

AD-A136 466

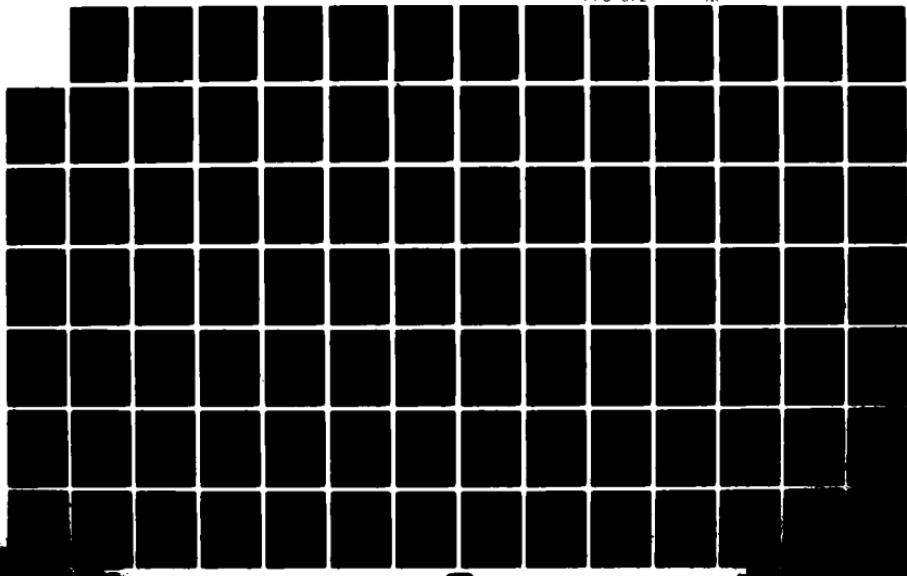
VERIFICATION OF HARDWARE DESIGN CORRECTNESS: SYMBOLIC  
EXECUTION TECHNIQUE.. (U) STANFORD UNIV CA COMPUTER  
SYSTEMS LAB W E CORY JUN 83 TR-83-241 DAAG29-80-K-0046

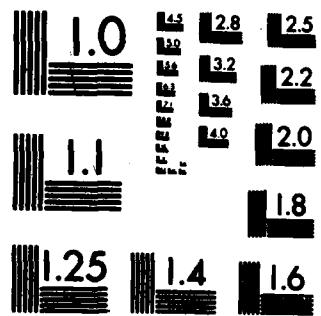
12

UNCLASSIFIED

F/G 9/2

NI





MICROCOPY RESOLUTION TEST CHART  
NATIONAL BUREAU OF STANDARDS-1963-A

# **COMPUTER SYSTEMS LABORATORY**

DEPARTMENTS OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE  
STANFORD UNIVERSITY · STANFORD, CA 94305



A 1 3 6 4 6 6

## **VERIFICATION OF HARDWARE DESIGN CORRECTNESS: SYMBOLIC EXECUTION TECHNIQUES AND CRITERIA FOR CONSISTENCY**

Warren E. Cory

**Technical Report No. 83-241**

DTIC June 1983  
ELECTED  
DEC 30 1983 D  
H

This research was supported by ARO Contract No. DAAG29-80-K-0046.

DISTRIBUTION STATEMENT A  
Approved for public release

DTIC FILE COPY

83 12 30 054

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM RECIPIENT'S CATALOG NUMBER
1. REPORT NUMBER <b>17200.7-EL</b>	2. GRANT ACCESSION NO. <b>A1364PD</b>	3. TYPE OF REPORT & PERIOD COVERED <b>Technical</b>
4. TITLE (and Subtitle) <b>Verification of Hardware Design Correctness: Symbolic Execution Techniques and Criteria for Consistency</b>	5. PERFORMING ORG. REPORT NUMBER	
7. AUTHOR(s) <b>Warren E. Cory</b>	8. CONTRACT OR GRANT NUMBER(s) <b>DAAG29 80 K 0046</b>	
9. PERFORMING ORGANIZATION NAME AND ADDRESS <b>Stanford University Stanford, CA 94305</b>	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS	
11. CONTROLLING OFFICE NAME AND ADDRESS <b>U. S. Army Research Office Post Office Box 12211 Research Triangle Park, NC 27709</b>	12. REPORT DATE <b>JUN 83</b>	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)	13. NUMBER OF PAGES <b>217</b>	
16. SECURITY CLASS. (of this report) <b>Unclassified</b>		
15a. DECLASSIFICATION/DOWNGRADING SCHEDULE		
16. DISTRIBUTION STATEMENT (of this Report) <b>Approved for public release; distribution unlimited.</b>		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from 16a) <b></b>		
18. SUPPLEMENTARY NOTES <b>The view, opinions, and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation</b>		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) <b></b>		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) <b>Symbolic executions of event-driven simulations tend to split into many cases. The substitution of non-selective trace for selective trace eliminates this case splitting for some elements. Case merging using conditional expressions is a more general technique for eliminating case splitting. New criteria for the comparison of two programs apply to both deterministic and nondeterministic models, with and without don't-cares. They allow hierarchical verification and facilitate the verification of programs with loops. A modification of a top-down hierarchical design methodology allows partitioned verification even when the design introduces new dependencies.</b>		



**VERIFICATION  
OF HARDWARE DESIGN CORRECTNESS:  
SYMBOLIC EXECUTION TECHNIQUES  
AND CRITERIA FOR CONSISTENCY**

Warren E. Cory

Technical Report No. 83-241

June 1983

Computer Systems Laboratory  
Departments of Electrical Engineering and Computer Science  
Stanford University  
Stanford, California 94305

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	<input type="checkbox"/>
By	
Distribution	
Priority Codes	
Ref ID	
Comments	

*AC*

**Abstract**

Symbolic executions of event-driven simulations tend to split into many cases. The substitution of non-selective trace for selective trace eliminates this case splitting for some elements. Case merging using conditional expressions is a more general technique for eliminating case splitting.

New criteria for the comparison of two programs apply to both deterministic and nondeterministic models, with and without don't-cares. They allow hierarchical verification and facilitate the verification of programs with loops. A modification of a top-down hierarchical design methodology allows partitioned verification even when the design introduces new timing details.

**Key Words and Phrases:** event-driven simulation with selective trace, non-selective trace, case analysis, case merging, graph models for behavior, hierarchical partitioned (piecewise) verification

© Copyright 1983  
by  
Warren Eugene Cory

Scripture text on dedication page: *The Holy Bible, New International Version.*  
Copyright © 1978 by New York International Bible Society. Used by permission of  
New York International Bible Society.

## Abstract

This thesis investigates two aspects of the functional verification problem in digital hardware design. First, It discusses the reduction of case splitting in the symbolic execution of event-driven simulations. Second, it defines criteria for consistency between two hardware descriptions.

Symbolic executions of event-driven simulations tend to split into many subcases due to the event detection mechanism. The substitution of non-selective trace for selective trace eliminates this case splitting for combinational elements. In general, an inductive assertion analysis or a symbolic execution test can determine whether the substitution is correct for a particular element. Case merging using conditional expressions is a more general technique for eliminating case splitting, but it tends to generate complex symbolic expressions. Case merging is most likely to be beneficial for functional verification when the simulation relation is sparse and the symbolic expression complexity grows slowly.)

Milner and Brand defined criteria for the comparison of two programs to establish their consistency. This thesis extends their work by defining a criterion that applies to both deterministic and nondeterministic models, with and without don't-cares. Hierarchical verification is demonstrated for the new criterion. This criterion is also modified to facilitate the verification of programs with loops. The modified criterion is shown to retain intuitive notions of correctness captured by the first criterion.

Finally, a modification of the hierarchical design methodology allows the application of Brand's theorem for partitioned verification even when the design introduces new timing details. The introduction of new timing details is performed separately from the introduction of new structural detail.

## Acknowledgments

Willem vanCleemput, my adviser, patiently gave his support and encouragement. Early suggestions by John Hennessy substantially improved the thesis. Later, John Hennessy and Robert Gray served on the reading committee. Robert Shostak at SRI International provided a theorem prover for an experimental symbolic execution system which yielded some of the ideas in Chapters 4 and 5. Edward Rice and John Meadors encouraged and advised me in my high school and undergraduate days. Thank you, everyone.

A National Science Foundation Graduate Fellowship provided financial assistance for my first three years of graduate study. Willem vanCleemput and Robert Dutton supported me primarily by ARO Contract No. DAAG29-80-K-0046, with earlier support from JSEP Contract Nos. N-00014-85-C-0601 and DAAG29-79-C-0047 and gift funds from Tektronix, Inc.

Finally, I wish to thank my friends. You gave me your unswerving, caring support through good times and bad. And it was only by the considerable grace of God that I could aspire to undertake this task, let alone finish it.

**In the beginning was the Word, and the Word was with God, and the Word was God. He was with God in the beginning.**

**Through him all things were made; without him nothing was made that has been made. In him was life, and that life was the light of men. The light shines in the darkness, but the darkness has not understood it.\***

**To Mom and Dad**

---

**\* Scripture text: *The Holy Bible, New International Version*. Copyright © 1978 by New York International Bible Society. Used by permission of New York International Bible Society.**

# Table of Contents

<b>1. Introduction</b>	<b>1</b>
1.1. Hardware Descriptions and Simulation	1
1.2. Hardware Verification	2
1.3. Functional Verification	3
1.3.1. Analysis of a Hardware Description	4
1.3.2. Comparison of Analysis Results for Two Descriptions	5
1.4. Organization of the Thesis	6
<b>2. Previous Work</b>	<b>9</b>
2.1. Symbolic Execution for Software Testing	9
2.2. Functional Verification of Hardware Designs	11
2.2.1. Comparison of Specification and Design	11
2.2.2. Comparison of Manual and Automated Designs	14
2.2.3. Comparison of Manual and Automated Specifications	15
2.2.4. Conclusion	16
2.3. Program Analysis with Conditional Expressions	17
2.4. Criteria for Consistency	18
2.5. Remaining Problems	19
2.5.1. Symbolic Execution	20
2.5.2. Consistency	21
<b>3. Symbolic Event-driven Simulation</b>	<b>23</b>
3.1. Review of Symbolic Execution	24
3.1.1. Symbolic Computations	24
3.1.2. Case Splitting	26
3.2. Model for Event-driven Simulation	31
3.3. Event-driven Simulation With Case Splitting	35
<b>4. Non-selective Trace</b>	<b>39</b>
4.1. Combinational Process	40
4.2. Applicability of Non-selective Trace	43
4.2.1. Generalized Process of Interest	44
4.2.2. Program State	45
4.2.3. Characteristics of the Generalized Process	46

4.2.4. Comparison of the Two Processes	48
4.2.5. Use of Assertions	51
4.2.6. Another Generalized Process	54
4.3. Non-selective Trace Summary	57
<b>5. Case Merging</b>	<b>61</b>
5.1. Case Merging Examples	62
5.1.1. Case Merging with Single Entry Point	62
5.1.2. Case Merging with Multiple Entry Points	66
5.1.3. Summary of Examples	69
5.2. Analysis of Acyclic Flow Graphs	69
5.2.1. Generality of the Construction	70
5.2.2. The Construction	71
5.2.3. Use of Constructed Process	75
5.2.4. Application of Construction to Routines	76
5.3. Analysis of Flow Graphs with Loops	77
5.3.1. New Outline for Case Merging with Loops	78
5.3.2. Selection of Basic Blocks for Execution	82
5.4. Optimizations	86
5.5. Casc Merging Summary	89
5.6. Symbolic Execution Summary	91
<b>6. Symbolic Execution Examples</b>	<b>93</b>
6.1. Symbolic Expression Complexity	94
6.2. Case Splitting and Case Merging Complexity	97
6.3. Blackjack Example	99
6.3.1. Blackjack Descriptions	99
6.3.2. Simulation Relation	100
6.3.3. Symbolic Case Splitting Results	103
6.3.4. Symbolic Case Merging Results	104
6.3.5. Conclusion to First Example	107
6.4. Associative Memory Example	108
6.5. Conclusion	111
<b>7. Criteria for Consistency</b>	<b>115</b>
7.1. Overview	115
7.2. Notation	117
7.3. Simulation between Deterministic Programs	117
7.4. Simulation between Nondeterministic Programs	118
7.5. Program Models	122
7.6. Consistency Criteria	125
7.7. Implied Implementation vs. Implied Specification	127
7.7.1. Deterministic Programs without Don't-Cares	127

7.7.2. Deterministic Programs with Don't-Cares	127
7.7.3. Nondeterministic Programs without Don't-Cares	132
7.7.4. Nondeterministic Programs with Don't-Cares	133
7.8. Deterministic Model Reconsidered	134
7.9. Closure Again	137
7.9.1. Stopping Points in Loops	137
7.9.2. New Criterion for Consistency	138
7.10. Relation between Extended and Simple Criteria	141
7.10.1. Preview	141
7.10.2. New Relation	143
7.10.3. Suitability of New Relation	144
7.10.4. Proof that New Relation Satisfies Simple Criterion	151
7.11. Hierarchical Verification	158
7.12. Summary	165
<b>8. Partitioned Verification</b>	<b>167</b>
8.1. Review of Partitioned Verification	168
8.1.1. Brand's Program Model	169
8.1.2. Definitions	170
8.1.3. Theorem for Partitioned Verification	171
8.2. Partitioned Verification Problem	172
8.3. Modification of the Design Methodology	175
8.4. Summary	178
<b>9. Summary</b>	<b>181</b>
9.1. Contributions	181
9.1.1. Non-selective Trace	181
9.1.2. Case Merging	183
9.1.3. Consistency of Hardware Descriptions	184
9.2. Suggestions for Further Research	186
9.2.1. Case Merging	186
9.2.2. Criteria for Consistency	186
9.2.3. Symbolic Execution for Functional Verification	188
9.2.4. Incorporation of Other Techniques for Verification	188
9.2.5. Strengthening the Weak Links	189
<b>Appendix A. Notation</b>	<b>191</b>
A.1. Formulae	191
A.2. Relations	191
<b>Appendix B. Notes on Proof of Theorem 8.1</b>	<b>193</b>
B.1. Counterexamples	194
B.2. Comments on Proof of Lemma 2	196

<b>Appendix C. Symbolic Execution of Blackjack Example</b>	<b>199</b>
C.1. Symbolic Execution with Case Splitting	199
C.2. Symbolic Execution with Case Merging	208

## List of Figures

<b>Figure 1-1:</b> Partitioned comparison of graph models	6
<b>Figure 2-1:</b> Symbolic execution example	10
<b>Figure 3-1:</b> Some control constructs for case splitting execution	31
<b>Figure 3-2:</b> Outline of event-driven simulator	33
<b>Figure 3-3:</b> Process constructs using pause	34
<b>Figure 3-4:</b> Modified equivalent event-driven simulator	36
<b>Figure 4-1:</b> NOR-gate process	40
<b>Figure 4-2:</b> Generalized process for non-selective trace	44
<b>Figure 4-3:</b> Modified generalized process	50
<b>Figure 4-4:</b> Final generalized process	50
<b>Figure 4-5:</b> Modified NOR-gate process	51
<b>Figure 4-6:</b> NOR-gate process with assertions	52
<b>Figure 4-7:</b> Semantics of pause for VCG	53
<b>Figure 4-8:</b> Generalized process for testing without assertions	54
<b>Figure 4-9:</b> Generalized process with assertions	55
<b>Figure 4-10:</b> Three paths for VCG	55
<b>Figure 4-11:</b> Symbolic execution for test of path (a)	56
<b>Figure 5-1:</b> Edge-triggered D flip-flop model	62
<b>Figure 5-2:</b> Flow graph for D flip-flop process	63
<b>Figure 5-3:</b> Transformed version of D flip-flop	64
<b>Figure 5-4:</b> Process with multiple entry points	66
<b>Figure 5-5:</b> Flow graph for <i>OddReg</i> process	67
<b>Figure 5-6:</b> Modified process with multiple entry points	68
<b>Figure 5-7:</b> Algorithmically generated process for case merging	75
<b>Figure 5-8:</b> Case merging outline for process without loops	77
<b>Figure 5-9:</b> Equivalent case merging version using basic block set	78
<b>Figure 5-10:</b> Case merging outline for process with loops	79
<b>Figure 5-11:</b> Graph with loop not broken by dismissal point	83
<b>Figure 5-12:</b> Priority assignment for graphs with loops	86
<b>Figure 5-13:</b> Example for optimizations	88
<b>Figure 5-14:</b> Second example for optimizations	89
<b>Figure 6-1:</b> Blackjack finite state machine specification	99
<b>Figure 6-2:</b> Blackjack machine implementation	101

Figure 6-3: Simulation relation for Blackjack descriptions	102
Figure 6-4: Case splitting results for one path	103
Figure 6-5: Case merging results for one variable	105
Figure 6-6: Final case merging symbolic program counter	106
Figure 6-7: Associative memory description	108
Figure 6-8: Second associative memory description	109
Figure 6-9: Simulation relation for associative memory	109
Figure 6-10: Case splitting symbolic execution results	110
Figure 6-11: Case merging results for Figure 6-7	111
Figure 6-12: Case merging results for Figure 6-8	112
Figure 7-1: Functional verification problem	119
Figure 7-2: Example motivating use of closure	121
Figure 7-3: Example - Specification program	129
Figure 7-4: Example - Implementation program	129
Figure 7-5: Example - Relations defining model	130
Figure 7-6: Four relations for test for consistency	130
Figure 7-7: Incorrect programs satisfying Eq. (7.8)	131
Figure 7-8: Original example motivating use of closure	138
Figure 7-9: Condition for simplified definition	139
Figure 7-10: Second example motivating use of closure	140
Figure 7-11: Incorrect verification by Eq. (7.18)	140
Figure 7-12: Properties of new simulation relation	145
Figure 7-13: Three types of terminal vertices	146
Figure 7-14: Reachability of terminal vertex for case (a)	146
Figure 7-15: Impossibility of two distinct terminal vertices	148
Figure 7-16: Connectivity between states in equivalence class	153
Figure 7-17: Proof of Eq. (7.26): first case	155
Figure 7-18: Proof of Eq. (7.26): second case	156
Figure 7-19: Need for Eq. (7.33)	160
Figure 7-20: Need for Eq. (7.34)	160
Figure 7-21: Lemma asserting existence of path in $P_1$	162
Figure 8-1: Changes in data type	168
Figure 8-2: Timing diagrams for write cycle	173
Figure 8-3: Graph models for high-level write cycle	174
Figure 8-4: Graph models for low-level write cycle	174
Figure 8-5: Initial system specification	176
Figure 8-6: Specification with translators	177
Figure 8-7: Interface design refinement	178
Figure B-1: First counterexample to Lemma 2	194
Figure B-2: Second counterexample to Lemma 2	195
Figure B-3: Three cases in proof of Lemma 2	196

This thesis investigates two aspects of the verification problem. First, it considers methods for reducing case splitting in the symbolic execution of event-driven simulations. Second, it discusses criteria for the comparison of two hardware descriptions.

## Chapter 1

### Introduction

This thesis investigates the verification of electronic digital hardware designs. The two areas of concern are the symbolic execution of hardware descriptions and criteria for establishing the correctness of hardware descriptions.

#### 1.1. Hardware Descriptions and Simulation

Hardware descriptions are textual or graph descriptions often similar to computer programs and commonly used to describe the behavior or structure of a proposed hardware design. A verification procedure establishes that the design, expressed by a hardware description, satisfies design objectives given in the specification. The specification itself may be represented by a hardware description.

A computer simulation predicts the behavior of a hardware description for specific input data. Event-driven simulation is a particular type of simulation which represents a design as a collection of independent components with some means of communication between them. In such a simulation, a component remains inactive unless some condition arises (the "event").

Symbolic execution is a modification of conventional simulation that tests the hardware description ideally for all possible input values simultaneously. Such an execution may split into many separate cases ("case splitting").

#### 1.2. Hardware Verification

The hardware verification problem can be divided into three areas: functional verification, performance verification, and design rule checking. The goal of functional verification is to establish that a proposed hardware design will correctly perform its intended function without regard to performance constraints such as speed, power, etc. Performance verification checks the design against these additional performance requirements. A design rule check tests the physical properties of the final physical design.

Although the boundaries between these three types of verification are not sharp, we can characterize the problem types by the inputs typically used in the verification process.

In functional verification, the inputs are usually two hardware descriptions of the hardware system to be built. Each description completely specifies the system behavior. One description is a system specification, while the other is a proposed design. The verification determines that the second description is a correct implementation of the first.

Alternatively, the input for functional verification may be a single hardware description annotated with assertions, where the assertions alone constitute a

complete specification of how the hardware should behave. The verification checks that the hardware description satisfies its assertions. For example, Floyd's inductive assertion method is one technique for performing this test [FI67, HK76, Ho69, LC79].

In performance verification, a typical input is a hardware description annotated with assertions. The assertions specify behavioral requirements, but do not completely specify the desired system behavior. For example, they may establish speed or power requirements. At a more abstract level, an assertion about a resource arbiter might require that if a system requests a resource infinitely often while the resource is available, then eventually it will be granted the resource, even in the presence of competing requests.

Design rule checking is like performance verification, except that the hardware description and assertions are structural (physical) rather than behavioral. Of course, behavioral considerations can motivate physical assertions ("design rules").

