Automatable Verification of Sequential Consistency*

[Extended Abstract]

Anne E. Condon and Alan J. Hu
Department of Computer Science
University of British Columbia
2366 Main Mall
Vancouver, B.C. V6T 1Z4
Canada
(condon,ajh)@cs.ubc.ca

ABSTRACT

Sequential consistency is a multiprocessor memory model of both practical and theoretical importance. Designing and implementing a memory system that efficiently provides a given memory model is a challenging and error-prone task, so automated verification support would be invaluable. Unfortunately, the general problem of deciding whether a finite-state protocol implements sequential consistency is undecidable. In this paper, we identify a restricted class of protocols for which verifying sequential consistency is decidable. The class includes all published sequentially consistent protocols that are known to us, and we argue why the class is likely to include all real sequentially consistent protocols. In principle, our method can be applied in a completely automated fashion for verification of all implemented protocols.

Categories and Subject Descriptors

B.3.3 [**Performance Analysis and Design Aids**]: [Formal Models]; C.0 [**Computer Systems Organization**]: General—*systems specification methodology*

General Terms

theory, verification

Keywords

sequential consistency, memory model, model checking

1. INTRODUCTION

The memory model of a shared memory multiprocessor is a specification of how memory will behave from the programmer's perspective. Memory systems use intricate finite-state protocols to implement the desired memory model. These protocols are notoriously difficult to design and debug — because the primary objective is performance rather than simplicity — making them natural targets for formal verification.

Sequential consistency is a memory model introduced by Lamport [9]. A memory system is sequentially consistent iff there always exists an interleaving of the program orders of all the processors such that each load returns the value of the most recent store to the same address. Sequential consistency is important both as a practical memory model that provides intuitive ease-of-programming while allowing efficient hardware optimizations (e.g. [8]) and also as an extensively studied memory model that can be used to understand other, more relaxed models (e.g. [1]).

Ideally, we would like an algorithm that inspects a finite-state protocol and determines automatically whether or not the protocol provides sequential consistency. Unfortunately, the general problem of deciding sequential consistency of a finite-state protocol is undecidable [3].

Real protocols, however, might not be fully general, suggesting that the undecidability result may not be relevant in practice. Suppose we can characterize a class of protocols with the following properties: membership in the class is decidable, all members of the class are sequentially consistent, and all real protocols that implement sequential consistency belong to the class. This paper proposes such a class, thereby reducing automatic verification of real, sequentially consistent protocols to testing for membership in the protocol class.

The basis for our verification method is a graph-based definition of sequential consistency that arises in the work of Gibbons and Korach [6]. For an execution trace of a protocol, they define a constraint graph with a node for each load and store operation in the trace. The graph has four kinds of edges: edges that enforce program order for each processor, edges that provide a total order over all store nodes to each memory location, edges from each store node to every load node that gets its value from that store, and forced edges from each load node to the store node that follows in the total store order the store node from which the load got its

*The authors were supported in part by research grants from the National Science and Engineering Research Council of Canada.

value. A protocol is sequentially consistent if and only if all of its traces have acyclic constraint graphs.

To perform automatic formal verification using this formulation of sequential consistency, we must provide an automatic way to construct the constraint graph and verify that it is acyclic for all possible executions of the protocol. In practice, this suggests that the construction and checking of the constraint graph must be done in (hopefully small) finite state, so that automatic verification based on finite-state model checking [5] is possible.

The remainder of this paper addresses these problems. In Section 3, we introduce a graph description notation tailored to describe constraint graphs, and a finite-state checker to verify that a graph so described is acyclic. We describe how the graph description notation and checker can be used to verify sequential consistency. In Section 4, we show how real protocols can be annotated with finite-state information, to obtain a finite state observer which generates a description of the constraint graph. Our method of generating this description characterizes a class of protocols for which sequential consistency is decidable, and we argue why all real protocols are likely to belong to this class. Finally, we derive size bounds on the finite-state observer, suggesting that our method is at the edge of what is currently feasible for automatic verification tools.

1.1 Related Work

There has been considerable work over the years on verifying memory system protocols and memory models. For brevity, we mention here only closely related work, pertaining to finite-state verification of protocols with respect to sequential consistency.

Plakal et al. [12] introduce a verification approach based on logical clocks and apply it to a directory based protocol. Our approach is inspired by the logical clocks approach, but in contrast to logical clocks, which are unbounded, our approach reduces verification to a language inclusion problem between finite state automata.

Henzinger et al. [7] propose a very similar approach to ours, using a finite-state observer to reorder loads and stores to construct a witness of sequential consistency. Because of the finite-state limit on reordering, the method is too restrictive to handle most real protocols. One could view our approach as a generalization of theirs that handles all realistic protocols. We note that Henzinger et al. prove very strong results for protocols in their restrictive class, namely that it is sufficient to reduce verification of a protocol with arbitrarily large parameters (number of processors, number of blocks, number of values per block) to a fixed-parameter problem. In contrast, our method applies to verification of only fixed-parameter protocols.

Nalumasu et al. [11] propose the Test Model-Checking technique, in which a protocol is checked against various predefined finite-state automata that test certain memory model properties. These tests can be considered to be finite-state observers. By combining these tests, it is possible to verify memory models that are close to, but not identical to, sequential consistency. Determining exactly how these test combinations relate to sequential consistency and to the class of protocols we can handle is an open question.

At a recent, informal workshop, Qadeer proposed an approach for automatically verifying that a memory protocol implements a memory model [13]. The basic idea is to identify and formalize many assumptions that typically hold of real protocols and real memory models. In the presence of these assumptions, one can generate a finite-state witness automatically. The protocol class we can verify is much more general than Qadeer's, which cannot handle Afek et al's Lazy Caching protocol [2], for example. On the other hand, his complexity bounds (on the size of the finite state witness) are better than ours, and he considers memory models other than just sequential consistency. We believe the two approaches are complementary: Qadeer's approach can be generalized by adopting our model; our approach can be made more efficient by exploiting Qadeer's assumptions.

