

Predicate Abstraction for Relaxed Memory Models

Andrei Dan¹, Yuri Meshman², Martin Vechev¹, and Eran Yahav²

¹ ETH Zurich

{andrei.dan,martin.vechev}@inf.ethz.ch

² Technion

{yurime,yahave}@cs.technion.ac.il

Abstract. We present a novel approach for predicate abstraction of programs running on relaxed memory models. Our approach consists of two steps.

First, we reduce the problem of verifying a program P running on a memory model M to the problem of verifying a program P_M that captures an abstraction of M as part of the program.

Second, we present a new technique for discovering predicates that enable verification of P_M . The core idea is to extrapolate from the predicates used to verify P under sequential consistency. A key new concept is that of cube extrapolation: it successfully avoids exponential state explosion when abstracting P_M .

We implemented our approach for the x86 TSO and PSO memory models and showed that predicates discovered via extrapolation are powerful enough to verify several challenging concurrent programs. This is the first time some of these programs have been verified for a model as relaxed as PSO.

1 Introduction

One approach for efficiently utilizing multi-core architectures, used by major CPU designs (e.g., [27,28,20]), is to define architectural *relaxed (weak) memory models (RMMs)* [14]. Some of those relaxations can be modeled using one or more per-processor FIFO buffers, where a store operation adds values to the buffer and a flush operation propagates the stored value to main memory. Programs running under those models exhibit unique caveats and verifying their correctness is challenging.

The Problem. Given a program P , a specification S and a memory model M , we would like to answer whether P satisfies S under M , denoted as $P \models_M S$.

Unfortunately, even for finite-state programs, automatic verification under relaxed memory models is a hard problem. The problem is either undecidable or has a non-primitive recursive complexity for stronger models such as x86 TSO and PSO (see [4] for details). It is therefore natural to explore the use of abstraction for verification of such programs.

Predicate abstraction [16] is a widely used approach for abstract interpretation [9]. Since predicate abstraction has been successfully applied to verify a wide range of sequential and concurrent programs (e.g., [6,13,15]), we are interested in the question:

how to apply predicate abstraction to programs running on relaxed models?

Given a program P and a vocabulary (set of predicates) $V = \{p_1, \dots, p_n\}$ with corresponding boolean variables $\hat{V} = \{b_1, \dots, b_n\}$, standard predicate abstraction (e.g. [16,6]) constructs a boolean program $\mathcal{BP}(P, V)$ that conservatively represents the behaviors of P using only boolean variables from \hat{V} . When considering predicate abstraction in the context of relaxed memory models, two key challenges need to be addressed: (i) *soundness*: the boolean program must faithfully abstract the behaviors of P running on model M ; (ii) *predicate discovery*: there should be a mechanism for automatically discovering predicates that enable successful verification of P running on memory model M .

Soundness. Under sequential consistency (SC), predicate abstraction is sound and we know that $\mathcal{BP}(P, V) \models_{SC} S$ implies $P \models_{SC} S$. Unfortunately, we observed this does not hold for relaxed memory models (see Section 2.2).

Intuitively, the problem is as follows: under sequential consistency, a shared variable has only one value — the value stored in main memory. Predicates used for predicate abstraction can therefore refer to that shared value and relate it to the values of other shared variables or thread-local variables. In contrast, for a program running on a relaxed model, threads may observe *different values for the same shared variable* as they have their own local buffered copies. This means that in a relaxed model, one cannot directly apply classic predicate abstraction, as the variables used in predicates are assumed to refer to a single value at a time.

Predicate Discovery. A key challenge with predicate abstraction is to discover a set of predicates that enable verification. Following classic abstraction refinement, one would start with a program that is to be verified on a particular relaxed model together with an initial set of predicates. Then, proceed to iteratively apply refinement until we find a set of predicates under which the program verifies (or the process times out).

We take a different approach to predicate discovery for programs running on RMMs. In our approach, we first obtain the predicates that enable verification of the program on sequential consistency (SC). Then, we automatically *extrapolate* from these SC predicates to produce a new set of predicates that can be used as a basis for verification on the relaxed model.

Our Approach. Given a program P , a specification S and a memory model M , our approach consists of the following steps:

1. *verify under SC*: find a set of predicates V , sufficient to verify P under sequential consistency, i.e., a set V such that $\mathcal{BP}(P, V) \models_{SC} S$.
2. *reduce to SC*: automatically construct a new program P_M such that if $P_M \models_{SC} S$ then $P \models_M S$. The program P_M contains an abstraction of the store buffers used in M .
3. *discover new predicates*: automatically compute a new set of predicates V_M that are used for predicate abstraction of P_M . This is a challenging step and the key idea is to leverage the verification of P under SC. We present two approaches: predicate extrapolation which discovers new predicates based on the predicates in V and cube extrapolation which discovers new predicates based on both V and $\mathcal{BP}(P, V)$.

4. *construct a new boolean program*: given the new program P_M and the new predicates V_M , automatically construct a boolean program $\mathcal{BP}(P_M, V_M)$ such that $\mathcal{BP}(P_M, V_M) \models_{SC} S$ ensures that $P_M \models_{SC} S$, which in turn guarantees that $P \models_M S$. Here, cube extrapolation enables us to build $\mathcal{BP}(P_M, V_M)$ *without* suffering from the usual problem of exponential search.
5. *check*: whether $\mathcal{BP}(P_M, V_M) \models_{SC} S$.

Main Contributions

- We provide a novel approach for predicate abstraction of programs running on relaxed memory models, extrapolating from the predicate abstraction proof of the same program for sequential consistency.
- One of our insights is that the predicates used to verify P under SC can be automatically *extrapolated* to discover new predicates for verification of the program with M encoded in it, P_M . We present two approaches for discovering new predicates called predicate extrapolation and cube extrapolation.
- We instantiated our approach for the x86 TSO and PSO memory models. We implemented our approach and applied it to verify several challenging concurrent algorithms (both finite and infinite-state) under these models. We show that extrapolation is powerful enough to verify these algorithms and in particular, cube extrapolation enables verification of Lamport’s Bakery algorithm, which otherwise (without cube extrapolation) times out when building the boolean program.

2 Overview

In this section, we give an informal overview of our approach using simple examples.

2.1 Motivating Example

Fig. 1 shows an implementation of an infinite state alternating bit protocol (ABP) with two concurrent threads. We use capitalized variable names to denote global shared variables, and variable names in lowercase to denote local variables. In this program, global shared variables `MSG` and `ACK` have an initial value 0. We use this algorithm as our illustrative example, additional examples are discussed in Section 6.

Specification. When executing on a sequentially consistent memory model, this program satisfies the invariant:

$$((lRCnt = lSCnt) \vee ((lRCnt + 1) = lSCnt))$$

Here, the local variable $lRCnt$ is the local counter for the receiver thread containing the number of received messages. Similarly, local variable $lSCnt$ is the local counter for the sender thread containing the number of sent messages.