### 1.3. Functional Verification

This thesis considers the functional verification problem in which we compare two distinct behavioral hardware descriptions to establish that one is consistent with the other. In this thesis, we most often use the term "consistent" rather than "equivalent." The term "equivalent" suggests a condition that may be stronger than necessary for a design to be correct. A design can be correct (consistent with the original specification) without being equivalent to the original specification, due to

don't-cares, optimizations, and changes in the representation of data ("level of abstraction").

We can divide the functional verification process into two steps:

1. Analyze each description to characterize its behavior.

2. Compare the resulting behavioral characterizations.

This thesis considers each of these steps independently, as described in the following sections.

#### 1.3.1. Analysis of a Hardware Description

One method of analyzing a hardware description is to simulate it for all possible input values, recording the results. Such exhaustive simulation is intractable for almost all descriptions. However, symbolic execution is a popular modification that, in many cases, allows the exhaustive simulation of a description with a small number of simulations [CM80, EG77, HK76, Ki76a].

Previous research with symbolic execution has not applied it to event-driven simulations. A symbolic execution of an event-driven simulation may still be intractable because the execution splits into too many cases. This thesis will define a simple model for event-driven simulation similar to that supported by SABLE [Hi80]. It then investigates ways of reducing the numbers of cases in symbolic executions of this model.

### 1.3.2. Comparison of Analysis Results for Two Descriptions

Given a characterization of two behavioral hardware descriptions, we need a criteria for comparing the results to establish whether the proposed implementation is consistent with the specification. Milner and Braßd both suggested criteria for consistency using a graph model to abstractly represent behavior [BR79, Mi71]. This thesis considers the problem further using a similar model. The following observations apply:

homomorphism of the graph models is too strong a condition for consistency. Due to don't-cares, optimizations, and differences in levels of abstraction, we do not expect that the graph models for the specification and the implementation will be identical.

It is possible that a criterion for consistency may not be commutative. That is, graph models  $A$  and  $B$  could be consistent assuming model  $A$  corresponds to the specification, while failing to be consistent if  $B$  is the specification. Therefore in the discussion we always identify which graph corresponds to the specification.

A desirable property of a criterion for consistency is that it should be usable in a hierarchical, partitioned fashion. The desired hierarchical property is that the criterion should be transitive. If graph  $B$  is consistent with graph  $A$ , and graph  $C$  is consistent with  $B$ , then we would like to conclude that  $C$  is consistent with  $A$ . This facilitates a hierarchical design strategy in which the final design is generated by a series of refinements.

The desired partitioned property is that the criterion should be applicable in a

piecewise fashion. Suppose graph  $A$  consists of two subgraphs  $A1$  and  $A2$  representing two components in the system. Suppose that graph  $B$  is similarly partitioned. Then to prove  $B$  consistent with  $A$ , it should suffice to demonstrate that  $B1$  is consistent with  $A1$  and  $B2$  is consistent with  $A2$ . Figure 1-1 illustrates the situation.

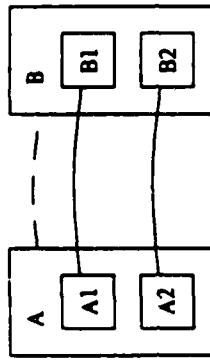


Figure 1-1: Partitioned comparison of graph models

## 1.4. Organization of the Thesis

Chapter 2 discusses previous work in symbolic execution, its use for hardware verification, and criteria for consistency. This chapter motivates the desirability of further work in the areas noted above.

Chapters 3 through 6 and chapters 7 and 8 form two independent sections, discussing symbolic execution and criteria for consistency, respectively. Either section may be read without consulting the other.

Chapter 3 reviews symbolic execution in greater detail and defines a model for event-driven simulation. Chapter 4 identifies simulations in which the numbers of

cases in symbolic executions can be reduced by substituting non-selective trace for selective trace. Chapter 5 describes case merging, a more general technique for reducing the number of cases. Case merging tends to generate complex symbolic expressions, increasing the cost of the analysis; Chapter 6 presents two examples and considers the cost and benefits of case merging.

Chapter 7 presents a criterion for consistency applicable to both deterministic and nondeterministic models with and without don't-cares. It demonstrates the transitive property for this criterion and extends it to facilitate the verification of programs with loops. Chapter 8 reviews Brand's result for partitioned verification, noting a problem that arises when the level of abstraction is changed in a design. It then suggests a method that allows partitioned verification even when the level of abstraction changes.

Chapter 9 summarizes the entire thesis. Appendix A defines some notational conventions. Appendix B corrects a proof from Brant's report [Br73]. Appendix C presents the complete results of the symbolic execution of the first example in Chapter 6.

unknown value. Operators will take symbolic operands and return symbolic results.

The advantage of symbolic execution is that it tests the program for all possible values of the symbolic variables, while a normal execution tests or specific (hopefully "representative") cases.

## Chapter 2

### Previous Work

Section 2.1 reviews King's work with symbolic execution. Section 2.2 discusses previous work in hardware design verification, including but not limited to the application of symbolic execution. Section 2.3 identifies an aspect of Cheatham's work that could be applied to symbolic execution. Section 2.4 reviews Milner's and Brand's criteria for consistency of programs. Section 2.5 then identifies remaining problems that are addressed in this thesis.

#### 2.1. Symbolic Execution for Software Testing

James King developed symbolic execution, a technique for testing programs [K775, K776a]. He implemented the technique in HONEY, a symbolic execution system for a subset of PL/I [H75].

At any given instant during the normal execution of a program, the program's variables have known, constant values. Arithmetic and Boolean operators will take constant operands and return constant results.

In symbolic execution, by contrast, the program's variables have symbolic values. For example, the initial value of program variable *x* might be denoted by the symbolic variable "%x". A symbolic variable represents a constant but

Figure 2-1 shows a simple example contrasting the two types of execution.

Figure 2-1a gives three assignments which exchange the values of two distinct integer variables. Figure 2-1b shows a sample normal execution. This test verifies one of the possible executions of this code; the assignments correctly exchange the values 24 and 101.

```
x := x - y;
y := x + y;
x := y - x;
```

(a) Code to exchange two integers

```
initial           x = 24          y = 101
after x := x - y; x = -77        y = 101
after y := x + y; x = -77        y = 24
after x := y - x; x = 101        y = 24
```

(b) Normal execution

```
initial           x = %x          y = %y
after x := x - y; x = %x - %y   y = %y
after y := x + y; x = %x - %y   y = %x
after x := y - x; x = %y        y = %x
```

(c) Symbolic execution

Figure 2-1: Symbolic execution example

Figure 2-1c shows a symbolic execution of the three assignments. The initial values of *x* and *y* are the symbolic variables "%x" and "%y". The final result shows that this code will exchange the values of *x* and *y*, for all possible values of "%x" and "%y".

This basic approach is complicated by the presence of conditional branching and loops in most programs. When a conditional branch is encountered, the single symbolic execution may split into several executions, one for each path at the conditional branch, to test all possible execution sequences for the program.

Symbolic execution may also be applied to formal software verification. The standard verification technique, the inductive assertion method introduced by Floyd [Fl67, HK76, Ho69, LG79], requires the generation and proof of "verification conditions." James King discussed four methods for generating the verification conditions [Ki76b]. Two of the methods, symbolic execution and computation of weakest precondition, yield the least complex results. Symbolic execution is somewhat more cumbersome to implement than weakest precondition, but symbolic execution is attractive because of its suitability for informal testing during program development, as discussed above.

Chapter 3 reviews symbolic execution in greater detail, giving particular attention to the handling of conditional branches.

## 2.2. Functional Verification of Hardware Designs

### 2.2.1. Comparison of Specification and Design

John Deninger noted that software verification techniques should be applicable to computer hardware description languages for the functional verification of hardware designs [Da79]. Efforts have been made to apply the inductive assertion method and symbolic execution to the comparison of two hardware descriptions.

Systems developed by Patterson (*STRUM*, [Pa76]) and Pitchumani and Stabler [PS82] used the inductive assertion method for hardware verification. Patterson's system was tailored for the verification of microcode specified in a high-level notation. Pitchumani and Stabler's system was tailored for the verification of designs expressed at the register transfer level (*RTL*). In these systems, the input program represented a proposed design, while assertions embedded in this program represented the specification. The verification procedure demonstrated that the program was consistent with its assertions.

This approach uses different notations for the specification (assertions) and the implementation (ordinary program statements). By contrast, the High Level Design and Verification System (*HLDVS*) [EG77, Ol78] designed by Ofek et al. uses a single notation for the specification and implementation. This is desirable for supporting a hierarchical design methodology. In a top-down design, for example, this would allow the implementation in one cycle of the design to become the specification in the next cycle.

The notation used by *HLDVS* is the non-procedural *RTL language*. *ICD* (Language for Computer Design) [EG76]. An *ICD* file may contain several alternate descriptions (models) of a single system or component. The verification procedure guarantees that two models are equivalent. *HLDVS* uses symbolic execution to carry out the verification: Symbolically execute both models and compare the symbolic results.

When using the same notation for specification and implementation, a more typical approach is to completely separate the two descriptions. W.C. Carter et

at [CE77, CJ78, JC76, LJ77] and Steve Crucher et al. [CM80, va79, vM78] each use this approach for the verification of microcode in microprogrammed machines. One description gives the specification, describing the desired architecture of the computer being designed. The second description, using the same notation, describes the proposed host machine. (The two systems also require a third description specifying the microcode; this could be considered as an initialization section for the second description.)

The separation of the two descriptions poses a new problem. In the ICD file, alternate models would manipulate the same variables, so the correspondence between the two models was self-evident. However, when the alternate models are in separate files, each has its own set of variables. This is more flexible, allowing the same data to be represented in different ways, at different levels of abstraction, in alternate models. But it also obscures the correspondence between the several models.

To solve this problem, a third input, the "simulation relation," specifies the correspondence between a specification and the implementation. Formally, a simulation relation is a set of program state pairs  $\langle s_g, s_i \rangle$ , where  $s_g(s_i)$  is an assignment of values to variables (including program counter and the system stack) for the specification (implementation). If  $\langle s_g, s_i \rangle$  is in the simulation relation, then  $s_g$  and  $s_i$  are corresponding states.

Given a specification, proposed implementation, and simulation relation, the verification procedure then is as follows:

1. Start symbolic executions of both descriptions. The starting points are a pair of corresponding states given in the simulation relation.

2. Continue each simulation until it reaches the next state which appears in the simulation relation.

3. Check. The two descriptions should again be in a pair of corresponding states as defined in the simulation relation.

4. Repeat the above test for all pairs of states in the simulation relation.

Crocker's Microcode Verification System uses ISPS [BB79] for the descriptions of the target machine, host machine, and microcode. Carter's Microprogram Certification System uses LSS (Language for Symbolic Simulation [CE77]) for these descriptions. Carter notes that this system could be applied to different problems, e.g. the verification of machine-code implementations of algorithms.

### 2.2.2. Comparison of Manual and Automated Designs

An alternate verification technique is to generate a correct design from a specification automatically. The resulting design is by no means guaranteed to be optimal. However, the generated design should be functionally equivalent to a better design created manually. A proof that a manual design is equivalent to the automated design shows that the manual design is functionally identical to the original specification. J.P. Roth at the IBM T.J. Watson Research Center and N. Kawato et al. at Fujitsu Laboratories in Japan use this approach.

Roth at IBM requires that the input hardware description be a regular algorithm to facilitate the automatic design generation. This is guaranteed by the syntax of the input language PL/R, which is approximately a subset of PL/I. A program KTRAN translates this description to a hardware design in a particular technology [HR74].

A designer may create his or her own design which is superior to the automatically generated design. A one-to-one correspondence must exist between the inputs, outputs, and states (registers) of the two designs; i.e., the designs must be isomorphic. If this is the case, then a program VERIFY will compare the designs [Ro73, Ro75, Ro77]. VERIFY will find that the two designs are equivalent if their state transition diagrams are identical and if their outputs are identical for all states and inputs.

Under this restricted definition of equivalence, the verification procedure reduces to a comparison of acyclic combinational circuitry. This comparison is np-hard. VERIFY uses a heuristic, the Consistency or line justification step of the D-Algorithm [BF76, Ro66], to perform the comparison. The heuristic attempts to take advantage of the structure of the combinational circuitry to make the comparison more efficient.

Kawai et al. at Fujitsu extended DFK [DD75] for use in a similar verification system for synchronous designs [KS79].

### 2.2.3. Comparison of Manual and Automated Specifications

S. Leinwand and T. Landau proposed the following approach: Given an algorithmic specification and a manual design, automatically generate a second algorithmic description from the manual design and then compare the algorithmic descriptions [LL79].

The manual design contains many implementation details. In the translation of this design, these details are removed, leaving only a description of the desired

function of the circuit; hence the term "functional abstraction" used by Leinwand. He notes that if the manual design is correct, then the behavior can always be abstracted in such a way that the resulting algorithmic description is identical to the original specification.

The functional abstraction proceeds in incremental steps. Leinwand has identified at least eight levels of abstraction from final design back to function specification:

1. Final design - basic components and interconnections.
2. Semic design with feedback cutpoints identified.
3. Boolean expressions for all inputs, outputs, and feedback cutpoints.
4. Guarded commands, / trigger / signal < expression, where the trigger may be sensitive to signal transitions.
5. Guarded commands with signal vectors identified.
6. Nonprocedural KTL (register transfer language).
7. Procedural KTL.
8. Abstract system specification (system responses to valid inputs).

Leinwand has been successful in abstracting behavior to the level of simple KTL.

### 2.2.4. Conclusion

The comparison of a manual design with an automated design naturally requires the automatic generation of a correct design. This technique is heavily technology-dependent, requiring a separate compiler for each target technology. Furthermore, the approach still requires the comparison of two different descriptions. Roth and Kuwayai handled this problem by requiring that the manual

design and automated design be isomorphic, reducing the problem to the comparison of combinational logic. This requirement prevents many optimizations, including simplifications due to don't-care conditions.

Functional abstraction requires the ability to distinguish implementation details from functional behavior, to allow the elimination of the irrelevant details.

Functional abstraction also imposes a particular hierarchy of levels of abstraction.

The disadvantages of these approaches justify continued interest in the direct comparison of a proposed implementation against the specification.

### 2.3. Program Analysis with Conditional Expressions

Cheatam et al. sought to derive by static analysis a compact closed-form symbolic expression of the behavior of a program [CH79]. One problem is that the program may contain many alternate paths which must be combined to obtain the single result. Cheatam suggested the use of conditional expressions for this purpose.

For example, consider the statement

***If C1 then X := C2 else X := C3;***

Rather than giving two sets of program results, one each for C1 TRUE or FALSE, we can generate the single conditional value

***X = If C1 then C2 else C3***

If X is Boolean, this can be written

***X = C1 and C2 or not C1 and C3***

The generation of the conditional expressions is directed by a flow graph similar to that used by Aho and Ullman for dataflow analysis [AU77].

17

18

Although Cheatam's work was not directed toward symbolic execution or functional verification, it is mentioned here because of its applicability to symbolic execution. The use of conditional expressions will permit the combination of multiple paths that arise when a symbolic execution cannot uniquely resolve a conditional branch.

### 2.4. Criteria for Consistency

Of course, the comparison of two descriptions to establish their consistency requires a criterion for consistency. In Section 2.2.1, this problem was simplified because both descriptions were deterministic with no don't-care conditions. Allowing nondeterministic descriptions, don't-cares, or loops complicates the issue.

Milner [Mi71] and Brund [Br78] have proposed criteria for consistency between programs. Both authors use graph models to represent programs. A vertex represents a program state, while an edge represents the change from one state to another effected by execution of the program. Two graphs  $P_s = \langle D_s, F_s \rangle$  and  $P_i = \langle D_i, F_i \rangle$  represent the specification and implementation programs to be compared. The simulation relation is  $R_{si} \subseteq D_s \times D_i$ .

Milner requires  $P_s$  and  $P_i$  to be deterministic programs with no don't-cares. That is,  $F_s(F_i)$  is a function defined for all states in  $D_s(D_i)$ . In this case, the criterion for consistency is

$$R_{si} F_i \subseteq F_s R_{si}$$

(Appendix A describes notational conventions for relations.) For deterministic programs with no don't-cares, this definition is equivalent to the test procedure described in Section 2.2.1.

Milner has also demonstrated the transitivity of this definition, given suitable restrictions on the simulation relation: If  $P_2$  is consistent with  $P_1$  and  $P_3$  is consistent with  $P_2$ , then  $P_3$  is consistent with  $P_1$ . This feature would be important for hardware verification in a hierarchical design, where  $P_2$  and  $P_3$  would represent successive refinements in the design of  $P_1$ .

Brand assumes  $P_s$  and  $P_i$  are nondeterministic models with no don't-cares. A simplified version of his criterion for consistency is

$$R_{si} F_i^* \subseteq F_s^* R_{is}$$

The presence of closure in this definition facilitates the use of symbolic execution for verification in cases where one program uses a loop while the other one does not.

Brand demonstrated how the verification of  $P_i$  against  $P_s$  by  $R_{si}$  can be partitioned. Suppose  $P_s$  consists of two concurrent processes  $P_s^1 = \langle D_s, F_s^1 \rangle$  and  $P_s^2 = \langle D_s, F_s^2 \rangle$ ;  $P_s = \langle D_s, (F_s^1 \cup F_s^2) \rangle$ . Similarly  $P_i$  has two concurrent processes  $P_i^1$  and  $P_i^2$ . Brand showed that under certain restrictions, we can verify the consistency of  $P_i^1$  with  $P_s^1$  and the consistency of  $P_i^2$  with  $P_s^2$ , separately; the verification of each component individually then implies the consistency of  $P_i$  with  $P_s$ .

Chapters 7 and 8 review Milner's and Brand's work in more detail.

## 2.5. Remaining Problems

### 2.5.1. Symbolic Execution

None of the previous work with symbolic execution applied it to event-driven simulation. Event-driven simulation is attractive for hardware simulation for several reasons:

- It supports the partitioning of a system into independent components which communicate with each other by means of triggering events.
- It is useful over a range of levels of abstraction, from the architectural level down to the gate level.
- It supports both synchronous and asynchronous simulation.

Event-driven simulation is not peculiar to a particular implementation, design, or language style. For example, event-driven simulations could be defined for either procedural or non-procedural languages.

An attempt to symbolically execute an event-driven simulation may lead to an excessive number of conditional branches due to the event-detection mechanism. The standard treatment of conditional branches splits a single execution into multiple executions; this results in a high number of separate executions.

It is desirable to avoid splitting a single execution into separate executions if possible. In some cases, it is correct to do this by simplifying the event-detection mechanism, substituting non-selective trace for selective trace (Chapter 4). In other cases, Cheatham's approach may be useful.

Cheatham was able to analyze programs in textual order (top to bottom), but

this is not adequate for the symbolic execution of event-driven simulations. This application requires analysis of the flow graph to organize the execution of the program (Chapter 5). A second analysis uncovers loop structures which were obvious in Chaitin's language.

The use of conditional expressions is likely to increase the complexity of the symbolic expressions generated in a symbolic execution. The added cost of analyzing the more complex expressions can offset the reduction in the number of executions. A consideration of this tradeoff is desirable (Chapter 6).

### 2.5.2. Consistency

Milner's and Brand's criteria for consistency both apply only to program models not allowing don't-care conditions. The addition of don't-cares changes some considerations. A single criterion that applies to both deterministic and nondeterministic program models, with and without don't-cares, is desirable; properties which could be proven about this criterion would then apply to the several model types.

The presence of closure in Brand's definition facilitates verification in some cases. Unfortunately, it also allows the verification of incorrect designs. Closure must be added in a limited way to preserve intuitive notions of correctness while still allowing verification of programs with loops.

Milner's result concerning the transitivity of program consistency assumes strict partitions on the vertex (data) space and restrictions on the simulation relation. These assumptions, reasonable for the terminating programs with which Milner was

concerned, are inappropriate for programs that represent hardware specifications. A more general result is required to facilitate hierarchical verification. Chapter 7 addresses each of these issues.

Finally, Brand's result on partitioned verification includes restrictions on accesses to variables by more than one component. This in turn implies that  $P_1$  and  $P_2$  must use similar timing protocols. This prevents partitioned verification when  $P_1$  introduces a new level of abstraction with new timing details. It is desirable to allow partitioned verification even in this case (Chapter 8).

bar "—" at the right margin. This bar will be preceded by the figure number (with dash), section number (with period), or page number (with "pg") against which the current figure should be compared. Sometimes, lines with small changes will not be marked to focus attention on the more significant differences.

## Chapter 3

### Symbolic Event-driven Simulation

**Verification of design correctness by exhaustive simulation is almost always intractable, even for simple designs.** Symbolic simulation addresses this problem by testing many, different cases simultaneously. (Symbolic simulation is the symbolic execution of a hardware description.) Each distinct case is characterized by a unique assignment of values to the symbols used in the simulation.

This chapter reviews symbolic execution. A more complete discussion was written by Hantler and King [HK76]. This chapter also describes an event driven simulation model to which symbolic execution can be applied.

Section 3.1.1 reviews the symbolic computation of expressions. Section 3.1.2 describes case splitting and presents a suitable set of functions for implementing a case splitting execution. Section 3.2 defines a simple model representing event-driven simulation. Finally, Section 3.3 discusses the drawbacks of symbolically executing an event-driven simulation.