At the same workshop, we presented a preliminary version of the ideas that evolved into this work [10]. The general approach was the same as in this paper, but the underlying model for recording and checking constraints was different, resulting in wildly impractical finite-state size bounds. In subsequent work [4], we demonstrated that the method does allow verification, using current model-checking tools, of the sequential consistency of a substantial cache protocol, provided that some human insight is used to reduce the complexity of the observer. In contrast, the present paper presents a revised theoretical framework that encompasses a broader class of protocols, yet allows proving much stronger complexity bounds, suggesting that this work will apply to more protocols and be fully automatable in practice. We do not have experimental results yet, but are hopeful given our previous experiences.

2. DEFINITIONS

2.1 Protocols

A **protocol** $\mathcal{P}$ is a tuple $(p, b, v, Q, q_0, \mathcal{A} \cup \mathcal{A}', \delta \cup \delta', \perp)$. The constants p, b, and v specify the number of processors, memory blocks, and data values in the protocol. The symbol $\perp$ denotes the initial value of each block. The set of states is Q, of which q_0 is the initial state. The set $\mathcal{A}$ is the set of all actions of the protocol that are LD and ST operations, namely actions of the form $\mathrm{LD}(P, B, V)$ and $\mathrm{ST}(P, B, V)$, where $1 \leq P \leq p$, $1 \leq B \leq b$, and $1 \leq V \leq v$. For notational convenience, we use *'s to denote sets of LD and ST actions over all values of a parameter: e.g., $\mathrm{ST}(*, B, V)$ denotes the set $\{\mathrm{ST}(P, B, V) \mid 1 \leq P \leq p\}$. Thus, $\mathcal{A} = \mathrm{ST}(*, *, *) \cup \mathrm{LD}(*, *, *)$. $\mathcal{A}'$ is the set of actions of the protocol other than LD and ST operations. Corresponding to $\mathcal{A}$ and $\mathcal{A}'$ there are two transition relations, δ and δ', with $\delta \subseteq Q \times \mathcal{A} \times \mathcal{Q}$ and $\delta' \subseteq Q \times \mathcal{A}' \times \mathcal{Q}$.

A sequence of actions $A_1, A_2, \ldots, A_k$ is a **protocol run** if there is a sequence of states $q_0, q_1, q_2, \ldots, q_k$ such that for all j, with $1 \leq j \leq k$, the transition $(q_{j-1}, A_j, q_j) \in \delta \cup \delta'$. A **protocol trace** is the subsequence of a protocol run that includes only the actions in $\mathcal{A}$ (i.e., the ST and LD operations). Two protocols P and P' are **equivalent** if the set of traces of P equals the set of traces of P'. Note that the runs and traces of a protocol are finite, so our theory uses regular automata rather than ω-automata.

2.2 Sequential Consistency

Intuitively, a serial trace is one in which each LD returns the value of the most recent (prior to the LD) ST to the same block. If there were no prior STs to that block, the load must return $\perp$. Formally, a trace $T = t_1, t_2, \ldots, t_k$ is a **serial trace** if for all blocks B and values V, for all $1 \leq j \leq k$:

$$(t_j \in \mathrm{LD}(*, B, V)) \Rightarrow \left(\begin{array}{c} (V = \perp) \wedge \forall i_{<j}[t_i \notin \mathrm{ST}(*, B, *)] \\ \vee \\ \exists h_{<j}[t_h \in \mathrm{ST}(*, B, V) \wedge \forall i_{h<i<j}(t_i \notin \mathrm{ST}(*, B, *)] \end{array} \right).$$

A **reordering** of a trace of length k is simply a permutation Π of the numbers from 1 to k. Let $\Pi = \pi(1), \pi(2), \ldots \pi(k)$ be a reordering of a trace T. Let $T' = t_{\pi(1)}, t_{\pi(2)}, \ldots t_{\pi(k)}$. Π is called a **serial reordering** and T' is the corresponding serial trace if Π and T' have the following two properties. First, Π preserves the "per processor" order of T, i.e., for all processors P, if t_a and t_b are operations of processor P then $a < b$ if and only if $\pi^{-1}(a) < \pi^{-1}(b)$. Second, T' must be a serial trace.

A protocol is **sequentially consistent** if all of its traces have a serial reordering.

3. VERIFICATION USING CONSTRAINT GRAPHS

In our method for verifying that a protocol is sequentially consistent, a finite-state observer watches a protocol as it executes and gathers information about how to reorder the trace. The observer presents this information, in the form of a finite-state constraint graph, to a checker. A key task of the checker, which is also finite state, is to ensure that the graph is acyclic. Verification reduces to proving that the checker accepts all constraint graphs generated by the observer. See Figure 1. Overall, the method exploits the "less is more" principle: a total reordering of a trace is too much to be collected and checked with a finite number of states, but partial information about the reordering is sufficient to deduce sequential consistency.

We first define sequential consistency using graph-theoretic notation. Application of this definition to protocol verification requires a finite state method for testing if a graph is acyclic. In Section 3.2, we identify a class of graphs for which this test can be done. We describe the finite state cycle-checker in Section 3.3. We combine everything into our verification method in Section 3.4.

