
LOGICAL CONCURRENCY CONTROL FROM SEQUENTIAL PROOFS

JYOTIRMOY DESHMUKH^a, G. RAMALINGAM^b, VENKATESH-PRASAD RANGANATH^c,
AND KAPIL VASWANI^d

^a University of Texas at Austin
e-mail address: jyotirmoy@cerc.utexas.edu

^{b,c,d} Microsoft Research, India
e-mail address: grama@microsoft.com, rvprasad@microsoft.com, kapilv@microsoft.com

ABSTRACT. We are interested in identifying and enforcing the *isolation requirements* of a concurrent program, i.e., concurrency control that ensures that the program meets its specification. The thesis of this paper is that this can be done systematically starting from a sequential proof, i.e., a proof of correctness of the program in the absence of concurrent interleavings. We illustrate our thesis by presenting a solution to the problem of making a sequential library thread-safe for concurrent clients. We consider a sequential library annotated with assertions along with a proof that these assertions hold in a sequential execution. We show how we can use the proof to derive concurrency control that ensures that any execution of the library methods, when invoked by concurrent clients, satisfies the same assertions. We also present an extension to guarantee that the library methods are linearizable or atomic.

1. INTRODUCTION

A key challenge in concurrent programming is identifying and enforcing the *isolation requirements* of a program: determining what constitutes undesirable interference between different threads and implementing concurrency control mechanisms that prevent this. In this paper, we show how a solution to this problem can be obtained systematically from a *sequential proof*: a proof that the program satisfies a specification in the absence of concurrent interleaving.

Problem Setting. We illustrate our thesis by considering the concrete problem of making a sequential library safe for concurrent clients. Informally, given a sequential library that works correctly when invoked by any sequential client, we show how to synthesize concurrency control code for the library that ensures that it will work correctly when invoked by any concurrent client.

1998 ACM Subject Classification: D.1.3, D.2.4, F.3.1.

Key words and phrases: concurrency control, program synthesis.

^a Work done while at Microsoft Research India.

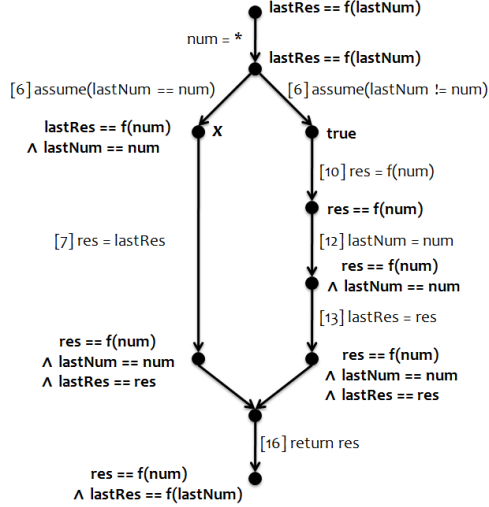
Part I: Ensuring Assertions In Concurrent Executions. Consider the example in Figure 1(a). The library consists of one procedure `Compute`, which applies an expensive function f to an input variable `num`. As a performance optimization, the implementation caches the last input and result. If the current input matches the last input, the last computed result is returned.

```

1: int lastNum = 0;
2: int lastRes = f(0);
3: /* @returns f(num) */
4: Compute(num) {
5:   /* acquire (1); */
6:   if(lastNum==num) {
7:     res = lastRes;
8:   } else {
9:     /* release (1); */
10:    res = f(num);
11:    /* acquire (1); */
12:    lastNum = num;
13:    lastRes = res;
14:  }
15:  /* release (1); */
16:  return res;
17: }

```

(a)



(b)

Figure 1: (a) Procedure `Compute` (excluding Lines 5,9,11,15) applies a (side-effect free) function f to a parameter `num` and caches the result for later invocations. Lines 5,9,11,15 contain a lock-based concurrency control generated by our technique. (b) The control-flow graph of `Compute`, its edges labeled by statements of `Compute` and nodes labeled by proof assertions.

This procedure works correctly when used by a sequential client, but not in the presence of concurrent procedure invocations. E.g., consider an invocation of `Compute(5)` followed by concurrent invocations of `Compute(5)` and `Compute(7)`. Assume that the second invocation of `Compute(5)` evaluates the condition in Line 6, and proceeds to Line 7. Assume a context switch occurs at this point, and the invocation of `Compute(7)` executes completely, overwriting `lastRes` in Line 13. Now, when the invocation of `Compute(5)` resumes, it will erroneously return the (changed) value of `lastRes`.

In this paper, we present a technique that can detect the potential for such interference and synthesize concurrency control to prevent the same. The (lock-based) solution synthesized by our technique for the above example is shown (as comments) in Lines 5, 9, 11, and 15 in Figure 1(a). With this concurrency control, the example works correctly even for concurrent procedure invocations while permitting threads to perform the expensive function f concurrently.

The Formal Problem. Formally, we assume that the correctness criterion for the library is specified as a set of assertions and that the library satisfies these assertions in any execution

of any sequential client. *Our goal is to ensure that any execution of the library with any concurrent client also satisfies the given assertions.* For our running example in Figure 1(a), Line 3 specifies the desired functionality for procedure `Compute`: `Compute` returns the value $f(\text{num})$.

Logical Concurrency Control From Proofs. A key challenge in coming up with concurrency control is determining what interleavings between threads are safe. A conservative solution may reduce concurrency by preventing correct interleavings. An aggressive solution may enable more concurrency but introduce bugs.

The fundamental thesis we explore is the following: a proof that a code fragment satisfies certain assertions in a sequential execution precisely identifies the properties relied on by the code at different points in execution; hence, such a sequential proof clearly identifies what concurrent interference can be permitted; thus, a correct concurrency control can be systematically (and even automatically) derived from such a proof.

We now provide an informal overview of our approach by illustrating it for our running example. Figure 1(b) presents a proof of correctness for our running example (in a sequential setting). The program is presented as a control-flow graph, with its edges representing program statements. (The statement “`num = *`” at the entry edge indicates that the initial value of parameter `num` is unknown.) A proof consists of an invariant $\mu(u)$ attached to every vertex u in the control-flow graph (as illustrated in the figure) such that: (a) for every edge $u \rightarrow v$ labelled with a statement s , execution of s in a state satisfying $\mu(u)$ is guaranteed to produce a state satisfying $\mu(v)$, (b) The invariant $\mu(\text{entry})$ attached to the entry vertex is satisfied by the initial state and is implied by the invariant $\mu(\text{exit})$ attached to the exit vertex, and (c) for every edge $u \rightarrow v$ annotated with an assertion φ , we have $\mu(u) \Rightarrow \varphi$. Condition (b) ensures that the proof is valid over any *sequence of* executions of the procedure.

The invariant $\mu(u)$ at vertex u indicates the property required (by the proof) to hold at u to ensure that a sequential execution satisfies all assertions of the library. We can reinterpret this in a concurrent setting as follows: when a thread t_1 is at point u , it can tolerate changes to the state by another thread t_2 as long as the invariant $\mu(u)$ continues to hold from t_1 ’s perspective; however, if another thread t_2 were to change the state such that t_1 ’s invariant $\mu(u)$ is broken, then the continued execution by t_1 may fail to satisfy the desired assertions.

Consider the proof in Figure 1(b). The vertex labeled x in the figure corresponds to the point before the execution of Line 7. The invariant attached to x indicates that the proof of correctness depends on the condition $\text{lastRes} == f(\text{num})$ being true at x . The execution of Line 10 by another thread will not invalidate this condition. But, the execution of Line 13 by another thread can potentially invalidate this condition. Thus, we infer that, when one thread is at point x , an execution of Line 13 by another thread should be avoided.

We prevent the execution of a statement s by one thread when another thread is at a program point u (if s might invalidate a predicate p that is required at u) as follows. We introduce a lock ℓ_p corresponding to p , and ensure that every thread holds ℓ_p at u and ensure that every thread holds ℓ_p when executing s .

Our algorithm does this as follows. From the invariant $\mu(u)$ at vertex u , we compute a set of predicates $\text{pm}(u)$. (For now, think of $\mu(u)$ as the conjunction of predicates in $\text{pm}(u)$.) $\text{pm}(u)$ represents the set of predicates required at u . For any edge $u \rightarrow v$, any predicate p that is in $\text{pm}(v) \setminus \text{pm}(u)$ is required at v but not at u . Hence, we acquire the lock for p along

this edge. Dually, for any predicate that is required at u but not at v , we release the lock along the edge. As a special case, we acquire (release) the locks for all predicates in $\text{pm}(u)$ at procedure entry (exit) when u is the procedure entry (exit) vertex. Finally, if the execution of the statement on edge $u \rightarrow v$ can invalidate a predicate p that is required at some vertex, we acquire and release the corresponding lock before and after the statement (unless it is already a required predicate at u or v). Note that our approach conservatively assumes that any two statements in the library may be simultaneously executed by different threads. If an analysis can identify that certain statements cannot be simultaneously executed (by different threads), this information can be exploited to improve the solution, but this is beyond the scope of this paper.

Our algorithm ensures that the locking scheme does not lead to deadlocks by merging locks when necessary, as described later. Finally, we optimize the synthesized solution using a few simple techniques. E.g., in our example whenever the lock corresponding to `lastRes == res` is held, the lock for `lastNum == num` is also held. Hence, the first lock is redundant and can be eliminated.

Figure 1 shows the resulting library with the concurrency control we synthesize. This implementation satisfies its specification even in a concurrent setting. The synthesized solution permits a high degree to concurrency since it allows multiple threads to compute f concurrently. A more conservative but correct locking scheme would hold the lock during the entire procedure execution.

A distinguishing aspect of our algorithm is that it requires only local reasoning and not reasoning about interleaved executions, as is common with many analyses of concurrent programs. Note that the synthesized solution depends on the proof used. Different proofs can potentially yield different concurrency control solutions (all correct, but with potentially different performance).

We note that our approach has a close connection to the Owicki-Gries [18] approach to computing proofs for concurrent programs. The Owicki-Gries approach shows how the proofs for two statements can be composed into a proof for the concurrent composition of the statements *if the two statements do not interfere with each other*. Our approach detects potential interference between statements and inserts concurrency-control so that the interference does not occur (permitting a safe concurrent composition of the statements).

Implementation. We have implemented our algorithm, using an existing software model checker to generate the sequential proofs. We used the tool to successfully synthesize concurrency control for several small examples. The synthesized solutions are equivalent to those an expert programmer would use.

Part II: Ensuring Linearizability. The above approach can be used to ensure that concurrent executions guarantee desired safety properties, preserve data-structure invariants, and meet specifications (e.g., given as a precondition/postcondition pair). Library implementors may, however, wish to provide the stronger guarantee of linearizability with respect to the sequential specification: *any concurrent execution of a procedure is guaranteed to satisfy its specification and appears to take effect instantaneously at some point during its execution*. In the second half of the paper, we show how the techniques sketched above can be extended to guarantee linearizability.

Contributions. We present a technique for synthesizing concurrency control for a library (e.g., developed for use by a single-threaded client) to make it safe for use by concurrent clients. However, we believe that the key idea we present – a technique for identifying and realizing isolation requirements from a sequential proof – can be used in other contexts as well (e.g., in the context of a whole program consisting of multiple threads, each with its own assertions and sequential proofs).

Sometimes a library designer may choose to delegate the responsibility for concurrency control to the clients of the library and not make the library thread-safe¹. Alternatively, library implementers could choose to make the execution of a library method appear atomic by wrapping it in a transaction and executing it in an STM (assuming this is feasible). These are valid options but orthogonal to the point of this paper. Typically, a program is a software stack, with each level serving as a library. Passing the buck, with regards to concurrency control, has to stop somewhere. Somewhere in the stack, the developer needs to decide what degree of isolation is required by the program; otherwise, we would end up with a program consisting of multiple threads where we require every thread’s execution to appear atomic, which could be rather severe and restrict concurrency needlessly. The ideas in this paper provide a systematic method for determining the isolation requirements. While we illustrate the idea in a simplified setting, it should ideally be used at the appropriate level of the software stack.

