

Optimizing Symbolic Model Checking for Constraint-Rich Models

Bwolen Yang, Reid Simmons, Randal E. Bryant, and David R. O'Hallaron

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
{bwolen, reids, bryant, droh}@cs.cmu.edu

Abstract. This paper presents optimizations for verifying systems with complex time-invariant constraints. These constraints arise naturally from modeling physical systems, e.g., in establishing the relationship between different components in a system. To verify constraint-rich systems, we propose two new optimizations. The first optimization is a simple, yet powerful, extension of the conjunctive-partitioning algorithm. The second is a collection of BDD-based macro-extraction and macro-expansion algorithms to remove state variables. We show that these two optimizations are essential in verifying constraint-rich problems; in particular, this work has enabled the verification of fault diagnosis models of the Nomad robot (an Antarctic meteorite explorer) and of the NASA Deep Space One spacecraft.

1 Introduction

This paper presents techniques for using symbolic model checking to automatically verify a class of real-world applications that have many time-invariant constraints. An example of constraint-rich systems is the symbolic models developed by NASA for on-line fault diagnosis [15]. These models describe the operation of components in complex electro-mechanical systems, such as autonomous spacecraft or robot explorers. The models consist of interconnected components (e.g., thrusters, sensors, motors, computers, and valves) and describe how the *mode* of each component changes over time. Based on these models, the Livingstone diagnostic engine [15] monitors sensor values and detects, diagnoses, and tries to recover from inconsistencies between the observed sensor values and the predicted modes of the components. The relationships between the modes and sensor values are encoded using symbolic constraints. Constraints

Effort sponsored in part by the Advanced Research Projects Agency and Rome Laboratory, Air Force Materiel Command, USAF, under agreement number F30602-96-1-0287, in part by the National Science Foundation under Grant CMS-9318163, and in part by grants from the Intel Corporation and NASA Ames Research Center. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Advanced Research Projects Agency, Rome Laboratory, or the U.S. Government.

between state variables are also used to encode interconnections between components. We have developed an automatic translator from such fault models to SMV (Symbolic Model Verifier) [10], where mode transitions are encoded as transition relations and state-variable constraints are translated into sets of time-invariant constraints.

To verify constraint-rich systems, we introduce two new optimizations. The first optimization is a simple extension of the conjunctive-partitioning algorithm. The other is a collection of BDD-based macro-extraction and macro-expansion algorithms to remove redundant state variables. We show that these two optimizations are essential in verifying constraint-rich problems. In particular, these optimizations have enabled the verification of fault diagnosis models for the Nomad robot (an Antarctic meteorite explorer) [1] and the NASA Deep Space One (DS1) spacecraft [2]. These models can be quite large, with up to 1200 state bits.

The rest of this paper is organized as follows. We first briefly describe symbolic model checking and how time-invariant constraints arise naturally from modeling (Section 2). We then present our new optimizations: an extension to conjunctive partitioning (Section 3), and BDD-based algorithms for eliminating redundant state variables (Section 4). We then show the results of a performance evaluation on the effects of each optimization (Section 5). Finally, we present a comparison to prior work (Section 6) and some concluding remarks (Section 7).

2 Background

Symbolic model checking [5, 6, 10] is a fully automatic verification paradigm that checks temporal properties (e.g., safety, liveness, fairness, etc.) of finite state systems by symbolic state traversal. The core enabling technology for symbolic model checking is the use of the Binary Decision Diagram (BDD) representation [4] for state sets and state transitions. BDDs represent Boolean formulas canonically as directed acyclic graphs such that equivalent sub-formulas are uniquely represented as a single subgraph. This uniqueness property makes BDDs compact and enables dynamic programming to be used for computing Boolean operations symbolically.

To use BDDs in model checking, we need to map sets of states, state transitions, and state traversal to the Boolean domain. In this section, we briefly describe this mapping and motivate how time-invariant constraints arise. We then finish with definitions of some additional terminology to be used in the rest of the paper.