3.1 A Graph-Based Definition of Sequential Consistency

A *constraint graph* G for a trace T records ordering constraints on the operations in T that must be obeyed for T to have a serial reordering. The nodes of G are labeled by operations of T. Nodes are numbered by consecutive integers, starting from 1, according to their order in the trace. Edges of G include program order edges, along with *inheritance edges*, which indicate from which ST operation a LD inherits its value; *ST order edges*, which provide a total ordering of all ST nodes to the same block, and *forced edges*,

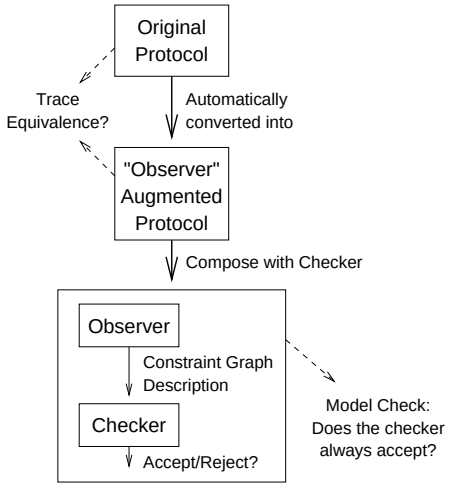

Figure 1: Verification Method Overview. The Observer is simply the original protocol augmented with reordering information. Automatic creation of the observer is discussed in Section 4. The observer generates a description of a constraint graph, which is checked by a finite-state checker. The same checker is used for all protocols. Constraint graphs and the checker are described in Section 3. The trace equivalence check can be omitted in practice because the observer is created in a non-interfering way from the original protocol.

which force the constraint that on any path from a ST node to a LD node that inherits its value, there is no other ST node to the same block. More precisely, edges of G must satisfy the following **edge annotation constraints:**

1. Each edge may be annotated as an inheritance, program order, ST order, or forced edge. An edge may have zero or more annotations.

2. For each processor P, if u nodes of G are labeled by operations of P then G has $u-1$ program order edges that define a total order on these u operations, consistent with trace order.

3. For each block B, if u nodes of G are labeled by ST operations to B, then G has $u-1$ ST order edges that define a total order on these u operations.

4. Each node labeled by $\mathrm{LD}(P, B, V)$ has one incoming inheritance edge from a $\mathrm{ST}(P', B, V)$ node (where P may equal P').

5. For all nodes i, j, and k such that there is a ST order edge from i to k and an inheritance edge from i to j, there is a forced edge on some path from j to k. Specifically, if j is labeled by $\mathrm{LD}(P, B, V)$ then there is either a forced edge directly from j to k or there is a (program order) path from j to another node j' labeled by $\mathrm{LD}(P, B, V)$, and a forced edge from $j'tok$. Similarly, for each node j labeled by a $\mathrm{LD}(P, B, \perp)$ operation, there is a forced edge on the path to the first node in the ST order for block B.

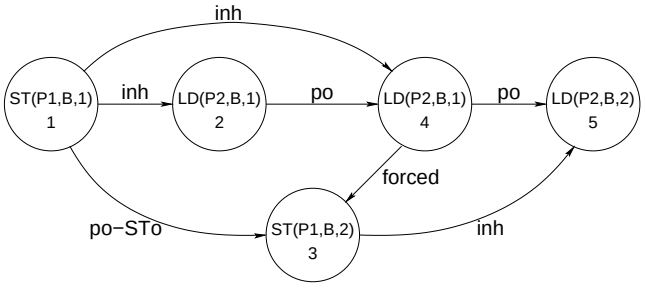

Figure 2: A Constraint Graph. Edge labels indicate inheritance (inh), program order (po), store order (STo), or "forced" edges. The inheritance edge from node 1 to node 4 and the store order edge from node 1 to node 3 forces an edge from node 4 to node 3, which prevents trace orders in which the LD in node 4 does not get its value from the most recent ST.

The graph has no annotated edges other than those specified in 2-5 above. An example of a constraint graph is given in Figure 2. The following claim is implicit in the work of Gibbons and Korach [6] and follows directly from the definition of constraint graph.

CLAIM 3.1. *A trace T has a serial reordering if and only if some constraint graph for T is acyclic.*

3.2 Node Bandwidth Bounded Graphs

For verification purposes, we are interested in constraint graphs (with ordered nodes) that are *node bandwidth bounded*. We denote the set $\{1, 2, \ldots, i\}$ by $\mathbf{N}_i$. We say that a graph with node set $\mathbf{N}_n$ is *k-node bandwidth bounded* if for all i, at most k nodes in $\mathbf{N}_i$ have edges to or from nodes in the set $\mathbf{N}_n - \mathbf{N}_i$. For example, the graph in Figure 2 is 3-node-bandwidth bounded. Note that node bandwidth boundedness is a property of both the graph and a fixed node ordering. Also, note that our definition differs from the usual edge-based notion of graph bandwidth, e.g., the number of edges between nodes in $\mathbf{N}_i$ and $\mathbf{N}_n - \mathbf{N}_i$ may be unbounded. For brevity, we omit the word "node" and simply refer to bandwidth bounded graphs.

We will represent a directed, k-bandwidth bounded graph G as a string, in a way that facilitates a finite state test that a graph is acyclic. For later convenience, nodes and edges of G may have labels from some finite alphabets $\mathcal{A}$ and $\mathcal{E}$, respectively. (In our application, $\mathcal{A}$ will be the set of trace operations, and symbols in $\mathcal{E}$ will denote the edge annotations of section 3.1.) Intuitively, our graph description notation simply lists nodes by number and edges as pairs of node numbers, with additional labels (if any) immediately following the node or edge to which they belong. A naive approach numbers all nodes and lists them in order. For example, the graph in Figure 2 corresponds to the description:

> 1, ST($P1, B, 1$), 2, LD($P2, B, 1$), (1,2), inh, 3, ST($P1, B, 2$), (1,3), po-STo, 4, LD($P2, B, 1$), (1,4), inh, (2,4), po, (4,3), forced, 5, LD($P2, B, 2$), (3,5), inh, (4,5), po

Our approach is like the naive approach, but is finite-state by exploiting k-bandwidth boundedness. In our approach, node numbers are not used directly to identify nodes and edges. Rather, each node has an ID (identification number) between 1 and $k+1$. When all edges in or out of the node with ID i have been listed, this node may be removed from the identification scheme and i can be used to identify another node. The graph in Figure 2 is 3-bandwidth bounded, so we can describe it as:

> 1, ST($P1, B, 1$), 2, LD($P2, B, 1$), (1,2), inh, 3, ST($P1, B, 2$), (1,3), po-STo, 4, LD($P2, B, 1$), (1,4), inh, (2,4), po, (4,3), forced, 1, LD($P2, B, 2$), (3,1), inh, (4,1), po

More formally, with respect to some fixed k and symbol alphabets $\mathcal{A}$ and $\mathcal{E}$, we define a *node descriptor* to be a symbol in $\mathbf{N}_{k+1}$, possibly followed by a symbol in $\mathcal{A}$ (that is, a node ID possibly followed by a node label) and an *edge descriptor* to be a symbol of the form (i, j) where $i, j \in \mathbf{N}_{k+1}$, possibly followed by a symbol in $\mathcal{E}$. A sequence of node descriptors and edge descriptors is a *k-graph descriptor*. Testing if a string is a proper graph descriptor (does not have two consecutive symbols from $\mathcal{A}$, for example), is easily done in finite state.