In practice, full specifications are rarely available. We believe that our technique can be used even with lightweight specifications or in the absence of specifications. Consider our example in Fig. 1. A symbolic analysis of this library, with a harness representing a sequential client making an arbitrary sequence of calls to the library, can, in principle, infer that the returned value equals $f(\text{num})$. As the returned value is the only observable value, this is the strongest functional specification a user can write. Our tool can be used with such an inferred specification as well.

Logical interference. Existing concurrency control mechanisms (both pessimistic as well as optimistic) rely on a data-access based notion of interference: concurrent accesses to the same data, where at least one access is a write, is conservatively treated as interference. A contribution of this paper is that it introduces a more logical/semantic notion of interference that can be used to achieve more permissive, yet safe, concurrency control. Specifically, concurrency control based on this approach permits interleavings that existing schemes based on stricter notion of interference will disallow. Hand-crafted concurrent code often permits “benign interference” for performance reasons, suggesting that programmers do rely on such a logical notion of interference.

2. THE PROBLEM

In this section, we introduce required terminology and formally define the problem. Rather than restrict ourselves to a specific syntax for programs and assertions, we will treat them abstractly, assuming only that they can be given a semantics as indicated below, which is fairly standard.

¹ This may be a valid design option in some cases. However, in examples such as our running example, this could be a bad idea.

2.1. The Sequential Setting. *Sequential Libraries.* A library \mathcal{L} is a pair (\mathcal{P}, V_G) , where \mathcal{P} is a set of procedures (defined below), and V_G is a set of variables, termed *global* variables, accessible to all and only procedures in \mathcal{P} . A procedure $P \in \mathcal{P}$ is a pair (G_P, V_P) , where G_P is a control-flow graph with each edge labeled by a primitive statement, and V_P is a set of variables, referred to as *local* variables, restricted to the scope of P . (Note that V_P includes the formal parameters of P as well.) To simplify the semantics, we will assume that the set V_P is the same for all procedures and denote it V_L .

Every control-flow graph has a unique entry vertex N_P (with no predecessors) and a unique exit vertex X_P (with no successors). Primitive statements are either **skip** statements, assignment statements, **assume** statements, **return** statements, or **assert** statements. An **assume** statement is used to implement conditional control flow as usual. Given control-flow graph nodes u and v , we denote an edge from u to v , labeled with a primitive statement S , by $u \xrightarrow{S} v$.

To reason about all possible sequences of invocations of the library's procedures, we define the *control graph* of a library to be the union of the control-flow graphs of all the procedures, augmented by a new vertex w , as well as an edge from every procedure exit vertex to w and an edge from w to every procedure entry vertex. We refer to w as the *quiescent* vertex. Note that a one-to-one correspondence exists between a path in the control graph of the library, starting from w , and the execution of a sequence of procedure calls. The edge $w \rightarrow N_P$ from the quiescent vertex to the entry vertex of a procedure P models an arbitrary call to procedure P . We refer to these as *call edges*.

Sequential States. A procedure-local state $\sigma_\ell \in \Sigma_\ell^s$ is a pair (pc, σ_d) where pc , the program counter, is a vertex in the control graph and σ_d is a map from the local variables V_L to their values. (We use “ s ” as a superscript or subscript to indicate elements of the semantics of sequential execution.) A global state $\sigma_g \in \Sigma_g^s$ is a map from global variables V_G to their values. A library state σ is a pair $(\sigma_\ell, \sigma_g) \in \Sigma_\ell^s \times \Sigma_g^s$. We define Σ^s to be $\Sigma_\ell^s \times \Sigma_g^s$. We say that a state is a *quiescent state* if its pc value is w and that it is an *entry state* if its pc value equals the entry vertex of some procedure.

Sequential Executions. We assume a standard semantics for primitive statements that can be captured as a transition relation $\rightsquigarrow_s \subseteq \Sigma^s \times \Sigma^s$ as follows. Every control-flow edge e induces a transition relation $\overset{e}{\rightsquigarrow}_s$, where $\sigma \overset{e}{\rightsquigarrow}_s \sigma'$ iff σ' is one of the possible outcomes of the execution of (the statement labeling) the edge e in state σ . The edge $w \rightarrow N_P$ from the quiescent vertex to the entry vertex of a procedure P models an arbitrary call to procedure P . Hence, in defining the transition relation, such edges are treated as statements that assign a non-deterministically chosen value to every formal parameter of P and the default initial value to every local variable of P . Similarly, the edge $X_P \rightarrow w$ is treated as a **skip** statement. We say $\sigma \rightsquigarrow_s \sigma'$ if there exists some edge e such that $\sigma \overset{e}{\rightsquigarrow}_s \sigma'$.

A *sequential execution* is a sequence of states $\sigma_0 \sigma_1 \cdots \sigma_k$ where σ_0 is the initial state of the library and we have $\sigma_i \rightsquigarrow_s \sigma_{i+1}$ for $0 \leq i < k$. A sequential execution represents the execution of a sequence of calls to the library's procedures (where the last call's execution may be incomplete). Given a sequential execution $\sigma_0 \sigma_1 \cdots \sigma_k$, we say that σ_i is the *corresponding entry state* of σ_j if σ_i is an entry state and no state σ_h is an entry state for $i < h \leq j$.

Sequential Assertions. We use **assert** statements to specify desired correctness properties of the library. **Assert** statements have no effect on the execution semantics and are equivalent

to **skip** statements in the semantics. Assertions are used only to define the notion of *well-behaved* executions as follows.

An **assert** statement is of the form **assert** θ where, θ is a 1-state assertion φ or a 2-state assertion Φ . A 1-state assertion, which we also refer to as a predicate, makes an assertion about a single library state. Rather than define a specific syntax for assertions, we assume that the semantics of assertions are defined by a relation $\sigma \models_s \varphi$ denoting that a state σ satisfies the assertion φ .

1-state assertions can be used to specify the invariants expected at certain program points. In general, specifications for procedures take the form of two-state assertions, which relate the input state to output state. We use 2-state assertions for this purpose. The semantics of a 2-state assertion Φ is assumed to be defined by a relation $(\sigma_{in}, \sigma_{out}) \models_s \Phi$ (meaning that state σ_{out} satisfies assertion Φ with respect to state σ_{in}). In our examples, we use special input variables v^{in} to refer to the value of the variable v in the first state. E.g., the specification “ $x == x^{in} + 1$ ” asserts that the value of x in the second state is one more than its value in the first state.

Definition 2.1. A sequential execution is said to satisfy the library’s assertions if for any transition $\sigma_i \xrightarrow{e}_s \sigma_{i+1}$ in the execution, where e is labelled by the statement “**assert** θ ”, we have (a) $\sigma_i \models_s \theta$ if θ is a 1-state assertion, and (b) $(\sigma_{in}, \sigma_i) \models_s \theta$ where σ_{in} is the corresponding entry state of σ_i , otherwise. A sequential library satisfies its specifications if every execution of the library satisfies its specifications.

2.2. The Concurrent Setting. *Concurrent Libraries.* A concurrent library \mathcal{L} is a triple (\mathcal{P}, V_G, Lk) , where \mathcal{P} is a set of concurrent procedures, V_G is a set of global variables, and Lk is a set of global locks. A concurrent procedure is like a sequential procedure, with the extension that a primitive statement is either a sequential primitive statement or a locking statement of the form **acquire**(ℓ) or **release**(ℓ) where ℓ is a lock.

Concurrent States. A concurrent library permits concurrent invocations of procedures. We associate each procedure invocation with a thread (representing the client thread that invoked the procedure). Let T denote an infinite set of thread-ids, which are used as unique identifiers for threads. In a concurrent execution, every thread has a private copy of local variables, but all threads share a single copy of the global variables. Hence, the local-state in a concurrent execution is represented by a map from T to Σ_ℓ^s . (A thread whose local-state’s pc value is the quiescent point represents an idle thread, i.e., a thread not processing any procedure invocation.) Let $\Sigma_\ell^c = T \rightarrow \Sigma_\ell^s$ denote the set of all local states.

At any point during execution, a lock lk is either free or held by one thread. We represent the state of locks by a partial function from Lk to T indicating which thread, if any, holds any given lock. (A lock that is not held by any thread will not be in the domain of the partial function.) Let $\Sigma_{lk}^c = Lk \hookrightarrow T$ represent the set of all lock-states. Let $\Sigma_g^c = \Sigma_g^s \times \Sigma_{lk}^c$ denote the set of all global states. Let $\Sigma^c = \Sigma_\ell^c \times \Sigma_g^c$ denote the set of all states. Given a concurrent state $\sigma = (\sigma_\ell, (\sigma_g, \sigma_{lk}))$ and thread t , we define $\sigma[t]$ to be the sequential state $(\sigma_\ell(t), \sigma_g)$.

Concurrent Executions. The concurrent semantics is induced by the sequential semantics as follows. Let e be any control-flow edge labelled with a sequential primitive statement, and t be any thread. We say that $(\sigma_\ell, (\sigma_g, \sigma_{lk})) \xrightarrow{(t,e)}_c (\sigma'_\ell, (\sigma'_g, \sigma_{lk}))$ iff $(\sigma_t, \sigma_g) \xrightarrow{e}_s (\sigma'_t, \sigma'_g)$ where

$\sigma_t = \sigma_\ell(t)$ and $\sigma'_\ell = \sigma_\ell[t \mapsto \sigma'_t]$. The transitions corresponding to lock acquire/release are defined in the obvious way. We say that $\sigma \rightsquigarrow_c \sigma'$ iff there exists some (t, e) such that $\sigma \xrightarrow{(t,e)}_c \sigma'$.

A *concurrent execution* is a sequence $\sigma_0\sigma_1\cdots\sigma_k$, where σ_0 is the initial state of the library and $\sigma_i \xrightarrow{\ell_i}_c \sigma_{i+1}$ for $0 \leq i < k$, where the label $\ell_i = (t_i, e_i)$ identifies the executing thread and executed edge. We say that $\ell_0 \cdots \ell_{k-1}$ is the *schedule* of this execution. A sequence $\ell_0 \cdots \ell_m$ is a *feasible schedule* if it is the schedule of some concurrent execution. Consider a concurrent execution $\sigma_0\sigma_1\cdots\sigma_k$. We say that a state σ_i is a *t-entry-state* if it is generated from a quiescent state by thread t executing a call edge. We say that σ_i is the *corresponding t-entry state* of σ_j if σ_i is a t -entry-state and no state σ_h is a t -entry-state for $i < h \leq j$.

We note that our semantics uses sequential consistency. Extending our results to support weaker memory models is future work.

Interpreting Assertions In Concurrent Executions. In a concurrent setting, assertions are evaluated in the context of the thread that executes the corresponding `assert` statement. We say that state σ satisfies a 1-state assertion φ in the context of thread t_i (denoted by $(\sigma, t_i) \models_c \varphi$) iff $\sigma[t_i] \models_s \varphi$. For any 2-state assertion Φ , we say that a given pair of states $(\sigma_{in}, \sigma_{out})$ satisfies Φ in the context of thread t (denoted by $((\sigma_{in}, \sigma_{out}), t) \models_c \Phi$) iff $(\sigma_{in}[t], \sigma_{out}[t]) \models_s \Phi$.