Predicate Abstraction under Sequential Consistency. A traditional approach to predicate abstraction is shown in Fig. 2(a). To verify that ABP satisfies its specification under SC, we instantiate predicate abstraction with the following predicates:

$$(Msg = 0), (Ack = 0), (lSt = 0), (lAck = 0) \\ (lMsg = 0), (lRSt = 0), (lRCnt = lSCnt), ((lRCnt + 1) = lSCnt)$$

```

initially: Msg = Ack = 0
Sender (thread 0):
1 lAck = Ack;
2 if ((lAck = 0 & lSSt = 0)
   | (lAck != 0 & lSSt != 0))
3
4     if (lSSt != 0) lSSt = 0;
5     else lSSt = 1;
6     lSCnt++;
7     Msg = lSSt;
8 goto 1;

Receiver (thread 1):
1 lMsg = Msg;
2 if ((lMsg = 0 & lRSt != 0)
   | (lMsg = 0 & lRSt != 0))
3
4     lRSt = lMsg;
5     lRCnt++;
6     Ack = lRSt;
7
8 goto 1;

```

Fig. 1. An alternating bit protocol example with two threads

The result of predicate abstraction using these predicates is a concurrent boolean program that conservatively represents all behaviors of the original ABP program. In Fig. 2(a), this boolean program is denoted as the oval named Boolean Program B.

In the worst case, the construction of the concurrent boolean program, following standard predicate abstraction techniques (e.g., [16,6,15]) involves an exponential number of calls to an underlying theorem prover. A critical part of the construction of the boolean program is searching the “cubes” — conjunctions of predicates — that imply a certain condition. This search is exponential in the number of predicates.

An Informal Overview of PSO. In the partial-store-order (PSO) memory model, each thread maintains a private buffer (a sequence) for each global variable. When the thread writes to a global variable, the value is enqueued into the buffer for that variable. Non-deterministically, the values can be dequeued from the buffer and written to main memory. When the thread reads from a global variable, it first checks if the buffer is empty and if so, it reads as usual from main memory. Otherwise, it reads the last value written in the buffer. The model is further equipped with a special *fence* instruction which can be executed by a thread to empty all buffers for that thread and write the most recent value in each buffer to the corresponding shared location. In our example, thread 0 maintains one buffer, the buffer for global variable *Msg* and thread 1 also maintains one buffer, the buffer for global variable *Ack*. For instance, the store $Msg = lSSt$ leads to the value of *lSSt* being enqueued into the private buffer for thread 0. This value can be flushed to main memory at any point in the execution, non-deterministically.

An Informal Overview of TSO. In the total-store-order (TSO) memory model, each thread maintains a single private buffer for all global variables. Similarly to PSO, values stored to global variables are first written to the buffer, and then are non-deterministically flushed to main memory. A *fence* instruction empties the thread’s private buffer.

The challenges we address are: (i) how to verify programs such as ABP under relaxed memory models such as x86 TSO and PSO, and (ii) how to deal with the exponential complexity of standard predicate abstraction in our setting of RMM.




Fig. 2. Predicate abstraction: (a) classic algorithm, (b) with predicate extrapolation and cube extrapolation. Here, a rectangle shape represents an algorithm, while an oval shape represents input-output information (data).

2.2 Standard (Naive) Predicate Abstraction under RMM Is Unsound

We now consider the following scheme for predicate abstraction: (i) construct a boolean program directly from the program of Fig. 3 using standard predicate abstraction; (ii) execute the boolean program using PSO semantics. This scheme is simple, but unfortunately, it is unsound. We now show that following this scheme, we can produce a boolean program that successfully verifies our assertion. This is unsound because we know the existence of an assertion violation, namely the one shown in Fig. 4(a).

```

initial: X=Y=0
Thread 0:
1 X = Y+1
2 fence(X)
Thread 1:
1 Y = X+1
2 fence(Y)
assert (X≠Y)
    
```

Fig. 3. Unsoundness example

To capture the assertion breach, we need to keep track of the relationship between X and Y . Consider the predicates:

$$P_1 : X = Y, P_2 : X = 1, P_3 : Y = 1, P_4 : X = 0, P_5 : Y = 0$$

Each P_i has been assigned a boolean variable B_i (with a buffer for each thread) where the effect of a thread writing to X or Y will write to a local buffer of $X = Y$ of that thread, and the effect of flushing X or Y will also flush the buffer associated with $X = Y$. Unfortunately this approach is insufficient. The problem is shown in Fig. 4(b). When thread 0 updates X to 1, and thread 1 updates Y to 1, the predicate $X = Y$ will be updated to *false* in both (denoted as F) and stored in the store buffer of each thread

(a) Concrete			(b) Predicate Abstraction		
Thread 0	Thread 1	Global	Thread 0	Thread 1	Global
(X,Y)	(X,Y)	(X,Y)	(X=Y,X=1,Y=1,X=0,Y=0)	(X=Y,X=1,Y=1,X=0,Y=0)	(X=Y,X=1,Y=1,X=0,Y=0)
(0,0)	(0,0)	(0,0)	(T, F, F, T, T)	(T, F, F, T, T)	(T, F, F, T, T)
T0: X = Y+1	(1,0)	(0,0)	(F, T, F, F, T)	(T, F, F, T, T)	(T, F, F, T, T)
T1: Y = X+1	(1,0)	(0,0)	(F, T, F, F, T)	(F, F, T, T, F)	(T, F, F, T, T)
T0: flush(X)	(1,0)	(1,0)	(F, T, F, F, T)	(F, T, T, F, F)	(F, T, F, F, T)
T1: flush(Y)	(1,1)	(1,1)	(F, T, F, F, T)	(F, T, T, F, F)	(F, T, T, F, F)

Fig. 4. (a) an error trace for Fig. 3 under PSO. Thread i are values observed by thread i , Global are the values in global memory. (b) Values of predicates in a boolean program corresponding to the program of Fig. 3 under PSO semantics.

(since neither of the two threads can see the update of the other). When the value of X is flushed from the buffer of thread 0, our setting flushes the value of the predicate $X = Y$ (F) to main memory. Similarly, when Y is flushed in thread 1, the value of $X = Y$ (F) is flushed. The main memory state after these two flushes is inconsistent, as it has $X = 1$ set to T, $Y = 1$ set to T and $X = Y$ set to F.

2.3 Predicate Abstraction for Relaxed Memory Models

In Fig. 2(b), we illustrate the ingredients and flow of our approach for solving the verification problem under relaxed memory models. The figure contains two approaches for adapting predicate abstraction to our setting called Predicate Extrapolation and Cube Extrapolation (which includes Predicate Extrapolation). Next, we discuss the steps of our approach.

Step 1: Verify P under SC. The first step is to verify the program P under sequential consistency using standard predicate abstraction as outlined earlier in Fig. 2(a). Once the program is verified, we can leverage its set of predicates as well as its boolean program in the following steps.

Step 2: Construct the reduced program P_M . This step is named “Reduction” in Fig. 2(b). To enable sound predicate abstraction of a program P under a relaxed memory model M , we first reduce the problem into predicate abstraction of a sequentially consistent program. We do so, by constructing a program P_M that conservatively represents the memory model M effects as part of the program.

The key idea in constructing P_M is to represent an abstraction of the store buffers of M as additional variables in P_M . Since the constructed program P_M represents (an abstraction of) the details of the underlying memory model, we can soundly apply predicate abstraction to P_M . The formal details of the reduction for x86 TSO and PSO are discussed later in the paper. Here, we give an informal description.

For PSO, it is sufficient to consider a program P_{PSO} where every global variable X in P is also associated with: (i) additional k local variables for each thread t : $x_{1,t}, \dots, x_{k,t}$, representing the content of a local store buffer for this variable in each thread t , (ii) a buffer counter variable $x_{cnt,t}$ that records the current position in the store buffer of X in thread t .

