission appendix. <http://www.cs.ucr.edu/%7Elesani/companion/nfm19/Appendix.pdf>, 2018.
28. Mohsen Lesani, Victor Luchangco, and Mark Moir. A framework for formally verifying software transactional memory algorithms. In *CONCUR*, 2012.
29. Mohsen Lesani and Jens Palsberg. Proving non-opacity. In *DISC, (LNCS 8205)*, 2013.
30. Mohsen Lesani and Jens Palsberg. Decomposing opacity. In *Distributed Computing*, volume 8784 of *Lecture Notes in Computer Science*, pages 391–405. 2014.
31. Hongjin Liang and Xinyu Feng. Modular verification of linearizability with non-fixed linearization points. In *PLDI '13*, pages 459–470, 2013.
32. João Lourenço and Gonçalo Cunha. Testing patterns for software transactional memory engines. In *Proceedings of the 2007 ACM Workshop on Parallel and Distributed Systems: Testing and Debugging*, PADTAD '07, pages 36–42, New York, NY, USA, 2007. ACM.
33. Nancy Lynch and Frits Vaandrager. Forward and backward simulations for timing-based systems. In J. W. de Bakker, W. P. de Roever, C. Huizing, and G. Rozenberg, editors, *Proceedings of Real-Time: Theory in Practice (REX Workshop, Mook, The Netherlands, June 1991)*, pages 397–446. Springer-Verlag (LNCS 600), 1992.
34. Nancy A. Lynch and Mark R. Tuttle. An introduction to input/output automata. *CWI Quarterly*, 2, 1989.
35. Chaiyasit Manovit, Sudheendra Hangal, Hassan Chafi, Austen McDonald, Christos Kozyrakis, and Kunle Olukotun. Testing implementations of transactional memory. In *Proceedings of the 15th International Conference on Parallel Architectures and Compilation Techniques*, PACT '06, pages 134–143, New York, NY, USA, 2006. ACM.
36. Alexander Matveev and Nir Shavit. Towards a fully pessimistic stm model, 2012.
37. Aleksandar Nanevski, Ruy Ley-Wild, Ilya Sergey, and Germán Andrés Delbianco. Communicating state transition systems for fine-grained concurrent resources. In *Proceedings of the 23rd European Symposium on Programming Languages and Systems - Volume 8410*, pages 290–310, New York, NY, USA, 2014. Springer-Verlag New York, Inc.

38. Peter W. O'Hearn. Resources, concurrency, and local reasoning. *Theor. Comput. Sci.*, 375(1-3):271–307, 2007.
39. J. O'Leary, B. Saha, and Mark R. Tuttle. Model checking transactional memory with spin. In *Distributed Computing Systems, 2009. ICDCS '09. 29th IEEE International Conference on*, pages 335–342, 2009.
40. Vasu Singh. Runtime verification for software transactional memories. In *Proceedings of the First International Conference on Runtime Verification, RV'10*, pages 421–435, Berlin, Heidelberg, 2010. Springer-Verlag.
41. Aaron Turon, Viktor Vafeiadis, and Derek Dreyer. Gps: Navigating weak memory with ghosts, protocols, and separation. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA '14*, pages 691–707, New York, NY, USA, 2014. ACM.
42. Viktor Vafeiadis and Matthew Parkinson. A marriage of rely/guarantee and separation logic. In *CONCUR*. 2007.