Definition 2.2. A concurrent execution π is said to satisfy an assertion “`assert θ` ” labelling an edge e if for any transition $\sigma_i \xrightarrow{(t,e)}_c \sigma_{i+1}$ in the execution, we have (a) $(\sigma_i, t) \models_c \theta$, if θ is a 1-state assertion, and (b) $((\sigma_{in}, \sigma_i), t) \models_c \theta$ where σ_{in} is the corresponding t -entry state of σ_i , otherwise. The execution is said to satisfy the library’s specification if it satisfies all assertions in the library. A concurrent library satisfies its specification if every execution of the library satisfies its specification.

Frame Conditions. Consider a library with two global variables x and y and a procedure `IncX` that increments x by 1. A possible specification for `IncX` is $(x == x^{in} + 1) \ \&\& \ (y == y^{in})$. The condition $y == y^{in}$ is `IncX`’s frame condition, which says that it will not modify y . Explicitly stating such frame conditions is unnecessarily restrictive, as a concurrent update to y by another procedure, when `IncX` is executing, would be considered a violation of `IncX`’s specification. Frame conditions can be handled better by treating a specification as a pair (S, Φ) where S is the set of all global variables referenced by the procedure, and Φ is a specification that does not refer to any global variables outside S . For our above example, the specification will be $(\{x\}, x == x^{in} + 1)$. In the sequel, however, we will restrict ourselves to the simpler setting and ignore this issue.

2.3. Goals. Our goal is: Given a sequential library \mathcal{L} with assertions satisfied in every sequential execution, construct $\hat{\mathcal{L}}$, by augmenting \mathcal{L} with concurrency control, such that every concurrent execution of $\hat{\mathcal{L}}$ satisfies all assertions. In Section 6, we extend this goal to construct $\hat{\mathcal{L}}$ such that every concurrent execution of $\hat{\mathcal{L}}$ is linearizable.

3. PRESERVING SINGLE-STATE ASSERTIONS

In this section we describe our algorithm for synthesizing concurrency control, but restrict our attention to single-state assertions.

3.1. Algorithm Overview. A *sequential proof* is a mapping μ from vertices of the control graph to predicates such that (a) for every edge $e = u \xrightarrow{t} v$, $\{\mu(u)\}t\{\mu(v)\}$ is a valid Hoare triple $\sigma_1 \models_s \mu(u)$ and $\sigma_1 \xrightarrow{e} \sigma_2$ implies $\sigma_2 \models_s \mu(v)$, and (b) for every edge $u \xrightarrow{\text{assert } \varphi} v$, we have $\mu(u) \Rightarrow \varphi$. Note that condition (a) requires $\{\mu(u)\}t\{\mu(v)\}$ to be partially correct. The execution of statement t in a state satisfying $\mu(u)$ does not have to succeed. This is required primarily for the case when t represents an **assume** statement. If we want to ensure that the execution of a statement t does not cause any runtime error, we can simply replace t by “**assert** p ; t ” where p is the condition required to ensure that t does not cause any runtime error.)

Note that the invariant $\mu(u)$ attached to a vertex u by a proof indicates two things: (i) any sequential execution reaching point u will produce a state satisfying $\mu(u)$, and (ii) any sequential execution from point u , starting from a state satisfying $\mu(u)$ will satisfy the invariants labelling other program points (and satisfy all assertions encountered during the execution).

A procedure that satisfies its assertions in a sequential execution may fail to do so in a concurrent execution due to interference by other threads. E.g., consider a thread t_1 that reaches a program point u in a state that satisfies $\mu(u)$. At this point, another thread t_2 may execute some statement that changes the state to one where $\mu(u)$ no longer holds. Now, we no longer have a guarantee that a continued execution by t_1 will successfully satisfy its assertions. The preceding paragraph, however, hints at the interference we must avoid to ensure correctness: when a thread t_1 is at point u , we should ensure that no other thread t_2 changes the state to one where t_1 's invariant $\mu(u)$ fails to hold. Any change to the state by another thread t_2 can be tolerated by t_1 *as long as the invariant $\mu(u)$ continues to hold*. We can achieve this by associating a lock with the invariant $\mu(u)$, ensuring that t_1 holds this lock when it is at program point u , and ensuring that any thread t_2 acquires this lock before executing a statement that may break this invariant. An invariant $\mu(u)$, in general, may be a boolean formula over simpler predicates. We could potentially get different locking solutions by associating different locks with different sub-formulae of the invariant. We elaborate on this idea below.

A *predicate mapping* is a mapping \mathbf{pm} from the vertices of the control graph to a set of predicates. A predicate mapping \mathbf{pm} is said to be a *basis* for a proof μ if every $\mu(u)$ can be expressed as a boolean formula (involving conjunctions, disjunctions, and negation) over $\mathbf{pm}(u)$. A basis \mathbf{pm} for proof μ is *positive* if every $\mu(u)$ can be expressed as a boolean formula involving only conjunctions and disjunctions over $\mathbf{pm}(u)$.

Given a proof μ , we say that an edge $u \xrightarrow{s} v$ *sequentially positively preserves* a predicate φ if $\{\mu(u) \wedge \varphi\}s\{\varphi\}$ is a valid Hoare triple. Otherwise, we say that the edge *may sequentially falsify* the predicate φ . Note that the above definition is in terms of the Hoare logic for our sequential language. However, we want to formalize the notion of a thread t_2 's execution of an edge falsifying a predicate φ in a thread t_1 's scope. Given a predicate φ , let $\hat{\varphi}$ denote the predicate obtained by replacing every local variable x with a new unique variable \hat{x} . We say that an edge $u \xrightarrow{s} v$ *may falsify* φ iff the edge may sequentially falsify $\hat{\varphi}$. (Note

that this reasoning requires working with formulas with free variables, such as \hat{x} . This is straightforward as these can be handled just like extra program variables.)

E.g., consider Line 13 in Fig. 1. Consider predicate $lastRes == f(num)$. By renaming local variable `num` to avoid naming conflicts, we obtain predicate `lastRes == f(nûm)`. We say that Line 13 *may falsify* this predicate because the triple $\{res == f(num) \wedge lastNum == num \wedge lastRes == f(nûm)\} \text{ lastRes} = \text{res} \{lastRes == f(nûm)\}$ is not a valid Hoare triple.

Let \mathbf{pm} be a positive basis for a proof μ and $\mathcal{R} = \cup_u \mathbf{pm}(u)$. For any program point u , if a predicate φ is in $\mathbf{pm}(u)$, we say that φ is *relevant* at program point u . In a concurrent execution, we say that a predicate φ is relevant to a thread t in a given state if t is at a program point u in the given state and $\varphi \in \mathbf{pm}(u)$. Our locking scheme associates a lock with every predicate φ in \mathcal{R} . The invariant it establishes is that a thread, in any state, will hold the locks corresponding to precisely the predicates that are relevant to it. We will simplify the initial description of our algorithm by assuming that distinct predicates are associated with distinct locks and later relax this requirement.

Consider any control-flow edge $e = u \xrightarrow{s} v$. Consider any predicate φ in $\mathbf{pm}(v) \setminus \mathbf{pm}(u)$. We say that predicate φ *becomes relevant*² at edge e . In the motivating example, the predicate `lastNum == num` becomes relevant at Line 12

We ensure the desired invariant by acquiring the locks corresponding to every predicate that becomes relevant at edge e *prior to statement s in the edge*. (Acquiring the lock after s may be too late, as some other thread could intervene between s and the acquire and falsify predicate φ .)

Now consider any predicate φ in $\mathbf{pm}(u) \setminus \mathbf{pm}(v)$. We say that φ *becomes irrelevant* at edge e . E.g., predicate `lastRes == f(lastNum)` becomes irrelevant once the false branch at Line 8 is taken. For every p that becomes irrelevant at edge e , we release the lock corresponding to p *after* statement s .

The above steps ensure that in a concurrent execution a thread will hold a lock on all predicates relevant to it. The second component of the concurrency control mechanism is to ensure that any thread acquires a lock on a predicate before it falsifies that predicate. Consider an edge $e = u \xrightarrow{s} v$ in the control-flow graph. Consider any predicate $\varphi \in \mathcal{R}$ that may be falsified by edge e . We add an acquire of the lock corresponding to this predicate before s (unless $\varphi \in \mathbf{pm}(u)$), and add a release of the same lock after s (unless $\varphi \in \mathbf{pm}(v)$).

Managing locks at procedure entry/exit. We will need to acquire/release locks at procedure entry and exit differently from the scheme above. Our algorithm works with the control graph defined in Section 2. Recall that we use a quiescent vertex w in the control graph. The invariant $\mu(w)$ attached to this quiescent vertex describes invariants maintained by the library (in between procedure calls). Any `return` edge $u \xrightarrow{\text{return}} v$ must be augmented to release all locks corresponding to predicates in $\mathbf{pm}(u)$ before returning. Dually, any procedure entry edge $w \rightarrow u$ must be augmented to acquire all locks corresponding to predicates in $\mathbf{pm}(u)$.

However, this is not enough. Let $w \rightarrow u$ be a procedure p 's entry edge. The invariant $\mu(u)$ is part of the library invariant that procedure p depends upon. It is important to ensure that when a thread executes the entry edge of p (and acquires locks corresponding to the

² Frequently it will be the case that the execution of statement s makes predicate φ true. This is true if every invariant $\mu(v)$ is a conjunction of the basis predicates in $\mathbf{pm}(u)$. Since we allow disjunctions as well, this is not, however, always true.

basis of $\mu(u)$) the invariant $\mu(u)$ holds. We achieve this by ensuring that any procedure that invalidates the invariant $\mu(u)$ holds the locks on the corresponding basis predicates until it reestablishes $\mu(u)$. We now describe how this can be done in a simplified setting where the invariant $\mu(u)$ can be expressed as the conjunction of the predicates in the basis $\mathbf{pm}(u)$ for every procedure entry vertex u . (Disjunction can be handled at the cost of extra notational complexity.) We will refer to the predicates that occur in the basis $\mathbf{pm}(u)$ of some procedure entry vertex u as *library invariant predicates*.

We use an *obligation* mapping $\mathbf{om}(v)$ that maps each vertex v to a set of library invariant predicates to track the invariant predicates that may be invalid at v and need to be reestablished before the procedure exit. We say a function \mathbf{om} is a valid obligation mapping if it satisfies the following constraints for any edge $e = u \rightarrow v$: (a) if e may falsify a library invariant φ , then φ must be in $\mathbf{om}(v)$, and (b) if $\varphi \in \mathbf{om}(u)$, then φ must be in $\mathbf{om}(v)$ unless e *establishes* φ . Here, we say that an edge $u \xrightarrow{s} v$ *establishes* a predicate φ if $\{\mu(u)\}_s\{\varphi\}$ is a valid Hoare triple. Define $\mathbf{m}(u)$ to be $\mathbf{pm}(u) \cup \mathbf{om}(u)$. Now, the scheme described earlier can be used, except that we use \mathbf{m} in place of \mathbf{pm} .

Locking along assume edges. Recall that we model conditional branching, based on a condition p , using two edges labelled “**assume** p ” and “**assume** $!p$ ”. Any lock to be acquired along an **assume** edge will need to be acquired before the condition is evaluated. If the lock is required along both **assume** edges, this is sufficient. If the lock is not required along all **assume** edges out of a vertex, then we will have to release the lock along the edges where it is not required.

Deadlock Prevention. The locking scheme synthesized above may potentially lead to a deadlock. We now show how to modify the locking scheme to avoid this possibility. For any edge e , let $\mathbf{mbf}(e)$ be (a conservative approximation of) the set of all predicates that may be falsified by the execution of edge e . We first define a binary relation \succrightarrow on the predicates used (i.e., the set \mathcal{R}) as follows: we say that $p \succrightarrow r$ iff there exists a control-flow edge $u \xrightarrow{s} v$ such that $p \in \mathbf{m}(u) \wedge r \in (\mathbf{m}(v) \cup \mathbf{mbf}(u \xrightarrow{s} v)) \setminus \mathbf{m}(u)$. Note that $p \succrightarrow r$ holds iff it is possible for some thread to try to acquire a lock on r while it holds a lock on p . Let \succrightarrow^* denote the transitive closure of \succrightarrow .