In this chapter and the next two chapters, we will often derive new algorithms by modifying existing algorithms. To facilitate the comparison of the new and old versions, we will adopt the convention of marking altered or new lines with vertical "

In the following discussion, the term "conventional execution" will refer to an execution in which all variables have specific constant (not symbolic) values.

#### 3.1. Review of Symbolic Execution

##### 3.1.1. Symbolic Computations

Denote by  $S_c$  the Cartesian product of the domains of the inputs provided during the course of a given execution of a program. For example, the domain for one scalar subrange input in a Pascal program may be  $\{0, 1, 2, 3, 4\}$ . In a conventional execution, each input is assigned a known constant value from its domain.

At any given time during a conventional execution, the program inputs and variables have *known constant values*. The program's operations take constant arguments and yield constant results. Each execution tests the program only for the particular values of the program inputs supplied; only for a single value of  $s \in S_c$ .

The motivation for symbolic execution is that each execution should test the program for many different values of the program inputs; ideally for all possible values. In a symbolic execution, some or all of the inputs may be represented by symbolic variables. In our discussion, symbolic variable names will always start with "%". A symbolic variable has a constant but unknown value.

Let subset  $S_s \subseteq S_t$  be the Cartesian product of the value ranges represented by the inputs in a symbolic execution. If an input  $x$  is a constant  $c$ , then the value range of  $x$  is just  $\{c\}$ . However, if  $x$  is assigned a symbolic variable  $\$y$ , then the value range of  $x$  is the entire domain of  $x$ .

In a symbolic execution, program inputs and program variables have symbolic values. That is, if a variable has domain  $R$  in a conventional execution, then its value in a symbolic execution is a function mapping  $S_s$  to  $R$ . The symbolic variables or constants assigned to the inputs are simple examples of such functions. The program's operations take symbolic arguments and yield symbolic results. A single execution tests the program for all values of  $s \in S_s$ .

For example, consider the statement " $x := y + z$ ". In a conventional execution,  $y$  and  $z$  may have values 5 and 7. The statement computes the result 12 and assigns it to  $x$ . This execution covers only the single case  $y = 5$  and  $z = 7$ .

In a symbolic execution, by contrast,  $y$  and  $z$  may have symbolic values  $\$a$  and  $\$b$ , where  $\$a$  and  $\$b$  are the symbolic inputs to the program. The statement computes the symbolic result  $\$a + \$b$  and assigns it to  $x$ . The execution covers all possible values of  $\$a$  and  $\$b$ .

Then  $S_s$  is the domain of inputs over which this execution tests the program. (Actually, this domain is constricted by a path condition, introduced below.) All symbolic values computed have domain  $S_s$ . Evaluation of the symbolic values for a particular  $s \in S_s$  yields the results of a conventional execution of the program using constant inputs  $s$ .

In a discussion of symbolic execution, the reader could interpret an expression in at least two ways:

- An expression to be evaluated as in a conventional execution, yielding a particular constant value.
- An expression to be evaluated as in a symbolic execution, yielding a symbolic value.

If the expression is Boolean, a third interpretation is possible:

- An assertion about the program, true for some or all  $s \in S_s$ .
- When the appropriate interpretation may be in doubt, we will identify expressions of the above types as being constant-valued expressions, symbolic expressions, or assertions, respectively.

In the following discussion, the first character of some symbolic execution functions and variables will be "/". This will avoid naming conflicts with user-declared identifiers.

### 3.1.2. Case Splitting

The fundamental problem with symbolic execution is its inability to resolve conditional branches, where a branch is taken only if a specified Boolean expression *Cond* is true. This expression has a symbolic value which may not simplify to *TRUE* or *FALSE*. In this case, the program cannot uniquely select a path to follow; for some values of  $s \in S_s$ , a conventional execution would follow the true branch, while for other values of  $s$ , the program would follow the false branch.

The traditional solution is to choose (interactively or automatically) either the

true or false branch, maintaining a *path condition* which gives the condition under which the assumed choice was correct. When the program reaches the end of one path, we may explore the other branch separately. Once a path splits at an unresolved conditional branch, the two paths do not rejoin. The program follows each path separately until the end of execution, and then analyzes the results for each path separately. We must eventually explore all possible choices.

This technique has been called *forking* or *case analysis* [K176a]. The following discussion will refer to this technique as *case splitting*. This emphasizes the fact that a case analysis splits a single program execution into several executions, one for each path. In the worst case, the total number of paths in a case splitting execution may increase exponentially as the number of conditional branch points in the program, the addition of each if statement potentially doubling the total number of paths.

Case splits may arise in other situations. For example, to simplify the internal representation of data structures, a symbolic program may require that all array subscripts evaluate to constants [K175]. If some subscript does not evaluate to a constant, the program chooses a constant value, again updating the path condition to reflect the choice made.

We now introduce one possible set of functions to support symbolic execution with case splitting. The program uses functions *!Decide* and *!Const* to select a path or constant value, introducing a case split if necessary. Function *!Decide* examines a symbolic Boolean expression, returning *TRUE* or *FALSE*. Function *!Const* examines a symbolic scalar expression, returning a constant scalar value. The remainder of this section describes path-condition *!Pc*, *!Decide*, and *!Const*, also introducing several auxiliary functions.

1. **!Pc:** Path condition *!Pc* is a Boolean function with domain  $S_s$ . We define  $\text{!Pc}(s) = \text{TRUE}$  iff the current symbolic execution has made the same conditional branches and assumed the same constant values as would a conventional execution with inputs  $s$ . Function *!Pc* then identifies the subset of  $S_s$  tested by this symbolic execution.

The remaining functions all take one argument  $f: S_s \rightarrow R$ , a symbolic expression, where  $R$  is any valid scalar type.

2. **!ConstP:** Predicate *!ConstP* is a Boolean function. It returns *TRUE* iff it recognizes argument  $f$  as being constant over the domain for which this symbolic execution is valid. That is, *!ConstP* returns *TRUE* iff it can find constant  $r \in R$  such that  $\forall s \in S_s : (\text{!Pc}(s) \supseteq f(s) = r)$ . The program will use *!ConstP* to test the symbolic value of an expression which must evaluate to a constant, e.g. a subscript. No case split is necessary if *!ConstP* is true. Function *!ConstP* can vary in sophistication. A simple version may simply check whether  $f$  is constant. A more complex version may attempt to determine that  $f$  is equivalent to a constant, or is constant for all values of  $s$  for which *!Pc*( $s$ ) is true.
3. **!Value:** If *!ConstP*( $f$ ) is true, then function *!Value* returns the value  $r \in R$  such that  $\forall s \in S_s : (\text{!Pc}(s) \supseteq f(s) = r)$ . Otherwise, *!Value* is undefined.
4. **!Choose:** Function *!Choose* returns a value  $r \in R$  such that  $(\text{!Pc} \text{ and } (f = r))$  is satisfiable; i.e., such that  $\exists s \in S_s : (\text{!Pc}(s) \wedge f(s) = r)$ . The program calls *!Choose* when *!ConstP*( $f$ ) is false, indicating the need for a case split. Function *!Choose* then does the following:

1. Save the current state of the program. This will allow us to return to this point later and select another path.

2. Choose the value of  $f$  to be assumed in this case. Function  $\text{!Choose}$  may request that the user make this choice.

3. Augment the path condition function  $\text{!Pc}$  to further constrain the subset of  $S_s$  for which the symbolic execution is now valid:

a. If  $f$  is Boolean and  $\text{!Choose}$  returns  $\text{TRUE}$ , then set

$$\text{!Pc} := \text{!Pc} \text{ and } f$$

b. If  $f$  is Boolean and  $\text{!Choose}$  returns  $\text{FALSE}$ , then set

$$\text{!Pc} := \text{!Pc} \text{ and not } f$$

c. Otherwise,  $f$  is not Boolean and  $\text{!Choose}$  returns  $r$ ; then set

$$\text{!Pc} := \text{!Pc} \text{ and } (f = r)$$

5.  **$\text{!Const}$ :** Function  $\text{!Const}$  returns the constant value to be assumed for an expression, initiating a case split and altering  $\text{!Pc}$  if necessary. Its definition follows:

```
If  $\text{!Const}(f)$ 
then  $\text{!Const} := \text{!Value}(f)$ 
else  $\text{!Const} := \text{!Choose}(f)$ 
```

For the remaining functions, the symbolic argument  $f$  is Boolean.

6.  **$\text{!IdTrue}$ :** Boolean function  $\text{!IdTrue}$  returns  $\text{TRUE}$  iff its argument  $f$  is identically true; that is,  $\forall s \in S_f : f(s)$ .

7.  **$\text{!Decide}$ :** Boolean function  $\text{!Decide}$  returns the Boolean constant value to be assumed for a conditional expression, initiating a case split and altering  $\text{!Pc}$  if necessary. Its definition follows:

29

30

```
if  $\text{!IdTrue}(\text{!Pc} \supseteq f)$ 
then  $\text{!Decide} := \text{TRUE}$ 
else if  $\text{!IdTrue}(\text{!Pc} \supseteq \text{not } f)$ 
then  $\text{!Decide} := \text{FALSE}$ 
else  $\text{!Decide} := \text{!Choose}(f)$ 
```

This definition corresponds to the use of  $\text{!Const}$  with a smart version of  $\text{!ConstP}$ .

The appeal of symbolic execution with case splitting is its ease of implementation. Given a conventional program, symbolic data types, and implementations of  $\text{!Const}$  and  $\text{!Decide}$ , the following modifications to the program allow a symbolic execution with case splitting:

1. Let user-declared variables have symbolic values rather than constant values.
2. Extend operators to allow symbolic computations as well as constant-valued computations. The form of an operation (symbolic or constant-valued) depends on the form of its operands.
3. Declare global variable  $\text{!Pc}$  which has a symbolic value, initially  $\text{TRUE}$ .
4. Replace subscripts and conditional expressions  $E_{\lambda pr}$  by  $\text{!Decide}(Expr)$  or  $\text{!Const}(Expr)$ .

As an example, Figure 3-1 shows the modification of a few Pascal control constructs to carry out symbolic execution with case splitting. The change is trivial for most cases. The change is nontrivial for the **for** statement because the conditional expression computation is implicit, not explicit as in the other cases.

<pre> <b>Normal execution</b>           Case splitting version <b>if Cond</b>                   <b>if !Decide(Cond)</b>     <b>then Statement1</b>          <b>then Statement1</b>     <b>else Statement2</b>          <b>else Statement2</b> <b>case Selector of</b>          <b>case !Cons(Selector)</b>     <b>CaseBody</b>                <b>CaseBody</b> <b>while Cond do Statement1</b>   <b>while !Decide(CS!Cond)</b>                             <b>do CS!Statement1</b> </pre>	<pre> <b>for VarId := Lo to Hi</b> <b>do Statement1</b> <b>begin</b> <b>VarId := Lo;</b> <b>!Boundn := Hi;</b> <b>if !Decide(VarId &lt;= !Boundn)</b>     <b>then goto Lj;</b>     <b>else goto Lk;</b> <b>Lj: VarId := Succ(VarId);</b> <b>Lj: Statement1;</b> <b>if !Decide(VarId &lt; !Boundn)</b>     <b>then goto Lj;</b> <b>Lk: VarId := %Undefined</b> <b>end</b> </pre>
--	---

$X := A[Expr]$

$X := A[!Cons(Expr)]$

**Figure 3-1:** Some control constructs for case splitting execution

combined [Hv79, Hv77]) with unit delays rather than variable, real (not integer) delays. Our model will avoid the SAMI: concepts of component type instantiations, subprocesses, and the event queue. The model also provides no mechanisms for driving, monitoring, or halting the simulation. These features, desirable or essential in an actual implementation, would only complicate the discussion without adding any insight here.

Figure 3-2 outlines an event-driven simulator in a notation similar to Pascal [JW79]. The simulation is an infinite loop with two major steps: (1) execute the processes concurrently in the cobegin statement, and (2) update array *Net* from *NewNet*. Intuitively, each process corresponds to a component in a hardware system, while *Net* and *NewNet* correspond to interconnections between the components. Array *WasUpdated*, not read in Figure 3-2, is used in the next chapter.

Processes, functions, and procedures may read only *Net* and *NetChanged* of the global variables. (Hereinafter, "procedures" will include functions.) They may write only *NewNet* and *Updated*, and only by use of the assign statement. The statement

assign *Expression* to *j*

does the following:

```

begin NewNet[j] := Expression;
  Updated[j] := TRUE; end

```

Processes and procedures may not pass globals as actual parameters to procedures.

Processes are like procedures, with the same restrictions, except for two differences. First, the process name is a number, not an alphanumeric identifier. Second, a process may use a pause statement, which is not available to procedures.

### 3.2. Model for Event-driven Simulation

The preceding section has reviewed symbolic execution for programs in general. We are concerned in particular with the symbolic execution of hardware descriptions using event-driven simulation. This section describes a model for event-driven simulations.

The model is similar to that used by SAMI: [Hi80], but it is much simpler. It is roughly equivalent to a SAMI: description (and its description and sim. description



Now impose the additional restriction that associated with each element  $NewNet[j]$  is a single process  $j$ , which is the only process ever permitted to write  $NewNet[j]$ . This implies that the final value of  $NewNet$  is also independent of the interleaving of the processes. (As a real example, the definition of  $SAT1$  specifies that the result of a simulation is undefined unless this last property holds [H180].)

We may then equivalently define the simulation using any interleaving we wish; in particular, an interleaving in which the process executions do not overlap. Such a definition appears in Figure 3-4.

The event-driven simulator now behaves not as a concurrent, nondeterministic program, but as a sequential, deterministic program. The symbolic execution techniques developed for sequential programs may be applied to this model for event-driven simulation.

### 3.3. Event-driven Simulation With Case Splitting

Case splitting tends to generate an excessive number of case splits for event-driven simulations. Typically, when a process is executed, it will not perform any computation unless the value of some input  $Net[j]$  has changed; this change in value is the driving event. For example, the case splitting version of a simple process may have the following form:

```
while TRUE do begin
  repeat pause until !Decide(NetChanged[j]);
  DoComputation
end
```

The determination of whether or not  $Net[j]$  has changed is a conditional branch which  $!Decide$  may not be able to resolve. Hence, each execution of a process that is waiting for a change in a net is likely to introduce a case split.

```
program EventDrivenSimulation;
var Net, NewNet: array [1..NumOfNets] of NetType;
  NetChanged, Updated, WasUpdated: array [1..NumOfNets] of Boolean;
  i: Integer;
```

<function and procedure definitions>  
<process definitions>

```
begin (* Main loop *)
  initialize NewNet;
  Net := NewNet;
  for i := 1 to NumOfNets do NetChanged[i] := FALSE;
  Updated := NetChanged;
  WasUpdated := Updated;
  initialize processes;
  while TRUE do begin
    cbegin resume process i coend;
    for i := 1 to NumOfProcs do
      if Updated[i]
        then begin
          NetChanged[i] := NewNet[i] <> Net[i];
          Net[i] := NewNet[i];
          Updated[i] := FALSE
        end
      else NetChanged[i] := FALSE
    end.
  end.
end.
```

Figure 3-4: Modified equivalent event-driven simulator

Moreover, in the worst case, the total number of paths increases as the product of the numbers of paths in each process. Consider a simulation with just two processes. Suppose that processes 1 and 2 split into  $M$  and  $N$  paths, respectively, during each execution.

After the first execution of process 1, the simulation has been split into  $M$  paths. For each of the  $M$  paths, the subsequent execution of process 2 will split the single path into  $N$  paths, for a grand total of  $MN$  paths after the first cycle. After  $i$  cycles, the single simulation will have split into  $(MN)^i$  separate simulations.

If each process splits into only two paths, e.g., when it has the simple form shown above, then after just four cycles the simulation splits into 256 separate paths. To exhaustively test this simulation over four cycles using cases splitting symbolic execution, we must save the complete state of the simulation 255 times and carry out 256 separate simulations.

Such performance makes symbolic execution with case splitting unsatisfactory for many event-driven simulations. The following chapters will consider alternatives to case splitting for symbolic execution.

## Chapter 4

### Non-selective Trace

The property that a process will not perform any computations unless some net  $Net[Y]$  receives a new value different from its old value is called *selective trace* [U16]. The selective trace often introduces an excessive number of case splits. This chapter shows an example in which case splitting can be reduced by the use of a non-selective trace. If a process uses non-selective trace, it will perform its computations whenever the net receives a new value, whether or not the value is different from the old value.

In the event-driven simulation model of Figure 3-4, the selective trace is typically implemented by

```
repeat pause until NetChanged[] or NetUpdated[k]
```

To use non-selective trace, we replace the references to  $Net[Changd]$  by references to  $WasUpdated$ , yielding the following:

```
repeat pause until WasUpdated[] or WasUpdated[k]
```

Section 4.1 justifies the correctness of this approach for combinational processes.

Section 4.2 considers the possibility of applying this approach to a more general process type. A brief outline is as follows:

- Section 4.2.1 describes the general process type of interest.

- Section 4.2.2 defines terminology for the discussion.
- Section 4.2.3 introduces some restrictions and derives important properties of the general process type.

- Section 4.2.4 compares the process using selective trace against the process using non-selective and derives a test to establish whether the latter may be substituted for the former.

- Section 4.2.5 illustrates how assertions and the inductive assertion verification technique help in applying the test from Section 4.2.4.
- Finally, Section 4.2.6 defines a more restrictive general process type and shows how it might be tested without the use of auxiliary assertions.

#### 4.1. Combinational Process

Figure 4-1 defines a NOR-gate process.  $Net[A]$  and  $Net[B]$  are the input nets;  $Net[Y]$  is the output net, written indirectly by assign. The nets are all assumed to be Boolean.

```
process 101; (* Nor *)
const A = 201; (* indices to Net *)
B = 202; (* index to NewNet *)
Y = 203; (* index to NewNet *)
begin
while TRUE do begin
    assign not (Net[A] or Net[B]) to Y;
    repeat pause until Net[Changd[A]] or Net[Changd[B]]
end
end; (* 101 - Nor *)
```

Figure 4-1: NOR-gate process

The main body of the process consists of a tight infinite loop which specifies

the behavior of a NOR-gate. This construction is typical of process definitions for combinational networks. Consider first a conventional simulation in which the nets are all constant-valued. The assign statement sets the output net  $Net[Y]$  to  $\text{not}(Net[A] \text{ or } Net[B])$ . The pause loop waits for a change in the values of the inputs. As long as no change occurs, the value of  $Net[Y]$  does not need to be altered, and the process need not perform any computation. As soon as either  $Net[A]$  or  $Net[B]$  changes its value, the process performs a new computation, returning to the assign statement to update  $Net[Y]$ . Then the process will again wait in the pause loop. Recall that  $Net[Y]$  is driven only by this process.

To modify this process for symbolic simulation with case splitting, the operator and variable types are extended to allow symbolic computations, and the condition in the pause loop becomes

$\text{!Decide}(NetChanged[A] \text{ or } NetChanged[B])$

The call to  $\text{!Decide}$  is required because array  $NetChanged$  has symbolic values. Function  $\text{!Decide}$  will introduce a case split if it cannot uniquely determine whether or not the values of nets  $A$  or  $B$  have changed.

This process uses selective trace. However, in combinational processes of this form, the use of a non-selective trace will still yield correct results. For example, consider an alternate conditional expression for the pause loop:

$(WasUpdated[A] \text{ or } WasUpdated[B])$

This expression replaces symbolic  $NetChanged$  by the constant-valued  $WasUpdated$ , so the call to  $\text{!Decide}$  is not necessary. Suppose that at some point in a symbolic simulation,  $Net[A]$  has a stable symbolic value  $\%ExprA0$ , while  $Net[B]$  has the value  $\%ExprB0$ . Then  $Net[Y]$  has the symbolic value  $\text{not}(\%ExprA0 \text{ or } \%ExprB0)$ .

due to a previous assignment by the process. When a new symbolic value  $\%ExprA1$  is assigned to  $Net[A]$ , the process using the non-selective trace leaves the pause loop and assigns  $\text{not}(\%ExprA1 \text{ or } \%ExprB0)$  to  $Net[Y]$ , whether or not  $\%ExprA1$  is equal to  $\%ExprA0$ .

For the symbolic simulation to be correct, it must predict the correct value of  $Net[Y]$  for all possible conventional simulations. If for some test case,  $\%ExprA1$  does not equal  $\%ExprA0$ , then the symbolic simulation took the correct action for that case and the result will be correct. If for some other test case,  $\%ExprA1$  equals  $\%ExprA0$ , then it is also true that

$$\forall s \in S_x : \\ (\text{not}(\%ExprA0 \text{ or } \%ExprB0))(s) \\ = (\text{not}(\text{symbolic expression})(s))$$

(The notation “(symbolic expression)(s)” denotes the constant value resulting from the evaluation of the symbolic expression for  $s \in S_x$ .) The symbolic simulation predicts no change in the value of  $Net[Y]$  for this case. This is also correct. The fact that the symbolic simulator executes an assign when a conventional simulator might not execute any assign does not result in any error in this example.