Let s be a k-graph descriptor. The graph G *represented by* s has a number of nodes equal to the number of node descriptors of s, with the ith node having the label (if any) of the ith node descriptor. Associated with each prefix s' of s is a set of *active* nodes which are associated with ID's, defined as follows. If s' has fewer than i node descriptors, then ID(i, s') is undefined. If s' has exactly i node descriptors with the ith having ID I, then ID$(i, s') = I$. Finally, suppose that s' has more than i node descriptors, with the ith having ID I. If no node descriptor in s' after the ith node descriptor has ID I, then ID$(i, s') = I$, else ID(i, s') is undefined. Now, the edges of G are defined as follows: for each prefix of the form $s', (I, I')$ of s, if for some pair (i, j) of nodes of G, ID$(i, s') = I$ and ID$(j, s') = I'$ then edge (i, j) is in G. Moreover, if $s', (I, I'), \beta$ is also a prefix of s for some $\beta \in \mathcal{E}$ then the edge (i, j) has label β.

A slightly extended notation for describing k-bandwidth bounded graphs will be useful later. Intuitively, in this extension, an active node may have more than one ID. This is useful, for example, when modeling the following situation: the value of a ST node in the constraint graph is in multiple cache locations of a finite state protocol, in which case it is convenient that these location addresses are the graph IDs for the ST node. For this purpose, we extend our graph descriptor strings to be sequences of node descriptors, edge descriptors and also symbols from the set $\{\text{add-ID}(I, I') \mid 1 \leq I, I' \leq k+1\}$. Intuitively, the add-ID(I, I') symbol causes the ID I' to be added to the node with ID I (and I' is no longer associated with any other node).

Such an *extended graph descriptor* represents a graph in which the nodes and node labels are defined just as for a valid string. To define the edges, for each node i, we define the ID-set of i with respect to s', denoted by ID-set(i, s'),

as follows. If s' has fewer than i node descriptors, then ID-set(i, s') is empty. If s' has exactly i node descriptors with the ith having ID I, then ID-set$(i, s') = \{I\}$. Next, suppose that s' has more than i node descriptors.

- If $s' = s'', I$ and $I \in$ ID-set(i, s''), then ID-set(i, s') is defined to be ID-set$(i, s'')-\{I\}$. (ID I is now being used to label another node, and so is no longer in the ID-set of the ith node.)
- If $s' = s''$, add-ID(I, I') and $I \in$ ID-set(i, s''), then ID-set(i, s') is defined to be ID-set$(i, s'')\cup\{I'\}$. (Add I' to the ID-set of node i.)
- If $s' = s''$, add-ID(I', I) with $I \neq I'$ and $I \in$ ID-set(i, s''), then ID-set(i, s') is defined to be ID-set$(i, s'')-\{I\}$. (Again, ID I is now being used to label another node, and so is no longer in the ID-set of the ith node.)
- Otherwise, ID-set(i, s') = ID-set(i, s''). (No change to the ID-set of the ith node.)

Then, the edges of G are defined as follows: for each prefix of the form $s'(I, I')$ of s, if for some pair (i, j) of nodes of G, $I \in$ ID-set(i, s') and $I' \in$ID-set(j, s') then edge (i, j) is in G. Any extended graph descriptor can be converted in finite state to a (standard) graph descriptor without add-ID symbols, so both types of descriptor can be used interchangeably.

3.3 Checking for Cycles in a Bandwidth Bounded Graph

CLAIM 3.2. *There is a finite state cycle-checker that, given as input a k-graph descriptor, accepts if and only if the string represents an acyclic graph.*

As node and edge descriptors are read off from the input string, the cycle-checker maintains a so-called *active graph* containing at most $k + 1$ nodes, in which each node has a unique ID. The checker ignores edge labels and, upon reading a node ID or edge pair, does the following:

- When a node ID, say I, is read, then if there is a node with ID I in the active graph, for all pairs of edges $(H, I), (I, J)$ in the active graph (where H, I, J refer to node IDs) a new edge (H, J) is added, if not already in the graph. Then, the node with ID I and all incident edges are removed from the graph. Finally, a new node with ID I is added to the graph.
- When edge (I, I') is read, an edge is added from node with ID I to the node with ID I'. If addition of this edge introduces a cycle in the graph, the automaton rejects.

If, upon reaching the end of the string, the checker has not rejected, it accepts. Correctness of the checker follows from the fact that the node removal plus edge contraction done in the first test of the checker preserves cycles in the graph.

3.4 Observer-Checker Verification Method

In our method for protocol verification, the *observer* generates the same set of traces as the protocol, but augments each trace with a description of a k-bandwidth bounded graph. Given a run of the observer, the *checker* checks that the graph is an acyclic constraint graph for the trace.

Let $\mathcal{P}$ be a protocol. Let $\mathcal{A}$ be the set of LD and ST operations of $\mathcal{P}$. An *observer* for $\mathcal{P}$ is itself a finite state protocol. The alphabet (set of actions) of an observer consists of the symbols used in a k-graph descriptor for some k, in which the node label set is $\mathcal{A}$ and the edge label set $\mathcal{E}$ is {**inh, po, forced, STo, po-STo, po-inh, po-forced**}, where **inh**, **po**, **STo** and **forced** indicate inheritance, program order, ST order and forced edges, respectively, and **po-STo, po-inh**, and **po-forced** denote edges with two annotations. Note that each run of an observer contains a trace as a subsequence, namely the subsequence of symbols from $\mathcal{A}$.

