## AD-A250 924

ELECT: JUN 31992
C

## Symbolic Model Checking

An approach to the state explosion problem
Kenneth L. McMillan
May, 1992
CMU-CS-92-131

Submitted to Carnegie Mellon University in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Computer Science

## (c) 1992 Kenneth L. McMillan

This research was supported in part by a scholarship grant from AT\&T Bell Laboratories and in part by a National Science Foundation grant. no. ('CR-8722633.

The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, cither expressed or implied, of the I'.S. government.


Keywords: formal verification, temporal logic, model checking, state explosion problem. symbolic model checking, binary decision diagrams. cache consistency, cache coherence, protocol verification. hardware description languages, SMV, asynchronous circuits. Petri nets. occurrence nets, process invariants.

## Carnegie Mellon

## School of Computer Science

| DOCTORAL THESIS <br> in the field of <br> Computer Science |
| :---: |
| SYMBOLIC MODEL CHECKING <br> An Approach to the State Explosion Problem <br> KENNETH McMILLAN |

Submitted in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy

ACCEPTED:


APPROVED:


Dedicated to the memory of William L. McMillan


#### Abstract

Finite state models of concurrent systems grow exponentially as the number of components of the system increases. This is known widely as the state explosion problem in automatic verification, and has limited finite state verification methods to small systems. To avoid this problem. a method called symbolic model checking is proposed and studied. This method avoids building a state graph by using Boolean formulas to represent sets and relations. A variety of properties characterized by least and greatest fixed points can be verified purely by manipulations of these formulas using Ordered Binary Decision Diagrams.

Theoretically, a structural class of sequential circuits is demonstrated whose transition relations can be represented by polynomial space OBDDs. though the number of states is exponential. This result is born out by experimental results on example circuits and systems. The most complex of these is the cache consistency protocol of a commercial distributed multiprocessor. The symbolic model checking technique revealed subtle errors in this protocol, resulting from complex execution sequences that would occur with very low probability in random simulation runs.

In order to model the cache protocol, a language was developed for describing sequential circuits and protocols at, various levels of abstraction. This language has a synchronous dataflow semantics, but allows nondeterminism and supports interleaving processes with shared variables. A system called SMV can automatically verify programs in this language with respect to temporal logic formulas, using the symbolic model checking technique.

A technique for proving properties of inductively generated classes of tinite state systems is also developed. The proof is checked zutomatically, but requires a user supplied process called a process invariant to act as an inductive hypothesis. An invariant is developed for the distributed cache protocol. allowing properties of syatems with an arbitrary number of processors to be proved.

Finally, an alternative method is developer for avoiding the state explosion in the case of asynchronous control circuits. This technique is based the unfolding of Petri nets. and is used to check for hazards in a distributed mutual exclusion cirout.


## Acknowledgments

To the extent that the work in this thesis is worthy of credit. Ed Clarke deserves a large portion of that credit, since his work is the foundation for almost everything in it. Ed got me started in the field by teaching me what I know about verification. and opening doors for me in the research community. His enthusiasm for my early, somewhat untutored tentatives in the field gave me the confidence to pursue my own ideas. Another important influence on this work and its author is Bob Kurshan. In more than once instance, his insistence on solving a particular problem led to a general solution in an unexpected way. What I now know about the process of research I learned from Ed Clarke and Bob Kurshan.

The work in this thesis also rests rather heavily on that of Randy Bryant. To a certain extent. the symbolic model checking technique resulted from the fortunate coincidence of my being at the same institution as Randy. Finding an application for this technique was also a more or less serendipitous occurrence. and for this I have to thank the people at Encore Computer Corporation, including Dyung Van Le. Jim Schwalbe and Drew Wilson, for taking an interest in formal verification. and generously allowing me to use their system as an application (and to publish it no less!).

I must also thank David Long and Jerry Burch, who contributed substantially to the ideas in this thesis. I owe special thanks to David Long, who is a font of information, and who helped me to solidify ideas in countless discussions. when no doubt he had more important things to do.

Thanks to Robert Brayton and the ('. ID group at the I'niversity of C'alifornia at Berkeley. who gave this thesis a very careful and intelligent reading, and to . Mlan Fisher. It was always a pleasure to have discussions with Allan. and no doubt some of his insights can be found interspersed in these pages. Of course. I am indebted to all the good people who make the (arnegie Mellon School of Computer Science and its facilities work. ('MI' provides a unique environment for graduate students, and I consider myself privileged to have been a part of it.

Finally, to Tracy Slein. the one indispensible person in the whole process - I can't thank sou emongh.

## Contents

1 Introduction ..... 11
1.1 Background ..... 12
1.1.1 Temporal logic ..... 12
1.1.2 Automata theoretic models ..... 14
1.2 Scope of the thesis ..... 16
1.3 Related research ..... 18
1.3.1 Reduction ..... 18
1.3.2 Induction ..... 20
1.3.3 Other symbolic methods ..... 20
2 Symbolic model checking ..... 23
2.1 Temporal logic ..... $-5$
2.1.1 Linear time ..... 26
2.1.2 Discrete time ..... 27
2.1.3 Branching time ..... 27
2.2 The temporal logic (TL ..... 29
2.2.1 Syntax and semantics of ("TL ..... ?!
2.2 .2 Fixed point characterization of (TL ..... 30
2.3 Symbolic ('TL model checking ..... 35
2.3.1 Quantified Boolean formulas ..... 35
2.3.2 Representing sets and relations ..... 36
2.3.3 CTL formulas ..... 38
2.3.t Binary Decision Diagrams ..... 10
2.f Examples ..... 1!
2.4.1 Synchronons state machines ..... $1!$
2.4.2 Asynchronous state machines ..... it
2.) (iraph width and OBDD ..... 62
-..5. 1 Bounded width circuits ..... 63
-.5.2 Bounded tree-width circuits ..... 66
2.6 Mu-('alculus model checking ..... 71
-.6.1 Applications of the Mu-C'alculus ..... 73
-.6.2 Symbolic algorithm ..... 77
2.7 Computing equivalence relations ..... 78
2.i. 1 State equivalence ..... 80
-. 7.1 Methods for functional composition ..... 83
2.i.3 Experimental results ..... 3.5
2. 3 Related research ..... 38
3 The SMV system ..... 97
:3.1 An informal introduction ..... 98
3.2 The input language ..... 108
3.2.1 Lexical conventions ..... 108
3.2.2 Expressions ..... 108
3.2.3 Declarations ..... 110
3.2.4 Modules ..... 114
3.2.5 Identifiers ..... 116
3.2.6 Processes ..... 117
3.2.7 Programs ..... 117
3.3 Formal semantics ..... 118
3.3.1 The model ..... 118
3.3.2 Expressions ..... 119
3.3.3 . Issignments and definitions ..... 120
3.3.1 Variable declarations ..... 121
3.3.5 Renaming ..... 121
3.3.6 Parallel composition ..... 122
3.3.7 Instantiation ..... $12 \cdot 2$
3.3.5 Specifications ..... 123
4 A Distributed Cache Protocol ..... 125
1.1 The Protocol ..... 127
1.1.1 Processors ..... 128
1.1.2 The local DIC interface ..... 13:3
1.I.3 The Remote I'IC interface ..... 137
1.1.1 Protorol example ..... 138
4.2 Verifying the protocol ..... $1+1$
4.2.1 Freedom from deadlock ..... 141
4.2.2 Correcting the deadlock ..... 145
4.2.3 Sequential consistency ..... 147
t.2.t Correctness of diagnostics ..... 150
4.3 Evaluation ..... 1.52
5 Induction and model checking ..... 155
5.1 The general framework ..... 156
j.).1 Induction in other models ..... 1.59
3. 2 Induction and SMV ..... 1.59
5.2.1 Proving compositionality ..... 1:59
5.2.2 Computing simulation relations ..... 163
5.2.3 Induction and S.MV. 1 ..... 164
5. 3 Example: The (iigamax protocol ..... 165
5. 4 Related research ..... 168
6 A partial order approach ..... 171
6.1 The unfolding operation ..... 172
6.2 Application example ..... 182
6.3 Deadlock and occurrence nets ..... 185
6.t Relation to AI techniques ..... 189
6.5 Evaluation ..... 190
7 Conclusion ..... 191
A Semantics of SMV. 1 ..... 195

1. 1 The model ..... 195
1.2 Expressions ..... 196
1.3 . Issignments and defintions ..... 197
A. $\frac{1}{\text { Variable derlarations }}$ ..... 198
.1. 5 Renaming ..... 198
2. 6 Parallel composition ..... 199
A. $\bar{i}$ Instantiation ..... 199
1.8 Programs and interleaving ..... 200
...) Specifications ..... 201

## Chapter 1

## Introduction

There are several practical reasons for applying formal verification methods to computer systems. The most obvious is the high cost of correcting errors in digital designs. This cost has been increasing with the rising level of integration in digital circuit technology. It can be decreased to an extent in application specific designs by the use of programmable device technologies, but at least for the present, programmable logic has distinct disadvantages in performance and area. Thus, there is a growing demand for design methodologies that can yield correct designs on the first fabrication run. Design errors that are discovered before fabrication can also be quite costly. however. in terms of the engineering effort required to correct the error. and the resulting impact on derrlopment schedules. At present. the best tools available to engineers for finding errors before fabrication are simulators, which model the behavior of a system for predetermined or random input patterns. The engineer using simulation is faced with two ill-characterized and increasingly intractable problems. The first is creating a set of input patterus that are sulficient to expose any incorrect behavior of the system. and the second is determining the correct output of the system under these conditions. to be compared with the simulated output. Inreased density of integration has allowed higher level functions such as network protocols to be implemented in hardware. and as a result, the problems of simulation have become critical. What seems to be needed is a precise vet understandable wav of specifying correct behavior. and an exhaustive methorl of determining that the system model satisties
this specification for all input patterns. This is the meaning of formal verification.

A formal verification framework has three basic elements - a mathematical model of the system to be verified. a formal language $\therefore$ : framing the correctness problem, and a methodology for proving the statement of correctness. One characteristic that many automatic verification methodologies have in common is that they require an exhaustive search of the state space of the model. Owing to simple combinatorics. the size of this state space can be. and usually is. exponential in the size of the system being modeled. This exponential growth in the state space. known as the state explosion problem is the limiting factor in applying automatic verification methodologies to large systems.

This thesis is directed toward solutions for the state explosion problem. This is essentially a question of methodology. but before we can discuss methodology. we need to discuss some of the models and formalisms that are commonly used in formal verification of hardware.

### 1.1 Background

The problem of hardware verification is in some ways similar to. and in other ways different from the problem of proving correctness of programs. Digital systems are most similar to what Pnueli has characterized as reactive programs [Pnu86], in that they receive input and produce output in a continuous interaction with their enviroument. rather than computing a single result and halting. In additi $n$, the behavior of digital systems is concurrent in the extreme. sincr every gate in the system is simultaneously evaluating its output as a function of its inputs.

### 1.1.1 Temporal logic

For reasoning about concurrent. reactive programs. Puueli proposed the use of a formal system originally studied by philosophers. called temporal logic [Pnuī. Pnu86. MP81. Kro87]. In a temporal logic, the usual operators of propositional logic are augmented by to nist operators. which are nsed to form assertions about changes in time. ()ne can
assert. for example, that if proposition $p$ holds in the present, then proposition $q$ holds at some instant in the future, or at some instant in the past. The temporal modalities can be combined to express fairly complex statements about past. present and future. For example "if $p$ holds in the present. then at some instant in the future, $p$ will have held in the past." A temporal system provides a complete set of axioms and inference rules for proving all validities in the logic for a given model of time. such as partially ordered time. linearly ordered time, dense time. and even branching time.

Temporal logic is powerful enough to define a semantics for programs which captures not only the traditional before and after conditions of Floyd-Hoare style program proving, but also a wide variety of temporal properties of programs. such as termination. possible termination, termination under fair scheduling of concurrent processes. etc. [CE8la. BAMP8I]. In the hardware area. Malachi and Owicki used temporal logic to give a concise specification of the conditions necessary for an asynchronous circuit to be speed independent [MO81]. Bochmann used temporal logic to give a semantics for self timed circuits, and used this system to verify a corrected version of an arbiter circuit [Sei80a]. Formal proofs of this kind are extremely tedious and difficult, however. and computationally intractable to automate. To simplify the hand proof. Bochmann used a somewhat oversimplified semantics for the circuit elements (neglecting gate delay) and as a result. missed a bug in the design. which was demonstrated by Dill [DCS6].

A more practical application of temporal logic in hardware verification. called mordel cheching, was introduced by ('larke and Emerson [CE81b] and independently by Quielle and Sifakis [QS81]. Instead of proving the validity of a logical formula for all morels. a model checker determines the truth value of the formula in a specific finite model. For branching time logic. the model checking problem is romputationally tractable, even though the validity problem is intractable. Here an important distinction between hardware and software sustems comes into play - hardware systems are finite-state. This allows the proof procedure to be automated insing model checking, while maintaining the formal elegance of temporal logic for specifying correct behavior.

The methorl of 'larke and Emerson first builds a complete state graph of the ssstem from a lescription in an appropriate langmage.

The truth value of a formula in the logic is determined by an algorithm which propagates formulas in this state graph until a fixed point is reached. Besides being fast and fully automatic, this technique has the advantage that it can produce state sequences as counterexamples when the formula being checked is false. This has made it possible to find bugs in a number of small but fairly subtle circuit designs [BCDM86, BCD86]. including the one verified by Bochmann.

For linear time logic, there is a decision procedure that translates a formula into an automaton by means of a tableau construction [RU71, CE81b. BAMP81]. This construction is similar the the semantic tableaux method of constructing proofs in standard logic [Smu68]. Each state in the tableau is associated with a set of formulas which are true in that state. Since the number of states in the tableau is exponential in the size of the formula. the method is not practical for proofs about very large systems. However, the tableau method can be used in a model checking framework, yielding an algorithm which is exponential in the size of the formula but linear in the size of the model [LP85].

### 1.1.2 Automata theoretic models

An alternative to the temporal logic framework is to cast the correctness problem in terms of a relation between the external or observable behaviors of two processes. One way to define this relation is by considering the set of all possible sequences of communications between processes. For example. in the $L$-automata model of Kurshan [Kur86], these sequences are defined by the language of an $\omega$-automaton. Correctness is framed as the containment of the language of one automaton in the language of another. This asymmetric relation makes it possible to "underspecify" a system. that is. to leave some choices open to the designer. The use of automata on infinite strings makes it possible to express liveness properties. For instance. one can easily construct an automaton whose language is the set of all infinite strings such that every time a message is sent on some channel. one is eventually received. Language containment between $\omega$-automata can be established by an algorithm which searches for cycles in the state space of a product automaton.

Van de Snepscheut [riS8:3] and Dill [Dil88] have used trace the-
ory to model speed independent circuits [Sei80b]. A trace is simply a history of the communications between a process and its environment. The trace sets of two process can be combined in a way which models communication between the two processes by synchronizing signals sent and received on the same channel. Dill's system is a circuit algebra which has both a structural interpretation (describing the physical connection of wires) and a trace theoretic interpretation (describing the communications along those wires). The actual trace sets are defined by the languages of finite automata (in this case. automata on finite strings, hence liveness cannot be modeled). A relationship called conformance between two processes determines when one process can safely be substituted for the other in all environments. Conformance can be tested by a polynomial algorithm which searches the state space of a finite automaton derived from the two processes.

In the Calculus of ('ommunicating Systems (C'CS) [Mil80]. Milner takes a different approach in which external behavior is modeled by a tree rather than a set of sequences. The way CCS models communication is not well suited to modeling hardware. since in CCS a signal cannot be sent until a receiver is ready to receive it. In hardware. a receiver cannot generally prevent a signal from being sent. Also, in C'C'S. communication is always between two processes, while in hardware signals are often broadcast to many receivers. A calculus specialized to circuits called CIRCAL [Mil83] was developed to remedy these problems. The notion of correctness in process calculi is called observational equivalence, meaning that an observer cannot ristinguish between two processes by any experiment. This notion of correctness is extremely strict. since it doesn't allow the specifier to leave any choice up to the designer regarding the externally visible behaviors. Observational equivalence can be proved by establishing a relation called bisimulation between the two processes. For finite state processes, there is an polynomial time algorithm for bisimulation which is very similar to the coarsest partitioning algorithms used for state machine minimization [NH84].

All of these methods can be viewed as variations on the theory of finite automata. tailored for modeling certain properties of a particular class of systems. In fact. the automata theoretic approach is not very far from the temporal logic approash in practice. The difference
is mostly a question of notations, since the tableau method provides a way of translating a temporal logic formula into an automaton. Although temporal logic is not as expressive as automata in characterizing classes of sequences. it has been shown by Wolper that temporal logic can be extended using right linear grammars to make it as expressive as automata without increasing the complexity of the decision procedure [Wols:3]. (larke and Kurshan have also proposed a branching time logic in which the temporal operators are defined by finite w-antomata [C(GKs)].

What all of the above systems have in common is that correctness. once formalized, can be determined by an algorithm that searches the entire state space of a finite state model. Such methods have the advantage of being fully antomatic, but invariably suffer from the state explosion problem.

### 1.2 Scope of the thesis

This thesis explores methods of state space search that avoid the state explosion problem by not explicitly representing the states of the model. To do this. some revolutionary new techniques are borrowed from the area of switching function analysis. In this domain. a combinational explosion also arises. since the number of input combinations to a Boolean function is exponential in the number of inputs. . New techniques for Boolean comparison avoid this problem by representing Boolean functions with a reduced form of decision graph called an Ordered Binary Decision Diagram (OBDD) [Brys6]. These decision graphs provide a compact canonical form for Boolean finctions. To apply this idea to temporal verification. we observe that if the state of a sestem is representerl by a rector of Boolean variables. then a set of states can be represented by a Boolean function which returns true for all states in the set. Similarly. a relation $r$ Ry between states can be represented by a Boolean function of two sets of variables. one set representing $x$ and the wher representing $y$. In this way a model checking algorithm ran be developed which uses OBDDs to represent sets and relations. Borrowing teminology from Bryant. this technique is called aymbolic monel cheching. since symbotic variables are used to represent the com-
ponents of the system state rather than numeric values. Using symbolic model checking, we can can automatically verify some regularly structured systems with literally astronomical numbers of states.

The principle contributions of this work are detailed below:
The symbolic model checking method. A technique is developed for state space search using Ordered Binary Decision Diagrams. We show that any algorithm that can be expressed in a fixed point logic called the Mu-Calculus can be computed using this method. These include algorithms for all of the correctness notions enumerated above. including CTL model checking ( with fairness constraints), the linear time tableau method. conformance. observational equivalence. language containment for $\omega$-automata, Mealy machine equivalence, and others. From a theoretical point of view. a structural class of sequential circuits is identified whose transition relations can be represented by a polynomially bounded OBDDs. This theoretical result is born out by experiments on classes of regularly structured circuits. for which the time used by the symbolic model checking method is found to be polynomially bounded in the circuit size. In addition. some experiments are reported using symbolic model checking to compute the equivalence relation between states of a finite state machine. Several techniques are advanced which improve the efficiency of this computation in practice.

The SMV system. A symbolic model checking system called SMV is presented. This system permits the automatic verification of programs written in a specialized language for describing concurrent finite state systems and protocols. This language is somewhat similar to LISTRE [CHPPS7] in its synchronous dataflow semantics. but has several unique aspects. For example. it allows systems to be morleled nondeterministically for purposes of abstraction. it allows arbitrary interleaving of concurrent processes, and it allows programs to be annotated with assertions in branching time temporal logic.

Formal verification of the Encore Gigamax cache consistency protocol. The cache consistency protocol of a distributed shared-memory multiprocessor called the Encore (iigamax is modeled in the SWV language and verified using the symbolic model checker. Running in minutes. the symbolic model checker discovered errors in this system which were not discovered by simulation. in spite of the very large state space of the morlel [MS'91]. This experiment shows that the model checking
technique can be used effectively in an industrial setting for highly complex systems. It also sheds light on issues involved in modeling such protocols as finite state systems, and the kinds of errors that can be found by model checking that are not likely to be found by simulation.

Induction over processes. A partially automated method of induction is described for proving properties of parameterized classes of designs. The method applies to a variety of process models. requiring of the model only rertain simple algebraic properties. The SMV' system is extended to support proof by induction. allowing some properties of the Gigamax cache protocol to be verified for configurations of arbitrary size.

Verification using occurrence nets. An alternative method for avoiding the state explosion is examined. This technique avoids considering all of the possible interleavings of concurrent actions by using a partially ordered representation of behavior called an occurrence net [.NPW81]. This method is used to verify that a design for an asynchronous distributed mutual exclusion circuit is hazard free (this example is also used for the symbolic model checking method). Using this technique, we also find empirically that the run time is polynomial in the number of components of the system, while the number of states is exponential.

### 1.3 Related research

Since the state explosion problem is ubiquitous in the verification of computer systems and protocols. many researchers in the area have studied it.

### 1.3.1 Reduction

The most common approach is based on reduction - reducing the correctness problem to a similar problem in a smaller state space. This is generally done by replacing processes in the model by smaller processes that have similar or identical communication behavior. The most general framework for this kind of reduction is that of Limshan [Kurst].
 to simplify mot only the internal state of a process, inn also its external
communications. In this methodology. one generally builds a hierarchy of reductions, in which processes at the lowest level are reduced, then combined at the next higher level and further reduced, etc. Kurshan advocates building this hierarchy from the top clown. so that the most abstract models can be verified before details are filled in at the next lower level.

I hierarchical approach was also taken by Dill in his trace theoretic system for speed independent circuits [Dil88]. In this case, the reduction is obtained mostly by hiding internal signals of a module. There is no provision for abstracting the signals by which the module communicates with its environment. That is. communication always remains at the same level. that of digital signal transitions.

The reduction approach is generally not automatic. Usually, the reduced process is obtained in an ad hoc manner. and the validity of the reduction is then tested automatically. Some methods have been proposed for obtaining reduced processes automatically, however. For example. in a method called compositional model checking, a state minimization procedure is used to obtain a reduced process that is equivalent to the original process with respect to observation via its inputs and outputs [CLM89b. CLM89a]. This reduction preserves the truth value of all formulas in a suitable logic. Graf and Steffen have also studied minimization with respect to a suitable notion of equivalence as a reduction technique [GS91]. Minimization techniques are fairly strict in terms of the required relation between the original and reduced processes. however. As a result. the reduction that can be obtained using these techniques is not generally as great can be obtained using more Hexible but unautomated methods.

The symbolic model checking technique is not really an alternative to reduction methods. but is complementary to them. In general. the larger the state space that can be searched automatically, the less the need for reduction. For example. Dill used a reduction (constructed by hand) to verify a spepd independent distributed mutual exclusion ring rircuit [Dil88]. I sing symbolic model checking, there is no need for a reduction - the verification time is polynomial in the size of the ring (cf. chapter 2). On the other hand. symbolic model checking techniques can be used to implement the validity test for reductions (df. chapter j). hence the two techniques can be combined.

### 1.3.2 Induction

In systems of many identical processes, it is sometimes possible to reduce an arbitrary number of processes to a single process while retaining certain properties of interest. For example, Browne. Clarke and Grumberg proposed a reduction technique of this sort which preserves the truth value of formulas in a suitably restricted logic with process quantifiers [BCG86]. Unfortunately, the reduction, a form of bisimulation, had to be established by hand. There was no automated way of checking it. Kurshan and McMillan proposed an inductive method of establishing the reduction that could be checked automatically [K.M89]. The method is also less restrictive in terms of the properties that can be proved since it does not rely on bisimulation. This method is used in chapter 5. A similar method was described independently by Wolper and Lovinfosse [WL89]. Another inductive technique has been described by Shtadler and Grumberg [SG89]. This technique is somewhat more flexible in that it treats networks generated by context free grammars. but is limited to bisimulation as a reduction technique. A more detailed comparison of these methods can be found in chapter $\mathbf{5}$.

### 1.3.3 Other symbolic methods

Coudert and Madre have described a method for verifying finite state machines using Ordered Binary Decision Diagrams which is similar to symbolic model checking [CBM89]. The symbolic model checking technique was developed in 1987. The technique of Coudert and Madre appears to have been developed two vears later [Cou91] but independently. There are several differences of approach between the two methods. Symbolic model checking is directed mostly toward proving temporal properties of finite state systems. whereas Coudert and Madre have concentrated mostly on proving equivalence of deterministic Mealy machines (though they also discuss temporal logic [CMB91]). Testing Mealy machine equivalence is useful. for example. when one is mapping a design from one technology to another, but is a tairly limited form of verification, since the specification is at the same level of detail as the implementation. Nso. in this work. we consider the performance of algorithms mostly in terms of asymptotic behavior for
regularly structured classes of systems, while Coudert and Madre have considered mostly a set of benchmark circuits for synthesis. This makes it difficult to determine how well their technique scales with circuit size. Finally, Coudert and Madre have not as yet reported any results for testing equivalence of two different implementations of the same finite state machine. In practice. they have only used symbolic techniques to generate the set of reachable states of a finite state machine. This information is useful for test generation and sequential synthesis [TSL ${ }^{+} 90$ ]. but these experiments provide no information about how well the technique works for verification. On the other hand. the svmbolic model checking technique has been applied to the verification of an industrial design for a distributed cache consistency protocol (cf. chapter 4). . more detailed description of the work of Coudert and Madre and others using OBDDs for sequential circuit verification. can be found in chapter 2 .

## Chapter 2

## Symbolic model checking

As mentioned in the introduction. a formal verification system has several basic elements. First. we require a model. A model is an imaginary universe, or more generally, a class of possible imaginary universes. To make our model meaningful, we require a theory that predicts some or all of the possible observations that might be made of the model. An observation generally takes the form of the truth or falsehood of a predicate, or statement about the model. Finally, to verify something meaningful about the model, we require a methodology for proving statements that are true in the theory.

In program proving, the universe is a totally imaginary one. driven by mechanisms (the compiler and hardware) of which the programmer has no knowledge. The logician is free to assign any semantics at all to programs. provided the compiler writer and hardware designer agree to implement them. This makes program proving an artificial science. in the sense that our theory is true because we say it is. In contrast. a hardware verification system requires a model of a real physical system. The underlying physical mechanism is still invisible to us (we can only postulate its existence). but we can empirically construct a model which predicts the necessary observations with a sufficient degree of accuracy for our purposes (the verification of digital circuits). It turns out that the required degree of accuracy is not very large. Though guite accurate models are possible. using partial differential equations to describe the time evolution of tields and particle densities, a suitable design style makes it possible to consider only the digital (one or zero) value of
voltages. ignoring entirely the exact voltage within the digital ranges, and the time it takes to switch from one range to another. Depending on the design style (eg.. synchronous or self timed), different models may be appropriate. In certain rare cases, we may have to use differential equations to model the analog behavior of circuits (for example, when metastability arises). In this thesis. though, we will consider only fairly abstract models of circuits as finite state machines. Thus, we return to the science of the artificial. wherein we choose the theory to suit our needs. but with the understanding that a method exists for translating our models into real systems.

The kind of theory that emerges for the model depends to a large extent on the kind of experiments the observer is able to perform. For example. in traditional program proving systems. the observer is allowed to set up the initial state of the program, wait for the program to terminate, and then examine the final state. The theory of this model can be expressed in a kind of before-and-after logic whose axioms determine the semantics of programs. For example. in Floyd-Hoare logic [Hoa69], the formula

$$
\{\text { true }\} x:=y\{x=y\}
$$

is an axiom: for any initial condition, after the program $x:=y$ terminates. $x$ and $y$ have the same value. The fact that no other variables change value in the process can also be expressed as an axiom:

$$
\{z=a\} x:=!\{z=a\}
$$

provided neither $z$ nor $a$ depend on $x$.
In this system. if the program fails to terminate (diverges), the observer must simply wait forever. it.. no observation is possible. One might ask whether waiting forever is not itself an observation. that is. should it not be possible to state in the semantics that a given program terminates or doesn t terminate for a given initial condition? This point can be argued either way for programs (since knowing that a program terminates before infinity is not very practical information). However. for digital systems (or reactive systems in general). it is clear that simple before and after conditions are not a sufficient theory; first of all termination for these systems is not well defined. and moreover the meaning of what these systems are supposed to do is inseparably
linked with the evolution of events in time [Pnu77]. ${ }^{1}$ What we need is a formal theory in which we can reason about temporal aspects of a system's behavior.

### 2.1 Temporal logic

Temporal logic (or tense logic) is a system devised by philosophers expressly for making statements about changes in time [Bur84]. In temporal logic. the formula $F q$ is true in the present if $q$ is true at some moment in the future. Similarly $P q$ is true in the present if $q$ is true at some moment in the past. These tense operators, $F$ and $P$, have duals which are generally given their own names. The formula $C_{i} q$ is equivalent to $\neg F \neg q$, meaning that $q$ is true at every moment in the future. The formula $H q$ is equivalent to $\neg P \neg q$, meaning that $q$ is true at every moment in the past. These operators can give surprisingly concise expressions of sentences with complex tense structures. For example. $q \Rightarrow F P_{q}$ can be interpreted as "if $q$ holds in the present, then at some time in the future $q$ will have held in the past".

The usual model theoretic semantics given to temporal logic (and other modal logics) is the so-called possible worlds semantics. A frame in this semantics consists of a class $S$ of states through which the system evolves. and a relation < representing temporal order. A model is a frame with a valuation $L$. which assigns truth or falsehood to every atomic proposition (propositional letter) in every state. ${ }^{2}$ The truth or falsehood of temporal formulas is relative to the present state. For example. the formula $F q$ is true in state $s$ iff there exists a state $t$ such that $p$ is true in state $t$ and $s<t$. Similarly, $P_{q}$ is true in state siff there exists a state $t$ such that $p$ is true in state $t$ and $t<s$. Votice that a temporal formula acts like an open sentence. with one free parameter representing the present state. Thus it defines a class of states in which the formula is true. Similarly a state defines a class of formulas which are true in that state.

[^0]The choice of axioms in the logic can be used to characterize the temporal ordering relation $<$. For example, the folluwing axioms (in addition to the propositional tautologies) exactly characterize those frames whose < relation is a partial order (transitive and antisymmetric) [Burst]:

$$
\begin{gather*}
G(p \Rightarrow q) \Rightarrow(G p \Rightarrow G q)  \tag{2.1}\\
H(p \Rightarrow q) \Rightarrow(H p \Rightarrow H q)  \tag{2}\\
p \Rightarrow G P p  \tag{2.3}\\
p \Rightarrow H F p \tag{2.4}
\end{gather*}
$$

One inference rule (in addition to modus ponens) is required: by temporal generalization. if $\alpha$ is provable. we infer that $G \alpha$ and $H \alpha$ (that is, a tautology must hold true at all times. or perhaps, the rules of sound inference (lo not change with time). By specializing this system slightly, we can obtain logics characterizing a variety of models of time. including linear time, discrete time, and branching (non-deterministic) time. All of these results can be found in [Bur84].

### 2.1.1 Linear time

We usually think of time as a linearly ordered set, measuring it either with the real numbers or the natural numbers. A frame is linearly ordered if. in addition to being partialiy ordered. it is total. ie.. for all states s.t. either $s<t$. $s=t$. or $t<s$. The temporal frames in which $<$ is a linear order can be characterized by simply adding the following two axioms to the basic set (they are time reversal duals):

$$
\begin{align*}
& \left(F P_{q}\right) \Rightarrow\left(P_{q} \vee q \vee F_{q}\right)  \tag{2.5}\\
& \left(P F_{q}\right) \Rightarrow\left(P_{q} \vee q \vee F q\right) \tag{2}
\end{align*}
$$

Linear temporal logic is usually extended by the until operator and the sincf operator. Informally. $p l^{\prime} q$ states that $p$ will hold at some moment in the finture. until which time $q$ will hold at all moments. Similarly. $p$ i states that $p$ held at some moment in the past. since which time $q$ has held at all moments. More precisely, $p I^{\prime} q$ is true in state if if there is some state $t$ such that $s<t$ and $q$ is tme in state $t$. and for all $\because<u<t, p$ is true in state $u$.

## ?. . 2 Discrete time

It is common in engineering to model time as a discrete sequence (measured by the integers). Discrete dynamics are commonly used. for example, in signal processing and synchronous digital systems. A discrete frame is one in which every state has an immediate successor and an immediate predecessor. The linear discrete frames can be characterized by adding the following two axioms to those for linear time logic:

$$
\begin{align*}
& p \wedge H p \Rightarrow F H p  \tag{2.7}\\
& p \wedge G p \Rightarrow P G p \tag{2.8}
\end{align*}
$$

It is useful in a discrete linear temporal logic to define a next time temporal operator. The formula $\overline{X q}$ is true in state $s$ when there is an immediate successor of $x$ in which $q$ is true. A state $t$ is an immediate successor of $s$ if $s<t$ and there does not exist a state $u$ such that $s<u<t$. Thus. $X q$ is exactly equivalent to false $U q$, so its addition does not increase the expressiveness of the logic.

### 2.1.3 Branching time

I branching frame is one in which the temporal order $<$ defines a tree which branches toward the future. Thus, every instant has a unique past. but an indeterminate future. This is an inherently nondeterministic model of time. and hence is well suited. for example, for defining a semantics of non-deterministic programs. A frame is tree ordered when for all states s.t. $u$. if $t<s$ and $u<s$ then $t<u$. $t=u$ or $t>u$. In other words. the past of every state is linearly ordered. The tree ordered frames can be characterized by simply dropping (2.6) from the axioms of linear time logic.

Though pure tense logic can exactly characterize the branching time frames. it leaves something to be desired in expressing properties of non-deterministic programs. For example. it is common in defining the semantics of these programs to say that a program aborts iff it must inevitably abort. This finctionality ran be implemented by backtracking. Similarly, a non-deterministic Turing machine terminates if it may possibly terminate. These notions of inevitability and possibility are
not represented in an ordinary tense logic. They can be incorporated, however. by combining notions from temporal logic and modal logic.

We would like to interpret the branching structure of time as meaning that each instant of time has many possible futures, and that as time evolves from present to future, these possibilities are reduced. Thus, in the past. there existed possible futures which are now precluded. This interpretation gives rise to notions of necessity (inevitability) and possibility in tense logic [Tho84]. We think of the truth or falsehood of tense formulas as being relative to a given branch of the tree ordered frame (one possible evolution of time into the future). A branch is defined as a maximal linearly ordered set of states. We will write q[ $s, b]$ if $q$ holds in state $s$ in branch $b$. Thus. $F q[s . b]$ iff there exists a state $t$ in $b$ such that $s<t$ and $q[t, b]$. Similarly, $P q[s . b]$ iff there exists a state $t$ in $b$ such that $t<s$ and $q[t . b]$. The notion that $q$ is necessarily true is represented by the formula $A q$. We will say $A q[s, b]$ iff for all branches $b^{\prime}$ containing $s, q\left[s, b^{\prime}\right]$. The notion that $q$ is possibly true is represented by the formula $E q$. We will say $E q[s, b]$ iff for some branch $b^{\prime}$ containing $s, q\left[s, b^{\prime}\right]$. Notice that $A$ and $E$ provide a kind of second order quantification over maximal linearly ordered subsets. ${ }^{3}$

According to this semantics for modal branching time logic, there may be possibilities in the past that are foreclosed in the present. For example. $q \Rightarrow H A F q$ is not valid. The fact of $q$ in the present does not imply the necessity of $q$ in the past. Thus, modal branching time logic might be termed the logic of regret. The logic can also express useful semantic properties of non-deterministic programs [BAMPS1]. For example. if \& represents the fact of a program terminating, then inevitable termination is expressed by the formula $A F_{q}$ (necessarily in the future $q$ ). Possible termination is expressed by $E F q$ (possibly in the future q). If the proposition $p$ represents a correct output of the program. then (inevitable) partial correctness is expressed by the formula $A G(q \Rightarrow p)$ (necessarily invariantly, termination implies correctness). The somewhat odd but definable notion of possible partial correctness is expressed by $E(i(q \Rightarrow p)$. Note that $P q$. A $P q$ and $E P q$ are all $\log$ ically equivalent. since the past of a state is the same for any branch.

[^1]Also note that $A$ and $E$ are dual. since $A q$ is equivalent to $\neg E \neg q$.

### 2.2 The temporal logic CTL

The temporal logic CTL is a subset of modal branching time logic defined by Clarke and Emerson [CE81b]. The acronym stands for Computation Tree Logic. ${ }^{4}$ In ('TL. temporal operators occur only in pairs consisting of $A$ or $E$. followed by $F . G . C^{\prime}$ or $X$. Thus, past time operators are not allowed. and tense operators cannot be combined directly with the propositional connectives.

### 2.2.1 Syntax and semantics of CTL

The syntax of CTL formulas is given as follows:

1. Every atomic proposition is a CTL formula.
2. If $f$ and $g$ are C.TL formulas, then so are

$$
\neg f,(f \wedge g), A Y f, E X f, A(f U g), E(f(f g)
$$

The remaining operators are viewed as being derived from these according to the following rules:

$$
\begin{aligned}
f \vee g & =-(\neg f \wedge \neg g) \\
A F g & =A\left(\text { true } l^{\prime} g\right) \\
E F g & =E\left(\text { true } l^{\prime} g\right) \\
H(f f & =-E(\text { true } C \neg f) \\
E(; f & =-1(\text { true } C \neg f)
\end{aligned}
$$

The truth or falsehood of formulas is defined with respect to a Kripke model. but in a slightly non-standard way. For ('TL. the model is a triple ( $S . R . L$ ). where $S$ is the set of states. $R$ is the transition relation and $L$ is the valuation. The trausition relation is the set of

[^2]all pairs ( $s, t$ ) such that $t$ is an imediate successor of $s$. A branching model (a.k.a. computation tree) can be obtained by starting at a designated state $s$ and unwinding the graph $(S, R)$ into an infinite tree (provided every state has at least one successor). The semantics for (TL given below is equivalent to the standard semantics with respect to this infinite tree. ${ }^{5}$

A path of a model $K=(S, R, L)$ is an infinite sequence of states $\left(s_{0}, s_{1}, s_{2} \ldots\right) \in S^{\omega}$ such that each successive pair of states $\left(s_{i}, s_{i+1}\right)$ is an element of $R$. Every path is maximal linearly ordered subset of the tree structure unwound from $s_{0}$.

The notation $K . s \models f$ means that the formula $f$ is true in state $s$ of Kiripke model $K$. In the sequel, where the model is unambiguous, we will write simply $s \vDash f$. The interpretation of a CTL formula $f$ with respect to a Kripke model $K$ is given below, by recursion over the structure of formulas:

$$
\begin{aligned}
& s \vDash p \text { iff } L(s)(p) \text {, where } p \text { is an atomic proposition } \\
& s \vDash \neg f \text { iff } s \not \models f \\
& s \vDash f \wedge g \quad \text { iff } \quad s \vDash f \text { and } s \vDash g \\
& s_{0} \models A X f \text { iff for all paths }\left(s_{0}, s_{1}, \ldots\right), s_{1} \models f \\
& s_{0}=E X f \quad \text { iff } \quad \text { for some path }\left(s_{0}, s_{1}, \ldots\right), s_{1} \models f \\
& s_{0} \vDash A(f U g) \quad \text { iff } \text { for all paths }\left(s_{0}, s_{1}, \ldots\right) \text {, for some } i \text {. } \\
& s_{i} \models g \text { and } \\
& \text { for all } j<i . s_{j} \models f \\
& \rightarrow=E\left(f l^{\prime} g\right) \quad \text { iff for some path }\left(s_{0}, s_{1} \ldots\right) \text {. for some } i \text {. } \\
& s_{i} \vDash g \text { and } \\
& \text { for all } j<i, \bullet_{j} \vDash f
\end{aligned}
$$

### 2.2.2 Fixed point characterization of CTL

Emerson and Clarke [CE8la] showed that various branching time properties of programs can be characterized as extremal fixed points of appropriate continuous functionals. Later. they introduced the logic CTL, and showed that its operators can be characterized in this way [C'ESIb]. This characterization led to an efficient algorithm for the model checking problem - determining whether a given ('TL formula is satisfied in

[^3]a given state of a finite Kripke model.
To obtain the fixed point characterization, we will identify each C'TL formula $f$ with $\{s \mid s \vDash f\}$. the set of states in which the formula is true. In this way. for example, true denotes the empty set. false denotes $S$, and every subset of $S$ represents an equivalence class of formulas. ${ }^{\text {b }}$ Let $\mathcal{P}(S)$ be the set of subsets of $S . \mathcal{P}(S)$ forms a lattice under union and intersection. This lattice is ordered by set inclusion. where $P \subseteq Q$ if and only if $P \cup Q=Q$. A functional $\tau[Y]$ is a formula with one uninterpreted propositional letter $Y$. This defines a function $\mathcal{P}(S) \rightarrow \mathcal{P}(S)$. where $\tau(P)$ is obtained by taking $P$ for $Y$ in $\tau$. By definition:

1. $\tau$ is monotonic when $P \subseteq Q$ implies $\tau(P) \subseteq \tau(Q)$.
2. $\tau$ is $\cup$-continuous when $P_{1} \subseteq P_{2} \subseteq \cdots$ implies $\tau\left(\cup_{i} P_{i}\right)=\cup_{i} \tau\left(P_{i}\right)$.
3. $\tau$ is $\cap$-continuous when $P_{1} \supseteq P_{2} \supseteq \cdots$ implies $\tau\left(\cap_{i} P_{i}\right)=\cap_{i} \tau\left(P_{i}\right)$.