The non-selective trace is preferable to case splitting. With case splitting, the process must call  $\text{!Decide}$  to determine whether or not  $\%ExprA1$  is different from  $\%ExprA0$ , possibly introducing a case split where none is needed. With the non-selective trace, the simulator avoids the call to  $\text{!Decide}$ . This reduces the number of calls to  $\text{!IfTrue}$ . It also potentially reduces the total number of simulations, since a case split requires a separate simulation for each path to test all the paths.

Note that even the references to *WasUpdated* in the repeat condition are unnecessary. A weaker test is adequate. All that is required of the substitute condition *Cond* is

$$\forall s \in S; (NetChanged[A] \text{ or } NetChanged[B])[s] \supset Cond[s]$$

It would be correct to replace the repeat statement by

*repeat pause until TRUE*

or just "pause". But this completely defeats the purpose of event-driven simulation: avoiding simulation of all logic in every time step.

Conversely, a stronger condition can be used for symbolic simulation. A possible symbolic repeat condition is

$$(WasUpdated[A] \text{ or } WasUpdated[B]) \text{ and } \neg !IdTrue(\neg (NetChanged[A] \text{ or } NetChanged[B]))$$

Since *NetChanges[i]*  $\supset$  *WasUpdated[i]* during process execution, this is equivalent

$$\text{to } \neg !IdTrue(\neg (NetChanged[A] \text{ or } NetChanged[B]))$$

This condition uses the np-hard *!IdTrue* test, as the case splitting version does by calling *!Decide*, but it never introduces a case split. This more complex condition may be desirable for combinational logic in a feedback path; it could prevent continuous generation of symbolic expressions when no actual changes of net values are occurring.

#### 4.2. Applicability of Non-selective Trace

The above example showed a process in which case splitting could be completely eliminated simply by replacing selective trace by non-selective trace, with no other changes. This section characterizes some processes for which this is possible.

#### 4.2.1. Generalized Process of Interest

We will let *process /* be the process of interest. Assume that at some point in process *i*, we have the code in Figure 4-2a. The second pause only logically follows *S*. The process may actually branch to some other pause after *S*; for example, the one in the repeat loop. Statement *S* exits only through its bottom and does not include any pause.

(• Label *L/I* is a dummy for reference \*)

```
repeat pause; L/I : until Csel:
  S;
  pause;
```

(a)

```
repeat pause; L/I : until Cnow:
  S;
  pause;
```

(b)

Figure 4-2: Generalized process for non-selective trace

The subscript "sel" denotes "selective trace"; the branch condition *C<sub>sel</sub>* contains references to *NetChanged*. Denote by *C<sub>now</sub>* the expression *C<sub>sel</sub>* with all references to *NetChanged* replaced by references to *WasUpdated*. Subscript "now" means "non-selective trace." Figure 4-2b then shows the non-selective version of process *i*. Note that statement *S* may access *NetChanged*, and these accesses are not changed in the non-selective version. However, we will require that the branch condition *C<sub>now</sub>* and statement *S* never introduce case splits.

The question of interest is this: Under what circumstances are the code segments in parts (a) and (b) of Figure 4-2 equivalent? If the two versions are equivalent, then we can substitute version (b) for version (a). This eliminates case splitting in this process, since *C<sub>now</sub>* and *S* never introduce case splits.

#### 4.2.2. Program State

Let  $s$  denote a program state in a conventional execution. The program state includes the values of the global variables, the main program counter, the program stack, plus the following for each process  $i$ :

- program counter  $Lc_i$  for process  $i$ .
  - local variables  $Locs_i$ ,
  - a flag  $Run_i$  denoting whether or not the process is blocked.
- A `cobegin` sets  $Run_j$  to *TRUE*, while `pause` resets it to *FALSE*. Note the distinction between the program state space and  $S_s$ , above;  $S_s$  is the space of initial and input values covered by a symbolic execution.  $S_s$  is pertinent for symbolic executions, but now we are discussing conventional executions.

We will write " $(expression)Xs$ " to denote the value of *expression* evaluated in state  $s$ .

Any program state in which the main program counter points to the top of the `while` loop in Figure 3-4 is a "start state." Inspection of Figure 3-4 shows that in every start state  $s$  reached in an actual simulation, the following assertion holds:

$$\forall i: NetChanged[K_3] \supset WasUpdated[K_3] \quad (4.1)$$

This also always holds during the `cobegin` statement.

We now define the following functions for the non-selective process in Figure 4-2b.

- $P_i(s)$  is the start state obtained from start state  $s$  by one pass through the main loop in Figure 3-4.

**4.2.3. Characteristics of the Generalized Process**

Let  $s$  denote a program state in a conventional execution. The program state includes the values of the global variables, the main program counter, the program stack, plus the following for the next. Process  $i$  is the process of interest, Figure 4-2b.

- $P_i(s)$  is the state obtained from  $s$  by the combined actions of all remaining processes.

- $F_i(s)$  is the state obtained from  $s$  by action of statement  $S$  (plus any `else` to a `pause` following  $S$ ) in Figure 4-2b.

- $U(s)$  is the state obtained from  $s$  by the loop in Figure 3-4 that updates `Net`, `NetChanged`, and `Updated`.

- " $Var[Seq]$ " denotes the sequence  $s \in Seq : Var[s]$ .
- " $\{w_i\}$ " denotes the set of indices for those nets which may be written by process  $i$ .
- " $\{r_i\}$ " denotes the set of indices for those nets which may be read by process  $i$ .

Note that  $P(s) = U(P_s(P_i(s))) = U(P_i(P_s(s)))$ .

#### 4.2.3. Characteristics of the Generalized Process

Following the rule from Chapter 3, neither  $C_{sel}$  nor  $S$  may access `WasUpdated`. This reflects the philosophy that process  $i$  cannot detect the updating of a net by another process when that update does not change the value of the net. Of course,  $C_{non}$  violates this restriction, but we wish Figure 4-2b to be equivalent to 4-2a, so that the resulting computation will still be independent of `WasUpdated`, conforming to this philosophy.

The following assertion is true for all program states  $s$  and net numbers  $j$  during process execution:

$$C_{sel} \langle s \rangle (FALSE \setminus NetChanged[j]) \supset C_{sel} \langle s \rangle (TRUE \setminus NetChanged[j]). \quad (4.2)$$

If this did not hold, then for some  $s$  and  $j$  we would have  $C_{sel} \langle s \rangle (FALSE \setminus NetChanged[j])$  true while  $C_{sel} \langle s \rangle (TRUE \setminus NetChanged[j])$  is false. This means that for this case, the absence of a change in the value of a net would be the triggering event. We will not admit this as a reasonable triggering event.

This requirement on  $C_{sel}$  implies that for all states  $s$  arising during process execution, we have

$$C_{sel} \langle s \rangle \supset C_{non} \langle s \rangle \quad (4.3)$$

To prove this, assume it were not true. Then for some  $s$ ,  $C_{sel} \langle s \rangle$  is true but  $C_{non} \langle s \rangle$  is false. Suppose  $C_{sel}$  refers to just one element of  $NetChanged$ ,  $NetChanged[j]$ .

Then

$$C_{non} = C_{sel} (WasUpdated[j] \setminus NetChanged[j])$$

Since  $C_{non} \langle s \rangle$  is not equal to  $C_{sel} \langle s \rangle$ , we know  $WasUpdated[j|Ks]$  is not equal to  $NetChanged[j|Ks]$ . Eq. (4.1) then implies  $NetChanged[j|Ks]$  is false while  $WasUpdated[j|Ks]$  is true. Now we have

$$TRUE = C_{sel} \langle s \rangle = C_{sel} (FALSE \setminus NetChanged[j]) \setminus Ks$$

$$FALSE = C_{non} \langle s \rangle = C_{sel} (TRUE \setminus NetChanged[j]) \setminus Ks$$

But this contradicts Eq. (4.2). An inductive argument proves the contradiction when  $C_{sel}$  refers to more than one element of  $NetChanged$ .

Finally, condition  $C_{sel}$  must have no side effects. This is because the process should not change any variables until the triggering event described by  $C_{sel}$  occurs.

#### 4.2.4. Comparison of the Two Processes

Consider the conventional execution (with constant, not symbolic, values) of the two code segments in Figure 4-2 from the first pause during one `begin` statement. We wish the results to be the same. We must define precisely the way in which the results must be "the same." The following are true for all start states reachable in a conventional simulation as defined in Figure 3-4:

- $Updated[i]$  is false for all  $i$ .

- $Net = NewNet$ .

- Main program counter points to top of main loop.

- Program stack is empty (except for globals).

- $Run_j$  is false for all  $j$ .

In addition, we have the following characteristics of the simulation:

- By definition, the results of a simulation are independent of global variables  $WasUpdated$  and  $i$ . Therefore the values of these variables in start states are immaterial.

- Of all the processes' local variables, process  $i$  can modify only  $Lc_i$ ,  $Run_i$ , and  $Loc_i$ .

- Process  $i$  can modify only elements  $Net[\{w_i\}]$  of  $Net$ .

- If two simulations always generate the same values of  $Net$ , then the values of  $NetChanged$  will necessarily also be identical.

Then the only program variables we need to test in a comparison of the two processes are global variables  $Net[\{w_i\}]$  and local variables  $Lc_i$  and  $Locs_i$ .

Returning to Figure 4-2, the actions of the two processes are clearly identical if

$C_{sel}$  and  $C_{non}$  are equal. If  $C_{sel}$  and  $C_{non}$  are not equal, then Eq. (4.3) implies that  $C_{sel}$  is false while  $C_{non}$  is true. This is the only interesting case in the comparison of the two processes: The process using selective tract returns to the first pause without modifying any variables; the other process executes  $S$ .

The following criterion for the equivalence of the two code segments in Figure 4-2 summarizes the above considerations:

✓ valid start states s |  $Lc_i(s) = LI$ :

$(C_{non} \langle s \rangle \text{ and not } C_{sel} \langle s \rangle) \supset$

Now make the following observations:

1)  $P(s) = U(P_o(P_i(s)))$

2)  $NewNei[w_i]KU(P_o(P_i(s))) = NewNei[w_i]KP_o(P_i(s))$

$= NewNei[w_i]KP_i(s)$

3)  $Lc_i \langle U(P_o(P_i(s))) \rangle = Lc_i \langle P_o(P_i(s)) \rangle = Lc_i \langle P_i(s) \rangle$

4)  $Loc_i \langle U(P_o(P_i(s))) \rangle = Loc_i \langle P_o(P_i(s)) \rangle = Loc_i \langle P_i(s) \rangle$

5)  $C_{non}$  implies  $P_i(s) = F_i(s)$

Eq. (4.4) is then equivalent to Eq. (4.5):

✓ valid start states s |  $Lc_i(s) = LI$ :

$(C_{non} \langle s \rangle \text{ and not } C_{sel} \langle s \rangle) \supset$

$((NewNei[w_i])KF_i(s)) = Nei[w_i]Ks \rangle \text{ and}$

$(Lc_i \langle F_i(s) \rangle = LI) \text{ and}$

$(Loc_i \langle F_i(s) \rangle = Loc_i \langle s \rangle)$

Now, since the action of statement  $S$  must be independent of *WasUpdated*, then Eq. (4.5) simplifies to Eq. (4.6):

✓ valid start states s |  $Lc_i(s) = LI$ :

$((not C_{sel} \langle s \rangle) \supset ((NewNei[w_i])KF_i(s)) = Nei[w_i]Ks \rangle) \text{ and}$

$(Lc_i \langle F_i(s) \rangle = LI) \text{ and}$

$(Loc_i \langle F_i(s) \rangle = Loc_i \langle s \rangle) = Loc_i \langle s \rangle)$

Then the code segments in Figure 4-2 must have the form in Figure 4-3, where Eq. (4.6) implies  $(not C_{sel} \langle s \rangle) \supset C2 \langle F_i(s) \rangle$ . ( $S'$  plus the computation and branch on  $C2$  in Figure 4-3 comprise statement  $S$  in Figure 4-2.) This implication, plus the requirement that Figure 4-3 not introduce a case split even if  $C_{sel}$  is not identically *TRUE* or *FALSE*, implies that  $C2$  is *TRUE* (or may be optimized to *TRUE*).

Then the final form implied for the original code segments is as shown in Figure 4-4.

```
repeat repeat pause; LI : until Csel;  
S'  
until not C2;  
pause;
```

(a)

(b)

Figure 4-3: Modified generalized process

```
while TRUE do begin  
  repeat pause; LI : until Csel;  
  S  
end;  
repeat pause; LI : until Cnon;
```

(a)

(b)

Figure 4-4: Final generalized process

$(Loc_i \langle F_i(s) \rangle = Loc_i \langle s \rangle) = Loc_i \langle s \rangle)$

(4.5)

In conclusion, given a code segment of the form shown in Figure 4-4a satisfying the properties given in Section 4.2.3, we can replace it by the code in Figure 4-4b if Eq. (4.7) is satisfied:

$\forall$  valid start states  $s \mid Lc_i(s) = LI$ :

$$\begin{aligned} (\text{not } C_{sel}(s)) \supset ((\text{NewNet}[\{w_j\}]Kf_j(s)) = \text{Net}[\{w_j\}]Ks) \text{ and} \\ (Locs_i(F_i(s)) = Locs_i(s))) \end{aligned} \quad (4.7)$$

Eq. (4.7) is identical to Eq. (4.6) except that the reference to  $Lc_i$  is removed. This is because the control structure of the final generalized process guarantees that  $Lc_i$  will have the proper value. The new version (Figure 4-4b) introduces no case splits. The example in the next section shows that assertions may be useful for properly identifying the set of "valid start states" in Eq. (4.7).

#### 4.2.5. Use of Assertions

Now reconsider the example in Figure 4-1. This process definition is transformed in Figure 4-5 so that it conforms to the pattern in Figure 4-2a.

```
process 101; (* Nor *)
  const A = 201; (* indices to Net *)
  B = 202;
  Y = 203; (* index to NewNet *)
begin
  assert TRUE;
  assign not (Net[A] or Net[B]) to Y;
repeat pause; LI : until NetChanged[A] or NetChanged[B];
  assign not (Net[A] or Net[B]) to Y
end (* 101 - Nor *)
```

Figure 4-5: Modified NOR-gate process

In this example,  $C_{sel}$  is

51

52

(NetChanged[A] or NetChanged[B])

We wish to replace this by  $C_{non}$ .

(WasUpdated[A] or WasUpdated[B])

To justify this, we must prove that Eq. (4.7) holds for any "valid" start state  $s$  with  $Lc_i(s) = LI$ . (Note that the loop satisfies the restrictions from Section 4.2.3, including Eq. (4.2).)

Unfortunately, it is trivial to choose a start state  $s$  that does not satisfy this condition; simply set every element of arrays  $Net$  and  $NetChanged$  to FALSE. Then  $C_{sel}(s)$  is false, but

$$\begin{aligned} \text{NewNet}[Y]Kf_i(s) &= \text{not } (\text{Net}[A]Ks) \text{ or } \text{Net}[B]Ks) = 1 \\ &\neq \text{Net}[Y]Ks \end{aligned}$$

We must add assertions to the process to constrain the set of valid start states. Figure 4-6 shows a possible solution.

```
process 101; (* Nor *)
  const A = 201; (* indices to Net *)
  B = 202;
  Y = 203; (* index to NewNet *)
begin
  assert TRUE;
  assign not (Net[A] or Net[B]) to Y;
while TRUE do begin
  repeat pause;
    LI : assert (not (NetChanged[A] or NetChanged[B])) C
      (Net[Y] = not (Net[A] or Net[B]))
    until NetChanged[A] or NetChanged[B];
    assign not (Net[A] or Net[B]) to Y
  end
end. (* 101 - Nor *)
```

Figure 4-6: NOR-gate process with assertions

With assertions, the process in Figure 4-6 clearly satisfies Eq. (4.7). The proof

that  $C_{sel}$  can be replaced by  $C_{non}$  has become a conventional software verification problem; that is, the proof of the verification conditions implied by the assertions.

It is desirable to perform the verification condition generation (VCG) for this process alone, without also having to analyze the other processes and the main program (Figure 3-4). To do this, define pause by Figure 4-7 during VCG. For example, if we use symbolic execution for VCG, we symbolically execute the code in Figure 4-7 for each pause, introducing a new, unique symbolic variable upon each reference to RandomValue.

```
for i := 1 to NumOfNets do begin
  if i  $\notin$  {w} then assign RandomValue to i;
  NetChange[i] := NewNet[i]  $\Delta$  Net[i];
  Net[i] := NewNet[i];
  Updated[i] := FALSE;
end
```

Figure 4-7: Semantics of pause for VCG

An inductive assertion analysis of the process in Figure 4-6, using the pause semantics in Figure 4-7, proves that the assertions are always satisfied. The assertion at  $L_f$  in turn implies that this process satisfies Eq. (4.7). This enables us to replace the references to NetChanged by references to WasUpdated in the repeat condition. The resulting process never introduces case splits.

This section has shown how assertions and inductive assertion verification can be applied to prove Eq. (4.7), and thereby permit the application of the non-selective trace.

#### 4.2.6. Another Generalized Process

In this application, assertions may tend to be cumbersome and redundant. The loop assertion in Figure 4-6 repeats the entire computation performed by the loop. This section describes a sufficient test of Eq. (4.7) that will avoid the generation of these assertions for some processes.

Figure 4-8 shows the new generalized process of interest. Figure 4-9 shows the same process, modified to conform to the form of Figure 4-4, with assertions added. Period "..." denotes the current state. The term " $X < . >$ " denotes  $(NewNet[w_i]) \langle F_i(. ) \rangle = Net[w_i]) \langle . >$  and  $(Locs_i \langle F_i(. ) \rangle = Locs_i \langle . >)$ , which is taken from Eq. (4.7).

```
begin
  while TRUE do begin
    S;
    repeat pause until  $\bigvee_j \in L_f$  NetChanged[j]
  end
end
```

Figure 4-8: Generalized process for testing without assertions

Eq. (4.7) is satisfied if the assertions hold. Figure 4-10 shows the three paths that must be analyzed to verify the assertions. In this figure, " $Cond$ " denotes  $\bigvee_j \in L_f$  NetChanged[j]. The assume statement conjoins the indicated condition with  $!P_C$ . Consider a treatment of the three paths by symbolic simulation.

Path (a): Symbolically execute path (a). At the end of the path, we must prove

```

begin
  assert TRUE;
  S;
  while TRUE do begin
    repeat pause;
      L[i]: assert (not Vj ∈ {ri} NetChanged[j]K.j)
      ⊃ X<.>;
    until Vj ∈ {ri} NetChanged[j];
  end
end

```

Figure 4-9: Generalized process with assertions

```

assume TR'JE;
S;
assume TR''JE;
S;
pause;
assert (not Cond<.>) ⊃ X<.>;
(a)
assume (not Cond<.>) ⊃ X<.>;
assume Cond<.>;
S;
pause;
assert (not Cond<.>) ⊃ X<.>;
(b)
IPc<.> ⊃ ((not Cond<.>) ⊃ X<.>)
(c)

```

Figure 4-10: Three paths for VCG

*This is equivalent to*

*(IPc<.> and not Cond<.>) ⊃ X<.>*

Therefore, we conjoin (not Cond<.>) to the path condition /Pc, and then wish to show that X<.> holds. Condition X<.> includes references to F<sub>i</sub>(.), which we obtain by symbolically executing S. The complete test for path (a) is then as shown

in Figure 4-11. This test can be performed without first generating the assertions in Figure 4-9, although of course the test itself effectively generates the assertion. The symbolic execution of S is straightforward, since we have required that S introduce no case splits.

```

S;
assume TRUE;
S;
pause;
assume not Vj ∈ {ri} NetChanged[j]K.j;
S;
(* save current state in auxiliary var. $0 *)
S;
assert (NewNet[{wj}]K.$0) = Net[{wj}]K.$0) and
(Locxj<.) = Locsj(\$0);

```

Figure 4-11: Symbolic execution for test of path (a)

Path (b): Path (b) is identical to path (a) except for the initial path condition (the precondition). The precondition for path (b) implies the precondition for path (a), so path (b) is proven by the analysis of path (a). No further analysis is necessary.

Path (c): Consider a conventional execution of this path. The initial assumptions imply X<.>. Either the pause will change one or more of Net[{r<sub>i</sub>}], or it will not. If it does, then Cond is true following the pause, which implies that the final assertion is true. If the pause does not modify any of Net[{r<sub>i</sub>}], then X is true following the pause. This is because X was true before the pause, and the pause did not modify any of the inputs to process i. Then X true implies that the final assertion is again true.

No further analysis of path (c) is required; the definition of Cond guarantees that path (c) is always correct.

Then path (c) is always correct, while the correctness of path (a) implies the correctness of path (b). Hence only path (a) need be tested. The symbolic execution in Figure 4-11 tests path (a) without requiring prior generation of the assertions in Figure 4-9. For example, this test will verify that non-selective trace may be used in Figure 4-1, without requiring the specification of the assertions shown in Figure 4-6.

In conclusion, given a process definition of the type in Figure 4-8, where  $S$  never introduces case splits, we may apply the symbolic execution in Figure 4-11. If the final assertion is proven, then we may use non-selective trace instead of selective trace, substituting  $WasUpdated$  for  $NetChanged$  in

$$\forall j \in \{r_i\} NetChanged[j]$$

The resulting process definition is equivalent to the old, and it never introduces case splits.