The x86 TSO model maintains a single local buffer per process. This buffer is updated with stores to *any* of the global variables. However, we need additional

variables to capture information about which global variable is stored in the buffer. The lhs variables contain the index of which global variable is addressed for each buffer element. The other variables are similar to PSO: (i) k local variables for each thread t : $lhs_{1,t}, \dots, lhs_{k,t}$, representing the index of the global variable stored at a local store buffer in each thread t , (ii) k local variables for each thread t : $rhs_{1,t}, \dots, rhs_{k,t}$, representing the value content of a local store buffer in each thread t , (iii) a buffer counter variable cnt_t that records the current position in the store buffer of thread t .

Step 3: Discover new predicates for P_M After verifying P under SC and constructing P_M , the remaining challenge is to discover a sufficient set of predicates for verifying that P_M satisfies a given specification. One of our main insights is that for buffered memory models such as x86 TSO and PSO, the predicates (and boolean program) used for verifying P under SC can be automatically leveraged to enable verification of P_M . This step corresponds to the “Extrapolate” box. This step takes as input the set of predicates V that were successful for verifying P under SC and outputs a new set V_M .

Predicates for the motivating example. Next, we illustrate via our running example how to generate new predicates under PSO (the process under x86 TSO is similar).

Consider again the ABP algorithm of Fig. 1 and the 8 predicates listed earlier in Section 2.1. Following the structure of the additional variables in P_M , we can introduce additional predicates by cloning each predicate over a global variable X into new predicates over store-buffer variables $x_{1,t}, \dots, x_{k,t}$. For example, assuming $k = 1$, in addition to $Msg = 0$, we introduce $Msg_{1,t0} = 0$.

To keep track of the buffer size for each buffered variable, we introduce additional predicates. For instance, for a global variable Msg , to register possible values for $Msg_{cnt,t0}$, assuming $k = 1$, we introduce $Msg_{cnt,t0} = 0$ and $Msg_{cnt,t0} = 1$. Note that we do not introduce predicates such as $Msg_{cnt,t1} = 0$ and $Msg_{1,t1} = 0$ as thread 1 always accesses Msg via main memory. Another predicate we could have added is $Msg_{1,t0} = Msg$. This predicate is not needed for the verification of ABP, however, we observe that such predicates can greatly reduce the state space of the model checker. In Section 5, we discuss rules and optimizations for generating new predicates for the PSO memory model and in [10] we present the details for the x86 TSO model (which are similar).

Finally, to ensure soundness, we add a designated flag *overflow* to track when the buffer size grows beyond our predetermined bound k . Overall, with a bound of $k = 1$, from 8 predicates used for sequential consistency we generate 15 predicates for PSO:

$$(Msg = 0), (Ack = 0), (lSSt = 0), (lAck = 0), (lMsg = 0), (lRSt = 0), (lRCnt = lSCnt), \\ ((lRCnt + 1) = lSCnt), (Msg_{cnt,t0} = 0), (Msg_{cnt,t0} = 1), (Ack_{cnt,t1} = 0), \\ (Ack_{cnt,t1} = 1), (Msg_{1,t0} = 0), (Ack_{1,t1} = 0), (overflow = 0)$$

Cube Search Space. A critical part of the predicate abstraction algorithm is finding the weakest disjunction of cubes that implies a given formula (see [13,6] for details). This is done by exploring the cube search space, typically ordered by cube size. Because the search is exponential in the number of predicates, most previous work on predicate abstraction has bounded the cube space by limiting the maximal cube size to 3.

The translation of SC predicates to PSO predicates implies a polynomial (in the buffer limit k , number of threads and number of shared variables) growth in the number

of predicates. Unfortunately, since cube search space is exponential in the number of predicates, exploration of the PSO cube space can sometimes be prohibitively expensive in practice.

For example, for ABP running on sequential consistency, the total number of cubes is $3^8 = 6561$. Here, 8 is the maximal sized cube which is the same as the number of predicates. And 3 means that we can use the predicate directly, its negation or the predicate can be absent. However, for PSO, the total number of cubes is 3^{15} exceeding 14 million cubes! If we limit cube size to 4, the SC cube search space becomes bounded by $\sum_{i=1}^4 2^i \binom{8}{i} = 288$, and the PSO cube space to be explored becomes bounded by $\sum_{i=1}^4 2^i \binom{15}{i} = 25630$.

The situation is worsened as the cube space is explored for every abstract transformer computation. Further, while in previous work, which mostly targets sequential programs, limiting the cube size to 3 seemed to work, with concurrency, where one needs to capture correlations between different threads, it is possible that we need a cube size of 4 or greater. As the number of predicates increases, directly exploring cubes of size 4 or more, even with standard optimizations, becomes infeasible (our experiments confirm that).

Reducing the PSO Cube Space using SC Cubes. One of our main insights is that we can leverage the boolean program (the proof) under sequential consistency to simplify reasoning under relaxed memory models. Technically, we realize this insight by using the cubes from the boolean program under SC in order to guide the search in the cube space under the weak memory model.

In Fig. 2(b), this step is denoted under Cube Extrapolation where in addition to the steps in Predicate Extrapolation, we also extract the cubes that appear in the boolean program of P under SC.

For example, for ABP, we examine the boolean program $\mathcal{BP}(ABP, V)$ where V are the eight predicates listed earlier in Section 2.1, and observe the following cubes:

$$\begin{aligned} c_1 &= (lSt = 0) \wedge (lAck = 0) \\ c_2 &= (lSt = 0) \wedge \neg(lAck = 0) \\ c_3 &= (lMsg = 0) \wedge \neg(lRSt = 0) \\ c_4 &= \neg(lSt = 0) \wedge \neg(lAck = 0) \\ c_5 &= \neg(lSt = 0) \wedge (lAck = 0) \end{aligned}$$

Since these cubes do not use the two global variables Msg and Ack , it stands to reason that the same cubes would be obtained from the PSO cube space exploration, which is indeed the case.

In this simple example, the above cubes from the predicate abstraction under SC could be used directly for the predicate abstraction under PSO, *without* needing to search for these cubes. Of course, there are cases where SC cubes do contain buffered variables (such as Msg or Ack) and in that case we need to extrapolate from these cubes in order to obtain useful cubes under PSO (see Section 5).

Building the Boolean Program. An enabling factor with cube extrapolation is that it changes the way we build the boolean program. We no longer require exponential search over the cube search space. In Section 4, we show how to use the cubes under SC as constraints over the cube space to reduce the size of the cube space to be explored under

$\llbracket X = r \rrbracket_k^t$	$\llbracket r = X \rrbracket_k^t$	$\llbracket \text{fence} \rrbracket_k^t$	$\llbracket \text{flush} \rrbracket_k^t$
<pre> if $x_{cnt,t} = k$ then abort("overflow") $x_{cnt,t} = x_{cnt,t} + 1$ if $x_{cnt,t} = 1$ then $x_{1,t} = r$... if $x_{cnt,t} = k$ then $x_{k,t} = r$ </pre>	<pre> if $x_{cnt,t} = 0$ then $r = X$ if $x_{cnt,t} = 1$ then $r = x_{1,t}$... if $x_{cnt,t} = k$ then $r = x_{k,t}$ </pre>	<pre> \triangleright for each $X \in Gvar$ generate: assume ($x_{cnt,t} = 0$) \triangleright end of generation </pre>	<pre> while * do \triangleright for each $X \in Gvar$ generate: if $x_{cnt,t} > 0$ then if * then $X = x_{1,t}$ if $x_{cnt,t} > 1$ then $x_{1,t} = x_{2,t}$... if $x_{cnt,t} = k$ then $x_{(k-1),t} = x_{k,t}$ $x_{cnt,t} = x_{cnt,t} - 1$ \triangleright end of generation </pre>

Fig. 5. PSO Translation Rules: each sequence of statements is atomic

the weak memory model. Technically, the idea is to lift the cubes under SC to buffered counterparts by a translation similar to the way in which we extrapolated predicates under SC (described above).