When the set $S$ is finite, every increasing chain of subsets has a maximum element, and every decreasing chain has a minimum element. As a result, in the finite case. monotonicity implies both $\cup$-continuity and $\cap$-continuity.

A fixed point of $\tau$ is any $P$ such that $\tau(P)=P$. Tarski [Tar55] showed that a monotonic functional always has a least and a greatest fixed point with respect to inclusion ordering:

Theorem 1 (Tarski-Knaster) Whenever $\tau[Y]$ is monotonic. it has a least fixed point. denoted $\mu Y . \tau[Y]$ and a greatest fired point. denoted $\nu Y . \tau[Y]$. When $\tau[Y]$ is also U-rontinuous, $\mu Y . \tau[Y]=U_{1 \geq 0} \tau^{\prime}$ (false). When $\tau[Y]$ is also ח-rontinuous. $\nu Y . \tau[Y]=\cap_{i \geq 0} T^{\prime}($ true $)$.

We can now characterize the (TL operators in terms of fixed points of appropriate functionals:

Theorem 2 (Clarke-Emerson) Provided $S$ is finite.

[^4]1. $E F p=\mu Y \cdot(p \vee E X Y)$
$\therefore E G p=\nu Y(p \wedge E X Y)$
2. $E\left(q l^{\prime} p\right)=\mu Y(p \vee(q \wedge E . Y Y))$

There is a standard algorithm for computing the least [greatest] fixed point of a monotonic functional. This is done by starting with false [true] and iterating the functional until a fixed point is reached. as shown below. Assuming $S$ is finite. this procedure terminates in at most $|S|+1$ iterations with the least [greatest] fixed point of $\tau[Y]$ :
to compute $\mu Y . \tau[Y]$ \{or $\nu Y . \tau[Y]\}$ :
let $Y=$ false: $\{$ or $Y=$ true $\}$
do
let $Y^{\prime \prime}=Y . Y=\tau[Y]$
until $Y^{\prime \prime}=Y$ :
return $Y$
Theorem 3 Given a finite set $S$. and a monotonic functional $\tau[Y]$, the standard fixed point algorithm computes $\mu \mathrm{Y} . \tau[Y]\{$ or $\nu Y . \tau[Y]\}$ in at most $|S|+1$ iterations.

Proof. Since $\tau$ is monotonic. $\tau^{0}[$ false $] \subseteq \tau^{1}[$ false $] \subseteq \tau^{2}[$ false $] \cdots$. The longest strictly increasing chain of subsets of $S$ has length $|S|+1$. Hence. there must be an $i$ such that $0 \leq i \leq|S|$ and $r^{\prime}[$ false $]=$ $\tau^{+1}$ [false] (otherwise there would be a strictly increasing chain of length $|S|+2$ ). Hence. the algorithm terminates after at most $|S|+1$ iterations. For any such $i$. $U_{1} \geq 0^{\prime} \tau^{j}\{$ false $]=\tau^{2}[$ false $]$. Hence, by theorem 1. $\mu \zeta . \tau[Y]=\tau^{\prime}$ [false].

For the greatest fixed point. substitute true for false. こ for $\Xi$, and decreasing for increasing in the above argument.

Having a fixed point characterization of the (TL operators allows us to use the standard fixed point algorithm to determine the set of states of a given model in which a ('TL formula is true. .Is an example. consider computing $E F p$ in the following Kripke model: ${ }^{7}$

[^5]

Since $|S|=4$, the number of iterations required to produce the fixed point is at most 4 . Therefore. let us compute $r^{i}[$ false $]$ for $i=1 \ldots 4$. where $\tau[Y]=p \vee E . Y Y$. At ter the first iteration, we have $\tau^{1}[$ false $]=$ $p \vee E$.Xfalse $=p:$


After the second iteration. we have $\tau^{2}[$ false $]=p \vee E . Y p$ :


After the third iteration, we have $\tau^{3}[$ false $]=p \vee E X(p \vee E X p)$ :

which is a fixed point. since the next iteration, $\tau^{4}$ [false] produces the same result. Notice that at each iteration $i$, we have the set of states so such that there exists a path $\left(s_{0}, s_{1}, s_{2} \ldots\right)$ where $p$ is true at some state less than i. This algorithm can be thought of as a backward breadth first search of the graph. In the end, we have labeled exactly the set of states on a path to a state labeled with $p$.

As a second example. consider computing $E$ (ip in the following Kripke model:


Ster the first iteration. we have $\tau^{1}[$ triej $]=p \wedge$ EStrue $=p$ :


After the second iteration, we have $\tau^{2}[$ true $]=p \wedge E X p:$


After the thicd iteration. we have $\tau^{3}[$ true $]=p \wedge E . Y(p \wedge E . Y p)$ :


This is the greatest fixed point. since the next iteration $\tau^{4}[t r u e]$ produces the same result. Notice that at iteration $i$, we have the set of states such that there exists a path of length $i$ where every state satisfies $p$. When we reach a fixed point. every state in the set has a successor in the set satisfying $p$. hence for every state in the set. there exists an infinite path where $p$ is always true.

The operators $E . Y, E\left(I^{Y}\right)$ and $E G$ are actually sufficient to characterize the entire logic. since the remaining operators can be derived from these three according to the following rules:

$$
\begin{aligned}
E F p & =E\left(\text { true } r^{\prime} p\right) \\
\cdot X p & =\neg E \cdot \neg \neg p \\
\mid(; p & =\neg E F \neg p \\
\left.\|_{q} \mid p\right) & =\neg E\left(\neg p I^{\prime}-q \wedge \neg p\right) \vee E(; \neg p)
\end{aligned}
$$

For this reason. in the sequel. we will consider only the operators ES. Ei $l^{\prime}$, and ECi. However. for completeness. here are the fixed point characterizations of the remaining operators:

$$
\begin{aligned}
H(p & =\mu Y(p \wedge A X Y) \\
H q I p) & =\mu Y \cdot p \vee(q \wedge A Y Y)
\end{aligned}
$$

The fixed point characterization provides an effective algorithm for the model checking problem. In fact, a more efficient algorithm exists, based on breadth first search and the calculation of strongly connected components in the graph ( $S, R$ ) [CES86]. Both of these algorithms suffer from the state explosion problem. however: it is necessary to construct the complete state graph of the system being modeled before model checking can be applied. Since the number of states of a system grows exponentially in the number of its components, these algorithms can only be applied to small systems.

### 2.3 Symbolic CTL model checking

In the previous section. we equated a ('TL formula with the set of states in which the formula is true. We showed how the CTL operators can thus be characterized as fixed points of certain monotonic functionals in the lattice of subsets, and how these fixed points can be computed iteratively. In this section. we equate sets and relations with Boolean formulas. and show how set theoretic operations such as union: intersection and image can be characterized in terms of Boolean operations. This allows the CTL model checking algorithm to be implemented using well developed automatic techniques for manipulating Boolean formulas. Since the state graph is symbolically represented by a Boolean formula. there is no need to actually construct it as an explicit data structure. Hence, the state explosion problem can be avoided.

### 2.3.1 Quantified Boolean formulas

Quantified Boolean Formulas (QBF) are an extension of propositional logic allowing quantifiers over propositional variables. Given a set $V$ of propositional variables. QBF (V) is the least set of formulas such that

1. true and false are formulas.
2. every variable in $I^{\circ}$ is a formula.
3. if $p$ and $q$ are formulas, then so are $p \vee q$ and $-p$. and
t. if $p$ is a formula and $r$ is in 1 . then $\exists r$. $p$ is a formula.

A truth assignment is a function $V \rightarrow\{$ false, true $\}$. We equate each QBF formula with the set of truth assignments that satisfy the formula. Thus. true represents the set of all truth assignments. false the empty set. and a propositional variable $v$ represents the set of all truth assignments a such that $a(v)=$ true. In addition.

1. $a \in(p \vee q)$ if and only if $a \in p$ or $a \in q$.
?. $a \in(\neg p)$ if and only if $a \notin p$, and
2. $a \in(\exists v \cdot p)$ iff $a(v \leftarrow$ true $) \in p$ or $a(v \leftarrow$ false $) \in p$.

It is useful to define an operator for QBF that substitutes a formula for a variable. If $p$ and $q$ are QBF formulas, and $v$ is a variable, then let $a \in p(u \leftarrow q)$ if and only if $a(v \leftarrow(a \in q)) \in p$. Note that quantification can be defined in terms of substitution. since $\exists v . p=$ $p(v \leftarrow$ false $) \vee p(v \leftarrow$ true $)$.

Quantification and substitution can also be defined for vectors of variables. If $W=\left(w_{1}, \ldots, w_{k}\right)$ is an $n$-tuple of propositional variables and $Q=\left(q_{1}, \ldots q_{n}\right)$ an $n$-tuple of formulas. then let

1. $a \in \exists W$. $p$ iff for some $b: W \rightarrow B, a\left(w_{i} \leftarrow b\left(w_{i}\right)\right) \in p$ and
$2 . a \in p(W \leftarrow Q)$ iff $a\left(w_{i} \leftarrow\left(a \in q_{i}\right)\right) \in p$.

### 2.3.2 Representing sets and relations

The state of a concurrent system is generally modeled as vector where each component represents the state of one component of the system. For the moment. let us make the simplifying assumption that all of the state components are Boolean valued. as is generally the case in digital systems. I state of the system can therefore be viewed as a truth assignment to a set of propositional variables $V=\left\{v_{1}, c_{2} \ldots \ldots v_{n}\right\}$. Finder this interpretation. every (QBF formula over the set of state variables 1 denotes a set of states. if.. the set of truth assignments which satisty the formula. For example. if we have two state variables " and $b$. then the formmala $a \vee b$ represents all the states in which $"$ is true or $b$ is true.

In order to represent a binary relation with a QBF formula, we introduce two ordered sets of variables $V=\left\{v_{1}, \ldots, v_{2}\right\}$ and $V^{\prime \prime}=$ $\left\{v_{1}^{\prime}, \ldots v_{2}^{\prime}\right\}$. The set $V$ represents the left argument of the relation. and the set $V^{\prime}$ represents the right argument. By this arrangement. a QBF formula $R$ over the variables $V \cup V^{\prime}$ stands for a binary relation $R^{\prime}$. the set of pairs $(x, y)$ in $(V \rightarrow B)^{2}$ such that

$$
\begin{equation*}
(x ., y) \in R^{\prime} \text { iff } x\left(v_{i}^{\prime} \leftarrow y\left(v_{i}\right)\right) \in R \tag{.2.9}
\end{equation*}
$$

Is an example. if we have two state variables. $a$ and $b$. then the QBF formula $a \wedge b^{\prime}$ represents all ordered pairs of states such that $a$ is true in the first state. and $b$ is true in the second state.

Using this representation. we can express a variety of standard set theoretic operations in terms of the QBF connectives. For example, the union of two sets represented by $A$ and $B$ is $A \vee B$, their intersection is $A \wedge B$ and the complement of $A$ is $\neg A$.

The image $R^{\prime}(Q)$ of a set $Q$ via a binary relation $R^{\prime}$ is the set of all $y$ such that for some $x \in Q,(x, y) \in R^{\prime}$. If $R$ is a QBF formula representing a relation $R^{\prime}$, and $Q$ is a QBF formula representing a set. then

$$
\begin{equation*}
R^{\prime}(Q)=(\exists V .(R \wedge Q))\left(V^{\prime} \leftarrow V\right) \tag{2.10}
\end{equation*}
$$

We can prove this by simply expanding the definitions of the QBF operators, as follows:

$$
\begin{gathered}
\eta \sqsubseteq(\exists l:(R \wedge Q))\left(V^{\prime \prime} \leftarrow V\right) \\
\text { iff } \\
!\left(r^{\prime}-y\left(r_{i}\right)\right) \in \exists V:(R \wedge Q) \\
\text { iff } \\
\text { exists } x: V-B \text { s.l. } y\left(v_{i}^{\prime} \leftarrow y\left(v_{i}\right)\right)\left(v_{i} \leftarrow x\left(v_{i}\right)\right) \in R \wedge Q \\
\text { iff } \\
\text { exists } x: l \rightarrow B \text { s.t. }(x . y) \in R^{\prime} \text { and } x \in Q \\
\text { iff } \\
y \in R^{\prime}(Q)
\end{gathered}
$$

As an example. len ()$=\| \vee b$ and $R=a \wedge b$. Then

$$
R(Q)=\left(\exists 1:\left(\left(a \wedge b^{\prime}\right) \wedge(a \vee b)\right)\left(l^{\prime \prime}-1\right)\right.
$$

$$
\begin{aligned}
& =\left(\exists V \cdot\left(a \wedge b^{\prime}\right)\right)\left(V^{\prime} \leftarrow V\right) \\
& =b^{\prime}\left(V^{\prime} \leftarrow V\right) \\
& =b
\end{aligned}
$$

The inverse inage $R^{-1}(Q)$ of a set $Q$ via a binary relation $R$ is the set of all $x$ such that for some $y \in Q, R(x, y)$. If $R$ is a QBF formula representing a relation, and $Q$ is a QBF formula representing a set. then

$$
\begin{equation*}
R^{\prime-1}(Q)=\exists V \cdot\left(R \wedge Q\left(V \leftarrow V^{\prime \prime}\right)\right) \tag{2.11}
\end{equation*}
$$

This can be shown by a derivation similar to the one above.

### 2.3.3 CTL formulas

We now have the necessary mechanics to represent Kripke structures using QBF formulas, and to characterize the CTL operators over these symbolically represented Kripke structures using QBF operators. In fact, it is only necessary to characterize the CTL operator E.K. since the logical connectives have identical meanings in both logics. and the remaining CTL operators have already been characterized as fixed points of functionals using only $E X$ and the logical operators.

To represent a Kiripke structure symbolically, we will assume two sets of variables $V=\left\{v_{1}, \ldots, v_{n}\right\}$ and $V^{\prime \prime}=\left\{v_{1}^{\prime} \ldots, v_{n}^{\prime}\right\}$. and a QBF formula $R$ on $V \cup V^{-1}$ to represent the transition relation. This induces a Kripke structure $h_{i: V^{\prime}, R}^{\prime}=\left(S, R^{\prime}, L\right)$ where

1. The state set $S$ is the set of truth assignments $I \rightarrow B$.

2 . The transition relation $R^{\prime}$ is the relation representer by the formula $R$. according to (2.9).
3. The valuation $L$ yields the truth value of each variable $r$, in each states. That is. for all $r_{i} \in V . L(s)\left(v_{i}\right)=s\left(v_{i}\right)$.

The complete procedure for symbolic model checking is characterized by the following theorem:

Theorem 4 Let $V^{\circ}=\left\{c_{1} \ldots \ldots c_{n}\right\}$ and $V^{\prime \prime}=\left\{r_{1}^{\prime} \ldots \ldots l_{n}^{\prime}\right\}$ be disjoint sets of rarmbles. lit R be a QBF formula on l! l", and let kiva.n be
the induced Kripke model. In this model. for all CTL formulas $p$ and $q$ :

$$
\begin{align*}
& s \vDash p \text { iff }  \tag{2.12}\\
& s \in p . \text { where } p \in V  \tag{2.13}\\
& s \vDash p \vee q \text { iff } s \in(p \vee q)  \tag{2.14}\\
& s \vDash \neg p \text { iff }: \in(\neg p)  \tag{2.15}\\
& s \vDash E . Y p \text { iff }: \in\left(\exists V^{\prime} .\left(R \wedge p\left(V \leftarrow V^{\prime}\right)\right)\right)  \tag{2.16}\\
& s \vDash E(q l p) \text { iff }: \in \in \mu Y .(p \vee(q \wedge E X Y))  \tag{2.17}\\
& s \vDash E C p \text { iff } s \in \nu Y .(p \wedge E X Y) .
\end{align*}
$$

Proof. The first three are trivial matters of definition. For ( $\because .1 \overline{5}$ ) . when we equate a formula witi the set of states satisfying it. $E . X p$ is just $R^{\prime-1}(p)$, which is equal to $\exists V^{\prime \prime} .\left(R \wedge p\left(V \leftarrow V^{\prime}\right)\right)$. The last two are just theorem 2 .

The above theorem shows that we can solve the model checking problem - ie., determining whether a given state in a symbolically represented Kripke structure $K_{V \cdot V^{\prime} \cdot R}^{\prime}$ satisfies a formula $f$ - purely by manipulations of Boolean formulas. A key point is that the Kripke structure itself is never built. Instead it is symbolically represented by a QBF formula. As an example. consider a system with one state variable $b$. Let the transition relation be represented by the formula $R=b \vee b^{\prime}$. and let state $\rightarrow$ be ( $b \leftarrow$ false). The induced Kripke structure $\mathscr{l}_{\left\{\left\{^{\prime},\left\{b^{\prime}\right\}, R\right.\right.}$ is depicted below:


Let ss say we want to determine whether or not $s \in E . Y \neg b$. According to theorem 4 , we can evaluate the formula $E X \neg b$ as follows:

$$
\begin{aligned}
E . X \neg b & =\exists b^{\prime} \cdot\left(R \wedge(\neg b)\left(b \leftarrow b^{\prime}\right)\right) \\
& =\exists b^{\prime} \cdot\left(\left(b \vee b^{\prime}\right) \wedge\left(\neg b^{\prime}\right)\right) \\
& =\exists b^{\prime} \cdot\left(b \wedge \neg b^{\prime}\right) \\
& =b
\end{aligned}
$$

Hence by theorem $4 . s \models E . Y \neg b$ iff the assignment ( $b \leftarrow$ false) satisfies b. which is ialse.

Now consider the problem of whether or not $s \vDash E F b$. Using the standard fixed point algorithm, we get

$$
\begin{aligned}
\left.\tau^{1} \text { [false }\right] & =b \vee E . \mathrm{false} \\
& =b \\
\tau^{2}[\text { false }] & =b \vee E \cdot \mathrm{~V} b \\
& =b \vee \exists b^{\prime} \cdot\left(\left(b \vee b^{\prime}\right) \wedge b^{\prime}\right) \\
& =\text { true } \\
\tau^{3}[\text { alse }] & =b \vee E . Y \text { true } \\
& =\text { true }
\end{aligned}
$$

A fixed point is reached after two iterations. Hence. $: \notin E F b$ iff the truth assignment ( $b \leftarrow$ false) satisfies true, which is true.

Note that when computing least (or greatest) fixed points of $\tau .2^{n}+1$ iterations are required in the worst case, where $n$ is the number of propositional state variables. This is the length of the longest possible strictly increasing (or decreasing) chain of subsets of $S$ (not including the empty set), plus one extra iteration to detect the fixed point. In practice. however. the number of iterations required to reach a fixed point can be quite small.

### 2.3.4 Binary Decision Diagrams

It should be clear that to make the symbolic model checking technigue practical. an efficient automated method for manipulating Boolean formulas is required. Fortunately, a variety of such techniques have been developed for the purpose of synthesizing digital circuits or comparing the functionality of digital circuits. These techniques may involve applying a set of rewriting rules to convert a given formula into a normal form. Ntematively. a data structure may be used to represent the formula as a Boolean fiunction." For example. a Boolean lunction may be

[^6]represented by a truth table, or by a set of "cubes" which cover the truth table of the function, or by a binary decision tree. Each representation has associated procedures for applying Boolean operations. Any method of manipulating Boolean formulas that can implement the operations $p \wedge q, p \vee q, \neg p, \exists V^{\prime} . p$ and $p\left(V \leftarrow V^{\prime}\right)$ can be used in symbolic model checking. By far the most effective method known to date. however, is the Ordered Binary Decision Diagram method developed by Bryant [Bry86].

Ordered Binary Decision Diagrams are a form of reduced decision graph that give compact canonical representation for Boolean formulas. They have been used extensively for comparison of switching functions $\left[\mathrm{BBB}^{+} 87\right.$. FB89]. The OBDD canonical representation for a Boolean function can be derived by reducing a related structure called an ordered decision tree. In an ordered decision tree, the value of the function is obtained by descending the tree from the root to a leaf. At each node along the path, one descends to the left child if the value of the variable labeling the node is 0 , and to the right child the value is 1 . Each leaf of the tree is labeled with a value 0 or 1 which gives the result of the function. The tree is said to be ordered if the variables always occur in the same order along any path from root to leaf. In this case, reading the leaves from left to right, one obtains the truth table of the function.

As an example. an ordered decision tree for the function $a \wedge b \vee c \wedge d$ is clepicted in figure 2.I.

The canonical OBDD form is a directed acyclic graph which can be obtained from the ordered decision tree by the following two steps:

1. Combine any isomorphic subtrees into a single tree.
2. Eliminate any nodes whose left and right children are isomorphic.

Steps 1 and 2 can be applied in a bottom up fashion, to yield the canonical OBDD representation in linear time. Bryant called this operation Reduce. The size of the resulting graph is strongly dependent
$V=\left\{c_{1} \ldots . v_{n}\right\}$ induces a function $f:\{0.1\}^{n}-\{0.1\}$ in the obvions way: $f\left(x_{1} \ldots \ldots x_{n}\right)=1$ iff the truth assignment $\left(r_{i}-x_{i}\right)$ satisfies $f$. The two views are equivalent. but the functional representation seems to be more standard in the context of Bootean manipulation.


Figure 2.1: Ordered Decision Tree
on the order of the variables. This variable ordering, however. is the key to obtaining the reduced form. This is what distinguishes OBDDs from the more general class of Binary Decision Diagrams described by Akers [Ake78].

As an illustration of reduction to canonical form. consider the the ordered decision tree of figure 2.1. The three nodes marked "*" are roots of isomorphic subtrees. Thus, they can be combined into a single subtree. In addition. from the node marked "+", one arrives at the same subtree when descending to the left or right (ie.. independently of the value of $b$ ). hence this vertex does not affect the value of the function and may be eliminated. The result of applying the Reducf operation to the tree of of figure 2.1 is depicted in depicted in figure 2.2. Note the signifirant rerluction in the number of vertices. resulting essentially from redundancy in the truth table of the function.

The canonical OBDDs are a subclass of DAGs (directed acyclic graphs) where each leaf is labeled by 0 or 1 . and each non-leat is labeled by a variable. It is most convenient to define this class inductively. by building large DAG from smaller ones. For this reason. we will number the variables from the bottom up.' In the sequel the term dimension will be used to denote the highest variable index occurring in a DAG.

[^7]

Figure 2.2: Ordered Binary Decision Diagram

We will simultaneously define the class of DAGs which are canonical OBDDs and the functions they denote, by induction on the dimension:

Definition 1 Let $V$ be an n-tuple $\left(v_{1}, v_{2}, \ldots, v_{n}\right)$ of variables. The class $O B D D(V)$ consists of the terminals 0 and 1 , and a collection of triples in $\mathcal{S} \times O B D D(\mathcal{S}) \times O B D D(\mathcal{S})$ called non-terminals. With each tement $p$ of $O B D D(V)$. we associate a dimension $d_{p}$. where $0 \leq d_{p} \leq$ n. and a Boolean function $f_{p}: B^{n} \rightarrow B$. The class $O B D D(\mathcal{S})$ is the least such that. for all $x \in\{0.1\}^{n}$ :

1. $0 \in O B D D(1)$. $d_{10}=0$. and $f_{0}\left(x^{\circ}\right)=0$.
$\therefore \quad 1 \in O B D D(V) . d_{1}=1$. and $f_{1}(x)=1$.
2. If 1 and $h$ art distinct elements of $O B D D(V)$. where $d_{l}<i \leq n$ fand $d_{h}<i$. then the triple $r=\left(r_{t}, l, h\right)$ is also in $O B D D\left(V^{\circ}\right)$. $l_{r}=i$ and

$$
f_{r}(x)= \begin{cases}f_{i}(x) & \text { if } x_{i}=0 \\ f_{h}(x) & \text { if } x_{i}=1\end{cases}
$$

With regard to canonicity, the salient aspects of the above definition are that a triple ( $r, l . l /$ ) is a canonical OBDD only if $l$ and $h$ are distimet and $i$ is qreater than the dimensions of $/$ and $h$ (the variable
ordering requirement). ${ }^{10}$ One important consequence of this is that $f_{p}$, the function represented by a DAG $p$, does not depend on any variables of index greater than $d_{p}$ :

Lemma 1 For all $p \in O B D D(V)$, for all $d_{p}<i \leq n . f_{p}\left(v_{i} \leftarrow 0\right)=$ $f_{p}\left(v_{i} \leftarrow 1\right)$.

Proof. By induction over $d_{p}$. We assume the statement of the theorem holds for all $q$ such that $d_{\eta}<d_{p}$. The terminal cases. $p=0$ and $p=1$ are trivial. For the non-terminal case. let $p=\left(v_{j}, l, h\right)$. where $j<i$. Now consider two cases. $v_{j}=0$ and $v_{j}=1$. In the first case. $f_{p}\left(v_{i} \leftarrow 0\right)=f_{l}\left(v_{i} \leftarrow 0\right)$ and $f_{p}\left(v_{i} \leftarrow 1\right)=f_{l}\left(v_{i} \leftarrow 1\right)$. These are equal by inductive hypothesis. since $d_{l}<i$. The other case. $c_{j}=1$ is similar, with $f_{h}$ for $f_{l}$.

It is not difficult to show that OBDDs canonically represent the Boolean functions. That is, each Boolean function is represented by exactly one OBDD. We show first that there are no two distinct OBDDs representing the same function, and second, that every Boolean function is represented by some OBDD. The following theorem is essentially due to Bryant [Bry86], although the formalization is different, and as a result. it is hoped, the proof is substantially simpler.

Theorem 5 (Bryant) If $p$ and $p^{\prime}$ are elements of $O B D D(V)$. then $f_{p}=f_{p^{\prime}}$ implies $p=p^{\prime}$

Proof. By simultaneous induction over $d_{p}$ and $d_{p}$. We assume the statement of the theorem holds for all $q$ and $q^{\prime}$. where $d_{q}<d_{p}$ and $d_{q^{\prime}}<d_{p^{\prime}}$. Suppose that $f_{p}=f_{p^{\prime}}$ :

Consider first the case where $d_{p}=d_{p^{\prime}}$. Either $p$ and $q$ are both terminals. (in which case $p=p^{\prime}=0$ or $p=p^{\prime}=1$ ) or they are both non-terminals. $p=\left(v_{i}, l, h\right)$ and $p^{\prime}=\left(v_{i}, l^{\prime}, h^{\prime}\right)$. For non-terminals, we

[^8]have $\left.f_{l}=f_{p}\left(v_{i} \leftarrow 1\right)\right)=f_{p^{\prime}}\left(l_{i^{\prime}} \leftarrow 0\right)=f_{l^{\prime}}$, and similarly $f_{h}=f_{p}\left(v_{i} \leftarrow\right.$ 1) $=f_{p^{\prime}}\left(v_{i^{\prime}} \leftarrow 1\right)=f_{h^{\prime}}$. Hence. by induction. $l=l^{\prime}$ and $h=h^{\prime}$. so $p=p^{\prime}$.

Second. consider the case where $d_{p}>d_{p^{\prime}}$. It follows that $p$ is a non-terminal ( $\left.c_{i}, l, h\right)$. Further, by the previous lemma. $f_{p^{\prime}}\left(v_{i} \leftarrow 0\right)=$ $f_{p^{\prime}}\left(v_{i} \leftarrow 1\right)$. Therefore. $f_{p}\left(v_{i} \leftarrow 0\right)=f_{p}\left(v_{i} \leftarrow 1\right)$, so $f_{l}=f_{h}$. $\mathrm{B} y$ induction. then, $l=h$. This is a contradiction, however, since if $l$ and $h$ are not distinct. then $p$ is not in $\operatorname{OBDD}(V)$.

A symmetric argument applies to the case $d_{p}<d_{p^{\prime}}$.

Theorem 6 Given a function $f: B^{n} \rightarrow B$. there exists $p \in O B D D(V)$ such that $f_{p}=f$.

Proof. By induction on $i$, the greatest number such that $f\left(c_{i} \leftarrow\right.$ $0) \neq f\left(v_{i} \leftarrow 1\right) . \quad$ By inductive hypothesis, there exist $q$ and $r$ in $\operatorname{OBDD}(V)$ such that $f_{q}=f\left(v_{i} \leftarrow 0\right)$ and $f_{r}=f\left(v_{i} \leftarrow 1\right)$. Further. $q$ and $r$ are distinct. since $f\left(v_{i} \leftarrow 0\right) \neq f\left(v_{i} \leftarrow 1\right)$. Thus, let $p=\left(v_{i}, q, r\right)$.

Because each function is represented by a unique OBDD. testing two OBDDs for functional equality can be accomplished in constant time. This property of OBDDs is useful for determining when a fixed point has been reached in the standard fixed point algorithm.

## The Apply algorithm

Bryant describes an algorithon called Apply. which applies an arbitrary Boolean operation - to two OBDDs. The operation - can be any of the 16 Boolean functions of two variables - $A p p l y$ computes the natural extension of to two Boolean functions. Given two non-terminal OBDDs $p$ and $q$. the tpply algorithm breaks the problem of computing $r=p \bullet q$ into two subproblems on the children of $p$ and $q$.

Take first the aspe where $d_{p}=d_{q}$. Let $p=\left(c_{i}, l_{p}, h_{p}\right)$ and $q=$ ( $b_{1}, l_{q}, h_{q}$ ). It is rasily shown that
$\bullet r\left(r_{i}-0\right)=p\left(r_{2}-0\right) \bullet y\left(r_{1}-0\right)=I_{p} \bullet l_{1}$, and

- $r\left(v_{i} \leftarrow 1\right)=p\left(v_{i} \leftarrow 1\right) \bullet q\left(v_{i} \leftarrow 1\right)=h_{p} \bullet h_{q}$.

Thus. we create two subproblems $l=l_{p} \bullet l_{q}$ and $h=h_{p} \bullet h_{q}$. On the other hand. suppose that $p=\left(v_{i}, l_{p}, h_{p}\right)$ and $q=\left(v_{j}, l_{q}, h_{q}\right)$, where $i>j$. In this case. $q\left(v_{i} \leftarrow 0\right)=q\left(v_{i} \leftarrow 1\right)=q$. So.
$\bullet r\left(v_{i} \leftarrow 0\right)=p\left(c_{i} \leftarrow 0\right) \bullet q=l_{p} \bullet q$ and
$\bullet r\left(v_{i} \leftarrow 1\right)=p\left(v_{i} \leftarrow 1\right) \bullet q=h_{p} \bullet q$.
Therefore, we create two subproblems $l=l_{p} \bullet q$ and $h=h_{p} \bullet q$. The remaining case. $i<j$, is symmetric.

The subproblems are solved recursively to obtain $l=r\left(c_{i} \leftharpoondown 0\right)$ and $h=r\left(c_{1} \leftarrow 1\right)$. From these two cofactors, we can derive $r$. If $l$ and $h$ are equal. then $r=h=l$. If they are distinct, then $r=\left(v_{i}, l . h\right)$. Finally, if $p$ and $q$ are both terminals. Apply simply uses the truth table for $\bullet$

Since each subproblem of dimension $d$ can generate two subproblems of dimension $d-1$, it might seem that this algorithm is exponential. It can be made polynomial, however, by applying dynamic programming. Notice that each subproblem is determined by a pair of OBDDs $p^{\prime}$ and $q^{\prime}$ which are descendants of $p$ and $q$ respectively. Hence, the maximum number of distinct subproblems is the product of the size of $p$ and the size of $q$. By keeping a hash table of triples $(p, q, r)$, we can reduce the number of recursive calls to $|p| \cdot|q|$. Bryant shows that this upper bound is tight. since there exist functions $p$ and $q$ for which the size of $r$ is $|p| \cdot|q|$.

## The (ompose algorithm

Bryant also gives an algorithm called ('ompose which computes $p\left(c_{i} \leftarrow\right.$ $q$ ), where $p$ and $q$ are OBDDs. and $v_{i}$ is a variable. The algorithm is easily adapted for simultaneous substitution of a vector of variables. Hence. given that

$$
\begin{equation*}
\exists v_{i}, p=p\left(c_{i} \leftarrow 0\right) \vee p\left(c_{i} \leftarrow 1\right) . \tag{2.18}
\end{equation*}
$$

the compose procedure could be used to implement both the variable substitution operation $p\left(V^{\circ} \leftarrow\left[^{\prime \prime}\right)\right.$ and the existential quantitication
operation $\exists V^{\prime}$.p needed for symbolic model checking. On the other hand. a much more efficient procedure can be obtained by combining the quantification and conjunction operations in the expression for $E . X p$ into a single OBDD operation computing $\exists V^{\prime} \cdot(p \wedge q)$. Applying the quantifiers in a bottom-up fashion as the conjunction subproblems are solved results in a substantial reduction in the size of the intermediate results by reducing the number of variables.

## The AndExists algorithm

This algorithm, which we will call AndExists is basically a modification of Apply. Let $r$ be the OBDD representing the function $\exists l^{\prime \prime} .(p \wedge q)$. We compute $r$ by generating subproblems $l$ and $h$ in the same manner as if using the Apply algorithm for $\bullet=\wedge$. When the results of the subproblems are obtained. if the leading variable $n_{i}$ is a component of $V^{\prime}$. the result is $r=l \vee h$ (see equation 2.18). This result is obtained by calling Apply with $\bullet=V$. On the other hand. if $v_{i}$ does not occur in $V^{\prime}$. then the result is the same as for Apply: if $l=h$, then $r=l=h$. else $r=\left(v_{i}, l, h\right)$.

The motivation for this algorithm is to avoid producing the entire OBDL for $p \wedge q$. which has $\rho_{n}$ variables. where $n$ is the number of state variables of the model. This is done by applying existential quantification to the results of subproblems as soon as they become available. yielding a result with only $n$ variables. Empirically, this provides a substantial savings in space.

As in the Apply algorithm. a table of triples ( $p, y . r$ ) is used to avoid resolving previonsly computed subproblems. The maximum size of this table is $|p| \cdot|q|$. However, unlike in the Apply algorithm. the recursive ralls cannot be executed in constant time. This is becanse each call may require a $V$ operation to be performed. It present. the author is unaware of a bound on the complexity of AudErists better than $O\left(|p| \cdot|q| \cdot 2^{2 n}\right)$. which is simply the number of $\vee$ problems to be solved $||p| \cdot| q \mid$ in the worst case $)$ imes the square of the largest possible OBDD size. $2^{n}$. In practice. this number of operations has not been observed. so one might conjecture that there is a tighter bound. It seems unlikely. that a polynomial bound will be found. however. since it is easily shown That if vector existemtial quantification on OBDDS can be computed in


Figure 2.3: Variable ordering for 3-SAT reduction
polynomial time, then $P=. N P$.
The proof of this is by reduction from 3-SAT, as follows: Let $f=$ $t_{1} \wedge t_{2} \wedge \cdots t_{k}$ be a 3-SAT formula, that is, $t_{i}=\left(x_{i} \vee y_{i} \vee z_{i}\right)$, where $x_{i}$, $y_{1}$ and $z_{1}$ are positive or negative literals. The OBDD representation of each $t_{1}$ has no more than 3 non-terminals. Now introduce new variables $V^{\prime \prime}=\left(v_{1}^{\prime}, v_{2}^{\prime}, \ldots . v_{k}^{\prime}\right)$. corresponding to the terms of $f$, and let

$$
f^{\prime}=\bigvee_{1 \leq i \leq k}\left(\neg t_{i} \wedge v_{i}^{\prime} \wedge \bigwedge_{1 \leq j<i} \neg v_{j}^{\prime}\right)
$$

For a suitable variable ordering, the OBDD representing $f^{\prime}$ has no more than $t k$ non-terminals (see figure 2.3 ). hence can be built in polynomial time. The formula $f$ is satisfiable iff $\exists V^{\prime \prime} . f^{\prime} \neq 1$. Thus, if $\exists l^{\prime \prime} . f^{\prime}$ can be computed in polynomial time, then $\mathrm{P}=\mathrm{NP}$.

Is an aside, it is not difficult (though a bit tedious) to show that the symbolic ('TL model checking problem is PSPAC'E-complete. To show PSPAC'E-hardness, one starts with a polynomial space bounded Turing machine. introduces a sufficient number of Boolean variables to encode the entire tape. plus the pointer and the finite control, then expresses the transition relation of the entire system as a QBF formula. Co show that the problem is in PSPACE. one can show that the problem can be reduced to satisfiability of a QBF formula of polynomial size. using the "iterative squaring" technique of Burch. et al. $\left[\mathrm{BC}^{\prime} \mathrm{M}^{+} 90\right]$. Details are left to the mader.

### 2.4 Examples

Although the worst case complexity of symbolic model checking is high (using OBDDs or other Boolean function representations), in practice the worst case complexity is rarely achieved. and the symbolic technique can in some cases be dramatically more efficient than previous methods. .Is an illustration of this. let's look at two hardware examples - a synchronous fair bus arbiter, and an asynchronous distributed mutual exclusion ring circuit (the one studied by David Dill in his thesis [Dil88] and designed by Alain Martin (Mar85]).

### 2.4.1 Synchronous state machines

For a synchronous finite state machine. the transition relation can be given as a conjunction of Boolean formulas. each determining the new state of one register as a function of its old state and the inputs. Let $V=\left\{v_{1}, v_{2}, \ldots, v_{n}\right\}$ be a set of Boolean variables representing the state of the registers in the circuit. and let $W=\left\{w_{1}, w_{2}, \ldots w_{m}\right\}$ be a set of variables representing the values of the inputs to the circuit. For all $i=1 \ldots n$, let $f_{i}[V, W]$ define the value of register $i$ in the next state, in terms of $V$ and $W$. The transition relation of the state machine can be expressed as a Boolean formula in the following form:

$$
\begin{equation*}
R=\bigwedge_{i=1}^{n} R_{i} . \quad \text { where } R_{i}=\left(v_{i}^{\prime} \Longleftrightarrow f_{i}[V . W]\right) \tag{2.19}
\end{equation*}
$$

In general. for morlels of senchronous systems. the transition relation is a conjunction of formulas representing the individual components of the system. since transitions of the components are simultaneous. The outputs of the state machine can be given as Boolean functions of the inputs and registers. These functions can be substituted for atomic propositions in ('TL formulas. so there is no need to introduce variables to represent the outputs.

As an example of a synchronous state machine. we will consider a synchronous bus arbiter circuit. The purpose of the bus arbiter is to grant access on each clock cycle to a single client among a number of rlients contending for the use of a bus (or other resource). The


Figure 2.4: Cell of synchronous arbiter circuit
inputs to the circuit are a set of request signals $r e q_{0} \ldots r e q_{k-1}$. and the outputs are a set of acknowledge signals $a c k_{0} \ldots a c k_{k-1}$. Normally, the arbiter asserts the acknowledge signal of the requesting client with the lowest index. However. as requests become more frequent. the arbiter is designed to fall back on a round robin scheme, so that every requester is eventually acknowledged. This is done by circulating a token in a ring of arbiter cells. with one cell per client. The token moves once every clock cycle. If a given client's request persists for the time it takes for the token to make a complete circuit, that client is granted immediate access to the bus.

The basic cell of the arbiter is depicted in figure 2.4.1. This cell is repeated $k$ times. as shown in figure 2.1.1. Each cell has a request input and an acknowledge output. The grant output of cell $i$ is passed to cell $i+1$. and indicates that no clients of index less than or equal to $i$ are requesting. Hence, a cell may assert its acknowledge output if its grant input is asserted. Each cell has a register $T$ which stores a one when the token is present. The $T$ registers form a circular shift register which shifts up one place each clock cycle. Each cell also has a register $W$ (for "waiting") which is set to one when the request input is asserted and the token is present. The register remains set while the request persists. until the token returns. At this time, the cell's override and acknowlerlge ontputs are assertel. The override ismal propagates


Figure 2.5: Configuration of the synchronous arbiter circuit
through the cells below, negating the grant input of cell 0 , and thus preventing any other cells from acknowledging at the same time. The circuit is initialized so that all of the $W$ registers are reset and exactly one $T$ register is set.

The desired properties of the arbiter circuit are:

1. No two acknowledge outputs are asserted simultaneously
2. Every persistent request is eventually acknowledged
3. Acknowledge is not asserted without request

Expressed in ('TL. they are:

1. $\wedge_{i \neq f} A G \neg\left(a c k, \wedge a_{j} k_{j}\right)$
2. $\wedge_{t} \cdot \operatorname{AGAF}\left(\mathrm{req}_{2} \Rightarrow \mathrm{ack}_{1}\right)$
3. $\wedge_{i} A\left(i\left(\right.\right.$ ack,$\left.=r e q_{i}\right)$

Using the symbolic ('TL model checking procedure, we can determine whether the design has these properties, for a given mumber of
cells. Figure 2.6 plots the performance of the symbolic model checking procedure for this example in terms of several measures: the size of the transition relation in OBDD nodes. the total run time (on a Sun3. running an implementation in the C' language), and the maximum number of OBDD nodes used at any given time. ${ }^{11}$ We observe that as the number of cells in the circuit increases. the size of the transition relation increases linearly (in section 2.5 . we will prove a theorem that shows why this is the case). The execution time is well fit by a quadratic curve. The number of reachable states. however. explodes exponentially (note the logarithmic scale on the reachable states axis).