2.1 Representing State Sets and Transitions

In the symbolic model checking of finite state systems, a state typically describes the values of many components (e.g., latches in digital circuits) and each component is represented by a *state variable*. Let $V = \{v_1, \dots, v_n\}$ be the set of state variables in a system, then a state can be described by assigning values to all the variables in V . This valuation can in turn be written as a Boolean formula that is true exactly for the valuation as $\bigwedge_{i=1}^n (v_i == c_i)$, where c_i is the value assigned to the variable v_i , and the “==” represents the equality operator in a predicate (similar to the C programming language). A set of states can be represented as a disjunction of the Boolean formulas

that represent the states. We denote the BDD representation for a set of states S by $S(V)$.

In addition to the set of states, we also need to map the system's state transitions to the Boolean domain. We extend the above concept of representing a set of states to representing a set of ordered-pairs of states. To represent a pair of states, we need two sets of state variables: V the set of *present-state variables* for the first tuple and V' the set of *next-state variables* for the second tuple. Each variable v in V has a corresponding next-state variable v' in V' . A valuation of variables in V and V' can be viewed as a state transition from one state to another. A transition relation can then be represented as a set of these valuations. We denote the BDD representation of a transition relation T as $T(V, V')$.

In modeling finite state systems, the overall state transitions are generally specified by defining the valid transitions for each state variable. To support non-deterministic transitions of a state variable, the expression that defines the transitions evaluates to a set, and the next-state value of the state variable is non-deterministically chosen from the elements in the set. Hereafter, we refer to an expression that evaluates to a set either as a *set expression* or as a *non-deterministic expression* depending on the context, and we use the bold font type, as in \mathbf{f} , to represent such expression. Let \mathbf{f}_i be the set expression representing state transitions of the state variable v_i . Then the BDD representation for v_i 's transition relation T_i can be defined as $T_i(V, V') := (v'_i \in \mathbf{f}_i(V))$. For synchronous systems, the BDD for the overall state transition relation T is $T(V, V') := \bigwedge_{i=0}^n T_i(V, V')$. Detailed descriptions on this formulation, including mapping of asynchronous systems, can be found in [5].

2.2 Time-Invariant Constraints and Their Common Usages

In symbolic model checking, time-invariant constraints specify the conditions that must always hold. More formally, let C_1, \dots, C_l be the time-invariant constraints and let $C := C_1 \wedge C_2 \wedge \dots \wedge C_l$. Then, in symbolic state traversal, we consider only states where C is true. We refer to C as the *constrained space*.

To motivate how time-invariant constraints arise naturally in modeling complex systems, we describe three common usages. One common usage is to make the same non-deterministic choice across multiple expressions in transition relations. For example, in a master-slave model, the master can non-deterministically choose which set of idle slaves to assign the pending jobs, and the slaves' next-state values will depend on the choice made. To model this, let \mathbf{f} be a non-deterministic expression representing how the master makes its choice. If \mathbf{f} is used multiple times, then each use makes a non-deterministic choice independently of other uses. Thus, to ensure that the same non-deterministic choice is seen by the slaves, a new state variable u is introduced to record the choice made, and u is then used to define the slaves' transition relations. This recording process is expressed as the time-invariant constraint $u \in \mathbf{f}$.

Another common usage is for establishing the interface between different components in a system. For example, suppose two components are connected with a pipe of a fixed capacity. Then, the input of one component is the minimum of the pipe's capacity and the output of the other component. This relationship is described as a time-invariant constraint between the input and the output of these two components.

Third common usage is specific uses of generic parts. For example, a bi-directional fuel pipe may be used to connect two components. If we want to make sure the fuel flows only one way, we need to constrain the valves in the fuel pipe. These constraints are specified as time-invariant constraints. In general, specific uses of generic parts arise naturally in both the software and the hardware domain as we often use generic building blocks in constructing a complex system.

In the examples above, the use of time-invariant constraints is not always necessary because some these constraints can be directly expressed as a part of the transition relation and the associated state variables can be removed. However, these constraints are used to facilitate the description of the system or to reflect the way complex systems are built. Without these constraints, multiple expressions will need to be combined into possibly a very complicated expression. Performing this transformation manually can be error-prone. Thus it is up to the verification tool to automatically perform these transformations and remove unnecessary state variables. Our optimizations for constraint-rich models is to automatically eliminate redundant state variables (Section 4) and partition the remaining constraints (Section 3).