An observer $\mathcal{O}$ for $\mathcal{P}$ is a *witness for* $\mathcal{P}$ if (i) the set of traces of $\mathcal{O}$ equals the set of traces of $\mathcal{P}$, and (ii) each run of $\mathcal{O}$ describes an acyclic constraint graph (as defined in section 3.1).

Checking property (i) can easily be reduced to the language equivalence problem for finite state automata. In practice, this check is trivial by construction, since the observer is a noninterfering augmentation of the protocol. The *checker* is a finite state automaton that checks property (ii). In what follows, assume that k (the bandwidth bound) is fixed. The alphabet of the checker equals that of the observer. Given as input a run r of observer $\mathcal{O}$, the checker does the following:

- Run the cycle-checker for k-bandwidth bounded graphs on r. If the cycle-checker rejects, then reject. Otherwise, r is an acyclic, k-bandwidth bounded graph.
- If not already rejected, check that edges satisfy the edge annotation properties listed in section 3.1. If so, then accept else reject.

By the definition of a witness in section 3.1, the checker accepts if and only if r describes an acyclic constraint graph. Also the checker has a finite number of states since the cycle checker of section 3.3 does, and the edge annotation checks are easy to do with a finite number of states. We now have the following theorem.

THEOREM 3.1. *Let $\mathcal{P}$, $\mathcal{O}$ be protocols. If $\mathcal{O}$ is a witness for $\mathcal{P}$, then $\mathcal{P}$ is sequentially consistent. Moreover, testing whether $\mathcal{O}$ is a witness for $\mathcal{P}$ can be reduced to the language inclusion problem for finite state automata.*

4. VERIFICATION OF REAL-WORLD PROTOCOLS

We claim that every real-world sequentially consistent protocol has a finite state witness observer and that the observer can be generated automatically from the protocol. To provide intuition that supports this claim, we first argue informally that a weaker property holds for real-world sequentially consistent protocols, namely that the witness graph

corresponding to each protocol run is bandwidth bounded. Later in this section we make this intuition precise, and also show the stronger property that the witness graph corresponding to each run is not only bandwidth bounded but can be generated in finite state from the run.

Let R be a run of a protocol and let R_1 be a prefix of R. Let R_2 be the corresponding suffix of R, so that $R = R_1R_2$. We need to show that if we view the operations of R as nodes of a constraint graph, the number of operations of R_1 with edges to operations of R_2 is bounded. We consider each type of edge in turn. It is easy to see that at most p operations of R_1 have program order edges to operations of R_2, namely the last operation in each processor's program order, if any.

We next consider inheritance edges; here we appeal to our understanding of how real-world sequentially consistent protocols work. These protocols create "views" of a block via ST operations, then copy these views into various protocol storage locations (such as queues, network message packets, or caches of other processors) where they can be read via the LD operation, and eventually delete or overwrite views. Multiple views of a block may exist in the protocol state. For example, one processor may do a ST to a block, thus creating a new view, while stale views of the block still exist in other caches. We call a ST operation of R_1 *inh-active* if one or more copies of the value (view) written by that ST is stored in the protocol state upon completion of run R_1. If a ST is inh-active, its value may be inherited by LDs in R_2. A key point is that, since the protocol is finite-state, only a constant number of STs of R_1 can be inh-active. Moreover, in real-world protocols, LDs of R_2 that inherit their values from STs of R_1 can only do so from STs of R_1 that are inh-active, because these LDs obtain their values from storage locations of the protocol.

Third, we consider ST order edges. Again, we appeal to a property of real-world protocols here, namely that for all runs, for each block B, the order of STs to B in the run is in fact the same as the order of the STs in the corresponding serial reordering. Thus, if we call ST nodes of R_1 with no outgoing ST order edge *STo-active* nodes, the number of STo-active nodes is at most the number of blocks b of the protocol. (Our class of verifiable protocols will actually be defined in section 4.2 to encompass protocols that do not satisfy this per-block real-time ST reordering property.)

Finally, we consider forced edges. The only LD nodes of R_1 that may have forced edges to STs of R_2 are those LDs which inherit their values from STo-active STs of R_1. For each STo-active operation S of R_1 and each processor P, at most one LD of processor P in R_1 need have a forced edge to a node in R_2, namely the last LD in P's program order that inherits its value from S. (This follows from edge constraint 5 of section 3.1.) Call such a LD operation a *forced-active* LD. Thus, the number of forced-active LDs of R_1 is at most pb. In addition, there may be ST nodes of R_1 that have incoming forced edges from LD nodes in R_2. Call these *forced-active* STs. Each forced-active ST is the immediate successor of an inh-active ST in ST order; thus, the number of forced-active STs is bounded by the number of inh-active STs, and therefore is bounded.

In section 4.1 we define a class of protocols for which the inheritance edges of a constraint graph can be generated in finite state. Protocols in this class have two properties, motivated by our informal arguments above. First, on a LD transition, the value of the LD is obtained from a known storage location of the protocol. Second, by tracking the movement of data among protocol storage locations, it is possible to automatically infer which ST conferred its value to each storage location. Then in section 4.2 we describe conditions under which the ST order edges of a constraint graph can be generated in finite state. In section 4.3, we define a class Γ of protocols that simultaneously satisfy the conditions of sections 4.1 and 4.2. We show that for protocols in Γ, the forced edges of a protocol run can also be generated in finite state, and conclude that such protocols have finite state observers.

4.1 Tracking Labels for Protocols

When a LD is performed by a protocol, how can we tell from which ST it inherits its value? We need to know from which storage location l of the protocol the LD gets its value, and which ST operation conferred its value to location l. We now descibe protocols with *tracking labels* which provide an automatic way to infer this knowledge. While real protocol descriptions do not explicitly have tracking labels, for all sequentially consistent protocols known to us, with an appropriate protocol description language the labeling could be generated automatically from the protocol description.

