

Automatically Validating Temporal Safety Properties of Interfaces

Thomas Ball and Sriram K. Rajamani

Software Productivity Tools
Microsoft Research
<http://www.research.microsoft.com/slam/>

Abstract. We present a process for validating temporal safety properties of software that uses a well-defined interface. The process requires only that the user state the property of interest. It then automatically creates abstractions of C code using iterative refinement, based on the given property. The process is realized in the SLAM toolkit, which consists of a model checker, predicate abstraction tool and predicate discovery tool. We have applied the SLAM toolkit to a number of Windows NT device drivers to validate critical safety properties such as correct locking behavior. We have found that the process converges on a set of predicates powerful enough to validate properties in just a few iterations.

1 Introduction

Large-scale software has many components built by many programmers. Integration testing of these components is impossible or ineffective at best. Property checking of interface usage provides a way to partially validate such software. In this approach, an interface provides a set of properties that all clients of the interface should respect. An automatic analysis of the client code then validates that it meets the properties, or provides examples of execution paths that violate the properties. The benefit of such an analysis is that errors can be caught very early in the coding process.

We are interested in checking that a program respects a set of *temporal safety* properties of the interfaces it uses. Safety properties are the class of properties that state that “something bad does not happen”. An example is requiring that a lock is never released without first being acquired (see [21] for a formal definition). Given a program and a safety property, we wish to either validate that the code respects the property, or find an execution path that shows how the code violates the property.

In this paper, we show that safety properties of system software can be validated and invalidated using model checking, without the need for user-supplied annotations (invariants) or user-supplied abstractions. As no annotations are required, we use model checking to compute fixpoints automatically over an abstraction of the C code. We construct an appropriate abstraction by (1) obtaining an initial abstraction from the property that needs to be checked, and (2) refining this abstraction using an automatic refinement algorithm.

We model abstractions of C programs using *boolean programs* [3]. Boolean programs are C programs in which all variables have boolean type. Boolean programs contain all the control-flow constructs of C program, procedures, and procedure calls with call-by-value parameter passing. Each boolean variable in a boolean program has an interpretation as a predicate over the infinite state space of the C program. Our experience shows that our refinement algorithm finds boolean program abstractions that are precise enough to validate properties. Furthermore, if the property is violated, the process of searching for a suitable boolean program abstraction leads to a manifestation of the violation.

We present the SLAM toolkit for checking safety properties of system software, and report on our experience in using the toolkit to check properties of Windows NT device drivers. Given a safety property to check on a C program, the SLAM process has the following phases: (1) abstraction, (2) model checking, and (3) predicate discovery. We have developed tools to support each of these phases:

- C2BP, a tool that transforms a C program P into a boolean program $\mathcal{BP}(P, E)$ with respect to a set of predicates E over the state space of P [1, 2];
- BEBOP, a tool for model checking boolean programs [3], and
- NEWTON, a tool that discovers additional predicates to refine the boolean program, by analyzing the feasibility of paths in the C program.

The SLAM toolkit provides a fully automatic way of checking temporal safety properties of system software. Violations are reported by the SLAM toolkit as paths over the program P . It never reports spurious error paths. Instead, it detects spurious error paths and uses them to automatically refine the abstraction (to eliminate these paths from consideration). Since property checking is undecidable, the SLAM refinement algorithm may not converge. However, in our experience, it usually converges in a few iterations. Furthermore, whenever it converges, it gives a definite “yes” or “no” answer.

The worst-case run-time complexity of the SLAM tools BEBOP and C2BP is linear in the size of the program’s control flow graph, and exponential in the number of predicates used in the abstraction. We have implemented several optimizations to make BEBOP and C2BP scale gracefully in practice, even with a large number of predicates. The NEWTON tool scales linearly with path length and number of predicates.

We applied the SLAM toolkit to check the use of the Windows NT I/O manager interface by device driver clients. There are on the order of a hundred rules that the clients of the I/O manager interface should satisfy. We have automatically checked properties on device drivers taken from the Microsoft Driver Development Kit¹. While checking for correct use of locks, we found that the SLAM process converges in one or two iterations to a boolean program that is sufficiently precise to validate/invalidate the property. We also checked a data-dependent property, which requires keeping track of value-flow and aliasing, using four iterations of the SLAM tools.

The remainder of this paper is organized as follows. Section 2 gives an overview of the SLAM approach by applying the tools to verify part of an NT device driver. Sections 3, 4 and 5 give brief descriptions of the three tools that compose the SLAM toolkit and explain how they work in the context of the running example. Section 6 describes our experience applying the tools to an NT device driver. Section 7 discusses related work and Section 8 concludes the paper.

2 Overview

This section introduces the SLAM refinement algorithm and then applies this algorithm to a small code example, extracted from a PCI device driver. The SLAM toolkit handles a significant subset of the C language, including pointers, structures, and procedures (with recursion and mutual recursion). A limitation of our tools is that they assume a *logical*

¹ The code of the device drivers we analyzed is freely available from <http://www.microsoft.com/ddk/W2kDDK.htm>

model of memory when analyzing C programs. Under this model, the expression $p + i$, where p is a pointer and i is an integer, yields a pointer value that points to the same object pointed to by p . That is, we treat pointers as references rather than as memory addresses. Note that this is the same basic assumption underlying most points-to analysis, including the one that our tools use [10].