We define an equivalence relation \rightleftharpoons on \mathcal{R} as follows: $p \rightleftharpoons r$ iff $p \succrightarrow^* r \wedge r \succrightarrow^* p$. Note that any possible deadlock must involve an equivalence class of this relation. We map all predicates in an equivalence class to the same lock to avoid deadlocks. In addition to the above, we establish a total ordering on all the locks, and ensure that all lock acquisitions we add to a single edge are done in an order consistent with the established ordering. (Note that the ordering on the locks does not have to be total; a partial ordering is fine, as long as any two locks acquired along a single edge are ordered.)

Improving The Solution. Our scheme can sometimes introduce *redundant* locking. E.g., assume that in the generated solution a lock ℓ_1 is always held whenever a lock ℓ_2 is acquired. Then, the lock ℓ_2 is redundant and can be eliminated. Similarly, if we have a predicate φ that is never falsified by any statement in the library, then we do not need to acquire a lock for this predicate. We can eliminate such redundant locks as a final optimization pass over the generated solution.

Using Reader-Writer Locks. Note that a lock may be acquired on a predicate φ for one of two reasons in the above scheme: either to “preserve” φ or to “break” φ . These are similar to read-locks and write-locks. Note that it is safe for multiple threads to simultaneously

hold a lock on the same predicate φ if they want to “preserve” it, but a thread that wants to “break” φ needs an exclusive lock. Thus, reader-writer locks can be used to improve concurrency, but space constraints prevent a discussion of this extension. However, since it is unsafe for a thread that holds a read-lock on a predicate φ to try to acquire a write-lock φ , using this optimization also requires an extension to the basic deadlock avoidance scheme.

Specifically, it is unsafe for a thread that holds a read-lock on a predicate φ to try to acquire a write-lock φ , as this can lead to a deadlock. Hence, any acquisition of a lock on a predicate φ (to preserve it) should be made an exclusive (write) lock if along some execution path it may be necessary to promote this lock to a write lock before the lock is released.

Generating Proofs. The sequential proof required by our scheme can be generated using verification tools such as SLAM [2], BLAST [11, 12] or Yogi [10]. Predicate abstraction [2] is a program analysis technique that constructs a conservative, finite state abstraction of a program with a large (possibly infinite) state space using a set of predicates over program variables. Tools such as SLAM and BLAST use predicate abstraction to check if a given program P satisfies a specification ϕ . The tools start with a simple initial abstraction and iteratively refine the abstraction until the abstraction is rich enough to prove the absence of a concrete path from the program’s initial state to an error state (or a real error is identified).

For programs for which verification succeeds, the final abstraction produced, as well as the result of abstract interpretation using this abstraction, serve as a good starting point for constructing the desired proof. The final abstraction consists of a predicate map \mathbf{pm} which maps each program point to a set of predicates and as well as a mapping from each program statement to a set of abstract predicate transformers which together define an abstract transition system. Furthermore, abstract interpretation utilizing this abstraction effectively computes a formula $\mu(u)$ over the set of predicates $\mathbf{pm}(u)$ at every program point u that conservatively describes all states that can arise at program point u .

The map μ constitutes a proof of sequential correctness, as required by our algorithm, and the predicate map \mathbf{pm} is a valid basis for the proof. The map \mathbf{pm} can be extended into a positive basis for the proof easily enough. Since a minimal proof can lead to better concurrency control, approaches that produce a “parsimonious proof” (e.g., see [12]) are preferable. A parsimonious proof is one that avoids the use of unnecessary predicates at any program point.

3.2. Complete Schema. We now present a complete outline of our schema for synthesizing concurrency control.

- (1) Construct a sequential proof μ that the library satisfies the given assertions in any sequential execution.
- (2) Construct positive basis \mathbf{pm} and an obligation mapping \mathbf{om} for the proof μ .
- (3) Compute a map \mathbf{mbf} from the edges of the control graph to \mathcal{R} , the range of \mathbf{pm} , such that $\mathbf{mbf}(e)$ (conservatively) includes all predicates in \mathcal{R} that may be falsified by the execution of e .
- (4) Compute the equivalence relation \rightleftharpoons on \mathcal{R} .
- (5) Generate a predicate lock allocation map $\mathbf{lm} : \mathcal{R} \rightarrow \mathcal{L}$ such that for any $\varphi_1 \rightleftharpoons \varphi_2$, we have $\mathbf{lm}(\varphi_1) = \mathbf{lm}(\varphi_2)$.

- (6) Compute the following quantities for every edge $e = u \xrightarrow{s} v$, where we use $\mathbf{lm}(X)$ as shorthand for $\{ \mathbf{lm}(p) \mid p \in X \}$ and $\mathbf{m}(u) = \mathbf{pm}(u) \cup \mathbf{om}(u)$:

$$\begin{aligned} \mathit{BasisLocksAcq}(e) &= \mathbf{lm}(\mathbf{m}(v)) \setminus \mathbf{lm}(\mathbf{m}(u)) \\ \mathit{BasisLocksRel}(e) &= \mathbf{lm}(\mathbf{m}(u)) \setminus \mathbf{lm}(\mathbf{m}(v)) \\ \mathit{BreakLocks}(e) &= \mathbf{lm}(\mathbf{mbf}(e)) \setminus \mathbf{lm}(\mathbf{m}(u)) \setminus \mathbf{lm}(\mathbf{m}(v)) \end{aligned}$$

- (7) We obtain the concurrency-safe library $\widehat{\mathcal{L}}$ by transforming every edge $u \xrightarrow{s} v$ in the library \mathcal{L} as follows:

- (a) $\forall p \in \mathit{BasisLocksAcq}(u \xrightarrow{s} v)$, add an `acquire`($\mathbf{lm}(p)$) before s ;
- (b) $\forall p \in \mathit{BasisLocksRel}(u \xrightarrow{s} v)$, add a `release`($\mathbf{lm}(p)$) after s ;
- (c) $\forall p \in \mathit{BreakLocks}(u \xrightarrow{s} v)$, add an `acquire`($\mathbf{lm}(p)$) before s and a `release`($\mathbf{lm}(p)$) after s .

All lock acquisitions along a given edge are added in an order consistent with a total order established on all locks.

3.3. Correctness. We now present a formal statement of the correctness claims for our algorithm. Let \mathcal{L} be a given library with a set of embedded assertions satisfied by all sequential executions of \mathcal{L} . Let $\widehat{\mathcal{L}}$ be the library obtained by augmenting \mathcal{L} with concurrency control using the algorithm presented in Section 3.2. Let μ , \mathbf{pm} , and \mathbf{om} be the proof, the positive basis, and the obligation map used to generate $\widehat{\mathcal{L}}$.

Consider any concurrent execution of the given library \mathcal{L} . We say that a thread t is *safe* in a state σ if $(\sigma, t) \models_c \mu(u)$ where t 's program-counter in state σ is u . We say that thread t is *active* in state σ if its program-counter is something other than the quiescent vertex. We say that state σ is *safe* if every active thread t in σ is safe. Recall that a concurrent execution is of the form: $\sigma_0 \xrightarrow{\ell_0} \sigma_1 \xrightarrow{\ell_1} \dots \sigma_n$, where each label ℓ_i is an ordered pair (t, e) indicating that the transition is generated by the execution of edge e by thread t . We say that a concurrent execution is safe if every state in the execution is safe. It trivially follows that a safe execution satisfies all assertions of \mathcal{L} .

Note that every concurrent execution π of $\widehat{\mathcal{L}}$ corresponds to an execution π' of \mathcal{L} if we ignore the transitions corresponding to lock acquire/release instructions. We say that an execution π of $\widehat{\mathcal{L}}$ is safe if the corresponding execution π' of \mathcal{L} is safe. The goal of the synthesized concurrency control is to ensure that only safe executions of $\widehat{\mathcal{L}}$ are permitted.

We say that a transition $\sigma \xrightarrow{(t,e)} \sigma'$ preserves the basis of an active thread $t' \neq t$ whose program-counter in state σ is u if for every predicate $\varphi \in \mathbf{pm}(u)$ the following holds: if $(\sigma, t') \models_c \varphi$, then $(\sigma', t') \models_c \varphi$. We say that a transition $\sigma \xrightarrow{(t,e)} \sigma'$ ensures the basis of thread t if either $e = x \rightarrow y$ is not the procedure entry edge or for every active thread $t' \neq t$ whose program-counter in state σ is u and for every predicate $\varphi \in \mathbf{pm}(u)$ none of the predicates in $\mathbf{pm}(y)$ are in $\mathbf{om}(u)$.

We say that a transition $\sigma \xrightarrow{(t,e)} \sigma'$ is *basis-preserving* if it preserves the basis of every active thread $t' \neq t$ and ensures the basis of thread t . A concurrent execution is said to be *basis-preserving* if all transitions in the execution are basis-preserving.

Lemma 3.1. (a) Any basis-preserving concurrent execution of \mathcal{L} is safe. (b) Any concurrent execution of $\widehat{\mathcal{L}}$ corresponds to a basis-preserving execution of \mathcal{L} .

Proof. (a) We prove that every state in a basis-preserving execution of \mathcal{L} is safe by induction on the length of the execution.

Consider a thread t in state σ with program-counter value u . Assume that t is safe in σ . Thus, $(\sigma, t) \models_c \mu(u)$. Note that $\mu(u)$ can be expressed in terms of the predicates in $\mathbf{pm}(u)$ using conjunction and disjunction. Let SP denote the set of all predicates φ in $\mathbf{pm}(u)$ such that $(\sigma, t) \models_c \varphi$. Let σ' be any state such that $(\sigma', t) \models_c \varphi$ for every $\varphi \in SP$. Then, it follows that t is safe in σ' . Thus, it follows that after any basis-preserving transition every thread that was safe before the transition continues to be safe after the transition.

We now just need to verify that whenever an inactive thread becomes active (representing a new procedure invocation), it starts off being safe. We can establish this by inductively showing that every library invariant must be satisfied in a given state or must be in $\mathbf{om}(u)$ for some active thread t at vertex u .

(b) Consider a concurrent execution of $\widehat{\mathcal{L}}$. We need to show that every transition in this execution, ignoring lock acquires/releases, is basis-preserving. This follows directly from our locking scheme. Consider a transition $\sigma \xrightarrow{(t,e)} \sigma'$. Let $t' \neq t$ be an active thread whose program-counter in state σ is u . For every predicate $\varphi \in \mathbf{pm}(u) \cup \mathbf{om}(u)$, our scheme ensures that t' holds the lock corresponding to φ . As a result, both the conditions for preserving basies are satisfied. \square

Theorem 3.2. (a) Any concurrent execution of $\widehat{\mathcal{L}}$ satisfies every assertion of \mathcal{L} . (b) The library $\widehat{\mathcal{L}}$ is deadlock-free.

Proof. (a) This follows immediately from Lemma 3.1.

(b) This follows from our scheme for merging locks and can be proved by contradiction. Assume that a concurrent execution of $\widehat{\mathcal{L}}$ produces a deadlock. Then, we must have a set of threads t_1 to t_k and a set of locks ℓ_1 to ℓ_k such that each t_i holds lock ℓ_i and is waiting to acquire lock $\ell_{i \oplus 1}$, where $i \oplus 1$ denotes $(i \bmod k) + 1$. In particular, t_i must hold lock ℓ_i because it wants a lock on some predicate p_i , and must be trying to acquire lock $\ell_{i \oplus 1}$ because of some predicate $q_{i \oplus 1}$. Thus, we must have $q_i \rightleftharpoons p_i$ and $p_i \rightsquigarrow q_{i \oplus 1}$ for every i . This implies that all of p_i and q_i must be in the same equivalence class of \rightleftharpoons and, hence, ℓ_1 through ℓ_k must be the same, which is a contradiction (since we must have $k > 1$ to have a deadlock). \square

As mentioned earlier, our synthesis technique has a close connection to Owicki-Gries [18] approach to verifying concurrent programs. An alternative approach to proving Theorem 3.2(a) would be to construct a suitable Owicki-Gries style proof for the library. We believe that this is doable.