2.3 Symbolic State Traversal

To reason about temporal properties, the *pre-image* and the *image* of the transition relation are used for symbolic state traversal, and time-invariant constraints are used to restrict the valid state space. Based on the BDD representations of a state set S and the transition relation T , we can compute the *pre-image* and the *image* of S , while restricting the computations to the constrained space C , as follows:

$$\text{pre-image}(S)(V) := C(V) \wedge \exists V'. [T(V, V') \wedge (S(V') \wedge C(V'))] \quad (1)$$

$$\text{image}(S)(V') := C(V') \wedge \exists V. [T(V, V') \wedge (S(V) \wedge C(V))] \quad (2)$$

One limitation of the BDD representation is that the monolithic BDD for the transition relation T is often too large to build. A solution to this problem is the *conjunctive partitioning* [5] of the transition relation. In conjunctive partitioning, the transition relation is represented as a conjunction $P_1 \wedge P_2 \wedge \dots \wedge P_k$ with each conjunct P_i represented by a BDD. Then, the pre-image can be computed by conjuncting with one P_i at a time, and by using *early quantification* to quantify out variables as soon as possible. The early-quantification optimization is based on the property that sub-formulas can be moved out of the scope of an existential quantification if they do not depend on any of the variables being quantified. Formally, let V'_i , a subset of V' , be the set of variables that do not appear in any of the subsequent P_j 's, where $1 \leq i \leq k$ and $i < j \leq k$. Then the pre-image can be computed as

$$\begin{aligned} p_1 &:= \exists V'_1. [P_1(V, V') \wedge (S(V') \wedge C(V'))] & (3) \\ p_2 &:= \exists V'_2. [P_2(V, V') \wedge p_1] \\ &\vdots \\ p_k &:= \exists V'_k. [P_k(V, V') \wedge p_{k-1}] \\ \text{pre-image}(S)(V) &:= C(V) \wedge p_k \end{aligned}$$

The determination and ordering of partitions (the P_i 's in above) can have significant performance impact. Commonly used heuristics [7, 11] treat the state variables' transition relations (T_i 's) as conjuncts. The ordering step then greedily schedules the partitions to quantify out more variables as soon as possible, while introducing fewer new variables. Finally, the ordered partitions are tentatively merged with their predecessors to reduce the number of intermediate results. Each merged result is kept only if the resulting graph size is less than a pre-determined limit.

The conjunctive partitioning for the image computation is performed similarly with present-state variables in V being the quantifying variables instead of next-state variables in V' . However, since the quantifying variables are different between the image and the pre-image computation, the resulting conjuncts for image computation is typically very different from those for pre-image computation.

2.4 Additional Terminology

We define the *ITE* operator (if-then-else) as follows: given arbitrary expressions f and g where f and g may both be set expressions, and Boolean expression p , then

$$ITE(p, f, g)(X) := \begin{cases} f(X) & \text{if } p(X); \\ g(X) & \text{otherwise;} \end{cases}$$

where X is the set of variables used in expressions p , f , and g . We define a *care-space optimization* as any algorithm *care-opt* that has following properties: given an arbitrary function f where f may be a set expression, and a Boolean formula c , then

$$care-opt(f, c) := ITE(c, f, d),$$

where d is defined by the particular algorithm used. The usual interpretation of this is that we only *care* about the values of f when c is true. We will refer to c as the *care space* and $\neg c$ as the *don't-care space*. The goal of care-space optimizations is to heuristically minimize the representation for f by choosing a suitable d in the don't-care space. Descriptions and a study of some care-space optimizations, including the commonly used *restrict* algorithm [6], can be found in [13].

3 Extended Conjunctive Partitioning

The first optimization is the application of the conjunctive-partitioning algorithm on the time-invariant constraints. This extension is derived based on two observations. First, as with the transition relations, the BDD representation for time-invariant constraints can be too large to be represented as a monolithic graph. Thus, it is crucial to represent the constraints as a set of conjuncts rather than a monolithic graph.

Second, in constraint-rich models, many *quantifying variables* (variables being quantified) do not appear in the transition relation. There are two common causes for this. First, when time-invariant constraints are used to make the same non-deterministic choices, new variables are introduced to record these choices (described as the first example in Section 2.2). In the transition relation, these new variables are used only in

their present-state form. Thus, their corresponding next-state variables do not appear in the transition relation, and for the pre-image computation, these next-state variables are parts of the quantifying variables. The other cause is that many state variables are used only to establish time-invariant constraints. Thus, both the present- and the next-state version of these variables do not appear in the transition relations.