2.1 Refinement Algorithm

We wish to check if a temporal safety property φ is satisfied by a program P . We assume that the program P has been instrumented to result in a program P' such that P satisfies φ iff the label **ERROR** is not reachable in P' . In particular, the instrumentation takes the form of calls to a finite state machine (FSM) transition function, written in C. The parameters to the function encode the events/data that determine the FSM’s next state. The transition function simply switches on the current state of the machine (kept in global variables) and its formal parameters, to decide which state comes next. The label **ERROR** in this function reflects the finite state machine moving into a reject state. This is known in the model checking community as a “product automaton construction” and is a fairly standard way to encode safety properties.²

Let i be a metavariable that records the SLAM iteration count. In the first iteration ($i = 0$), we start with a set of predicates E_0 that capture the state of the FSM. Let $state$ be the global variable representing the state of the FSM and let $D(state)$ be its domain. Without loss of generality, let x be the single formal parameter of the transition function and $D(x)$ be its domain. Then the set E_0 is given as: $E_0 = \{(state = s) \mid s \in D(state)\} \cup \{(x = f) \mid f \in D(x)\}$. Let E_i be some set of predicates over the state of P' , the instrumented version of P . Then iteration $i + 1$ of SLAM is carried out using the following steps:

1. Apply C2BP to construct the boolean program $\mathcal{BP}(P', E_i)$, which has the same control-flow graph as P' .
2. Apply BEBOP to check if there is a path p_i in $\mathcal{BP}(P', E_i)$ that reaches the **ERROR** label. If BEBOP determines that **ERROR** is not reachable, then the property φ is valid in P , and the algorithm terminates.
3. If there is such a path p_i , then we use NEWTON to check if p_i is feasible in P . There are two outcomes:
 - “yes”: the property φ has been invalidated in P , and the algorithm terminates with an error path p_i (a witness to the violation of φ).
 - “no”: NEWTON finds a set of predicates F_{i+1} that explain the infeasibility of path p_i in P .
4. Let $E_{i+1} = E_i \cup F_{i+1}$, and $i = i + 1$, and proceed to the next iteration.

As stated before, this algorithm is potentially non-terminating. However, when it does terminate, it provides a definitive answer.

2.2 Example

Figure 1(a) presents a snippet of C code from a PCI device driver that processes interrupt request packets (irps). Of interest here are the calls the code makes to acquire and

² We have created a low-level specification language called SLIC (Specification Language for Interface Checking) that can be used to generate such an instrumented C program, which will be the topic of a future paper.

<pre> void example() { do { //get the write lock KeAcquireSpinLock(&devExt->writeListLock); nPacketsOld = nPackets; request = devExt->WriteListHeadVa; if(request && request->status){ devExt->WriteListHeadVa = request->Next; // release the lock KeReleaseSpinLock(&devExt->writeListLock); irp = request->irp; if(request->status > 0){ irp->IoS.Status = STATUS_SUCCESS; irp->IoS.Information = request->Status; } else { irp->IoS.Status = STATUS_UNSUCCESSFUL; irp->IoS.Information = request->Status; } SmartDevFreeBlock(request); IoCompleteRequest(irp, IO_NO_INCREMENT); nPackets++; } } while (nPackets != nPacketsOld); // release the lock KeReleaseSpinLock(&devExt->writeListLock); } </pre>	<pre> typedef {Locked, Unlocked} STATETYPE; typedef {Acq, Rel} MTYPE; stateType state = Unlocked; FSM(m : MTYPE){ if ((state==Unlocked) && (m==Acq)) A: state = Locked; else if ((state==Locked) && (m==Rel)) B: state = Unlocked; else ERROR: ; } void example() { do { C: KeAcquireSpinLock(&devExt->writeListLock); FSM(Acq); nPacketsOld = nPackets; request = devExt->WriteListHeadVa; if(request && request->status){ D: devExt->WriteListHeadVa = request->Next; KeReleaseSpinLock(&devExt->writeListLock); FSM(Rel); irp = request->irp; if(request->status > 0){ irp->IoS.Status = STATUS_SUCCESS; irp->IoS.Information = request->Status; } else { irp->IoS.Status = STATUS_UNSUCCESSFUL; irp->IoS.Information = request->Status; } E: SmartDevFreeBlock(request); IoCompleteRequest(irp, IO_NO_INCREMENT); nPackets++; } } while (nPackets != nPacketsOld); F: KeReleaseSpinLock(&devExt->writeListLock); FSM(Rel); } </pre>
(a) Program P	(b) Instrumented Program P'

Fig. 1. (a) A snippet of device driver code P and the (b) instrumented code P' that checks proper use of spin locks.

release spin locks (`KeAcquireSpinLock`, `KeReleaseSpinLock`). Figure 1(b) shows the program, instrumented to check that locks are properly acquired and released using a finite state machine with two states `Locked` and `Unlocked`. The procedure `FSM` implements the transition function of the state machine, as described before.

<pre> decl {state==Locked}, {state==Unlocked}; void FSM({m==Acq},{m==Rel}) begin if ({state==Unlocked}&&{m==Acq}) A: {state==Locked}, {state==Unlocked} := 1,0; else if ({state==Locked}&&{m==Rel}) B: {state==Locked}, {state==Unlocked} := 0,1; else ERROR: skip; fi end void example() // void example() { begin // do // do { skip; // KeAcquireSpinLock(...); C: FSM(1,0); // FSM(Acq); skip; // nPacketsOld = nPackets; skip; // request = devExt->...; // if (*) then // if(request && request->status){ D: skip; // devExt->WriteListHeadVa = ...; skip; // KeReleaseSpinLock(...); FSM(0,1); // FSM(Rel); skip; // irp = request->irp; if (*) then // if(request->status > 0){ skip; // irp->IoS.Status = ... skip; // irp->IoS.Information = ...; else // } else { skip; // irp->IoS.Status = ...; skip; // irp->IoS.Information = ...; fi // } E: skip; // SmartDevFreeBlock(request); skip; // IoCompleteRequest(... skip; // nPackets++; end // } while (*) // } while (nPackets != nPacketsOld); skip; // KeReleaseSpinLock(...); F: FSM(0,1); // FSM(Rel); end // } </pre>	<pre> // see code in left pane void example() begin do C: FSM(1,0); b := 1; skip; skip; if (*) then D: skip; skip; FSM(0,1); skip; if (*) then skip; skip; else skip; skip; fi E: skip; skip; b := choose(0,b); fi while (!b) skip; F: FSM(0,1); end </pre>
(a) Boolean program $\mathcal{BP}(P', E_0)$	(b) Boolean program $\mathcal{BP}(P', E_1)$

Fig. 2. The two boolean programs created while checking the code from Figure 1(b). (The second boolean program also contains the state machine function and global state variable, but we omit it to enhance the clarity of the figure). See the program text for the definition of the choose function.