4. HANDLING 2-STATE ASSERTIONS

The algorithm presented in the previous section can be extended to handle 2-state assertions via a simple program transformation that allows us to treat 2-state assertions (in the original program) as single-state assertions (in the transformed program). We augment the set of local variables with a new variable \tilde{v} for every (local or shared) variable v in the original program and add a primitive statement \mathcal{LP} at the entry of every procedure, whose execution essentially copies the value of every original variable v to the corresponding new variable \tilde{v} .

Let $\underline{\sigma}'$ denote the projection of a transformed program state σ' to a state of the original program obtained by forgetting the values of the new variables. Given a 2-state assertion Φ ,

Library	Description
<i>compute.c</i>	See Figure 1
<i>reduce.c</i>	See Figure 3
<i>increment.c</i>	See Figure 4
<i>average.c</i>	Two procedures that compute the running sum and average of a sequence of numbers
<i>device_cache.c</i>	One procedure that reads data from a device and caches the data for subsequent reads [7]. The specification requires quantified predicates.
<i>server_store.c</i>	A library derived from a Java implementation of Simple Authentication and Security Layer (SASL). The library stores security context objects for sessions on the server side.

Table 1: Benchmarks used in our evaluation.

let $\tilde{\Phi}$ denote the single-state assertion obtained by replacing every v^{in} by \tilde{v} . As formalized by the claim below, the satisfaction of a 2-state assertion Φ by executions in the original program corresponds to satisfaction of the single-state assertion $\tilde{\Phi}$ in the transformed program.

Lemma 4.1.

- (1) *A schedule ξ is feasible in the transformed program iff it is feasible in the original program.*
- (2) *Let σ' and σ be the states produced by a particular schedule with the transformed and original programs, respectively. Then, $\sigma = \sigma'$.*
- (3) *Let π' and π be the executions produced by a particular schedule with the transformed and original program, respectively. Then, π satisfies a single-state assertion φ iff π' satisfies it. Furthermore, π satisfies a 2-state assertion Φ iff π' satisfies the corresponding one-state assertion $\tilde{\Phi}$.*

Synthesizing concurrency control. We now apply the technique discussed in Section 3 to the transformed program to synthesize concurrency control that preserves the assertions transformed as discussed above. It follows from the above Lemma that this concurrency control, used with the original program, preserves both single-state and two-state assertions.

5. IMPLEMENTATION

We have built a prototype implementation of our algorithm. Our implementation takes a sequential library and its assertions as input. It uses a pre-processing phase to combine the library with a harness that simulates the execution of any possible sequence of library calls to get a complete C program. (This program corresponds to the control graph described in Section 2.) It then uses a verification tool to generate a proof of correctness for the assertions in this program. We use the predicate-abstraction based software verification tool Yogi described in [3] to generate the required proofs. We modified the verifier to emit the proof from the final abstraction, which associates every program point with a boolean formula over predicates. It then uses the algorithm presented in this paper to synthesize concurrency control for the library. It utilizes the theorem prover Z3 [5] to identify the statements in the program whose execution may falsify relevant predicates.

We used a set of benchmark programs to evaluate our approach (Table 1). We also applied our technique manually to two real world libraries, a device cache library [7], and a C implementation of the Simple Authentication and Security Layer (SASL). The proofs for the device cache library and the SASL library require quantified predicates, which were beyond the scope of the verifier we used.

In all these programs, the concurrency control scheme we synthesized was identical to what an experienced programmer would generate. The concurrency control we synthesized required one lock for all libraries, with the exception of the SASL library, where our solution uses two locks. Our solutions permit more concurrency as compared to a naive solution that uses one global lock or an atomic section around the body of each procedure. For example, in case of the server store library, our scheme generates smaller critical sections and identifies a larger number of critical sections that acquire different locks as compared to the default implementation. For these examples, the running time of our approach is dominated by the time required to generate the proof; the time required for the synthesis algorithm was negligible.

The source code for all our examples and their concurrent versions are available online at [1]. Note that our evaluation studies only small programs. We leave a more detailed evaluation of our approach as future work.

6. CONCURRENCY CONTROL FOR LINEARIZABILITY

6.1. The Problem. In the previous section, we showed how to derive concurrency control to ensure that each procedure satisfies its sequential specification even in a concurrent execution. However, this may still be too permissive, allowing interleaved executions that produce counter-intuitive results and preventing compositional reasoning in clients of the library. E.g., consider the procedure `Increment` shown in Fig. 2, which increments a shared variable `x` by 1. The figure shows the concurrency control derived using our approach to ensure specification correctness. Now consider a multi-threaded client that initializes `x` to 0 and invokes `Increment` concurrently in two threads. It would be natural to expect that the value of `x` would be 2 at the end of any execution of this client. However, this implementation permits an interleaving in which the value of `x` at the end of the execution is 1: the problem is that both invocations of `Increment` individually meet their specifications, but the cumulative effect is unexpected³.

This is one of the difficulties with using pre/post-condition specifications to reason about concurrent executions.

One solution to this problem is to apply concurrency control synthesis to the code (library) that calls `Increment`. The synthesis can then detect the potential for interference between the two calls to `Increment` and prevent them from happening concurrently. Another possible solution, which we explore in this section, is for the library to guarantee a stronger correctness criteria called *linearizability* [13]. Linearizability gives the illusion that in any concurrent execution, (the sequential specification of) every procedure of the library appears to execute *instantaneously* at some point between its call and return. This illusion allows clients to reason about the behavior of concurrent library compositionally using its

³ We conjecture that such concerns do not arise when the specification does not refer to global variables. For instance, the specification for our example in Fig. 1 does not refer to global variables, even though the implementation uses global variables.


```

1 int x = 0;
2 //@ensures x == xin + 1 ∧ returns x
3 Increment() {
4   int tmp;
5   acquire(l(x==xin)); tmp = x; release(l(x==xin));
6   tmp = tmp + 1;
7   acquire(l(x==xin)); x = tmp; release(l(x==xin));
8   return tmp;
9 }

```

Figure 2: A non-linearizable implementation of the procedure `Increment`

sequential specifications. In this section, we show how our approach presented earlier for synthesizing a *logical concurrency control* can be adapted to derive concurrency control mechanisms that guarantee linearizability.

6.1.1. *Linearizability.* We now extend the earlier notation to define linearizability. Linearizability is a property of the library’s externally observed behavior. A library’s interaction with its clients can be described in terms of a *history*, which is a sequence of events, where each event is an *invocation* event or a *response* event. An invocation event is a tuple consisting of the procedure invoked, the input parameter values for the invocation, as well as a unique identifier. A response event consists of the identifier of a preceding invocation event, as well as a return value. Furthermore, an invocation event can have at most one matching response event. A complete history has a matching response event for every invocation event. Note that an execution, as defined in Section 2, captures the internal execution steps performed during a procedure execution. A history is an abstraction of an execution that captures only procedure invocation and return steps.

A sequential history is an alternating sequence $inv_1, r_1, \dots, inv_n, r_n$ of invocation events and corresponding response events. We abuse our earlier notation and use $\sigma + inv_i$ to denote an entry state corresponding to a procedure invocation consisting of a valuation σ for the library’s global variables and a valuation inv_i for the invoked procedure’s formal parameters. We similarly use $\sigma + r_i$ to denote a procedure exit state with return value r_i . Let σ_0 denote the value of the globals in the library’s initial state. Let Φ_i denote the specification of the procedure invoked by inv_i . A sequential history is *legal* if there exist valuations σ_i , $1 \leq i \leq n$, for the library’s globals such that $(\sigma_{i-1} + inv_i, \sigma_i + r_i) \models_s \Phi_i$ for $1 \leq i \leq n$.

A complete interleaved history H is *linearizable* if there exists some legal sequential history S such that (a) H and S have the same set of invocation and response events and (b) for every return event r that precedes an invocation event inv in H , r and inv appear in that order in S as well. An incomplete history H is said to be linearizable if the complete history H' obtained by appending some response events and omitting some invocation events without a matching response event is linearizable.

Finally, a library \mathcal{L} is said to be linearizable if every history produced by \mathcal{L} is linearizable.

The concept of a *linearization point* is often used in explanations and proofs of correctness of linearizable algorithms. Informally, a linearization point is a point (or control-flow edge) inside the procedure such that the procedure appears to execute atomically when it executes that point. Our eventual goal is to parameterize our synthesis algorithm with a

linearization point specification (a description of the point or points we wish to serve as the linearization point). In this paper, however, we treat the procedure entry edge as the linearization point and will refer to it as the linearization point.

6.1.2. Implementation As A Specification and Logical Serializability. The techniques we present in this section actually guarantee linearizability with respect to the given sequential implementation (i.e., treating the sequential implementation as a sequential specification). In particular, this approach guarantees that the concurrent execution will return the same values as some sequential execution. (The word *atomicity* is sometimes used to describe this behavior.) Such an approach has both advantages as well as disadvantages. The advantage is that the technique is more broadly applicable, in practice, as it does not require a user-provided specification. The disadvantage is that, in theory, the sequential implementation may be more restrictive than the intended specification. Hence, preserving the sequential implementation behavior may unnecessarily restrict concurrency.

The properties of atomicity and linearizability relate to the externally observed behavior of the library (i.e., the behavior as seen by clients of the library). The implementation technique we use also guarantees certain properties about the internal (execution) behavior of the library, which we explain now.

Recall that an execution, as defined in Section 2, captures the internal execution steps performed during a procedure execution while a history is an abstraction of an execution that captures only procedure invocation and return steps.

Recall that every transition $\sigma \xrightarrow{(t,e)} \sigma'$ is labelled by a pair (t, e) , indicating that the transition was created by the execution of edge e by thread t . We refer to a pair of the form (t, e) as a *step*. A schedule ζ is a sequence of steps ℓ_1, \dots, ℓ_k . We say that a schedule ℓ_1, \dots, ℓ_k is *feasible* if there exists an execution $\sigma_0 \xrightarrow{\ell_1} \sigma_1 \dots \xrightarrow{\ell_k} \sigma_k$, where σ_0 is the initial program state. Given an execution $\pi = \sigma_0 \xrightarrow{\ell_1} \sigma_1 \dots \xrightarrow{\ell_k} \sigma_k$, the sub-schedule of t in π is the sequence $\ell_{s_1}, \dots, \ell_{s_n}$ of steps executed by t in π .

A procedure invocation t_1 is said to precede another procedure invocation t_2 in an execution if t_1 completes before t_2 begins.

Two complete executions are said to be observationally-equivalent if they consist of the same set of procedure invocations and for each procedure invocation the return values are the same in both executions. An execution π_1 is said to be a permutation of another execution π_2 if for every thread (procedure invocation) t the sub-schedule of t in π_1 and π_2 are the same. An execution π_1 is said to be topologically consistent with another execution π_2 if for every pair of procedure invocations t_1 and t_2 , if t_1 precedes t_2 in π_1 then t_1 precedes t_2 in π_2 as well.

Our goal is to synthesize a concurrency control mechanism that permits only executions that are observationally-equivalent, topologically consistent, permutations of sequential executions. We note that this concept is similar to various notions of serializability [25] (commonly used in database transactions). The new variant we exploit may be thought of as *logical serializability*: corresponding points in the compared executions satisfy equivalence with respect to certain predicates of interest, as determined by the basis.