#### 4.3. Non-selective Trace Summary

The use of selective trace in event-driven simulations often introduces a large number of case splits in a symbolic simulation. Sometimes, non-selective trace may be substituted for selective trace without altering the function performed by the process. The non-selective trace does not introduce case splits.

Non-selective trace may be substituted for selective trace in the process type typically used for combinational elements. The selective trace is used not to alter the functionality of the process, but rather to improve its efficiency in event-driven simulations. An analysis shows that the weaker non-selective trace may be used instead.

This substitution is also possible in the following more complex process code, where condition  $C$  refers to  $NetChanged$  (selective trace):

```
while TRUE do begin
  repeat pause; L1; until C;
  S;
end;
```

If  $C$  and  $S$  satisfy certain restrictions including Eq. (4.7),

$$\begin{aligned} \forall \text{ valid start states } s \mid L_C(s) &= L_I: \\ (\text{not } C(s)) \supset ((NewNet[\{w_i\}] \& F_i(s)) = Net[\{w_i\}] \& s) \text{ and} \\ (Loc_i \& F_i(s)) &= Loc_i(s)) \end{aligned} \quad (4.7)$$

then  $WasUpdated$  may be substituted for  $NetChanged$  in  $C$ . In this equation,

- $L_C$  is the program counter for this process,
- $Loc_i$  denotes the local variables for this process,
- $\{w_i\}$  is the set of numbers of nets that can be written by this process, and
- $F_i$  is a function describing the action of statement  $S$ .

The new process which accesses  $WasUpdated$  never introduces case splits. An inductive assertion analysis may be helpful in properly defining the set of "valid start states" so that Eq. (4.7) can be proven.

Alternately, a symbolic execution test is possible for processes of the following form:

```
while TRUE do begin
  S;
  repeat pause until  $\forall j \in \{r_i\} NetChanged[j]$ 
end
```

Here,  $\{r_i\}$  is the set of numbers of nets which can be read by this process. If the process satisfies the test, then we can substitute  $WasUpdated$  for  $NetChanged$  in the repeat condition. This test has the advantage of not requiring the use of assertions.

Our discussion did not conjecture or prove whether the test is necessary as well as sufficient.

59

Sometimes, a condition stronger than the simple non-selective condition is preferable. For example, in circuits with feedback, a stronger condition might prevent endless symbolic execution during periods of no actual activity in the circuit. This chapter gave one such stronger condition, which involves more computation than the simpler non-selective condition, but never introduces case splits.

Previous work did not apply symbolic execution to event-driven simulation. In particular, it did not investigate the substitution of non-selective trace for selective trace to reduce case splitting.

60

## Chapter 5

### Case Merging

In the previous chapter, we eliminated case splitting by replacing selective trace by non-selective trace. This is appropriate only in limited cases. In most cases, the process definition arising from this substitution is not equivalent to the original.

This chapter describes case merging, a more general technique for reducing case splitting. A case splitting simulation generates a new simulation for each alternative path at a conditional branch. A case merging simulation, by contrast, combines the effects of all possible alternatives into a single simulation.

The approach is to generate a new process equivalent to the original. The new process executes each basic block [AU77] from the original, subject to a path condition for each block. The new process includes the effects of conditional branches in the computation of conditional values for variables. The transformed process will never introduce case splits.

Cheatham et al. first suggested the use of conditional expressions to combine the results of several paths for a similar static program analysis [CH79]. This chapter applies this idea to the symbolic execution of event-driven simulations with the aid of additional analyses of a process's flow graph.

Section 5.1 shows case merging for a couple examples. Section 5.2 describes the application of case merging in general for processes in which every loop is broken by a pause. Section 5.3 presents the more complex analysis required if a process includes loops not broken by pause. Section 5.4 discusses a few optimizations which help to reduce expression complexity in symbolic execution.

The elimination of case splits does not occur without cost. Typically, the symbolic expressions which are generated by case merging are far more complex than those generated by case splitting. The next chapter will discuss this problem.

### 5.1. Case Merging Examples

#### 5.1.1. Case Merging with Single Entry Point

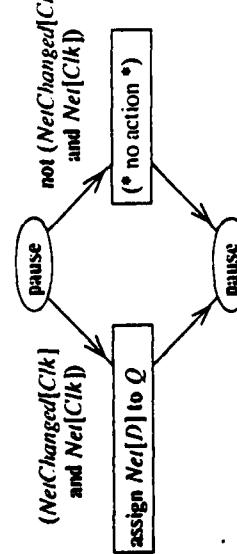
Figure 5-1 shows a process that models an edge-triggered D flip-flop. In a conventional simulation, the pause loop waits until  $Net[CK]$  is true immediately after a change in  $Net[CK]$ . Hence, the pause waits for a positive transition in  $Net[CK]$ . When this edge is detected, the flip-flop copies the D input to net Q and then waits for the next positive  $CK$  edge.

```
process 102: (* Dff *)
  const D = 210; (* indices to Net *)
  CK = 211;
  Q = 212; (* index to NewNet *)
  begin
    while TRUE do begin
      repeat pause until NetChanged[CK] and Net[CK];
      assign Net[D] to Q
    end
  end;
end: (* 102 - Dff *)
```

Figure 5-1: Edge-triggered D flip-flop model

This process does not satisfy the restrictions from the previous chapter. If the statement **assign  $Net[D]$  to  $Q$**  were executed when the **repeat** condition was **false**, it could alter **NewNet[Q]**, violating Eq. (4.7). Therefore, we cannot simply substitute **WasUpdated** for **NetChanged** in the **pause** loop to eliminate case splits. Instead, we will transform the process to eliminate case splits by another method.

**Figure 5-2** shows a graph representation of the process. Vertices represent basic blocks as defined by Aho and Ullman [AU77], with the added provision that a **pause** itself constitutes a separate basic block that appears twice in the graph, once as an entry point and once as an exit point. Conditions on edges out of a vertex denote the conditions under which those edges are followed. This graph is the flow graph defined by Aho and Ullman for dataflow analysis [AU77].



**Figure 5-2:** Flow graph for D flip-flop process

The action of the model in Figure 5-2 on **Net[Q]** can be characterized by the statement

**assign ( $AuxI$  and  $Net[D]$ ) or ( $\neg AuxI$  and  $Net[C/k]$ ) to  $Q$**

where  $AuxI$  is **NetChanged[C/k]** and **Net[C/k]**. The method for deriving this result from the flow graph is covered in the more general discussion later. This example is simple enough that the result is apparent by inspection. Figure 5-3 then shows the transformed version of the D flip-flop process. In this transformed version, the **pause** loop condition could be any condition  $C_x$  such that

$$(NetChanged[C/k] \text{ and } Net[C/k]) \supseteq C_x$$

**Figure 5-3** uses

$$C_x = WasUpdated[C/k]$$

This process is equivalent to that in Figure 5-1, and it never causes a case split. The results of two parallel paths thus are combined into a single simulation.

```

process 102; (* Dff *)
  const D = 210; (* indices to Net *)
  C/k = 211;
  Q = 212; (* index to NewNet *)
  var AuxI: Boolean;
begin
  while TRUE do begin
    repeat pause until WasUpdated[C/k];
    AuxI := NetChanged[C/k] and Net[C/k];
    assign (AuxI and Net[D]) or (not AuxI and Net[Q]) to Q
  end
begin, (* 102 - Dff *)
  
```

**Figure 5-3:** Transformed version of D flip-flop

The claim that the process in Figure 5-3 never introduces case splits depends on the value of **WasUpdated[C/k]** not being symbolic. However, case merging itself appears to put this condition in jeopardy: One path might make an assignment to  $\neg I$  while a parallel path makes no such assignment. If these two paths are

combined into a single simulation by case merging, then the value of  $WasUpdated[i]$  should be symbolic, denoting the condition under which a value was assigned.

For example, the assign in Figure 5-3 sets  $Updated[Q]$  (which becomes  $WasUpdated[Q]$ ) to TRUE, when it would be more accurate to set  $Updated[Q]$  to  $Aux[1]$ . This is because, in Figure 5-2, the process makes an assignment to net  $Q$  if and only if  $Aux[1]$  is true.

The solution to this problem is to recall that the simulation results are by definition independent of  $WasUpdated$ . Therefore we can modify the management of  $WasUpdated$ . The only requirement is that we must have

$$NetChanged[j] \supseteq WasUpdated[j]$$

This requirement is due to our use of  $WasUpdated$  above and in the previous chapter to reduce case splitting. Our strategy will be to set  $Updated[j]$  to TRUE if any path included in a case merging simulation performs an assign to net  $j$ .

Finally, note that the first time the process in Figure 5-1 is called, its execution proceeds from the beginning of the process, not from the pause. Our discussion above did not take this into account. Fortunately, the result is satisfactory, because on this first entry, we know (from Figure 3-4) that  $NetChanged[Clock]$  is false. Therefore the action taken by the process starting from its beginning is the same as if it started from the pause. In both cases, the process goes to the pause with no modifications to  $D$ . In the next example, we will not rely on such a coincidence.

### 5.1.2. Case Merging with Multiple Entry Points

The example in Figure 5-4 copies  $Net[A]$  and  $Net[B]$  to the output net on rising and falling edges of net  $X$ , respectively. This is similar to the double-edge-triggered flip-flop suggested by Unger [Un81]. This description will demonstrate a complication in the elimination of case splitting.

```
process 103; (* OddReg *)
  const A = 220;           B = 221;
                        X = 222;           Y = 223;
  begin
    while TRUE do begin
      repeat pause until NetChanged[X] and Net[X];
      assign Net[A] to Y;
      repeat pause until NetChanged[X] and not Net[X];
      assign Net[B] to Y
    end
  end; (* 103 - OddReg *)
```

Figure 5-4: Process with multiple entry points

The numbered comments in Figure 5-4 mark places at which the process may begin execution in a cobegin, namely its beginning and after a pause. In Figure 5-1, the process has only one such entry point, namely the pause. (Recall that in the D flip-flop, starting from the beginning was equivalent to starting from the pause, so there was effectively only one entry point.)

Figure 5-5 shows the flow graph for this process. By convention, the beginning of a process (the first entry point) appears as a single basic block with no predecessors. The presence of multiple entry and exit points complicate the analysis. We will introduce an extra array  $LPC$  containing one Boolean element for each vertex (basic block) in the graph. The number labeling each vertex in

Figure 5-5 denotes the corresponding element of  $Lpc$ . Note that each pause gets two such indices; one for its function as an entry and one for its function as an exit. Figure 5-6 then shows the transformed process, described below.

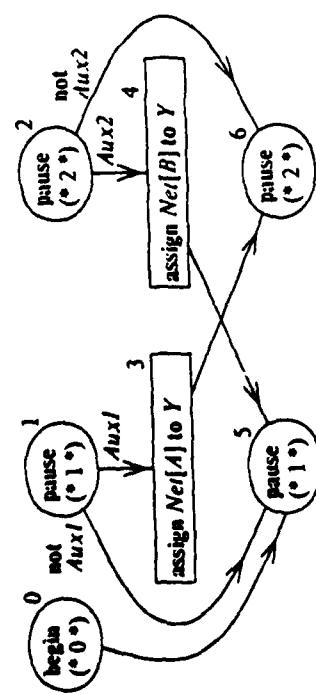


Figure 5-5: Flow graph for *OddReg* process

```

process 103; (* OddReg *)
  const A = 220;      B = 221;
  X = 222;          Y = 223;
  var Aux1, Aux2: Boolean;
  Lpc: array [0..6] of Boolean;

begin
  Lpc[0] := TRUE;
  Lpc[1] := FALSE;
  Lpc[2] := FALSE;
  while TRUE do begin
    Aux1 := NetChanged[X] and Net[X];
    Aux2 := NetChanged[X] and not Net[X];
    Lpc[3] := Lpc[1] and Aux1;
    Lpc[4] := Lpc[2] and Aux2;
    Lpc[5] := Lpc[0] or Lpc[1] and not Aux1 or Lpc[4];
    Lpc[6] := Lpc[3] or Lpc[2] and not Aux2;
    assign (Lpc[3] and Net[A]) or (Lpc[4] and Net[B]) or
    ((Lpc[0] or Lpc[1] and not Aux1 or Lpc[2] and not Aux2) and
    Net[Y]);
    to Y;
  end
end; (* 103 - OddReg *)

```

Figure 5-6: Modified process with multiple entry points

$Lpc[j]$  becomes true if and only if basic block  $j$  is to be run during this execution of the process.  $Lpc[j]$  may be computed from the values of  $Lpc$  for the predecessors of block  $j$ , taking into account any conditions associated with the edges. A topological sort on the graph yields an appropriate order for the processing of the blocks, since the graph is acyclic. The order in which the vertices are numbered happens to be a valid topological sort, and so is the order we use. The value assigned to net  $Y$  is then computed using the values of  $Lpc$  computed for the various basic blocks.

Now, the process in Figure 5-6 can be executed symbolically without case splitting. This is true since the only conditional branch expression in the process is  $WasUpdated[X]$ , and by convention  $WasUpdated$  will never have symbolic values in a symbolic simulation. (Recall discussion near end of previous section.)

$Lpc[5..6]$  become the initial values of  $Lpc[1..2]$  for the next execution of the process. The pause loop condition may be any function  $C_x$  such that  $(Aux1 \text{ or } Aux2) \supset C_x$ .

We choose  $C_x$  equal to  $WasUpdated[X]$ .

Note that in a symbolic simulation,  $L_{Pc}$  may have symbolic (not constant) values. This implies that in a case merging symbolic simulation, we cannot uniquely identify the path executed by a process during a given cobegin call. By the same token, we can assign symbolic initial values to  $L_{Pc}$ , yielding a more general symbolic execution.

A later section will discuss the derivation of the assignments to  $L_{Pc}$  and net  $Y$  in Figure 5-6.

### 5.1.3. Summary of Examples

The above case merging examples cannot eliminate case splitting simply by substituting non-selective trace for selective trace. Rather, each example combines the results of all possible paths into a single simulation. The second example shows that in general, the simulator may have to combine paths with different starting and stopping points, since the path followed by the process is not uniquely identified.

The following sections will describe the application of this technique to achieve symbolic simulation without case splitting for a large class of process types. The major area of concern is how to characterize the behavior of the process given the flow graph.

## 5.2. Analysis of Acyclic Flow Graphs

The generation of a flow graph like those in Figures 5-2 and 5-5 is straightforward. The graph will be acyclic if every loop in the process is broken by a place. If the graph is acyclic, we can construct a new process according to the following procedure, illustrated with the *Old/Reg* example. This construction yields

an equivalent process of a form similar to that in Figure 5-6. This resulting process can be executed symbolically without case splitting.

The strategy will be to construct a process that executes every basic block from the original process once. Each basic block will use the results from any preceding basic blocks. Since every basic block is executed, there will be no need for conditional branches to decide which blocks to execute. The removal of the conditional branches eliminates case splitting. The only remaining conditional branch will be in the pause loop.

### 5.2.1. Generality of the Construction

This construction is applicable to any process satisfying two restrictions:

- Every loop is broken by a pause.
- During symbolic execution, every subscript evaluates to a constant.

If the first restriction is not satisfied, then a more general analysis is required. This will be considered later in this chapter. The second restriction allows us to ignore the problems of case splitting due to subscript evaluation. If it is not satisfied, we can rewrite the code to satisfy the restriction, or we can use more complex symbolic expressions that permit non-constant subscripts.

For example, in the expression

$Reg[Reg.Spec]$

the subscript *RegSpec* must evaluate to a constant. However, it is common for *RegSpec* to be part of a symbolically specified machine instruction that does not evaluate to a constant.

If the subscript range is small, then the subscripted reference can be replaced by a case statement which extracts the correct value. If there were only two registers, we could write the following:

```
case RegSpec of
  0: RSubR := Reg[0];
  1: RSubR := Reg[1];
end;
```

Subsequent references to  $Reg[RegSpec]$  would then be replaced by " $RSubR$ ".

If the subscript range is not small, then the subscript must evaluate to a constant, a serious limitation in some cases. We can lift this restriction by using more powerful symbolic representations of data structures, such as those used by Luckham and Suzuki for Pascal [LS79] (originally proposed by McCarthy for a different application [Mc62]). This in turn would require more powerful symbolic expression manipulation.

### 5.2.2. The Construction

This section describes the generation of code equivalent to the process represented by the flow graph, but with no conditional branches except in the pause loop. The result of this construction for *OddReg* appears in Figure 5-7.

Every vertex (basic block) in the flow graph will have its own copy of every variable that the process can modify. In *OddReg*, the only such variables are *NewNet[Y]* and *Updated[Y]*. According to the convention noted earlier, we will not compute a symbolic value for *Updated[Y]* in symbolic simulations (it will be set true), so ignore it here. Let *NewY[0..6]* be the array of copies of *NewNet[Y]* for the several basic blocks.

In addition to the multiple copies of modified variables, we also have Boolean array *LPr*, which holds a "local path condition" for each basic block. Initially, the element of *LPr* corresponding to the beginning of the process (the first entry point) is *TRUE*, while the elements for the other entry points are *FALSE*.

The remaining code to be generated becomes the body of a while *TRUE* do loop.

Consider the blocks in any order that is a valid topological sort. The reason this order is important is that basic block *j* needs the final values of variables from its immediate predecessors to initialize its own variables. Ordering the construction by a topological sort ensures that these final values will be available. The topological sort is of complexity  $O(V, E)$ , where  $V$  is the number of vertices (basic blocks) and  $E$  is the number of edges (branches between basic blocks) [AH74].

For each basic block *j* in order, generate the following three sections of code:

1. Code to compute *LPr[j]*.

If block *j* is an entry point (no predecessors), then *LPr[j]* is already set and need not be computed. Otherwise, let  $\{b_j\}$  be the numbers of the predecessors of *j*. Let  $C(k, j)$  be the condition on the edge from block *k* to *j*. If no such condition is given for edge  $\langle k, j \rangle$ , then  $C(k, j)$  is *TRUE*. Then *LPr[j]* is given by the following:

$$LPr[j] := \text{or}_{k \in \{b_j\}} (LPr[k] \text{ and } C(k, j))$$

For example, *LPr[6]* in Figure 5-5 is found by

$$LPr[6] := LPr[3] \text{ or } LPr[2] \text{ and not Aux2}$$

73            If  $C(k,j)$  reads any variables modified by this process, it reads the copies local to block  $k$ .

## 2. Code to initialize the local copies of variables for block $j$ .

If block  $j$  is an entry point, then the local copies of variables are copied from the actual variables. For example, the initialization for block 1 in Figure 5-5 is simply

$$\text{NewY}[1] := \text{NewNet}[Y]$$

(To observe the restriction against reading  $\text{NewNet}$ , substitute  $\text{Net}$  for  $\text{NewNet}$  here; the two are equal at this point.)

Otherwise, block  $j$  must initialize its local variables from the variables in its predecessor blocks. If block  $j$  has more than one predecessor, then the values computed will be conditional, denoting several possible alternatives computed in different paths. For example, the value computed for  $\text{NewY}[6]$  in Figure 5-5 is  $\text{NewY}[6] := LP[3] \text{ and } \text{NewY}[3] \text{ or } LP[2] \text{ and not } \text{Net}[2]$ . Notice the similarity between this assignment and the computation of  $LP[6]$ .

## 3. Code to execute block $j$ .

The code from block  $j$  can be used directly, with two modifications: (i) Any references to variables that can be modified by this process are replaced by references to the copies local to this basic block. (ii) The assign statement is a special case of the first modification. This reference to  $\text{NewNet}$  and  $Updated$  is replaced by a simple assignment to the local copy of  $\text{NewNet}$ . For example, the code for block 4 in Figure 5-5 becomes the following:

74

$$\text{NewY}[4] := \text{Net}[B]$$

We must remember to set the appropriate elements of  $Updated$  to  $TRUE$  later.

Following the above generation of code for all the basic blocks, generate code to merge the results from the several exit blocks. This merging is similar to the initialization of local variables in a block with predecessors, except now the variables being written are the actual variables, not local copies. Various assign statements will update  $\text{NewNet}$  and set appropriate elements of  $Updated$  to  $TRUE$ .

Next, generate code to copy the values of  $LP_C$  for the exit blocks back to the corresponding elements for the entry blocks. The element of  $LP_C$  for the initial entry point is set false here.

Finally, generate a pause loop that terminates this call to the process. The pause loop condition may be any condition that is guaranteed to be satisfied whenever any of the pause loop conditions from the original process is satisfied. The condition also should never cause a case split in a symbolic simulation. A condition satisfying both of these requirements can always be found; in the worst case, the condition can be  $TRUE$ .