Based on this observation, we can improve the early-quantification optimization by pulling out the quantifying variables (V'_0) that do not appear in any of the transition relations. Then, these quantifying variables (V'_0) can be used for early quantification in conjunctive partitioning of the constrained space (C) where the time-invariant constraints hold. Formally, let Q_1, Q_2, \dots, Q_m be the partitions produced by the conjunctive partitioning of the constrained space C , where $C = Q_1 \wedge Q_2 \wedge \dots \wedge Q_m$. For the pre-image computation, Equation 3 is replaced by

$$\begin{aligned} q_1 &:= \exists W'_1. [Q_1(V') \wedge S(V')] \\ q_2 &:= \exists W'_2. [Q_2(V') \wedge q_1] \\ &\vdots \\ q_m &:= \exists W'_m. [Q_m(V') \wedge q_{m-1}] \\ p_1 &:= \exists V'_1. [P_1(V, V') \wedge q_m] \end{aligned}$$

where W'_i , a subset of V'_0 , is the set of variables that do not appear in any of the subsequent Q_j 's, where $1 \leq i \leq m$ and $i < j \leq m$. Similarly, this extension also applies to the image computation.

4 Elimination of Redundant State Variables

Our second optimization for constraint-rich models is targeted at reducing the state space by removing unnecessary state variables. This optimization is a set of BDD-based algorithms that compute an equivalent expression for each variable used in the time-invariant constraints (*macro extraction*) and then globally replace a suitable subset of variables with their equivalent expressions (*macro expansion*) to reduce the total number of variables.

The use of macros is traditionally supported by language constructs (e.g., DEFINE in the SMV language [10]) and by simple syntactic analyses such as detecting deterministic assignments (e.g., $a == f$ where a is a state variable and f is an expression) in the specifications. However, in constraint-rich models, the constraints are often specified in a more complex manner such as *conditional* dependencies on other state variables (e.g., $p \Rightarrow (a == f)$ as conditional assignment of expression f to variable a when p is true). To identify the set of valid macros in such models, we need to combine the effects of multiple constraints. For these models, one drawback of syntactic analysis is that, for each type of expression, syntactic analysis will need to add a template to pattern match these expressions. Another more severe drawback is that it is difficult for syntactic analysis to estimate the actual cost of instantiating a macro. Estimating this cost is important because reducing the number of variables by macro expansion can sometimes result in significant performance degradation caused by large increases

in other BDD sizes. These two drawbacks make the syntactic approach unsuitable for models with complex time-invariant constraints.

Our approach uses BDD-based algorithms to analyze time-invariant constraints and to derive the set of possible macros. The core algorithm is a new assignment-extraction algorithm that extracts assignments from arbitrary Boolean expressions (Section 4.1). For each variable, by extracting its assignment form, we can determine the variable’s corresponding equivalent expression, and when appropriate, globally replace the variable with its equivalent expression (Section 4.2). The strength of this algorithm is that by using BDDs, the cost of macro expansion can be better characterized since the actual model checking computation is performed using BDDs.

Note that there have been a number of research efforts on BDD-based redundant variable removal. To better compare our approach to these previous research efforts, we postpone the discussion of this prior work until Section 6, after describing our algorithms and the performance evaluation.

4.1 BDD-Based Assignment Extraction

The assignment-extraction problem can be stated as follows: given an arbitrary Boolean formula f and a variable v (where v can be non-Boolean), find \mathbf{g} and h such that

- $f = (v \in \mathbf{g}) \wedge h$,
- \mathbf{g} does not depend on v , and
- h is a Boolean formula and does not depend on v .

The expression $(v \in \mathbf{g})$ represents a non-deterministic assignment to variable v . In the case that \mathbf{g} always returns a singleton set, the assignment $(v \in \mathbf{g})$ is deterministic. A solution to this assignment-extraction problem is as follows:

$$\begin{aligned}
 h &= \exists v. f \\
 \mathbf{t} &= \bigcup_{k \in K_v} ITE(f|_{v \leftarrow k}, \{k\}, \emptyset) \\
 \mathbf{g} &= \text{restrict}(\mathbf{t}, h)
 \end{aligned} \tag{4}$$

where K_v is the set of all possible values of variable v , and *restrict* [6] is a care-space optimization algorithm that tries to reduce the BDD graph size (of \mathbf{t}) by collapsing the don’t-care space ($\neg h$). The BDD algorithm for the $\bigcup_{k \in K_v}$ operator is similar to the BDD algorithm for the existential quantification with the \forall operator replaced by the \cup operator for variable quantification. A correctness proof of this algorithm can be found in the technical-report version of this paper [17].