6.1.3. Terminology. In this section, we will use a modified notion of basis introduced in Section 3.

The key idea we explore in this paper is that of precisely characterizing what is *relevant* to a thread at a particular point and using this information to derive a concurrency control solution. In the previous sections, we captured the relevant information as an invariant or set of predicates (the basis). In this section, we will find it necessary to mark certain values (e.g., the value of a variable at a particular program point) as relevant as well. In order to seamlessly reason about such relevant values (e.g., of type integer) along with relevant predicates, we utilize *symbolic* predicates to encode the relevance of values.

In the sequel, note that we consider two predicates to be equal only if they are syntactically equal.

A symbolic predicate is one that utilizes auxiliary (logical) variables. As an example, given program variable \mathbf{x} and a logical variable w , we will make use of predicates such as “ $\mathbf{x} = w$ ”. Such symbolic predicates can be manipulated just like normal predicates (e.g., in computing weakest-precondition). Conceptually, such a symbolic predicate can be interpreted as a short-hand notation for the (possibly infinite) family of predicates obtained by replacing the logical variable w by every possible value it can take. Thus, if \mathbf{x} and w are of type T , then the above symbolic predicate represents the set of predicates $\{\mathbf{x} = c \mid c \in T\}$. Note that this set of predicates captures the value of \mathbf{x} : i.e., we know the value of every predicate in this set iff we know the value of \mathbf{x} . This trick lets us use the symbolic predicate “ $\mathbf{x} = w$ ” to indicate that the value of \mathbf{x} is relevant to a thread (and, hence, should not be modified by another thread).

Given any predicate φ , let φ^* denote the set of predicates it represents (obtained by instantiating the logical variables in φ as explained above). (Thus, for a non-symbolic predicate φ , $\varphi^* = \{\varphi\}$.) We say that φ_1 and φ_2 are equivalent if $\varphi_1^* = \varphi_2^*$. E.g., if w ranges over all integers, then “ $\mathbf{x} = w$ ” and “ $\mathbf{x} = w + 1$ ” are equivalent predicates. Predicate equivalence can be used to simplify a set of predicates or a basis. Given a set of predicates S , let S^* represent the set of predicates $\cup\{\varphi^* \mid \varphi \in S\}$. If $S_1^* = S_2^*$, then it is safe, in the sequel, to replace the set S_1 by the set S_2 in a basis. This may be critical in creating finite representations of certain basis.

We say that a predicate φ is covered by a set of predicates S if φ can be expressed as a boolean formula over the predicates in S using conjunctions and disjunctions.

Recall that a *predicate mapping* is a mapping \mathbf{pm} from the vertices of the control graph to a set of predicates.

We say that a predicate mapping \mathbf{pm} is *wp-closed* if for every edge $e = u \xrightarrow{s} v$ and for every $\varphi \in \mathbf{pm}(v)$, (a) If e is not the entry edge of a procedure, then the weakest-precondition of φ with respect to s , $wp(s, \varphi)$ is covered by $\mathbf{pm}(u)$, and (b) If e is the edge $w \rightarrow N_P$ from the quiescent vertex to the entry vertex of P , then φ' is covered by $\mathbf{pm}(w)$, where φ' is obtained by replacing the occurrence of any procedure parameter x_i by a new logical variable x'_i .

Finally, we say that a predicate mapping is *closed* if it is wp-closed and if for every vertex u and every predicate φ in $\mathbf{pm}(u)$, the negation of φ is also in $\mathbf{pm}(u)$. The later condition helps us reuse the algorithm description from Section 3 in spite of some differences in the context.

Without loss of generality, we assume that each procedure P_j returns the value of a special local variable ret_j .

6.2. The Synthesis Algorithm. We now show how our approach can be extended to guarantee linearizability or atomicity. We use a few tricky cases to motivate the adaptations we use of our previous algorithm.

We start by characterizing non-linearizable interleavings permitted by our earlier approach. We classify the interleavings based on the nature of linearizability violations they cause. For each class of interleavings, we describe an extension to our approach to generate additional concurrency control to prohibit these interleavings. Finally, we prove correctness of our approach by showing that all interleavings we permit are linearizable.

6.2.1. Delayed Falsification. The first issue we address, as well as the solution we adopt, are not surprising from a conventional perspective. (This extension is, in fact, the analogue of two-phase locking: i.e., the trick of acquiring all locks before releasing any locks to avoid interference.) Informally, the problem with the **Increment** example can be characterized as “dirty reads” and “lost updates”: the second procedure invocation executes its linearization point later than the first procedure invocation but reads the original value of x , instead of the value produced by the the first invocation. Dually, the update done by the first procedure invocation is lost, when the second procedure invocation updates x . From a logical perspective, the second invocation relies on the invariant $x == x^{in}$ early on, and the first invocation breaks this invariant later on when it assigns to x (at a point when the second invocation no longer relies on the invariant). This prevents us from reordering the execution to construct an equivalent sequential execution (while preserving the proof). To achieve linearizability, we need to avoid such “delayed falsification”.

The extension we now describe prevents such interference by ensuring that instructions that may falsify predicates and occur after the linearization point appear to execute atomically at the linearization point. We achieve this by modifying the strategy to acquire locks as follows.

- We generalize the earlier notion of *may-falsify*. We say that a path *may-falsify* a predicate φ if some edge in the path may-falsify φ . We say that a predicate φ *may-be-falsified-after* vertex u if there exists some path from u to the exit vertex of the procedure that does not contain any linearization point and may-falsify φ .
- Let \mathbf{mf} be a predicate map such that for any vertex u , $\mathbf{mf}(u)$ includes any predicate that may-be-falsified-after u .
- We generalize the original scheme for acquiring locks. We augment every edge $e = u \xrightarrow{S} v$ as follows:
 - (1) $\forall \ell \in \mathbf{lm}(\mathbf{mf}(v)) \setminus \mathbf{lm}(\mathbf{mf}(u))$, add an “**acquire**(ℓ)” before S
 - (2) $\forall \ell \in \mathbf{lm}(\mathbf{mf}(u)) \setminus \mathbf{lm}(\mathbf{mf}(v))$, add an “**release**(ℓ)” after S

This extension suffices to produce a linearizable implementation of the example in Fig. 2.

6.2.2. Return Value Interference. We now focus on interference that can affect *the actual value returned by a procedure invocation*, leading to non-linearizable executions.

Consider procedures **IncX** and **IncY** in Fig. 3, which increment variables x and y respectively. Both procedures return the values of x and y . However, the postconditions of **IncX** (and **IncY**) do not *specify anything about the final value of* y (and x respectively). Let us assume that the linearization points of the procedures are their entry points. Initially, we have $x = y = 0$. Consider the following interleaving of a concurrent execution of the two procedures. The two procedures execute the increments in some order, producing the state with $x = y = 1$. Then, both procedures return $(1, 1)$. This execution is non-linearizable because in any legal sequential execution, the procedure executing second is obliged to return a value that differs from the value returned by the procedure executing first. The

<pre> int x, y; IncX() { acquire($l_{x==x^{in}}$); x = x + 1; (ret₁₁, ret₁₂) = (x, y); release($l_{x==x^{in}}$); } IncY() { acquire($l_{y==y^{in}}$); y = y + 1; (ret₂₁, ret₂₂) = (x, y); release($l_{y==y^{in}}$); } </pre> <p style="text-align: center;">(a)</p>	<pre> int x, y; @ensures $x = x^{in} + 1$ @returns (x, y) IncX() { [$ret'_{11} == x + 1 \wedge ret'_{12} == y$] $\mathcal{LP} : x = x^{in}$ [$x == x^{in} \wedge ret'_{11} == x + 1 \wedge ret'_{12} = y$] x = x + 1; (ret₁₁, ret₁₂) = (x, y); [$x == x^{in} + 1 \wedge ret'_{11} == ret'_{11}$ $\wedge ret'_{12} == ret'_{12}$] } </pre> <p style="text-align: center;">(b)</p>	<pre> int x, y; IncX() { acquire(l_{merged}); x = x + 1; (ret₁₁, ret₁₂) = (x, y); release(l_{merged}); } IncY() { acquire(l_{merged}); y = y + 1; (ret₂₁, ret₂₂) = (x, y); release(l_{merged}); } </pre> <p style="text-align: center;">(c)</p>
---	--	---

Figure 3: An example illustrating return value interference. Both procedures return (x, y) . ret_{ij} refers to the j^{th} return variable of the i^{th} procedure. Figure 3(a) is a non-linearizable implementation synthesized using the approach described in Section 3. Figure 3(b) shows the extended proof of correctness of the procedure `IncX` and Figure 3(c) shows the linearizable implementation.

left column in Figure 3 shows the concurrency control derived using our approach with previously described extensions. This is insufficient to prevent the above interleaving. This interference is allowed because the specification for `IncX` allows it to change the value of y arbitrarily; hence, a concurrent modification to y by any other procedure is not seen as a hindrance to `IncX`.

To prohibit such interferences within our framework, we need to determine whether the execution of a statement s can potentially affect the return-value of another procedure invocation. We do this by computing a predicate $\phi(ret')$ at every program point u that captures the relation between the program state at point u and the value returned by the procedure invocation eventually (denoted by ret'). We then check if the execution of a statement s will break predicate $\phi(ret')$, treating ret' as a free variable, to determine if the statement could affect the return value of some other procedure invocation.

Formally, we assume that each procedure returns the value of a special variable ret . (Thus, “`return exp`” is shorthand for “ $ret = exp$ ”.) We introduce a special auxiliary variable ret' . We say that a predicate map \mathbf{pm} covers return statements if for every edge $u \rightarrow v$ labelled by a return statement “`return exp`” the set $\mathbf{pm}(u)$ covers the predicate $ret' == ret$. (See the earlier discussion in Section 6.1.3 about such symbolic predicates and how they encode the requirement that the value of ret at a return statement is relevant and must be preserved.)

By applying our concurrency-control synthesis algorithm to a closed basis that covers return statements, we can ensure that no return-value interference occurs.

The middle column in Figure 3 shows the augmented sequential proof of correctness of `IncX`. The concurrency control derived using our approach starting with this proof is shown in the third column of Fig. 3. The lock l_{merged} denotes a lock obtained by merging locks corresponding to multiple predicates simultaneously acquired/released. It is easy to see that this implementation is linearizable. Also note that if the shared variables y and x were

```

1 int x, y;
2 //@ensures y = yin + 1
3 IncY() {
4   [true]  $\mathcal{LP} : y^{in} = y$ 
5   [y == yin] y = y + 1;
6   [y == yin + 1]
7 }

1 //@ensures x < y
2 ReduceX() {
3   [true]  $\mathcal{LP}$ 
4   [true] if (x ≥ y) {
5     [true] x = y - 1;
6   }
7   [x < y]
8 }

```

Figure 4: An example illustrating interference in control flow. Each line is annotated (in square braces) with a predicate the holds at that program point.

not returned by procedures `IncX` and `IncY` respectively, we will derive a locking scheme in which accesses to `x` and `y` are protected by different locks, allowing these procedures to execute concurrently.

6.2.3. Control Flow Interference. An interesting aspect of our scheme is that it permits interference that alters the control flow of a procedure invocation if it does not cause the invocation to violate its specification. Consider procedures `ReduceX` and `IncY` shown in Fig. 4. The specification of `ReduceX` is that it will produce a final state where $x < y$, while the specification of `IncY` is that it will increment the value of `y` by 1. `ReduceX` meets its specification by setting `x` to be `y - 1`, but does so *only if* $x \geq y$.

Now consider a client that invokes `ReduceX` and `IncY` concurrently from a state where $x = y = 0$. Assume that the `ReduceX` invocation enters the procedure. Then, the invocation of `IncY` executes completely. The `ReduceX` invocation continues, and does nothing since $x < y$ at this point.