To obtain polynomial performance for this example. it was necessary to add a wrinkle to the symbolic model checking algorithm. In the first experiment it was found that although most of the specification was checked quickly, the time required to check property 2 for cell 0 doubled each time a cell was added. The reason for this is rather remarkable. Consider a function called Rotate, which returns true for a pair of $n$ bit binary numbers when one number can be obtained from the other by a rotation of $j$ bits. There is no variable ordering which yields an efficient OBDD for this function for all $j .{ }^{12}$ In fact, a very similar function occurs in computing the set of states satisfying the formula $A F$ (req ${ }_{0} \Rightarrow$ ack $_{0}$ ), where the two binary numbers are given by the $W$ and $T$ registers respectively. Note that. for cell 0 . request implies acknowledge exactly when no other cell has both $W$ and $T$ registers set. The $T$ registers rotate once per clock cycle. Thus. re $\mathrm{l}_{0} \Rightarrow$ ack $_{0}$ is necessarily true $j$ steps in the future exactly when there is no other cell $i$ for which $W_{i} \wedge T_{i-\jmath \operatorname{modk}}$. The OBDD representing this set of states grows exponentially in the number of cells.

This illustrates a fairly general phenomenon: circuits tend to be "well behaved" in the part of their state space which is reachable from the initial state. but not elsewhere. In the case of the synchronous arbiter. only states with one $T$ register set are reachable. However.

[^9]
### 2.4. EXAMPLES



Figure 2.6: Performance - synchronous arbiter example.
the symbolic model checking technique considers all states. including states with multiple tokens. A good solution to this problem in general is first to compute the set of reachable statis. and then to restrict all of computations of the CTL model checking algorithm to those states. Since the reachable states are closed under the transition relation. this has no effect on the truth value obtained for formulas at the initial state. In particular, this solves the problem of the bus arbiter circuit, since in its reachable state space, the $T$ registers cannot represent an arbitrary binary number.

The set of reachable states is the least fixed point of

$$
\tau[Y]=I \vee R(Y)
$$

where $l$ is the set of initial states. Applying the standard fixed point algorithm in this case effectively yields a forward breadth first search of the state space. By computing the reachable states first and then using this set to restrict the (TL model checking algorithm. we obtain the polynomial run time results described above. This technique is also used for other experiments described in the sequel. unless otherwise noted.

### 2.4.2 Asynchronous state machines

In an asynchronous state machine. there is no global clock to which all state changes are synchronized. This makes designing correct asynchronous circuits considerably more challenging than designing correct synchronous circuits. We will consider two plausible models of asynchronous state machines. In the first. which we will rall the simulta"rous morlel, any or all state variables may change state in a given transition. Each state component makes an independent and nondeterministic choice regarding whether to change value or not. In the second morlel. which we will call the interleaving model, ouly one state component changes value in a given transition. The choice of which component changes value is non-deterministic. ${ }^{13}$ In cither model. we

[^10]consider an asynchronous state machine composed of $n$ gates. We will use state variable $v_{i}$ to stand for the output of gate $i$. and $f_{i}[V, W]$ to represent the function computed by gate $i$ (where $V$ is the set of state variables. and $W$ the set of inputs).

In the simultaneous model. the transition relation can be represented by a formula in the form:

$$
\begin{equation*}
R=\bigwedge_{1 \leq i \leq n} R_{i}, \text { where } R_{i}=\left(v_{i}^{\prime} \Longleftrightarrow f_{i}[V, W]\right) \vee\left(v_{i}^{\prime} \Longleftrightarrow v_{i}\right) . \tag{2.20}
\end{equation*}
$$

For any transition and any state variable $c_{i}$, either the new value of $c_{i}$ is determined by $f_{i}[V, W]$. or it is the same as the old value. Note that this differs from the synchronous model (2.19) in which every state variable is reevaluated at every transition.

In the interleaving model. the transition relation can be represented by a formula in the form:

$$
\begin{equation*}
\bigvee_{i} R_{i}, \text { where } R_{i}=\left(v_{i}^{\prime} \Longleftrightarrow f_{i}[V, W]\right) \wedge\left(\wedge_{j \neq i}\left(v_{j}^{\prime} \Longleftrightarrow v_{j}\right)\right) \tag{2.21}
\end{equation*}
$$

In any transition, for some state variable $v_{i}$, the new value of $v_{i}$ is determined by $f_{i}[V, W]$, and the remaining variables keep their old value. Note that in this case, the transition relation is represented by a disjunction of component relations rather than a conjunction.

In general. for models of parallel processes whose actions interleave arbitrarily, the transition relation is disjunctive. If this is the case. we ran make an easy optimization in the symbolic model checking technique: we observe that the set of states reachable by one step of the system is the union of the sets of states reachable by one step of each individual component. This is reflected in the fact that existential quantification distributes over disjunction. Thus:

$$
\begin{aligned}
E . X p & \left.=\exists V^{\prime} \cdot\left(1 V R_{t}\right) \wedge p\left(V^{\prime}-V^{\prime}\right)\right) \\
& =V \exists V^{\prime} \cdot\left(R_{t} \wedge p\left(V \leftarrow V^{\prime}\right)\right)
\end{aligned}
$$

Fing this equality, we can avoid computing the transition relation of the system and instead nse only the transition relations of the indi-


Figure 2.7: One cell of the DME circuit
vidual processes. This technique is called early quantification ${ }^{14}$ - by rearanging the computations, we apply quantification before the logical disjunction operation. Heuristically, quantification tends to reduce OBDD size. since it reduces the number of variables. Hence. the size of the intermediate results is usually reduced (though the final result is the same).

Our example of an asynchronous state machine is the distributed mutual exclusion (DME) circuit of Alain Martin [Mar85]. It is a speedindependent circuit [Sei80b] and makes use of special two-way mutual exclusion circuits as components. Figure 2.7 is a diagram of a single cell of the distributed mutual-exclusion ring. The circuit works by passing a token around the ring, via the request and acknowledge signals $l r$ and $l a$ on the left and $r r$ and $r a$ on the right. A user of the DME gains exclusive access to the resource via the request and acknowledge signals ur and ua.

The specifications of the DME circuit are as follows:

1. Vo two users are acknowledged simultaneously:
$\because$. In acknowledgment is not output without a request.
2. An acknowledgment is not removed while a request persists.

[^11]4. All requests are eventually acknowledged.

We will consider only the first specification, regarding mutual exclusion. The others are easily formulated in CTL, although the last requires the use of fairness constraints (see section 2.6.1) to guarantee that all gate delays are finite. The formalization of the mutual exclusion specification is

$$
\bigwedge_{i \neq j} A G \neg\left(u a_{i} \wedge u a_{j}\right)
$$

Now let's look at the performance of the symbolic model checking algorithm in checking this formula. for both a simultaneous and an interleaving model of the circuit. For the interleaving model. we use the early quantification technique. Figure 2.3 plots the relative performance for the simultaneous model (method 1) and the interleaving model (method 2 ). Part (a) shows the run time as a function of the number of DME cells, part (b) shows the total storage used (measured in OBDD nodes) and part (c) shows the number of nodes used to represent the transition relation. For the moment, disregard the curves for method 3. The experiment was run for up to 7 cells of the simultaneous model (limited by space) and up to 10 cells of the interleaving model (limited by time). Part (b) of the figure shows the substantial space advantage of the interleaving model, and from part (c), we can see that most of the difference is accounted for by the savings in representing the transition relation using early quantification. In both cases, the space used is linear in the number of cells. However. we note that the increase in run time appears to be cubic for the simultaneous model, but quartic for the interleaving model. It would seem that if enough storage were available to continue the curve for method 1 . the two curves would meet in the neighborhood of 10 cells.

The different asymptotic performance for the simultaneous and interleaving models can be understood by looking at the OBDDs that occur in the fixed point iterations computing the reacbable states. Figure 2.9 plots the size of the largest such OBDD for each method. We ran see clearly that the size is increasing linearly for the simultaneous model. but quadratically for the interleaving model. This is a phenomenon which occurs generally when comparing simultaneous is interleaving mordels. It ran be understood by considering a tery simple
system composed of $n$ processes. each with states 0 and 1 . and each alternating non-deterministically between these two states. If we start the system with all processes in state 0 , what do we observe after $k$ steps? In the simultaneous case. after one step, all possible states are reachable. In the interleaving case. however, after $k$ steps, all global states with at most $k$ l's are reachable. This is a symmetric function. As Bryant noted [Brys6], all symmetric functions can be represented by a quadratic size OBDDs. The symmetry results from the fact that in an interleaving model. exactly one state component changes in a given transition. and the choice is arbitrary. In general, after $k$ steps of such a model. the number of steps taken by each state component sums to $k$. Hence. in the set of states reachable after $k$ steps, there is an induced correlation between the states of otherwise independent processes.

The simultaneous model appears to be inferior to the interleaving model from a symbolic model checking point of view. owing to the large amount of space required to represent the transition relation. Most of this. however. can be attributed to a phenomenon we observed in the previous example: systems tend to be well behaved only in their reachable state space. In the symbolic model checking technique, we represent the transition relation over the entire state space. Although representing only the reachable transitions might be more efficient, we seem to be caught in Catch 22: we need to represent the transition relation to compute the set of reachable states. We can avoid this problem by incrementally computing only as much of the transition relation as is necessary to compute the next iteration of the fixed point algorithm. Recall that the reachable state set is the least fixed point of $\tau[Y]=I \vee R\left(Y^{*}\right)$. By rearranging the fixed point computation slightly. we only need represent $R$ correctly for those transitions ( $x, y$ ), where $x$ is on the "frontier" of the search:

$$
\begin{aligned}
\tau^{+1}(\text { false }) & =\tau^{i}(\text { false }) \vee R\left(\tau^{i}(\text { false })\right) \\
& =\tau^{i}(\text { false }) \vee R\left(\tau^{i}(\text { false })-\tau^{i-1}(\text { false })\right)
\end{aligned}
$$

It each iteration. we can reevaluate the formula $R$ over the set of states $\tau^{i}$ (false) $-\tau^{-1}$ (false). This can be done by restricting each subformula using either the logical and or using the Restrict operator of Coudert. Madre and Berthet (see section 2. §). This results in a sequence
of approximations to the transition relation which are substantially more compact than the complete transition relation. although we must reevaluate $R$ at each iteration, rather that evaluating it once at the beginning. We will call this method 3 .

In part (a) of figure 2.8. we see that the time used by this method. while still cubic. is a substantial improvement over the previous method for the simultaneous model (method 1). More importantly, the space used is dramatically improved, allowing a model with a larger number of cells to be checked. The method overtakes the interleaving model in run time at about 8 cells. owing to its better asymptotic performance.

Figure 2.10 plots ihe number reachable states as a function of the number of cells (the numbers are indistinguishable for the two models). The number of reachable states grows exponentially in the number of cells. though not as rapidly as the total number of states. which is $2^{18 n}$. The key point is that for all three methods. the space and time necessary for the symbolic model checking method is polynomial in the number of cells. Thus. the state explosion problem has been avoided. The overall time complexity of $O\left(n^{3}\right)$ for the simultaneous model derives from three factors: a linear increase in the transition relation OBDD, a linear increase in the state set OBDDs obtained in the fixed point iterations. and a linear increase in the number of iterations. For the interleaving model. the quadratic increase in the state set OBDDs results in an overall $O\left(n^{4}\right)$ time complexity. On the other hand. the number of reachable states increases roughly a factor of ten with each added cell.

It is not immediately clear that either the interleaving or simultaneous model is preferable in general. Interleaving models seem to be better when the number of astnchronons processes is small, and simultaneons when the number is large. The cache consistency protocol of chapter $t$ is an example of a large system with a fairly small number of complex asynchronous processes. This is an appropriate application of an interleaving model.

The polynomial performance of the symbolic model checking algorithm, in spite of the exponential increase in states, makes it possible to analyze fairly large instantiations of the two example circuits the synchronous arbiter and the D.ME circuit). It should be possible to werify these and similat cirmits for any reasonable fixed mumer of cells. This


Figure 2.s: Performance for DME circuit example


Figure 2.9: State set size for DME circuit example


Figure 2.10: Reachable states for DME circuit example
begs the question - how many cells do we need to analyze to be guaranteed that the design is correct for any number of cells? Intuitively, for sufficiently large $n$. a sequence of $n+1$ cells should be equivalent in some sense to a sequence of $n$ cells. But in what sense equivalent? This problem is dealt with in chapter 5 , where we consider induction over processes.

### 2.5 Graph width and OBDDs

In this section. we consider the asymptotic growth of OBDDs representing certain topological classes of circuits. This analysis explains some of the performance results of the previous section.

In 1989. Berman proved a bound on the OBDD size needed to represent circuits of bounded width. A circuit has bounded width if its elements can be arranged in a linear order such that any cut through the order crosses at most a bounded number of wires $v$. called the width of the circuit. There exists a variable ordering such that the OBDD size is bounded by $n 2^{w}$. where $n$ is the number of primary inputs of the circuit. This result applies only if the order is "topological". meaning essentially that the direction of all the wires follows the ordering. Here. this result is generalized. to show that if $w_{f}$ bounds the number wires through any cut in the forward direction. and $w_{r}$ bounds the number in the reverse direction, then the OBDD size is bounded by $n \cdot 2^{\omega_{t}} 2^{\omega_{r}}$. For the case where $w_{r}=0$. this is the same as Berman's result. I'sing this result. we can linearly bound the OBDD representation for the transition relation of circuits like the arbiter and the D.ME ring, which have linear arrangements with a bounded number of wires through any rrosss section.

Fujita states that tree circuits using only I.ND. OR and XOR gates have linearly bounded OBDD representations [F.MK90]. Here, we show that a more general class of circuits with bounded "tree width" and arbitrary function elements have polynomially bounded OBDDs. The essence of the argument is to show that these circuits can be arranged in a linear order with a width that is logarithmic in the number of gates. This yields a bound on the OBDD size which is polynomial in the mumber of gates.

### 2.5.1 Bounded width circuits

Let $L=(G,<)$ be a linear order on the gates of a circuit. We classify the primary inputs and outputs of the circuit as special instances of gates in order to simplify the definitions, and assume that the primary output is at the top of the order. Given an order $L$. we will say that the forward cross section of the circuit at gate $g$ is the set of wires connected to an output of some gate $g_{1}$ and an input of some gate $g_{2}$ such that $g_{1} \leq g$ and $g<g_{2}$. The reverse cross section is the set of wires connected to an output of some gate $g_{1}$ and an input of some gate $g_{2}$ such that $g_{2} \leq g$ and $g<g_{1}$. We assume that no wire is connected to the outputs of two distinct gates. so these two sets are disjoint. We also assume that there are no cycles in the circuit. to insure that the circuit computes a function. The order $L$ is said to be topological when all of the reverse cross sections are empty.

The forward width of the circuit under order $L$. denoted $w_{f}$, is the maximum size of the forward cross section at any gate $g$. Similarly, the reverse width of the circuit under order $L$, denoted $w_{r}$ is the maximum size of the reverse cross section at any gate $g$.

The cross section of an OBDD at level $i$ is the set of nodes labeled with variable $v_{i}$. Note that in this section, we will number the variables of the OBDD from the top down, since this makes the proofs simpler. The width $w_{p}$ of an OBDD $p$ is the maximum size of any cross section of $p$. The size of an OBDD is the sum of the sizes of its cross sections. Thus, the OBDD size if bounded by $n \cdot w_{p}$. where $n$ is the number of variables.

It is easily shown that the size of the cross section of an OBDD at level $i$ is the number of distinct functions

$$
f_{p x}\left(c_{1} \ldots \ldots c_{n}\right)=f_{p}\left(x_{1} \ldots \ldots c_{1-1} \cdot c_{1}, \ldots c_{n}\right)
$$

which depend on $r_{i}$. where $x=\left(r_{1} \ldots \ldots r_{1-1}\right)$ is a Boolean vector and $f_{p}$ is the function represented by $p$. This observation leads to the following theorem bounding the size of an OBDI) in terms of the forward and reverse widths of the circuit it represents:

Theorem 7 If a ciruit computing function f has forward width wf "and renerse width wefor some linear orver L. then there is an (OBI)D)


Figure 2.11: Proof of bounded width theorem
$p$ representing function $f$ of size bounded by $n 2^{w^{w} 2^{w_{r}}}$, where $n$ is the number of inputs of the circuit.

Proof. Associate the variables $v_{1}, \ldots, v_{n}$ of the OBDD with the inputs of the circuit, such that for all $i \leq j, v_{i} \leq v_{j}$. We can bound the size of the $i$ th cross section of the resulting OBDD as follows. Let $x=\left(x_{1} \ldots \ldots x_{t-1}\right)$ be a Boolean vector. Split the circuit in half by choosing any gate $g$ such that $v_{i-1} \leq g<v_{i}$, letting $Y$ be the forward cross section at $y$ and $Z$ the reverse cross section. This situation is depicted in figure 2.11 . For any given value of $x, Y$ is a function of $Z$. and this finction determines $f_{r}\left(v_{i} \ldots \ldots v_{n}\right)$. The number of Boolean functions with $|Z|$ inputs and $|Y|$ outputs is $2^{|Y| 2^{|Z|}}$ (to see this. count the number of entries in the truth table). This bounds the total number of distinct functions $f_{r}$. which in turn bounds the width of the OBDD representing $f$ at level $i$. We know that $\left|Y^{\prime}\right| \leq w_{j}$ and $|Z| \leq w_{r}$. Thus. the overall OBDD size is bounded by $n \cdot 2^{\omega_{f} 2^{\omega_{r}}}$.

This bond is linear in the number of inputs. exponential in the forward width and doubly exponential in the reverse width. The double exponential appears to be necessary. This can be shown using the
"hidden weighted bit" function of [Bry91] as a counterexample. This circuit can be ordered in such a way that between any two inputs there is a cross section with $O\left(\log _{2} n\right)$ wires in each direction, yet there is an exponential lower bound on its OBDD size. If we could bound the OBDD size with a single exponential in both the forward and reverse widths, the OBDD size would be $O\left(n 2^{k \log _{2} n}\right)=O\left(n^{k+1}\right)$ where $k$ is a constant.

The theorem is concerned with a single output of a combinational circuit, but it can also be applied to the transition relation of a sequential circuit. To do this. we simply transform the sequential circuit into a combinational circuit which computes the transition relation of the sequential circuit. This is clone by adding a pair of inputs $v_{i}$ and $v_{i}^{\prime}$ to represent the old and new ralues of each state component. Since the transition relation of the circuit is the conjunction of the transition relations of its components. we can do this while increasing the width of the circuit by only one wire in the forward direction as depicted in figure 2.12. Thus, for bounded width sequential circuits (even with wires in both directions), the size of the OBDD representing the transition relation is linear in $n_{i}+n_{s}$, where $n_{i}$ is the number of inputs and $n_{s}$ is the number of state components. The synchronous arbiter circuit and the DME circuit of the previous section provide experimental confirmation of this.

We have shown for a certain structural class of circuits that the representation of the transition relation is linearly bounded in the size of the circuit. We should note that in the symbolic model checking algorithm, we also use OBDDs to represent the set of states labeled with a given ('TL formula. 「nfortunately, we cannot expect to polynomially bound the size of the OBDDs representing these sets based purely on structural considerations. The simplest example of this is probably a circuit that inputs a binary number. stores one copy of it. then serially. rotates the original by an arbitrary number of bits. This circuit has the simplest structure we might hope for that has any communication at all between the components, yet there is no variable order which vields a compact OBDD for the reachable state set of this circuit. since it implements the rotate function. The same argument would apply to a serial multiplier circuit. In general. if a circuit computes a function serially witich cannot be represented by a compact OBDD. then we


Figure 2.12: Computing a conjunctive transition relation
cannot expect the symbolic model checking algorithm using OBDDs to be efficient.

### 2.5.2 Bounded tree-width circuits

In the previous section. we considered the OBBD representation of circuits whose gates can be arranged in a sequence with a bounded number of wires in each cross section. Now we consider the slightly more general class of circuits which can be can be arranged in a tree with a bounded width property. This is not to say that the topology of the circuit must be a tree: rather. it must be possible to lay a spanning tree over the circuit in such a way that the width of the circuit across any arc of the spanning tree is bounded. This notion of bounded treewidth is defined as follows.

Let $T=1\left(\begin{array}{r} \\ .\end{array}<\right)$ be a tree order over the gates of a circuit. where $g^{\prime}<g$ iff $g^{\prime}$ is a descendant of $g$. Let $b$ be the branching degree of $T$ (ie.. the maximum number of children of any gate). As before, the forward rooss section at node $g$ is the set of wires connecting an output of $g_{1}$
and an input of $g_{2}$ such that $g_{1} \leq g$ and $g<g_{2}$. Similarly, the reverse cross section of $T$ at node $g$ is the set of wires connecting an output of $g_{1}$ and an input of $g_{2}$ such that $g_{2} \leq g$ and $g<g_{1}$. The forward width of the tree $w_{f}$ is the size of the largest forward cross section. while the reverse width $w_{r}$ is the size of the largest reverse cross section.

For the moment. let us consider the case $w_{r}=0$. and let the width $w$ stand for the forward width:

Lemma 2 For any topological tree order $T=(G,<)$, with width $w$ and branching degree $b>1$. there is a topological linear order $L=\left(G,<^{\prime}\right)$, with width $w^{\prime} \leq w(b-1) \log _{2}\left|G^{\prime}\right|$.

Proof. By induction over $\left|C_{i}\right|$. the number of gates. The base rase. $|G|=1$. is trivial. Assume the theorem holds for all circuits of size less than $|G|$. Let $g$ be the root of the tree. and let $G_{1} \ldots \ldots G_{k}$ be the subtrees of the root, where $k \leq b$, and $\left|G_{1}\right| \leq \cdots \leq\left|G_{k}\right|$. By inductive hypothesis. there exist linear orders $L_{i}=\left(G_{i}^{\prime},<_{i}\right)$ of width $w_{i} \leq v(b-1) \log _{2}\left|G_{i}\right|$, for all $1 \leq i \leq k$. Let $L=\left(G^{\prime},<^{\prime}\right)$ be the extension of these orders such that $G_{k}<^{\prime} \cdots<^{\prime} G_{1}<^{\prime} g$, as depicted in figure 2.13. The width $w^{\prime}$ of $L$ is bounded by $\max _{1 \leq i \leq k}\left(w_{i}+(k-i) w\right)$. Therefore for some $i$.

$$
w^{\prime} \leq w_{i}+(k-i) w
$$

In the case $k=i$. we have

$$
\left.\left.w^{\prime} \leq w_{k} \leq w^{\prime} b-1\right) \operatorname{log_{2}}\left|C_{k}\right| \leq w_{1} b-1\right) \log _{2}\left|C_{i}\right|
$$

()therwise. $i<k$ and

$$
\begin{aligned}
w^{\prime} & \leq w\left(k-i+(b-1) \log _{2}\left|G_{i}^{\prime}\right|\right) \\
& \leq w(b-1) \log _{2}\left(2^{\frac{k-1}{b-1}}\left|G_{i}^{\prime}\right|\right)
\end{aligned}
$$

Here we note that $\left|G_{i}\right| \leq\left\{\sum_{i \leq s \leq i}\left|(i, \mid) /(k-i+1) \leq\left|G_{i}\right| /(k-i+1)\right.\right.$. so

$$
u^{\prime}=w(b-1) \log _{2}\left(\left.\frac{2^{\frac{k-1}{b-1}}}{k-i+1} \right\rvert\,\left(: 0^{\prime}\right)\right.
$$



Figure 2.13: Arrangement of bounded width tree
We note that since $i \geq 1$ and $k \leq b . k-i \leq b-1$. Therefore $2^{\frac{k-1}{b-1}} \leq 2$. Further, since $i<k, k-i+1 \geq 2$. Thus $\frac{2^{k-1}}{k-i+1} \leq 1$. Therefore.

$$
w^{\prime} \leq w(b-1) \log _{2}\left|G_{i}^{\prime}\right|
$$

The theorem says that from any topological tree order of width $w$ wr can derive a linear order of width $w^{\prime} \leq w(b-i) \log _{2}|C i|$. It follows be the previous theorem that the OBDD size is bounded by $n 2^{n^{\prime}} \leq n 2^{n+4 n-1)} \log _{2}|G|=n\left|G_{0}\right|^{n+(b-1)}$. where $n$ is the number of primary inputs. This bound is polynomial in the size of the circuit for a fixed width and branching factor.
. $\begin{aligned} & \text { ow we turn to the question of tree orders that are not topological }\end{aligned}$ (if.. bounded tree-width circuits with both forward and reverse wires). In this case. a logarithmic bound on the width of the linear order $L$ is not sufficient. because the OBDD size can be doubly exponential in the number of it er erse wires.

We ran still obtain a polynomial bound in $n$. however. by converting a tree ordered cirnit with reverse wises into a functionally equivalent
tree ordered circuit with only forward wires:
Lemma 3 If $T=((i,<)$ is a tree order over a circuit computing function $f$. with forward width $w_{f}$ and reverse width $w_{r}$. then there is a circuit computing $f$ with topological tree order $T^{\prime}=\left(C_{r}^{\prime},<^{\prime}\right)$ of forward width $w_{f}^{\prime} \leq w_{f} 2^{w_{r}}$.

Proof. Consider $H$. a subtree rooted at gate $h$. letting $Y$ be the forward cross section at $h$. and $Z$ the reverse cross section at $h$. Let $h_{1}, \ldots . h_{k}$ be the children of $h$. and let $Y_{1} \ldots \ldots Y_{k}$ and $Z_{1} \ldots \ldots Z_{k}$ be their respective forward and reverse cross sections. This situation is depicted in figure 2.14. Let the output functions computed by $H$ be

$$
Y=f\left(Z . Y_{1}, \ldots . Y_{k}\right)
$$

and for $1 \leq i \leq k$, let

$$
\begin{aligned}
Z_{i} & =r_{i}\left(Z . Y_{1} \ldots \ldots Y_{k}\right) \\
Y_{i} & =f_{i}\left(Z_{i}\right)
\end{aligned}
$$

We show by induction over $|H|$ that there exists a tree circuit $H^{\prime}$ of forward width $w_{f}^{\prime} \leq w_{f} 2^{w_{r}}$ and reverse width $w_{r}^{\prime}=0$. computing the functions

$$
f_{r}=f\left(x . Y_{1} \ldots . . Y_{k}\right) \text {, for } x \in\left\{0 .[ \}^{|Z|}\right.
$$

Note that $f_{r}$ is simply row $r$ in the truth table for $\gamma$. Since there are $y^{|2|}$ possible values of $x$. and $f_{r}$ has $|Y|$ components. the number of outputs of $H^{\prime}$ is $\mid \xi^{2 / 2 /}$.

By inductive haporlesis. there exist rercuits $I_{1}^{\prime}$ for $1 \leq 1-k$. satisfying the width bound and romputing the functions

$$
l_{i r}=f_{i}(x) \text { for } x \in\{0,1\}^{\left|Z_{z}\right|}
$$

Now. let $h^{\prime}$ be a gate computing $f_{r}$ accorting to the following sistem of equations:

$$
\begin{aligned}
f_{r} & \left.=j 1 r \cdot f_{1,1} \ldots \ldots f_{k r_{k}}\right) \\
r & =r \cdot f_{1,1} \ldots \ldots f_{k s_{k}} \mid \text { for } 1,1, k
\end{aligned}
$$



Figure 2.14: A non-topological tree order
Let $H^{\prime}$ be the tree ordered circuit obtained by taking $h^{\prime}$ as the root. and $H_{1}^{\prime} \ldots \ldots H_{k}^{\prime}$ as the children of the root. The reverse width at the root is 0 . since $f_{i x}$ does not depend on $Z$. and the forward width at the root is $|Y| 2^{|z|}$. Hence, using the inductive hypothesis. $u u_{r}^{\prime}=0$ and $w_{f}^{\prime} \leq w_{f} 2^{w_{r}}$. If $h$ is the root node of $G$. then $H^{\prime}$ computes the same function as $G$.

This gives us the following theorem. bounding the OBDD size for tree ordered circuits with both forward and reverse wires:

Theorem 8 If a rircuit $C_{i}$ computing function $f$ has formard width w; und ifefrse width $w_{r}$ for some tree order $T$ of branching degree $h>1$. then there is un ()BDD representing function $f$ of size bounded


Proof. According to lemma 3. for any tree ordered circuit of forward width $w_{f}$ and reverse width $w^{\prime}$.. we can construct a iopological tree ordered cirruit of width $w \leq w_{f} 2^{w_{r}}$. which computes the same function. By lemma 2. this circuit has a topological linear order $L$ of width at most $w^{\prime} \leq m(b-1) \log _{2}|\boldsymbol{r}|$. By theorem 7 . there is an OBDD for the circuit of size bounded by

$$
\begin{aligned}
& n 2^{n u^{\prime}} \leq \pi 2^{2(t h-1) \log _{2} \mid(G)} \\
& \because 112^{n+2^{n r}\left(h-11 \log _{2} i(i)\right.}
\end{aligned}
$$

$$
=n|C|^{\omega_{f} 2^{w_{r}}(b-1)}
$$

Hence. in the case of bounded tree width circuits (of a fixed branching degree), we also find that the OBDD size can be bounded polynomially in the size of the circuit. In this case. the exponent of $n$ is related to both the width and the branching factor. Clearly, for this bound to be of any practical interest. $w_{f}$ must be small, and $w_{r}$ must be very small. .Vonetheless, the theorem demonstrates a more general topological class of circuits with asymptotically compact OBDDs than was previously known.

### 2.6 Mu-Calculus model checking

The Mu-Calculus [Parit] is a logic based on extremal fixed points that is strictly more expressive than CTL, ${ }^{15}$ and can also express a variety of properties of transition systems. such as reachable state sets. state equivalence relations, and language containment between automata. A symbolic model checking algorithm for this logic allows all of these properties to be computed using OBDDs [ $\mathrm{BCM}^{+} 90$ ].

The Mu-C'alculus augments the ordinary predicate calculus in two ways. First. it allows terms to stand for relations. If $f$ is a formula in which variables $x$ and $y$ are free. then $f$ characterizes a relation the set of all pairs $(x, y)$ satisfying $f$. This relation is denoted in the Mu-('alculus by the term ..r.y.f. Second. the Mu-('alculus allows us 10 express least and greatest fixed points. If $T$ is a term, and $Y$ is a relational (predicate) symbol. then $T$ is said to be formally monotonic in $Y$ if $Y$ always occurs under an even number of negations in $\tau$. In this rase. $\tau$ has least and greatest fixed points with respect to $Y$. which are denoted $\mu Y$. - and $\nu Y$. r. . A fixed point of $\tau$ with respect to $Y$ is a relation which yields itself when substituted for all free occurrences of $\gamma$ in $\tau$.

[^12]A structure in the Mu-Calculus consists of a set $\mathcal{D}$ (the clomain). a valuation $\varphi$ for the individual symbols $\{a . b, c, \ldots\}$ and a valuation $\psi$ for the relational symbols $\{A, B, C \ldots\}$. The valuations assign an element from the domain to each individual symbol. and a set of $n$-tuples from the domain to each relational symbol. The meaning of an $n$-ary term $\tau$ is a set of $n$-tuples which we will denote $\tau[0, \Delta]$. The 0 -ary terms will be called simply propositions. and denote truth values.

The terms of the Mu-Calculus are the least set such that:

1. Every relational symbol is a term.

2 . If $\tau$ is an $n$-ary term and $\left(v_{1} \ldots \ldots v_{n}\right)$ are individual symbols. then $r\left(r_{1} \ldots \ldots r_{n}\right)$ is a proposition.
3. If $\tau_{1}$ and $\tau_{2}$ are $n$-ary terms. then so are $\neg \tau_{1}$ and $\left(\tau_{1} \vee \tau_{2}\right)$.

1. If $p$ is a proposition. and $v$ is an individual symbol. then $\exists r \cdot p$ is a proposition.

5 . If $p$ is a proposition and $\left(v_{1}, \ldots v_{n}\right)$ are individual symbols, then $\lambda v_{1}, \ldots . v_{n}, p$ is an $n$-ary term.
6. If $\tau$ is an $n$-ary term and $Y$ is an $n$-ary relational symbol, where $\tau$ is formally monotonic in $Y$, then $\mu Y . \tau$ and $\nu Y . \tau$ are $n$-ary terms.
$\bar{i}$. The usual abreviations are used for $\wedge . \Rightarrow . \forall . \operatorname{tc}$.
The semantics of Mu-Calculus terms are defined as follows:

1. $R[0,0]=\theta(R)$, where $R$ is a relational symbol.


2. ( $\exists r, p)[0.0]$ is true if for some $x \in \mathcal{D} . p[\Phi(r-x), \Leftrightarrow]$ is true.
3. $\left(\lambda_{1} \ldots \ldots r_{n} \cdot p\right)[0.1 \cdot]$ is the set of $n$-tuples $\left(x_{1} \ldots \ldots r_{n}\right) \leqslant D^{n}$ such that pol: - $r, x, \cdots$ is tme.
(i. $\{\mu\}$. $T\left[0.4\right.$. where $T$ is an 1 -ary term, is the least set $S \subseteq D^{n}$ such that $\cdot=T\left[0.0\left(Y \leftarrow Y^{\prime}\right)\right] .(\nu Y . T)[0.0]$ is the greatest such $\therefore$.

### 2.6.1 Applications of the Mu-Calculus

The Mu-Calculus is quite expressive, as can be seen by the following compendium of applications. To begin with, given a binary relation $R$. the image of a set $Q \subseteq \mathcal{D}$ via $R$ is

$$
R(Q)=\lambda y \cdot \exists x \cdot(R(x, y) \wedge Q(x))
$$

The set reachable from $Q$ in any number of steps of $R$ (including 0 ) is

$$
R^{*}(Q)=\mu Y \cdot(Q \vee R(Y))
$$

The transitive (irreflexive) closure of the relation $R$ is

$$
R^{+}=\mu \digamma \cdot[R \vee \lambda x . z \cdot \exists y \cdot(Y(x, y) \wedge Y(y, z))]
$$

## CTL and fairness constraints

The interpretation of the operators of CTL in a Kripke model ( $\mathcal{D}, R, L$ ) can be characterized in the Mu-Calculus as follows:

$$
\begin{aligned}
E X p & =\lambda x \cdot \exists y \cdot(R(x, y) \wedge p(y)) \\
E F p & =\mu Y \cdot(p \vee E X Y) \\
E G p & =\nu Y \cdot(p \wedge E X Y) \\
E\left(q L^{\prime} p\right) & =\mu Y \cdot(p \vee(q \wedge E X Y))
\end{aligned}
$$

In addition to these standard operators. we can also characterize the ("TL operators umder fairness constraints. I fairness constraint in its simplest form is a condition that is assumed to hold infinitely often along all computation paths. Such conditions can be used to enforce fair scheduling of processes and access to resources. They are not directly expressible in ("TL. since the tense operators $F$ and (; cannot be directly rombined. Instead. we restrict the path quantifiers of ('TL to apply only to those paths along which each formula in a set (' holds infinitely often. To distinguish these constrained path quantifiers from orrlinary path quantifiers. we subseript them with ('. Thus. Ic $f$. where $r$ ' is a set of ('TL formulas and $f$ is a linear formula. means that for all paths. if each formula of $($ is true intintely often. then $f$ is true.

Similarly: the formula $E_{C} f$ means that there exists a path such that each formula of (' is true infinitely often and $f$ is true. Here. we consider only the ('TL operators with existential path quantifiers. since the operators with universal quantifiers can be derived from these.

The formula $E_{C} G p$ is true when there is some path in which $p$ is true in every state. and each element of $C$ is true infinitely often. Let

$$
T[Y]=p \wedge E . Y \bigwedge_{c \in C} E(Y \zeta(Y \wedge c))
$$

We argue as follows that $E_{C} C i p$ is the greatest fixed point of r. First. if $Y$ is a tixed point, then every state in $Y$ satisfies $p$, and further. has a nontrivial path remaining in $Y$ which leads to a state satisfying each fairness constraint. Hence. a looping path can be constructed satisfying each intinitely often without exiting $Y$. Thus $Y \subseteq E_{C} G p$. On the other hand. suppose $Y^{\prime}=E_{c}$ Gip. Since every state in $Y$ has a path touching each fairness constraint infinitely, as does each state along that path. it follows that every state in $Y$ can reach every fairness constraint without exiting $Y$. Thus $Y \subseteq \tau[Y]$. Therefore, $E_{C} G_{p} p$ is the greatest fixed point of $\tau$. The set of states satisfying $E_{C} C_{i} p$ is expressed in the Mu-Calculus as

$$
\nu Y .\left(p \wedge E X \bigwedge_{c \in C} E\left(Y Y^{-}(Y \wedge c)\right)\right)
$$

The remaining operators under fairness constraints can be characterized in terms of Er. Cip as follows:

$$
\begin{aligned}
E_{C} \cdot \mathrm{Y} p & =E \cdot\left(p \wedge E_{C} \cdot(; \text { true })\right. \\
E_{r} \cdot F p & =E F\left(p \wedge E_{r} \cdot(; \text { true })\right. \\
E_{r}\left(\eta I^{\prime} p\right) & =E\left(q l^{\prime}\left(p \wedge E_{C}(; \text { true })\right)\right.
\end{aligned}
$$

Emerson and Lai [ELS6] give a characterization in the Mu-('alculus of ('TL. under a more general class of faimess constraints. Each constraint in this sheme reguires that one condition holds infinitely often or a second condition holds finitely often (for example. either acknowledge holds infinitele sem. or request holds timitely often).

## Simulation relations

Two states $x$ and ! of a Kripke structure are said to be bisimular if:

1. $x$ and ! agree on the atomic propositions.
$\because$ erery successor of $x$ is bisimular to a successor of ! and
2. every successor of ! 1 is bisimular to a successor of $r$.

Two states are bisimular if and only if they satisfy the same set of (TL formulas $\left[B C\left(N_{1}\right]\right.$. If $\left.1 a_{1}, a_{2} \ldots \ldots \|_{4}\right)$ are the atomic propositions. then the bisimulation relation ran be expressed in the Mu-('alculus as follows:

$$
\begin{aligned}
\text { Bisim }= & \nu Y \cdot \lambda \cdot x^{\prime} \cdot y \cdot \bigwedge_{1 \leq \leq \leq k}\left(u_{i}(x) \Longleftrightarrow a_{i}(y)\right) \\
& \wedge \forall x^{\prime} \cdot\left(R\left(x \cdot x^{\prime}\right) \Rightarrow \exists y^{\prime} \cdot\left(R\left(y, y^{\prime}\right) \wedge Y\left(\cdot x^{\prime} \cdot y^{\prime}\right)\right)\right) \\
& \left.\wedge \forall y^{\prime} \cdot\left(R\left(y \cdot y^{\prime}\right) \Rightarrow \exists x^{\prime} \cdot\left(R\left(x \cdot x^{\prime}\right) \wedge Y\left(x^{\prime} \cdot y^{\prime}\right)\right)\right)\right)
\end{aligned}
$$

where we have. as usual. identified each atomic proposition with the set of states in which is is true. There is also an asymmetric notion of -imulation we say hat a state $r$ simulates a state $!$ if:

1. rand! agree on the chmir propositions.
Z. very - uccenor of 1 i- immated by a sucessor of $x$.
 $r$ in a dialect of (TI. called $7-1$ "TL, which allows only miversal path grantitiers pelan! $^{16}$ Testing bismulation and simulation relations can be used as a form of wertication. or it can be nsed to test abstractions wed in compositional motel cherking terhnigues ('L.Ma9a, (iL.91'. The atue idea can easily be extended to ststems with labeled transtions.
[^13]
## Language containment

The Mu-C'alculus can can express the relation of language containment between two deterministic $\omega$-automata. For the sake of simplicity, we consider only deterministic Büchi automata. which are not complete for the class of $\omega$-regular languages. but it is not substantially more difficult to handle more general classes of deterministic automata. such as Street automata.

I finite deterministic Büchi automaton consists of a set of states $K$. an initial state $p_{0} \in K$, an alphabet $\Sigma$. a set of transitions $\nu \subseteq$ $K \times \Sigma \times K$. and an acceptance set $B \subseteq K$. The transition relation is such that. for any state $p$ and symbol $\sigma$. there is exactly one $q$ for which $\Delta(p, \sigma, q)$. The automaton accepts an infinite sequence $\sigma \in \Sigma^{\omega}$ iff the sequence of states $p$. where $\Delta\left(p_{i}, \sigma_{i}, p_{i+1}\right)$ holds for all $i$, passes through the acceptance set $B$ infinitely often. The set of sequences accepted by an automaton $M$ is called the language of $M$ and denoted $\mathcal{L}(M)$.

To determine whether the language of a Büchi automaton $M$ is contained in the language of a Büchi automaton $M^{\prime}$ (with the same alphabet). we define a Kripke structure representing the product of $M$ and $M^{\prime}$. and write a formula in CTL which is true if and only if every sequence accepted by $M$ is also accepted by $M^{\prime}$ [CDK90]. This formula can be evaluated using its Mu-Calculus characterization.