The question we wish to answer is: is the label **ERROR** reachable in the code in Figure 1(b)? The following sections apply the algorithm given above to show that **ERROR** is unreachable.

2.3 Initial Boolean Program

The first step of the algorithm is to generate a boolean program from the C program and the set of predicates E_0 that define the states of the finite state machine. We represent our

abstractions as *boolean programs*. The syntax and semantics of boolean program was defined in [3]. Boolean programs are C programs in which the only allowed types are `bool` (with values 0 and 1) and `void`. Boolean programs also allow control non-determinism, through the conditional expression “*”, as shown later on.

For our example, the set E_0 consists of four predicates: two *global* predicates, $(state = Locked)$ and $(state = Unlocked)$, and two *local* predicates over the formal parameter m to the procedure `FSM`, $(m = Acq)$ and $(m = Rel)$. These four predicates and the C program of Figure 1(b) are input to the C2BP (C to Boolean Program) tool to create the boolean program $\mathcal{BP}(P', E_0)$, shown in Figure 2(a). This program has two global variables, $\{state==Locked\}$ and $\{state==Unlocked\}$, and the procedure `FSM` has two formal parameters, $\{m==Acq\}$ and $\{m==Rel\}$.³ For every statement s in the C program and predicate $e \in E_0$, the C2BP tool determines the effect of statement s on predicate e . For example, consider the assignment statement “`state = Locked;`” at label `A` in program P' of Figure 1(b). This statement makes the predicate $(state = Locked)$ true and the predicate $(state = Unlocked)$ false. This is reflected in the boolean program $\mathcal{BP}(P', E_0)$ by the parallel assignment statement

```
{state==Locked}, {state==Unlocked} := 1,0;
```

The translation of the boolean expressions in the conditional statements of the C program results in the obvious corresponding boolean expressions in the `FSM` procedure in the boolean program. Control non-determinism is used to conservatively model the conditions in the C program that cannot be abstracted precisely using the predicates in E_0 .

Many of the assignment statements in the C program are abstracted to the **skip** statement (no-op) in the boolean program. The C2BP tool uses Das’s points-to analysis [10] to determine whether or not an assignment statement through a pointer dereference can affect a predicate e . In our example, the points-to analysis shows that no variable in the C program can alias the address of the global `state` variable (or the formal parameter m of procedure `FSM`).⁴

We say that the boolean program $\mathcal{BP}(P', E_0)$ *abstracts* the program P' , since every feasible execution path p of the program P' also is a feasible execution path of $\mathcal{BP}(P', E_0)$.

2.4 Model Checking The Boolean Program

The second step of our process is to determine whether or not the label `ERROR` is reachable in the boolean program $\mathcal{BP}(P', E_0)$. We use the BEBOP model checker to determine the answer to this query. In this case, the answer is “yes”. Like most model checkers, the BEBOP tool produces a (shortest) path leading to the error state. In this case, the shortest path to the error state is the path that goes around the loop twice, acquiring the lock twice without an intermediate release, as given by the following error path p of labels:

[C, A, E, C, ERROR]

³ Boolean programs permit a variable identifier to be an arbitrary string enclosed between “{” and “}”. This is helpful for giving boolean variables names to directly represent the predicates in the C program that they represent.

⁴ The analysis also shows that the procedures `SmartDevFreeBlock`, and kernel procedures `IoCompleteRequest`, `KeAcquireSpinLock`, and `KeReleaseSpinLock` cannot affect these variables so the calls to them are removed.

2.5 Predicate Discovery over Error Path

Because the C program and the boolean program abstractions have identical control-flow graphs, the error path p in $\mathcal{BP}(P', E_0)$ produced by BEBOP is also a path of program P . The question then is: does p represent a feasible execution path of P ? That is, is there some execution of program P that follows the path p ? If so, we have found a real error in P . If not, path p is a spurious error path.

The NEWTON tool takes a C program and a (potential) error path as an input. It then uses verification condition generation (VCGen) to determine if the path is feasible. The answer may be “yes” or “no”.⁵

If the answer is “yes”, then an error path has been found, and we report it to the user. If the answer is “no” then NEWTON uses a new algorithm to identify a small set of predicates that “explain” why the path is infeasible.

In the running example, NEWTON detects that the path p is infeasible, and returns a single predicate $nPackets = nPacketsOld$ as the explanation for the infeasibility. This is because the predicate $nPackets = nPacketsOld$ is required to be both true and false by path p . The assignment of `nPacketsOld` to `nPackets` makes the predicate true, and the loop test requires it to be false at the end of the **do-while** loop for the loop to iterate, as specified by the path p .

2.6 The Second Boolean Program

In the second iteration of the process, the predicate $nPackets = nPacketsOld$ is added to the set of predicates E_0 to result in a new set of predicates E_1 . Figure 2(b) shows the boolean program $\mathcal{BP}(P', E_1)$ that C2BP produces. This program has one additional boolean variable (b) that represents the predicate $nPackets = nPacketsOld$. The assignment statement `nPackets = nPacketsOld`; makes this condition true, so in the boolean program the assignment `b:=1`; represents this assignment. Using a theorem prover, C2BP determines that if the predicate is true before the statement `nPackets++`, then it is false afterwards. This is captured by the assignment statement in the boolean program

```
b := choose(0,b);
```

The `choose` function is defined as follows:

```
bool choose(pos, neg)
begin
  if (pos) then return 1;
  elsif (neg) then return 0;
  elsif (*) then return 0;
  else return 1; fi
end
```

The `pos` parameter represents positive information about a predicate while the `neg` parameter represents negative information about a predicate. The `choose` function is never

⁵ Since underlying decision procedures in the theorem prover are incomplete, “don’t know” is also a possible answer. In practice, the theorem provers we use have been able to give a “yes” or “no” answer in every example we have seen so far.

called with both parameters evaluating to true. If both parameters are false then there is not enough information to determine whether the predicate is definitely true or definitely false, so 0 or 1 is returned, non-deterministically.