Figure 4 shows a sequential proof and the concurrency control derived by the scheme so far, assuming that the linearization points are at the procedure entry. A key point to note is that `ReduceX`'s proof needs only the single predicate $x < y$. The statement $y = y + 1$ in `IncY` does *not falsify* the predicate $x < y$; hence, `IncY` does not acquire the lock for this predicate. This locking scheme permits `IncY` to execute concurrently with `ReduceX` and affect its control flow. While our approach guarantees that this control flow interference will not cause assertion violations, proving linearizability in the presence of such control flow interference, in the general case, is challenging (and an open problem).

We now describe how our technique can be extended to prevent control flow interference, which suffices to guarantee linearizability.

We ensure that interference by one thread does not affect the execution path another thread takes. We say that a basis \mathbf{pm} covers the branch conditions of the program if for every branch edge $u \xrightarrow{s} v$, the set $\mathbf{pm}(u)$ covers the assume condition in s . If we synthesize concurrency control using a *closed* basis \mathbf{pm} that covers the branch conditions, we can ensure that no control-flow interference happens.

In the current example, this requires predicate $x \geq y$ to be added to the basis for `ReduceX`. As a result, `ReduceX` will acquire lock $l_{x \geq y}$ at entry, while `IncY` will acquire the same lock at its linearization point and release the lock after the statement $y = y + 1$. It is easy to see that this implementation is linearizable.

6.2.4. *The Complete Schema.* In summary, our schema for synthesizing concurrency control that guarantees linearizability is as follows.

First, we determine a closed basis for the program that covers all return statements and branch conditions in the program. (Such a basis is the analogue of the proof and basis used in Section 3. An algorithm for generating such a basis is beyond the scope of this paper. Such a basis can be computed by iteratively computing weakest-preconditions, but, in the general case, subsumption and equivalence among predicates will need to be utilized to simplify basis sets to ensure termination.) We then apply the extended concurrency control synthesis algorithm described in Section 6.2.1.

6.3. **Correctness.** The extensions described above to the algorithm of Sections 3 and 4 for synthesizing concurrency control are sufficient to guarantee linearizability, as we show in this section.

Let σ be a program state. We define $\text{TBP}(\sigma, t)$ to be the set $\{\varphi \in (\mathbf{pm}(u))^* \mid (\sigma, t) \models_c \varphi\}$ where t 's program-counter in state σ is u . (See Section 6.1.3 for the definition of S^* for any set of predicates S .)

Lemma 6.1. *Let \mathbf{pm} be a wp-closed predicate map. Consider transitions $\sigma_1 \xrightarrow{(t,e)}_c \sigma_2$ and $\sigma_3 \xrightarrow{(t,e)}_c \sigma_4$. If $\text{TBP}(\sigma_1, t) \supseteq \text{TBP}(\sigma_3, t)$, then $\text{TBP}(\sigma_2, t) \supseteq \text{TBP}(\sigma_4, t)$.*

Proof. Let e be the edge $u \xrightarrow{S} v$. Note that for every predicate $\varphi \in \mathbf{pm}(v)$, the weakest-precondition of φ with respect to the statement S can be expressed in terms of the predicates in $\mathbf{pm}(u)$ using conjunction and disjunction (by definition of a wp-closed predicate map). The result follows. \square

Consider any concurrent execution π_1 produced by a schedule ξ . We assume, without loss of generality, that every procedure invocation is executed by a distinct thread. Let t_1, \dots, t_k denote the set of threads which complete execution in the given schedule, ordered so that t_i executes its linearization point before t_{i+1} . We show that ξ is linearizable by showing that ξ is equivalent to a sequential execution of the specifications of the threads t_1, \dots, t_k executed in that order.

Let ξ_i denote a *projection* of schedule ξ consisting only of execution steps by thread t_i . Let ζ denote the schedule $\xi_1 \cdots \xi_k$.

Lemma 6.2. *ζ_k is a feasible schedule. Furthermore, for any corresponding execution steps $\sigma_j \xrightarrow{(t,e)}_c \sigma_{j+1}$ and $\sigma'_k \xrightarrow{(t,e)}_c \sigma'_{k+1}$ of the two executions, we have $\text{TBP}(\sigma_j, t) \supseteq \text{TBP}(\sigma'_k, t)$.*

Proof. Proof by induction over the execution steps of ζ .

The claim is trivially true for the first step of ζ , since the initial state is the same in both executions.

Now, consider any pair of “candidate” successive execution steps $\sigma'_{k-1} \xrightarrow{(t,e')}_c \sigma'_k \xrightarrow{(t,e)}_c \sigma'_{k+1}$ of ζ . That is, we assume, from our inductive hypothesis, that the first execution step above is feasible, but we need to establish that the second step is a feasible execution step.

Let $\sigma_{m-1} \xrightarrow{(t,e')}_c \sigma_m$ and $\sigma_j \xrightarrow{(t,e)}_c \sigma_{j+1}$ be the two corresponding execution steps in the original execution.

Our inductive hypothesis guarantees that $\text{TBP}(\sigma_m, t) \supseteq \text{TBP}(\sigma'_k, t)$. But any concurrent execution is guaranteed to be interference-free. Hence, it follows that $\text{TBP}(\sigma_j, t) \supseteq \text{TBP}(\sigma_m, t)$. Hence, it follows that $\text{TBP}(\sigma_j, t) \supseteq \text{TBP}(\sigma'_k, t)$.

Now, if e is a conditional branch statement labelled with the statement “assume φ ”, then we must have $(\sigma_j, t) \models_c \varphi$. It follows that $(\sigma'_k, t) \models_c \varphi$. (This follows because we use a basis that covers all branch conditions.) Thus, the second candidate execution step of ζ is indeed a feasible execution step.

It then follows from Lemma 6.1 that $\text{TBP}(\sigma_{j+1}, t) \supseteq \text{TBP}(\sigma'_{k+1}, t)$.

Now, consider any pair of successive execution steps $\sigma'_{k-1} \xrightarrow{(t_{h-1}, e')} \sigma'_k \xrightarrow{(t_h, e)} \sigma'_{k+1}$ of ζ . Thus, we consider the first step executed by thread t_h after thread t_{h-1} executes its last step.

Note that $\text{TBP}(\sigma'_k, t_{h-1}) = \text{TBP}(\sigma'_k, t_h)$ (since none of the basis predicates at the quiescent vertex involve thread-local variables).

Let $\sigma_p \xrightarrow{(t_{h-1}, e)} \sigma_{p+1}$ denote the corresponding, last, execution step performed by t_{h-1} in the interleaved execution. Let $\sigma_j \xrightarrow{(t_h, e)} \sigma_{j+1}$ denote the corresponding, first, execution step performed by t_h in the interleaved execution. By the inductive hypothesis, $\text{TBP}(\sigma_p, t_{h-1}) \supseteq \text{TBP}(\sigma'_k, t_{h-1})$.

Note that in the interleaved execution p may be less than or greater than j : t_{h-1} may or may not have completed execution by the time t_h performs its first execution step. Yet, we can establish that $\text{TBP}(\sigma_j, t_h) \supseteq \text{TBP}(\sigma'_k, t_{h-1})$. This is because no thread can execute a step that will change the value of any predicate in $\text{TBP}(\sigma_p, t_{h-1})$ between the last step of t_{h-1} and the first step of t_h (no matter how these two steps are ordered during execution). \square

Lemma 6.3. *For $t \in \{t_1, \dots, t_k\}$, the value returned by procedure invocation t in π_1 is the same as the value returned by t in the sequential execution π_2 corresponding to schedule ζ .*

Proof. Let $\sigma_j \xrightarrow{(t, e)} \sigma_{j+1}$ and $\sigma'_k \xrightarrow{(t, e)} \sigma'_{k+1}$ denote the execution of the return statements by t in π_1 and π_2 respectively. It follows from Lemma 6.2 that $\text{TBP}(\sigma_j, t) \supseteq \text{TBP}(\sigma'_k, t)$. Suppose that t returns a value c in the sequential execution π_2 . Note that we use a basis that covers all return statements. Hence, the predicate $c == \text{ret}$ must be in $\text{TBP}(\sigma'_k, t)$. It follows that $c == \text{ret}$ must be in $\text{TBP}(\sigma_j, t)$ as well. Hence, t returns c in π_1 as well. \square

Theorem 6.4. *Given a library \mathcal{L} that is totally correct with respect to a given sequential specification, the library $\widehat{\mathcal{L}}$ generated by our algorithm is linearizable with respect to the given specification.*

Proof. Follows immediately from Lemma 6.3. \square

The above theorem requires *total correctness* of the library in the sequential setting. *E.g.*, consider a procedure P with a specification **ensures** $\mathbf{x}==0$. An implementation that sets \mathbf{x} to be 1, and then enters an infinite loop is *partially* correct with respect to this specification (but not totally correct). In a concurrent setting, this can lead to non-linearizable behavior, since another concurrent thread can observe that \mathbf{x} has value 1, which is not a legally observable value *after* procedure P completes execution.

6.4. Discussion. In this section, we have presented a logical approach to synthesizing concurrency control to ensure linearizability/atomicity. In particular, we use predicates to describe what is *relevant* to ensure correctness (or desired properties). Predicates enable us to describe relevance in a more fine-grained fashion, creating opportunities for more concurrency.

We believe that this approach is promising and that there is significant scope for improving our solution and several interesting research directions worth pursuing. Indeed, some basic optimizations to the scheme presented may be critical to getting reasonable solutions. One example is an optimization relating to frame conditions, hinted at in Section 2. As an example, assume that $x > 0$ is an invariant that holds true in between procedure invocations in a sequential execution. (Thus, this is a library invariant.) A procedure that neither reads or writes x will, nevertheless, have the invariant $x > 0$ at every program point (to indicate that it never breaks this invariant). Our solution, as sketched, will require the procedure to acquire a lock on this predicate and hold it during the entire procedure. However, this is not really necessary, and can be optimized away. In general, the invariant or the basis at any program point may be seen as consisting of two parts, the *frame* and the *footprint*. The footprint relates to predicates that are relevant and/or may be modified by the procedure, while the frame simply indicates predicates that are irrelevant and left untouched by the procedure. We need to consider only the footprint in synthesizing the concurrency control solution. We leave fleshing out the details of such optimizations as future work.

We conjecture that the extensions presented in this section to avoid control-flow interference is not necessary to ensure linearizability. Indeed, note that if we can ensure that any concurrent execution is observationally equivalent and topologically equivalent to some sequential execution, this is sufficient. Our current technique ensures that the concurrent execution is also a permutation of the sequential execution: i.e., every procedure invocation follows the same execution path in both the concurrent and sequential execution. However, our current proof of correctness relies on this property. Relaxing this requirement is an interesting open problem.

We believe that our technique can be adapted in a straight-forward fashion to work with linearization points other than the procedure entry (as long as the linearization point satisfies certain constraints). Different linearization points can potentially produce different concurrency control solutions.

We also believe that with various of these improvements, we can synthesize the solution presented in Fig. 1 as a linearizable and atomic implementation, starting with no specification whatsoever.

7. RELATED WORK

Synthesizing Concurrency Control: Vechev et al. [24] present an approach for synthesizing concurrency control for a concurrent program, given a specification in the form of assertions in the program. This approach, Abstraction Guided Synthesis, generalizes the standard counterexample-guided abstraction refinement (CEGAR) approach to verification as follows. The algorithm attempts to prove that the concurrent program satisfies the desired assertions. If this fails, an interleaved execution that violates an assertion is identified. This counterexample is used to either refine the abstraction (as in CEGAR) or to restrict the program by adding some atomicity constraints. An atomicity constraint indicates that a context-switch should not occur at a given program point (thus requiring the statements immediately preceding and following the program point to be in an atomic-block) or is a disjunction of such constraints. Having thus refined either the abstraction or the program, the algorithm repeats this process.