4.2 Macro Extraction and Expansion

In this section, we describe the elimination of state variables based on macro extraction and macro expansion. The first step is to extract macros with the algorithm shown in Figure 1. This algorithm extracts macros from the constrained space (C), which is represented as a set of conjuncts. It first uses the assignment-extraction algorithm to

extract assignment expressions (line 5). It then identifies the deterministic assignments as candidate macros (line 6). For each candidate, the algorithm tests to see if applying the macro may be beneficial (line 7). This test is based on the heuristic that if the BDD graph size of a macro is not too large and its instantiation does not cause excessive increase in other BDDs' graph sizes, then instantiating this macro may be beneficial. If the resulting right-hand-side \mathbf{g} is not a singleton set, it is kept separately (line 9). These \mathbf{g} 's are combined later (line 10) to determine if the intersection of these sets would result in a macro (lines 11-13).

```

extract_macros( $C, V$ )
  /* Extract macros for variables in  $V$  from
  the set  $C$  of conjuncts representing the constrained space */
1   $M \leftarrow \emptyset$  /* initialize the set of macros found so far */
2  for each  $v \in V$ 
3     $N \leftarrow \emptyset$  /* initialize the set of non-singletons found so far */
4    for each  $f \in C$  such that  $f$  depends on  $v$ 
5       $(\mathbf{g}, h) \leftarrow \text{assignment-extraction}(f, v)$  /*  $f = (v \in \mathbf{g}) \wedge h$  */
6      if ( $\mathbf{g}$  always returns a singleton set) /* macro found */
7        if (is-this-result-good( $\mathbf{g}$ ))
8           $M \leftarrow \{(v, \mathbf{g})\} \cup M$ 
9        else  $N \leftarrow \{\mathbf{g}\} \cup N$ 
10      $\mathbf{g}' \leftarrow \bigcap_{\mathbf{g} \in N} \mathbf{g}$ 
11     if ( $\mathbf{g}'$  always returns a singleton set) /* macro found */
12       if ((is-this-result-good( $\mathbf{g}'$ )))
13          $M \leftarrow \{(v, \mathbf{g}')\} \cup M$ 
14  return  $M$ 

```

Fig. 1. Macro-extraction algorithm. In lines 7 and 12, “is-this-result-good” uses BDD properties (such as graph sizes) to determine if the result should be kept.

After the macros are extracted, the next step is to determine the instantiation order. The main purpose of this algorithm (in Figure 2) is to remove circular dependencies. For example, if one macro defines variable v_1 to be $(v_2 \wedge v_3)$ and a second macro defines v_2 to be $(v_1 \vee v_4)$, then instantiating the first macro results in a circular definition in the second macro ($v_2 = (v_2 \wedge v_3) \vee v_4$) and thus invalidates this second macro. Similarly, the reverse is also true. To determine the set of macros to remove, the algorithm builds a dependence graph (line 1) and breaks circular dependencies based on graph sizes (lines 2-4). It then determines the ordering of the remaining macros based on the topological order (line 4) of the dependence graph.

Finally, in the topological order, each macro (v, \mathbf{g}) is instantiated in the remaining macros and in all other expressions (represented by BDDs) in the system, by substituting the variable v with its equivalent expression \mathbf{g} .

```

order_macros( $M$ )
/* Determine the instantiation order of the macros in set  $M$  */
/* first build the dependence graph  $G = (M, E)$  */
1  $E = \{(x, y) | x = (v_x, \mathbf{g}_x) \in M, y = (v_y, \mathbf{g}_y) \in M, \mathbf{g}_y \text{ depends on } v_x\}$ 
/* then remove circular dependences */
2 while there are cycles in  $G$ ,
3    $M_C \leftarrow$  set of macros that are in some cycle
4   remove the macro with largest BDD size in  $M_C$ 
5 return a topological ordering of the remaining macros in  $G$ 

```

Fig. 2. Macro-ordering algorithm.