In a protocol with tracking labels, each state of the protocol records block values in at most L locations for some constant L (in caches, queues, and memory where blocks are stored). The tracking labels are of two types.

- Each transition in δ (where δ is the set of transitions on LD and ST operations) is labeled by a location identifier $l \in [1, L]$. Intuitively, the operation is read from or written to location l. Formally, the LD/ST tracking function is a mapping $f : \delta \to [1, L]$.
- For each transition t in δ' (where δ' is the set of transitions on actions other than LD and ST operations) and each $l \in [1, L]$, the copy tracking label, $c_l(t)$, indicates whether the value stored in location l is unchanged by the transition t or whether it has been copied from another location, namely $c_l(t)$. Formally, for each l, there is a copy tracking function $c_l : \delta' \to [1, L]$ (with $c_l(t) = l$ if the value is unchanged).

Intuitively, for every run R and location l of a protocol $\mathcal{P}$ with tracking labels, the ST index of l with respect to R is either 0 or is the index of the ST operation from which location l inherits its value upon completion of run R. Formally, the *ST index*, denoted by **ST-index**(R, l), can be defined inductively using the tracking labels as follows.

1. If $|R| = 0$ then **ST-index**$(R, l) = 0$.
2. If $R = R', A$, if the transition t taken on A is a ST operation with tracking label l, and if A is the ith trace operation of R, then **ST-index**$(R, l) = i$. Otherwise,

if A is not a LD or ST operation then **ST-index**(R, l) = **ST-index**$(R', c_l(t))$. Otherwise, **ST-index**(R, l) = **ST-index**(R', l).

Example: An example to illustrate ST indexes and tracking labels is given in Figure 3. This example describes a run of an extremely simple protocol with two processors, $P1$ and $P2$, and three blocks, $B1, B2$, and $B3$. Each processor has two cache locations in which values of blocks can be stored (part (a) of the figure). Thus, there are four locations in all: $P1$'s locations are numbered 1 and 2, and $P2$'s locations are numbered 3 and 4. In the illustration, each location contains information about which block is being stored there, if any, and what its value is. Thus, block $B1$ with value 1 is stored (by $P2$) in location 3, whereas location 2 is undefined.

The location values reflect the protocol state at the end of the run R given in part (b) of the figure. R is of length four and has three ST operations and one "Get-Shared" operation. The Get-Shared operation causes the value of $B1$ stored in location 1 by $P1$ after the first action of R to be copied to location 3 of $P2$; it is reminiscent of how values of blocks can be shared or copied in real protocols, albeit highly simplified. The tracking label of each transition corresponding to each action in run R is also given. The first operation of R, ST$(P1, B1, 1)$ has tracking label 1, indicating that $B1$'s value is written in location 1. The second operation, ST$(P2, B2, 2)$, has tracking label 4; thus $B2$'s value is written into location 4. The third action of R is not a LD or ST operation and so there are four copy tracking labels $c_1, \ldots, c_4$ associated with this action, one per location. Note that $c_3 = 1$ since the value now stored in location 3 is copied from location 1, but $c_i = i$ for $i = 1, 2$, and 4, since the contents of locations 1, 2, and 4 are unchanged by the Get-Shared action. The last operation of R, ST$(P1, B3, 3)$, has tracking label 1 indicating that block $B1$ is overwritten by $B3$ in location 1. Thus, upon completion of run R, the ST index of each location is given by part (c) of the figure. □

Let R', LD(P, V, B) be a prefix of R in which the LD(P, V, B) operation is the jth trace operation of R. Intuitively, if the LD operation gets its value from location l and location l inherits its value from the ith trace operation of R (which must be a ST operation), then (i, j) is an inheritance edge. More precisely, let t be the transition taken on the LD operation, and let the tracking label of t be l. Then, if **ST-index**(R', l) $\neq 0$ the edge (**ST-index**$(R', l), j)$ is an *inheritance edge* of R.

For any run R of a protocol with tracking functions f and c_l, $1 \leq l \leq L$, let the *inheritance graph* of R with respect to these tracking functions be the graph whose nodes are the trace operations of R, numbered by their order in R, and whose edges are the inheritance edges of R. This graph is L-bandwidth bounded, where L is the total number of locations in a state of the protocol. This is because, for any prefix R' of R, at most L ST operations are "active", in the sense that they are indexed in the set {**ST-index**(R', l)} and thus may be in future inheritance edges. Indeed, we have the following claim.

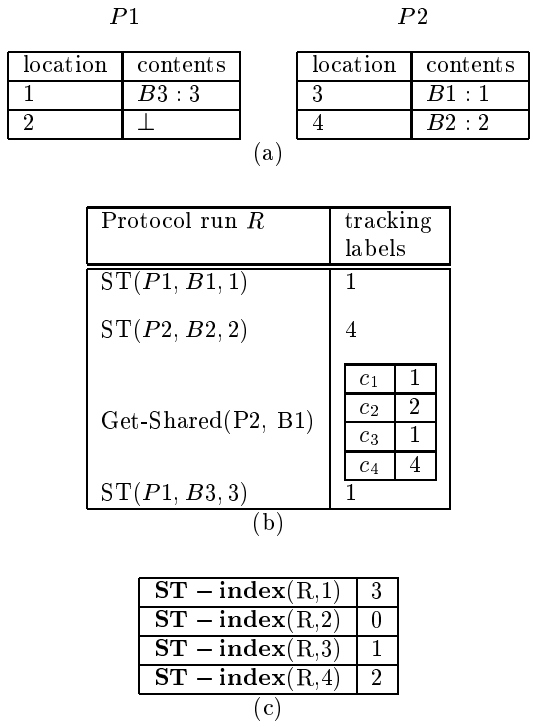

Figure 3: ST Index Example. Part (a) depicts the state of four protocol locations, where locations 1 and 2 correspond to cache lines of processor $P1$ and locations 3 and 4 correspond to cache lines of processor $P2$. Location 2 is empty, and each of the other locations stores the value of one of blocks $B1$, $B2$, or $B3$. Part (b) lists an example run R of length 4, in which the Get-Shared action copies block $B1$ from location 1 to location 3. Also, the tracking labels of each transition corresponding to the actions of R are given. The state of the protocol in part (a) represents the state upon completion of run R. Part (c) Lists the ST-index of each location with respect to run R.