The product is defined by its transition relation $R$, and set of initial states so. Let

$$
\begin{aligned}
& \text { 1. } R=\lambda . s, s^{\prime} . r \cdot r^{\prime} . \exists \sigma .\left(\Delta(s, \sigma \cdot r) \wedge \Delta^{\prime}\left(s^{\prime} \cdot \sigma . r^{\prime}\right)\right) . \\
& \because . r_{0}=\lambda . . s^{\prime} \cdot\left(\left(. s=p_{0}\right) \wedge\left(s^{\prime}=p_{0}^{\prime}\right)\right) .
\end{aligned}
$$

There is a sequence in the language of $I /$ but not in the language of $W^{\prime}$ if and only if there is an path of the product passing through $B$ intinitely often. but not through $B^{\prime}$ infinitely often. Thus. $\mathcal{L}(M) \subseteq$ $\mathcal{L}\left(I^{\prime}\right)$ iff

$$
\therefore_{0} \Rightarrow A\left(i _ { 1 } \left\{_{\left\{1 s . s^{\prime}, B(s)\right\}} F \text { A.s. } s^{\prime} . B^{\prime}\left(. s^{\prime}\right)\right.\right.
$$

Another possible approach to the language containment problem makes use of the transitive closure of the transition relation. First, we remowe from the product structure all transitions that b: gin or end
with a state in $B^{\prime}$. That is. let

$$
T=\lambda s . s^{\prime} . r . r^{\prime}\left[R\left(. s . s^{\prime} . r . r^{\prime}\right) \wedge \neg B^{\prime}\left(s^{\prime}\right) \wedge \neg B^{\prime}\left(r^{\prime}\right)\right]
$$

The transitive closure of this relation is
$T^{+}=\mu Q\left[\lambda s, s^{\prime} . r, r^{\prime}\left[T\left(\ldots, \therefore^{\prime}, r^{\prime}, r^{\prime}\right) \vee \exists u, u^{\prime}\left[Q\left(\ldots, s^{\prime} \cdot u . u^{\prime}\right) \wedge Q\left(u, u^{\prime}, r, r^{\prime}\right)\right]\right]\right]$
This is the set of all pairs $(x . y)$ of states of the product such that $x$ can reach ! without passing through $B^{\prime}$. This holds for the pair $(x . x)$ if and only if $x$ is on a crele not passing through $B^{\prime}$. If there is any such $x$ in $B$. and $x$ is reachable. then there is a path passing through $B$ but not $B^{\prime}$ intinitely often. hence there is a sequence in $\mathcal{L}(M)$. but not in $\mathcal{L}^{\prime}\left(M^{\prime}\right)$. The conterse is also true. Hence. $\mathcal{L}(M) \subseteq \mathcal{L}\left(. M^{\prime}\right)$ if and only if $-E F \lambda s . s^{\prime}$. $\left(T^{+}\left(\mathrm{s}, s^{\prime} \ldots . s^{\prime}\right) \wedge B(s)\right)$. The $E F$ operator can also be evaluated using the transitive closure. since

$$
E F p=\lambda x \cdot\left(p(x) \vee \exists y \cdot\left(R^{+}(x, y) \wedge p(y)\right)\right)
$$

### 2.6.2 Symbolic algorithm

By devising a symbolic model checking procedure for the Mu-Calculus. we can quickly establish symbolic algorithms for all of the above properties. If we assume that the domain is $B^{k}$. a symbolic model checking algorithm is easily established. by translating formulas into a Boolean Mu-l alculus where the domain is just $B=\{$ false, true $\}$. This is , tone by replacing every individual symbol $a$ by a $k-t u p l e ~ o f ~ i n d i v i d u a l ~ s y m-~$
 In the Boolean Mn- ${ }^{\prime}$ alculns we can represent terms bey Boolean formulas by introducing a new set of dumme individual symbols $d_{1} . d_{2} \ldots$. to represent relational parameters. An $n$-ary term $\tau$ is represented by a formula ric.j such that

$$
\begin{aligned}
& \left.1 . r_{1} \ldots . r_{n}\right)=-\{0.1] \\
& \text { iff }
\end{aligned}
$$

(iisen $\because$ wer can compur the formula repersenting a lerm in the Boolean


1. The value of relational variable $A$ is $w(A)$.
2. The logical connectives and quantifiers are evaluated by the corresponding QBF operations.
3. The value of an $n$-ary term $\lambda\left(v_{1}, \ldots . v_{n}\right) . \tau$ is

$$
\tau[u]\left(v_{1} \leftarrow d_{1} \ldots . v_{n} \leftarrow d_{n}\right)
$$

4. The value of the proposition $\tau\left(v_{1} \ldots \ldots v_{n}\right)$. where $\tau$ is an $n$-ary term. is $\tau\left[\ddot{c}^{\prime}\right]\left(d_{1} \leftarrow v_{1} \ldots \ldots d_{n} \leftarrow v_{n}\right)$.
j. The $n$-ary relational terms $\mu Y . \tau$ and $\nu Y$. $\tau$ are evaluated using the standard fixed point algorithm.

Because the variables are Boolean valued, we can implement all of the above using the operations of QBF. with OBDDs as our representation. The symbolic Mu-C'alculus model checking algorithm is shown in pseudo-code form in figure 2.15 . Using this algorithm. any quantity that can be characterized in the Mu-Calculus can be computed using the symbolic model checking technique, with the possibility that a combinational explosion can be reduced or avoided. This also allows us to use the expressive powers of the Mu-Calculus in describing and manipulating symbolic algorithms, with the understanding that the translation from Mu-('alculus to a symbolic algorithm is merely mechanical.

### 2.7 Computing equivalence relations

In this section. we consider the problem of computing a symbolic representation of the equivalence relation between the states of two finite state machines. or between states of the same machine. In the former rase. the relation can be used to determine the erguivalence of the two machines. while in the latter case. as Lin it al. have observed [LTV90]. the self equivalence relation can be used in optimizing the logic or register nsage of the machine.

```
function \(\operatorname{eval}\left(\tau . w^{\prime}\right)\)
    case
        \(\tau\) a relational variable: return \((\tau)\)
        \(r=\neg p\) : return \(-\operatorname{ral}(p, \because)\)
        \(\tau=p \vee q\) : return eval \(\left(p, \iota^{\cdot}\right) \vee\) eval \(\left(q, \iota^{\circ}\right)\)
        \(\tau=\exists w . q:\) return \(\exists w . \operatorname{eval}\left(p, v^{\prime}\right)\)
        \(\tau=\mu Y . p\) : return fixedpoint \((Y . p, \nu(Y-\) false \()\) )
        \(\tau=\nu Y: p:\) return fixedpoint \((Y, p, v(Y-\operatorname{true}))\)
    end case
end function
function fixedpoint(Y. \(\mathrm{P} . \mathfrak{C l}^{\prime}\) )
    \(Y^{\prime \prime}=\operatorname{eval}(p . \because)\)
    if \(Y^{\prime \prime}=!\cdot(Y)\) then return \(\zeta^{\prime \prime}\)
```



```
end function
```

Figure 2.f.s: Simbotic Mu-('alculus morlel thecking algorithm:

### 2.7.1 State equivalence

We use a standard notion of the equivalence of states of finite Mealy machines. Two states are equivalent if and only if for all input sequences. they yield the same output sequence. The following is an alternate characterization: equivalence is the greatest relation between states such that if $x$ is equivalent to $y$, then for all inputs. the output in state $r$ is equal to the output in state $y$. and the successor state of $x$ is equivaleut to the successor state of $y$. Let $\delta(x, z)$ be the function which determines the next state, as a function of current state $x$ and current input $z$. and let $f(x, z)$, be the function that determines the current output. In the Mu-Calculus. the equivalence relation $R_{\nu}$ is

$$
\begin{equation*}
R_{\nu}=\nu R . \lambda . x, y . \forall z \cdot(\neg(x, z)=\gamma(y, z) \wedge R(\delta(x, z), \delta(y, z))) \tag{2.22}
\end{equation*}
$$

Using the standard fixed point approach. we can evaluate this relation by a sequence of approximations $R_{0}, R_{1} \ldots \ldots$ where $R_{i}$ is the set of state pairs which are equivalent for all input sequences of length $i$. This sequence is characterized by the recurrence

$$
\begin{equation*}
R_{1}=\lambda \cdot x, y \cdot \forall z \cdot(\gamma(x, z)=\gamma(y, z)) \tag{2.23}
\end{equation*}
$$

and

$$
\begin{equation*}
R_{i+1}=\lambda \cdot x \cdot y \cdot \forall z \cdot\left(R_{i}(x, y) \wedge R_{i}(\delta(x, z) \cdot \delta(y, z))\right) \tag{2.24}
\end{equation*}
$$

This is simply the standard $O\left(n^{2}\right)$ algorithm for computing state equivalence of Mealy machines. The problem of determining whether two Mealy machines are equivalent in their initial states can be approached in two wass either their equivalence relation can be computed. or the state space of their product can be exhausted by a forward search. The number of iterations required for the former approach can be substantially less. however. In the trivial case of an $n$-bit counter. the number of iterations in the forward search is is exponential in $n$. while one step suffices to reach a fixed point in the equivalence calculation. since all states are distinguished by their outputs.

It is immerliately seen that the crucial step in calculation 2.21 is the substitution of vector functions $\delta(x, z)$ and $\partial(y, z)$ into $R$. The most obvious way to accomplish this is to use Bryant's ('omposf algorithm. Somenther possible methods are introduced in this section. ('omputing
the OBDD representation for a composition of functions is an NP-hard problem (cf., section 2.3.t). thus we expect no good general solutions to the problem. Inother tractability issue is whether the approximations to the equivalence relation can be compactly represented using OBDDs. There is no guarantee of this, of course. but there is some reason to believe. a priori, that it may often be the case. First of all. for single-machine equivalence. if all distinct states are distinguishable. then the equivalence relation is the identity relation. which can be represented as by a linear size ()BDD. provided the component variables of $x$ and $y$ are interleaved in the variable ordering. It also seems plausible that the equivalence relation will otten be simply a logical conjunction of independent relations, each corresponding to some morlular component of the system. In this case, the OBDD representation will also be compact. provided the variable ordering conforms to the modular structure of the machine. In any case. we will see examples of fairly complex machines whose equivalence relations are expressed compactly in OBDD form.

## Algorithm using restrictions

Because of the basic difficulty of computing compositions of OBDDs. it is useful to have some restrictions on the result in order to be able to solve the problem. Fortumately, the decreasing nature of the series of approximations defined in $2.2 t$ provides a constraint on the result of the substitution. since each approximation $R_{t+1}$ is strictly contained in $R$. We can use this fact be rewriting $2.2 t$ as
 and Berthet (CBMSO). This operation produces a function which agrees with $R_{,}\left(\delta\left(x .=1 . \delta(y .=1)\right.\right.$ wer the set $R_{\text {}}$. attempting to minimize the OBDD) size. The restriction can be nsed to varying adrantage. depending on the algorithm nsed for substitution.

## Iterative abstraction algorithnı

Another way to provide a restriction on the equivalence relation is tiret tolime the ernitalenceretation of all abst racted marhine. Wie cheose the
abstraction in such a way that the equivalence relation of the abstract machine is strictly weaker than the equivalence relation of the original machine. Thus, we can compute the equivalence relation of the abstract machine first. and use it as a restriction in computing the equivalence relation of the original machine. In particular, we can abstract the machine by choosing a subset $V$ of the state variables, and at each approximation quantifying out the remaining variables existentially. That is. let $V^{\prime \prime}$ be the complement of $V$, and let

$$
\begin{equation*}
R_{1}^{V}=\lambda x, y, \exists V^{c} \cdot \forall z \cdot(\gamma(x, z)=\gamma(y, z)) \tag{2.26}
\end{equation*}
$$

and

$$
\begin{equation*}
R_{i+1}^{V}=,, x, y, \exists V^{c}, \forall z\left(R_{i}^{V}(x, y) \wedge\left(R_{i}^{V}(\delta(x, z), \delta(y, z)) \downarrow R_{i}^{V}\right)\right) \tag{2.27}
\end{equation*}
$$

It is trivial to see that each approximation in the series $R^{\mathfrak{V}}$ is strictly weaker than the corresponding approximation in $R$. It follows that $R_{\lrcorner}^{V}$, the greatest fixed point. is weaker than $R_{\omega}$. Therefore, we can restrict the entire calculation of $R_{\nu}$ to only those state pairs satisfying $R_{\lrcorner}^{V}$. In addition. we can use a series of subsets $V_{1} \subset V_{2} \subset \cdots \subset V_{k}$ where $V_{k}$ is the set of all state variables. restricting the first approximation in each series $R^{l_{i}}$ to the equivalence relation for the previous subset. Thus, we let

$$
\begin{equation*}
R_{1}^{V_{j}}=R_{\mu}^{V_{j-1}} \wedge \lambda x \cdot y \cdot \exists V_{j}^{c} \cdot \forall z \cdot(\gamma(x, z)=\gamma(y, z)) \tag{2.28}
\end{equation*}
$$

and

We will refer to this as the iterative abstraction algorithm for computing the erpitalence relation. By adding only a few variables to each successive subset. we can in some cases obtain fairly strong restriction. which allows the substitution to be computed more efficiently. In other cases the equivalence relation obtained for the abstracted machine may be trivial. since abstracting out the variables in $1^{\circ}$ may result in all machine states appearing equivalent at the outputs. This is especially likely if the abstracted variables hold important control information that emables machine registers to be observed at the outputs. . Vonetheless. we can show cases where this incremental approach is greatly more efficient than the basic algorithm.

### 2.7.2 Methods for functional composition

This section considers methods for substituting functions for variables in OBDDs. This operation is referred to by Bryant as Compose. It is the syntactic mechanism corresponding to functional composition. As nuch. it has a number of applications apart from finding the equivalence relation of finite state machines, including the evaluation of ("TL formulas (BF89bj. Most of the algorithms presented here for this purpose have been modified to take as an extra argument a restriction on the reoult. in the hope that efficiency can be obtained by combining these two operations. We consider the problem of calculating $f\left(g_{1} \ldots \ldots g_{n}\right) \downarrow R$. where $f . g_{1} \ldots \ldots g_{n}$ and $R$ are all Boolean functions.

## "bottom-up" substitution

This is the method originally proposed by Bryant for his Compose algorithm. but with a restriction on the result. In this method. we view each OBDD node in $f$ as a gate. which computes the function "if $v_{2}$ then $h$ else $l "$ or equivalently, $\left(\neg v_{i} \wedge l\right) \vee\left(v_{i} \wedge h\right)$. Having substituted the functions $g_{1} \ldots, g_{n}$ for the variables in $l$ and $h$, we can then compute the result for $f$ mising the standard $\vee$ and $\wedge$ operators. The basic bottom-up algorithm is

```
function bettom-1p( f. IV)
    if f}\mathrm{ is a leat then retum f
    if botwom-up(f.l) has abradw been solved then remurn old wolution
    else ff is a triple (r.0.fiffolt
        I}=\mathrm{ bottom-mpl fl. R 
    h=bottom-up( / /h. 盾)
    return|-y, \/1!!:h):R
end
```

Xote that the restriction uperator is nsed at each step 10 simplifs the rocult. Sime rath mbpeothem is solved onle moe the mumber of


## Domain partitioning

The domain partitioning strategy is so named because it divides the problem into two subproblems by partitioning the domain of the functions $y_{1} \ldots \ldots g_{n}$ according to the value of one of the variables. The operation proceeds in several steps.

First. we observe that if any of $g_{1} \ldots \ldots g_{n}$ are constants, we can immediately substitute these values into $f$, since substitution by a constant is a linear time operation which can only reduce the size of the OBDD. We use the fact that if $g_{i}=c$. where $c$ is 0 or 1 . then

$$
\begin{equation*}
f\left(g_{1}, \ldots g_{n}\right)=f\left(v_{i} \leftarrow c\right)\left(g_{1} \ldots \ldots g_{n}\right) \tag{2.30}
\end{equation*}
$$

Vext. we olserve that we can eliminate any argument position on which $f$ loes not depend. thus obtaining a smaller problem with the same result. We can determine the set of variables on which $f$ depends in linear time. since $f$ clepends on $v_{i}$ if and only if $c_{i}$ appears in some node in $f$.

If at this point the function $f$ has been reduced to a constant. we are done. Otherwise, we split the problem into two cases and recurse. We choose the first variable $v_{i}$ occurring in $g_{1} \ldots \ldots g_{n}$, and apply Shannon’s expansion, obtaining two subproblems

$$
\begin{aligned}
l & =f\left(g_{1}\left(c_{i} \leftarrow 0\right) \ldots \ldots g_{n}\left(c_{t} \leftarrow 0\right)\right) \\
h & =f\left(g_{1}\left(c_{t} \leftarrow 1\right) \ldots \ldots g_{n}\left(r_{i} \leftarrow 1\right)\right.
\end{aligned}
$$

Is in wher (OBDD algorithms, the result is an OBDD $r=(r, l . h)$. provided $l \neq h$. otherwise $r=l=h$. Veedtess to say, we use a hash table. caching the results of subproblems so that the same subproblem is not olved twice. With caching, the complexity of the aigorithm is ( $\left.0|f| \times \prod_{!} \mid\right)$.

Making use of the restriction $R$ in this algorithm is straightforward. If $R=0$. the result can be any function at all. so we simply return 0 . Each time we partition the problem into subproblems, we also split $R$ into two cases. $B\left(r_{1}-0\right)$ and $R\left(r_{1}-1\right)$. The restriction has the effect of cultime off the recursion each time a 0 leat is reached in $R$.

## Sequential substitution

This is perhaps the simplest approach to substitution; it transforms a simultaneous substitution problem into a sequence of substitutions. This is done by replacing each variable $v_{i}$ in the OBDD for $f$ by a new variable $v_{i}^{\prime}$. Having done this. it is safe to perform the substitutions of each function $g_{i}$ for $v_{i}^{\prime}$ in any order. since none of the functions $g_{1}, \ldots g_{n}$ depends on any variable being substituted. Substitution of a function for a single variable can be accomplished as follows:

$$
\begin{equation*}
f\left(v_{i}^{\prime} \leftarrow g_{i}\right)=\exists v_{i}^{\prime} \cdot\left[\left(v_{i}^{\prime} \Longleftrightarrow g_{i}\right) \wedge f\right] \tag{2.31}
\end{equation*}
$$

This approach can also make effective use of a restriction. The restriction operator operator may in fact be applied after each substitution step if clesired. potentially reducing the size of the intermediate results. In the case of the iterative abstraction algorithm. the fact that some of the variables in the result will later be quantified out existentially can also be put to use. We can move the existential quantifiers for these variables inside the conjunction, thus quantifying the abstracted variables out of the term ( $v_{i}^{\prime} \Longleftrightarrow g_{i}$ ) before applying the conjunction. This may weaken our result somewhat. since $[\exists x . a] \wedge[\exists x . b]$ is weaker than $\exists x .[a \wedge b]$, but it can produce significant reductions in the size of the intermediate results. The final result of the equivalence algorithm is unchanged. since it is computed with no variables abstracted.

### 2.7.3 Experimental results

This section presents the results of applying the various equivalence relation algorithms to several example state machines, with a range of complexity. The results are compared to published results for the same circuits by Tonati t al. (computing only the reachable states) and Lin et al. (computing the equivalence relation). In all cases. it is self-equivalence that is calculated. It would be interesting to have some results in this sertion on ralculating the state equivalence relation between two different implementations of a given machine. but unforthately such examples were lacking. The three different approarhes to OBDD substitution are compared. for each example. Where posnible the direct aloorithm is need. otherwise. the iterative abstraction

| machine | mtd | result <br> (nodes) | b-u <br> (secs) | d-p <br> (secs) | seq <br> (secs) | Touati <br> (secs) | Lin <br> (secs) |
| :--- | :---: | ---: | ---: | ---: | ---: | ---: | ---: |
| sbc | iter | 361 | 2054 | $>10 \mathrm{~K}$ | 2415 | 2903 |  |
| cpb32 | iter | 95 | 4.5 .5 | 14.4 | 54.4 | 14.1 | 12.10 |
| key | iter | 167 | 342 | 17.38 | 884 | 5706 | 17.5 .20 |
| minmax10 | dir | 89 | 197 | 2.55 | 204 |  |  |
| minmax.30 | dir | 59 | 3.0 | 4.5 | 6.7 |  |  |
| minmax 30 | dir | 89 | 6.2 | 9.0 | 15.7 |  |  |

Table 2.1: Equivalence calculation times
algorithm is used. For example, the state equivalence relation for the machine sbc was calculated using iteration. but could not be calculated directly. Table 2.1 gives the total execution time in seconds, while table 2.2 gives the total number of OBDD nodes used. ${ }^{17}$ The latter numbers are not very reliable. since they depend to some extent on arbitrary choices about when to scavenge unused cells and cache entries. However, if the available memory limit of 190.000 nodes is exceeded, it is certain that all of the nodes in use were necessary for the computation. since all available nodes were scavenged when the memory limit was reached. The columns give the following information: the name of the circuit. the method used (direct or iterative) the size of the equivalence relation. and the time or space needed for each of the three substitution methods (bottom-up. domain partitioning, and sequential). The times are for a LISP implementation running on a 1 MIP minicomputer. The final two columns give the results obtained by Tonati +1 al. and Lin et al. for the same circuit. These times are for (' language implementations running on a DEC. 5400 and IBM R6000 respectivel:

It would seem that for the circuits cpb. key and minmax. which have regular structures with no control registers. there is no clear choice as to

[^14]| machine | mtd | result <br> (nodes) | b-u <br> (nodes) | d-p <br> (nodes) | seq <br> (nodes) |
| :--- | :---: | ---: | ---: | ---: | ---: |
| sbc | iter | 361 | 22184 |  | 34609 |
| cpb32 | iter | 95 | 8202 | 4225 | 11295 |
| key | iter | 167 | 13328 | 24868 | 11563 |
| minmax10 | dir | 89 | 16589 | 17815 | 17566 |
| minmax20 | dir | .99 | 8538 | 3492 | 9190 |
| minmax30 | dir | 39 | 11952 | 11886 | 10001 |

Table 2.2: Equivalence calculation space
which substitution algorithm is best. The bottom-up algorithm tends to provide the best performance with the least memory usage. but there are a number of exceptions. The machine sbc, which is somewhat more complex, is a more interesting case. Here bottom-up and sequential both provide fairly efficient solutions, although the iterative method was required in both cases to solve the problem. The domain partitioning approach fails to terminate after 10,000 seconds. In the first stage of the iterative algorithm, domain partitioning produced over 100.000 subproblems for a final result of approximately 100 nodes. Obviously. many different subproblems with identical results are being solved. The difficulty is that there is no easy way to identify equivalent problems. It is worth mentioning the the limit on the size of the rache for this method was 5000 entries. With an unbounded cache. the performance of the algorithm may be much better (a matter of theoretical interest at best. since an unbounded cache cannot be provided). It should also be noted that the results for minmax are somewhat anomalous. since the 10 -bit version seems to be substantially more complex than the 30 and 30 -bit cases. This is explained by the fact that the output functions of these clifferent versions were not the same. In the $20-$ and $30-$ bit versions. the outputs appear to depend only on the "last" register. and not the "min" and "max" registers. It is also interesting to observe that for minmax 10 . not all of the states are distinguishable. that is. the equivalence relation is not the identity.

Comparing these results to those of Touati et al., it is interesting to note that the self-equivalence relation can be computed in less time than the reachable states for sbc and key (taking into account the difference in machines speeds of roughly a factor 10 , the equivalence method seems to be about one order of magnitude faster for sbc, and two orders of magnitude faster for key). Of course, the information obtained by the two methods is not the same. It seems. however. that in some cases where the set of reachable states is not obtainable, the equivalence computation may still provide useful information for logic optimization. The results of Lin $e t$ al. seem to be roughly comparable for the machines key and cpb32 (again, taking into account the difference in machine speeds). It is not clear from the Lin et al. article which substitution method was used. since two were mentioned. The one benchmark for which the iterative method was required to produce a result was sbc, but unfortunately Lin et al. do not report a figure for this machine. Also. because of the fact the the 20 - and 30 -bit versions of minmax had modified output functions. it is not possible to compare figures for this benchmark. As a result of these ambiguities. it difficult to draw conclusions about the effectiveness of the iterative abstraction method, except to say that in one case (sbc) it was the only method that successfully computed the equivalence relation.

### 2.8 Related research

The author first experimented with the use of OBDDs to represent sets of states and transition relations in 1987 . building the first symbolic model checker for CTL. Various heuristic improvements to the basic technique were developed. including the OBDD algorithm combining existential quantification and conjunction (cf. section 2.3.4.4) and the technique of early quantification for disjunctive transition relations (cf. section 3..4.2). Extending this work. Burch. C'larke. Long, Mc:Millan. Dill and Hwang described a symbolic model checking procedure for the propositional Mu-C'alculus. which could be used for a variety of purposes. including CTL model checking, testing various process equivalences. testing language containment of $\omega$-automata. and checking satistiability of LTL formulas $\left[\mathrm{BCM}^{+} 90\right]$. In 1989 , the author used

### 2.8. RELATED RESEARCH

the model checking technique to verify the cache consistency protocol of the Encore Gigamax multiprocessor (see chapter 4). In the process, a model checking system called SMV was developed, along with an associated description language (see chapter 3).

In 1989, the idea of using the OBDD representation for verification of finite state machines appears to have been independently developed by Coudert, Madre and Berthet [CBM89], who used it in their PRIA.M system for testing equivalence of finite state machines. They represent a finite state machine by a pair of vector Boolean functions. The function $\delta(v, w)$ yields the next state vector as a function of the current state vector $v$ and the input vector $w$. The function $\lambda(v, w)$ yields the output vector as a function of $v$ and $w$. The equivalence of two state machines is tested by creating a combined machine in which both machines receive the same input vector. and the output is a single bit which is true if and only if the output vectors of the two machines are equal. The reachable states of this combined machine are computed. If in all reachable states the output is true, the two machines are equivalent, since no input sequence can produce differing output sequences from the two machines.

The set of reached states is computed as the limit of an increasing series of approximations. starting with the initial state. The set of states reachable in one step from a set $S$ is computed by a function called Imag, where $\operatorname{Imag}(\delta . S)=\{s \mid \exists v, w: v \in S, \delta(v, w)=s\}$. Most of Coudert. Madre and Berthet's efforts are applied to computing the Imag function without resort to representing the transition relation as an OBDD, which they claim is generally intractable. Their approach begins by reducing the problem of computing the image of a set via a function. to computing the range of a function. This is done using an OBDD operation called Constrain. The Constran operator takes two Boolean functions $f$ and $g$. and returns a function $f^{\prime}=$ (onstrainf $f . g$ with the following property: for all $x^{\prime} . f^{\prime}\left(x^{\prime}\right)=f(x)$, where $x$ is the nearest Boolean vector to $x^{\prime}$ (according to a suitable distance metric) such that $g(x)=1$. If we let $\delta^{\prime}=$ Constraini $\left.\delta . S^{\prime}\right)$. then the image of $i$ via $\delta$ is just the range of $\delta^{\prime}$.

Coudert and Madre suggest two methods for computing the range of $n^{\prime}$. The first is called range partitioning. In this approach. we pick the lowest remaining variable in the ordering (call it $v_{i}$ ), and. and divide the problem into two subproblems, depending on the out put of function
$\delta_{i}^{\prime}$. Thus.

$$
\begin{aligned}
& \left(\operatorname{Range}\left(\delta^{\prime}\right)\right)\left(v_{i} \leftarrow 0\right)=\operatorname{Range}\left(\operatorname{Constrain}\left(\delta^{\prime}, \neg \delta_{i}^{\prime}\right)\right) \\
& \left(\operatorname{Range}\left(\delta^{\prime}\right)\right)\left(v_{i} \leftarrow 1\right)=\operatorname{Range}\left(\operatorname{Constrain}\left(\delta^{\prime}, \delta_{i}^{\prime}\right)\right)
\end{aligned}
$$

Note that for any function $f$,

$$
\begin{aligned}
\operatorname{Constrain}(f, f) & =1 \text { and } \\
\operatorname{Constrain}(f, \neg f) & =0
\end{aligned}
$$

so each recursion effectively eliminates one component function of $\delta^{\prime}$. The recursion terminates when all of the components of $\delta^{\prime}$ are constants.

The other approach. called domain partitioning, is to divide into subproblems based on the value of one of the inputs to $\dot{\partial}^{\prime}$. Thus.

$$
\operatorname{Range}\left(\delta^{\prime}\right)=\operatorname{Range}\left(\delta^{\prime}\left(v_{i}-0\right)\right) \vee \operatorname{Range}\left(\delta^{\prime}\left(v_{i}-1\right)\right)
$$

Again, the recursion terminates when all of the components of $\delta^{\prime}$ are constants.

Both of these strategies are special cases of a general strategy where one chooses a cover. which is a pair of functions $h_{1}$ and $h_{2}$ such that $h_{1} \vee h_{2}=1$. and then computes the recursion

$$
\left.\operatorname{Ranye}\left(\delta^{\prime}\right)=\operatorname{Range}\left(\text { Constrain }\left(\delta^{\prime}, h_{1}\right)\right) \vee \text { Rangel Constrain }\left(\delta^{\prime} . h_{2}\right)\right)
$$

In the case of range partitioning. $h_{1}=0_{0}^{\prime}$ and $h_{2}=-力_{:}^{\prime}$. In the case of domain partitioning, $h_{1}=v_{i}$ and $h_{2}=-v_{i}$. It is suggested that other covers mav be usefiul as well. As with other OBDD techniques, a table of pairs ( $x^{\prime}$. Runge( $\delta^{\prime}$ )) is kept to avoid solving the same subproblem twice. This table is not as effective as the in the case of the standard OBDD operations. however. since the number of possible subproblems is exponential in the number of state variables. ('oudert and Madre suggest several optimizations for increasing the hit rate in this table.

A firther optimization introduced b" ('ondert and Madre is to use an OBDD finction called Restrict to reduce the size of the reached state set before appiving the Imay operator. The Restrict operator takes two functions $f$ and $g$. and produces a function $f^{\prime}=\operatorname{Restract} f . y$ ) such that for all $x$. if $\boldsymbol{y}(x)=1$, then $f^{\prime}(x)=f(x)$. otherwise $f^{\prime}(x)$ is arbitrary.

Usually (but not always), the size of $f^{\prime}$ is less than the size of $f$. We note that if $R_{i}$ is the set of states reachable after $i$ steps of the machine, then

$$
\begin{aligned}
R_{i+1} & =R_{i} \vee \operatorname{Imag}\left(\delta, R_{i}\right) \\
& =R_{i} \vee \operatorname{Imag}\left(\delta, \operatorname{Restrict}\left(R_{i}, \neg R_{i-1}\right)\right)
\end{aligned}
$$

As a result, the size of the arguments of Imagcan sometimes be reduced using Restrict.

Coudert and Madre report experimental results for computation of the set of reachable states for a variety of small sequential circuits (mostly ISCAS ${ }^{18}$ sequential benchmark circuits). Computing the set of reached states can be useful for generating test patterns or "don"t care" condition for logic optimization [TSL ${ }^{+} 90$ ]. Unfortunately. Coudert and Madre do not use their techniques to actually test the equivalence of two state machines, so it is unknown whether the technique is useful for this purpose. They have not studied the asymptotic performance of their techniques for classes of circuits, so it is not possible to determine whether their optimizations yield asymptotic improvements.

A variant on the symbolic model checking technique for CTL was proposed by Bose and Fisher [BF89b]. Their technique, which is limited to deterministic finite state machines, also represents the transition relation of the machine by a vector of Boolean functions $\delta$. and uses Bryant’s Compose operation to compute EXp $=p\left(v_{i} \leftarrow\right.$ $n$ :l. They do not report experimental results using this technique for practical circuits. A similar terhnique was proposed by Coudert and Madre CMB91!.

Other researchers have proposed techniques to avoid constructing the transition relation. For example, Burch. Clarke and Long use early quantification (cf. section 2.4.2) for both disjunctive and conjunctive transition relations (BC'LOLb. BC'L')laj. They use the term "partitioned ransition relations." for this. The technique is somewhat limited in the rase of conjunctive transition relations. because existential quantifica:ion only distributes over conjunction in the special case when one of the conjuncts does not depend on the variable being quantified. Vev-

[^15]ertheless, there are cases where the support of the component relations is sufficiently disjoint to make this technique effective.

The basic technique is the following: assume we wish to compute $\exists v . \wedge_{i} f_{i}$, where $v=\left(v_{1}, \ldots, v_{k}\right)$ is a vector of variables and $f=\left(f_{1}, \ldots, f_{m}\right)$ is a vector of Boolean functions. Since conjunction is associative and commutative, we can combine these functions in any order we choose. In addition, if at any time there is a variable occurring in only one function, we can quantify that variable out, since $\exists w .(p \wedge q)$ is equivalent to $(\exists w . p) \wedge q$ when $q$ does not depend on $w$. Since quantification tends to reduce OBDD size by reducing the number of variables. the strategy is to combine the functions in such an order that variables can be quantified out as soon as possible.

Burch Clarke and Long use a fixed order determined by the user for combining the functions. They show that this is quite effective for pipelined data path circuits. and an asynchronous stack circuit, improving the asymptotic performance as the circuit size increases. For the DME circuit. the asymptotic performance of this method was not as good as a method using a disjunctive transition relation. but it can be more efficient for small rings. ${ }^{19}$ It was found most efficient to group the components of the transition relation and combine each group in advance, thus avoiding some computation at each step.

For disjunctive transition relations (interleaving models). Burch. Clarke and Long introduce a modified search order that tends to reduce the representation of the reached state set. In a breadth first search. the representation of this set is complicated by the fact that the after $n$ steps. the number of steps taken by each process is constrained to sum to $n$. This produces an artificial correlation between the states of otherwise independent processes (cf. section 2.1.2). To counter this. one can modify the search order. searching first all of the states reachable by transitions of one subset of the system processes, then the next. and repeating this process until a fixed point is reached. This technique. called "modified breadth first search". was effective in reducing the OBDDs representing the reached state sets for an asynchronous stack circuit. but was found not to be as effective as the "conjunctive partitioning" method. For the DME circuit. the modified breadth first

[^16]
### 2.8. RELATED RESEARCH

search method was faster up to about 16 cells, but had slower asymptotic performance. The grouping of processes into subsets was manual.

Another OBDD based technique for computing the reachable states of a machine was introduced by Touati et al. [TSL +90 ]. They also use a conjunction of component relations to represent the transition relation, along with early quantification. However, they combine this technique with the Constrain operation of Coudert et al.. This reduces the problem of computing the image of a set via a relation to that of computing the codomain of a relation. A series of approximations $A_{i}$ to the reachable states is computed, such that

$$
A_{i+1}=A_{i} \vee \lambda y . \exists x .\left(\bigwedge_{j} \operatorname{Constrain}\left(R_{j}, A_{i}\right)\right)(x, y)
$$

where $R$ is a vector of component relations, each relation determining the new state of one state variable. Touati et al. find this technique to be superior to using the transition relation directly and to using the Imag operation of Coudert et al. for computing the reachable states of the benchmark circuits minmax and sbc, somewhat slower for key, and roughly the same for cpb.32.4. It would be interesting to know for the cases where an improvement was obtained, how much was due to the use of Constrain and how much to the use of early quantification. Touati et al. have also suggested partitioning complex next-state functions into the composition of a sequence of smaller functions. This could be useful for circuits containing multipliers, or other functions which have no compact OBDD representation.

Touati. Brayton and Kurshan report a technique for testing language containment of $\dot{\mu}$-automata using OBDDs [TBK91]. They use the L-automaton model of Kurshan [Kur86], and an algorithm similar to the one described in section 2.6 .1 using the transitive closure of the transition relation. No experimental results using this technique are available.

Another way that one can test equivalence of two finite state machines is by computing the equivalence relation on states. as described in section 2.7. Lin et al. also describe OBDD based algorithms for computing this relation [LT, $90 \mid$. A comparison of the methods can be found in section 2.7. Lin et al. describe how the equivalence relation can be used for computing "don't care" conditions for logic optimiza-
tion. In a later paper [LN91], Lin shows how this relation (represented as an OBDD) can be used for state minimization, using an operator which takes an equivalence relation and returns a relation which maps every state onto the least element of its equivalence class.

Bryant and Seger have taken an an approach to formal verification using OBDDs based on symbolic simulation [Bry88, BBS90. BS90]. The symbolic simulator is similar to an ordinary logic simulator, except that the inputs are symbolic values (variables) rather than numeric values, and the outputs are given as symbolic functions in terms of these variables. These functions are represented by OBDDs. The simulation method gains a great deal in efficiency by using an abstract interpretation of the circuit model. This abstraction uses a lattice consisting of the three values 0,1 and X . where X is the least upper bound of 0 and 1 . The circuit operations such as AND and OR are abstracted in such a way as to be monotonic with respect to this lattice. Therefore. the output of the abstract simulation is always an upper bound on the output of the concrete simulation. In many cases, a large number of the inputs and initial values of state variables can be replaced by X without sacrificing the particular circuit property being proved. The art in this technique is to decompose the specification in such a way that each part can be verified using only a small number of symbolic values and X everywhere else. The simulation technique is limited to a logic with only next-time operators. These formulas can be verified using symbolic simulations of finite execution sequences. This rules out proving properties such as liveness. fairness or deadlock freedom. but allows satety properties to be proved using invariants.

Bose and Fisher have demonstrated a technique for using representation functions to verify sequential circuits using OBDDs [BF39a]. A representation function maps each state of the implementation to $=$ state of the specification (which is also a circuit). Symbolic simulation techniques can be used to show a kind of single step equivalence between the implementation and specification vis it vis this relation. As in the method of Bryant and Seger, this proof can be decomposed into parts in such a way that each part requires only a small number of symbolic variables. with the remaining circuit nodes initialized to X. Typically. an invariant is also required. since the single step equivalence only holds over the reachable state space of the implementation. This
technique is also limited in that it cannot prove liveness or deadlock properties.

Long and Grumberg have introduced an abstraction technique using OBDDs which is more general than simply introducing $X$ values [CGL92]. Their technique uses an OBDD to express the relation between the abstract and concrete domains. The abstract transition relation is automatically derived using OBDD techniques from the concrete transition relation. This can be done in a compositional way to reduce the number of symbolic variables that are required. A variety of abstractions have been put to use in this way. For example, a binary number can be represented by its remainders modulo a set of relatively prime numbers. This has allowed the use of the Chinese remainder theorem to prove the correctness of a multiplier circuit. In another case. a single bit was used to represent whether a given binary number in a circuit is equal to a given symbolic binary value. In this way the entire function of the arithmetic unit was abstracted away, allowing a data pipeline circuit with 6464 -bit registers to be verified. This abstraction technique is quite general. and is closely related to more classical abstraction techniques [Kur37]. The difference is that function graph methods are used to actually compute the abstract transition relation. rather than giving this relation a priori.

## Chapter 3

## The SMV system

The SMV system is a tool for checking finite state systems against specifications in the temporal logic CTL. The input language of SMVV is designed to allow the description of finite state systems that range from completely synchronous to completely asynchronous, and from the detailed to the abstract. One can readily specify a system as a synchronous Mealy machine, or as an asynchronous network of abstract, nondeterministic processes. The language provides for modular hierarchical descriptions, and for the definition of reusable components. Since it is intended to describe finite state machines, the only basic data types in the language are finite scalar types. Static, structured data types can also be constructed. The logic CTL allows a rich class of temporal properties, including safety. liveness, fairness and deadlock freedom. to be specified in a concise syntax. SMV uses the OBDD-based symbolic model checking algorithm to efficiently determine whether specifications expressed in CTL are satisfied.

The primary purpose of the SMV input language is to provide a symbolic description of the transition relation of a finite Kripke structure. Any propositional formula can be used to describe this relation. This provides a great deal of flexibility, and at the same time a certain danger of inconsistency. For example. the presence of a logical contradiction can result in a deadlock - a state or states with no successor. This can make some specifications vacuously true. and makes the description unimplementable. While the model checking process can be used to check for deadlocks, it is best to avoid the problem
when possible by using a restricted description style. The SMV system supports this by providing a parallel-assignment syntax. The semantics of assignment in SMV is similar to that of single assignment data flow languages. A program can be viewed as a system of simultaneous equations, whose solutions determine the next state. By checking programs for multiple assignments to the same variable, circular dependencies, and type errors. the compiler insures that a program using only the assignment mechanism is implementable. Consequently, this fragment of the language can be viewed as a hardware description language, or a programming language. The SMV system is by no means the last word on symbolic model checking techniques, nor is it intended to be a complete hardware description language. It is simply an experimental tool for exploring the possible applications of symbolic model checking to hardware verification.

### 3.1 An informal introduction

Before delving into the syntax and semantics of the language, let us first consider a few simple examples that illustrate the basic concepts. Consider the following short program in the language.

```
MODULE main
VAR
    request : boolean;
    state : {ready,busy};
ASSIGN
    init(state) := ready;
    next(state) := case
                                state = ready & request : busy;
                                1 : {ready,busy};
            esac;
SPEC
    AG(request -) AF state = busy)
```

The input file describes both the model and the specification. The model is a Kripke structure, whose state is defined by a collection of state variables. which may be of Boolean or scalar type. The variable
request is declared to be a Boolean in the above program, while the variable state is a scalar. which can take on the symbolic values ready or busy. The value of a scalar variable is encoded by the compiler using a collection of Boolean variables. so that the transition relation may be represented by an OBDD. This encoding is invisible to the user. however.