Applying BEBOP to the new boolean program shows that the label `ERROR` is not reachable. In performing its fixpoint computation, BEBOP discovers that the following loop invariant holds at the end of the **do-while** loop:

$$\begin{aligned} & (state = Locked \wedge nPackets = nPacketsOld) \\ \vee & (state = Unlocked \wedge nPackets \neq nPacketsOld) \end{aligned}$$

That is, either the lock is held and the loop will terminate (and thus the lock needs to be released after the loop), or the lock is free and the loop will iterate. The combination of predicate abstraction of C2BP and the fixpoint computation of BEBOP has found this loop-invariant over the predicates in E_1 . This loop-invariant is strong enough to show that the label `ERROR` is not reachable.

3 C2BP: A Predicate Abstractor For C

C2BP takes a C program P and a set $E = \{e_1, e_2, \dots, e_n\}$ of predicates on the variables of P , and automatically constructs a boolean program $\mathcal{BP}(P, E)$. [1, 2] The set of predicates E are pure C boolean expressions with no function calls. The boolean program $\mathcal{BP}(P, E)$ contains n boolean variables $V = \{b_1, b_2, \dots, b_n\}$, where each boolean variable b_i represents a predicate e_i . Each variable in V has a *three-valued* domain: **false**, **true**, and $*$.⁶ The program $\mathcal{BP}(P, E)$ is a *sound* abstraction of P because every possible execution trace t of P has a corresponding execution trace t' in B . Furthermore, $\mathcal{BP}(P, E_0)$ is a precise abstraction of P with respect to the set of predicates E_0 , in a sense stated and shown elsewhere [2]. Since $\mathcal{BP}(P, E)$ is an abstraction of P , it is guaranteed that an invariant I discovered (by BEBOP) in $\mathcal{BP}(P, E)$, as boolean combinations of the b_i variables, is also an invariant in the C code, where each b_i is replaced by its corresponding predicate e_i .

C2BP determines, for every statement s in P and every predicate $e_i \in E$, how the execution of s can affect the truth value of e_i . This is captured in the boolean program by a statement s_B that conservatively updates each b_i to reflect the change. C2BP computes s_B by (1) first computing the weakest precondition of e_i , and its negation with respect to s , and (2) strengthening the weakest precondition in terms of predicates from E , using a theorem prover.

We highlight the technical issues in building a tool like C2BP:

- **Pointers:** We use an alias analysis of the C program to determine whether or not an update through a pointer dereference can potentially affect an expression. This greatly increases the precision of the C2BP tool. Without alias analysis, we would have to make very conservative assumptions about aliasing, which would lead to invalidating many predicates.
- **Procedure calls:** Since boolean programs support procedure calls, we are able to abstract procedures modularly. The abstraction process for procedure calls is challenging, particularly in the presence of pointers. After a call, the caller must conservatively update local state that may have been modified by the callee. We provide a sound and precise approach to abstracting procedure calls that takes such side-effects into account.

⁶ The use of the third value $*$, is encoded using control-nondeterminism as shown in the `choose` function of Section 2. That is, “ $*$ ” is equivalent to “`choose(0,0)`”.

- **Precision-efficiency tradeoff.** C2BP uses a theorem prover to strengthen weakest pre-conditions in terms of the given predicate set E . Doing this strengthening precisely requires $O(2^E)$ calls to the theorem prover. We have explored a number of optimization techniques to reduce the number of calls made to the theorem prover. Some of these techniques result in an equivalent boolean program, while others trade off precision for computation speed. We are also investigating using other decision procedures, such as those embodied in the Omega test [25] and PVS [23].

Complexity. The runtime of C2BP is dominated by calls to the theorem prover. In the worst-case, the number of calls made to the theorem prover for computing $\mathcal{BP}(P, E)$ is linear in the size of P and exponential in the size of E . In practice, we find that the complexity is cubic in the size of E .

4 BEBOP: A Model Checker for Boolean Programs

The BEBOP tool [3] computes the set of reachable states for each statement of a boolean program using an interprocedural dataflow analysis algorithm in the spirit of Sharir/Pnueli and Reps/Horwitz/Sagiv [29, 26]. A state of a boolean program at a statement s is simply a valuation to the boolean variables that are in scope at statement s (in other words, a bit vector, with one bit for each variable in scope). The set of reachable states (or invariant) of a boolean program at s is thus a set of bit vectors (equivalently, a boolean function over the set of variables in scope at s).

BEBOP differs from typical implementations of dataflow algorithms in two crucial ways. First, it computes over sets of bit vectors at each statement rather than single bit vectors. This is necessary to capture correlation between variables. Second, it uses binary decision diagrams [4] (BDDs) to implicitly represent the set of reachable states of a program, as well as the transfer functions for each statement in a boolean program. BEBOP also differs from previous model checking algorithms for finite state machines, in that it does not inline procedure calls, and exploits locality of variable scopes for better scaling. Unlike most model checkers for finite state machines, BEBOP naturally generalizes to handle recursive and mutually recursive procedures. BEBOP uses an explicit control-flow graph representation, as in a compiler, rather than encoding the control-flow with BDDs, as done in most symbolic model checkers. It computes a fixpoint by iterating over the set of facts associated with each statement, which are represented with BDDs.

Complexity. The worst-case complexity of BEBOP is linear in the size of the programs control-flow graph, and exponential in the maximum number of boolean variables in scope at any program-point. We have implemented a number of optimizations to reduce the number of variables needed in support of BDDs. For example, we use live variable analysis to find program points where a variable becomes dead and then eliminate the variable from the BDD representation. We also use a global MOD/REF analysis of the boolean program in order to perform similar variable eliminations.