We can then pack the extrapolated cubes as new predicates provided to the predicate abstraction procedure for the weak memory model, and limit cube size to 1. Limiting the maximal cube size to 1 turns the process of cube search from exponential to linear (in the number of predicates and cubes obtained from extrapolation). This is key to making predicate abstraction tractable for relaxed memory models.

In Fig. 2(b), to reduce clutter, both Predicate Extrapolation and Cube Extrapolation lead to the same Predicate Abstraction box. However, it is important to note that the process of predicate abstraction for Cube Extrapolation *does not* perform exponential search while the process for predicate extrapolation does. As we will see in the experimental results, Predicate Extrapolation is sufficient for simpler programs, while Cube Extrapolation is necessary for more complex programs.

Step 4: Model checking. Once the boolean program is built (either via Predicate or Cube Extrapolation), the final step is to model check the program. This step is exactly the same as in the standard case of traditional predicate abstraction shown in Fig. 2(a).

3 Reduction

In this section, we describe a translation that transforms a program running on a relaxed memory model into a program with no weak memory model effects. The basic idea is to translate the store buffers in the semantics into variables that are part of the program. This translation enables us to leverage classic program analysis techniques such as predicate abstraction. Further, because the translation is parametric on the size of the buffer, it allows us to tailor the abstraction to the buffer size required by the particular concurrent algorithm. We show the process for the PSO memory model, the process for x86 TSO is similar and is shown in [10].

Reduction: PSO to SC. The reduction takes as input a thread identifier (whose statements are to be handled), as well as a bound on the maximum buffer size k for the per-variable buffer. Here, k denotes the maximum number of writes to global variables without a fence in-between the writes. While the translation presented here uses a fixed k for all global variables, we can easily use different k 's for different variables.

The translation considers in turn every statement that involves global variables. We introduce a translation function, which takes as input a statement, a thread identifier, and a bound on the maximum buffer size and produces a new statement as output:

$$\llbracket \cdot \rrbracket \in Stmt \times Thread \times \mathbb{N} \rightarrow Stmt$$

As a shorthand we write $\llbracket S \rrbracket_k^t$ for $\llbracket S, t, k \rrbracket$. $\llbracket S \rrbracket_k^t$ denotes the statement obtained from translating statement S in thread t with buffer bound k .

To perform the translation, the entries of the per-variable buffer are translated into thread-local variables and a local counter is introduced to maintain its depth. That is, for each global variable X and thread t , we introduce the following local variables:

- buffer content variables: $x_{1..t}, \dots, x_{k..t}$, where k is the maximum size of the buffer.
- a buffer counter variable: $x_{cnt..t}$.

Fig. 5 presents the translation of the three program code statements and the memory subsystem statement (flush). In the translation, the newly generated sequence of statements is atomic.

Store to a global variable $\llbracket X = r \rrbracket_k^t$: The store to a global variable X first checks if we are about to exceed the buffer bound k and if so, the program aborts. Otherwise, the counter is increased. The rest of the logic checks the value of the counter and updates the corresponding local variable. The global variable X is not updated and only local variables are involved.

Load from a global variable $\llbracket r = X \rrbracket_k^t$: The load from a global variable X checks the current depth of the buffer and then loads from the corresponding local variable. When the buffer is empty (i.e., $x_{cnt..t} = 0$), the load is performed directly from the global store. We do not need to check whether the buffer limit k is exceeded as that is ensured by the global store.

Fence statement $\llbracket fence \rrbracket_k^t$: For *each* shared variable X , the fence statement waits for the buffer of X to be empty (flush instructions to be executed). The fence has no effect on X .

Flush action $\llbracket flush \rrbracket_k^t$: The flush action is translated into a loop with a non-deterministic exit condition (we use $*$). New statements are introduced for each global variable X . If the buffer counter for the variable is positive, then it non-deterministically decides whether to update the global variable X or to continue the iteration. If it has decided to update X , the earliest write (i.e. $x_{1..t}$) is stored in X . The contents of the local variables are then updated by shifting: the content of each $x_{i..t}$ is taken from the content of the successor $x_{(i+1)..t}$ where $1 \leq i < k$. Finally, the buffer count is decremented. The composite statement inside the while loop is generated for each global variable. To ensure a faithful translation of the flush action, the whole newly generated statement is placed after *each* statement of the resulting program. The atomic statements are translated directly, without change (not shown in the figure).

The translation extends naturally to a sequence of statements and to programs with n concurrent threads: $\llbracket P \rrbracket_k = \llbracket S \rrbracket_k^1 \parallel \dots \parallel \llbracket S \rrbracket_k^n$, leading to the following theorem:

Theorem 1 (Soundness of Translation). *For a given program, P and a safety specification S , if $P \not\models_{pso} S$ then there exists a $k \in \mathbb{N}$ such that $\llbracket P \rrbracket_k \not\models_{sc} S$.*

From the theorem it follows that if $\llbracket P \rrbracket_k \models_{sc} S$ then $P \models_{pso} S$. When we successfully verify the program with a given k , it is guaranteed that no execution of the program ever requires a buffer of size larger than k . If the program does have an execution which exceeds k , then during verification we will encounter overflow and can attempt a higher value of k . That is, if we verify the program for a certain bound k , then the algorithm is correct for any size of the buffer greater or equal to k . In our experience, most concurrent algorithms exhibit low values for k as typically they use fences after a small number of global stores.

4 Predicate Abstraction for Relaxed Memory Models

In this section we describe how predicate abstraction is used to verify concurrent programs running on relaxed memory models. The central point we address is how to discover the predicates necessary for verification under the relaxed model from the predicates and the proof that was successful for verification of the program under sequential consistency (SC).

4.1 Predicate Abstraction

Predicate abstraction [16] is a special form of abstract interpretation that employs cartesian abstraction over a given set of predicates. Given a program P , and vocabulary (set of predicates) $V = \{p_1, \dots, p_n\}$ with corresponding boolean variables $\hat{V} = \{b_1, \dots, b_n\}$, predicate abstraction constructs a boolean program $\mathcal{BP}(P, V)$ that conservatively represents behaviors of P using only boolean variables from \hat{V} (corresponding to predicates in V). We use $[b_i]$ to denote the predicate p_i corresponding to the boolean variable b_i . We similarly extend $[b]$ to any boolean function b .