5 Evaluation

5.1 Experimental Setup

The benchmark suite used is a collection of 58 SMV models gathered from a wide variety of sources, including the 16 models used in a BDD performance study [16]. Out of these 58 models, 37 models have no time-invariant constraints, and thus our optimizations do not get triggered and have no influence on the overall verification time. Out of the remaining 21 models, 10 very small models (< 10 seconds) are eliminated. On the remaining 11 models, our optimizations have made non-negligible performance impact on 7 models. In Figure 3, we briefly describe these 7 models. Note that some of these models are quite large, with up to 1200 state bits.

Model	# of State Bits	Description
acs	497	the altitude-control module of the NASA DS1 spacecraft
ds1-b	657	a buggy fault diagnosis model for the NASA DS1 spacecraft
ds1	657	corrected version of <i>ds1-b</i>
futurebus	174	FutureBus cache coherency protocol
nomad	1273	fault diagnosis model for an Antarctic meteorite explorer
v-gate	86	reactor-system model
xavier	100	fault diagnosis model for the Xavier robot

Fig. 3. Description of models whose performance results are affected by our optimizations.

We performed the evaluation using the Symbolic Model Verifier (SMV) model checker [10] from Carnegie Mellon University. Conjunctive partitioning was used only when it was necessary to complete the verification. In these cases (including *acs*, *nomad*, *ds1-b*, and *ds1*), the size limit for each partition was set to 10,000 BDD nodes. For the remaining cases, the transition relations were represented as monolithic BDDs. The performance of all the benchmark models were measured in the following four settings:

Base: no new optimizations except that the constrained space C is represented as a conjunction with each conjunct’s BDD graph size limited to 10,000

nodes (i.e., optimizations in Section 3 are used without the “early quantification on the constrained space C ” optimization). Without this partitioning, the BDD representation of the constrained space could not be constructed for 4 models.

Quan: same as the **Base** case with the addition of the “early quantification on the constrained space” optimization (Section 3).

SynMacro: same as the **Quan** case with the addition of a syntactic analysis that pattern matches deterministic assignment expressions ($v == f$, where v is a state variable and f is an expression) as macros and expands these macros.

BDDMacro: all the optimizations are turned on; i.e., same as the **SynMacro** case with the addition of BDD-based assignment extraction to extract macros.

The evaluation was performed on a 200MHz Pentium-Pro with 1 GB of memory running Linux. Each run was limited to 6 hours of CPU time and 900 MB of memory.

5.2 Results

Figure 4 shows the impact of our optimizations for the 7 models whose results changed by more than 10 CPU seconds and 10% from the **Base** case. For all benchmarks, the time spent by our optimizations is very small (< 5 seconds or $< 5\%$ of total time) and is included in the running time reported.

The overall impact of our optimizations is shown in the rightmost column of Figure 4. These results demonstrate that our optimizations have significantly improved the performance for 2 cases (with speedups up to 74) and have enabled the verification of 4 cases. For the *v-gates* model, the performance degradation (speedup = 0.7) is in the computation of the reachable states from the initial states. Upon further investigation, we believe that it is caused by the macro instantiation, which increases the graph size of the transition relation from 122-thousand to 476-thousand nodes. This case demonstrates that reducing the number of state variables does not always improve performance.

Model	Base sec	Quan sec	SynMacro sec	BDDMacro sec	Base / BDDMacro speedup
acs	<i>m.o.</i>	32	76	7	<i>enabled</i>
ds1-b	<i>m.o.</i>	321	138	54	<i>enabled</i>
ds1	<i>m.o.</i>	<i>m.o.</i>	<i>t.o.</i>	37	<i>enabled</i>
futurebus	1410	53	35	19	74.2
nomad	<i>m.o.</i>	<i>t.o.</i>	7801	633	<i>enabled</i>
v-gates	36	35	53	50	0.7
xavier	16	5	1	2	8.0

Fig. 4. Performance impact of each optimization. The *m.o.*'s and *t.o.*'s are the results that exceeded the 900-MB memory limit and the 6-hour time limit, respectively.