5 NEWTON: A Predicate Discoverer

NEWTON takes a C program P and an error path p from a boolean program B as inputs. It is assumed that the boolean program B was produced by running C2BP on P with some

```

s1:  nPacketsOld = nPackets;
s2:  request = devExt->WriteListHeadVa;
s3:  assume(!request);
s4:  assume(nPackets != nPacketsOld);

```

loc.	value dep.	condition-set dep.
1. <i>nPackets</i> :	α ()	
2. <i>nPacketsOld</i> :	α (1)	

after s1

loc.	value dep.	condition-set dep.
1. <i>nPackets</i> :	α ()	
2. <i>nPacketsOld</i> :	α (1)	
3. <i>devExt</i> :	β ()	
4. $\beta \rightarrow WriteListHeadVa$:	γ (3)	
5. <i>request</i> :	γ (3,4)	

after s2

loc.	value dep.	condition-set dep.
1. <i>nPackets</i> :	α ()	$!(\gamma)$ (5)
2. <i>nPacketsOld</i> :	α (1)	
3. <i>devExt</i> :	β ()	
4. $\beta \rightarrow WriteListHeadVa$:	γ (3)	
5. <i>request</i> :	γ (3,4)	

after s3

loc.	value dep.	condition-set dep.
1. <i>nPackets</i> :	α ()	$!(\gamma)$ (5)
2. <i>nPacketsOld</i> :	α (1)	$(\alpha \neq \alpha)$ (1,2)
3. <i>devExt</i> :	β ()	
4. $\beta \rightarrow WriteListHeadVa$:	γ (3)	
5. <i>request</i> :	γ (3,4)	

after s4

Fig. 3. A path of four statements and four tables showing the state of NEWTON after simulating each of the four statements.

set of predicates. For the purposes of describing NEWTON, we can identify the path p as a sequence of assignments and assume statements (every conditional is translated into an assume statement).

The internal state of NEWTON has three components: (1) *store*, which is a mapping from locations to symbolic expressions, (2) *conditions*, which is a set of predicates, and (3) a *history* which is a set of past associations between locations and symbolic expressions. The high-level description of NEWTON is given in Figure 4. The functions LEval and REval evaluate the l-value and r-value of a given expression respectively. NEWTON maintains an internal dependency of each element in the state with the elements in *store*, to be used in Phase 3. It also uses symbolic constants for unknown values. We illustrate these using an example. Consider a path with the following four statements:

```

s1:  nPacketsOld = nPackets;
s2:  request = devExt->WriteListHeadVa;
s3:  assume(!request);
s4:  assume(nPackets != nPacketsOld);

```

This path is a projection of the error path from BEBOP in Section 2. Figure 3 shows four states of NEWTON, one after processing each statement in the path. The assignment `nPacketsOld = nPackets` is processed by first introducing a symbolic constant α for the variable `nPackets`, and then assigning it to `nPacketsOld`. The assignment `request = devExt->WriteListHeadVa` is processed by first introducing a symbolic constant β for the value of `devExt`, then introducing a second symbolic constant γ for the value of `β ->WriteListHeadVa`, and finally assigning γ to `request`. The conditional `assume(!request)` is processed by adding the predicate $!(\gamma)$ to the condition-set. The dependency list for this predicate is (5) since its evaluation depended on entry 5 in the store. Finally, the conditional

```

Input: A sequence of statements  $p = s_1, s_2, \dots, s_m$ .
store := null map;
history := null set;
conditions := null set;
/* start of Phase 1 */
for  $i = 1$  to  $m$  do {
  switch(  $s_i$  ) {
    " $e_1 := e_2$ " :
      let  $lval = \text{LEval}(store, e_1)$  and
      let  $rval = \text{REval}(store, e_2)$  in
        if( $store[lval]$  is defined)
           $history := history \cup \{(lval, store[lval])\}$ 
           $store[lval] := rval$ 
    "assume( $e$ )" :
      let  $rval = \text{REval}(store, e)$  in
         $conditions := conditions \cup \{rval\}$ 
        if( $conditions$  is inconsistent){
          /*Phase 2 */
          Minimize size of  $conditions$ 
            while maintaining inconsistency
          /*Phase 3 */
           $predicates :=$  all dependencies of  $conditions$ 
          Say "Path  $p$  is infeasible"
          return( $predicates$ )
        }
  }
}
Say "Path  $p$  is feasible"
return

```

Fig. 4. The high-level algorithm used by NEWTON

```

VOID
SerialDebugLogEntry(IN ULONG Mask, IN ULONG Sig,
  IN ULONG_PTR Info1, IN ULONG_PTR Info2, IN ULONG_PTR Info3)
{
  KIRQL irql;
  irql = KeGetCurrentIrql();
  if (irql < DISPATCH_LEVEL) {
    KeAcquireSpinLock(&LogSpinLock, &irql);
  } else {
    KeAcquireSpinLockAtDpcLevel(&LogSpinLock);
  }
  // other code (deleted)
  if (irql < DISPATCH_LEVEL) {
    KeReleaseSpinLock(&LogSpinLock, irql);
  } else {
    KeReleaseSpinLockFromDpcLevel(&LogSpinLock);
  }
  return;
}

```

Fig. 5. Code snippet from serial-port driver.

`assume(nPackets != nPacketsOld)` is processed by adding the (inconsistent) predicate ($\alpha \neq \alpha$) to the condition-set, with a dependency list (1,2). At this point, the theorem prover determines that the condition-set is inconsistent, and NEWTON proceeds to the Phase 2.

Phase-2 removes the predicate $!(\gamma)$ from the condition store, since the remaining predicate ($\alpha \neq \alpha$) is inconsistent by itself. Phase-3 traverses store entries 1 and 2 from the dependency list. A post processing step then determines that the symbolic constant α can be unified with the variable `nPackets`, and NEWTON produces two predicates: ($nPacketsOld = nPackets$) and ($nPacketsOld \neq nPackets$). Since one is a negation of the other, only one of the two predicates needs to be added in order for the path to be ruled out in the boolean program. Though no symbolic constants are present in the final set of predicates in our example, there are other examples where the final list of predicates have symbolic constants. C2BP is able to handle predicates with symbolic constants. We do not discuss these details here due to want of space. The *history* is used when a location is overwritten with a new value. Since no location was written more than once in our example, we did not see the use of *history*. NEWTON also handles error paths where each element of the path is also provided with values to the boolean variables from BEBOP, and checks for their consistency with the concrete states along the path.