Next we explain how to construct $\mathcal{BP}(P, V)$. A literal is a boolean variable or its negation. A *cube* is a conjunction of literals, the size of a cube is the number of literals it contains. The concrete (symbolic) domain is defined as formulae over the predicates p_1, \dots, p_n . The abstract domain is a disjunctions of cubes over the variables b_1, \dots, b_n . The abstraction function α maps a formula φ over predicates from V to the weakest disjunction of cubes d such that $[d] \Rightarrow \varphi$.

The abstract transformer of a statement st w.r.t. a given vocabulary V can be computed using weakest-precondition computation and performing implication checks using a theorem prover:

$$b_i = \text{choose}(\alpha(wp(st, p_i)), \alpha(wp(st, \neg p_i)))$$

where

$$\text{choose}(\varphi_t, \varphi_f) = \begin{cases} 1, & \varphi_t \text{ evaluates to true;} \\ 0, & \text{only } \varphi_f \text{ evaluates to true;} \\ \star, & \text{otherwise.} \end{cases}$$

Different predicate abstraction techniques use different heuristics for reducing the number of calls to the prover.

Input: Vocabulary V , Statement st , Maximum cube size k Output: Abstract transformer for st over predicates from V	Input: RMM predicates V_{rmm} , RMM cubes C_{rmm} , Statement st , Output: Abstract transformer for st over predicates from $V_{rmm} \cup C_{rmm}$
<pre> function COMPUTETRANSFORMER(V, st, k) for each $p \in V$ do $\psi_p^+ = \psi_p^- = false$ $\varphi^+ = wp(st, p)$ $\varphi^- = wp(st, \neg p)$ for each $i = 1 \dots k$ do $cubes^+ = \text{BUILDBOUNDEDCUBES}(V, i, \psi_p^+)$ if $cubes^+ = \emptyset$ then break $\psi_p^+ = \text{COMPUTEAPPROX}(cubes^+, \varphi^+, \psi_p^+)$ for each $i = 1 \dots k$ do $cubes^- = \text{BUILDBOUNDEDCUBES}(V, i, \psi_p^-)$ if $cubes^- = \emptyset$ then break $\psi_p^- = \text{COMPUTEAPPROX}(cubes^-, \varphi^-, \psi_p^-)$ $\psi(p) = \text{choose}(\psi_p^+, \psi_p^-)$ </pre> <p style="text-align: center;">(a)</p>	<pre> function COMPUTETRANSFORMER(V_{rmm}, C_{rmm}, st) for each $p \in V_{rmm}$ do $\psi_p^+ = \psi_p^- = false$ $\varphi^+ = wp(st, p)$ $\psi_p^+ = \text{COMPUTEAPPROX}(V_{rmm} \cup C_{rmm}, \varphi^+, \psi_p^+)$ $\varphi^- = wp(st, \neg p)$ $\psi_p^- = \text{COMPUTEAPPROX}(V_{rmm} \cup C_{rmm}, \varphi^-, \psi_p^-)$ $\psi(p) = \text{choose}(\psi_p^+, \psi_p^-)$ </pre> <p style="text-align: center;">(b)</p>

Fig. 6. Computing abstract transformers: (a) classical predicate abstraction; (b) using extrapolation. COMPUTEAPPROX is shown in Fig. 7.

```

function COMPUTEAPPROX( $cubes, \varphi, \psi$ )
  for each  $c \in cubes$  do
    if  $c \Rightarrow \varphi$  then
       $\psi = \psi \vee c$ 
  return  $\psi$ 
                
```

Fig. 7. Predicate abstraction - helper function

Fig. 6 (a) shows a standard predicate abstraction algorithm in the spirit of [6]. The algorithm takes as input a statement st and a set of predicates (vocabulary) V . It then computes an abstract transformer for st using combinations of predicates from V . The algorithm works by computing an update formula $\psi(p)$ for every predicate p . The update formula is constructed as a choice between two sub-formulae, ψ_p^+ which holds when p should be set to *true*, and ψ_p^- which holds when p should be set to *false*.

The function BUILDBOUNDEDCUBES builds cubes of size i over predicates from V , checking that cubes of size i are not subsumed by previously generated cubes in ψ_p^+ or ψ_p^- . This function is standard, and we do not list it here due to space restrictions.

There are other algorithms that can be used here, such as the Flanagan&Qadeer's [15], or the one of Das et al. [11]. However, our focus is not on these optimizations, but on leveraging information from the verification of the *SC* program to discover the new predicates for verification under the relaxed memory model.

4.2 Predicate Extrapolation: From Predicates under SC to Predicates under Relaxed Model

Given a program which successfully verified under SC, our first approach to verifying the program under x86 TSO or PSO is to:

- Reduce the program as described in Section 3. That is, given a statement st of the original program, obtain a new statement st_{ps0} or st_{tso} from the translation.
- Compute $V_{ps0} = \text{EXTRAPOLATEPSO}(V)$ for PSO as discussed in Section 5, or for TSO: $V_{tso} = \text{EXTRAPOLATETSO}(V)$ (the TSO extrapolation is discussed in [10] and is similar to PSO). This extrapolates from the set of input predicates under SC and derives new predicates under x86 TSO or PSO.
- Invoke $\text{COMPUTETRANSFORMER}(V_{ps0}, st_{ps0}, k)$ to build the abstract transformer under PSO. Similarly, invoke $\text{COMPUTETRANSFORMER}(V_{tso}, st_{tso}, k)$ for x86 TSO. Here, k is the appropriate buffer bound. The function $\text{COMPUTETRANSFORMER}$ is shown in Fig. 6 (a).

That is, with this approach, the entire predicate abstraction tool chain remains the same except we change the input to the function for computing abstract transformers to be the new predicates V_{ps0} and the new statement st_{ps0} .

4.3 Cube Extrapolation: From SC Proof to PSO Predicates

As we will see in our experimental results, predicate extrapolation is effective only in some cases. The problem is that on more complex programs, the cube search space increases significantly meaning the function $\text{COMPUTETRANSFORMER}$ as described in Fig. 6 (a) times out. Next, we discuss another approach for computing abstract transformers under the relaxed memory model.

Core Idea. The core idea is that cubes generated during SC predicate abstraction capture invariants that are important for correctness under SC, but the same relationships between variables can be extrapolated to relationships between variables in the relaxed setting. Based on this observation, we extrapolate from these cubes similarly to the way we extrapolated from the predicates under SC. We use the function $C_{ps0} = \text{EXTRAPOLATEPSO}(C)$ where C denotes the cubes under SC. The newly generated extrapolated cubes C_{ps0} are then used as predicates for the verification.

The key point is that the cube search over C_{ps0} is now limited to cubes of size 1! The steps are as follows:

- Compute $V_{ps0} = \text{EXTRAPOLATEPSO}(V)$
- Compute $C_{ps0} = \text{EXTRAPOLATEPSO}(C)$
- Invoke $\text{COMPUTETRANSFORMER}(V_{ps0}, C_{ps0}, st)$ as shown in Fig. 6 (b) and taking as input extrapolated predicates and extrapolated cubes together with the statement.

The process for the x86 TSO model is identical.

Search vs. Extrapolation. In contrast to the standard $\text{COMPUTETRANSFORMER}$ of Fig. 6 (a), the algorithm of Fig. 6 (b) *does not* perform exhaustive search of the cube space. In particular, it does not take as input the parameter k . That is, our new way of building transformers is based on extrapolating from a previous (SC) proof.