The above construction carried out for  $OutReg$  yields the process in Figure 5-7.

```

process 103; (* OddReg *)
  const A = 220;      B = 221;
  X = 222;          Y = 223;
  var LPc, NewY: array [0..6] of Boolean;
  LPc[0]:= TRUE;
  LPc[1]:= FALSE;
  while TRUE do begin
    NewY[0]:= Net[Y];
    NewY[1]:= Net[X];
    NewY[2]:= Net[Y];
    NewY[3]:= LPc[1] and NetChanged[X] and Net[X];
    LPc[3]:= NewY[1];
    NewY[3]:= Net[A];
    NewY[4]:= LPc[2] and NetChanged[X] and not Net[X];
    LPc[4]:= NewY[2];
    NewY[4]:= Net[B];
    NewY[4]:= LPc[0] or LPc[1] and not (NetChange[X] and Net[X]);
    LPc[5]:= LPc[0] or LPc[1];
    or LPc[4];
    NewY[5]:= LPc[0] and NewY[0] or
    LPc[1] and not NetChanged[X] and Net[X] and NewY[1] or
    LPc[4] and NewY[4];
    LPc[6]:= LPc[3] or LPc[2] and not (NetChange[X] and Net[X]);
    NewY[6]:= LPc[3] and NewY[3] or
    LPc[2] and not (NetChange[X] and not Net[X]) and NewY[2];
    assign LPc[5] and NewY[5] or LPc[6] and NewY[6] to Y;
    LPc[0]:= FALSE;
    LPc[1]:= LPc[5];
    repeat pause until WasUpdated[X]
    end
  end; (* 103 - OddReg *)

```

**Figure 5-7:** Algorithmically generated process for case merging

Alternately, we may symbolically execute this process independently and use the symbolic results to define a new version of the process. In brief, this is done by (i) initializing all program variable values (including  $LPc$ ) to distinct symbolic variables, (ii) symbolically executing the body of the loop once, and (iii) generating assignment statements for the several variables from their resulting symbolic values.

The latter option has the advantages of perhaps eliminating many intermediate computations and allowing a one-time optimization of the symbolic expressions generated by this process. The process given in Figure 5-6 represents a mixture of these two options. The computations for array  $LPc$  in Figure 5-6 are taken directly from Figure 5-7; the only change is the use of auxiliary variables  $Aux1$  and  $Aux2$ . But the computation of  $NewNet[Y]$  in Figure 5-6 is a substantial abbreviation of the computation in Figure 5-7.

#### 5.2.4. Application of Construction to Routines

The above construction can also be applied to function and procedure ("routine") definitions as well as to process definitions. While a process has entry and exit points at each pause plus an entry point at its beginning, a routine will have just one entry point, at its beginning, and one exit point, at its end. The construction then is the same as for processes, except that the entry element of the routine's  $LPc$  array is set equal to the value of  $LPc$  for the calling basic block.

#### 5.2.3. Use of Constructed Process

A process resulting from the above construction, for example the one in Figure 5-7, can be used as the final form of the process. It can be executed symbolically without case splitting. In the symbolic execution, array  $LPc$  will have symbolic values which can be considered as being path conditions for their associated basic blocks.

As with processes, the result of the construction can itself be the final version, or it can be symbolically executed to generate a shorter final version. In the latter case, one must take care to account for possible anomalous parameter bindings, such as the binding of one variable to two distinct call-by-reference formats in the same call.

### 5.3. Analysis of Flow Graphs with Loops

Figure 5-8 shows an outline of the process generated by the construction of Section 5.2.2. The generated process executes each basic block from the original process exactly once. This construction is not adequate for processes with loops not broken by pause. Basic blocks in such loops may be executed more than once in one call to the process. We need a mechanism that will allow the execution of basic blocks more than once.

```

init elements of LPc for entry basic blocks;
while TRUE do begin
  for each basic block j, taking the blocks in topological order,
    do begin
      set LPc[j];
      init Vars[j] from Vars[{bj}];
      execute block j
    end;
    update vars modified by process, including assigns;
    copy LPc elements for entry blocks from elements for exit blocks;
  repeat pause until Condition
  end;

```

**Figure 5-8:** Case merging outline for process without loops

One possibility is to analyze the loop to determine a closed-form representation of its computation. An inductive assertion analysis or the solution of recurrence relations [CH79] are possible approaches. This is an open research problem which we will not consider here.

A second possibility is to modify the outline of Figure 5-8 to handle loops.

The remainder of the section discusses this approach.

#### 5.3.1. New Outline for Case Merging with Loops

In Section 5.2.2, the choice of the next basic block is statically determined by the structure of the generated process. Figure 5-9 shows an equivalent outline that uses a dynamic mechanism instead. This figure marks lines that differ from the outline of Figure 5-8. *BBSel* is a set of integers. We have  $j \in BBSel$  when one of *j*'s predecessor blocks has been executed. This indicates that block *j* should be executed. Set *BBSel* is managed as a priority queue, where the priority of each block is determined by its position in a topological sort. Recall the motivation for topological sort from Section 5.2.2.

```

init elements of LPc for entry basic blocks;
while TRUE do begin
  init BBSet to include entry basic block numbers,
  while BBSet includes numbers of blocks that are not exit blocks
    do begin
      choose and remove non-exit block j from BBSet;
      set LPc[j];
      init Vars[j] from Vars[{bj}];
      execute block j;
      add to BBSet the numbers of successor blocks
    end;
  update vars modified by process, including assigns;
  copy LPc elements for entry blocks from elements for exit blocks;
  repeat pause until Condition
end;

```

**Figure 5-9:** Equivalent case merging version using basic block set

Set *BBSel* initially contains only the numbers of the entry blocks. When a block is selected, the simulator removes its number from *BBSel* and adds the numbers of its successors. Execution continues until *Condition* of exit blocks.

The outline of Figure 5-9 still cannot handle processes with loops. Figure 5-10 presents a modified outline that corrects the problems noted below.

```

1 init elements of  $L.Pc$  for entry basic blocks;
2 while  $TRUE$  do begin
3   init non-entry  $L.Pc$  elements to  $FALSE$ ;
4   init  $Vars$  elements to  $FALSE$  (empty);
5   init  $BBSel$  to include entry basic block numbers;
6   while  $BBSel$  includes numbers of blocks that are not exit blocks
7     do begin
8       choose and remove non-exit block  $j$  from  $BBSel$ ;
9       set  $UseL.Pc$ ;
10      if  $UseL.Pc$  is satisfiable then begin
11        init  $UseVars$  from  $Vars[\{b_j\}]$ ;
12        adjust  $Vars[\{b_j\}]$ ;
13         $\forall k \in b_j \text{ do } L.Pc[k] := L.Pc[k] \text{ and not } C(k,j)$ ;
14        execute block  $j$ ;
15        add to  $BBSel$  the numbers of successor blocks;
16         $Vars[j] := Vars[j] \text{ or } UseVars \text{ and } UseL.Pc$ ;
17         $L.Pc[j] := L.Pc[j] \text{ or } UseL.Pc$ 
18      end;
19    ut late vars modified by process, including assigns;
20    copy  $L.Pc$  elements for entry blocks from  $L.Pc$  for exit blocks;
21    repeat pause until Condition
22    end;
23

```

if  $UseL.Pc$  then begin ...

However, we want to be able to symbolically execute this process without case splitting. This explains the use of the more complex test:

if  $UseL.Pc$  is satisfiable then begin ...

The end of this section makes further observations on the use of this test.

Second, the presence of loops implies that a basic block  $j$  can be executed before one of its predecessors  $b_j$  has been executed. In this case,  $L.Pc[b_j]$  should be FALSE. Line 3 in Figure 5-10 initializes  $L.Pc$  to ensure this is true.

Third, the presence of loops implies that a basic block  $j$  may be executed more than once. Suppose block  $k$  is a predecessor of block  $j$ . If condition  $L.Pc[k]$  and  $C(k,j)$  is included in the first execution of block  $j$ , then it should not also be included in subsequent executions of block  $j$ . In Figure 5-10, line 13 makes the necessary adjustment to  $L.Pc[k]$  to prevent this problem.

Finally, the second execution of a basic block should not destroy the results of its first execution. This implies that the basic block requires two local copies of variables that can be modified. In Figure 5-10,  $UseVars$  is the copy of variables for the current computation, while  $Vars$  accumulates and saves the results from previous computations for later use. Similarly, each basic block has two copies of the path condition,  $UseL.Pc$  and  $L.Pc$ . But since only one basic block is executed at a time, then a separate copy of  $UseVars$  and  $UseL.Pc$  is not required for each basic block.  $UseVars$  and  $UseL.Pc$  need not be arrays.

**Figure 5-10:** Case merging outline for process with loops

First, the outlines of Figures 5-8 and 5-9 include no provision for omitting the execution of a basic block  $j$  which has  $UseL.Pc \equiv FALSE$ . (The use of  $UseL.Pc$  instead of  $L.Pc[j]$  is explained later.) This is merely inefficient when the original process is loop-free. But if the process has loops, this problem prevents the termination of the loop. Line 10 in Figure 5-10 makes the necessary check. For conventional execution,  $UseL.Pc$  is always  $TRUE$  or  $FALSE$ , so the following test would suffice:

**Line 4** initializes the elements of *Vars* to reflect the fact that no results have been accumulated yet. Lines 9 and 11 set *UseLPc* and *UseVars*, not *LPc[j]* and *Vars[j]*, for current use. The adjustment of *Vars[{b<sub>j</sub>}]* in line 12 is similar to the adjustment of *LPc[{b<sub>j</sub>}]* in line 13, discussed above. Lines 16 and 17 merge the results *UseVars* and *UseLPc* into *Vars[j]* and *LPc[j]*.

This completes the modification of the case merging outline to support processes with loops. The only remaining detail is the selection of the next basic block from *BBSel*: execution, covered in the next section. The remainder of this section discusses further the satisfiability test for *UseLPc* in line 10.

Consider a symbolic execution of this process in which *UseLPc* is satisfiable but not identically true. Since *UseLPc* is satisfiable, block *j* will execute. Now we are not introducing a case split even though the condition is not identically true. This implies that this symbolic execution represents some conventional executions which would not execute block *j* at this point. The symbolic execution must predict no changes due to the block for such cases.

Block *j* modifies *UseLPc* and *UseVars*. However, each block initializes these variables itself before using them, so any modifications to *UseLPc* and *UseVars* by block *j* are immaterial. Block *j* modifies *LPc[j]* and *Vars[j]*, but it computes conditional values, so that (*not UseLPc*) implies that *LPc[j]* and *Vars[j]* do not change. Similarly, the conditional adjustment to *Vars[{b<sub>j</sub>}]* and *LPc[{b<sub>j</sub>}]* in lines 12 and 13 ensure that these variables are not changed when *UseLPc* is false. For *LPc[k]*, *k*  $\in \{b\}$ , we have

$$\begin{aligned} (\text{not } UseLPc) &\Rightarrow \text{not } (LPc[k] \text{ and } C(k, j)) \\ &\Rightarrow (LPc[k] \supset \text{not } C(k, j)) \\ &\Rightarrow ((LPc[k] \text{ and not } C(k, j)) \equiv LPc[k]) \end{aligned}$$

The argument for *Vars[{b<sub>j</sub>}]* is similar.

The other issue concerning the satisfiability test on *UseLPc* is its effect on the generality of this approach to case merging. For a loop to terminate during symbolic execution, *LPc[j]* for every block *j* in the loop must eventually become unsatisfiable. This implies that the loop must unfold during symbolic execution; that is, there must be a known finite bound on the number of passes through the loop. A similar restriction applies to loops implemented by recursive routine calls; otherwise, a routine may call itself endlessly, never terminating. These restrictions limit the generality of processes to which case merging as described in this section can be applied.

Note that this technique for symbolically executing loops is likely to generate complex symbolic expressions, especially if the loop condition does not resolve to *TRUE* or *FALSE* on each pass through the loop.

### 5.3.2. Selection of Basic Blocks for Execution

For loop-free processes, a topological sort yields appropriate priorities for use in Figure 5-10. However, in processes with loops, a topological sort cannot be found. In this case, it is desirable to find a priority assignment which will tend to minimize the numbers of times basic blocks are executed. Given a flow graph for a process, a possible heuristic is to remove edges until the graph is acyclic. Then carry out a topological sort on the graph to determine block priorities.

The guiding principle for the removal of edges is as follows: In the topological sort, all vertices (basic blocks) in the body of the loop should precede vertices for code following the loop. This ensures that when the process is in a loop, it will completely exit the loop before proceeding from the loop. The choice of edges removed from the graph should ensure that any topological sort will satisfy this property.

For example, in Figure 5-11a, the simulator should select block  $s$  from BBSet only if none of blocks  $w$ ,  $x$ ,  $y$ , or  $z$  remain in BBSet. Figure 5-11b shows the graph from Figure 5-11a with edge  $\langle w, x \rangle$  removed. All of vertices  $w$ ,  $x$ ,  $y$ , and  $z$  will precede vertex  $s$  in any topological sort of a graph containing this subgraph.

the fact that Tarjan's algorithm to find strongly connected components [AH74] directly yields a topological sort when applied to a loop-free graph.

1. Delete all self-loops  $\langle v, v \rangle$  in the flow graph for the process. Call the resulting graph  $G_{pr}$ .

2. Find the strongly connected components in the  $G_{pr}$ . Assign each component a unique number. Let  $\text{COMP}(v)$  denote the number of the component containing vertex  $v$ . The component numbers should be ordered so that  $\text{COMP}(v) \geq \text{COMP}(w) \text{ iff } v \rightarrow^* w$ . Tarjan's algorithm for finding strongly connected components is  $O(V, E)$  (linear on the number of vertices and edges) [AH74]. It also easily yields component numbers satisfying the above ordering.

3. Any component which contains more than a single vertex has a loop. For each such component  $i$ , do steps 4 and 5.
4. Let  $m$  be the largest integer such that  $i > m$  and there exists an edge  $\langle v, w \rangle$  in  $G_{pr}$  with  $v$  in component  $i$  and  $w$  in component  $m$ . Component  $m$  is a closest successor of component  $i$ , in the sense that if  $v \rightarrow^* x \rightarrow^* w$ , then vertex  $x$  is either in component  $i$  or in component  $m$ .

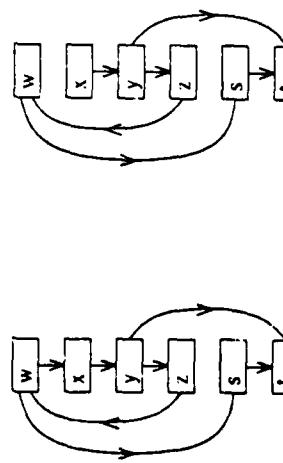
(a) Graph with loop

(b) Loop broken

**Figure 5-11:** Graph with loop not broken by dismissal point

The following algorithm describes the assignment of vertex (block) priorities in graphs containing loops. It selectively removes edges until the graph is loop-free, at which point a topological sort gives the vertex priorities. This algorithm exploits

5. Choose a vertex  $s$  in component  $i$  such that there exists an edge  $\langle s, u \rangle$  with  $\text{COMP}(i) = m$ . Remove from the graph all edges  $\langle s, u \rangle$  for which  $\text{COMP}(u) = i$ .



6. Repeat steps 2 through 5 until step 2 finds no strongly connected component with more than a single vertex. When this occurs, the component numbering generated by step 2 is the priority assignment to be used for  $G_{pr}$ : the higher the value of  $\text{COMP}(v)$ , the higher the priority of vertex (basic block)  $v$ .

We can apply this procedure to the graph in Figure 5-11a. This graph has no self-loops. After step 2, component number 1 has vertex  $t$ , component 2 has vertex  $s$ , and component 3 has vertices  $w$ ,  $x$ ,  $y$ , and  $z$ . Only component 3 has more than one vertex. For this component, the closest successor is component 2. Edge  $\langle w, s \rangle$  is incident out of component 3 and into component 2. Edge  $\langle w, x \rangle$  is incident out of  $w$  inside component 3, so we remove that edge. The first execution of steps 2 through 5 then yields the graph in Figure 5-11b.

The next execution of step 2 does not find any component with more than a single vertex. The vertex ordering generated by step 2 is (highest priority to lowest)  $w, x, y, z, t$ . This ordering guarantees that the simulator will not select label  $s$  or  $t$  unless none of  $w, x, y$ , and  $z$  are in the queue.

This procedure can be summarized as shown in Figure 5-12. Each pass through the repeat loop takes  $O(V, E)$  time. In the worst case,  $O(V)$  passes through the repeat are required. This is so because

- each execution of step 5 will cause vertex  $s$  to be a separate component in the next pass,
- therefore each execution of step 5 increases the number of components by at least one,
- the maximum number of components is  $O(V)$ , and

- an example can be constructed in which  $O(V)$  passes are actually required.

```

begin (* priority assignment *)
remove self-loops < v, v > from graph, yielding graph  $G_{pr}$ ;
repeat
    find strongly connected components of  $G_{pr}$ ;
    numbering components so that
    component  $j$  is reachable from a different component  $i$ 
    only if  $i > j$ ;
    tag each vertex with its component number;
    for each component  $i$  with more than one vertex do begin
        find edge <  $s_i, t_i$  > with  $s$  in component  $i$ ,
         $t$  in component  $m$ , a closest successor of component  $i$ ;
        remove all edges <  $s_i, u$  > for which  $u$  is in component  $i$ 
    end
    until no component has more than one vertex;
current component numbering provides desired basic block priorities
end; (* priority assignment *)

```

Figure 5-12: Priority assignment for graphs with loops

The assignment of priorities to labels is then  $O(V^2, E)$  in the worst case, although probably much faster if the control flow is structured and loops are not nested deeply. In structured process descriptions, fast (linear) algorithms might take advantage of the regular structure present in the original process definition, entirely bypassing the generation of the flow graph.

#### 5.4. Optimizations

The new processes generated by a case merging construction will compute complex symbolic expressions. Therefore optimization of the constructed processes will be important. Whereas optimizations typically aim to reduce time or space requirements for code, optimizations of the case merging processes for symbolic

simulation would aim to reduce the complexity of symbolic expressions, since the comparison of symbolic Boolean expressions is np-hard, for example.

Following is a brief description of some possible optimizations. The list is non-exhaustive. Algorithmic application of the optimizations is not discussed. Opportunities for these optimizations may be found by heuristics or dataflow analysis.

In Figure 5-13, the case merging construction would generate

$$LPc[5] := LPc[4] \text{ or } LPc[2] \text{ and not } Aux2 \text{ or } LPc[3]$$

But from the figure, it is clear that

$$LPc[5] := LPc[1]$$

is an equivalent and more desirable alternative. Similarly, if the basic blocks have local copies  $QX[j]$  of variable  $X$ , but blocks 2, 3, and 4 do not modify  $X$ , then the assignment

$$\begin{aligned} QX[5] &:= LPc[4] \text{ and } QX[4] \text{ or } LPc[2] \text{ and not } Aux2 \text{ and } QX[2] \\ &\quad \text{or } LPc[3] \text{ and } QX[3] \end{aligned}$$

can be optimized to

$$QX[5] := QX[1]$$

Subtler possibilities exist. If a variable  $T$  is written only in block 1 and read only in blocks 2, 3, 4, and/or 5, then local copies of  $T$  are not required at all.

In Figure 5-14, assume the blocks have local copies of variable  $Z$ . According to the case merging construction, block 3 computes the following conditional expressions:

$$\begin{aligned} LPc[3] &:= LPc[1] \text{ and } Aux1 \text{ or } LPc[2]; \\ QZ[3] &= LPc[1] \text{ and } Aux1 \text{ and } QZ[1] \text{ or } LPc[2] \text{ and } QZ[2]; \end{aligned}$$

Subsequently, block 5 will compute a conditional value for  $QZ[5]$ :

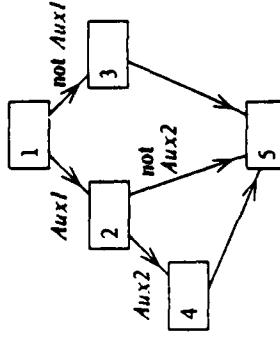


Figure 5-13: Example for optimizations

$$QZ[5] := LPc[3] \text{ and } Aux2 \text{ and } QZ[3] \text{ or } LPc[4] \text{ and } QZ[4];$$

Now, if block 3 did not modify  $Z$ , then we can substitute for  $LPc[3]$  and  $QZ[3]$ ,

finding the following:

$$\begin{aligned} QZ[5] &:= (LPc[1] \text{ and } Aux1 \text{ or } LPc[2]) \text{ and } Aux2 \text{ and } \\ &\quad (LPc[1] \text{ and } Aux1 \text{ and } QZ[2]) \text{ or } \\ &\quad (LPc[2] \text{ and } QZ[2]) \text{ or } \\ &\quad LPc[4] \text{ and } QZ[4]; \end{aligned}$$

From this we see that adding  $(LPc[3] \text{ and } Aux2)$  with  $QZ[3]$  is a redundant operation in the computation of  $QZ[5]$ . ( $Aux2$  is redundant because only one path goes from block 3 to block 5; compare to the computation of  $QX[5]$  in Figure 5-13, with two paths from block 2 to block 5.) Another example of this appears in the computation of the value assigned to  $Y$  in Figure 5-7. The case merging construction for processes with loops aggravates this problem, since  $Vars$  values are always conditional expressions (line 16 in Figure 5-10).

which may execute some basic blocks from the original process more than once. The new process contains a conditional branch in the processing of each block in addition to the one in the pause loop. The new branch conditions test for the *satisfiability* of an expression, and so do not introduce case splits in a symbolic execution. A symbolic execution of this new process is guaranteed to terminate only if all loops in the original process unfold during symbolic execution. This restricts the generality of this approach for processes with loops.