Complexity. The number of theorem-prover calls made by NEWTON on a path p is $O(p*n)$, where p is the length of the path, and n is the number of predicates in the boolean program B .

6 NT Device Drivers: Case Study

This section describes our experience applying the SLAM toolkit to check properties of Windows NT device drivers. We checked two kinds of properties: (1) Locking-unlocking sequences for locks should conform to allowable sequences (2) Dispatch functions should

either complete a request, or make a request pending for later processing. In either case, a particular sequence of Windows NT specific actions should be taken.

The two properties have different characteristics from a property-checking perspective.

- The first property depends mainly on the program’s control flow. We checked this property for a particular lock (called the “Cancel” spin lock) on three kernel mode drivers in the Windows NT device driver tool kit. We found two kinds situations where spurious error paths led our process to iterate. With its inter-procedural analysis and detection of simple variable correlations, the SLAM tools were able to eliminate all the spurious error paths with at most one iteration of the process. In all the drivers, we started with 5 predicates from the property specification FSM and added at most one predicate to rule out spurious error paths.
- The second property is data-dependent, requiring the tracking of value flow and alias relationships. We checked this property on a serial port device driver. It took 4 iterations through the SLAM tools and a total of 30 predicates to validate the property.

The drivers we analyzed were on the order of a thousand lines of C code each. In each of the drivers we checked for the first property, the SLAM tools ran in under a minute on an 800MHz Pentium PC with 512MB RAM. For the second property on the serial driver, the total run time for all the SLAM tools was about 3 minutes to complete all the four iterations.

6.1 Property 1

We checked for correct lock acquisition/release sequences on 3 kernel mode drivers: MCA-bus, serial-port and parallel-port. The SLAM tools validated the property on MCA-bus and parallel-port drivers without iteration. However, interprocedural analysis was required for checking the property, as calls to the acquire and release routines were spread across multiple procedures in the drivers. Furthermore, in the serial-port driver, the SLAM tools found one false error path in the first iteration, which resulted in the addition of a single predicate. Then the property was validated in the second iteration. The code-snippet that required the addition of the predicate is shown in Figure 5. The snippet shows that the code has a dependence on the interrupt request level variable (`irq1`) that must be tracked in order to eliminate the false error paths.

6.2 Property 2

A dispatch routine to a Windows NT device driver is a routine that the I/O manager calls when it wants the driver to perform a specific operation (e.g, read, write etc.) The dispatch routine is “registered” by the driver during ‘when it is initialized. A dispatch routine has the following prototype:

```
NTSTATUS DispatchX(IN PDEVICE_OBJECT DeviceObject,  
                 IN PIRP  irp)
```

The first parameter is a pointer to a so called “device object” that represents the device, and the second parameter is a pointer to a so called “I/O request packet”, or “irp” that contains information about the current request.

<pre> if (status != STATUS_PENDING) { Irp->Status = status; IoCompleteRequest(Irp, 0); } </pre>	<pre> if (status != STATUS_PENDING) { Irp->Status = status; if (Irp==i) { EMIT_FSM(ASSIGN, status); } IoCompleteRequest(Irp, 0); if (Irp==i) { EMIT_FSM(CALL_IOCOMPLETE, status); } } </pre>
--	---

Fig. 6. Code snippet from a driver (left) and the instrumentation added (right).

Property P1. All dispatch routines must either process the irp immediately (call this *option A*, or queue the irp for processing later (call this *option B*). Every irp has to be processed under one of these two options. If the driver chooses option A, then it has to do the following actions in sequence:

1. Set `irp->IoS.status` to `STATUS_PENDING`
2. Call the kernel function `IoMarkIrpPending(irp)`
3. Queue the irp into the driver's internal queue using the kernel function `IoStartPacket(irp)`
4. Return `STATUS_PENDING`

If the driver chooses option B, then it has to do the following actions in sequence:

1. Set the `irp->IoS.status` to some return code other than `STATUS_PENDING` (such as `STATUS_SUCCESS`, `STATUS_CANCELLED` etc.)
2. Call `IoCompleteRequest(irp)`
3. Return the same status code as in step 1.

Note that this is a partial specification for a dispatch routine—just one of several properties that the dispatch routine must obey. We first coded up the above property as a finite state machine with a transition function named `EMIT_FSM` that takes two parameters: an action (such as `CALL_IOCOMPLETE`, `CALL_QUEUEIRP`, etc.) and a status (such as `STATUS_PENDING`, etc).

Instrumenting the driver code. In order to check if the driver code satisfies the property we added instrumentation code to the driver. At the entry point to the driver, we store the value of the irp in a new global, `gIrp`. Every time a kernel function `IoCompleteRequest(irp)`, `IoMarkIrpPending(irp)`, or `IoStartIrp(irp)` is called, we check if the `irp` parameter is the same as `gIrp`, and if so we add a call to `EMIT_FSM` with the appropriate message as the second parameter. Every time a variable of type `PIRP` has the status field assigned, we check if the `irp` parameter is the same as `gIrp`, and if so we add a call to `EMIT_FSM` with the status as the first parameter, and `ASSIGN` as the second parameter. Figure 6 shows a sample code snippet from a driver and the instrumentation we add.⁷

Checking the instrumented driver. The initial set of predicates described the FSM includes 17 predicates: (1) 5 predicates to keep track of the global variable `fsmState` (2)

⁷ We remind the reader that in the future, we plan to have a tool that will generate such instrumentation automatically from a high-level specification of the property. For now, our ability to analyze properties is limited mainly by the need to hand-instrument the property into the code of interest.

4 predicates to keep track of the global variable `fsmStatus` (3) 3 predicates to keep track of the formal parameter `m` of `EMIT_FSM`, and (4) 5 predicates to keep track of the formal parameter `s` of function `EMIT_FSM`.