The transition relation of the Kiripke structure, and its initial state (or states), are determined by a collection of parallel assignments (a system of simultaneous equations), which are introduced by the keyword ASSIGN. In the above program, the initial value of the variable state is set to ready. The next value of state is determined by the current state of the system by assigning it the value of the expression

```
case
    state = ready & request : busy;
    1 : {ready,busy};
esac;
```

The value of a case expression is determined by the first expression on the right hand side of a (:) such that the condition on the left hand side is true. Thus, if state $=$ ready \& request is true. then the result of the expression is busy. otherwise, it is the set \{ready, busy\}. When a set is assigned to a variable. the result is a non-deterministic choice among the values in the set. Thus. if the value of status is not ready. or request is false (in the current state). the value of state in the next state can be either ready or busy. .Von-deterministic choices are useful for describing systems which are not yet fully implemented (ie.. where some design choices are left to the implementor), or abstract models of complex protocols. where the value of some state variables cannot be completely determined.
. Notice that the variable request is not assigned in this program. This leaves the SMV system tree to choose any value for this variable. giving it the characteristics of an unconstrained input to the system.

The specification of the system appears as a formula in CTL under the keyword SPEC. The SWV model checker verifies that all possible initial states satisty the specification. In this case. the specification is that invariantly if request is true. then inevitably the value of state is busy.

The following program illustrates the definition of reusable modules and expressions. It is a model of a 3 bit binary counter circuit. Notice that the module name "main" has special meaning in SMV, in the same way that it does in the C programming language. The order of module definitions in the input file is inconsequential.

```
MODULE main
VAR
    bitO : counter_cell(1);
    bit1 : counter_cell(bit0.carry_out);
    bit2 : counter_cell(bit1.carry_out);
SPEC
    AG AF bit2.carry_out
MODULE counter_cell(carry_in)
VAR
    value : boolean;
ASSIGN
    init(value) := 0;
    next(value) := value + carry_in mod 2;
DEFINE
    carry_out := value & carry_in;
```

In this example, we see that a variable can also be an instance of a user defined module. The module in this case is counter_cell, which is instantiated three times, with the names bit0, bit1 and bit2. The counter cell module has one formal parameter carry_in. In the instance bit0, this formal parameter is given the actual value 1. In the instance bit1. carryin is given the value of the expression bito.carry-out. This expression is evaluated in the context of the main module. However, an expression of the form $a . b$ denotes component $b$ of module $a$. just as if the module $a$ were a data structure in a standard programming language. Hence, the carry_in of module bit1 is the carry_out of module bito. The keyword DEFINE is used to assign the expression value \& carry-in to the symbol carry_out. Definitions of this type are useful for describing Mealy machines. They are analogous to macro definitions. but notice that a symbol can be referenced before it is defined.

The effect of the DEFINE statement could have been obtained by declaring a variable and assigning its value, as follows:

## VAR

carry_out : boolean;
ASSIGN
carry_out := value \& carry_in;
Votice that in this case. the current value of the variable is assigned. rather than the next value. Defined symbols are sometimes preferable to variables. however, since they don't require introducing a new variable into the OBDD representation of the system. The weakness of defined symbols is that they cannot be given values non-deterministically. Another difference between defined symbols and variables is that while variables are statically typed. definitions are not. This may be an advantage or a disadvantage. depending on your point of view.

In a parallel-assignment language. the question arises: "What happens if a given variable is assigned twice in parallel?" More seriously: "What happens in the case of an absurdity. like a $:=a+1$; (as opposed to the sensible next (a) $:=a+1 ;$ )?" In the case of SMV. the compiler detects both multiple assignments and circular dependencies. and treats these as semantic errors, even in the case where the corresponding system of equations has a unique solution. Another way of putting this is that there must be a total order in which the assignments can be executed which respects all of the data dependencies. The same logic applies to defined symbols. As a result. all legal S.MV programs are realizable.

By default. all of the assignment statements in an SMV program are executed in parallel and simultaneously. It is possible, however. to define a collection of parallel processes. whose actions are interleaved arbitrarily in the execution sequence of the program. This is useful for describing communication protocols, asynchronous circuits. or other systems whose actions are not synchronized (including synchronous circuits with more than one clock). This technique is illustrated by the following program. which represents a ring of three inverting gates.

```
MODULE main
VAR
    gate1 : process inverter(gate3.output);
```

```
    gate2 : process inverter(gate1.output);
    gate3 : process inverter(gate2.output);
SPEC
    (AG AF gate1.out) & (AG AF !gatel.out)
MODULE inverter(input)
VAR
    output : boolean;
ASSIGN
    init(output) := 0;
    next(output) := !input;
```

A process is an instance of a module which is introduced by the keyword process. The program executes a step by non-deterministically choosing a process, then executing all of the assignment statements in that process in parallel. It is implicit that if a given variable is not assigned by the process. then its value remains unchanged. Because the choice of the next process to execute is non-deterministic. this program models the ring of inverters independently of the speed of the gates. The specification of this program states that the output of gatel oscillates (ie., that its value is infinitely often zero, and infinitely often 1). In fact. this specification is false, since the system is not forced to execute every process infinitely often, hence the output of a given gate may remain constant, regardless of changes of its input.

In order to force a given process to execute infinitely often. we can use a fairness constraint. A fairness constraint restricts the attention of the model checker to those execution paths along which a given CTL formula is true infinitely often. Each process has a special variable called running which is true if and only if that process is currently executing. By adding the declaration

## FAIRNESS

running
to the module inverter, we can effectively force every instance of inverter to execute infinitely often. thus making the specification true.

One advantage of using interleaving processes to describe a system is that it allows a particularly efficient OBDD representation of the transition relation. We observe that the set of states reachable by
one step of the program is the union of the sets of states reachable by each individual process. Hence, rather than constructing the transition relation of the entire system, we can use the transition relations of the individual processes separately and the combine the results (cf. section 2.4.2). This can yield a substantial savings in space in representing the transition relation.

The alternative to using processes to model an asynchronous circuit would be to have all gates execute simultaneously, but allow each gate the non-deterministic choice of evaluating its output, or keeping the same output value. Such a model of the inverter ring would look like the following:

```
MODULE main
var
    gate1 : inverter(gate3.output);
    gate2 : inverter(gate2.output);
    gate3 : inverter(gate1.output);
SPEC
    (AG AF gatel.out) & (AG AF !gatel.out)
MODULE inverter(input)
VAR
    output : boolean;
ASSIGN
    init(output) := 0;
    next(output) := !input union output;
```

The union operator allows us to express a nondeterministic choice between two expressions. Thus, the next output of each gate can be either its current output, or the negation of its current input - each gate can choose non-deterministically whether to delay or not. . Is a result. the number of possible transitions from a given state can be as high as $2^{n}$. where $n$ is the number of gates. This sometimes (but not alwavs) makes it more expensive to represent the transition relation. The relative advantages of interleaving and simultaneous models of asynchronons systems are discussed in section 2.4.2.

As a second example of processes. the following program uses a variable semaphore to implement mutual exclusion between two asynchronous processes. Each process has four states: idle. entering.
critical and exiting. The entering state indicates that the process wants to enter its critical region. If the variable semaphore is zero, it goes to the critical state, and sets semaphore to one. On exiting its critical region, the process sets semaphore to zero again.

MODULE main
VAR
semaphore : boolean;
proc1 : process user;
proc2 : process user;
ASSIGN
init(semaphore) := 0 ;
SPEC
AG ! (proc1.state $=$ critical \& proc2.state $=$ critical $)$
MODULE user
VAR
state : \{idle, entering, critical,exiting\};
ASSIGN
init(state) := idle;
next(state) :=
case
state $=$ idle : \{idle,entering\};
state $=$ entering \& !semaphore : critical;
state $=$ critical : \{critical,exiting\};
state $=$ exiting : idle;
1 : state;
esac;
next(semaphore) :=
case
state $=$ entering : 1 ;
state $=$ exiting : 0 ;
1 : semaphore;
esac;
FAIRNESS
running
If any specification in the program is false. the SMV model checker attempts to produce a counterexample. proving that the specification is false. This is not always possible. since formulas preceded by existential
path quantifiers cannot be proved false by a showing a single execution path. Similarly, subformulas preceded by universal path quantifier cannot be proved true by a showing a single execution path. In addition, some formulas require infinite execution paths as counterexamples. In this case, the model checker outputs a looping path up to and including the first repetition of a state.

In the case of the semaphore program. suppose that the specification were changed to

```
AG (proc1.state = entering -> AF proci.state = critical)
```

[n other words. we specify that if proc1 wants to enter its critical region, it eventually does. The output of the model checker in this case is shown in figure 3.1. The counterexample shows a path with proc1 going to the entering statc. followed by a loop in which proc2 repeatedly enters its critical region and the returns to its idle state. with proc1 only executing only while proc2 is in its critical region. This path shows that the specification is talse. since proci never enters its critical region. Note that in the printout of an execution sequence. only the values of variables that change are printed. to make it easier to follow the action in systems with a large number of variables.

Although the parallel assignment mechanism should be suitable to most purposes. it is possible in S.MV to specify the transition relation directly as a propositional formula in terms of the current and next values of the state variables. . Iny current/next state pair is in the transition relation if and only if the value of the formula is one. Similarly. it is possible to give the set of initial states as a formula in terms of only the current state variables. These two functions are accomplished by the TRANS and INIT statements respectively. As an example. here is a description of the three inverter ring using only TR.ANS and [NIT:

```
MODULE main
VAR
    gate1 : inverter(gate3.output);
    gate2 : inverter(gate1.output);
    gate3 : inverter(gate2.output);
SPEC
```

    (AG AF gatel.out) \& (AG AF !gatel.out)
    ```
specification is false
AG (proci.state = entering -> AF proc1.s... is false:
.semaphore = 0
.procl.state = idle
.proc2.state = idle
next state:
[executing process .proc1]
next state:
.proc1.state = entering
AF procl.state = critical is false:
[executing process .proc2]
next state:
[executing process .proc2]
.proc2.state = entering
next state:
[executing process .proci]
.semaphore = 1
.proc2.state = critical
next state:
[executing process .proc2]
next state:
[executing process .proc2]
.proc2.state = exiting
next state:
.semaphore = 0
    proc2.state = idle
```

Figure 3.1: Model checker output for semaphore example

```
MODULE inverter(input)
VAR
    output : boolean;
INIT
    output = 0
TRANS
    next(output) = !input I next(output) = output
```

According to the TRANS declaration, for each inverter, the next value of the output is equal either to the negation of the input, or to the current value of the output. Thus, in effect, each gate can choose nondeterministically whether or not to delay. The use of TRANS and INIT is not recommended. since logical absurdities in these declarations can lead to unimplementable descriptions. For example, one could declare the logical constant 0 (false) to represent the transition relation, resulting in a system with no transitions at all. However, the flexibility of these mechanisms may be useful for those writing translators from other languages to SMV.

To summarize. the SMV language is designed to be flexible in terms of the styles of models it can describe. It is possible to fairly concisely describe synchronous or asynchronous systems, to describe detailed deterministic models or abstract nondeterministic models. and to exploit the modular structure of a system to make the description more concise. It is also possible to write logical absurdities if one desires to. and also sometimes if one does not desire to. using the TRANS and INIT declarations. By using only the parallel assignment mechanism, however. this problem can be avoided. The language is designed to exploit the capabilities of the symbolic model checking technique. As a result the available data types are all static and finite. No attempt has been made to support a particular model of communication between concurrent processes (eg., synchronous or asynchronous message passing). In addition. there is no explici support for some features of communicating process models such as sequential composition. Since the full generality of the symbolic model checking technique is available through the SMV language. it is possible that translators from various languages. process models; and intermediate formats could be created. In particular. existing silicon compilers could be used to translate high level languages with rich feature sets into a low level form (such as a Mealy machine)
that could be readily translated into the SMV language.

### 3.2 The input language

This section describes the various constructs of the SMV input language, and their syntax.

### 3.2.1 Lexical conventions

An atom in the syntax described below may be any sequence of characters in the set $\{A-2, a-z, 0-9, \ldots,-\}$, beginning with an alphabetic character. All characters in a name are significant, and case is significant. Whitespace characters are space. tab and newline. Any string starting with two dashes ("--") and ending with a newline is a comment. A number is any sequence of digits. Any other tokens recognized by the parser are enclosed in quotes in the syntax expressions below.

### 3.2.2 Expressions

Expressions are constructed from variables, constants. and a collection of operators, including Boolean connectives, integer arithmetic operators, and case expressions. The syntax of expressions is as follows.

```
expr ::
```

| atom | ; ; a symbolic constant |
| :---: | :---: |
| \| number | ; ; a numeric constant |
| id | ;; a variable identifier |
| \| "!" expr | ; l logical not |
| \| expr1 "\&" expr2 | ;; logical and |
| 1 expr1 "\|" expr2 | ; l logical or |
| \| expr1 "->" expr2 | ; logical implication |
| \| expr1 "<->" expr2 | ; logical equivalence |
| \| expri "=" expr2 | ;; equality |
| \| expr1 "く" expr2 | ; less than |
| \| expri ">" expr2 | ; greater than |
| \| expr1 "く=" expr2 | ; ; less that or equal |
| \| expri ">=" expr2 | ; greater than or equal |
| 1 expr $1^{\prime \prime}+{ }^{\text {" expr2 }}$ | ;; integer addition |

```
| expr1 "-" expr2 ;; integer subtraction
| expr1 "*" expr2 ;; integer multiplication
| expri "/" expr2 ; integer division
| expr1 "mod" expr2 ;; integer remainder
| "next" "(" id ")" ;; next value
| set_expr ;; a set expression
| case_expr ;; a case expression
```

An id. or identifier. is a symbol or expression which identifies ait object, such as a variable or defined symbol. Since an id can be an atom, there is a possible ambiguity if a variable or defined symbol has the same name as a symbolic constant. Such an ambiguity is flagged by the compiler as an error. The expression next ( $x$ ) refers to the value of identifier $x$ in the next state (see section 3.2.3). The order of parsing precedence from high to low is

```
*,/
+,-
mod
=,\langle,\rangle,\langle=,>=
!
&
l
->,\langle->
```

Operators of equal precedence associate to the left. Parentheses may be used to group expressions.

A case expression has the syntax

```
case_expr ::
    "case"
        expr_al ":" expr_b1 ";"
        expr_a2 ":" expr_b2 ";"
    "esac"
```

I case expression returns the value of the first expression on the right hand side. such that the corresponding condition on the left hand side is true. Thus. if expral is true. then the result is expr_b1. Otherwise. if expr_a2 is true, then the result is expr_b2. etc. If none of
the expressions on the left hand side is true, the result of the case expression is the numeric value 1 . It is an error for any expression on the left hand side to return a value other than the truth values 0 or 1 .

A set expression has the syntax

```
set_expr ::
    "{" val1 "," val2 "," ... "}"
    | expr1 "in" expr2 ;; set inclusion predicate
    | expr1 "union" expr2 ;; set union
```

$A$ set can be defined by enumerating its elements inside curly braces. The elements of the set can be numbers or symbolic constants. The inclusion operator tests a value for membership in a set. The union operator takes the union of two sets. If either argument is a number or symbolic value instead of a set. it is coerced to a singleton set.

### 3.2.3 Declarations

## The VAR declaration

A state of the model is an assignment of values to a set of state variables. These variables (and also instances of modules) are declared by the notation

```
decl :: "VAR"
    atom1 ":" type1 ";"
    atom2 ":" type2 ";"
```

The type associated with a variable declaration can be either Boolean. scalar. or a user defined module. A type specifier has the syntax

```
type :: boolean
    | "{" val1 "," val2 "," ... "}"
    | atom [ "(" expr1 "," expr2 "," ... ")" ]
    | "process" atom [ "(" expr1 "," expr2 "," ... ")" ]
val :: atom | number
```

A variable of type boolean can take on the numerical values 0 and 1 (representing false and true. respectively). In the case of a list of
values enclosed in set brackets (where atoms are taken to be symbolic constants), the variable is a scalar which can take any of these values. Finally, an atom optionally followed by a list of expressions in parentheses indicates an instance of module atom (cf. section 3.2.4). The keyword process causes the module to be instantiated as an asynchronous process (cf. section 3.2.6).

## The ASSIGN declaration

An assignment declaration has the form

```
decl :: "ASSIGN"
    dest1 ":=" expri ";"
    dest2 ":=" expr2 ";"
dest :: atom
    | "init" "(" atom ")"
    | "next" "(" atom ")"
```

On the left hand side of the assignment. atom denotes the current value of a variable. init(atom) denotes its initial value. and next (atom) denotes its waiue in the next state. If the expression on the right hand side evaluates to an integer or symbolic constant. the assignment simply means that the left hand side is equal to the right hand side. On the other hand. if the expression evaluates to a set. then the assignment means that the left hand side is contained in that set. It is an error if the value of the expression is not contained in the range of the variable on the left hand side.

In order for a program to be implementable. there must be some order in which the assignments can be executed such that no variable is assigned after its value is referenced. This is not the case if there is a circular dependency among the assignments in any given process. Hence. such a condition is an error. In addition. it is an error for a variable to be assigned more than once simultaneously. To be precise. it is an error if:

1. the next or current value of a variable is assigned more than once in a given process. or
2. the initial value of a variable is assigned more than once in the program, or
3. the current value and the initial value of a variable are both assigned in the program, or
t. the current value and the next value of a variable are both assigned in the program. or
4. there is a circular dependency, or
5. the current value of a variable depends on the next value of a variable.

## The TRANS declaration

The transition relation $R$ of the model is a set of current state/next state pairs. Whether or not a given pair is in this set is determined by a Boolean valued expression. introduced by the TRANS keyword. The syntax of a TRANS declaration is

```
decl :: "TRANS" expr
```

It is an error for the expression to yield any value other than 0 or 1 . If there is more than one TRANS declaration. the transition relation is the conjunction of all of TRANS declarations.

## The INIT declaration

The set of initial states of the model is determined by a Boolean expression under the INIT keyword. The syntax of a INIT declaration is

```
decl :: "INIT" expr
```

It is an error for the expression to contain the next() operator. or to vield any value other than 0 or 1 . If there is more than one INIT declaration. the initial set is the conjunction of all of the INIT declarations.

## The SPEC declaration

The system specification is given as a formula in the temporal logic CTL, introduced by the keyword SPEC. The syntax of this declaration is

```
decl :: "SPEC" ctlform
```


## A CTL formula has the syntax

```
ctlform ::
    expr ;; a Boolean expression
    | "!" ctlform ;; logical not
    | ctlform1 "&" ctlform2 ;; logical and
    | ctlform1 "|" ctlform2 ;; logical or
    | ctlform1 "->" ctlform2 ;; logical implies
    | ctlform1 "<->" ctlform2 ;; logical equivalence
    | "E" pathform ;; existential path quantifier
    | "A" pathform ;; universal path quantifier
```

The syntax of a path formula is

```
pathform ::
```

```
"X" ctlform ;; next time
    "F" ctlform ;; eventually
    "G" ctlform ;; globally
    ctlform1 "U" ctlform2 ;; until
```

The order of precedence of operators is (from high to low)

```
E,A,X,F,G,U
!
&
|
->,<->
```

Operators of equal precerlence associate to the left. Parentheses may be used to group expressions. It is an error for an expression in a ('TL formula to contain a next() operator or to return a value other than 0 or L . If there is more than one SPEC declaration, the specification is the conjunction of all of the SPEC declarations.

## The FAIR declaration

A fairness constraint is a CTL formula which is assumed to be true infinitely often in all fair execution paths. When evaluating specifications, the model checker considers path quantifiers to apply only to fair paths. Fairness constraints are declared using the following syntax:

```
decl :: "FAIR" ctlform
```

A path is considered fair if and only if all fairness constraints declared in this manner are true infinitely often.

## The DEFINE declaration

In order to make descriptions more concise, a symbol can be associated with a commonly used expression. The syntax for this declaration is

```
decl :: "DEFINE"
    atom1 ":=" expr1 ";"
    atom2 ":=" expr2 ";"
```

When every an identifier referring to the symbol on the left hand side occurs in an expression. it is replaced by the value of the expression on the right hand side (not the expression itself). Forward references to defined symbols are allowed. but circular definitions are not allowed. and result in an error.

### 3.2.4 Modules

$\lambda$ module is an encapsulated coltection of declarations. Once defined. a module can be reused as many times as necessary. Modules can also be parameterized. so that each instance of a module can refer to different data values. A module can contain instances of other modules. allowing a structural hierarchy to be built. The syntax of a module is as follows.

```
module ::
    [ "OPAQUE" ]
    "MODULE" atom [ "(" atom1 "," atom2 "," ... ")" ]
    decl1
    decl2
```

The optional keyword OPAQUE is explained in the section on identifiers. The atom immediately following the keyword MODULE is the name associated with the module. Module names are drawn from a separate name space from other names in the program, and hence may clash with names of variables and definitions. The optional list of atoms in parentheses are the formal parameters of the module. Whenever these parameters occur in expressions within the module. they are replaced by the actual parameters which are supplied when the module is instantiated.

A instance of a module is created using the VAR declaration cf. section 3.2.3). This declaration supplies a name for the instance, and also a list of actual parameters. which are assigned to the formal parameters in the module definition. In actual parameter can be any legal expression. It is an error if the number of actual parameters is different from the number of formal parameters. The semantics of module instantiation is similar to call-by-relerence. For example, consider the following program fragment:

VAR
a : boolean;
b : foo(a);
MODULE $f \circ 0(x)$
ASSIGN
$x \quad:=1$;
The variable a is assigned the value 1. Now consider the following program:

DEFINE

```
    a := 0;
VAR
    b : bar(a);
MODULE bar(x)
DEFINE
    a := 1;
    y := x;
```

In this program. the value assigned to $y$ is 0 . Using a call-by-name (macro expansion) mechanism, the value of $y$ would be 1 , since a would be substituted as an expression for $x$.

Forward references to module names are allowed, but circular references are not. and result in an error.

### 3.2.5 Identifiers

An id, or identifier. is an expression which references an object. Objects are instances of modules, variables, and defined symbols. The syntax of an identifier is as follows.
id : :

```
atom
| id "." atom
```

An atom identifies the object of that name as clefined in a VAR or DEFINE declaration. If $a$ identifies an instance of a module, then the expression $a . b$ identifies the component object named $b$ of instance $a$. This is precisely analogous to accessing a component of a structured data type. .iote that an actual parameter of module instance a can identify another module instance $b$. allowing $a$ to access components of b. as in the following example:

## VAR

a : foo (b);
b : bar(a);
MODULE $\mathrm{fOO}(\mathrm{x})$
DEFINE
$c:=x . p \mid x . q ;$
MODULE bar ( $x$ )
VAR
p : boolean;
$q$ : boolean;
Here. the value of $c$ is the logical or of $p$ and $q$. If the keyword OPAQUE appears before a module definition. then the variables of an in-
stance of that module are not externally accessible. Thus, the following program fragment is not legal:

```
VAR
    a : foo();
DEFINE
    b := a.x;
OPAQUE MODULE fOO()
VAR
    x : boolean;
```


### 3.2.6 Processes

Processes are used to model interleaving concurrency, with shared variables. A process is a module which is instantiated using the keyword process (cf. section 3.2.3). The program executes a step by nondeterministically choosing a process, then executing all of the assignment statements in that process in parallel. simultaneously. Each instance of a process has special variable Boolean associated with it called running. The value of this variable is 1 if and only if the process instance is currently selected for execution. The rule for determining whether a given variable is allowed to change value when a given process is executing is as follows: if the next value of a given variable is not assigned in the currently executing process. but is assigned in some other process. then the next value is the same as the current value.

### 3.2.7 Programs

The syntax of an S:NV program is

```
program ::
    module1
    module2
```

There must be one module with the name main and no formal parameters. The module main is the one instantiated by the compiler.

### 3.3 Formal semantics

In this section we assign a formal semantics to SMV programs. In essence, a program is viewed as a system of equations whose solutions determine the transition relation and initial states of a Kripke structure. In fact, this semantics assigns meaning to some programs which are not actually accepted by the compiler due to the rules regarding multiple assignments and circular dependencies. Here, we define a semantics for a subset of the language which does not include the process keyword. This subset will be called SMV.0. The semantics of S.MV. 0 is syntax directed - the denotation of a program is a function of the denotations of its syntactic components. It is also compositional with regard to bisimulation and simulation. as we will prove in chapter 5. This makes it possible to use compositional proof methods for verifying S.MV. 0 programs. including induction over the structure of programs. The semantics for S.MV.I. which includes the process keyword. is given in appendix $A$.

### 3.3.1 The model

The set $N$ of names. is the set of all character strings made up of the letters. the digits. the underscore and the minus sign characters. beginning with a letter. The store $L=L_{V} \cup L_{H}$ is made up of two clisjoint. countably infinite sets of locations $L_{V}$ and $L_{H}$. We will call the former the visible locations. and the latter the hidden locations. The set of locations $L$ is defined recursively. It is the least set such that

1. if $n \in \lambda$. then $n \in L_{V}$. and
2. if $l \in L_{v}$ and $n \in . V$. then $l . n \in L_{v}$. and
3. if $l \in L_{l}$, then $l \in L_{H}$.

The set of values $V$ is the union of the integers in the range $\left[-2^{31} \cdot 2^{31}-1\right]$ and $V$. the set of names. $A$ state $x: L \rightarrow V$ is a function from locations to values. Let $S=L \rightarrow V$ be the set of all possible states.

If $p$ is a declaration, then its denotation $\llbracket p \rrbracket$ is a triple $(T . l . R)$. The $T$ component is a partial function from $L$ to the finite subsets of $I$.

If $l$ is a location, then $T(l)$, when defined, is the type of $l$ - the set of values that can be assigned to location $l$. The component $I \subseteq S$ is the set of initial states. Finally, the component $R \subseteq S \times S$ is the transition relation.

In the following sections, we define the denotations of the various kinds of declarations. We then define a composition operator \| which gives the denotation of a program in terms of its declarations.

### 3.3.2 Expressions

An expression denotes a function from states to finite subsets of $V$. according to the following rules:

1. If $v$ is a value. then $[c](x)=\{c\}$.
2. If $l$ is a location, then $\mathbb{I} l(x)=\{x(l)\}$.
3. If $e_{1}, \epsilon_{2}$ are expressions. and $o$ is one of

$$
+.-. * . / . \bmod .\rangle .\rangle=.\langle.\langle=,=, \& \varepsilon . \mid,-\rangle .\langle-\rangle
$$

then

$$
\llbracket e_{1} \circ e_{2} \rrbracket(x)=\left\{\llbracket o \rrbracket\left(v_{1} \cdot v_{2}\right) \mid v_{1} \in \llbracket e_{1} \rrbracket(x), v_{2} \in \llbracket e_{2} \rrbracket(x)\right\}
$$

4. If $e$ is an expression. then

$$
\mathbb{T} e](x)=\{\rrbracket!\square(u) \mid v \in[e](x)\}
$$

5. If $e_{1}, \epsilon_{2}$ are expressions.

$$
\left[\epsilon_{1} \text { union } \epsilon_{2}\right](x)=\llbracket c_{1} \rrbracket \cup \llbracket e_{2} \rrbracket
$$

6. If $\epsilon_{1}, \epsilon_{2}$ are expressions.

$$
\left[e_{1} \text { in } \epsilon_{2} \mathbb{Z}(x)=\left[e_{1} \mathbb{1} \subseteq \mathbb{[} e_{2} \mathbb{1}\right.\right.
$$

The functions denoted by $+,-, *, /$ are the usual functions of arithmetic modulo $2^{32}$. The function denoted by mod is the positive remainder of division mod $2^{32}$. The functions denoted by the relational operators $>,>=$. < and <= return 0 when the relation is false and 1 when the relation is true. and are defined for numeric values only. For nonnumeric values. they return $\perp$. The equality operator $=$ is defined for all values. and returns 0 when they are unequal. and 1 when they are equal. The functions denoted by the Boolean operators are \& (for and). I (for or).! (for not). $\rightarrow$ (for implies) and $\Leftrightarrow>$ (for logical equivalence) are defined only for the values 0 and 1 . and return $\perp$ otherwise.

### 3.3.3 Assignments and definitions

There is no semantic difference between assignments and definitions. If $l$ is a location. and $t$ is an expression. then the assignment $l:=f$ : denotes a triple ( $I . I . R$ ). where

1. $T=\emptyset$
2. $I=$.
3. $R=\left\{(x, y) \in x^{2} \mid l(x) \in \mathbb{I} \mathbb{\}(x)\right\}$

The assignment next $(l):=\ell$ : denotes a triple ( $T . I, R$ ) where

1. $T=\emptyset$
2. $I=\breve{C}$
3. $R=\left\{(x . y) \in s^{2} \mid l(y) \in \llbracket e \rrbracket(x)\right\}$

The assignment init(l) :=e: denotes a triple (T. $I . R$ ) where

1. $T=\eta$
2. $I=\{r \doteq: \|(x) \in \llbracket \in \rrbracket(x)\}$
3. $R=4$

### 3.3.4 Variable declarations

If $l$ is an identifier and $v_{1}, v_{2}, \ldots, v_{n}$ are values, then

$$
\operatorname{VAR} l:\left\{v_{1}, v_{2}, \ldots, v_{n}\right\}
$$

denotes a triple ( $T, I, R$ ) where

1. $T=\left\{\left(l .\left\{v_{1}, v_{2}, \ldots, c_{n}\right\}\right)\right\}$
2. $\left.I=\left\{x \in S \mid x(l) \in\left\{c_{1} \cdot v_{2}, \ldots . c_{n}\right\}\right)\right\}$
3. $\left.R=\left\{(x, y) \in S \mid x(l), y(l) \in\left\{v_{1}, c_{2}, \ldots . v_{n}\right\}\right)\right\}$

### 3.3.5 Renaming

Let $\sigma: L \rightarrow L$ be a function from locations to locations. This in turn induces a map $\Phi$ on states. such that for all states $x$ and locations $l$.

$$
\Phi(x)(l)=x(o(l)) .
$$

If $M=(T . I . R)$. then let $\phi(. M)=\left(T^{\prime} . I^{\prime} . R^{\prime}\right)$ where

1. $T^{\prime}(\phi(l))=T(l)$.
2. $I^{\prime}=\{x \mid \Phi(x) \in I\}$ and
3. $R^{\prime}=\{(x, y) \mid(\Phi(x), \Phi(y)) \in R\}$.

Vote that the definition of $T$ does not make sense if $\odot$ maps two locations with different types onto the same location. In this case. $\mathscr{O}(M)$ is a type error. There are two rules regarding the renaming function $o$ which must be respected to allow compositional reasoning about SMIV programs. These are:

1. A hidden location cannot be renamed to a visible location. and
2. Two distinct locations cannot be mapped to the same hidden location.

These rules are respected by the SWLV. 0 semantics. Notice that it is allowable to rename visible locations to hidden locations. In this way. we can accomplish both hiding and renaming with the same operator.

### 3.3.6 Parallel composition

The parallel composition of two processes $M_{1} \|_{1} M_{2}$ is formed in two steps. First, a renaming is applied to map the hidden locations of $M_{1}$ and $M_{2}$ onto disjoint spaces. Then the union of the type functions $T$ and the intersections of the initial sets $I$ and the transition relations $R$ are taken. (learly, this does not make sense if the $T$ components do not agree on the type of some location, since the union would not be a function. Formally, let $M_{1}=\left(T_{1}, I_{1}, R_{1}\right)$ and $M_{2}=\left(T_{2}, I_{2}, R_{2}\right)$. Let $n_{1}$ and $n_{2}$ be two distinct names. For $i \in 1,2$, let $\varphi_{i}(l)=. n_{i} . l$ for all $l \in L_{H}$ and $\phi_{i}(l)=l$ otherwise, and let $M_{i}^{\prime}=\varnothing\left(M_{i}\right)$. The parallel composition $M=M_{1} \| M_{2}$ is defined as follows:

1. $T=T_{1}^{\prime} \cup T_{2}^{\prime}$
2. $I=I_{1}^{\prime} \cap I_{2}^{\prime}$
3. $R=R_{1}^{\prime} \cap R_{2}^{\prime}$

If $d_{1}, d_{2}, \ldots . d_{k}$ are cleclarations, then $\llbracket d_{1} d_{2} \ldots d_{k} \rrbracket$ is the parallel composition

$$
\llbracket d_{1} \rrbracket\left\|\llbracket d_{2} \rrbracket\right\| \cdots \| \llbracket d_{k} \rrbracket
$$

### 3.3.7 Instantiation

Suppose that module $A$ is defined as follows:

$$
\text { MODULE } A\left(n_{1}, n_{2}, \ldots, n_{k}\right) D
$$

where $n_{1}, n_{2} \ldots \ldots n_{k}$ are distinct names and $D$ is a sequence of declarations. Let $r$. $l_{1}, l_{2}, \ldots . l_{k}$ be visible locations. Let $\oplus$ be a renaming. such that. for all $l \in L_{V}$.

1. for all $1 \leq i \leq h: ~ o\left(n_{i}\right)=l_{i}$, and $\mathscr{O}\left(n_{i} \cdot l\right)=l_{i} . l$.
2. for all $n \in . V-\left\{n_{1}, n_{2}, \ldots, n_{k}\right\}, \mathscr{D}(n)=\cdots . n$. and $\Phi(n . l)=r . n . l$.
3. $\omega(. l)=. l$

Then $\mathbb{\operatorname { V A R } r} \boldsymbol{r}: A\left(l_{1}, l_{2}, \ldots . l_{k}\right) ; \rrbracket=\varnothing(D)$.
On the other hand. suppose that $A$ is defined as follows:

$$
\text { OPAQUE MODULE } A\left(n_{1}, n_{2}, \ldots, n_{k}\right) D
$$

where $n_{1}, n_{2}, \ldots, n_{k}$ are distinct names and $D$ is a sequence of declarations. Let $r . l_{1}, l_{2} \ldots \ldots l_{k}$ be visible locations. Let $m_{1}$ and $m_{2}$ be distinct names in $N$. Let $\mathscr{O}$ be a cenaming. such that. for all $t \in L_{V}$.

1 . for all $\mathrm{l} \leq i \leq k: \propto\left(n_{i}\right)=l_{i}$. and $\Phi\left(n_{i} \cdot l\right)=l_{i}, l$,
2. for all $n \in N-\left\{n_{1}, n_{2} \ldots \ldots i i_{k}\right\}$, $O(n)=m_{1} \cdot n$. and $\propto(n . l)=$ .mı.n.l.
3. $\propto(. l)=. m_{2} . l$

Then $\left[\operatorname{VaR} r: A\left(l_{1}, l_{2}, \ldots, l_{k}\right): \rrbracket=M i\right.$.

### 3.3.8 Specifications

Each program is associated with a Kripke structure which determines the truth value of (TL formulas in the specification. The atomic propositions in this case are all the Boolean valued expressions. The Kripke structure assoriated with a program whose denotation is the triple $(T . I . R)$ is a hripke model $\left.K^{\prime}=1, S, L^{\prime}\right)$ where

1. 3 is the set of states defined above.
2. $R$ is the transition melation and
3. if $\epsilon$ is an expression. then

$$
L^{\prime}(e)=\{x \in S \mid[e](x)=\{1\}\}
$$

The specification is a formula $f$ in ("TL with fairness constraints. It is satisfied exactly when $K . v_{0}=f$ for all $\boldsymbol{r}_{0} \in I$.

## Chapter 4

## A Distributed Cache Protocol

In this chapter, we look at an application of the SMIV symbolic model checker to a cache consistency protocol developed at Encore Computer C'orporation for their Cigamax distributed multiprocessor [MS91]. This protocol is of interest as a test case for automatic verification for two reasons. First. it is not a theoretical exercise, but a real design. which is driven by considerations of performance and economics. as well as the usual constraints of industrial design, such as compatibility with existing hardware and software. Second. this protocol is a good example of a system where random simulation methods are ineffective in finding resign errors.

The Gigamax is a distributed. shared memory multiprocessor. in which the processors are grouped into clusters. Each cluster has a local bus, and uses bus snooping [ $A B 86]$ to maintain cache consistency within the cluster. In addition. each cluster has an interface called a [IC'. which links the cluster into a net,work. The UIC keeps the caches in the cluster consistent with the rest of the network by acting as both a bus snooper and a bus master on behalf of the remote clusters, using a table which keeps track of the remote status of all cache blocks from the local main memory. This allows it to intervene in bus transactions which affect remotely owned blocks. and to send appropriate invalidation or call back requests to the network. The network is organized into a hierarchy, as depicted in figure 1.1. The global bus, at the top of the


Figure 4.1: Gigamax memory architecture
hierarchy, has one UIC connected to each cluster. These UICs record the status of all cache blocks which are present in the corresponding cluster. This eliminates the need for directory pointers in main memory. at the possible expense of a bottleneck in the global bus.

Protocols such as this are difficult to debug using simulation. in part because the order of events such as cache misses and message arrivals in various parts of the system is unpredictable. Subtle errors sometimes require a long sequence of such events to manifest themselves. Since the number of such sequences is combinatoric, the probability of such a sequence occurring in a random simulation rapidly vanishes as the sequence length increases. Nevertheless. for the design process to stabilize. it is necessary to provide timely information about errors to the design team. since the greater the delay in discovering an error. the greater is the disruption required to fix it. Ideally, a protocol should be error free belore a hardware (or software) implementation is considered. Otherwise. the options for fixing the errors will be greatly limited by cost considerations, and the likelihood of the design change introducing
other errors will be high.
For this reason, we will consider the verification of the Gigamax protocol at a high level of abstraction. neglecting many admittedly important details of the implementation, such as the widespread use of pipelining, or the link level protocol that communicates messages between clusters. The basic method for building an abstract model of a protocol is to introduce nondeterminism in those cases where the level of detail of the model is insufficient to uniquely determine the outcome of an event. or where design decisions have been left open. We will make a note of places in the model where nondeterminism has been used in this way. and in what way the state of an implementation might correspond to the state of the abstract model.

### 4.1 The Protocol

The purpose of a cache consistency system is to provide the illusion to the programmer of a distributed computer that all processors in the system have access to a shared global store. This illusion must be provided despite the fact that the physical storage is distributed. To reduce the latency of access to the distributed main storage. each processor is provided with a local cache - a semi-associative store. which holds a collection of memory blocks recently used by the processor. The time required to access to this store is less than to access main storage. An access to a memory biock stored in the cache is called a hit. while an access to a memory block not stored in the cache is called a miss. A miss requires an access to main storage (which may be remote). to retrieve the required memory block and enter it in the cache. This may result in the replacement of another block in the cache. to make room for the block being entered in the cache. If the replaced block has been modified while in the cache. it must be returned to main storage. This is called a copy back operation.

The first cache consistency protocols for multiprocessors were called bus snooping protocols [AB86]. They required that the processors in the system be connected by a bus. or other broadcast medium. In a bus snooping system. each time a memory access occurs over the bus. all of the caches are chocked to determine whether they contain the
addressed block. If the block is present in a cache, a change in status may be required. For example, if the block is present and modified, the access must be stalled while the modified data are copied back to main storage. In a more sophisticated protocol, the cache with the modified data may supply the data directly to the requesting cache, without the intermediary of main storage. In case of a memory access caused by an attempt to modify the data, all caches in which the block is present must invalidate. that is, remove the block from cache storage. This insures that all cached copies of the block remain consistent.

The Gigamax protocol uses bus snooping techniques to maintain consistency of the caches within a single cluster. The main difference between the Gigamax snooping protocol and those described in [AB86] is that the Cigamax uses a split transaction bus. This means that a processor accessing memory over the bus first places a request on the bus, and then frees the bus for other transactions while awaiting a response. The bus snooping technique is not practical for large scale multiprocessors. because the broadcast medium quickly becomes saturated. For this reason. the Gigamax uses a message passing protocol to maintain consistency between clusters. The split transaction bus protocol allows trathc to continue on the bus while messages are in transit in the network.

The terminology used in the sequel is changed somewhat from the Encore terminology, and the protocol is somewhat simplified to make the presentation clearer. The basic protocol is preserved. however. including a subtle error which was discovered by the SMV system. The following is a description of the protocol. first in English. then in the SMV input language. In the model, we consider only the status of a single memory block. This is our first use of abstraction. and results in nondeterminism in several places in the model.

### 4.1.1 Processors

Each memory block stored in each cache has an associated state. which can be either invalid. shared, or owned. Alternative names for these states would be absent. present, and modified, respectively. The shared state indicates that there may be other processors which have this block stored in their cache. Therefore, a block in the shared state can be
read by the processor, but not written. since writing might result in an inconsistency between two caches. The owned state indicates that no other processors have this block in their cache, and that the data in the cache have been modified. Therefore, a block in the owned state can be both read and written by the processor. The invalid state indicates that the block is not present in the cache. Therefore. the block cannot be read or written by the processor.