CLAIM 4.1. *Let $\mathcal{P}$ be a protocol with L locations and tracking functions $f, \{c_l\}$. There is a finite state automaton that, given a run R of $\mathcal{P}$, generates a descriptor of the inheritance graph of R.*

The generator generates the graph while executing the protocol on run R, and outputs an extended graph descriptor. Upon transition $t = (q, A, q')$, the generator does the following:

- If A is a ST operation and t has tracking label l then output "l, A". (Recall that this adds a new node to the graph with ID l and label A.)
- For each l, if $c_l(t) \neq l$ then output "add-ID$(c_l(t), l)$". (Intuitively, the ST node with ID $c_l(t)$ is being copied to location l, so l is added to the set of IDs for this ST node. More generally, the number of IDs of a ST node equals the number of copies of the ST in the protocol state.)

- If A is a LD operation and t has tracking label l then output "$L+1, A, (l, L+1)$, inh". (This causes a new node with ID $L+1$, labeled A, to be added to the graph, and an inheritance edge to be added into A.)

4.2 Finite State ST Reordering

Let R be a run of protocol $\mathcal{P}$. A *ST order graph* for R is a graph whose nodes are the trace operations of R, numbered by their order in R. As in section 3.1, for each block B, if there are u ST operations to B in R then there are $u-1$ ST order edges in the graph which define a total order on these u operations.

A *ST order generator* for $\mathcal{P}$ is a finite state automaton that, given run R as input, generates a k-graph descriptor that describes the ST order graph, for some k. Moreover, the number of states of the automaton is at most the number of states of $\mathcal{P}$.

Protocols implemented in practice have the *real-time ST reordering* property that for all traces, for each block B, the trace order of STs to B is in fact the same as the corresponding serial reordering. Thus, the ST order generator is trivial. One well-known protocol that does require non-trivial (but still finite state) ST ordering is the lazy caching protocol of Afek et al. [2], but this protocol has not been implemented in a real machine.

4.3 The Γ Protocol Class

Let $\mathcal{P}$ be a protocol. Let $f, \{c_l, 1 \leq l \leq L\}$ be tracking functions and let $\mathcal{G}$ be a ST order generator. With respect to $f, \{c_l\}$, and $\mathcal{G}$, for each run R of $\mathcal{P}$, let $W(R)$ be the graph whose nodes are the trace operations of R. The edges of $W(R)$ are the inheritance edges of the inheritance graph with respect to f and $\{c_l\}$, the ST order edges given by $\mathcal{G}$, the forced edges implied by these inheritance and ST order edges, and the program order edges given by the order of operations in R.

DEFINITION 4.1. *A protocol $\mathcal{P}$ belongs to the class Γ if for some tracking functions $f, \{c_l\}$ and some ST order generator $\mathcal{G}$, for all runs R of $\mathcal{P}$, the graph $W(R)$ is an acyclic constraint graph.*

THEOREM 4.1. *Every protocol in Γ has a finite state witness observer.*

PROOF. We describe a finite state observer $\mathcal{O}$ that, given $\mathcal{P}$ in Γ, along with associated tracking functions $f, \{c_l\}$ and ST order generator $\mathcal{G}$, converts a run R of $\mathcal{P}$ into a descriptor for a constraint graph $W(R)$.

$\mathcal{O}$ adds each LD and ST operation of R to the graph as the operation is read. From Claim 4.1 and section 4.2, the inheritance and ST order edges can be generated in finite state. It is also trivial to generate the program order edges.

It remains to extend the observer so that forced edges are also generated. For this purpose, each node N' labeled by a LD(P, B, V) operation remains in the active graph maintained by the observer until one of the following events occurs. Let the inheritance edge to N' be from node N. (i) Another node, N'', labeled by LD(P, B, V) is added to the graph, along with inheritance edge (N, N''). Node N' can now be removed because there is a path of program order edges from the N' to N''. (ii) A ST order edge from N, say to node S, is present in the graph. In this case, a forced edge is added from N' to S.

The number of LD nodes that need to be in the active graph for the purpose of generating forced edges is bounded by p (the number of processors) times the number of ST nodes with no outgoing ST order edges. The latter number is bounded, since the ST order graph is bandwidth bounded. In addition, if ST node S has an incoming ST order edge (N, S) where the value of the ST labeling N may be read by future LDs, then S must be maintained in the active graph. The number of such ST nodes S is at most L.

Thus, the witness graph is bandwidth bounded, where the bound depends only on $\mathcal{G}$, L, p, and b and does not otherwise depend on R, and so the observer is finite state. □

To summarize, we have shown the following. Let $\mathcal{P}$ be a protocol for which tracking labels can be generated automatically and the real-time ST reordering property holds (or more generally, for which a ST order generator exists). Then, sequential consistency can be verified by an algorithm that first generates the observer from the protocol in a noninterfering fashion (so that the the set of traces of the observer equals those of the protocol) and then uses a model checker (based on our cycle-checker) to verify that every graph descriptor generated by the observer describes an acyclic constraint graph. Note that the checker is independent of the protocol.

4.4 Size of Observer

In order to apply our constraint graph method to the verification of a protocol, the major obstacle will be the size of the observer. In addition to the protocol state, the observer needs to maintain in its state a subgraph of the constraint graph that may have a number of nodes up to the bandwidth bound of that graph. Here, we describe an upper bound on the number of bits of extra state required by the observer, under reasonable assumptions.