C2BP generated a boolean program B_1 using these 17 predicates, and BEBOP found an error trace that led to the label `ERROR`. NEWTON analyzed this error trace, and came up with 3 more predicates. These predicates kept track of the value of a local variable where a status value was stored before being assigned into the `irp->status`. After iterating through C2BP with these predicates added, BEBOP found a second error trace, which passed through 2 function calls, and NEWTON came up with 4 more predicates to be added. These predicates kept track of the flow of the `irp` pointer through the function call, and a local variable of the called function where the status value was stored temporarily.

After one more iteration through C2BP with the additional predicates, BEBOP found a third error trace, which passed through 3 levels of function calls. This error trace was fairly complicated, and it involved the driver storing the `irp` pointer in a global structure, passing a pointer to the structure, and then retrieving the pointer two levels of function calls later. When fed with this error trace, NEWTON came up with 9 more predicates to be added that tracked this value flow.

In the fourth iteration BEBOP was able to validate the property on the boolean program produced by C2BP with all the predicates discovered thus far. It took 4 iterations through the tools and a total of 30 predicates to discover the right abstraction to validate this property. We found one bug in the fourth iteration, but it turned out to be a cut-and-paste error in our instrumentation process. After fixing it, the property passed.

7 Related Work

SLAM seeks a sweet spot between VCGen-based tools [16, 22, 5] that operate directly on the concrete semantics and model checking or data flow-analysis based tools [7, 18, 13, 11] that operate on abstractions of the program. We use VCGen-based approach on finite (potentially interprocedural) paths of the program, and use the knowledge gained to construct abstract models of the program. NEWTON uses VCGen on the concrete program, but as it operates on a single finite interprocedural path at a time, it does not require loop-invariants, or pre-conditions and post-conditions for procedures. C2BP also reasons about the statements of the C program using decision procedures, but does so only locally, one statement at a time. Global analysis is done only on the boolean program abstractions, using the model checker BEBOP. Thus, our hope is to scale without losing precision, as long as the property of interest allows us to do so, by inherently requiring a small abstraction for its validation or invalidation.

SLAM generalizes Engler et al.’s approach in three ways: (1) it is sound (modulo the assumptions about memory safety); (2) it permits interprocedural analysis; (3) it avoids spurious examples through iterative refinement (in some of the code Engler et al. report on, their techniques generated three times as many spurious error paths as true error paths, a miss rate of 300 percent.⁸) In fact, with a suitable definition of abstraction, and choice of initial predicates, the first iteration of the SLAM process is equivalent to performing Engler et al.’s approach interprocedurally.

⁸ Jon Pincus, who led the development of an industrial-strength error detection tool for C called PREFIX [5], observes that users of PREFIX will tolerate a false alarm rate in the range 25-50% depending on the application [24].

Constructing abstract models of programs has been studied in several contexts. Abstractions constructed by [13] and [19] are based on specifying transitions in the abstract system using a pattern language, or as a table of rules. Automatic abstraction support has been built into the Bandera tool set [12]. They require the user to provide finite domain abstractions of data types. Predicate abstraction, as implemented in C2BP is more general, and can capture relationships between variables. The predicate abstraction in SLAM was inspired by the work of Graf and Saidi [17] in the model checking community. Efforts have been made to integrate predicate abstraction with theorem proving and model checking [27]. Though our use of predicate abstraction is related to these efforts, our goal is to analyze software written in common programming languages. A predicate abstraction tool for Java has recently been reported in [31].

The SLAM tools C2BP and BEBOP can be used in combination to find loop-invariants expressible as boolean functions over a given set of predicates. The loop-invariant is computed by the model checker BEBOP using a fixpoint computation on the abstraction computed by C2BP. Prior work for generating loop invariants has used symbolic execution on the concrete semantics, augmented with widening heuristics [30, 32]. The Houdini tool guesses a candidate set of annotations (invariants) and uses the ESC/Java checker to refute inconsistent annotations until convergence [15].

Boolean programs can be viewed as abstract interpretations of the underlying program [8]. The connections between model checking, dataflow analysis and abstract interpretation have been explored before [28] [9]. The model checker BEBOP is based on earlier work on interprocedural dataflow analysis [29, 26]. Automatic iterative refinement based on error paths first appeared in [20], and more recently in [6]. Both efforts deal with finite state systems.

An alternative approach to static validation of safety properties, is to provide a rich type system that allows users to encode both safety properties and program annotations as types, and reduce property validation to type checking [14].

8 Conclusions

We conclude by summarizing the main contributions of our work:

- We have presented a fully automated methodology to validate/invalidate temporal safety properties of software interfaces. Our process does not require user supplied annotations, or user supplied abstractions. When our process converges, it always give a definitive “yes” or “no” answer.
- The ideas behind the SLAM tools are novel. The use of boolean programs to represent program abstractions is new. To the best of our knowledge, C2BP is the first automatic predicate abstraction tools to handle a full-scale programming language, and perform a sound and precise abstraction. BEBOP is the first model checker to handle procedure calls using an interprocedural dataflow analysis algorithm, augmented with representation tricks from the symbolic model checking community. NEWTON uses a path simulation algorithm in a novel way, to generate predicates for refinement.
- We have demonstrated that the SLAM tools converge in a few iterations on device drivers from the Microsoft DDK.

The SLAM toolkit has a number of limitations that we plan to address. The logical model of memory is a limitation, since it is not the actual model used by C programs. We plan to

investigate using a physical model of memory. We are working on a property specification language, and automatic instrumentation of the source code from the specification language. We are exploring theoretical guarantees we can give about the termination of our iterative refinement. We plan to evolve the SLAM tools by applying them to more code bases, both inside and outside Microsoft.

Acknowledgements

We thank Rupak Majumdar and Todd Millstein for their hard work in making the C2BP tool come to life. Thanks to Andreas Podelski for helping us describe the C2BP tool in terms of abstract interpretation. Thanks also to the members of the Software Productivity Tools research group at Microsoft Research for many enlightening discussions on program analysis, programming language and device drivers, as well as their numerous contributions to the SLAM toolkit.