5 Extrapolating Predicates: SC to PSO

In this section, we elaborate on how the function $preds_{ps0} = \text{EXTRAPOLATEPSO}(preds_{sc})$ operates. The operation $\text{EXTRAPOLATEPSO}(preds_{sc})$ for TSO is similar and is discussed in [10]. The function EXTRAPOLATEPSO computes the ingredients, that is, the new predicates $preds_{ps0}$, using which the final abstraction is built. We discuss the core reasons for introducing the new predicates. This reasoning is independent of predicate abstraction and can be used with other verification approaches.

Any abstraction for store buffer based RMMs such as PSO must be precise enough to preserve the following properties: (i) Intra-thread coherence: If a thread stores several values to shared variable X , and then performs a load from X , it should not see any value it has itself stored except the most recent one. (ii) Inter-thread coherence: A thread T_i should not observe values written to shared variable X by thread T_j in an order different from the order in which they were written. (iii) Fence semantics: If a thread T_i executes a fence when its buffer for variable X is non-empty, the value of X visible to other threads immediately after the fence should be the most recent value T_i wrote.

Fig. 8 shows a simple example in which a single thread stores two values into a shared variable X and then loads the value of X into $l1$. To successfully verify this program, the abstraction we use must be precise enough to capture intra-thread coherence.

Thread 1:

```

1  X=0;
2  X=1;
3  l1=X;
4  fence;
5  assert (X = l1);

```

5.1 Generic Predicates

For a shared variable X , if we use a buffer with a max size of 1, our translation adds the predicates: $X_cnt_t1 = 0$, $X_cnt_t1 = 1$ indicating the last location of the buffer for X which has been written to by thread 1, but not yet flushed. These predicates serve multiple purposes: (i) track the store buffers size; (ii) provide knowledge during `store` and `load` operations on where to write/read the value of X . (iii) preserve *Intra-thread coherence* in the abstraction.

Fig. 8. Intra-thread coherence example

Another predicate we generate is: $overflow = 0$. This predicate supplements the previously described predicates, giving indication of when the number of subsequent stores to a shared variable X , without a fence or a flush in between these stores, exceeds the limit k of our abstraction. This is crucial to ensure soundness of the abstraction.

The general extrapolation rule, which is independent of the verified program and of the input SC predicates $preds_{sc}$, is:

Rule 1. For a buffer size bound k , add the following predicates to $preds_{ps0}$:

- $\{\mathbf{V_cnt_T} = i \mid 1 \leq i \leq k, \mathbf{V} \in \text{Gvar}, \mathbf{T} \in \text{Tthread}\}$
- $overflow = 0$

5.2 Extrapolating from $preds_{sc}$

We now describe how to extrapolate from the predicates in the set $preds_{sc}$ in order to compute new predicates that become part of $preds_{ps0}$. The rule below ensures that the SC executions of the new program can be verified.

Rule 2. Update the set $preds_{pso}$ to contain the set $preds_{sc}$

Next, we would like properties on the values of a shared variable X captured by predicates in $preds_{sc}$ to also be captured for the buffered values of X . For example if $preds_{sc}$ contains $X = 0$, we add the predicate $X_{.1.t1} = 0$ for a buffer of X for thread T_1 . This can be seen in the example of Fig. 8 where we need to track that the buffered value of X is 0 at line 1. We summarize these observations in the following rule:

Rule 3. Update $preds_{pso}$ to contain the set $\bigcup_{p_{sc} \in preds_{sc}} lift(p_{sc})$

Here, $lift(p_{sc})$ generates from each SC predicate a set of PSO predicates where the original variables are replaced with buffered versions of the variables (for each buffered version of a variable and their combination).

In addition to the above rules, adding a predicate $X_{.1.t1} = X$ ensures that the shared value of X and the buffered value of X are in agreement (when the predicate is set to *true*). This reduces the time and space of model checking. Following similar reasoning, the predicate $X_{.1.t1} = X_{.2.t1}$ is also added.

Rule 4. For $\mathbf{V} \in Gvar$, $\mathbf{T} \in Thread$ and k the buffer bound, update $preds_{pso}$ to contain the sets:

- $\{\mathbf{V}_{.(i-1).\mathbf{T}} = \mathbf{V}_{.i.\mathbf{T}} \mid 2 \leq i \leq k\}$
- $\{\mathbf{V}_{.i.\mathbf{T}} = \mathbf{V} \mid 1 \leq i \leq k\}$

The above rules add both: generic predicates that are independent of $preds_{sc}$ as well as predicates that are extrapolated from $preds_{sc}$. But these rules may sometimes generate a larger set of predicates than necessary for successful verification. We now describe several optimizations that substantially reduce that number.

Rule 5 Read-only shared variables. If a thread t never writes to a shared variable X do not extrapolate the SC predicates referencing X to their PSO counterparts for t .

Rule 6 Predicates referencing a shared variable more than once. Replace all occurrences of the shared variable with the same buffered location.

For example, for $X \leq Y \wedge 0 \leq X$, where X is referred to more than once, we generate the new predicate $X_{.1.t1} \leq Y \wedge 0 \leq X_{.1.t1}$, but we do not generate the predicate $X_{.1.t1} \leq Y \wedge 0 \leq X_{.1.t2}$. The intuition is that the SC predicate captures information regarding the value of X at some point in the execution and when extrapolating the predicate to PSO, we need to capture that information regarding the shared value of X or its buffered value, yet the mixture of the two is redundant. Similarly for $Y2 \leq Y1$, we do not necessarily need to generate the predicate $Y2_{.1.t2} \leq Y1_{.1.t1}$.

Rule 7 Predicates referencing different shared variables. For a predicate referencing more than one shared variable, if it can be guaranteed that a fence will be executed between every two shared location writes, restrict to generating predicates that relate to one buffered location at most.

In summary, EXTRAPOLATEPSO is computed by applying the above seven rules for the predicates in V and C (both obtained from the verification under SC).

6 Experimental Evaluation

We implemented predicate abstraction based on predicate extrapolation (PE) and cube extrapolation (CE) as outlined earlier in our tool called CUPEX. Then, we thoroughly evaluated the tool on a number of challenging concurrent algorithms. All experiments were conducted on an Intel(R) Xeon(R) 2.13GHz with 250GB RAM. The key question we seek to answer is whether predicate extrapolation and cube extrapolation are precise and scalable enough to verify all of our (relaxed) programs.

6.1 Prototype Implementation

CUPEX works in two phases. In the first phase, given an input program, it applies abstraction and produces a boolean program. In the second phase, the boolean program is verified using a three-valued model checker for boolean programs. To reduce the cube search space, CUPEX uses optimizations such as bounded cube size search and cone of influence. For every assignment statement in the original program, it updates only the boolean variables corresponding to predicates which contain the assigned variable from the statement (subject to aliasing). The search in the cube space is performed in increasing cube size order, thus we find the weaker (smaller) cubes first. CUPEX uses Yices 1.0.34 as the underlying SMT solver.

The second phase relies on a three-valued model checker to verify the boolean program. Our model checker uses 3-valued logic for compact representation, and handles assume statements in a number of clever ways, performs partial concretization of assume conditions and merges states after updates.