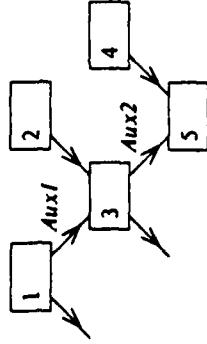


Figure 5-14: Second example for optimizations

### 5.5. Case Merging Summary

This chapter has described the symbolic execution of processes without case splitting in an event-driven simulation. The approach is to create a flow graph for the process of interest. In this graph, each path from an entry point to an exit point represents a possible execution path for the process during one *cobegin* statement. Then from the graph, generate a new process equivalent to the original. The discussion illustrated this approach with two examples, and then described it for the general case.

If the graph is loop-free, then the new process executes each basic block from the original process exactly once. The new process contains no conditional branches except for one in a pause loop that does not introduce a case split when symbolically executed. The new process can be symbolically executed without case splitting. It therefore combines all possible paths as represented by the several paths in the graph; hence the name "case merging."

If the graph is not loop-free, a different construction yields a new process

In the case of graphs with loops, a criterion is needed for choosing the order in which the basic blocks from the original process will be executed. A heuristic selectively removes edges from the flow graph until it is loop-free. A (static) priority is assigned to each basic block according to a topological sort on the graph. A priority queue manager then controls the selection of basic blocks for execution.

This chapter also described a few optimizations which may reduce the complexity of the symbolic expressions generated in a case merging simulation.

Cheatam, Holloway, and Townley used conditional expressions for case merging in another context [CHT79]. They built a similar graph model and used conditional expressions to characterize the effects of multiple paths in a static program analysis. This chapter added the following enhancements to their work:

- treatment of processes with multiple entry points (essentially coroutines); use of topological sort to order the analysis of such routines,
- extraction of loop structure in Section 5.3.2,
- optimizations from Section 5.4.

Cheatam et al. did not include a comparison of case merging and case splitting for symbolic execution. Such a comparison appears in the next chapter.

## 5.6. Symbolic Execution Summary

The last three chapters have discussed techniques for symbolically executing an event-driven simulation. The discussion has followed this outline:

Define a model for event-driven simulation. The model consists of concurrent processes which communicate with each other only through global variables in a controlled fashion. This protocol is defined so that the results of a simulation are independent of the interleaving of processes in the simulation. This follows the philosophy of SABLE.

Since simulation results are independent of process interleaving, we choose an interleaving in which processes do not overlap. Then we may consider symbolic execution of the simulation one process at a time. For some processes, a simple substitution of non-selective trace for selective trace eliminates case splitting without changing the functionality of the process. For more complex processes, the general case merging analyses from this chapter yield equivalent processes which do not introduce case splits when symbolically simulated.

In a symbolic simulation, the process created by a case merging analysis typically generates more complex symbolic expressions than those generated in a case splitting simulation. The next chapter discusses this problem further.

Section 6.4 presents a second example for which case merging is preferable. This example also includes two descriptions plus the simulation relation.

## Chapter 6

### Symbolic Execution Examples

**Case merging** seeks to improve the performance of symbolic execution by combining the results of all possible paths into a single execution. However, it typically does so at the cost of generating symbolic results far more complex than those produced in a case splitting symbolic execution. It is possible that the increased expression complexity can more than negate the reduction in the number of case analyses. This is true because of the difficulty of analyzing symbolic expressions.

This chapter presents qualitative criteria that suggest when case merging may be an effective approach for reducing the cost of a symbolic execution analysis. Section 6.1 discusses the complexity of generating and analyzing symbolic expressions. Section 6.2 discusses the complexity of case splitting and case merging symbolic executions.

Section 6.3 presents one example for which case splitting is preferable to case merging. This example includes two independent descriptions related by a simulation relation, as described in Section 2.2 for microcode verification. The complete results of the case splitting and case merging executions of the two descriptions appear in Appendix C.

### 6.1. Symbolic Expression Complexity

The factor favoring case merging over case splitting in symbolic execution is the reduction in the total number of paths. The factor favoring case splitting is the increased symbolic expression complexity introduced by case merging. Other factors remaining equal, we expect the computation time to increase linearly as the number of paths. The cost of the increased expression complexity in the analysis of the case merging results is more difficult to evaluate.

The fundamental operation in the analysis of the results of symbolic executions is the test that a Boolean assertion is identically true. We now consider the cost of this operation for simple Boolean expressions, universally quantified and not including arithmetic quantities or operations. Previously, we have not distinguished between a function and its symbolic representation. Now, we will let “string” denote the symbolic representation of a function.

Let  $\{0,1\}^n$  denote the set of  $n$ -bit vectors. Then

$$f: \{0,1\}^n \rightarrow \{0,1\}$$

is a Boolean function of  $n$  Boolean variables. Let  $A^*$  be the set of finite-length strings using the alphabet  $A$ . The following result indicates the potential for complex strings representing Boolean functions.

**Theorem 6.1:** Given relation *Rep*,  
 $Rep \subseteq \{(0,1)^n \rightarrow \{0,1\}\} \times A^*$ .

with the following properties:

$\forall f \in (\{0,1\}^n \rightarrow \{0,1\}) \text{ and } s \in A^* : ((f \text{ Rep } s) \text{ and } (f2 \text{ Rep } s))$

$$\supseteq (f1 = f2).$$

$\forall f \in (\{0,1\}^n \rightarrow \{0,1\}) : \exists s \in A^* \mid f \text{ Rep } s.$

Let  $A_{Rep}^*$  denote the set of strings  $s$  in  $A^*$  such that  $s$  is the shortest string for which  $f \text{ Rep } s$  for some function  $f$ , and let  $\alpha \geq 2$  be the number of symbols in  $A$ . Then the fraction of strings in  $A_{Rep}^*$  with length less than

$$\Gamma(2^n / \lg \alpha) \cdot i - 1$$

is less than  $\alpha^{-i}$ , for integer  $i \geq 0$ .  $\square$

The strings in  $A_{Rep}^*$  are the shortest symbolic-expression representations of the Boolean functions. Theorem 6.1 says that the shortest strings denoting most functions will have length exponential on the number of Boolean variables  $n$ . For example, assume twenty Boolean variables, with functions represented by bit strings ( $\alpha = 2$ ). All the distinct functions could be represented by distinct bit strings of length 1048576 ( $= 2^{20}$ ). But in any "compact" representation strategy, the fraction of shortest strings with length less than one million bits (about 95% of 1048576) must be less than  $3 \times 10^{-14623}$  ( $= \alpha^{-i} = 2^{-48575}$ )!

The proof of this theorem will not be given here. It is a simple counting argument identical to that used to prove that most bit strings have a high Kolmogorov complexity, a result stated by Martin-Lof [Maa66].

We can try several approaches to keep string length to a minimum. One is to use a notation that is fairly compact for most functions actually computed. We can also attempt simplification to make the strings less complex.

However, testing that an assertion is identically true using a compact expression notation is likely to be np-hard. This claim is motivated by the proof that the satisfiability problem for Boolean functions represented in conjunctive normal form is np-complete. (See for example the text of Aho et al [AH74].) The proof generates a string of polynomial length representing a particular function. Such a string is most likely to exist in expression notations which strive to be compact.

Furthermore, regardless of any clever compact encoding or simplification, Theorem 6.1 ensures that in the worst case, the string length will be exponential on the number of variables  $n$ . This means that the worst case time to check that a function is identically true is an exponential of an exponential of  $n$  (assuming the np-complete problem cannot be solved in less than exponential time)!

A second approach, which may be used in conjunction with the first, is to enforce the use of a canonical representation, which ensures that each function has a unique representation. That is, require relation  $Rep$  in Theorem 6.1 to be a function. The test that two strings denote the same function becomes a simple check that the strings themselves are identical. This takes time proportional to the length of the shorter string, which is exponential on the number of variables  $n$  in the worst case.

An enhancement to this approach is to enter the strings in a hashing table, so that identical strings actually occupy the same storage. The test that two strings denote the same function is then a check that the two strings are in fact the same string, an O(1) pointer comparison.

The hidden cost in the use of a canonical form is the cost of computing the canonical representation for a given function. This is np-hard for compact canonical forms, since the following procedure tests whether an expression is satisfiable, an np-complete problem for compact expression notations:

1. Given an arbitrary expression, compute its canonical representation.
2. Test whether or not the canonical representation is equal to the canonical representation for FALSE.

If a hashing table is used, then strings hashing to the same location must be compared, taking time exponential on  $n$  in the worst case. Then the use of canonical representations and hashing does not remove the np-hardness associated with compact notations or the worst-case exponential length of strings, although it may otherwise improve the performance.

## 6.2. Case Splitting and Case Merging Complexity

The previous section discussed the complexity of Boolean expressions of  $n$  Boolean variables. Now consider a case splitting symbolic execution that uses  $n$  symbolic variables, all Boolean. In the extreme case, this execution could split into  $2^n$  separate paths. In that event, each path corresponds to a distinct assignment of constant values to the symbolic variables. This corresponds to the exhaustive conventional execution of the program for all possible values of the  $n$  variables.

In conventional executions, all variables have constant values. The analysis of the results for each case is then  $O(1)$  (independent of  $n$ ). The total complexity of the exhaustive execution is  $O(2^n)$ .

If an actual symbolic execution split into  $2^n$  paths, the case splitting may require some overhead to derive the constant variable values associated with each path. Alternately, the symbolic execution might split into few paths but generate complex symbolic strings. In these cases, the total computational complexity could exceed  $O(2^n)$ .

Conversely, the case splitting symbolic execution may split into few paths while generating simple symbolic strings. In such cases, the total computational complexity would be less than  $O(2^n)$ .

In case merging symbolic executions, we can use a compact notation in an effort to make "real" strings (strings likely to be generated) be of length polynomial

on  $n$ .

Comparison of these strings will be np-hard, but we can use heuristics that

strive for polynomial performance in "real" problems.

If these heuristics are successful, then the total computational complexity of a

case merging analysis will be a polynomial of  $n$ . Also, if most variables are modified

only in a few basic blocks, then dataflow analysis can substantially simplify the

symbolic expression computation. But in the worst case, the complexity will be an

exponential of an exponential of  $n$ , much worse than the  $O(2^n)$  complexity noted

above for case splitting.

The large possible variations in the computational complexity of case splitting and case merging symbolic analyses prevent us from giving a simple quantitative criterion for preferring one method over the other in specific cases. However, the following sections present two examples which illustrate some possible qualitative criteria.

### 6.3. Blackjack Example

#### 6.3.1. Blackjack Descriptions

**Figure 6-1** gives the specification for the first example, a machine that plays Blackjack by dealer's rules. This description is adapted from a DDL description by Dietmeyer and Dueley [DD75]. Call the description  $P_s$ .

```
(* Blackjack machine specification *)
var   Hit, Broke, Stand, FF: Boolean;
      Score, CardBuf: Integer;
      YCard: Boolean;
      Value: Integer;

begin
  A: Hit := FALSE; Broke := FALSE; Stand := FALSE; YA :=
    FF := FALSE;
  Score := 0;
  B: Hit := TRUE;
  CardBuf := Value;
  if YCard then goto C else goto B;
  C: Hit := FALSE;
  if YCard then goto C else goto D;
  D: Score := Score + CardBuf;
  if (CardBuf ≠ 1) or FF then goto F else goto E;
  E: FF := TRUE;
  F: if Score < 17 then goto B else goto G;
  G: if Score < 22 then goto K else goto H;
  H: CardBuf := -10;
  if FF then begin FF := FALSE; goto D end
  else goto J;
  J: Broke := TRUE;
  if YCard then goto A else goto J;
  K: Stand := TRUE;
  if YCard then goto A else goto K
end; (* Blackjack machine *)
```

**Figure 6-1:** Blackjack finite state machine specification

**Figure 6-1** gives a proposed design,  $P_i$ , for this machine. To obtain  $P_i$ , we

interpret  $P_s$  as a finite state machine, with each *goto* representing a synchronous transfer to the next state. We assume that assignments to *Hit*, *Broke*, and *Stand* occur immediately upon entering the state, while all remaining assignments in the state do not occur until the machine exits the state.  $P_i$  then uses JK flip-flops to encode the state, combining states *F* and *G* into a single state.

#### 6.3.2. Simulation Relation

**Figure 6-3** is the simulation relation which describes how  $P_s$  and  $P_i$  should correspond to each other. For Boolean variables, it uses  $0 \equiv FALSE$  and  $1 \equiv TRUE$ . Subscripts "s" and "i" denote variables in  $P_s$  and  $P_i$ , respectively. To denote that a particular assignment of values to variables satisfies the simulation relation, we write

$$r_s R_{si} r_i$$

where  $r_s$  ( $r_i$ ) is the assignment of values for  $P_s$  ( $P_i$ ). Formally, the simulation relation is  $R_{si}$ , the set of all value assignments  $\langle r_s, r_i \rangle$  that satisfy the condition in Figure 6-3. We will say that a given value assignment  $r_s$  for  $P_s$  is a stopping point if there exists a value assignment  $r_i$  for  $P_i$  such that  $r_s R_{si} r_i$ . The same terminology applies for value assignments  $r_i$  for  $P_i$ .

Now let  $F_{si}$  be the function that describes the action of  $P_s$  mapping one stopping point  $r_s$  to another. That is, initialize  $P_s$  by a value assignment  $r_s$  that is a stopping point. Execute  $P_s$  until it reaches a new stopping point  $s_j$ . Then by definition  $s_j = F_{si}(r_s)$ .  $F_{si}$  is the corresponding function for  $P_i$ .

We now claim that  $P_i$  is a correct implementation of  $P_s$  by virtue of Milner's criterion [Mil71].

```
(* Blackjack machine design *)
var
  Hit, Broke, Stand, Q4, Q5, Q6, FF: Boolean;
  J1, K1, J2, K2, J3, K3: Boolean;
  J4, K4, J5, K5, J6, K6, JFF, KFF: Boolean;
  Score, CardBuf: Integer;
  YCrd: Boolean;
  Value: Integer;
begin
  while TRUE do begin
    J1 := not (Broke or Stand or Q4 or Q6) or
      Q5 and not Q6 and (Score < 17);
    K1 := YCrd;
    J2 := Q4 and not Q5 and not FF;
    K2 := YCrd;
    J3 := Q5 and not Q6 and (17 ≤ Score < 22);
    K3 := YCrd;
    J4 := Q5;
    K4 := not Q5 or Q6 or (Score < 22);
    J5 := Q6 and not YCrd or (Q4 and JFF);
    K5 := Q4 and not Q6;
    J6 := Hit and YCrd or Q4 and not Q5 and FF;
    K6 := ((CardBuf ≠ 1) or F7) and not Q4 and Q5;
    JFF := Q4 and Q6;
    KFF := not (Hit or Stand or Q5 or Q6);
    Score := if not (Hit or Broke or Stand or Q4 or Q6) then 0 else
      if (not Q4 and Q5) then Score + CardBuf else Score;
    CardBuf := if Hit then Value else
      if (Q4 and Q6) then 10 else
        if (Q4 and not Q6) then -10 else CardBuf;
  end;
end; (* Blackjack machine *)
```

Figure 6-2: Blackjack machine implementation

(\* "@" denotes concatenation

 $Lc \equiv$  location counter $State_i \equiv Hit_i @ Broke_i @ Stand_i @ Q4_i @ Q5_i @ Q6_i @$ 

$(Hil_s = Hil_j)$  and  $(Broke_s = Broke_j)$  and  $(Stand_s = Stand_j)$  and  
 $(FF_s = FF_j)$  and  $(Score_s = Score_j)$  and  $(CardBuf_s = CardBuf_j)$  and  
 $(YCrd_s = YCrd_j)$  and  $(Value_s = Value_j)$  and  
 $(LC_i = \text{top})$  and  
 $((LC_s = XA) \text{ and } (State_i = 000000)) \text{ or}$   
 $((LC_s = XB) \text{ and } (State_i = 100000)) \text{ or}$   
 $((LC_s = XC) \text{ and } (State_i = 000001)) \text{ or}$   
 $((LC_s = XD) \text{ and } (State_i = 000011)) \text{ or}$   
 $((LC_s = XE) \text{ and } (State_i = 000111)) \text{ or}$   
 $((LC_s = XF) \text{ and } (State_i = 000100)) \text{ or}$   
 $((LC_s = XH) \text{ and } (State_i = 000100)) \text{ or}$   
 $((LC_s = XK) \text{ and } (State_i = 001000))$

Figure 6-3: Simulation relation for Blackjack descriptions

forall assignments  $r_s$  for  $P_s$  and  $r_i$  for  $P_i$ :

$$(r_s R_{si} r_i) \supset (F_{si}(r_s) R_{si} F_{is}(r_i)) \quad (6.1)$$

The next chapter will consider criteria for correctness in detail; for now, the primary concern is the use of symbolic execution to test Eq. (6.1).

Symbolic execution helps verify Eq. (6.1) by testing many possible cases in parallel. We may use either case splitting or case merging to handle the conditional branches that arise during this verification. Appendix C outlines the symbolic execution of both descriptions using both case splitting and case merging. The

appendix lists the complete results of such executions. We now consider the utility of case merging vs. case splitting in this example.

### 6.3.3. Symbolic Case Splitting Results

Figure 6-4 shows the results of the case splitting symbolic execution for a single case. This path is typical of the seventeen cases. (Results for all cases appear in Appendix C.) Two facts are apparent: (i) The symbolic values of the variables are very simple; in most cases, just symbolic variables or constants. (ii) The test that these results satisfy the simulation relation is trivial.

( $\oplus$  “ $\oplus$ ” denotes concatenation)

<i>State</i> , $\equiv$ <i>Lc</i> , $\oplus$ <i>Hh</i> , $\oplus$ <i>Brake</i> , $\oplus$ <i>Stand</i> ,
<i>State</i> <sub>s</sub> , $\equiv$ <i>Hh</i> , $\oplus$ <i>B.oke</i> , $\oplus$ <i>Stand</i> , $\oplus$ <i>Q4</i> , $\oplus$ <i>Q5</i> , $\oplus$ <i>Q6</i> ,
<i>Pc</i> : ( <i>c</i> <sub>0</sub> $\neq$ 1) $\oplus$ <i>f</i> <sub>0</sub>
Initial: <i>State</i> <sub>s</sub> = <i>D</i> :000 <i>ff</i> <sub>s</sub> = <i>f</i> <sub>0</sub> <i>Score</i> <sub>s</sub> = <i>s</i> <sub>0</sub> <i>CardBuf</i> <sub>s</sub> = <i>c</i> <sub>0</sub>
Final: <i>State</i> <sub>s</sub> = <i>F</i> :000 <i>ff</i> <sub>s</sub> = <i>f</i> <sub>0</sub> <i>Score</i> <sub>s</sub> = <i>s</i> <sub>0</sub> + <i>c</i> <sub>0</sub> <i>CardBuf</i> <sub>s</sub> = <i>c</i> <sub>0</sub>

Initial:   *State*<sub>t</sub> = 000011           *ff*<sub>t</sub> = *f*<sub>0</sub>           *Score*<sub>t</sub> = *s*<sub>0</sub>           *CardBuf*<sub>t</sub> = *c*<sub>0</sub>

Final:    *State*<sub>t</sub> = 000110           *ff*<sub>t</sub> = *f*<sub>0</sub>           *Score*<sub>t</sub> = *s*<sub>0</sub> + *c*<sub>0</sub>   *CardBuf*<sub>t</sub> = *c*<sub>0</sub>

Figure 6-4: Case splitting results for one path

The second observation follows not only because of the first, but also because the initial value assignments *r*<sub>s</sub> and *r*<sub>t</sub> for each path guarantee that we have

$\forall$  value assignments *r*<sub>s</sub> and *r*<sub>t</sub> represented by this symbolic execution: *r*<sub>s</sub>  $R_{st}$  *r*<sub>t</sub>

Therefore, we need not prove Eq. (6.1):

$\forall$  represented value assignments *r*<sub>s</sub> and *r*<sub>t</sub>:

$$(r_s R_{st} r_t) \supseteq (F_{s,t}(r_s) R_{st} F_{t,s}(r_t)) \quad (6.1)$$

Rather, we need only show the simpler condition, Eq. (6.2):

$\forall$  represented value assignments *r*<sub>s</sub> and *r*<sub>t</sub>:

$$F_{s,t}(r_s) R_{st} F_{t,s}(r_t) \quad (6.2)$$

As noted above, this test is easy for the 17 paths in the case splitting symbolic execution.

Overall, we have the following statistics for *P*<sub>s</sub>: *P*<sub>s</sub> has 16 basic blocks. This happens to be the number of labels in the Blackjack description (15) plus the number of *then-else* clauses containing more than just a *test* (one). In the symbolic execution of the 17 paths by case splitting, one block is executed three times, no other block is executed more than twice, and the average number of executions of a basic block is less than 1.4. The simulation state must be saved and later restored eight times due to case splits.

Description *P*<sub>t</sub> has a single basic block. It is executed 17 times, once for each path from *P*<sub>s</sub>.

### 6.3.4. Symbolic Case Merging Results

Now consider the results of the case merging executions. Figure 6-5 shows the initialization of some pertinent variables and the symbolic results for variable *Hh*. (In Appendix C, the program counter of *P*<sub>s</sub> is represented in a distributed fashion by array *L.Pc*.) We assume that the initial value assignments satisfy *R<sub>st</sub>*.