```
MODULE cache-device
```

```
VAR
    state : {invalid,shared,owned};
```

```
DEFINE
    readable := ((state = shared) | (state = owned)) & !waiting;
    writable := (state = owned) & !waiting;
```

The split transaction bus snooping protocol works in the following way. At each bus cycle. the bus arbiter chooses a processor among the requesting processors to be the bus master. The remaining processors are referred to as slaves. The master issues a command on the bus, of which there are three basic types. A read command is a request for a given memory block, and is answered by a response command. A write command stores data in main memory. The write and response commands can be combined into a single command called a write-response. which has the simultaneous offect of supplying data to a requester and -toring it in main memory: Fach command also signals the next state that the bus master will enter. Thus. a read-owned command indicates that the bus master intends to modify the data. and a read-shared inrlicates that it does not. A write-shared indicates that the bus master is writing data. but maintaining a shared copy, while a write-invalid indicates that it is not keeping the block leg.. it is replacing it with another block). The basic commands. and their uses are summarized in table 4.1 . We note that no external command is required to go from the shared state to the invalid state. This occurs when a shared block is removed to make room for another block in the cache. Since our
model does not contain the states of any other blocks, we allow this replacement to occur nondeterministically, at any time. Thus we model any possible cache replacement policy.

A slave, observing a command on the bus, may decide to modify its state. For example, a slave observing a read-owned command changes its state to invalid, since the bus master, entering the owned state, will assume it has the only cached copy of the block. Correspondingly, a slave in the owned state observing a read-shared command will change to the shared state. A special command called invalidate is used to invalidate all caches in the system. A slave observing this command changes to the invalid state.

```
ASSIGN
    init(state).:= invalid;
    next(state) :=
        case
            abort : state;
            master :
                case
                    CMD = read-shared : shared;
                    CMD = read-owned : opned;
                CMD = urite-invalid : invalid;
                CMD = write-shared : shared;
                1 : state;
            esac;
        !master :
            case
                CMD = read-owned : invalid;
                CMD = invalidate & !waiting : invalid;
                CMD = read-shared & state = owned : shared;
                state = shared & !waiting : {shared,invalid};
                1 : state;
            esac;
        esac;
```

On receiving the command, each slave checks its own cache and indicates the state of the block in its own cache by asserting the signals reply-owned, and reply-waiting on the bus. These are wired or signals. meaning that the signal is observed to be asserted on the bus if one or

| from state | command | to state | cause |
| :---: | :---: | :---: | :---: |
| invalid | read-shared | shared | read miss |
| invalid <br> or shared | read-owned | owned | write miss |
| owned | write-invalid | invalid | copy-back |
| owned | write-resp-invalid | invalid | snoop read-owned |
| owned | write-shared | shared | write-through |
| owned | write-resp-shared | invalid | snoop read-shared |

Table 4.1: Summary of commands
more caches assert the signal. The reply-owned signal is asserted by a slave when the block is in the owned state in the slave's cache. Replywaiting is asserted when the slave has previously requested the block. and is waiting for a response. This signal will be discussed in more detail shortly. The process of looking up the slave's state and signaling on the bus is known as bus snooping. On observing a read command. a slave in the owned state sets a flag called snoop. This causes the cache to issue a write-response at a later bus cycle, supplying the data to the requester. and simultaneously storing it in main memory. When this happens. the snoop Hag is reset.

An additional reply signal called reply-stall may be asserted by any slave. including main storage. if the slave if not ready to respond to the command because some resource is busy. If reply-stall is asserted. the command is nullified.

```
DEFINE
    reply-owned := state = owned;
VAR
    snoop : boolean;
ASSIGN
    init(snoop) := 0;
    next(snoop) :=
```

```
case
    abort : snoop;
    state = owned & CMD = read-shared : 1;
    state = owned & CMD = read-owned : 1;
    CMD = response : 0;
    CMD = write-resp-invalid : 0;
    CMD = write-resp-shared : 0;
    1 : snoop;
esac;
```

After issuing a read command. the master releases the bus and waits for a response. During this time. a flag called waiting is set. Normally. if no slave asserts reply-owned. the response comes from main memory: If any slave asserts reply-owned. however. main memory is inhibited. allowing the slave to supply the data at a future cycle with a writeresponse command.

```
MODULE bus-device
VAR
    master : boolean;
    cmd : {idle,read-shared,read-owned,cty-read,write-invalid,
                write-shared,write-resp-invalid,write-resp-shared,
                invalidate,response};
    waiting : boolean;
    reply-stall : boolean;
ASSIGN
    init(waiting) := 0;
    next(waiting) :=
        case
            abort : waiting;
            master & CMD = read-shared : 1;
            master & CMD = read-owned : 1;
            CMD = response : 0;
            CMD = write-resp-invalid : 0;
            CMD = write-resp-shared : 0;
            1 : waiting;
        esac;
```

A slave which is waiting for a given cache block responds to any read command for that block by asserting reply-waiting. This nullifies the read command and forces the master to retry at a later cycle.

```
DEFINE
    reply-waiting := waiting;
    abort := REPLY-STALL
            | ((CMD = read-shared | CMD = read-owned)
                & REPLY-WAITING);
```

The commands which may be issued by a processor when it is bus master are a function $\mathrm{c}^{f}$ the state. For example, if the snoop flag is set. the processor mcy issue a write-response on the bus. From the owned state. a pr cessor may issue a write-invalid command in order to replace thn cache block with another. A processor in the shared state may issue a read-owned in case of a write miss. and a processor in the invalid state may issue either a read-shared or a read-owned command. in case of a read miss and write miss respectively.

```
MODULE processor(CMD,REPLY-OWNED,REPLY-WAITING,REPLY-STALL,DATA)
ISA bus-device
ISA cache-device
```

ASSIGN
cmd :=
case
master \& snoop \& state = invalid : write-resp-invalid;
master \& snoop \& state = shared : write-resp-shared;
master \& state $=$ owned \& !waiting : write-invalid;
master \& state $=$ shared \& !waiting : read-owned;
master \& state = invalid : \{read-shared,read-owned\};
1 : idle;
esac;

### 4.1.2 The local UIC interface

The I'IC' is the interface from one cluster to another. IIC's come in pairs. connected by a communication link. A EIC is said to be local
for a given memory block if that block is found in main storage on the same side of the link as the UIC. It is said to be remote if the memory block is found in a main memory on the other side of the link. Thus, for any memory block, one of the UICs in the pair is local, and the other remote. The UIC determines whether it is local or remote by address decoding. We consider the local case first. In this discussion. local refers to any part of the system on the bus side of the UIC, and remote refers to any part of the system on the link side of the UIC.

Viewed from the bus, the UIC behaves like a processor, with the capability to issue and respond to commands. The UIC"s cache records the state (but not the data) of all blocks of the local main storage that are present in remote caches. This allows the UIC to snoop the bus on behalf of remote caches. The UIC performs this function in exactly the same manner as the processors. The state of a block in the UIC changes with commands issued in the same manner as the state of cache blocks in processor caches.

The UIC receives command messages from from the link. and stores them in one of two queues. The low priority queue is for read commands. and the high priority queue is for all other commands. The depth of the queues is arbitrary. but for now. we consider queues of only one entry. A command in one of the queues is issued on the bus when the UIC becomes master. If both queues are non-empty. the command in the high priority queue is issued first. Provided the command is not aborted, the queue issuing the command is emptied. Since the UIC becomes bus master at nondeterministic intervals. the delay between the time a message arrives in the queue and is issued on the bus is arbitrary. This nondeterminism covers two abstractions made in the model. First, it allows for any amount of latency in the link level protocol, which is not modeled. Second. it allows the time to issue an arbitrary number of messages relating to other memory blocks that may be queued ahead of the one message that is modeled.

```
MODULE receiver
VAR
    hiq : {none,response,write-shared,write-resp-shared,
    write-invalid,write-resp-invalid,invalidate};
    loq : {none,read-owned,read-shared,cty-read};
```

```
ASSIGN
    cmd :=
        case
            master & !(hiq = none) : hiq;
            master & !(loq = none) : loq;
            1 : idle;
        esac;
    init(hiq) := none;
    next(hiq) :=
        case
            !master | abort : hiq;
        1 : none;
        esac;
    init(loq) := none;
    next(loq) :=
        case
            !master | abort | !(hiq = none) : loq;
        1 : none;
    esac;
```

The local UIC' can send commands to the link in response to commands observed on the local bus. Whenever a read command is sent to the link, it is entered in the remote EIC's low priority queue. If any other command is sent to the link. it is entered in the remote "IC's high priority queue. If the remote queue is full. the local bus cycle is stalled.

```
MODULE sender
DEFINE
    lopri := sending in {read-shared,read-owned,cty-read};
    hipri := sending in {invalidate,response,write-shared,
        write-invalid,write-resp-shared,urite-resp-invalid};
ASSIGN
    next(remote.hiq) :=
        case
            !abort & remote.hiq = none & hipri : sending;
            1 : remote.hiq;
        esac;
```

```
next(remote.loq) :=
    case
            labort & remote.loq = none & lopri : sending;
            1 : remote.loq;
    esac;
reply-stall :=
    (hipri & !(remote.hiq = none) |
        lopri & !(remote.loq = none)) union 1;
```

The local UIC sends command messages to the link in two cases. The first is to invalidate or call back cache blocks in remote caches. This occurs when the UIC is a slave and a read-owned or read-shared is received on the bus. If the UIC is in the owned state, the readowned or read-shared is forwarded to the link. This causes the remote cache in the owned state to issue a write-resp-invalid or write-respshared. returning the cache block to the local bus. If the UIC is in the shared state, and a read-owned is received on the bus. an invalidate command is forwarded to the link. This causes all remote caches to go to the invalid state. Note that this may allow a processor on the local bus to write before the invalidate command has reached all remote caches. This is a possible violation of strict consistency, which is tolerated for performance reasons. Hence, the protocol does not implement a strongly consistent memory model. The memory model which the protocol does support will be discussed in more detail in the next section.

The second case in which the local UC sends a command to the link is when the UIC has issued a read-shared or read-owned and is waiting for a response. In this case. if the UIC' is a slave and a response. write-resp-shared, or write-resp-invalidate is asserted on the bus. a response is sent to the link.

```
MODULE local-UIC(remote,CMD,REPLY-OWNED,REPLY-WAITING,
    REPLY-STALL,DATA)
ISA bus-device
ISA cache-device
ISA receiver
ISA sender
```

```
DEFINE
    sending :=
        case
            master : none;
            CMD = read-shared & state = ouned : read-shared;
            CMD = read-owned & state = owned : read-ouned;
            CMD = read-owned & state = shared : invalidate;
            CMD = write-resp-invalid & waiting : write-resp-invalid;
            CMD = write-resp-shared & waiting : write-resp-shared;
            CMD = response & waiting : response;
            1 : none;
        esac;
```


### 4.1.3 The Remote UIC interface

When the UIC is remote. it behaves as if it were a main storage device. It accepts read-shared. read-owned. write-shared. and writeinvalid commands from the bus. and forwards them to the local UIC' via the link. When the response arrives in the high priority queue. it issues the response on the bus. In addition. it can provide a special service to caches on the local side. If the remote I IC issues a read-shared or read-owned command. and there is no reply on the remote bus (ie.. no slave asserts reply-owned). it is assumed that the block was copied back to main storage while the read command was in transit. The remote [IC therefore sends the read rommand back to the local side. This operation is called a courtesy read. The courtesy read will cause the main store on the local bus to respond to the original requester.

```
MQDULE remote-UIC(remote,CMD,REPLY-OWNED,REPLY-WAITING,
    REPIY-STALL,DATA)
ISA bus-device
ISA receiver
ISA sender
DEFINE
    sending :=
    case
    master :
```

```
    case
    CMD = read-shared & !REPLY-OWNED : cty-read;
    CMD = read-owned & !REPLY-OWNED : cty-read;
    1 : none;
    esac;
!master :
    case
    CMD = read-shared & !REPLY-OWNED : read-shared;
    CMD = read-owned & !REPLY-OWNED : read-owned;
    CMD = write-resp-invalid & waiting : write-resp-invalid;
    CMD = write-resp-shared & waiting : write-resp-shared;
    CMD = write-resp-shared & !waiting : write-shared;
    CMD = write-shared : write-shared;
    CMD = write-invalid : write-invalid;
    1 : none;
    esac;
esac;
reply-omned := 0;
```

The text for the complete model in the SMIV language includes such details as an abstracted model of main storage and the cluster bus. which ties the above modules together. These are omitted here. Each cluster is modeled as an asynchronous process. Hence, the early quantification method for disjunctive relations can be used to avoid constructing the global transition relation (cf. section 2.4.2).

### 4.1.4 Protocol example

As an example of the protocol in operation. consider the sequence of events depicted in figures 4.2 and 4.3. In the figures. clusters 1 and 2 are both remote (ie., the memory block in question resides in some other cluster). The sequence begins when a read miss occurs in a processor in cluster 2 , while a processor in cluster 1 is in the owned state. It this point. the following sequence of events might occur:

1. The processor in cluster 2 issues a read-shared command on the bus. and sets its waiting Hag.
2. The UIC in cluster 2 sends the read-shared command up the link, storing it in the low priority queue of the global bus UIC for cluster 2.
3. The global bus UIC for cluster 2 issues the read-shared command on the global bus, entering the shared state, and setting its waiting flag.
4. Since the global bus IIC. for cluster 1 is in the owned state, it asserts reply-owned. sends the read-shared command down the link to cluster 1. enters the shared state, and sets its snoop flag.
5. The UIC in cluster 1 issues this read-shared command. entering the shared state and setting its waiting flag.
6. The processor in cluster 1 in the owned state asserts reply-owned. enters the shared state, and sets its snoop flag.
7. The processor in cluster l issues a write-resp-shared command. containing the block data. and clears its snoop flag.
8. The UIC in cluster 1 sends the write-resp-shared command up the link, storing it in the high priority queue of of the global bus IIC' for cluster 1. and clears its waiting flag.
9. The global bus ICIC for cluster 1 issues the write-response-shared command on the global bus. and clears its waiting flag.
10. (a) The global bus I.IC connected to main memory sends a writeshared command containing the block data and (b) The global bus UIC for cluster $\mathbf{2}$ sends a response command. clearing its waiting flag.
11. The IC in cluster 2 issues the response command.
12. The requesting processor in cluster 2 stores the data in its cache. and clears its waiting flag.


Figure 4.2: Protocol example


Figure 4.3: Protocol example (cont.)

### 4.2 Verifying the protocol

We now consider the problem of formal specification and verification of the protocol. The properties we will be concerned with are:

1. freedom from deadlock.
2. sequential consistency. and
3. local safety conditions. related to diagnostics.

Using the symbolic model checking technique. we can verify these properties automatically, despite the very large state space of the model. In fact, the model checker discovered a fairly subtle bug in the protocol an execution sequence leading to a deadlocked state.

### 4.2.1 Freedom from deadlock

We will say that the protocol is deadlocked if it reaches a state in which some processor is permanently blocked from receiving access to the given memory block. Thus. our definition of deadlock takes in situations that might also be called livelock. in which the system continues to loop infinitely, but without the possibility of making progress. We can express this property in (.TL with the following formula. which must hold for all processors:

$$
\begin{equation*}
A C_{1}^{\prime}(E F \text { readable } \wedge E F \text { writable }) \tag{4.1}
\end{equation*}
$$

In other words. it it alwavs possible that the memory block will become readable by the given processor. and always possible that it will become writable. We can check this property using SMV by adding the following specification to the processor module:

SPEC
AG(EF readable \& EF writable)
The specification turns out to be false, and as a counterexample. the model checker produces an execution trace leading to a deadlocked state. This is an actual bug in the original protocol which was found by
the model checker, but not in behavioral simulations. The complexity of the counterexample, and the unusual sequence of events that leads to the deadlock should give some indication of why this error would be unlikely to occur in random simulations. The time required to produce the counterexample was slightly under ten minutes running on a Sun 3/60.

The steps of the counterexample are depicted in figures 4.4 to 4.6 . Cluster 1 is the local cluster, and clusters 2 and above are remote clusters. We pick up the counterexample at a point where a processor in cluster 2 is in the owned state:

1. A read miss occurs in a processor in cluster 1 . This processor issues a read-shared command on the bus. It enters the shared state and sets its waiting flag.
2. Since the UIC in cluster 1 is in the owned state. it asserts replyowned, enters the shared state. and sends a read-shared command up the link. storing it in the low priority queue of the global bus UIC for cluster 1.
3. A processor in cluster 3 also issues a read-shared command. As a result, the global bus UIC for cluster 3 issues the read-shared command on the global bus, entering the shared state. and setting its waiting flag.
4. Since the global bus (IC for cluster - is in the owned state. it asserts reply-owned. sends a read-shared command down the link to cluster 2, enters the shared state. and sets its snoop flag.
5. The UIC in cluster 2 issues this read-shared command. entering the shared state and setting its waiting flag.
6. The processor in cluster -3 in the owned state asserts reply-owned. enters the shared state. and sets its snoop flag.
7. The processor in cluster 2 issues a write-resp-shared command. containing the block data. and clears its snoop Hag.
8. The UIC in cluster 2 sends the write-resp-shared command up the link, storing it in the high priority queue of of the global bus UIC for cluster 1, and clears its waiting flag.
9. The global bus UIC for cluster 2 issues the write-response-shared command on the global bus. and clears its waiting flag.
10. (a) The global bus [ IC' connected to main memory (cluster l) sends a write-shared command containing the block data and (b) The global bus UIC for cluster 3 sends a response command. clearing its waiting flag.
11. The UIC in cluster $I$ issues the write-shared command.
12. The block data are stored in main memory.
13. A processor in cluster 3 again issues a read-shared command. As a result, the global bus ['IC' for cluster 3 issues the read-shared command on the global bus. entering the shared state. and setting its waiting flag.
14. Since read-owned is not asserted. the VIC for cluster 1 sends the read-shared command down the link towards main memory.

At this point. the system is deadlocked. The original read-shared command sent in step $l$ in cluster $l$ is still in the low priority queue at th global bus level. but is stalled by the waiting flag set in the global UIC for cluster 3. Similarly, the read-shared command sent by cluster 3 is in the low priority queue in the cluster 1 CIC. but is stalled hy the waiting flag of the original requester. This is an example of the classic deadlock situation which occurs when two processes attempt to obtain locks on two resources (in this case two buses) in different orders. Vonetheless. the sequence of events that lead to this situation were sufficiently complex that the designers did not anticipate that the situation could occur. and simulations did not produce it. In fact. the deadlock situation was found at a search depth of thirteen transitions. It each step in this sequence. there were several alternatives that might have averted the deadlock. Thus it is possible. but unlikely that this


Figure 4.4: Deadlork example


Figure 4.5: Deadlock example (cont.)


Figure 4.6: Deadlock example (cont.)
deadlock would be found by a random simulation run. or a simulation run based on address traces. ${ }^{1}$

The fact the the model checker was able to print out automatically an example of this deadlock highlights an important practical aspect of the technique. Counterexamples are of perhaps even greater value than a proof that the system is correct. since such a proof is based on the assumption that the system is correctly modeled. and the specification is correct and complete. I counterexample. however. provides an important clue as to where a bug in the system lies. and how it might be corrected.

### 4.2.2 Correcting the deadlock

The problem causing the deadlock is that the remote owner of the memory block can write the data back to main memory while a read

[^17]command from the local cluster is in transit to the remote cluster. The write command crosses the read command in the mail, so to speak. A remote request for the same block can then lock the global bus, leading to deadlock. The Encore engineers corrected the deadlock problem in the following way. The write command. when it reaches the local bus. is converted by the UIC into a write-response command. This supplies data to the local requester and frees the local bus. Unfortunately. it also leaves an orphan read command in the svstem. If a read command from the remote side is issued on the local bus, and a remote processor subsequently reaches the owned state. the orphan read command will disrupt th protocol. To prevent this, when the orphan read is issued. it is converted to a special command called echo-response. which is sent back to the local cluster. The UIC' in the local cluster stalls any commands on the local bus until the echo-response arrives. thus guaranteeing that the orphan read command is destroyed.

The corrected model satisfies the absence of deadlock specification. The performance of the SMV model checker in verifying this is plotted in figure 4.7 . for a model with 2 clusters. as the number of caches in each cluster is increased from 2 to 6 (thus. in the largest model, there are 12 caches and 4 UIC's). Part (a) shows the run time as a function of the number of caches per cluster. Part (b) shows the number of OBDD nodes used overall. and for representing the transition relation. Part (c) shows the number of reachable states of the model. Although the run time points are well fit by a quadratic curve. the actual asymptotic performance is most likely cubic. as in the case of the synchronous arbiter (cf. section 2.4.1). owing to lincar increases in the transition relation size. the number of fixed point iterations and the size of the OBDDs representing fixed point approximations.

Since the number of bus wires running between successive caches is fixed. we can apply theorem 7 to show that the transition relation OBDD size must grow linearly in the number of caches. The fact that the fixed point approximation OBDDs also grow linearly bears further examination. however. This phenomenon can be understood by considering the nature of the protocol. Imagine cutting a cluster bus in half. and consider how much information must be communicated from one half of the bus to determine whether a given state of the system is in the reachable set or not. In fact. this amomet is fixed. independent of
the number of caches on the bus, since we need only know if there are any caches in the shared state or the owned state on the other side of the cut, and not in particular which caches these are or how many. As a result, the number of OBDD nodes (representing the reached state set) at the level corresponding to our cut is bounded. ${ }^{2}$ This is characteristic of bus snooping protocols. and other protocols which are "loosely coupled", in the sense that one half of the system has bounded knowledge of the state of the other half of the system.

As part (c) of the figure shows. the number of states of the system increases exponentially with the number of caches per cluster. Despite this, the performance of the symbolic model checking algorithm is polynomial. Thus, for this particular model and specification, we have solved the state explosion problem.

### 4.2.3 Sequential consistency

When writing a formal specification for the Gigamax cache consistency protocol. we need to consider the model of a distributed memory which the Gigamax provides to the programmer. As mentioned previousiv, for performance reasons the protocol does not maintain strict consistency of the caches. A cache block in the shared state may be out of date for a short time while an invalidate message is traversing the network. This is tolerated, since maintaining strict consistency would require an acknowledgment of invalidation to be collected from all caches in the shared state before a cache block could be modified.

There are a number of distributed memory models that may be supported by such a system. A totally ordered model is one in which all processors observe all values written to the memory in the same order. For example. in a totally ordered model. if the processors write into a location the sequence of values $1,2.3, \ldots$ then all processors which read the location will observe any new values to be greater than or equal to all previous values. We will show that the Gigamax protocol has this property, for a one block system. In a partially ordered model. values written may in some cases be observed in a different order by different processors. Some guarantee of ordering is usually made. For

[^18]

Figure 4.7: Performance for cherking deadlock
example, all writes must be observed in the same order relative to special synchronization operations, or all writes by the same processor must be observed in the same order. The latter model is supported by the Gigamax protocol for writes to different cache blocks. Since our model of the protocol only describes the behavior one cache block, however, the model cannot be used to check this property.

Returning to the problem of total ordering of writes to the same block, it might seem at first that there is no "finite state" description of a protocol that writes an unbounded sequence of values. We can check the property, however. by using an abstraction. We do this by choosing a value $n$, and storing in the model only one bit of information - whether the data value is less than $n$ or greater than or equal to $n$. We then assume that the processors never write a value less than $n$ after a value greater than $n$ has been written. and we show that a processor never reads a value less than $n$ after reading a value greater than $n$. Since the value of $n$ is arbitrary, it follows that all processors read data values in non-decreasing order. satisfying the total ordering requirement. We now consider how to model the system using this abstraction. For each cache, we introduce a variable whose value is 0 when the data value is less than $n$ and 1 when the date value is greater than or equal to $n$. This variable may change whenever the block is writable. but may only change from 0 to 1 . since we assume the processors only increase the data value. The following SMV code models the data held in the processor's cache:

```
MODULE data-device
VAR
    data : boolean;
ASSIGN
    next(data) :=
        case
            !master & waiting & CMD in {response,write-resp-invalid,
                write-resp-shared} : DATA;
            writable : data union 1;
            1 : data;
        esac;
DEFINE
    data-enable := master & CMD in {response,write-resp-invalid,
```

```
write-resp-shared,write-invalid};
```

Additionally, we introduce variables to represent the values on the buses and the values in the high priority message queues. The low priority queues hold only requests, which have no data value.

We would now like to prove that this abstract model of the data path of the protocol satisfies the following specification in CTL. for all processors:

$$
\begin{equation*}
A G[(\text { readable } \wedge \text { dat } a \geq n) \Rightarrow A(i \neg(\text { readable } \wedge \text { dat } a<n)] \tag{4.2}
\end{equation*}
$$

In other words. if ever a value greater or equal to $n$ is observed. a value less than $n$ is never observed in the luture. We can check this using SMV by adding the following specification to the processor module:

```
SPEC
    AG(readable & data -> AG (readable -> data))
```

Figure 4.8 shows the performance of the symbolic model checking algorithm in verifying this formula. again for a model with 2 clusters. Part (a) of the figure shows the execution time. while part (b) shows the amount of storage used. Notice that although the execution times are roughly ten times those obtained for the model without data. they are still cubic in the number of processors per cluster.

### 4.2.4 Correctness of diagnostics

In addition to the above specifications. it was also particularly useful to check that the diagnostics built into the protocol never flagged an error under normal operation of the protocol. Errors are flagged be the diagnostic system in each processor subsestem whenever a command is observed on the bus which is inconsistent with the processors local state. Determining which command/static combinations are normal. and which are errors is difficult. and a number of errors of this type were found in the protocol using the model checking technique.


Figure 4.3: Performance for checking sequential consistency

### 4.3 Evaluation

In verifying the Gigamax model with respect to the formal specifications, the symbolic model checker was able to perform an exhaustive search of the model's state space without explicitly constructing the global state graph. As a result, the state explosion problem was avoided. In addition, the model checker exposed a number of subtle errors in the design that were not found in simulation. These errors were usually caused by events (eg.. cache misses and message arrivals) occurring out of the normal sequence anticipated by the designers. This type of error is difficult to find in random simulations. since the probability of a given sequence of random events occurring by pure chance is in inverse exponential proportion to the length of the sequence. Awe have seen. the sequences necessary to produce protocol errors can be quite long. As the design evolved to correct the errors found by model checking, the model was easily adapted. and quickly provided an analysis of any new errors introduced by design changes. This tend to amortize the initial effort required to produce the protocol model. The ability of the symbolic model checker to find errors quickly makes it easier to experiment with alternative designs, and also helps to build the designer's intuition about the behavior of the system. This is important. because designers tend to concentrate on normal sequences of events. and overlook the unusual sequences. The use of OBDDs in the symbolic model checker made it possible to check a model that would have been very time consuming. or perhaps impossible to check using earlier algorithms.

At this point. the technique has a number of limitations. (ne limitation is the use of OBDDs. For example. while we find the OBDD sizes growing polynomially in the number of caches in the Gigamax model. if we instead increase the number of cache blocks and leave thenumber of caches constant. we find the size of the OBDDs increasing, exponentially. As a result. it was extremely difficult to check specifications of a system with just two cache block a some runs took up to a week. and others never finished). In these cases, the size of the OBDDs representing the fixed point approximations became intractabis large. When this happens. techmiques such as early quantification that make the representation of the transition relation smaller are little nse.
since they do not effect the size of the OBDDs representing fixed point approximations.

Another major issue is implementation of the protocol. Clearly, verification of the protocol itself is important, since a correct protocol is a prerequisite for a correct implementation. This is, of course. only half the story. Techniques are also needed to insure that the verified protocol is implemented correctly in hardware. This can in fact be done. using a process of successive refinement of finite state systems that has been studied extensively by Kurshan [Kursi]. The work of Bose and Fisher [BF89a] is also an example of this. Unfortunately, the truth of CTL formulas containing existential quantifiers is not necessarily preserved by this kind of refinement. Thus. for example, though the high level protocol may be deadlock free, a specific implementation of the protocol may not be deadlock free ${ }^{3}$. In order for the implementation to preserve all CTL properties of the protocol. the two would have to be bisimular (cf. section 2.6.1). Since this is a very strong requirement. it is not clear that the protocol could in fact be implemented with this degree of accuracy in an efficient way. For essentially this reason. Grumberg and Long have studied the use of a subset of CTL using only universal path quantifiers for hierarchical reasoning [CLD91]. In any event. though checking the absence of deadlock specification was very useful in fincling bugs in the protocol, we must attach a special caveat to this result, since it does not guarantee that all reasonable implementations of protocol will be deadlock free.

Finally, there is the problem of verifving a model with a finite number of processors. when there is no finite limit on the number of processors that could in principle be added to the system. In practice, the intended maximum number of processors is approximately 100 . Even using the symbolic model checking technique, however. checking a system of 100 processors seems infeasible at present. and 1000 processors is out of the question. [o deal with systems with a very large num-

[^19]ber of identical components, we can apply methods of induction over processes. As in the case of successive refinement. induction methods are not fully automatic - some human input is required in the form of an inductive hypothesis. In the next chapter, we will deal with the problem of induction over processes.

## Chapter 5

## Induction and model checking

This chapter deals with the verification of systems that have an arbitrary number of similar components. arranged in some inductively defined structure. Systems of this type are commonplace - they occur in bus protocols and network protocols. I/O channels. and many other structures that are designed to be extensible by adding similar components. Atter using a model checking system to determine the correctness of a system configured with a fixed number of processors or other components. it is natural to ask whet her this number is enough in some sense to represent a system with any number of components. For example, a Cigamax system can be built by connecting some arbitrary number of cluster buses to a global bus. then filling each cluster bus with an arbitrary number of processor cards. It is practically impossible to verify using model checking methods alone that all possible configurations of the system satisfy the specifications, even given a physical bound on the number of cards in a backplane. However. by supplying an appropriate inductive hypothesis, we can in many cases reduce the problem of verifying a system of arbitrary size to one of verifying a system of fixed size. The inductive hypothesis can take the form of a finite state process.

### 5.1 The general framework

Induction over systems of processes can be put in a fairly general framework. which is independent of the mechanics of the process model. relying only on certain algebraic properties of the operators for combining processes. Let us assume that we have a collection of processes. and a collection of operators acting on processes. In a typical process model. we have some form of parallel composition operator. some form of operator for renaming signals, and perhaps a hiding operator. which makes a given signal invisible to the outside. The exact choice of operators is not material here. however. We require only that the operators be monotonic with respect to a reflexive transitive relation $\leq$ on processes. The idea of this order is that if $p \leq q$, the $p$ is in some sense more specific. or more deterministic, than $q$. The properties we wish to verify should be preserved as we descend the order.

As an example of induction on processes. suppose we have a parallel composition operator $\|$ on processes. which is monotonic with respect to a pre-order $\leq$. In this case. we can apply the following induction rule:

$$
\begin{gathered}
p \leq q \\
q \| p \leq q \\
p\|\cdots\| p \leq q
\end{gathered}
$$

Think of the inequalities $p \leq q$ and $q \| p \leq q$ as substitution rules If $p \leq q$. we can safely substitute $p$ for any occurrence of $q$ in a given term. in the sense that we will only make the cerm lesser in the partial order. Thus. we can always substitute $p$ for 4 on the lesser side of an inequality. For example. if $q \| p \leq q$. we have

$$
\begin{aligned}
q \| p & \leq q \\
(q \| p) \| p & \leq q \\
((q \| p) \| p) \| p & \leq q
\end{aligned}
$$

If $p \leq q$. we can substitute $p$ for $q$. giving us $p \| \cdots$ | $p \leq 4$. We call $q$ a process invariant. Other induction menes can be generated. based on other substitutions. For example, assume we have a parallel


Figure 5.1: Processes generated by safe substition
composition operator || and a renaming operator $\varphi$. both monotonic with respect to $\leq$. Then we have

$$
\frac{O(q)\|r\| p \leq q}{O(\phi(\cdots)\|r\| p)\|r\| p \leq q}
$$

Given a collection of substitution rules, we can inductively generate a class of processes from any process q. For example, figure ... 1 depicts the first few processes in the class generated by the above induction rule. Every process in the class is smaller than $q$ in the partial order. Thus, any properties of $q$ which are preserved as we descend the partial order are inherited by all the processes in the class. The key, then. is to choose a partial order that preserves the properties we are interested in verifying. The most straightforward way to do this is to choose a class of properties that we wish to preserve, and then define the partial order accordingly.

For example. suppose we wish to preserve all properties expressible in the logic CTL. In this case. the partial order we obtain is a degenerate one, which partitions the Kripke models into a set of incomparable equivalence classes. To see this. assume towards a contradiction that $p$ satisfies every (TL formula satisfied by $q$, and there is some formula $f$ satisfied by $p$ but not by $\%$. In this case. it follows that $y$ satisties $-f$. This implies. however. that $p$ satisfies $\neg f$. a contradiction. Hence. if $p \leq q$, then $p$ and $q$ satisfy the same set of CTL formulas. Since CTL characterizes Kripke models up to bisimulation [BCGS7]. it follows that $p$ and $q$ are bisimular. This is unfortunate. since we do not want our induction framework to apply only to classes of Kripke structures that are equivalent. In general. we would like to treat systems whose behavior becomes more specific as we add processes to the system.

One way to do this is to use a subset of the logic. For example. suppose we choose to preserve those formulas which use only universal
path quantifiers. This subset is called $\forall$-C'TL [GL91]. A formula in CTL is also in $\forall$-CTL if driving the negations in to the literals results in a formula without the $E$ path quantifier. Examples of $\forall$-CTL formulas are

$$
\begin{aligned}
A G \neg E G p & \equiv A G A F \neg p \\
\neg E G E X p & \equiv A F A X \neg p
\end{aligned}
$$

Examples of CTL formulas which are not in $\forall$-CTL are

$$
\begin{aligned}
A G \neg A F p & \equiv A G E(i \neg p \\
\neg E G A X p & \equiv A F E X \neg p
\end{aligned}
$$

C'learly, if a formula $f$ contains path quantifiers. then $f$ and $\neg f$ cannot both be in $\forall$-CTL. Grumberg and Long [GL91] have shown that if $p$ satisfies every $\forall$-CTL formula satisfied by $q$. then $q$ simulates $p$, and conversely. Simulation is easily shown to be reflexive and transitive. Thus simulation is a pre-order suitable for inductive proois of $\forall$-CTL formulas. Let $p \leq q$ iff $q$ simulates $p$. This gives us the following induction rule:

$$
\begin{gathered}
q \vDash f \quad(f \in \forall \cdot(\text { C'TL }) \\
p \leq q \\
q \| p \leq q \\
\hline p\|\cdots\| p \vDash l
\end{gathered}
$$

as well as other rules engendered be various systems of safe substitutions. Recall from section 2.6 .1 that simmlation is the greatest relation between the states of $q$ and the states of $p$ such that if $x$ simulates $y$. then:

1. $x$ and $y$ agree on the atomic propositions. and
2. every successor of $y$ is simulated by a successor of $x$.

A Kripke model $q$ simulates $p$ if every initial state of $p$ is simulated by. some initial state of $q$. Since this relation can be expressed as a greatest fixed point in the Mu-Calculus. it can be verified antomatically usine the symbolic model checking technique. The fact that simulation in not symmetric allows us more flexibility in constructing systems usine substitution rules than we would have using bisimulation.

### 5.1.1 Induction in other models

We can set up an induction framework for a variety of models by establishing a pre-order and a set of monotonic process operators. For models of concurrent automata (such as the $s / r$ model [Kur85]), the natural partial order is language containment.

In the $s / r$ model, a process is an automaton which accepts infinite strings over a Boolean algebra. There are two natural operators over this class of processes. The automaton product operation simulates parallel execution, while Boolean algebra homomorphisms can be used to induce a renaming or abstraction of the variables by which processes communicate. Kurshan shows that both of these operations respect the relation of language containment between automata [Kur86]. An example of induction in this framework can be found in [K.M89].

Induction can also be applied in process algebras tike CCS [Mil80] which are based on two-way synchronization. In this case. there is a variety of plausible process relations. including observational equivalence. weak observational equivalence. and a number of pre-order relations on processes. An induction example using the "may" pre-order for CC'S processes can also be found in [K.M89].

### 5.2 Induction and SMV

In induction framework can be set up for the SMV. 0 language. based on either simulation or bisimulation. This framework includes two kinds of process operators - the parailel composition operator \| and renamings operators hased on maps ofrom locations to locations. We will show that both are monotonic with ressect to simulation and bisimulation.

### 5.2.1 Proving compositionality

Recall that semantically, an SMV. 0 program denotes a triple (T.I. R). where $T$ assigns types to locations. $I$ is the set of initial states. and $R$ is the transition relation. There are two basic process operators provided by S.MV: instantiation and parallel composition. An instantiation results from a map on on locations (a renaming). This renaming
induces a map $\Phi$ from states of $\phi(M)$ to states of $M$, such that for all locations $l, \Phi(x)(l)=x(\phi(l))$. If $M=(T . I . R)$ is a process, then $\phi(M)=\left(T^{\prime}, I^{\prime}, R^{\prime}\right)$ where

1. $T^{\prime}(\phi(l))=T(l)$,
2. $I^{\prime}=\{x \mid \Phi(x) \in I\}$ and
3. $R^{\prime}=\{(x, y) \mid(\Phi(x), \Phi(y)) \in R\}$.

The rules for renaming require that no hidden location can be renamed to a visible location, and no two distinct locations can be renamed to the same hidden location. The following lemma and theorem show that this definition of renaming is a suitable operation for inductive reasoning using simulation or bisimulation:

Lemma 4 Let 0 be a legal renaming. and let $\Phi$ be the state map induced by $\phi$. If $x_{1}^{\prime}=\Phi\left(x_{1}\right)$ and $x_{1}^{\prime}$ agrees with $r_{2}^{\prime}$ on the visible locations. then there exists $x_{2}$ which agrees with $x_{1}$ on the visible locations. such that $x_{2}^{\prime}=\Phi\left(x_{2}\right)$.

Proof. Construct $x_{2}$ as follows: For every location $l$. if $l$ is in the range of $\phi$. then choose any $l^{\prime}$ such that $l=o\left(l^{\prime}\right)$. and let $x_{2}(l)=r_{2}^{\prime}\left(l^{\prime}\right)$. Otherwise, let $x_{2}(l)=x_{1}(l)$.

First we show that $x_{2}^{\prime}=\Phi\left(x_{2}\right)$. For all locations $l^{\prime \prime}$. if $\dot{\phi}\left(l^{\prime \prime}\right)$ is visible. then $l^{\prime \prime}$ must also be visible (since hidden locations cannot legally be renamed to invisible locations). Let $l=o\left(l^{\prime \prime}\right)$. Since $l$ is in the range of 0 . there is some visible $l^{\prime}$ such that $l=o\left(l^{\prime}\right)$ and $x_{2}(l)=x_{2}^{\prime}\left(l^{\prime}\right)$. Since $x_{1}^{\prime}$ and $x_{2}^{\prime}$ agree on the visible locations. $x_{2}^{\prime}\left(l^{\prime}\right)=x_{1}^{\prime}\left(l^{\prime}\right)$. Since $x_{1}^{\prime}=\Phi\left(x_{1}\right), x_{1}^{\prime}\left(l^{\prime}\right)=x_{1}(l)=x_{1}^{\prime}\left(l^{\prime \prime}\right)=x_{2}^{\prime}\left(l^{\prime \prime}\right)$. Thus $x_{2}^{\prime}\left(l^{\prime \prime}\right)=x_{2}\left(o\left(l^{\prime \prime}\right)\right)$. On the other hand. if $\phi\left(l^{\prime \prime}\right)$ is hidden. there exists no other location $l^{\prime \prime}$ such that $\Phi\left(l^{\prime}\right)=\phi\left(l^{\prime \prime}\right)$. since two locations cannot legally be renamed to the same hidden location. Therefore $x_{2}\left(0\left(l^{\prime \prime}\right)\right)=x_{2}^{\prime}\left(l^{\prime \prime}\right)$. Thus. by definition $x_{2}^{\prime}=\Phi\left(x_{2}\right)$.

Second. we show that $x_{1}$ and $x_{2}$ agree on the visible locations. Let / be any visible location. If $l$ is not in the range of 0 , then $x_{1}(l)=x_{2}(l)$ by construction. Otherwise. there is an $l^{\prime}$ such that $o\left(l^{\prime}\right)=l$ and $x_{2}(l)=$ $x_{2}^{\prime}\left(l^{\prime}\right)$. Since $l^{\prime \prime}$ must be visible, $x_{2}^{\prime}\left(l^{\prime}\right)=x_{1}^{\prime}\left(l^{\prime}\right)=x_{1}\left(o\left(l^{\prime}\right)\right)=x_{1}(l)$. $\quad$.

Theorem 9 Instantiation of SMV. 0 modules is monotonic with respect to simulation and bisimulation.

Proof. Imagine that we have two processes, $M_{1}$ and $M_{2}$, such that . $M_{2}$ simulates $M_{1}$.

First, let $x_{1}$ and $x_{2}$ be states of $\phi\left(M_{1}\right)$ and $\phi\left(M_{2}\right)$ respectively. We show by induction that if $x_{1}$ and $x_{2}$ agree on the value of all visible locations, and if $\Phi\left(x_{2}\right)$ simulates $\Phi\left(x_{1}\right)$, then $x_{2 .} n$-simulates $x_{1}$, for all $n$.

The basis case is trivial. since $x_{1} 0$-simulates $x_{2}$ exactly when $x_{1}$ and $x_{2}$ agree on the value of all visible locations.

For the induction step, let $\left(x_{1}, y_{1}\right)$ be any transition of $\phi\left(M_{1}\right)$. By definition, $\left(\Phi\left(x_{1}\right), \Phi\left(y_{1}\right)\right)$ is a transition of $: M_{1}$. Hence, there exists a transition $\left(\Phi\left(x_{2}\right), y_{2}^{\prime}\right)$ of $M_{2}$ such that $\Phi\left(y_{1}\right)$ simulates $y_{2}^{\prime}$. Since $\Phi\left(y_{1}\right)$ simulates $y_{2}^{\prime}, \Phi\left(y_{1}\right)$ and $y_{2}^{\prime}$ agree on the visible locations. By the lemma. there must exist $y_{2}$ which agrees with $y_{1}$ on the visible locations, such that $y_{2}^{\prime}=\Phi\left(y_{2}\right)$. By inductive hypothesis. $y_{2}(n-1)$-simulates $y_{1}$, therefore $x_{2} n$-simulates $x_{1}$.

Now we show that every initial state of $\phi\left(M_{1}\right)$ is simulated by some initial state of $O\left(M_{2}\right)$. $A$ state $x_{1}$ is initial in $O\left(M_{1}\right)$ exactly when $\Phi\left(x_{1}\right)$ is initial in $M_{1}$. Let $x_{1}^{\prime}=\Phi\left(x_{1}\right)$. If $x_{1}^{\prime}$ is initial in $M_{1}$. then it is simulated by some $x_{2}^{\prime}$ which is initial in.$M_{2}$. Since $x_{1}^{\prime}$ is simulated by $x_{2}^{\prime}$, they agree on the values of the visible locations. Hence. by the lemma. there exists $x_{2}$ which agrees with $x_{1}$ on the visible locations. such that $x_{2}^{\prime}=\Phi\left(x_{2}\right)$. By the above argument. $x_{1}$ is simulated by $x_{2}$. Therefore. $\Phi\left(M_{1}\right)$ is simulated by $o\left(M_{2}\right)$.

We can prove that renaming respects bisimulation by the same argument. applied symmetrically:

The parallel composition of SMV. 0 programs is formed by taking the union of the type functions $T$ and the intersections of the initial sets $I$ and the transition relations $R$. after renaming the hidden variables of each process onto disjoint spaces. We can show that this operation is also suitable for inductive reasoning:

Theorem 10 Parallel composition of SMV. 0 programs is monotonic with respect to simulation and bisimulation.

Proof. Imagine that we have four processes, $M_{1}, M_{2}, M_{1}^{\prime}$ and $M_{2}^{\prime}$. such that $M_{1}$ is simulated by $M_{2}$ and $M_{1}^{\prime}$ is simulated by $M_{2}^{\prime}$.

Let $\phi$ and $\phi^{\prime}$ be the renamings associated with parallel composition. These are identities over the visible locations, and map the hidden locations onto disjoint ranges. Let $\Phi$ and $\Phi^{\prime}$ be the induced state maps. (Thus, given a state of a parallel composition $M \| M^{\prime}, \Phi$ yields the corresponding state of $M$ and $\Phi^{\prime}$ yields the corresponding state of $M^{\prime}$.) Let $x_{1}$ be a state of $M_{1} \| M_{1}^{\prime}$ and let $x_{2}$ be a state of $M_{2} \| M_{2}^{\prime}$. We show by induction that if $\Phi\left(x_{1}\right)$ is simulated by $\Phi\left(x_{2}\right)$ and $\Phi^{\prime}\left(x_{1}\right)$ is simulated by $\Phi^{\prime}\left(x_{2}\right)$, then $x_{1} n$-simulates $x_{2}$, for all $n$.

For the base case, since $\Phi$ is the identity for the visible locations. $\Phi\left(x_{1}\right)$ agrees with $x_{1}$ on the visible locations, as does $\Phi\left(x_{2}\right)$ with $x_{2}$. Since $\Phi\left(x_{1}\right)$ is simulated by $\Phi\left(x_{2}\right)$, they also agree, therefore $x_{1} 0$ simulates $x_{2}$.

For the induction step, let $u_{1}=\Phi\left(x_{1}\right), u_{1}^{\prime}=\Phi^{\prime}\left(x_{1}\right), u_{2}=\Phi\left(x_{2}\right)$. $u_{2}^{\prime}=\Phi^{\prime}\left(x_{2}\right)$. Let $\left(x_{1}, y_{1}\right)$ be a transition of $M_{1} \| M_{1}^{\prime}$. and let $v_{1}=\Phi\left(y_{1}\right)$ and $v_{1}^{\prime}=\Phi^{\prime}\left(y_{1}\right)$. By definition. ( $\left.u_{1}, v_{1}\right)$ is a transition of $M_{1}$ and ( $u_{1}^{\prime}, v_{1}^{\prime}$ ) is a transition of $M_{1}^{\prime}$. Since $u_{1}$ is simulated by $u_{2}$ and $u_{1}^{\prime}$ is simulated by $u_{2}^{\prime}$, there must exist $v_{2}$ and $v_{2}^{\prime}$ such that ( $u_{2} \cdot v_{2}$ ) is a transition of $M_{2}$. $\left(u_{2}^{\prime}, v_{2}^{\prime}\right)$ is a transition of $M_{2}^{\prime} . v_{1}$ is simulated by $v_{2}$ and $c_{1}^{\prime}$ is simulated by $c_{2}^{\prime}$. Now we construct $y_{2}$. For all visible locations $l$. let $y_{2}(l)=y_{1}(l)$. For all hidden locations $l$ in the range of 0 . there is a unique $l^{\prime}$ such that $l=\phi\left(l^{\prime}\right)$, since a renaming cannot legally map distinct locations on to the same hidden location. Let $y_{2}(l)=c_{2}\left(l^{\prime}\right)$. Similarly, for all hidden locations $l$ in the range of $\Phi^{\prime}$. there is a unique $l^{\prime}$ such that $l=o^{\prime}\left(l^{\prime}\right)$. Let $y_{2}(l)=v_{2}^{\prime}\left(l^{\prime}\right)$. By this construction. $r_{2}=\Phi\left(y_{2}\right)$ and $c_{2}^{\prime}=\Phi^{\prime}\left(y_{2}\right)$. Hence. by inductive hypothesis. $x_{2}(11-1)$-simulates $y_{2}$. By definition. $\left(x_{2}, y_{2}\right)$ is a transition of $M_{2} \| M_{2}^{\prime}$. Therefore $x_{1} n$-simulates $x_{2}$.

Now we show that every initial state of $M_{1} \| M_{1}^{\prime}$ is simulated by an initial state of $M_{2} \| M_{2}^{\prime}$. Let $x_{1}$ be an initial state of $M_{1} \| . M_{1}^{\prime}$ and let $u_{1}=\Phi\left(x_{1}\right), u_{1}^{\prime}=\Phi^{\prime}\left(x_{1}\right)$. By definition. $u_{1}$ is initial in $M_{1}$ and $u_{1}^{\prime}$ is initial in $M_{1}^{\prime}$. Hence there exist $u_{2}$ and $u_{2}^{\prime}$, such that $u_{1}$ is simulated by $u_{2}$. $u_{1}^{\prime}$ is simulated by $u_{2}^{\prime} . u_{2}$ is initial in $M_{2}$ and $u_{2}^{\prime}$ is initial in $M_{2}^{\prime}$. We can construct $x_{2}$ such that $u_{2}=\Phi\left(x_{2}\right)$ and $u_{2}^{\prime}=\Phi^{\prime}\left(r_{2}\right)$ in the same manner as we constructed $y_{2}$ above. $B y$ the above argument. $x_{1}$ is simulated by $x_{2}$.

We can prove that parallel composition respects bisimulation by the same argument, applied symmetrically.

### 5.2.2 Computing simulation relations

Since the simulation and bisimulation relations can be expressed in the Mu-Calculus (cf.. section 2.6.1). they can be computed using the symbolic model checking technique. In this way, we can automatically test whether substituting a given module $p$ for another module $q$ is safe. in the sense of preserving all CTL or $\forall$ C'TL properties.

There are a few techniques that can improve the efficiency of this process. The simplest is to note that simulation between two states implies that they agree on the values of the visible locations. Therefore, there is no need to use separate OBDD variables to encode the visible locations of the two processes when representing the simulation relation. As with C'TL model checking, we can compute the reachable state space of the two programs. and use these sets to restrict the computation of the equivalence relation. In cases where the simulation relation cannot be computed. we can instead compute a stronger relation between the programs. which requires that all pairs of states which are simultaneously reachable (reachable along paths which agree on the visible locations) are 1 -simular. This relation can be tested by a forward search of the reachable state space of the composition of the two programs. In the case of deterministic programs (in which no two successors of a given state agree on all of the visible locaticns). this amounts to a test of string language containment [GL91]. In either approach. if the test fails. we can extract as a counterexample a pair of paths, such that all corresponding states are 0 -simular. and the last pair fails to be 1 -simular. This test can be used to formulate another guess for the process invariant. until a sound invariant is found.

SWV supports induction in the following way. Each hypothesis of an induction rule is of the form $p \leq q$, where $p$ and $q$ are modules. This is completely general. since module $p$ can be an arbitrary parallel romposition of instances of other modules. By inserting the declaration
in module $q$, we cause the SMV model checker to test whether $q$ simulates $p$ and if not, to produce a counterexample.

### 5.2.3 Induction and SMV. 1

Is it possible to extend the above framework to SMV.1, which includes interleaving processes? Unfortunately. the answer is no. Consider. for example. the following two modules. which are bisimular:

```
MODULE a
var
    x : boolean;
ASSIGN
    init(x) := 0;
    next(x) := 0;
```

MODULE b
x : boolean;
ASSIGN
init $(x)$ := 0 ;
next $(x)$ : $=x$;

Note, however, that if we substitute a for $b$ in the following program. the resulting program is not bisimular to the original:

```
MODULE main
VAR
    p : process b;
ASSIGN
    next(p.x) := 1;
```

This is because process main may intervene between steps of process p. changing the value of $p . x$ to 1 . In this state. which is not reachable in $a$ or $b$ alone, the two modules have different behaviors. Hence parallel composition in S.VV. 1 does not respect bisimulation (neither does it respect simulationi. This problem is a general feature of languages that support interleaving processes with shared variables. It is difficult. for example. to formulate a compositional mule hor the leadn-to operator of WNITY logic [CMB8]. For this reason. we will nse onty the S.MT: I subset for induction over processes.


Figure 5.2: Substitution generating processors on bus

### 5.3 Example: The Gigamax protocol

In this section, we formulate a safe substitution rule that generates an arbitrary number of Gigamax processors attached to a cluster bus. As our invariant, we will use a single processor module attached to the end of a cluster bus. The strategy will be to generalize this module by adding nondeterminstic choice until it is able to simulate itself with one additional processor attached, as viewed from the bus. The counterexamples produced by the model checker will provide clues as to how the proposed invariant should be generalized. After a correct invariant is obtained, we can use this invariant to prove properties of the protocol that hold independent of the number of processors on a cluster bus.

The general form of the substitution rule we use is depicted in figure 5.2. The SMV code representing the left hand side is shown in figure figure 5.3. and the code for the right hand side (the invariant) is shown in figure j.t. In our first guess for the invariant, we will use the original processor model from the previous chapter. Our approach will be to add behaviors (if.. non-determinism) to the processor model until we have a correct invariant.

Essentially, we are testing whether one processor can mimic the actions of two processors as seen from the bus. Checking this produces a counterexample in which one of the two processors reaches the owned state. then the second processor issues a read command. This behavior cannot be produced by a single processor. To fix this problem. we can modify the processor model so that a processor is allowed to issue a read command in the owned state. It then sets its own "snoop" flag. and enters the shared state on a read-shared. and the owned state on a read-owned. Testing this ne:v invariant produces another counterex-

```
MODULE rule(CMD,REPLY-OWNED,REPLY-WAITING,REPLY-STALL,
    cmd,reply-owned,reply-waiting,reply-stall,master)
VAR
    b : bus-connector(self,p);
    p : processor(CMD, REPLY-OWNED, REPLY-WAITING, REPLY-STALL);
    q: invariant(CMD, REPLY-OWNED, REPLY-WAITING, REPLY-STALL,
        b.cmd,b.reply-owned,b.reply-waiting,b.reply-stall,b.master);
MODULE bus-connector(a,b)
ASSIGN b.master := !a.master union 0;
DEFINE
    cmd :=
        case
            a.master : a.cmd;
            b.master : b.cmd;
            1 : idle;
        esac;
    reply-owned := a.reply-owned | b.reply-owned;
    reply-waiting := a.reply-waiting | b.reply-waiting1;
    reply-stall := a.reply-stall | b.reply-stall1;
    master := a.master | b.master;
```

Figure 5.3: Substituion rule for adding one processor

```
OPAQUE MODULE invariant(CMD,REPLY-OWNED,REPLY-WAITING,
    REPLY-STALL, cmd,reply-owned,reply-waiting,reply-stall,master)
SIMULATES rule
VAR
    b : bus-connector(self,p);
    p : processor(CMD, REPLY-OWNED, REPLY-WAITING, REPLY-STALL);
    t : bus-terminator(CMD,REPLY-OWNED,REPLY-WAITING,REPLY-STALL,
        b.cmd,b.reply-owned,b.reply-waiting,b.reply-stall,b.master);
MODULE bus-terminator(CMD,REPLY-OWNED,REPLY-WAITING,
    REPLY-STALL,cmd,reply-owned,reply-waiting,reply-stall,master);
ASSIGN
    CMD := cmd;
    REPLY-OWNED := reply-owned;
    REPLY-HAITING := reply-waiting;
    REPLY-STALL := reply-stall;
```

Figure 5. t : The invariant
ample in which the first processor reaches the owned state. then issues a read command (thus setting its snoop and waiting bits). then the second processor issues a read command. One processor alone cannot produce this behavior. since it cannot issue a second read command while its waiting flag is set. We modify the processor model to allow this behavior. Note that this is behavior is safe, since the second read command is blocked by the waiting flag which is already set. With this modification. we have a correct invariant.

Ving this invariant. we can check properties of the system in $\forall$ ('TL, using the invariant in place of the processors on the cluster buses. The substitution rule can be applied as many times as necessary to produce a system with an arbitrary number of processors while preserving all of the verified properties. We can also refer these properties back to our original morlel by showing that the generalized processor model simulates the original one. In order to verity properties such as deadlock freedom. however. which use existential path quantifiers. it would be necessary to prove bisimulation rather than simulation. This
means that it would not be possible to use the strategy of generalizing the original model until an invariant is reached. since the generalized model would not be bisumulation equivalent to the original model.

### 5.4 Related research

A number of methods have been proposed in the past for extending automatic verification to parameterized designs that have an arbitrary number of similar or identical processes.

The first to approach this question were Browne. Clarke and Grumberg [BCG86], who extended the logic ('TL to a logic called inde.xed CTL. This logic allows the restricted use of process quantifiers as in the formula $V_{i} f(i)$. which means that the formula $f$ holds for some process $i$. Restricting the use of these quantifiers and eliminating the next-time operator makes it impossible to write a formula which can distinguish the number of processes in a system. By establishing an appropriate equivalence between a system with $n$ processes and a system with $n+1$ processes. one can guarantee that all systems satisfy the same set of formulas in the indexed logic. This method was used to establish the correctness of a mutual exclusion algorithm by exhibiting a bisimulation relation between all $n$-process system and a 2 -processes system. and applying model checking to the $\because$-process system.

A disadvantage of the indexed (TL method is that the bisimulation relation must be proved in an ad hoc manner. Finite state methorls cannot be used to check it because it is a relation between a finitestate process and a process with an arbitrary number of states. (larke and Grumberg dealt with the problem of establishing a bisimulation by introducing the notion of a proctse closure $l^{\prime \prime}$. This process must be derived by hand. and have the property that $. V_{r} \| P^{*}$ is equivalent io $M_{r+1} \| P^{*}$ for some small $r$. This can be verified mechanically. Shtarller and Grumberg took this notion a step further by introducing network grammars to describe classes of finite state systems. This technique used an indexed form of linear temporal logic. and required that the processes on the left and right hand sides of each grammar rule be equivalent in an appropriate sense.

The requirement that all systems generated by the grammar be.
equivalent seems to be a rather strict limitation, however. The method of this chapter, which uses a partial order rather than an equivalence, was first proposed by Kurshan and McMillan [KM89], and simultaneously by Wolper and Lovinfosse [WL89]. Around the same time, Burch was also applying a șimilar idea to Dill's trace theory for speedindependent circuits. ${ }^{1}$

Another method for proving properties of systems of identical processes is due to German and Sistla [GS]. It uses a linear-time temporal logic for specifications (again. the next-time operator is not allowed) and is fully automatic. By means of a distinguished "control" process. it is possible to check some global properties (although process quantifiers are not present in the logic). ( $n$ fortunately, because the decision algorithm is doubly exponential in the process size, this method has not been applied in practice.

A system called GORMEL has been created by Marelly and Grumberg, implementing the techniques of [SG89]. GORMEL uses context free grammars to describe systems of processes. This is fairly similar to the use of module substutution rules in S.MV. There are a number of differences between the systems. however. GORMEL is oriented towards verification of distributed algorithms. It uses a model of transition systems with pairwise synchronized actions, as in CCS. This model is not well suited for clescribing digital systems - first because most signals in hardware are broadcast to more than one location. and second because many signals are exchanged back and forth between components of a system in a single clock cyrle. The difficulty of reducing this two way exchange of many signals to a single atomic action would make it extremely cumbersome to create a ('S-like model for a system like the (iigamax.

Another difference is in the logic - CORMEL uses an indexed version of LTL without next-time called LTL². Is in indexed CTL. it is not possible to nest process quantifiers. Because of the ability to use process quantifiers. it is possible to express some properties which are not expressible in CTL. for example that if a proposition $p$ is true in some process. then it is eventually true in all processes.

For the process relation. (;)RMEL uses a form of stuttering equiv-
alence rather than simulation. This places fairly strong requirements on the allowable grammar rules. In particular, it is not possible in such a system to take the approach taken here of successively generalizing a component process in order to obtain an invariant, since the required relation between the left and right hand sides of the grammar rule is a symmetric one. The GORMEL approach will work if the various systems generated by the grammar can be distinguished only by stuttering (arbitrary repetition of the same state labeling).

A final difference between the systems is, of course, that SMV is based on symbolic model checking methods. This is not clearly an advantage, however, since the state explosion problem may not be very severe for the small number of processes that tend to be involved in induction rules.

## Chapter 6

## A partial order approach

In this chapter, we consider an alternative to the symbolic model checkang method which is also aimed at avoiding the state explosion problem. A number of researchers have observed that the arbitrary interleaving of concurrent actions is a major contributor to the state explosion problem. and that substantial efficiencies could be obtained if the enumeration of all possible interleavings could be avoided. As a result. several have proposed verification algorithms based on partial orders [Val89. Val90. God90. CiW91. PLS9. PL90. YTK91]. The method presented here is based on unfolding a Petri net into an acyclic structure called an occurrence net. The notion of unfolding was introduced by .Vielsen. Plotkin and Winskel as a means for giving a concurrent semantics to nets, but in this case the goal is to avoid the state explosion problem. In algorithm is introduced for constructing the unfolding of a net. which terminates when the unfolded net represents all of the reachable states of the original net. The unfolding is therefore adequate for testing reachability (to be more precise, coverability) and deadlock properties. It is shown using an asynchronons circuit example that the unfolding can be polynomial in the circuit size while the state space is exponential. In contrast. the stubborn sets method of Valmari (Val89. Val90] and trace automaton method of Godefroid [God90. (iW9l! are ineffective in reducing the state explosion problem for asynchronous circuit models. because of the ubiquity of confusion in such models. In addition, because the unfolding method is fully automatic. it has a certain advantage over behavior machines
method of Probst [PL89. PL90]. which requires a pomset grammar describing the circuit's behavior to be constructed by hand.

### 6.1 The unfolding operation

Briefly. an occurrence net is a Petri net without backward conflict (two transitions outputting to the same place). and without cycles. Such a net can be obtained from an ordinary place/transition net by an unfolding process. Figure 6.1 shows an example of a net and part of its unfolding. Since the occurrence net it is acyclic and rooted. there is a natural well founded (partial) order on the transitions and places of the net. This order is called the dependency order. It is impossible for a transition of the occurrence net to fire miless all of its predecessors in the dependency order have fired.

The most important theoretical notion regarding occurrence nets is that of a configuration. A configuration represents a possible partial run of the net - it is any set of transitions that satisfies the following conditions:

1. If any transition is in the configuration, then so are all of its predecessors in the dependency order (a configuration is downward closed).
2. A configuration cannot contain two transitions in conffict. meaning that both input from the same place.

An example of a configuration is shown in tigure 6.2. with element: of the configuration filled in black. Two transitions in the figure are hatched in. Either of these transitions can be added to the biack set to form a new configuration. Adding any other transition would be illegal. however, since it would either violate downward closure or conflictfreeness.

In an unfolding, each transition corresponds to a transition of the original net. and each place corresponds to a place of the original net. We can associate each configuration of the mufolding with a statet marking) of the original net by simply identifying those places whose token. are produced but not consumed by the transitions in the contiguration.


Figure 6.1: 「nfolding example.


Figure 6.2: (onfiguration

This set is marked with black dots in figure 6.2. Mapping this set back onto the original net, we obtain the final state of the configuration.

The final theoretical notion we need regarding unfoldings is that of a local configuration. The local configuration associated with any transition consists of that transition and all of its predecessors in the dependency order (that is. the downward closure of the transition as a singleton). This is the set of transitions which necessarily are contained in any configuration containing the given transition. Note that a local configuration may not exist if this set contains two transitions in conflict.

We are now ready to consider the problem of building a fragment of the unfolding which is large enough to represent all of the reachable markings of the original net. Building the unfolding itself is straightforward. The process starts with a set of places corresponding to the initial marking of the original net. The unfolding is grown by finding a set of places in the unfolding which correspond to the inputs (preset) of a transition in the original net, then adding a new instance of that transition to the unfolding, as well as a new set of places corresponding to its outputs (postset). If the new transition has no conflicts in its local configuration (more precisely, if it has a local configuration) it is kept, otherwise it is discarded. This is because the existence of a conflict means that the new transition can occur in no configurations of the unfolding.

The key to termination of the unfolding is to identify a set of transitions of the unfolding to act as cutoff points. This set must have the following property: any configuration containing a cutoff point must be equivalent (in terms of final state) to some configuration containing no cutoff points. From this definition, it follows that any successor of a cutoff point can be safely omitted from the unfolding, without sacrificing any reachable markings of the original net. To see this. suppose we have built the unfolding only up to the cutoff points. in the sense that any new transition we can add must have a cutoff point as a predecessor. From this point on, any transition we add must be descended from some cutoff point. Thus. any configuration we might add to the unfolding must have the same tinal state as some configuration already present.

A sufficient condition for a transition to be a cutoff point is the
following: the final state of its local configuration is the same as that of some other transition whose local configuration is smaller. The proof of this statement is as follows: suppose there are two transitions $t_{1}$ and $t_{2}$, whose local configurations have the same final state. with that of $t_{2}$ being smaller. Now imagine a configuration $C_{1}$ (local or otherwise) containing $t_{1}$. We can obtain $C_{1}$ from the local configuration of $t_{1}$ by adding the transitions in the difference one at a time. in an order consistent with the dependency relation. According to our construction. at each step of this process, there is a corresponding transition we can add to the local configuration of $t_{2}$ leading to the same final state. Hence. we can build a configuration $C_{2}$ containing $t_{2}$ which has the same final state. but is at least one transition smaller then $C_{1}$, since we started from a smaller set. Thus if any configuration contains a cutoff point. it is equivalent to a smaller configuration. Configurations cannot be made arbitrarily small, however. so any configuration containing a cutoff point must be equivalent to a configuration not containing a cutoff point. Since all the reachable states are represented by configurations containing no cutoff points, it is unnecessary to build the unfolding beyond any cutoff point.

We can find the cutoff points by simply keeping a hash table of all transitions, indexed by the final state of the local configuration. If when generating a transition. we find in the table a transition with equivalent but smaller local configuration. we discard the new transition. We can show. as follows, that this process is guaranteed to terminate if the original net is bounded and finite. First. the depth of the unfolding must be bounded by the number of number of reachable markings. The depth of a given transition in the unfolding is the longest chain of predecessors of that transition. Each transition in this chain has a local configuration. and these local configurations form a chain of increasing size. If the depth of the given transition is greater than the number of reachable markings of the original net. then by the pidgeon-hole principle. two of these local configurations must have the same tinal state. This cannot be, however. since in this case one of the transitions in the chain would have been determined to be a cutoff point. If the original net is bounded. it has a finite number of reachable markings. hence the depth of the unfolding is bounderl. If the original net is timite. we can show by induction that the number of transitions at any given


Figure 6.3: Dining philosophers net.
repth in the unfolding is finite. Hence the total number of transitions generated by the unfolding process is finite.

As an example of termination. consider the net of figure 6.3. which represents the dining philosophers paradigm. In this scenario. there are $n$ concurrent processes (philonophers). each of which must acquire the use of two shared resources ( forks) in order to execute its critical section (eating spaghetti). The processes are organized in a ring. with each neighboring pair sharing one resource. Figure 6.1 shows the completed unfolding for the case of three philosophers ( $n=3$ ). The cutoff points are marked with an $\mathcal{X}$. The local contiguration of each of these transitions is equivalent to the empty configuration. We observe that the size of the unfolding is not only bounded. but is linear in the number of philosophers. while the number of states is exponential as shown in table 6.1.

Recall that in growing the untolding, it is necessary to enumerate all of the subsets of places which correspond to the inputs of transitions. The complexity of this is $O\binom{\prime \prime}{1}$. where $n$ is the size of the unfolding. and , is the largest number of inputs of any transition. This is.of course. bounded by $n^{\prime}$. which is polyomial given a fixed value of $i$. In practice.


Figure 6.4: Unfolding of the dining philosophers net.

| number <br> of philosophers | size of unfolding <br> (transitions) | number of <br> reachable states |
| :---: | :---: | :---: |
| 2 | 9 | 22 |
| 3 | 13 | 100 |
| 4 | 17 | 466 |
| .5 | 21 | 2164 |

Table 6.1: Vnfolding size and number of states for Dining Philosophers
however, the number of subsets which are considered can be reduced quite effectively, using the following two techniques. First, suppose we are enumerating the subsets: we need not add any place to the set if the result would not be contained in the set of inputs of any transition. Second whenever a place is added to the set. we can immediately eliminate from consideration all of the places which have a predecessor in conflict with a predecessor of the new element. since any transition with both places as inputs would be discarded. We add transitions to the net in order increasing size of the local configuration. so that we can use a hash table to determine whether or not each transition is a cutoff point. Thus. whenever a candidate for a transition in the untolding is generated. it is placed in a queue ordered by increasing local configuration size. The places of the net are enumerated by pulling the first element $t^{\prime}$ from this queue. twsting whether it is a cutoff point. and if not. generating places for its outputs. The procedure terminates when the queue of candidate transitions becomes empty. Figures 6.5 and 6.6 show a psetdo code implementation of this procedure. The pseudo-code is written somewhat inefficiently in places for simplicity.

In function Infold. the arguments $P . T$ and.$U_{1}$ are the places. transitions and initial marking of the original net. Each place in the unfolding is represented by a pair iplace. preds), where place is the corresponding place in the original net, and preds is the set of immerliate predecessor transitions in the miolding tnote that since there is no backward conflict the size of this set is at most one,. Each transition in the unfoldine is represented by a pair (trmes. preds), where trans: in the corresponding transition in the original net. and preds is the set of immediate predecessor places in the unfolding. The finction returns $P^{\prime}$ and $T^{\prime}$. the set of places and tansitions. respectively, of the untolding. There is also a quene ()' ot transitions to be expanded. and a hash table(Hash Table) used for wemtifving cutoff points.
('overability problems ran be solved using the motolding in the following way. Imagine wo have a wet of places in the original net. and we wish to determined whether this set can every be simultaneonstr marked. We simply ard a new transition to the net. whose inputs are the given set, and then construct the untolding. If the mbidine contains any instance of his new fransition the sen is comerable, wed ot herwise nut.

```
global \(P^{\prime}, T^{\prime}, Q^{\prime}\). Hash Table[]
function \(\operatorname{Unfold}\left(P, T, M_{0}\right)\)
    \(P^{\prime}=T^{\prime}=Q^{\prime}=\emptyset\) : clear Hash Table
    for each \(p \in M_{0}\) do
        add \(p^{\prime}=(p, \emptyset)\) to \(P^{\prime}\)
        GenTrans( \(\left.\left\{p^{\prime}\right\} . T\right)\)
    end for
    while the queue \(Q^{\prime}\) is not empty do
        pull the first \(t^{\prime}\) off of \(Q^{\prime}\)
        if not IsCutoffPoint? \(\left(t^{\prime}\right)\) do
            for each \(p\) in outputs of \(\operatorname{trans}\left(t^{\prime}\right)\) do
            add \(p^{\prime}=\left(p,\left\{t^{\prime}\right\}\right)\) to \(P^{\prime}\)
                GenTrans( \(\left\{p^{\prime}\right\}\). \(T^{\prime}\) )
            end for
        end if
    end while
    return \(\left(P^{\prime} . T^{\prime}\right)\)
end function
procedure GenTrans( \(S^{\prime} . T\) )
    if not exists \(t \in T\) such that place \(\left(S^{\prime}\right) \subseteq\) inputs of \(t\) then returu
    if Predecessors( \(\left(S^{\prime}\right)\) has forward conflict then return
    forall \(t \in T\) do if place( \(\left.S^{\prime}\right)=\) inputs of \(t\) then
        add \(t^{\prime}=\left(t . S^{\prime}\right)\) to set \(T^{\prime}\)
        insert \(t^{\prime}\) in \(Q^{\prime}\) in order of 'Loral( onfig(t \(\left.t^{\prime}\right)\)
    end for
    for all \(p^{\prime} \in P\) where \(p^{\prime}\) older than any member of \(S^{\prime}\) do
        GenTrans \(\left(S^{\prime} \cup p^{\prime}, T\right)\)
end procedure
```

Figure 6.): Pseudo-corle implementation of mifolding procedure