The remaining columns of Figure 4 show the impact of each optimization. The results show that by simply performing early quantification on the constraints (the **Quan** column), we have enabled the verification of *acs* and *ds1-b*, and achieved significant performance improvement on *futurebus* (speedup > 20). This is mostly due to the fact that a large number of variables can be pulled out of the transition relations and applied to conjunctive partitioning and early quantification of the time-invariant constraints (Figure 5(a)). With the addition of syntactic analysis for macro extraction (the **SynMacro** column), we are able to verify *nomad*. Finally, by adding BDD-based macro extraction (the **BDDMacro** column), we are able to verify *ds1*. The results in Figure 5(b) show that BDD-based macro extraction (**BDDMacro**) can be rather effective in reducing the number of variables, especially for the *acs*, *nomad*, *ds1-b*, and *ds1* models where > 150 additional BDD variables (i.e., > 75 state bits) are removed in comparison to using syntactic analysis (**SynMacro**).

Model	Total # of BDD Variables	CP Optimization # of BDD vars extracted		Macro Optimization # of BDD vars removed	
		image	pre-image	SynMacro	BDDMacro
acs	994	439	449	82	352
ds1-b	1314	550	566	148	492
ds1	1314	550	566	220	496
futurebus	348	58	110	12	18
nomad	2546	1121	1174	688	844
v-gates	172	0	17	16	16
xavier	200	69	86	64	116

(a)
(b)

Fig. 5. Effectiveness of each optimization. (a) Number of quantifying BDD variables that are pulled out of the transition relation for early quantification of the time-invariant constraints. These results are measured without macro optimizations. With macro optimizations, the corresponding results are basically the same as subtracting off the number of state variables removed. (b) The number of BDD variables removed by macro expansion. **Note:** the number of BDD variables is twice the number of state variables—one copy for the present state and one copy for the next state.

6 Related Work

There have been many research efforts on BDD-based redundant state-variable removal in both logic synthesis and verification. These research efforts all use the reachable state space (set of states reachable from initial states) to determine functional dependencies for Boolean variables (macro extraction). The reachable state space effectively plays the same role as a time-invariant constraint, because the verification process only needs to check specifications in the reachable state space,

Berthet et al. propose the first redundant state-variable removal algorithm in [3]. In [9], Lin and Newton describe a branch-and-bound algorithm to identify the maximum set of redundant state variables. In [12], Sentovich et al. propose new algorithms for latch removal and latch replacement in logic synthesis. There is also some work on detecting and removing redundant state variables while the reachable state space is being computed [8, 14].

From the algorithmic point of view, our approach is different from prior work in two ways. First, in determining the relationship between variables, the algorithms used to extract functional dependencies in previous work can be viewed as direct extraction of deterministic assignments to Boolean variables. In comparison, our assignment extraction algorithm is more general because it can also handle non-Boolean variables and extract non-deterministic assignments. Second, in performing the redundant state-variable removal, the approach used in the previous work would need to combine all the constraints first and then extract the macros directly from the combined result. However, for constraint-rich models, it may not be possible to combine all the constraints because the resulting BDD is too large to build. Our approach addresses this issue by first applying the assignment extraction algorithm to each constraint separately and then combining the results to determine if a macro can be extracted (see Figure 1).

Another difference is that in previous work, the goal is to remove as many variables as possible. However, we have empirically observed that in some cases, removing additional variables can result in significant performance degradation in overall verification time (slowdown over 4). To address this issue, we use simple heuristics (size of the macro and the growth in graph sizes) to choose the set of macros to expand. This simple heuristic works well in the test cases we tried. However, in order to fully evaluate the impact of different heuristics, we need to gather a larger set of constraint-rich models from a wider range of applications.

7 Conclusions and Future Work

The two optimizations we proposed are crucial in verifying this new class of constraint-rich applications. In particular, they have enabled the verification of real-world applications such as the Nomad robot and the NASA Deep Space One spacecraft.

We have shown that the BDD-based assignment-extraction algorithm is effective in identifying macros. We plan to use this algorithm to perform a more precise cone-of-influence analysis with the assignment expressions providing the exact dependence information between the variables. In general, we plan to study how BDDs can be used to further help other compile-time optimizations in symbolic model checking.