6.2 Results

We evaluated CUPEX on the following concurrent algorithms: Dekker’s mutual exclusion algorithm [12], Peterson’s mutual exclusion algorithm [26], Szymanski mutual exclusion algorithm [29], Alternating Bit protocol (already discussed in Section 2), an Array-based Lock-Free Queue (here, we verified its memory safety), Lamport’s Bakery algorithm [24] and the Ticket locking algorithm [3]. The first three algorithms are finite-state, while the last four are infinite-state. For each algorithm, we evaluated our tool for x86 TSO and PSO models. We ran tests with buffer bounds ranging from $k \in 1 \dots 3$. For each k , we tested various fence configurations. We present in the result tables values for $k = 1$, obtained for the minimal fence configurations which successfully verified.

Meaning of table columns. Our results are summarized in Table 1 and Table 2. The meaning of most table columns is self explanatory, but we elaborate on the following columns of Table 1:

- Build Boolean Program (i) # input preds: number of initial input predicates. For x86 TSO and PSO, these are obtained by extrapolating from the predicates in the # input preds column of the corresponding SC program. (ii) # SMT calls: total number of calls (in thousands) to the SMT solver required to build the boolean program (BP). (iii) time: time in seconds that it took to build the boolean program. We use T/O for timed out (keeps running after 10 hours). (iv) # cubes used: total

Table 1. Results for Predicate Extrapolation

algorithm	memory model	Build Boolean Program					Model check		
		# input preds	# SMT calls (K)	time (sec)	# cubes used	cube size	# states (K)	memory (MB)	time (sec)
Dekker	SC	7	0.7	0.1	0	1	14	6	1
	PSO	20	26	6	0		80	31	5
	TSO	18	22	5	0		45	20	3
Peterson	SC	7	0.6	0.1	2	2	7	3	1
	PSO	20	15	3	2		31	13	3
	TSO	18	13	3	2		25	11	2
ABP	SC	8	2	0.5	5	2	0.6	1	0.6
	PSO	15	20	4	5		2	3	1
	TSO	17	23	5	5		2	3	1
Szymanski	SC	20	16	3.3	1	2	12	6	2
	PSO	35	152	33	1		61	30	4
	TSO	37	165	35	1		61	31	5

number of cubes in the boolean program whose size is greater than 1, that is, these are cubes composed of 2 or more input predicates. (v) cube size: maximum cube size found in the boolean program. For instance, cube size 4 means that there exist cubes which combine 4 input predicates.

- Model check (i) # states: total number of states (thousands) explored by the underlying three-valued model checker.

Table 2 contains two additional columns in the Build Boolean Program column: (i) method used to build the boolean program: PE, CE or, for SC programs where no new predicates are generated, Trad (Traditional). (ii) # input cubes: These are obtained by extrapolating from the cubes in the # cubes used column of the corresponding SC program. These cubes are then added to the initial # input preds and CE is performed.

6.3 Observations

Sequential Consistency. We assume that the predicates for the original SC program are provided by the programmer or inferred using existing SC verification tools such as [17]. For our benchmarks we manually provided the SC predicates, since we focus on the relaxed memory verification.

Predicate Extrapolation. We first focus on the results of Table 1. Here, PE was quick enough to verify all programs under x86 TSO and PSO. For example for Dekker’s algorithm, even though there was a significant increase in the number of predicates (7 input predicates required to verify the program under SC yet under PSO, 20 predicates were needed), the newly generated predicates were precise enough to prove the program. A similar pattern can be observed for Peterson, ABP and Szymanski. For all four algorithms, a small cube size (at most 2) was sufficient for verification. Furthermore, the number of cubes of size 2 that are used in the boolean program is fairly small. Overall, these four programs require a small cube size to verify, at most 2, leading to quick building of the boolean program. For these programs, CE was unnecessary.

Table 2. Results for Predicate Extrapolation and Cube Extrapolation

algorithm	memory model	method	Build Boolean Program					Model check				
			# input preds	# input cubes	# SMT calls (K)	time (sec)	# cubes used	cube size	# states (K)	memory (MB)	time (sec)	
Queue	SC	Trad	7	-	20	5	50	4	1	2	1	
		PE	15	-	5,747	1,475	412		1	4	1	
	CE	99	-	98	17	99	11		6	2		
		16	-	11,133	2,778	412	12		4	1		
	TSO	CE	99	-	163	31	99		12	7	2	
		Trad	15	-	1,552	355	161		20	8	2	
Bakery	SC	PE	38	-	-	T/O	-	4	-	-	-	
		CE	422	-	9,018	1,773	381		979	375	104	
	TSO	PE	36	-	-	T/O	-		-	-	-	
		CE	422	-	7,048	1,386	383		730	285	121	
	SC	Trad	11	-	218	51	134		4	2	2	1
		PE	56	-	-	T/O	-			-	-	-
CE	622	-	15,644	2,163	380	193	123	40				
	48	-	-	T/O	-	-	-	-		-		
TSO	CE	622	-	6,941	1,518	582	71	67		545		

Cube Extrapolation. Next, we discuss the results of Table 2. We first observe that for Queue, under PSO, PE required around 6 million calls to the SMT solver and took a little over 24 minutes to complete in building the boolean program. Indeed, the combination of increased cube size (4) together with 15 initial input predicates significantly affected the running time for building the boolean program. Interestingly, we observe that CE, was able to reduce the number of calls to the SMT solver by a factor of 60 and reduce the running time by a factor of 80. Importantly, CE was precise enough to verify the program. Here we see that CE generated 99 new cubes which were extrapolated from the 50 SC cubes. The final program used exactly these 99 cubes, meaning that CE did not generate redundant cubes.

For both Bakery and Ticket, the benefits of CE are even more startling. With PE, building the boolean program fails under both x86 TSO and PSO due to a time out. However, CE massively reduced the number of SMT calls enabling successful generation of the boolean program. The set of cubes CE returned was fairly close to the set of cubes used during the boolean program verification. For instance, in Bakery under PSO, 422 input cubes were generated out of which 381 were used in the boolean program (fairly close to ideal).

It is worth noting that in all benchmarks we experimented on, the minimal fence placement required was different for x86 TSO and PSO.

Discussion. Next we note two observations from our approach which we encountered experimentally and which we believe are interesting items for future work.

First, when we directly applied CE to the Ticket algorithm, it took hours to verify for PSO. To solve this problem, we hypothesized that given a safety property, which does not reference buffered values, we may allow inconsistent values at buffered locations, and that inconsistency would be resolved when those values are flushed and before an

error state is reached. Therefore, to enable quicker verification, we first applied CE as usual, and then automatically removed all predicates referencing buffered values from the resulting cubes found in the boolean program after CE. Such a reduction preserves soundness while abstracting the proof. We note that although this approach worked for Ticket under PSO, when we tried it under x86 TSO this additional pass introduced too much imprecision and the program failed to verify (the table reports results for Ticket on PSO using this additional pass and on x86 TSO without this pass).

Second, for the Queue algorithm our initial successful SC proof was insufficient to extrapolate from. Portions of the program where the boolean program under SC lost precision due to abstraction were amplified by the extrapolation. For instance, where the SC proof used a predicate $Tail < Head$ which was unknown through parts of the SC proof with no adverse effects, the extrapolated proof propagated this uncertainty causing an error state to be reached. Adding $Tail \leq Head$ strengthened the SC proof and enabled successful extrapolation (this is the result we report in the Table).