```
function IsCutoffPoint?(t')
    C
    S
    L'= HashTable[HashFun(S1)]
    forall tt2 in L' do
        C
        if S1
    end for
    add tt1 to HashTable[HashFun(S'S
    return(0)
end function
function LocalConfig(t')
    return(Predecessors({t'})\capT')
end
function Predecessors(S')
    do
        S'=}\mp@subsup{S}{}{\prime}\cup\mathrm{ preds( }\mp@subsup{S}{}{\prime}
    until }\mp@subsup{S}{}{\prime}\mathrm{ unchanged
end function
function FinalState( (')
    let }\mp@subsup{S}{}{\prime}\mathrm{ be the set of all }\mp@subsup{p}{}{\prime}\in\mp@subsup{P}{}{\prime}\mathrm{ such that preds ( p
    return(place( }\mp@subsup{S}{}{\prime\prime}-\operatorname{preds(\mp@subsup{C}{=}{\prime})))
end function
```

Figure 6.6: Pseudo-code. continued.


Figure 6.7: Translation from circuit to net

### 6.2 Application example

We now consider a more realistic example than the dining philosophers - a speed-independent [Sei80b] circuit designed to implement a distributed mutual exclusion (DME) protocol. The circuit was designed by Alain Martin [Mar85] and has been analyzed using an abstracted trace theoretic model by Dill [Dil88]. It was also used as an example for the symbolic model checking merthod in section 2.4.2.

Networks of logic gates in sperd-independent circuits are readily. modeled by Petri nets. A network of 11 gates can be modeled by a Petri net of $O(n)$ places. When we model a network of gates as a Petri net. we introduce two places for each input of each gate. One represents the the input in a logic low state. while the other represents the input in a logic high state. Transitions in the Petri net correspond to rising or falling transitions of gate outputs. I rising transition of a gate output removes all the logic low tokens from the inputs to which it is connected. and places tokens on the correspondine logic high places.

As an example. figure 6.7 shows the net fragment representing an AND gate. When both inputs of the gate are at the logic high state. we can move a token from the place representing logic low at the output to the place representing logic high. Similarly: if either input is at the logic low state, we can move a token from the place representing logir high at the output to the place representing logic low.

A dynamic hazard ocrurs. for example, if the $A . N D$ gates output is enabled to rise while one of the inputs is emabled to fall. The problem of whether or not a dynamic hazard can ocour can thus be posed as a
coverability problem. Alternatively, since dynamic hazards correspond to dynamic conflicts in the unfolding, the problem can be solved by constructing the unfolding and examining it for dynamic conflicts. ie.. two transitions which are in conflict, and which may be simultaneously enabled. The DME circuit also uses special two-way mutual exclusion elements as components, which are immune to certain hazards. In checking the DME ring for hazards, we ignore conflicts between rising transitions of a mutual exclusion element's acknowledge outputs.

Table 6.2 shows the results of the occurrence net unfolding procedure (ONU), and a depth-first traversal of the state space (DFT), for the Petri net model of the circuit. for rings with one to five cells. The depth of the occurrence net unfolding for the case of 5 cells was 141 transitions. The number of transitions in the ONU increases quadratically in the number of cells. This is because as the number of cells in the ring increases. a request must be relayed through a greater number of stages in order to obtain the token. in the worst case. At the same time, the number of cells which are requesting also increases. The occurrence net therefore grows in both width and depth in proportion to the number of cells. As we increase the number of cells in the ring. the number of reachable global markings increases exponentially. For this reason. it was only possible to apply DFT to a system of five cells. before the available memory resources were exhausted. It is known. however. from using OBDD based methods, that the number of states increases asymptotically by slightly less than a factor ten for each added cell.

How do these results compare other methods for avoiding the state explosion problem? The trace theory approach of Dill [DilS8] required an abstract model of the arbiter cell to be created by hand. This reduces the state explosion problem. but does not entirely solve it. since even with the reduced model. the number of states still increases exponentially with the number of components. Probst [PL90] reports a method which requires quadratic space and time in the number of cells. but also is not fully automatic. The methorls of Valmari [Val89, Val90] and Godefroid [God!90. CW91] and Yonerla [YTK91] cannot be effectively applied to this example or to other speed independent circuits. because in all states. all enabled transitions are in contlict with some disabled transition. Thus no transition ran be statically guaranteed to be per-


Figure 6.3: Distributed mutual exclusion circuit
sistent. Experiments by Holger Schlingloff ${ }^{1}$ have confirmed this to be the case. It is possible. perhaps. that some more clever static analysis technique could be used to show that some transitions are persistent. in which case these methods could be applied to some effert.

Finally, in chapter 2. we saw that the symbolic model checkine method had cubic time complexity and lincar space complexity. using a simultaneous model. Burch and Long- have obtained $O\left(n^{-i}\right)$ time complexity for this circuit using symbolic model checking with a morlified search order (cf. section 2.3). This methori requires some hand optimization. however. In any event. it appears that the symbolic model checking method yields somewhat better asymptotic performance for the DME circuit. though both methorls effectively solve the state ex. plosion problem.

[^20]| number <br> of cells | size of unfolding <br> (transitions) | number of <br> reachable states |
| :---: | :---: | :---: |
| 1 | 23 | 22 |
| 2 | 125 | 502 |
| 3 | 313 | 6579 |
| 4 | 604 | 75172 |
| .$j$ | 1013 | 802425 |

Table 6.2: Performance of ONU and DFT on hazard-detection problem for the distributed mutual exclusion circuit

### 6.3 Deadlock and occurrence nets

Besides coverability, another interesting problem for Petri nets is the question of deadlock. A terminal marking of a Petri net is one in which no transitions are enabled. Reachability of a terminal (or deadlocked) state cannot be framed in terms of the coverability problem. However, since the unfolding represents all reachable markings, a net has a reachable terminal marking if and only if its unfolding has a reachable terminal marking. The problem of existence of a reachable terminal marking of an occurrence net is NP-complete. This is easily shown by reduction from 3-SAT. ${ }^{3}$ To see this consider the formula $\left(x_{1}+y_{1}+z_{1}\right)\left(x_{2}+y_{2}+z_{2}\right) \cdots\left(x_{n}+y_{n}+z_{n}\right)$ where each $x_{i}, y_{i}$ and $z_{1}$ is a positive or negative literal. Assume the formula has $m$ variables. Let the positive literals be $l_{1} \ldots \ldots l_{m}$, and the negative literals be $\bar{l}_{1} \ldots \ldots \bar{l}_{m}$. In polynomial time, we can construct a net which has a terminal marking if and only if the formula is satisfiable. The initial marking of the net is a set of places $\left\{v_{1} \ldots \ldots v_{m}\right\}$. There is a place representing each positive literal $l_{1}, \ldots, l_{m}$ and each negative literal $\bar{l}_{1}, \ldots . \bar{l}_{m}$. For each variable $v_{i}$, there is a transition from $v_{i}$ to $l_{i}$ and from $v_{i}$ to $\bar{l}_{i}$. For each conjunct $\left(x_{i}+y_{i}+z_{i}\right)$, there is a transition $c_{i}$, whose preset is $\left\{\bar{x}_{i}, \bar{y}_{i}, \bar{z}_{i}\right\}$. In other words. the transition $c_{i}$ is enabled to fire if and only if ( $x_{i}+y_{i}+z_{i}$ ) is false. Thus. some transition $c_{i}$ is enabled to fire if and only if the whole formula is false. The postset of each transition

[^21]

Figure 6.9: Reduction from 3-S.AT problem to a terminal marking problem.
$c_{1}$ is the single place $\{q\}$. and there is a transition from $\{q\}$ to $\{q\}$. Thus. if any $c_{i}$ fires. the net may never reach a terminal marking. As a result. there is a terminal marking of the net if and only if the tormula is satisfiable. For example. figure 6.9 shows the net constructed for the formula $(a+b+\bar{c})(b+c+\bar{d})$.

The reader may easily verify that the size of the unfolding of such a net (up to the cutoff points) is linear in the size of the original net. In fact. it is essentially the same net, except the the place q occurs " times in the unfolding. Since all reachable markings of the original net occur as configurations of the unfolding. the unfolding has a terminal marking if and only if the formula is satisfiable. Hence $3-5 . A T$ is P-time reducible to reachability of a terminal marking of an unfolding. Since the configuration representing the erminal marking can be guessed in $P$-time in the size of the unfolding, and also tested in $P$-time. it follow that the problem is in $\triangle P$. and hence.$\stackrel{P}{\text { P }}$-complete.

Interestingly: however the problem is readily solved in practice even

[^22]Figure 6.10: Procedure to detect terminal marking.
for very large unfoldings, using an algorithm based on techniques of constraint satisfaction search. The key observation which leads to this algorithm is that there is no terminal marking exactly when all configurations the unfolding can reach some configuration containing a cutoff point. This is simply because if there is no terminal marking, then all configurations can reach a configuration which is arbitrarily large. $A$ configuration $C^{\prime}$ can reach a configuration containing transition $t^{\prime}$ if and only if the union of $C^{\prime}$ and the local configuration of $t^{\prime}$ is a configuration. If it is not, then no set containing $C^{\prime}$ and $t^{\prime}$ is a configuration. If the union is not a configuration, we will say that $C^{\prime \prime}$ and $t^{\prime}$ are in conflict. Hence, there is a terminal marking if and only if there is a configuration which is in conflict with every cutoff point. The search for such a configuration can be carried out using branch and bound techniques. For example. if a configuration $C^{\prime}$ is in conflict with a cutoff point $t^{\prime}$, there must be a transition $t_{1}^{\prime} \in C^{\prime}$ which is in conflict with some transition in the local configuration of $t^{\prime}$. Such a transition $t_{1}^{\prime}$ will be called a spoiler of $t^{\prime}$.

There exists a configuration in conflict with all of the all of the cutoff points (equivalently, there exists a terminal marking) if and only if there exists a configuration containing a spoiler for every cutoff point. The set of spoilers contained in this contiguration will be called $I_{s}$. The algorithm of figure 6.10 uses branch and bound techniques to find such a set $T$, if one exists.

Note that in line 3 of the procedure, the cutoff point with the smallest number of spoilers is chosen so that the number of choices in line 5 is minimized. Whenever a spoiler for a given cutoff point is chosen to belong to $T_{s}$ in line 5 , everything in conflict with $T_{s}$ is eliminated from future consideration in line 7 . Note that the cutoff points in conflict with $T_{s}$ are also eliminated, which cuts down on the amount of future branching. Whenever there is a cutoff point with no remaining spoilers. the procedure backtracks. from line 4 to the most recent occurrence of line 5 where there are remaining choices. If there are no remaining choices. the procedure fails. Of course. when backtracking occurs. the the net is also returned to the state it was in at the point where execution is being resumed. This backtracking is easily implemented by keeping a stack of the remaining choices for $t^{\prime}$ in each iteration of the loop. and marking each transition in the net with the level of the stack at the time it was "removed". Interestingly, if the procedure terminates successfully, the remaining net has the property that every path leads to a terminal marking of the original net $\lambda$. This makes it straightforward to extract a path leading to a terminal marking.

Obviously. because of the backtracking. this procedure is exponential (as it must be. if $\mathcal{P} \neq \mathcal{V} \mathcal{P}$ ). However. this is only the worst case. The dining philosophers serve as an example of a case in which the exponential complexity is avoided. In fact. the procedure finds the terminal marking in time which is linear in the number of philosophers. This is easily seen by examining the unfolding of the Dining Philosophers net in figure 6.t. There is one cutoff point in this net for each process. Initially, each of these transitions has two spoilers, which correspond to the two resources required to enter the critical region being granted to the two neighboring processes. Regardless of which cutoff point is used first. the symmetry is then broken as the part of the net in conflict with one of the two spoilers is removed. This removes. in particular. the transition which granted one of the resources to the first philosopher. hence one of its neighbors now has only one spoiler. so there is only one choice available the next time line j is reached. After this spoiler is added to $T_{s}$. the remaining neighbor of the second philosopher now has only one spoiler. This process continues withon backtracking until it has come full circle and the terminal marking is found. Note that if the cutoff point with the fewest spoilers were not
chosen in line 3 , the procedure might have examined an exponential number of candidates for $T$, before a valid one was found.

In fact, using nets representing communication protocols as examples, this procedure has been successfully been applied to unfoldings with more than 3000 transitions and 1000 cutoff points, where some cutoff points had as many as 50 spoilers. It is clear that the branch and bound technique quickly narrows down the number of choices in these examples.

### 6.4 Relation to AI techniques

The occurrence net unfolding method, as applied to the coverability problem, was inspired by so-called "least commitment" strategies for AI planning problems, especially nonlinear planning techniques. Like these strategies, the method falls into the category of searching in the space of solutions (ie.. covering sequences for a given set), rather than the space of the problem (ie., the reachable markings). Each partially constructed unfolding represents some set of possible partial solutions which may be extended to a complete solution. As in other constraint satisfaction search methods. the method tries to eliminate as early as possible those partial solutions which cannot be extended to complete solutions. This is done in the unfolding procedure by the elimination of candidate transitions which have no local configuration, and also by the cutoff points. which effectively discard those partial solutions which cannot be extended to a lowest cost (ie., fewest transition) solution. This is done without unnecessarily committing to the order of independent transitions. This makes the unfolding method somewhat similar to constraint posting methods used in non-linear planning. Both methods construct fairly similar structures, although non-linear planners. such as NOAH [Sacī] only represent one partial solution, while the occurrence net represents all partial solutions. Non-linear planners attempt to detect conflicts and eliminate them by posting additional constraints on the solution or by modifying elements of the solution. Non-linear planners also use heuristics to guide them towards a solution. and hence sometimes overconstrain the solution space and require backtracking. For this reason, they are heuristically efficient. but would
not be suitable for exhausting the solution space, as is required in automatic verification methods. In fact, NOAH is not even guaranteed to find a solution where one exists. A later system called TWEAK [Cha87] does guarantee a solution where one exists, but fails to terminate if no solution exists. Since TWEAK does overconstrain the solution space. exhaustive search could only be achieved by backtracking, even if a suitable termination condition were found. The unfolding procedure never overconstrains the solution space, however. as conflicting transitions may coexist in the unfolding. Hence. it cloes not backtrack.

### 6.5 Evaluation

When is unfolding a suitable strategy for problems in automatic verification? The most promising application is hazard checking for asynchronous control circuits. In these circuits, the state explosion seems to derive almost entirely from arbitrary interleavings of concurrent transitions. In such cases. the unfolding method can have a considerable advantage over methods that search the entire state space. . .ote. however, that other methods based on partial orders are not necessarily effective in reducing the state explosion for these circuits. because of the aforementioned problem of determining when transitions of the net are persistent.

In general. any problem which can be posed in terns of roverability or cleadlock in a Petri net model is a possible application of the unfolding method. In addition. it is possible that heuristically efficient procedures can be found for deciding the existence of an infinite firing path in some $\boldsymbol{\omega}$-regular set, given an unfolding. In this case. specifications framed as linear time temporal logic formulas. or $\sim$-automata could be evaluated.

## Chapter 7

## Conclusion

What we have seen in the preceding chapters is that Ordered Binary Decision Diagrams can be used as a representation in a wide variety of automatic verification algorithms, in order to cope with the state explosion problem. This can be done in a unified way by representing the algorithms in the Mu-Calculus fixed point notation. For fairly diverse families of regularly structured systems, the C'TL model checking algorithm was observed to run in time and space which increased polynomially in the size of the system. while the number of reachable states increased exponentially. These results bear out a theoretical result bounding the OBDD representation of the transition relation for such systems. Standard automatic verification algorithms would be unsuitable for these examples because their complexity is proportional to the number of reachable states.

Using OBDD based techniques. and a language suitable for the abstract modeling of digital systems, it was possible to verify a fairly complex industrial design for a cache consistency protocol. finding a number of subtle errors in the process. The verification process is valuable not only because of the advantages of formalization and exhaustive checking, but because it can find protocol errors more quickly than simulation. despite the exponential growth in states as the model increases in size. The ability to isolate high level errors quickly shortens the loop between design and verification. making it possible to experiment more freely with alternative designs. and shortening the "critical path" from conceptualization to implementation.

By a technique of induction over processes. it is possible to prove properties of a protocol which are independent of the number of processes participating in the protocol. This type of proof requires a sufficient understanding of the protocol on the part of the designer to construct a process invariant. Invariants are difficult to find. but the symbolic model checker provides an aid in this process by producing counterexamples for unsound invariants. In the author's opinion. finding a process invariant for a protocol is not only of value as a proof technique .- the understanding of the protocol required to formulate the invariant can lead to simpler and more elegant protocols. This is another reason for formalizing and verifying a protocol before attempting to implement it.

The verification technique based on occurrence nets shows that OBDDs are not the only representation that can be used to avoid the state explosion problem. There are. in fact. certain advantages to the occurrence net based method for the example presented. since the memory usage is small. and no hev ristic technique is required to produce a variable ordering. Still. at this stage, the occurrence net methoc! is certainly not as well advanced as the symbolic model checking method.

There are several areas where the current work falls short of the goal of complete automatic verification of digital systems. In the case of the Gigamax protocol. an abstract model of the protocol was verified and not the actual implementation. Verification of the implementation would have been impossible due to a lack of formal models of the components of the system (ie.. standard devices, such as memories. registers. programmable logic. central processing units. ftc.). If such models were available from the manulacturers. in principle the method. described in chapter 5 could be used to show that the implementationi is simulated by the abstract model. Hierarchical reasoning of this kind has been extensively studied by Kurshan [Kiursi]. Vnfortunately. simulation does not preserve existential C'TL properties such as absence of deadlock. As mentioned previously. bisimulation equivalence. which preserves all C'TL properties. is too strong for this purpose. since the abstract models are necessarity non-deterministic, and the actmal implementation cannot 1 and shonld not exhibit this non-determinism. A practical technique of abstraction which preserves existential 'Tl. properties is needed if ex.stential properties are to be proved usime

## hierarchical reasoning.

There is also a need for heuristic strategies for generating process invariants in inductive proofs. Marelly and Grumberg view the design of the invariant as part of the design of the protocol. This is a useful point of view, but some automated help beyond the generation of counterexamples would be useful for this purpose.

Finally, this work concentrates on how to solve the verification problem, once it has been formalized as the satisfaction of a temporal logic formula by a finite model, or as an appropriate relation between finite automata. There is, of course, a wide range of issues involved in formalizing the problem in the first place. For example, there is the ever present danger that the specification itself is incorrect. In the case of the very simple CTL formulas used to specify the Gigamax protocol, this is perhaps not a severe problem. The abstraction that was used to create a model for checking the sequential consistency property was. however, not obviously correct.

In general, there is a clear need for complete mechanical checking that the implementation of a processor or protocol matches the intended architecture (user model). This requires first of all a definitive model of the architecture - something that is currently lacking even for standardized architectures in the public domain. Second there must be a well defined criterion for determining what is a valid implementation of the architecture. Loosely, an implementation of a processor is equivalent to an architecture model if for all programs. the two machines produce the same "answer". However, for many reasons, this equivalence cannot be directly stated in terms of equivalence of finite state machines. For one, most modern CPU architectures have no explicitly defined notion of input and output. It is not adequate to view input and output as the sequence of loads or stores observed at the memory interface, since this sequence will differ among implementations (especially if the implementations contain cache memories, which is often the case). Solutions to the formalization problem are needed. but cannot be obtained by studying theoretical models alone. It is necessary to carefully consider what verification means in an engineering sense. as well as a mathematical sense.

Despite the shortcomings of current verification technology, it is clear that there are at least small areas of the problem space for which
reasonable solutions exist, and these solutions can be put into practice to positive effect in an industrial setting. Those involved in verification research should perhaps take a closer look at engineering practice to determine how well the verification solutions match up with real engineering problems. This effort may lead not only to a more practical theory of formal verification, but also to a rich source of theoretical problems.

## Appendix A

## Semantics of SMV. 1

This appendix defines the semantics of programs in the language SMV.1. which includes the subset SMIV. 0 plus the process keyword. In SMV.1. we need to account for both the arbitrary interleaving of processes, and the rules regarding when a variable may change value as a result of executing a given process.

## A. 1 The model

The set $N$ of names, is the set of all character strings made up of the letters, the digits, the underscore and the minus sign characters. beginning with a letter. The store $L=L_{V} \cup L_{H}$ is made up of two disjoint. countably infinite sets of locations $L_{V}$ and $L_{H}$. We will call the former the visible locations, and the latter the hidden locations. The set of locations $L$ is defined recursively. It is the least set such that

1. if $n \in N$, then $n \in L_{V}$, and
2. if $l \in L_{V}$ and $n \in N$, then $l . n \in L_{V}$, and
3. if $l \in L_{V}$, then $l \in L_{H}$.

The set of values $V$ is the union of the integers in the range $\left[-2^{31}, 2^{31}-1\right]$ and $N$, the set of names. A state $x: L \rightarrow V$ is a function from locations to values. Let $S=L \rightarrow V$ be the set of all possible states.

If $p$ is a declaration, then its denotation $\lfloor p \rrbracket$ is a quadruple $(T, I, R, C)$. The $T$ component is a partial function from $L$ to the finite subsets of $V$. If $l$ is a location, then $T(l)$, when defined, is the type of $l$ - the set of values that can be assigned to location $l$. The component $I \subseteq S$ is the set of initial states. The component $R \subseteq S \times S$ is the transition relation. An asynchronous process is identified with a location $r$, which has the value 1 in a given state exactly when the process is executing in that state. The component $C \subseteq L \times L$ is the set of pairs ( $r . l$ ) such that process $r$ assigns the next value of location $l$.
A. 2 Expressions

An expression denotes a function from states to finite subsets of 1 . according to the following rules:

1. If $v$ is a value. then $\{v \rrbracket(x)=\{\cdot\}$.
2. If $l$ is a location, then $\llbracket l \rrbracket(x)=\{x(l)\}$.
3. If $\epsilon_{1}, e_{2}$ are expressions, and $\omega$ is one of

$$
+.-. *, / . \bmod .\rangle,\rangle=.\langle.\langle=.=. \& . \mid,-\rangle .\langle-\rangle
$$

then

$$
\llbracket e_{1} o e_{2} \rrbracket(x)=\left\{\llbracket o \rrbracket\left(v_{1} \cdot c_{2}\right) \mid r_{1} \in \llbracket \epsilon_{1} \rrbracket(x), v_{2} \in \llbracket \epsilon_{2} \rrbracket(x)\right\}
$$

4. If $\varepsilon$ is an expression, then

$$
\llbracket!e](x)=\{\mathbb{I}!(x) \mid c \equiv \mathbb{I}!\mathbb{d}(x)\}
$$

3. If $\epsilon_{1}, \epsilon_{2}$ are expressions.

$$
\llbracket e_{1} \text { union } \varepsilon_{2} \rrbracket(x)=\left[r_{1} \rrbracket \cup \llbracket t_{2} \rrbracket\right.
$$

6. If $\epsilon_{1}, \epsilon_{2}$ are expressions.

$$
\left.\llbracket t_{1} \text { in } \epsilon_{2} \rrbracket(x)=\llbracket t_{1}\right]=\left[t_{2} \rrbracket\right.
$$

The functions denoted by $+,-, *, /$ are the usual functions of arithmetic modulo $2^{32}$. The function denoted by mod is the positive remainder of division $\bmod 2^{32}$. The function denoted by the relational operators $>,>=,<$ and $<=$ return 0 when the relation is false and 1 when the relation is true, and are defined for numeric values only. For nonnumeric values, they return $\perp$. The equality operator $=$ is defined for all values. and returns 0 when they are unequal, and 1 when they are equal. The functions denoted by the Boolean operators are \& (for and), | (for or), ! (for not), $\rightarrow$ (for implies) and $\langle->$ (for logical equivalence) are defined only for the values 0 and 1 , and return $\perp$ otherwise.

## A. 3 Assignments and definitions

There is no semantic difference between assignments and definitions. If $l$ is a location. and $e$ is an expression, then the assignment $l:=e$ : denotes a quadruple ( $T, I, R . C^{\prime}$ ), where

1. $T=\emptyset$
2. $I=S$
3. $R=\left\{(x, y) \in S^{2} \mid l(x) \in \epsilon(x)\right\}$
4. $C=\emptyset$

The assignment next $(l):=\epsilon$ : denotes a triple ( $T, I, R$ ) where

1. $T=\emptyset$
2. $I=S$
3. $R=\left\{(x, y) \in s^{2} \mid x(\right.$ running $\left.) \Rightarrow l(y) \in e(x)\right\}$
4. $C^{\prime}=\{($ running, $l)\}$

The assignment init $(l):=\rho$ : denotes a triple ( $T, I, R$ ) where

1. $T=\emptyset$
2. $I=\{x \in S \mid l(x) \in+(x)\}$
3. $R=S^{2}$
b. $\quad \therefore=\theta$

## A. 4 Variable declarations

If $l$ is an identifier and $v_{1}, v_{2}, \ldots, v_{n}$ are values, then

$$
\operatorname{VAR} l:\left\{v_{1}, v_{2}, \ldots, v_{n}\right\} ;
$$

denotes a quadruple ( $T, I, R, C$ ) where

1. $T=\left\{\left(l,\left\{v_{1}, v_{2}, \ldots, v_{n}\right\}\right)\right\}$
2. $\left.I=\left\{x \in S \mid x(l) \in\left\{v_{1}, v_{2}, \ldots, v_{n}\right\}\right)\right\}$
3. $\left.R=\left\{(x, y) \in S \mid x(l), y(l) \in\left\{v_{1}, v_{2} \ldots \ldots v_{n}\right\}\right)\right\}$
4. $C=0$

## A. 5 Renaming

Let $\phi: L \rightarrow L$ be a function from locations to locations. This in turn induces a map $\Phi$ on states, such that for all states $x$ and locations $l$.

$$
\Phi(x)(l)=x(\circ(l)) .
$$

If $M=\left(T, I, R, C^{\prime}\right)$, then let $\Phi(M)=\left(T^{\prime} . I^{\prime} . R^{\prime} . C^{\prime}\right)$ where

1. $T^{\prime}(\phi(l))=T(l)$.
2. $I^{\prime}=\{x \mid \Phi(x) \in I\}$ and
3. $R^{\prime}=\{(x, y) \mid(\Phi(x), \Phi(y)) \in R\}$.
4. $C^{\prime}=\left\{(\phi(r) . \phi(l)) \mid(r . l) \in C^{\prime}\right\}$.

Note that the definition of $T$ does not make sense if o maps two locations with different types onto the same location. In this case. $O(M)$ is a type error.

## A. 6 Parallel composition

Let $M_{1}=\left(T_{1}, I_{1}, R_{1}, C_{1}\right)$ and $M_{2}=\left(T_{2}, I_{2}, R_{2}, C_{2}\right)$. Let $n_{1}$ and $n_{2}$ be two distinct names. For $i \in 1,2$, let $\phi_{i}(l)=. n_{i} . l$ for all $l \in L_{H}$ and $\dot{\phi}_{i}(l)=l$ otherwise, and let $M_{i}^{\prime}=\phi\left(M_{i}\right)$. The parallel composition $M=M_{1} \| M_{2}$ is defined as follows:

1. $T=T_{1}^{\prime} \cup T_{2}^{\prime}$
2. $I=I_{1}^{\prime} \cap I_{2}^{\prime}$
3. $R=R_{1}^{\prime} \cap R_{2}^{\prime}$
4. $C^{\prime}=C_{1}^{\prime \prime} \cup C_{2}^{\prime}$

If $d_{1}, d_{2}, \ldots . d_{k}$ are declarations. then $\left.\llbracket d_{1} d_{2} \ldots d_{k}\right]$ is the parallel composition

$$
\llbracket d_{1} \rrbracket\left\|\llbracket d_{2} \rrbracket\right\| \cdots \| \llbracket d_{k} \rrbracket
$$

## A. 7 Instantiation

Suppose that module $A$ is defined as follows:

$$
\text { MODULE } A\left(n_{1}, n_{2}, \ldots, n_{k}\right) D
$$

where $n_{1}, n_{2}, \ldots, n_{k}$ are distinct names and $D$ is a sequence of declarations. We first consider the variable declaration VAR $r: l\left(l_{1}, l_{2} \ldots \ldots l_{k}\right)$; where $r . l_{1}, l_{2}, \ldots . l_{k}$ are visible locations. Let $\Phi$ be a renaming. such that, for all $l \in L_{V}$.

1. for all $l \leq i \leq h: D\left(n_{i}\right)=l_{i}$, and $\phi\left(n_{i} \cdot l\right)=l_{i} . l$.
2. $\Phi(. l)=. l$
3. for all $n \in V-\left\{\right.$ running. $\left.n_{1}, n_{2}, \ldots . n_{k}\right\} . \phi(n)=r . n$. and $o(n . l)=$ r.n.l.
4. $\phi($ running $)=$ running, $D($ running.$l)=$ running. $l$

Then [VAR $\left.r: A\left(l_{1}, l_{2}, \ldots, l_{k}\right) ;\right]=\phi(D)$.
Now we consider the declaration

$$
\text { VAR } r: \text { process } A\left(l_{1}, l_{2}, \ldots, l_{k}\right)
$$

Let $\phi$ be a renaming, such that, for all $l \in L_{V}$,

1. for all $1 \leq i \leq k: \phi\left(n_{i}\right)=l_{i}$, and $\phi\left(n_{i} \cdot l\right)=l_{i} \cdot l$,
2. for all $n \in N-\left\{n_{1}, n_{2}, \ldots, n_{k}\right\} . \phi(n)=r . n$, and $\phi(n . l)=. m_{1} . n . l$.
3. $\phi(. l)=. l$

Then $\llbracket \operatorname{VAR} r:$ process $A\left(l_{1}, l_{2}, \ldots, l_{k}\right): \rrbracket=\Phi(D)$.

## A. 8 Programs and interleaving

Suppose that module main is defined as follows:
MODULE main $D$
where $D$ is a sequence of declarations and $\llbracket D \rrbracket=(T, I, R . C)$. The number of processes executing in state $x$ is

$$
n_{e}(x)=|\{r \mid(x(r)=1) \wedge \exists l:(r . l) \in C\}|
$$

The set of legal interleaving states is

$$
S_{I}=\left\{x \in S \mid n_{e}(x)=1\right\} .
$$

The set of states in which location $l$ is constrained to remain unchanged is

$$
U(l)=\{x \in S \mid[\exists r:(r . l) \in C \cdot] \wedge \neg \exists r:[(r . l) \in C \wedge(x(r)=1)]\}
$$

The denotation of the program is a triple ( $T . I^{\prime} . R^{\prime}$ ). where

1. $I^{\prime}=I \cap S_{I}$.
2. $R^{\prime}=\left\{(x, y) \in R \mid x \in \overleftrightarrow{\exists} \wedge \forall l \in L:\left(x \in C^{\prime}(l) \Rightarrow x(l)=y(l)\right)\right\}$.

## A. 9 Specifications

Each program is associated with a Kripke structure which determines the truth value of CTL formulas in the specification. The atomic propositions in this case are all the Boolean valued expressions. The Kripke structure associated with a program whose denotation is the quadruple $(T, I, R)$ is a Kripke model $K=\left(S, R . L^{\prime}\right)$ where

1. $S$ is the set of states defined above.
2. $R$ is the transition relation. and
3. if $e$ is an expression. then

$$
L^{\prime}(\epsilon)=\{x \in S \mid \mathbb{I} \rrbracket(x)=\{L\}\}
$$

The specification is a formula $f$ in (TL with fairness constraints. It is satisfied exactly when $K$. $s_{0} \vDash f$ for all $s_{0} \in I$.

## Bibliography

[AB86] J. Archibald and J. L. Baer. C'ache coherence protocols: Evaluation using a multiprocessor simulation model. ACM Transactions on Computer Systems, 4(4):273-298, 1986.
[Ake78] S. B. Akers. Binary decision diagrams. IEEE Trans. Computers, C-27(6):509-516. August 1978.
[BAMP81] M. Ben-Ari. Z. Manna, and A. Pnueli. The temporal logic of branching time. In ACM Symp. Principles of Programming Languages, pages 164-176, 1981.
$\left[\mathrm{BBB}^{+} 87\right]$ R. E. Bryant. D. Beatty, K. Brace, K. Cho, and T. Sheffler. COSMOS: $\lambda$ compiled simulator for MOS circuits. In ${ }^{3} 4$ th Design Automation Conference, 1987.
[BBS90] Derek L. Beatty. Randal E. Bryant, and Carl-Johan H. Seger. Synchronous circuit verification by symbolic simulation: An illustration. In Advanced Research in L'LSI: Proceedings of the Sixth MITT Conference. April 1990.
[BCD86] M. C. Browne. E. M. Clarke, and Dill. Automatic verification using temporal logic: Two new examples. In George J. Milne and P. A. Subramanyam, editors. Formal Aspects of VLSI Design. Proceedings of the 1985 Edinburgh Workshop on VLSI, pages 113-124. North-Holland. 1986.
[BCDM86] M. C. Browne. E. M. Clarke. D. L. Dill. and B. Mishra. Automatic verification of sequential circuits using temporal logic. IEEE Transactions on Computers. C'-35(12):10:351044. 1986.
[BCG86] M.C. Browne, E. M. Clarke, and O. Grumberg. Reasoning about networks with many identical finite state processes. In ACM Symp. Principles of Distribated Computing 5. 1986.
[BCG87] M.C. Browne, E. M. Clarke, and O. Grumberg. Characterizing kripke structures in temporal logic. Technical Report 87-104, Carnegie-Mellon University. Computer Science Department, 1987.
[BCL91a] Jerry R. Burch, Edmund M. Clarke. and David E. Long. Representing circuits more efficiently in symbolic model checking. In Proceedings of the 28th ACM/IEEE Design Automation Conference. San Francisco. California. June 1991.
[BCL91b] Jerry R. Burch, Edmund M. Clarke, and David E. Long. Symbolic model checking with partitioned transition relations. In A. Halaas and P. B. Denyer. editors, Proceedings of the IFIP International Conference on Very Large Scalf Integration. Edinburgh. Scotland. August 1991.
$\left[\mathrm{BCM}^{+} 90\right]$ J. R. Burch. E. M. Clarke. K. L. McMillan. D. L. Dill. and J. Hwang. Symbolic model checking: $10^{20}$ states and beyond. In Proceedings of the Fifth Annual Symposium on Logic in Computer Science. June 1990.
[BF59a] S. Bose and A. Fisher. Verifying pipelined hardware using symbolic logic simulation. In IFEE Internationa conference on Computer Design. 1989.
[BF89b] Soumitra Bose and Allan L. Fisher. Automatic verification of synchronous circuits using symbolic logic simulation and temporal logic. In Luc Claestn. editor. Procetdings of the IMEC-IFIP International Workshop on Applitd Formal Methods For Correct VLSI Design. pages 759-764. . $o v e m-$ ber 1989.
[Bil87] J. P. Billon. Perfect normal forms for discrete functions. Technical Report 37019. BILL. March 1987.
[Bry86] R. E. Bryant. Graph-based algorithms for boolean function manipulation. IEEE Transactions on Computers, C-35(8), 1986.
[Bry88] Randal E. Bryant. Verifying a static ram design by logic simulation. In Jonathan Allen and F. Thomson Leighton, editors, Advanced Research in VLSI: Proceedings of the Fifth MIT Conference, pages 335-349. MIT Press, 1988.
[Bry91] R. E. Bryant. On the complexity of VLSI implementations and graph representations of boolean functions with application to integer multiplication. IEEE Trans. Computers, 40(2):205-213, February 1991.
[BS90] Randal E. Bryant and C'arl-Johan Seger. Formal verification of digital circuits using symbolic ternary system models. In Robert Kurshan and Edmund M. Clarke, editors, Workshop on Computer-Aided Verification, New Brunswick, New Jersey, June 1990. C'enter for Discrete Mathematics and Theoretical Computer Science (DIMACS).
[Bur84] J. P. Burgess. Basic tense logoc. In D. Gabbay and F. Guenthner. editors. Handbook of Philosophical Logic. Volume II: Extensions of Classical Logic, pages 89-134. D. Reidel. 1984.
[CBM[89] Olivier Coudert, Christian Berthet, and Jean Christophe Madre. Verification of synchronous sequential machines based on symbolic execution. In Joseph Sifakis, editor. Automatic Verification Methods for Finite State Systems. International Workshop. Grenoble. France, volume 407 of Lecture Notes in Computer Science. Springer-Verlag, June 1989.
[CDK.90] E. M. Clarke. I. A. Draghicescu. and R. P. Kurshan. A unified approach for showing language containment and equivalence between various types of $\omega$-automata. In Fifteenth Colloquium on Trees in Algebra and Programming, volume

431 of Lecture Notes in Computcr Science, Copenhagen, Denmark, May 1990. Springer-Verlag.
[CE81a] E. M. Clarke and E. A. Emerson. Characterizing properties of parallel programs as fixpoints. In Seventh International Colloqium on Automata. Languages. and Programming. volume 85 of LNCS, 1981.
[CE81b] E. M. Clarke and E. A. Emerson. Synthesis of synchronization skeletons for branching time temporal logic. In Dexter Kozen, editor, Logic of Programs: Workshop. volume 1:31 of Lecture Notes in C'omputrr Scienct. Yorktown Heights. New York. May 1981. Springer-Verlag.
[CES86] E. M. Clarke, E. A. Emerson. and A. P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. ACM Transactions on Programıing Languages and Systems. S(?):2+4-263. 1986.
[CGK89] E. M. Clarke. O. Cirumberg. and R. P. Kurshan. A synthesis of two approaches for verifying finite state concurrent systems. In Logic at Botik ©9. Symposium on Logical Foundations of Computer Science. volume 363 of Lecturf Vote. in C'omputer Scienre. Springer-Verlag. July 1989.
[CGLy?] Edmund M. Clarke. Orna (irmmberg. and David E. Long. Model checking and abstraction. In Proceedings of the Annual AC'M Symposium on Principles of Programmony Languages. January 1992.
[ChaS7] D. Chapman. Planning for conjunctive goals. Artificerl Intelligence. 32:33:3-3:7. 1957.
[C.HPPS7] P. Caspi. N. Halbwachs. D. Pilaud. and J. .I. Plaice. Lustre. a declarative language for programming sunchronous. systems. In 14th A('./ Symp. on Principles of Promram. ming Languages. January 19si.
[CLM89a] E. M. Clarke, D. E. Long, and K. L. McMillan. Compositional model checking. In Proceedings of the Fourth Annual IEEE Symposium on Logic in Computer Science, 1989.
[CLM89b] E. M. Clarke, D. E. Long, and K. L. McMillan. A language for compositional specification and verification of finite state hardware controllers. In 9th International Symposium on Hardware Description Languages and their Applications, 1989.
[CM88] K. M. Chandy and J. Misra. Parallel program design : a foundation. Addison-Wesley, 1988.
[CMB91] Olivier Coudert. Jean Christophe Madre, and Christian Berthet. Verifying temporal properties of sequential machines without building their state graphs. In E. M. Clarke and R. P. Kurshan. editors. Computer Aided Verification. '90. volume 3 of DIMACS Series in Discrete Mathematics and Theoretical Computer Science. 1991.
[Cou91] O. Coudert. SIAM: une boite a outils pour la preave formelle de systèmes sequentiels. PhD thesis. Ècole Va tionale Supérieure des Télécommunications, 1991.
[DC86] D. L. Dill and E. M. Clarke. Automatic verification of asynchronous circuits using temporal logic. IEE Proceedings. Pt. E, 1333(5):276-282. September 1986.
[Dil88] D. Dill. Trace theory for automatic hierarchical verification of speed-independent circuits. Technical Report 88119. Carnegie Mellon Viniversity, Computer Science Dept. 1988.
[EL36] E. A. Emerson and C.-L. Lei. Efficient model checking in fragments of the propositional mu-calculus. In Proceedings of the Second Annual Symposium on Logic in Computer Science. IEEE C'omputer Society Press, June 1986.
[FB89] Allan L. Fisher and Randal E. Bryant. Performance of COSMOS on the IFIP workshop benchmarks. In Proceedings of IMEC Conference, 1989.
[FMK90] M. Fujita, Y. Matsunaga, and T Kakuda. Automatic and semi-automatic verification of switch-level circuits with temporal logic and binary decision diagrams. In ICCAD. pages 38-41, 1990.
[GL91] Orna Grumberg and David E. Long. Model checking and modular verification. In .J. ('. M. Baeten and J. F. Groote. editors. Proceedings of CO.VC'LR 91: 2nd International Conference on Concurrency Theory, volume 527 of Lecturf Notes in Computer Sciencr. Springer-Verlag. August 1991.
[God90] P. Godefroid. Using partial orders to improve automatic verification methods. In Workshop on Computer Aided Verification. 1990.
[GS] S. M. German and A. P. Sistla. Reasoning about svstems with many processes. (iTE Laboratories Inc.. Waltham. Massachusetts.
[GS91] Susanne Grat and Bernhard Steffen. Compositional minimization of finite state systems. In Robert Kurshan and Edmund M. Clarke, editors. (omputer-Aided litrification. Procecdings of the 1990 Workshop. volume 3 of DIM:A('S Serif.: in Discrete . Wathematios and Theoretical Computr Science. American Mathematical Society: 1991.
[GW91] P. Godefroid and P. Woiper. I partial approach to model checking. In LIC'S. 1991.
[Hoa69] C. A. R. Hoare. An axiomatic hasis for computer programming. (omm. AC.M. 12:56-50.50. October 1969.
[KC90] S. Kimura and E. M. Clarke. I parallel algorithm for constructing binary decision diagrams. In 1990 IEEE Intronttional Conference on ('omputer Design. September l990.
[KM89] R. Kurshan and K. L. McMillan. A structural induction theorem for processes. In ACM Symposium on Principles of Distributed Computing, Edmonton, Alberta, 1989.
[Kro87] Fred Kroger. Temporal Logic of Programs. Springer-Verlag, 1987.
[Kur85] R. P. Kurshan. Modelling concurrent processes. In Symp. in Applied Math. 31. pages 45-57. 1985.
[Kur86] R. P. Kurshan. Testing containment of $\omega$-regular languages. Technical Report 1121-861010-33-TM, Bell Laboratories. 1986.
[Kur87] R. P. Kurshan. Reducibility in analysis of coordination. In LNCS, volume 103, pages 19-39. Springer-Verlag, 1987.
[LN91] B. Lin and A. R. Newton. Efficient symbolic manipulation of equivalence relations and classes. In IMEC-IFIP International Workshop on Formal Methods in VLSI Design. Miami, Florida. 1991.
[LPS5] Orna Lichtenstein and Amir Pnueli. Checking that finite state concurrent programs satisfy their linear specification. In Conference Record of the Twelfth Annual ACM Symposium on Principles on Programming Languages. 1985.
[LTN90] B. Lin. H. J. Touati, and A. R. Newton. Don't care minimization of multi-level sequential logic networks. In $I C^{\circ}$ CAD, pages $414-417,1990$.
[Mar85] A. J. Martin. The design of a self-timed circuit for distributed mutual exclusion. In Henry Fuchs. editor. 198.5 Chapel Hill Conference on VLSI. pages 245-260. Computer Science Press. 1985.
[MB59] D. E. Muller and W. S. Bartky. A theory of asynchronous circuits. In The Annals of the Computation Laboratory of

Harvard University. Volume $X X I X$ : Proceedings of an International Symposium on the Theory of Switching, Part I, pages 204-243. Harvard University Press, 1959.
[Mil80] R. Milner. A Calculus of Communicating Systems, volume 92 of Lecture Notes in Computer Science. SpringerVerlag, 1980.
[Mil83] G. J. Milne. Circal, calculus for circuit descriptions. Integration, 1:121-160, 1983.
[MO81] Y. Malachi and S. S. Owicki. Temporal specifications of selftimed systems. In H. T. Kung. B. Sproull, and Ci. Steele. editors, VLSI Systems and Computations. 1981.
[MP81] Z. Manna and A. Pnueli. Verification of concurrent programs: the temporal framework. In R. S. Boyer and J. S. Moore, editors. The Correctness Problem in Computer Science, pages 215-273. 1981.
[MS91] K.L. McMillan and J. Schwalbe. Formal verification of the Encore Gigamax cache consistency protocol. In International Symposium on Shared Memory Multiprocessors. 1991.
[NH84] R. De Nicola and M. Hennessy: Testing equivalences for processes. Theoretical Computer Science. 34(3:3). 1984.
[NPW81] M. Nielsen. G. Plotkin. and G. Winskel. Petri nets. event structures and domains. part I. Theoretical Computer Science. 13:85-108. 1981.
[Par74] David Park. Finiteness is mu-ineffable. Theory of Computation Report No. 3. The University of Warwick. 197t.
[PL89] D. K. Probst and H. F. Li. Abstract specification. composition. and proof of correctness of delay-insensitive circuit. and systems. Technical report. ( oncordia University. Dept. of Computer Science. 1989.
[PL90] D. K. Probst and H. F. Li. Using partial order semantics to avoid the state explosion problem in asynchronous systems. In Workshop on Computer Aided Verification, 1990.
[Pnu77] A. Pnueli. The temporal semantics of concurrent programs. In 18th Symposium on Foundations of Computer Science, 1977.
[Pnu86] A. Pnueli. Applications of temporal logic to the specification and verification of reactive systems: A survey of current trends. In Lecture Notes in Computer Science, volume 224 , pages 510-584. Springer-Verlag, 1986.
[QS81] J. P. Quielle and J. Sifakis. Specification and verification of concurrent systems in CESAR. In Proceedings of the Fifth International Symposium in Programming, 1981.
[RUT1] N. Rescher and A. Urquhart. Temporal Logic. SpringerVerlag, 1971.
[Saci7] E. Sacerdoti. A Structure for Plans and Behavior. American Elsevier, 1977.
[Sei80a] C'. L. Seitz. Ideas about arbiters. Lambda, 10(14), 1980.
[Sei80b] C. L. Seitz. System timing. In Carver Mead and Lynn Conway, editors. Introduction to VLSI Systems. pages $213-$ 262. Addison-Wesley, 1980.
[SG89] Z. Shtadler and O. Grumberg. Network grammars. communication behaviors and automatic verification. In Workshop on Automatic Verification Methods for Finite State Systems. LNCS, pages 151-165. Springer-Verlag, 1989.
[Smu68] R. M. Smullyan. First Order Logic. Springer-Verlag, 1968.
[Taryos] A. Tarski. A lattice-theoretical fixpoint theorem and its applications. Pacific J. Math., 5:285-309. 1955.
[TBK91] H. J. Touati, R. K. Brayton, and R. P. Kurshan. Testing language containment for $\omega$-automata using BDD's. In IMEC-IFIP International Workshop on Formal Methods in VLSI Design, Miami, Florida, 1991.
[Tho84] R. H. Thomason. Combinations of tense and modality. In D. Gabbay and F. Guenthner, editors, Handbook of Philosophical Logic. Volume II: Extensions of Classical Logic. pages 89-134. D. Reidel, 1984.
[TSL+90] H. J. Touati, H. Savoj, B. Lin, R. K. Brayton. and A. Sangiovanni-Vincentelli. Implicit state enumeration of finite state machines using BDD's. In ICCAD. pages 130133. 1990.
[Val89] A. Valmari. Stubborn sets for reduced state space generation. In 10th Int. Conf. on Application and Theory of Petri Nets, 1989.
[Val90] A. Valmari. A stubborn attack on the state explosion problem. In Workshop on Computer Aided Verification. 1990.
[vdS83] Jan L. A. van de Snepscheut. Trace Theory and VLSI design. PhD thesis. Department of Computer Science. Eindhoven University of Technology. October 1983.
[WL89] P. Wolper and V Lovinfosse. Verifying properties of large sets of processes with network invariants. In Workshop on Automatic Verification Methods for Finite State S!ystem. LNC'S, pages 68-80. Springer-Verlag. 1989.
[Wol83] Pierre Wolper. Temporal logic can be more expressive. Information and Control. 36:72-99. 1983.
[YTK91] Tomohiro Yoneda. Yoshihiro Tohma. and Yintaka Kondo. Acceleration of timing verification method based on time
 1991.


[^0]:    ${ }^{1}$ The question of termination is in any event not undecidable for hardware systems. since they are not computation universal (only programs are).
    $\because$ These are usually ralled Kripke frames and Kripke models. after one of the first mathematicians to give a model theoretic interpretation of modal logic.

[^1]:    $3^{3}$ 'lassically, the symbol $\square$ is used to represent necessity. and $\rangle$ is used to represent possibility. The symbols $A$ and $E$ are used here for consistency with [BAMPYI].

[^2]:    ('TL is actually a subset of a more general temporal logic described in [' EX1al. adopting the syntax of $\{(3 A M P \times 1\}$.

[^3]:    ${ }^{5}$ With one additional distinction: in ('TL, the future is taken to include the present. Thus. if $p$ holds in the present, then so does $F p$.

[^4]:    ${ }^{\text {' }}$ This is essentially an algebraic interpretation of logic, where we embed the formulas of the logic in a Boolean algebra ( $\left.P\left(S^{\prime}\right), 0,1 . \cap, \cup,-\right)$, with n representing conjunction. U representing disjunction and - (set complement.) representing negation

[^5]:    We represent a Kripke model pictorially by drawing the eraph (.i. R) and labeling ath state with the atomic propositions which are true in that state.

[^6]:    ${ }^{*}$ Vormally, when we discussing switching functions, we think of a Bonlean formula as represented by a function rather than the set of satisfyine truth assignments , if the formula. A Boolean formula $f$ over an ondered set of variables

[^7]:    ' nfortunately, this is the opposite of the numbering adopted by Bryant. but it makes the pronts clearer.

[^8]:    ${ }^{10}$ There is an alternative formulation of OBDDs due to ('larke [KC 90 ] which does not require $l$ and $h$ to be distinct, but requires that $i=d_{l}+1=d_{h}+1$. In this case. the OBDD for a function $f$ is exactly the minimal DFA recognizing the language $\left\{x \in\{0 \text {.I }\}^{n} \mid f(x)=I\right\}$. Thinking of OBDDs as minimal DFAs can provide useful insights into the complexity of representing certain classes of functions as UBDDs.

[^9]:    ${ }^{11}$ The latter number should be regarded as being accurate only to within a factor of two. since the garbage collector in the implementation scavenges for unreferenced nodes only when the number of nodes doubles.
    ${ }^{12}$ This ran be shown using the technique of [Bry9l]. It is sufficient that for any variable order there is some rotation such that when the order is cut in half. information proportional to $n$ must be passed from one half to the other.

[^10]:    ${ }^{13} \mathrm{~A}$ discussion of which state machine morlel is more suitable for circuit design is beyond the scope of this work. In general. conditions would have to be imposed on either moitel in order to make it implementable in a given design style. For liscussion of asynchronous lesign terhniques. see [MB59. Seix(0b].

[^11]:    ${ }^{14}$ The tudErests algorithm of section 2.3.4. which combines conjunction and quantification in a bottom-up manner is also at example of early quantification.

[^12]:    "Fomerson and Lei [ELxfi] nave a motel checking algorithm for a somewhat differ-- nt wersion of the Mu- 'alculus. and showed that there are formulas in this logic that
    

[^13]:    
    

[^14]:    ${ }^{17}$ Actually function graphs with negated ares were used for this calculation [Bilst? hence the number of nodes mav be slightly smaller than what would i,. , ,biamed usmg ()BDDs.

[^15]:    'International Symposilim on ('ircutts and Systems

[^16]:    ${ }^{15}$ Personal communication.

[^17]:    ${ }^{\text {L }}$ In fact. the number of possible transitions from a given state ranges from 6 to 12. The probability of a random simulation run executing this trace is therefore in the range $0^{-13}=7.7 \times 10^{-11}$ to $12^{-13}=9.3 \times 10^{-15}$. The expected time for a random simulation to exhibit this behavior would be somewhere between 2.4 years and 29 millenia, assuming the simulation could be carried out at 10.000 steps per second.

[^18]:    ${ }^{2}$ For other applications of this kind of argument. see $[$ Bry 91$]$.

[^19]:    ${ }^{3}$ In fact. such a deadlock, involving an interaction between the memory and processor subsystems was known to the Encore engineers. The memory system. when busy. would stall any new repluests. but the stalled request would still remain in the memory system's pipeline for four clock cycles. Thus. when the processor retried the request four clock cyrles later it would be stalled again, and the proress would repeat indefinitely.

[^20]:    ${ }^{4}$ Personal communication
    ${ }^{2}$ Personal communication

[^21]:    ${ }^{3}$ Satisfiability of a Boolean forr ula in conjunctive normal form, with three literals in each conjunct.

[^22]:    let $B$ be the set of the cutoff points, $T_{s}=\varnothing$
    while $B$ is not empty do
    3 let $t$ the the element of $B$ with the fewest spoilers
    $t$ if $t$ has no spoilers. then backtrack
    5 choose an element $t^{\prime}$ from the spoilers of $t$
    6 add $t^{\prime}$ to $T_{s}$
    7 delete all transitions in conflict with $T_{s}$
    3 end do