Our work has the same high-level goal and philosophy as Vechev et al.: derive a concurrency control solution automatically from a specification of the desired correctness properties. However, there are a number of differences between the two approaches. Before we discuss these differences, it is worth noting that the concrete problem addressed by these two papers are somewhat different: while our work focuses on making a sequential library safe for concurrent clients, Vechev et al. focus on adding concurrency control to a given concurrent program to make it safe. Thus, neither technique can be directly applied to the other problem, but we can still observe the following points about the essence of these two approaches.

Both approaches are similar in exploiting verification techniques for synthesizing concurrency control. However, our approach decouples the verification step from the synthesis step, while Vechev et al. present an integrated approach that combines both. Our verification step requires only sequential reasoning, while the Vechev et al. algorithm involves reasoning about concurrent (interleaved) executions. Specifically, we exploit the fact that a sequential proof indicates what properties are critical at different program points (for a given thread), which allows us to determine whether the execution of a particular statement (by another thread) constitutes (potentially) undesirable interference.

We present a locking-based solution to concurrency control, while Vechev et al. present the solution in terms of atomic regions. Note that if our algorithm is parameterized to use a single lock (i.e., to map every predicate to the same lock), then the generated solution is effectively one based on atomic regions.

Raza et al. [19] present an approach for automatically parallelizing a program that makes use of a separation logic proof. This approach exploits the separation logic based proof to identify whether candidate statements for parallelization access disjoint sets of locations. Like most classical approaches to automatic parallelization, this approach too relies on a data-based notion of interference, while our approach identifies a logical notion of interference.

Several papers [9, 4, 8, 16, 14, 21] address the problem of inferring lock-based synchronization for atomic sections to guarantee atomicity. These existing lock inference schemes identify potential conflicts between atomic sections at the granularity of data items and acquire locks to prevent these conflicts, either all at once or using a two-phase locking approach. Our approach is novel in using a logical notion of interference (based on predicates), which can permit more concurrency.

[20] describes a sketching technique to add missing synchronization by iteratively exploring the space of candidate programs for a given thread schedule, and pruning the search space based on counterexample candidates. [15] uses model-checking to repair errors in a concurrent program by pruning erroneous paths from the control-flow graph of the interleaved program execution. [23] is a precursor to [24], discussed above, that considers the tradeoff between increasing parallelism in a program and the cost of synchronization. This paper allows users to specify limitations on what may be used as the guard of conditional critical regions (the synchronization mechanism used in the paper), thus controlling the costs of synchronization. [6] allows users to specify synchronization patterns for critical sections, which are used to infer appropriate synchronization for each of the user-identified region. Vechev *et al.* [22] address the problem of automatically deriving linearizable objects with fine-grained concurrency, using hardware primitives to achieve atomicity. The approach is semi-automated, and requires the developer to provide algorithm schema and insightful manual transformations. Our approach differs from all of these techniques in

exploiting a proof of correctness (for a sequential computation) to synthesize concurrency control that guarantees thread-safety.

Verifying Concurrent Programs: Our proposed style of reasoning is closely related to the axiomatic approach for proving concurrent programs of Owicki & Gries [18]. While they focus on proving a concurrent program correct, we focus on synthesizing concurrency control. They observe that if two statements *do not interfere*, the Hoare triple for their parallel composition can be obtained from the sequential Hoare triples. Our approach identifies statements that *may interfere* and violate the sequential Hoare triples, and then synthesizes concurrency control to ensure that sequential assertions are preserved by parallel composition.

Prior work on verifying concurrent programs [17] has also shown that attaching invariants to resources (such as locks and semaphores) can enable modular reasoning about concurrent programs. Our paper turns this around: we use sequential proofs (which are modular proofs, but valid only for sequential executions) to identify critical invariants and create locks corresponding to such invariants and augment the program with concurrency control that enables us to lift the sequential proof into a valid proof for the concurrent program.

8. LIMITATIONS, EXTENSIONS, AND FUTURE WORK

In this paper, we have explored the idea that proofs of correctness for sequential computations can yield concurrency control solutions for use when the same computations are executed concurrently. We have adopted simple solutions in a number of dimensions in order to focus on this central idea. A number of interesting ideas and problems appear worth pursuing in this regard, as explained below.

Procedures. The simple programming language presented in Section 2 does not include procedures. The presence of procedure calls within the library gives rise to a different set of challenges. Verification tools often compute procedure summaries to derive the overall proof of correctness. Our approach could use summaries as proxies for procedure calls and derives concurrency control schemes where locks are acquired and released only in the top-level procedures. A more aggressive approach could analyze the proofs bottom up and infer nested concurrency control schemes where locks are acquired and released in procedures that subsume the lifetimes of the corresponding predicates.

Relaxed Memory Models. The programming language semantics we use and our proofs assume sequential consistency. We believe it should be possible to extend the notion of logical interference to relaxed memory models. Under a relaxed model, reads may return more values compared to sequential consistent executions. Therefore, we may have to consider these additional behaviors while determining if a statement can interfere with (the proof of) a concurrent thread. We leave this extension for future work.

Optimistic Concurrency Control. Optimistic concurrency control is an alternative to pessimistic concurrency control (such as lock-based techniques). While we present a lock-based pessimistic concurrency control mechanism, it would be interesting to explore the possibility of optimistic concurrency control mechanisms that exploit a similar weaker notion of interference.

Choosing Good Solutions. This paper presents a space of valid locking solutions that guarantee the desired properties. Specifically, the locking solution generated is dependent on several factors: the sequential proof used, the basis used for the proof, the mapping from basis predicates to locks, the linearization point used, etc. Given a metric on solutions, generating a good solution according to the given metric is a direction for future work. E.g., one possibility is to evaluate the performance of candidate solutions (suggested by our framework) using a suitable test suite to choose the best one. Integrating the concurrency control synthesis approach with the proof generation approach, as done by [24], can also lead to better solutions, if the proofs themselves can be refined or altered to make the concurrency control more efficient.

Fine-Grained Locking. Fine-grained locking refers to locking disciplines that use an unbounded number of locks and associate each lock with a small number of shared objects (typically one). Programs that use fine-grained locking often scale better because of reduced contention for locks. In its current form, the approach presented in this paper does not derive fine-grained locking schemes. The locking schemes we synthesize associate locks with predicates and the number of such predicates is statically bounded. Generalizing our approach to infer fine-grained locking from sequential proofs of correctness remains an open and challenging problem.

Lightweight Specifications. Our technique relies on user-provided specifications for the library. Recently, there has been interest in lightweight annotations that capture commonly used correctness conditions in concurrent programs (such as atomicity, determinism, and linearizability). As we discuss in Section 1, we believe that there is potential for profitably applying our technique starting with such lightweight specifications (or even no specifications).

Class invariants. In our approach, a thread holds a lock on a predicate from the point the predicate is established to the point after which the predicate is no longer used. While this approach ensures correctness, it may often be too pessimistic. For example, it is often the case that a library is associated with class/object invariants that characterize the *stable* state of the library’s objects. Procedures in the library may temporarily break and then re-establish the invariants at various points during their invocation. If class invariants are known, it may be possible to derive more efficient concurrency control mechanisms that release locks on the class invariants at points where the invariants are established and re-acquire these locks when the invariants are used. Such a scheme works only when all procedures “co-operate” and ensure that the locks associated with the invariants are released only when the invariant is established.

REFERENCES

- [1] WYPIWYG examples. http://research.microsoft.com/en-us/projects/wypiwyg/wypiwyg_examples.zip, June 2009.
- [2] T. Ball and S. K. Rajamani. Bebop: A symbolic model checker for Boolean programs. In *SPIN 00: SPIN Workshop*, pages 113–130. 2000.
- [3] Nels E. Beckman, Aditya V. Nori, Sriram K. Rajamani, Robert J. Simmons, SaiDeep Tetali, and Aditya V. Thakur. Proofs from tests. *IEEE Trans. Software Eng.*, 36(4):495–508, 2010.
- [4] Sigmund Chereh, Trishul Chilimbi, and Sumit Gulwani. Inferring locks for atomic sections. In *Proc. of PLDI*, 2008.

- [5] Leonardo Mendonça de Moura and Nikolaj Bjørner. Z3: An efficient smt solver. In *TACAS*, pages 337–340, 2008.
- [6] Xianghua Deng, Matthew B. Dwyer, John Hatcliff, and Masaaki Mizuno. Invariant-based specification, synthesis, and verification of synchronization in concurrent programs. In *Proc. of ICSE*, pages 442–452, 2002.
- [7] Tyfun Elmas, Serdar Tasiran, and Shaz Qadeer. A calculus of atomic sections. In *Proc. of POPL*, 2009.
- [8] Michael Emmi, Jeff Fischer, Ranjit Jhala, and Rupak Majumdar. Lock allocation. In *Proc. of POPL*, 2007.
- [9] Cormac Flanagan and Stephen N. Freund. Automatic synchronization correction. In *Proc. of SCOOOL*, 2005.
- [10] Bhargav S. Gulavani, Thomas A. Henzinger, Yamini Kannan, Aditya V. Nori, and Sriram K. Rajamani. Synergy: A new algorithm for property checking. In *Proc. of FSE*, November 2006.
- [11] T. A. Henzinger, R. Jhala, R. Majumdar, and G. Sutre. Lazy abstraction. In *Proc. of POPL*, pages 58–70, 2002.
- [12] Thomas A. Henzinger, Ranjit Jhala, Rupak Majumdar, and Kenneth L. McMillan. Abstractions from proofs. In *Proc. of POPL*, pages 232–244, 2004.
- [13] Maurice P. Herlihy and Jeannette M. Wing. Linearizability: a correctness condition for concurrent objects. *Proc. of ACM TOPLAS*, 12(3):463–492, 1990.
- [14] Michael Hicks, Jeffrey S. Foster, and Polyvios Pratikakis. Lock inference for atomic sections. In *First Workshop on Languages, Compilers, and Hardware Support for Transactional Computing*, 2006.
- [15] Muhammad Umar Janjua and Alan Mycroft. Automatic correcting transformations for safety property violations. In *Proc. of Thread Verification*, pages 111–116, 2006.
- [16] Bill McCloskey, Feng Zhou, David Gay, and Eric A. Brewer. Autolocker: Synchronization inference for atomic sections. In *Proc. of POPL*, 2006.
- [17] Peter W. O’Hearn. Resources, concurrency, and local reasoning. *Theor. Comput. Sci.*, 375(1-3):271–307, 2007.
- [18] Susan Owicki and David Gries. Verifying properties of parallel programs : An axiomatic approach. In *Proc. of CACM*, 1976.
- [19] Mohammad Raza, Cristiano Calcagno, and Philippa Gardner. Automatic parallelization with separation logic. In *ESOP*, pages 348–362, 2009.
- [20] Armando Solar-Lezama, Christopher Grant Jones, and Rastislav Bodik. Sketching concurrent data structures. In *Proc. of PLDI*, pages 136–148, 2008.
- [21] Mandana Vaziri, Frank Tip, and Julian Dolby. Associating synchronization constraints with data in an object-oriented language. In *Proc. of POPL*, pages 334–345, 2006.
- [22] Martin Vechev and Eran Yahav. Deriving linearizable fine-grained concurrent objects. In *In Proc. of PLDI*, pages 125–135, 2008.
- [23] Martin Vechev, Eran Yahav, and Greta Yorsh. Inferring synchronization under limited observability. In *Proc. of TACAS*, 2009.
- [24] Martin T. Vechev, Eran Yahav, and Greta Yorsh. Abstraction-guided synthesis of synchronization. In *POPL*, pages 327–338, 2010.
- [25] Gerhard Weikum and Gottfried Vossen. *Transactional Information Systems: Theory, Algorithms, and the Practice of Concurrency Control*. Morgan Kaufmann, 2001.