References

1. T. Ball, R. Majumdar, T. Millstein, and S. K. Rajamani. Automatic predicate abstraction of C programs. In *PLDI 2001 (to appear)*.
2. T. Ball, A. Podelski, and S. K. Rajamani. Boolean and cartesian abstractions for model checking c programs. Technical Report (to appear, TACAS 2001) MSR Technical Report 2000-115, Microsoft Research, December 2000.
3. T. Ball and S. K. Rajamani. Bebop: A symbolic model checker for boolean programs. In *Proceedings of the 7th International SPIN Workshop (Lecture Notes in Computer Science No. 1885)*, pages 113–130. Springer-Verlag, September 2000.
4. R. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Transactions on Computers*, C-35(8):677–691, 1986.
5. W. R. Bush, J. D. Pincus, and D. J. Sielaff. A static analyzer for finding dynamic programming errors. *Software-Practice and Experience*, 30(7):775–802, June 2000.
6. E. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith. Counterexample-guided abstraction refinement. In *Proceedings of the 12th International Conference on Computer Aided Verification (LNCS No. 1855)*, pages 154–169. Springer, July 2000.
7. J. Corbett, M. Dwyer, J. Hatcliff, C. Pasareanu, Robby, S. Laubach, and H. Zheng. Bandera : Extracting finite-state models from Java source code. In *Proceedings of the 22nd International Conference on Software Engineering*, June 2000.
8. P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for the static analysis of programs by construction or approximation of fixpoints. In *Proceedings of the Fourth Annual Symposium on Principles of Programming Languages*. ACM Press, 1977.
9. P. Cousot and R. Cousot. Temporal abstract interpretation. In *Proceedings of the Twenty Seventh Annual Symposium on Principles of Programming Languages*. ACM Press, 2000.
10. M. Das. Unification-based pointer analysis with directional assignments. In *PLDI '00: Proceedings of the ACM SIGPLAN 2000 Conference on Programming Language Design and Implementation*, June 2000.
11. R. DeLine and M. Fähndrich. Enforcing high-level protocols in low-level software. In *PLDI 2001 (to appear)*.
12. M. Dwyer and L. Clarke. Data flow analysis for verifying properties of concurrent programs. In *Proceedings of the Second ACM SIGSOFT symposium on Foundations of Software Engineering*, pages 62–75, December 1994.
13. M. Dwyer, J. Hatcliff, R. Joehanes, S. Laubach, C. Pasareanu, Robby, W. Visser, and H. Zheng. Tool-supported program abstraction for finite-state verification. In *Proceedings of the 22nd International Conference on Software Engineering (to appear)*, June 2001.
14. D. Engler, B. Chelf, A. Chou, and S. Hallem. Checking system rules using system-specific, programmer-written compiler extensions. In *Proceedings of 4th Symposium on Operating System Design and Implementation*. Usenix Association, October 2000.

15. C. Flanagan, R. Joshi, and K. R. M. Leino. Annotation inference for modular checkers. *Information Processing Letters*, 2000. To appear.
16. C. Flanagan and J. B. Saxe. Generating compact verification conditions. In *POPL '2001 (to appear)*.
17. S. Graf and H. Saidi. Construction of abstract state graphs with PVS. In *CAV 97: Computer-aided Verification*, Lecture Notes in Computer Science 1254, pages 72–83. Springer-Verlag, 1997.
18. G. Holzmann. The Spin model checker. *IEEE Trans. on Software Engineering*, 23(5):279–295, May 1997.
19. G. Holzmann. Logic verification of ANSI-C code with Spin. In *Proceedings of the SPIN 2000 Workshop*, pages 131–147. Springer Verlag / LNCS 1885, Sep. 2000.
20. R. Kurshan. *Computer-aided Verification of Coordinating Processes*. Princeton University Press, 1994.
21. L. Lamport. Proving the correctness of multiprocess programs. *IEEE Transactions on Software Engineering*, SE-3(2):125–143, 1977.
22. G. Necula. Proof carrying code. In *Proceedings of the 24th Annual Symposium on Principles of Programming Languages*, pages 106–119. ACM Press, 1997.
23. S. Owre, S. Rajan, J. Rushby, N. Shankar, and M. Srivas. PVS: Combining specification, proof checking, and model checking. In R. Alur and T. A. Henzinger, editors, *Computer-Aided Verification, CAV '96*, volume 1102 of *Lecture Notes in Computer Science*, pages 411–414, New Brunswick, NJ, July/August 1996. Springer-Verlag.
24. J. Pincus. personal communication, October 2000.
25. W. Pugh. A practical algorithm for exact array dependence analysis. *Commun. ACM*, 35(8):102–114, August 1992.
26. T. Reps, S. Horwitz, and M. Sagiv. Precise interprocedural dataflow analysis via graph reachability. In *Proceedings of the 22nd ACM Symposium on Principles of Programming Languages*, pages 49–61, January 1995.
27. H. Saidi and N. Shankar. Abstract and model check while you prove. In *CAV 99: Computer-aided Verification*, Lecture Notes in Computer Science 1633, pages 443–454. Springer-Verlag, 1999.
28. D. Schmidt. Data flow analysis is model checking of abstract interpretation. In *Proceedings of the Twenty Fifth Annual Symposium on Principles of Programming Languages*, pages 38–48. ACM Press, 1998.
29. M. Sharir and A. Pnueli. Two approaches to interprocedural data dalow analysis. In S. Muchnick and N. Jones, editors, *Program Flow Analysis: Theory and Applications*, pages 189–233. Prentice-Hall, 1981.
30. N. Suzuki and K. Ishihata. Implementation of an array bound checker. In *Proceedings of the ACM Symposium on Principles of Programming Languages*. ACM, January 1977.
31. W. Visser, S. Park, and J. Penix. Using predicate abstraction to reduce object-oriented programs for model checking. In *The Third Workshop on Formal Methods in Software Practice*, pages 3–12. ACM, August 2000.
32. Z. Xu, B. P. Miller, and T. Reps. Safety checking of machine code. In *Proceedings of the ACM SIGPLAN '00 Conference on Programming Language Design and Implementation*, pages 70–82. ACM, June 2000.