Summary. For cubes of small size, 2 or less, with PE, CUPEX builds the boolean program quickly and is precise enough to verify the program. For larger cube sizes, PE takes too long to build the boolean program or times out. However, CUPEX with CE enables us to build the boolean program in reasonable time and critically, is precise enough to verify the program both for x86 TSO and PSO.

7 Related Work

There has been almost no work on automatically verifying *infinite-state* concurrent programs running on relaxed memory models. We briefly survey some of the more closely related work.

Model Checking for Relaxed Memory Models. The works of [25,21,22,19] describe explicit-state model checking under several memory models. In [7], instead of working with operational memory models and explicit model-checking, they convert programs into a form that can be checked against an axiomatic model specification. These approaches do not handle infinite-state programs. The work in [23] focuses on programs that are finite-state under SC but infinite-state under x86 TSO and PSO and suggests an abstraction to deal with the issue. Unfortunately, it also cannot handle general infinite-state programs (i.e., the program must be finite-state under SC).

The works of [2,5] present a code-to-code transformation which encodes the relaxed memory semantics into the program. Our approach goes beyond this transformation and tackles the difficulty of verifying the newly obtained relaxed program. These new programs are more difficult to verify because of the complexity added by the encoded semantics. Our approach solves this problem by learning from the proof under sequential consistency.

The work of [1] builds on [23] and handles infinite state programs (on x86 TSO) by applying both predicate abstraction and store buffers abstraction. Their approach discovers predicates via traditional abstraction refinement and does not reuse information from the proof under SC, while in our approach we present a technique which leverages an existing proof under SC in order to derive a new proof for a more relaxed program.

Further, we also handle a memory model (PSO) that allows for more behaviors and complexity than x86 TSO.

Lazy abstraction. The work of [18] introduces the concept of adjusting the level of abstraction for different sections of the verified program’s state space. This is achieved by applying on-the-fly refinement for search-tree sub-graphs. Their approach does not construct a boolean program during verification. However, encoding the weak memory semantics in the code and extrapolating from the SC proof are concepts applicable for extending lazy abstraction to relaxed memory models. The backwards counter-example analysis phase, which requires costly calls to the theorem prover, may in part be avoided by anticipating in each branch of the search tree which predicates are required.

8 Conclusion and Future Work

We introduced a novel approach for predicate abstraction of concurrent programs running on relaxed memory models such as x86 TSO and PSO. The essence of our approach is extrapolation: learning from an existing proof of the program under sequential consistency in order to obtain a proof for a more relaxed version of the program.

We implemented our extrapolation approach and successfully applied it to automatically verify several challenging concurrent algorithms for both x86 TSO and PSO. This is the first time some of these programs have been verified for a model as relaxed as PSO.

As future work, we plan to investigate how these techniques apply to other relaxed models, both hardware models such as Power, as well as software programming models such as [8].

References

1. Abdulla, P.A., Atig, M.F., Chen, Y.-F., Leonardsson, C., Rezine, A.: Automatic fence insertion in integer programs via predicate abstraction. In: Miné, A., Schmidt, D. (eds.) SAS 2012. LNCS, vol. 7460, pp. 164–180. Springer, Heidelberg (2012)
2. Alglave, J., Kroening, D., Nimal, V., Tautschnig, M.: Software verification for weak memory via program transformation. In: Felleisen, M., Gardner, P. (eds.) ESOP 2013. LNCS, vol. 7792, pp. 512–532. Springer, Heidelberg (2013)
3. Andrews, G.R.: Concurrent programming - principles and practice. Benjamin/Cummings (1991)
4. Atig, M.F., Bouajjani, A., Burckhardt, S., Musuvathi, M.: On the verification problem for weak memory models. In: POPL (2010)
5. Atig, M.F., Bouajjani, A., Parlato, G.: Getting rid of store-buffers in tso analysis. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 99–115. Springer, Heidelberg (2011)
6. Ball, T., Majumdar, R., Millstein, T., Rajamani, S.K.: Automatic predicate abstraction of C programs. In: PLDI (2001)
7. Burckhardt, S., Alur, R., Martin, M.M.K.: CheckFence: checking consistency of concurrent data types on relaxed memory models. In: PLDI (2007)
8. Burckhardt, S., Baldassin, A., Leijen, D.: Concurrent programming with revisions and isolation types. In: OOPSLA (2010)

9. Cousot, P., Cousot, R.: Abstract interpretation: A unified lattice model for static analysis of programs by construction of approximation of fixed points. In: POPL (1977)
10. Dan, A., Meshman, Y., Vechev, M., Yahav, E.: Predicate abstraction for relaxed memory models. Tech. rep
11. Das, S., Dill, D.L., Park, S.: Experience with Predicate Abstraction. In: Halbwachs, N., Peled, D.A. (eds.) CAV 1999. LNCS, vol. 1633, pp. 160–171. Springer, Heidelberg (1999)
12. Dijkstra, E.: Cooperating sequential processes, TR EWD-123. Tech. rep., Technological University, Eindhoven (1965)
13. Donaldson, A., Kaiser, A., Kroening, D., Wahl, T.: Symmetry-aware predicate abstraction for shared-variable concurrent programs. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 356–371. Springer, Heidelberg (2011)
14. Dubois, M., Scheurich, C., Briggs, F.A.: Memory access buffering in multiprocessors. In: ISCA (1986)
15. Flanagan, C., Qadeer, S.: Predicate abstraction for software verification. In: POPL (2002)
16. Graf, S., Saïdi, H.: Construction of abstract state graphs with PVS. In: Grumberg, O. (ed.) CAV 1997, vol. 1254, pp. 72–83. Springer, Heidelberg (1997)
17. Gupta, A., Popeea, C., Rybalchenko, A.: Threader: A constraint-based verifier for multi-threaded programs. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 412–417. Springer, Heidelberg (2011)
18. Henzinger, T.A., Jhala, R., Majumdar, R., Sutre, G.: Lazy abstraction. In: POPL (2002)
19. Huynh, T.Q., Roychoudhury, A.: Memory model sensitive bytecode verification. *Form. Methods Syst. Des.* (2007)
20. IBM. Power ISA v.2.05. (2007)
21. Jonsson, B.: State-space exploration for concurrent algorithms under weak memory orderings. *SIGARCH Comput. Archit. News* (2008)
22. Kuperstein, M., Vechev, M., Yahav, E.: Automatic inference of memory fences. In: FMCAD (2010)
23. Kuperstein, M., Vechev, M., Yahav, E.: Partial-coherence abstractions for relaxed memory models. In: PLDI (2011)
24. Lamport, L.: A new solution of Dijkstra’s concurrent programming problem. *Commun. ACM* (1974)
25. Park, S., Dill, D.L.: An executable specification and verifier for relaxed memory order. *IEEE Trans. on Computers* 48 (1999)
26. Peterson, G.L.: Myths about the mutual exclusion problem. *Inf. Process. Lett.* 12(3) (1981)
27. Sarkar, S., Sewell, P., Nardelli, F.Z., Owens, S., Ridge, T., Braibant, T., Myreen, M.O., Alglave, J.: The semantics of x86-cc multiprocessor machine code. In: POPL (2009)
28. SPARC International, Inc. The SPARC architecture manual (version 9). Prentice-Hall, Inc., Upper Saddle River, NJ, USA (1994)
29. Szymanski, B.K.: A simple solution to Lamport’s concurrent programming problem with linear wait. In: ICS (1988)