Acknowledgement

We thank Ken McMillan for discussions on the effects of macro expansion. We thank Olivier Coudert, Fabio Somenzi and reviewers for comments on this work. We are grateful to Intel Corporation for donating the machines used in this work.

References

1. BAPNA, D., ROLLINS, E., MURPHY, J., AND MAIMONE, M. The Atacama Desert trek - outcomes. In *Proc. of the 1998 International Conference on Robotics and Automation* (May 1998), pp. 597–604.
2. BERNARD, D. E., DORAIS, G. A., FRY, C., JR., E. B. G., KANEFISKY, B., KURIEN, J., MILLAR, W., MUSCETTOLA, N., NAYAK, P. P., PELL, B., RAJAN, K., ROUQUETT, N., SMITH, B., AND WILLIAMS, B. Design of the remote agent experiment for spacecraft autonomy. In *Proc. of the 1998 IEEE Aerospace Conference* (March 1998), pp. 259–281.
3. BERTHET, C., COUDERT, O., AND MADRE, J. C. New ideas on symbolic manipulations of finite state machines. In *1990 IEEE Proc. of the International Conference on Computer Design* (September 1990), pp. 224–227.
4. BRYANT, R. E. Graph-based algorithms for Boolean function manipulation. *IEEE Transactions on Computers* C-35, 8 (August 1986), 677–691.
5. BURCH, J. R., CLARKE, E. M., LONG, D. E., MCMILLAN, K. L., AND DILL, D. L. Symbolic model checking for sequential circuit verification. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 13, 4 (April 1994), 401–424.
6. COUDERT, O., AND MADRE, J. C. A unified framework for the formal verification of circuits. In *Proc. of the International Conference on Computer-Aided Design* (Feb 1990), pp. 126–129.
7. GEIST, D., AND BEER, I. Efficient model checking by automated ordering of transition relation partitions. In *Proc. of the Computer Aided Verification* (June 1994), pp. 299–310.
8. HU, A. J., AND DILL, D. L. Reducing BDD size by exploiting functional dependencies. In *Proc. of the 30th ACM/IEEE Design Automation Conference* (June 1993), pp. 266–71.
9. LIN, B., AND NEWTON, A. R. Exact redundant state registers removal based on binary decision diagrams. *IFIP Transactions A, Computer Science and Technology* A, 1 (August 1991), 277–86.
10. MCMILLAN, K. L. *Symbolic Model Checking*. Kluwer Academic Publishers, 1993.
11. RANJAN, R. K., AZIZ, A., BRAYTON, R. K., PLESSIER, B., AND PIXLEY, C. Efficient BDD algorithms for FSM synthesis and verification. Presented in the IEEE/ACM International Workshop on Logic Synthesis, May 1995.
12. SENTOVICH, E. M., AND HORIA TOMA, G. B. Latch optimization in circuits generated from high-level descriptions. In *Proc. of the International Conference on Computer-Aided Design* (November 1996), pp. 428–35.
13. SHIPLE, T. R., HOJATI, R., SANGIOVANNI-VINCENTELLI, A. L., AND BRAYTON, R. K. Heuristic minimization of BDDs using don't cares. In *Proc. of the 31st ACM/IEEE Design Automation Conference* (June 1994), pp. 225–231.
14. VAN EIJK, C. A. J., AND JESS, J. A. G. Exploiting functional dependencies in finite state machine verification. In *Proc. of European Design and Test Conference* (March 1996), pp. 266–71.
15. WILLIAMS, B. C., AND NAYAK, P. P. A model-based approach to reactive self-configuring systems. In *Proc. of the Thirteenth National Conference on Artificial Intelligence and the Eighth Innovative Applications of Artificial Intelligence Conference* (August 1996), pp. 971–978.
16. YANG, B., BRYANT, R. E., O'HALLARON, D. R., BIÈRE, A., COUDERT, O., JANSSEN, G., RANJAN, R. K., AND SOMENZI, F. A performance study of BDD-based model checking. In *Proc. of the Formal Methods on Computer-Aided Design* (November 1998), pp. 255–289.
17. YANG, B., SIMMONS, R., BRYANT, R. E., AND O'HALLARON, D. R. Optimizing symbolic model checking for constraint-rich models. Tech. Rep. CMU-CS-99-118, School of Computer Science, Carnegie Mellon University, March 1999.