First, we bound the bandwidth of the constraint graphs of a protocol $\mathcal{P}$ with L locations. We consider here the case that the protocol has real-time ST ordering, and that the value of a ST is stored in some protocol location at least until the ST following it in ST order has been done. In this case, with respect to a prefix of a run, at most L distinct ST nodes may be actively stored in protocol locations and thus may have future outgoing inheritance edges. Up to pb LD nodes may contribute to the bandwidth needed for generating forced edges. Nodes needed for generation of program order edges and ST order edges are already counted among these nodes, so the total bandwidth is bounded by $L + pb$.

For each active node of the constraint graph, the node label must be stored. This requires up to $\lg p + \lg b + \lg v + 1$ bits.

Here lg denotes the ceiling of log to the base 2; 1 bit indicates whether the label is a LD or ST, and parameters P, B, and V are represented using the other bits. Also, IDs for each ST node are needed, in order to generate inheritance edges. An addition $L \lg L$ bits are needed to store IDs.

Edges of the constraint graph must also be represented. If the active nodes are stored in a linear array, no extra storage is needed for edges. Roughly, this is because the nodes can be stored in an order consistent with the partial order of the constraint graph, so that graph edges can be inferred. For example, in the linear array order, a ST to block B is followed (not necessarily contiguously) by LD nodes that inherit its value, and no other ST to the same block separates them, so inheritance edges are completely determined by the linear order.

Thus, an upper bound on the number of bits of extra state needed by the observer (in addition to the protocol state) is $(L + pb)(\lg p + \lg b + \lg v + 1) + L \lg L$ bits. This upper bound is likely to be substantially less than the number of bits in the protocol itself. Real memory system protocols, however, are already roughly at the limits of current model checking tools, so any additional state is problematic in practice. Fortunately, some simple optimizations should help to reduce the size of the observer. For example, the value of a node is needed only to check that each LD gets the same value as the ST from which it supposedly inherits its value. This check can be done independently from the cycle-testing check, thereby saving $\lg v$ bits per node.

5. FUTURE WORK

Understanding how the size of the observer can be reduced, perhaps by imposing further assumptions on the class of protocols to be handled, is an important direction for future work from a practical point of view, and will help to relate this work to that of Qadeer [13]. Extending these techniques to other memory models is another important direction of this research.

Experimental results will be needed to assess the applicability of our results in practice. We intend to apply our techniques to substantial memory system protocols using model checking tools and explore means to combat state explosion.

An interesting theoretical question is whether the problem of testing sequential consistency is undecidable for protocols that are bandwidth bounded. The reduction used in the undecidability result of Alur et al. [3] exploits protocols that are not bandwidth bounded.

Finally, we note that our method can also be used for testing that a particular run of a protocol does not violate sequential consistency, building on the approach proposed by Gibbons and Korach [6]. The finite-state observer and checker could be simulated together with detailed implementation descriptions that are too complex for formal verification.

Acknowledgments

We thank Mark Hill, Dan Sorin, Manoj Plakal and the other members of the Wisconsin Multifacet group for sharing their insights and intuition about proving sequential consistency.

6. REFERENCES

[1] Sarita V. Adve and Kourosh Gharachorloo. Shared memory consistency models: A tutorial. *IEEE Computer*, pages 66–76, December 1996.

[2] Yehuda Afek, Geoffrey Brown, and Michael Merritt. Lazy caching. *ACM Transactions on Programming Languages and Systems*, 15(1), January 1993.

[3] Rajeev Alur, Ken McMillan, and Doron Peled. Model-checking of correctness conditions for concurrent objects. In *Eleventh Symposium on Logic in Computer Science*, pages 219–228. IEEE, 1996.

[4] Tim Braun, Anne E. Condon, Alan J. Hu, Kai S. Juse, Marius Laza, Michael Leslie, and Rita Sharma. Proving sequential consistency by model checking. Technical Report TR-2001-03, Department of Computer Science, University of British Columbia, April 2001.

[5] Edmund M. Clarke and E. Allen Emerson. Design and synthesis of synchronization skeletons using branching time temporal logic. In Dexter Kozen, editor, *Workshop on Logics of Programs*, pages 52–71, May 1981. Published 1982 as Lecture Notes in Computer Science Number 131.

[6] Phillip B. Gibbons and Ephraim Korach. Testing shared memories. *SIAM Journal on Computing*, 26(4):1208–1244, August 1997.

[7] Thomas A. Henzinger, Shaz Qadeer, and Sriram K. Rajamani. Verifying sequential consistency on shared-memory multiprocessor systems. In *Computer-Aided Verification: 11th International Conference*, pages 301–315. Springer, 1999. Lecture Notes in Computer Science Vol. 1633.

[8] Mark D. Hill. Multiprocessors should support simple memory-consistency models. *IEEE Computer*, pages 28–34, August 1998.

[9] Leslie Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *ACM Transactions on Computer*, 28(9):690–691, September 1979.

[10] Marius Laza, Rita Sharma, Anne Condon, and Alan J. Hu. Protocols for which proving sequential consistency is easy. In *Workshop on Formal Specification and Verification Methods for Shared Memory Systems*. Unpublished Proceedings, October 31, 2000. Workshop affiliated with FMCAD 2000, Austin, TX.

[11] Ratan Nalumasu, Rajnish Ghughal, Abdel Mokkedem, and Ganesh Gopalakrishnan. The 'test model-checking' approach to the verification of formal memory models of multiprocessors. In *Computer-Aided Verification: 10th International Conference*, pages 464–476. Springer, 1998. Lecture Notes in Computer Science Vol. 1427.

[12] M. Plakal, D. Sorin, A. Condon, and M. Hill. Lamport Clocks: Verifying a directory cache coherence protocol. In *Symposium on Parallel Algorithms and Architectures*, pages 67–76, 1998.

[13] Shaz Qadeer. On the verification of memory models of shared-memory multiprocessors. In *Workshop on Formal Specification and Verification Methods for Shared Memory Systems*. Unpublished Proceedings, October 31, 2000. Workshop affiliated with FMCAD 2000, Austin, TX.