Then, for the final results to satisfy  $R_u$ , the final values of  $Hit_s$  and  $Hit_i$  should be equal.

```

105          106
 $Lc_s = \text{if } ((\rho_0 = XJ) \text{ and } y_0) \text{ or } ((\rho_0 = XK) \text{ and } y_0)$ 
 $\quad \quad \quad \text{if } (\rho_0 = XA) \text{ or } ((\rho_0 = XB) \text{ and not } y_0) \text{ or }$ 
 $\quad \quad \quad (((\rho_0 = F) \text{ and } (s_0 < 17))$ 
 $\quad \quad \quad \text{if } ((\rho_0 = XB) \text{ and } y_0) \text{ or } ((\rho_0 = XC) \text{ and } y_0)$ 
 $\quad \quad \quad \text{if } ((\rho_0 = XC) \text{ and not } y_0) \text{ or } ((\rho_0 = E) \text{ or } ((\rho_0 = H) \text{ and } f_0))$ 
 $\quad \quad \quad \text{if } ((\rho_0 = D) \text{ and } (c_0 = 1) \text{ and not } f_0)$ 
 $\quad \quad \quad \text{if } ((\rho_0 = D) \text{ and } (c_0 \neq 1) \text{ or } f_0)$ 
 $\quad \quad \quad \text{if } ((\rho_0 = F) \text{ and } (s_0 \geq 17) \text{ and } (s_0 \geq 22))$ 
 $\quad \quad \quad \text{if } ((\rho_0 = H) \text{ and not } f_0) \text{ or } ((\rho_0 = XJ) \text{ and not } y_0)$ 
 $\quad \quad \quad \text{if } ((\rho_0 = F) \text{ and } (s_0 \geq 17) \text{ and } (s_0 \geq 22)) \text{ or }$ 
 $\quad \quad \quad ((\rho_0 = XK) \text{ and not } y_0)$ 
 $\quad \quad \quad \text{then } XK$ 

107
 $State_s \equiv Lc_s, @ Broke_s, @ Stand_s,$ 
 $State_i \equiv Hit_i, @ Broke_i, @ Stand_i, @ Q4_i, @ Q5_i, @ Q6_i,$ 
 $Initial: State_s = p_0 : h_0 : b_0 : 0$ 
 $Score_s = s_0$ 
 $YCd_s = y_0$ 
 $CardBuf_s = c_0$ 
 $Final: Hit_s = (\rho_0 = XB) \text{ or } ((\rho_0 = XB) \text{ and not } y_0) \text{ or } ((\rho_0 = F) \text{ and } (s_0 < 17)) \text{ or }$ 
 $(h_0 \text{ and not } b_0 \text{ and } p_0 \in \{XB,XC,XJ,XK\})$ 

```

```

Initial: State_i = h_0 : b_0 : l_0 : \mu_0 : w_0 : x_0      RF_i = f_0      Score_i = s_0
          CardBuf_i = c_0                            YCd_i = y_0

Final: Hit_i = ((not (b_0 or l_0 or u_0 or x_0) or
                  w_0 and not x_0 and (s_0 < 17)) and not h_0) or
                  not y_0 and h_0

```

**Figure 6-5:** Case merging results for one variable

Figure 6-6 shows the final symbolic value of  $Lc_s$  in the case merging symbolic execution of  $P_s$ . We make two observations on the case merging results.

First, the final case merging symbolic expressions are much more complex than those from the case splitting executions. The expressions for  $Hit_i$  are typical.

The final value of  $Lc_s$  is the most complex symbolic expression generated. Note that the condition on  $Lc_s$  in  $R_s$  is also the most complex condition in  $R_{si}$ , so the verification of the implementation of  $Lc_s$  will be especially complex.

105

106

```

 $Lc_s = \text{if } ((\rho_0 = XJ) \text{ and } y_0) \text{ or } ((\rho_0 = XK) \text{ and } y_0)$ 
 $\quad \quad \quad \text{if } (\rho_0 = XA) \text{ or } ((\rho_0 = XB) \text{ and not } y_0) \text{ or }$ 
 $\quad \quad \quad (((\rho_0 = F) \text{ and } (s_0 < 17))$ 
 $\quad \quad \quad \text{if } ((\rho_0 = XB) \text{ and } y_0) \text{ or } ((\rho_0 = XC) \text{ and } y_0)$ 
 $\quad \quad \quad \text{if } ((\rho_0 = XC) \text{ and not } y_0) \text{ or } ((\rho_0 = E) \text{ or } ((\rho_0 = H) \text{ and } f_0))$ 
 $\quad \quad \quad \text{if } ((\rho_0 = D) \text{ and } (c_0 = 1) \text{ and not } f_0)$ 
 $\quad \quad \quad \text{if } ((\rho_0 = D) \text{ and } (c_0 \neq 1) \text{ or } f_0)$ 
 $\quad \quad \quad \text{if } ((\rho_0 = F) \text{ and } (s_0 \geq 17) \text{ and } (s_0 \geq 22))$ 
 $\quad \quad \quad \text{if } ((\rho_0 = H) \text{ and not } f_0) \text{ or } ((\rho_0 = XJ) \text{ and not } y_0)$ 
 $\quad \quad \quad \text{if } ((\rho_0 = F) \text{ and } (s_0 \geq 17) \text{ and } (s_0 \geq 22)) \text{ or }$ 
 $\quad \quad \quad ((\rho_0 = XK) \text{ and not } y_0)$ 
 $\quad \quad \quad \text{then } XK$ 

 $State_s \equiv Lc_s, @ Broke_s, @ Stand_s,$ 
 $State_i \equiv Hit_i, @ Broke_i, @ Stand_i, @ Q4_i, @ Q5_i, @ Q6_i,$ 
 $Initial: State_s = p_0 : h_0 : b_0 : 0$ 
 $Score_s = s_0$ 
 $YCd_s = y_0$ 
 $CardBuf_s = c_0$ 
 $Final: Hit_i = (\rho_0 = XB) \text{ or } ((\rho_0 = XB) \text{ and not } y_0) \text{ or } ((\rho_0 = F) \text{ and } (s_0 < 17)) \text{ or }$ 
 $(h_0 \text{ and not } b_0 \text{ and } p_0 \in \{XB,XC,XJ,XK\})$ 

```

Second, we cannot use Eq. (6.2) to verify these symbolic results. This holds partially because the final value of  $Hit_i$  uses symbolic variables  $u_0$ ,  $w_0$ , and  $x_0$ . These appear only in the execution of  $P_i$  and are related to other variables by the left side of Eq. (6.1). However, this does not completely explain the problem.

$R_{si}$  implies the following equalities for  $u_0$ ,  $w_0$ , and  $x_0$ .

$$\begin{aligned} u_0 &\equiv \rho_0 \in \{E,F,H\} \\ w_0 &\equiv \rho_0 \in \{D,E,F\} \\ x_0 &\equiv \rho_0 \in \{XC,D,F\} \end{aligned}$$

If we substitute these values for  $u_0$ ,  $w_0$ , and  $x_0$ , then the symbolic executions of  $P_s$  and  $P_i$  both use the same set of symbolic variables, yet their equivalence still cannot be proven by Eq. (6.2).

### 6.3.5. Conclusion to First Example

In general, it is not necessarily true that a case splitting simulation will allow the use of Eq. (6.2). Conversely, in the Blackjack example, we could use Eq. (6.2) to verify the case merging results if we substituted expressions for  $h_0$ ,  $b_0$ , and  $t_0$  as well as for  $w_0$ ,  $w_0$ , and  $x_0$ . However, our expectation is that Eq. (6.2) is less likely to be valid in case merging analyses than it is in case splitting analyses. This aggravates the problem of increased expression complexity. For the Blackjack machine, a case splitting symbolic execution and analysis is likely to be less expensive than the corresponding case merging analysis.

What properties of the Blackjack example make case merging undesirable? The outstanding characteristic is the highly detailed simulation relation  $R_{S^*}$ . This is partially responsible for the following properties:

1. A minimum of case splitting occurs in the symbolic execution starting from any stopping point. In only one case does the execution split into as many as three paths (namely, starting from label  $F$ ).
2. Each path is short and simple. Most paths modify few variables.

This in turn simplifies the case splitting symbolic results.

On the other hand, the high level of detail in  $R_{S^*}$  causes the case merging results to be complex, with the computation of complicated conditional expressions.

A dataflow analysis allows most of the symbolic values in the Blackjack machine to be simplified, but the value of  $I.c_s$  remains complex. The condition on  $I.c_s$  in  $R_{S^*}$  is also complex. Finally, Eq. (6.2) is less likely to be valid in the case merging analysis when  $R_{S^*}$  is highly detailed.

### 6.4. Associative Memory Example

The above example suggests that case merging is not beneficial when the simulation relation is detailed, including a stopping point in most of the basic blocks of a description. Conversely, we might expect that case merging may be advantageous if we have a "sparse" simulation relation. This section considers a simple example to illustrate this point.

Figure 6-7 shows a simple associative memory which finds the appropriate address by a linear search. Figure 6-8 shows a similar description which specifies a tree structure for extracting the correct address. The simulation relation in Figure 6-9 defines the correspondence between the two versions. The subscripts "s" and "l" refer to Figures 6-7 and 6-8, respectively.

```
var M: array [0..3] of Integer;
  Data: Integer;
  found: Boolean;
  Addr: 0..3;
begin (* CAM *)
  Found := TRUE;
  Addr := 0;
(**) if M[0] = Data then (* nothing *)
else if M[1] = Data then Addr := 1
else if M[2] = Data then Addr := 2
else if M[3] = Data then Addr := 3
else Found := FALSE;
end; (* CAM *)
```

Figure 6-7: Associative memory description

The comparison of these descriptions by symbolic execution is similar to that for the Blackjack example. One difference is that only the stopping points with  $I.c = \text{top}$  are used as the starting point of a symbolic execution. This is because

```

var M : array [0..3] of Integer;
Data: Integer;
F1, F2, Found: Boolean;
A1, A2, Addr: 0..3;
begin (* CAM *)
  F1 := TRUE;
  (** if M[0] = Data then A1 := 0
    else if M[1] = Data then A1 := 1
    else F1 := FALSE;
  F2 := TRUE;
  A2 := 3;
  (** if M[2] = Data then A2 := 2
    else if M[3] ≠ Data then F2 := FALSE;
    if F1 then Addr := A1 else Addr := A2;
    Found := F1 or F2
  end; (* CAM *)

```

Figure 6-8: Second associative memory description

```

( (I.c = top) and (I.c = top) and
  (Datas = Datat) and (Ms = Mt) ) or
  ( (I.c = bottom) and (I.c = bottom) and
    ( (not Founds and not Foundt) or
      (Founds and Foundt and (Addrs = Addrt)) ) )

```

Figure 6-9: Simulation relation for associative memory

*I.c = bottom* represents the end of the procedure; the procedure is not an infinite loop as is the Blackjack machine. Figure 6-10 shows the results of case splitting simulation for the two descriptions.

In the symbolic execution of Figure 6-7, variables *I.c*, *M*, *Data<sub>s</sub>*, *Found<sub>s</sub>*, and *Addr<sub>s</sub>* have initial values **top**, *m<sub>0</sub>*, *d<sub>0</sub>*, *f<sub>0</sub>*, and *a<sub>0</sub>*, respectively. In Figure 6-8, variables *I.c*, *M*, *Data<sub>t</sub>*, *Found<sub>t</sub>*, *Addr<sub>t</sub>*, *F1*, *F2*, *A1*, and *A2*, have initial values **top**, *m<sub>0</sub>*, *d<sub>0</sub>*, *d<sub>1</sub>*, *f<sub>1</sub>*, *f<sub>2</sub>*, *a<sub>1</sub>*, and *a<sub>2</sub>*. The simulation relation implies the equality of the initial values of *I.c*, *M*, and *Data* in the two simulations.

	110	109	P <sub>C</sub>	Found <sub>t</sub>	Addr <sub>t</sub>	F1	F2	A1	A2	Found <sub>s</sub>	Addr <sub>s</sub>
<i>e<sub>0</sub></i>	<i>e<sub>2</sub></i>		<i>T</i>	0	<i>T</i>	<i>T</i>	0	2	<i>T</i>	0	
<i>e<sub>0</sub></i>	<i>n<sub>2</sub>e<sub>3</sub></i>		<i>T</i>	0	<i>T</i>	<i>T</i>	0	3	<i>T</i>	0	
<i>e<sub>0</sub></i>	<i>n<sub>2</sub>n<sub>3</sub></i>		<i>T</i>	0	<i>T</i>	<i>F</i>	0	3	<i>T</i>	0	
<i>n<sub>0</sub>e<sub>1</sub>e<sub>2</sub></i>			<i>T</i>	1	<i>T</i>	<i>T</i>	1	2	<i>T</i>	1	
<i>n<sub>0</sub>e<sub>1</sub>n<sub>2</sub>e<sub>3</sub></i>			<i>T</i>	1	<i>T</i>	<i>T</i>	1	3	<i>T</i>	1	
<i>n<sub>0</sub>e<sub>1</sub>n<sub>2</sub>n<sub>3</sub></i>			<i>T</i>	1	<i>T</i>	<i>F</i>	1	3	<i>T</i>	1	
<i>n<sub>0</sub>n<sub>1</sub>e<sub>2</sub></i>			<i>T</i>	2	<i>F</i>	<i>T</i>	<i>a<sub>t</sub></i>	2	<i>T</i>	2	
<i>n<sub>0</sub>n<sub>1</sub>n<sub>2</sub>e<sub>3</sub></i>			<i>T</i>	3	<i>F</i>	<i>T</i>	<i>a<sub>t</sub></i>	3	<i>T</i>	3	
<i>n<sub>0</sub>n<sub>1</sub>n<sub>2</sub>n<sub>3</sub></i>			<i>F</i>	3	<i>F</i>	<i>F</i>	<i>a<sub>t</sub></i>	3	<i>F</i>	0	

Figure 6-10: Case splitting symbolic execution results

The first column in Figure 6-10, *P<sub>C</sub>*, gives the final path condition for the nine paths in Figure 6-8. In this column, *c<sub>j</sub>* denotes *m<sub>0</sub>[j]* · *d<sub>0</sub>*, and *n<sub>j</sub>* denotes *m<sub>0</sub>[j] ≠ d<sub>0</sub>*.

The remaining columns specify final symbolic values of variables in the two simulations. The results satisfy the simulation relation in all nine paths. In each case, the test of the simulation relation is trivial; *Found* and *Addr* have constant values.

Figures 6-11 and 6-12 gives the symbolic results of case merging simulations for Figures 6-7 and 6-8. These results also verify that the simulation relation is satisfied.

This example differs from the Blackjack example in that the simulation relation is not detailed. It includes only one stopping point (excluding the exit point) in each program. Program P<sub>1</sub> uses a tree structure that, when generalized, splits the case splitting symbolic execution into a number of paths exponential on the size *n* of *M*, into  $3^{n/2}$  paths.

$\begin{aligned} \text{Found}_s &= e_0 \text{ or } n_0e_1 \text{ or } n_0n_1e_2 \text{ or } n_0n_1n_2e_3 \\ &= e_0 \text{ or } e_1 \text{ or } e_2 \text{ or } e_3; \\ \text{Addr}_s &= \text{if } e_0 \text{ then 0 else} \\ &\quad \text{if } n_0e_1 \text{ then 1 else} \\ &\quad \text{if } n_0n_1e_2 \text{ then 2 else} \\ &\quad \text{if } n_0n_1n_2e_3 \text{ then 3 else} \\ &\quad \text{if } n_0n_1n_2n_3 \text{ then 0}; \end{aligned}$
<p>Figure 6-11: Case merging results for Figure 6-7</p>

By contrast, with case merging we could use a simple canonical representation in which the expression complexity grows linearly as  $n$ . The total cost of the case merging analysis may be polynomial, depending on the cost of simplification. Also, Eq. (6.2) is valid with case merging here, even though this was not true for the Blackjack example.

This example does not represent a realistic application of symbolic execution. Due to the regular structure in both Figures 6-7 and 6-8, a program proof by induction on  $n$  would be preferable to a proof by symbolic execution. Nonetheless, this example illustrates characteristics which motivate a preference for case merging over case splitting: a simulation relation with little detail (few stopping points), in which the number of paths grows exponentially while the symbolic expression complexity grows linearly.

## 6.5. Conclusion

This chapter has presented two examples to illustrate the issues in the comparison of case splitting and case merging symbolic executions. The computational complexity of these two approaches varies over such a wide range

111      112

$\begin{aligned} F1_i &= \text{not}(n_0 \text{ and } n_1) = e_0 \text{ or } e_1; \\ AI_i &= \text{if } e_0 \text{ then 0 else} \\ &\quad \text{if } n_0e_1 \text{ then 1 else} \\ &\quad \text{if } n_0n_1 \text{ then } a_i; \\ F2_i &= \text{not}(n_2 \text{ and } n_3) = e_2 \text{ or } e_3; \\ A2_i &= \text{if } e_2 \text{ then 2 else} \\ &\quad \text{if } n_2 \text{ then 3;} \\ Addr_i &= \text{if } e_0 \text{ or } e_1 \text{ then (if } e_0 \text{ then 0 else} \\ &\quad \text{if } n_0e_1 \text{ then 1 else} \\ &\quad \text{if } n_0n_1 \text{ then } a_i \text{ else} \\ &\quad \text{if } n_0n_1e_2 \text{ then 2 else} \\ &\quad \text{if } n_2 \text{ then 3}) \\ &= \text{if } (e_0 \text{ or } e_1) \text{ and } e_0 \text{ then 0 else} \\ &\quad \text{if } (e_0 \text{ or } e_1) \text{ and } n_0e_1 \text{ then 1 else} \\ &\quad \text{if } (e_0 \text{ or } e_1) \text{ and } n_0n_1 \text{ then } a_i \text{ else} \\ &\quad \text{if } n_0n_1e_2 \text{ then 2 else} \\ &\quad \text{if } n_0n_1n_2 \text{ then 3} \\ &= \text{if } e_0 \text{ then 0 else} \\ &\quad \text{if } n_0e_1 \text{ then 1 else} \\ &\quad \text{if } n_0n_1e_2 \text{ then 2 else} \\ &\quad \text{if } n_0n_1n_2 \text{ then 3;} \\ \text{Found}_i &= e_0 \text{ or } e_1 \text{ or } e_2 \text{ or } e_3; \end{aligned}$
--

Figure 6-12: Case merging results for Figure 6-8

that we cannot make general quantitative statements about the superiority of one approach or the other. However, the examples illustrated qualitative properties that might suggest a preference in certain cases.

In the Blackjack example, the simulation relation was highly detailed, including many different stopping points in  $P_s$ . Case splitting symbolic executions

113                    114

split into few paths and allowed the  $t \cdot e$  of Eq. (6.2). But with case merging, the detailed simulation relation required the proof of Eq. (6.1) rather than the simpler Eq. (6.2). The symbolic expression complexity for  $L_C$ , and  $L_{C_s}$  was linear on the number of paths in the case merging analysis. Case splitting was preferred in this case.

In the associative memory example, the simulation relation had little detail, including only one stopping point (excluding the exit point). The case splitting executions split into an exponential number of paths. The case merging execution results grew in complexity as the log of the number of paths. Also, Eq. (6.2) was valid in this case merging analysis. For this example, case merging was preferable.

Implementation and experimentation with high-quality case splitting and case merging symbolic execution systems is desirable to gain more insight into the relative strengths of the two approaches, particularly when applied to event-driven simulations (an application not considered in these examples). In addition to the above considerations, case merging might be effective if we can guarantee an acceptable upper bound on the complexity of the symbolic expressions generated, for example by limiting the number of symbolic inputs in the program executions.

## Chapter 7

### Criteria for Consistency

#### 7.1. Overview

A criterion for consistency is a test by which we determine whether a proposed hardware design is a correct implementation of the given specification. The design and specification are expressed by hardware descriptions.

We will call hardware descriptions "programs" and represent them by graphs. Note however that these "programs" represent hardware, not software, and this may affect verification. For example, Milner was concerned with software, and some of his work stressed the final state of a program. But a correct hardware description might never terminate, so that a "final state" is of no importance.

Milner's criterion for consistency uses a restrictive program model which is not conveniently applicable for hardware design verification (Section 7.3). Brand's criterion verifies clearly incorrect designs (Section 7.4). Two issues in the attempt to improve these criteria are

1. characteristics of the program model (Section 7.5), and
2. intuitive understanding of correctness (Sections 7.6 - 7.8)

A single, simple criterion can be found which is appropriate for a variety of program

models (Section 7.8). This test captures intuitive notions of correctness, but is tedious to check by symbolic execution. A second criterion is easier to test, but more complex than the first (Section 7.9). Given two deterministic programs which satisfy the second definition of consistency, we can show that they also satisfy the first criterion under a modified (but still "reasonable") simulation relation (Section 7.10).

It is desirable to apply the criterion for consistency in a hierarchical, partitioned fashion. Hierarchical verification using the criterion from Section 7.8 is possible (Section 7.11). We know how to allow partitioned verification (1) only for an inappropriate criterion for consistency, (2) only for a nondeterministic program model, and (3) only assuming a property of the program model which is unreasonable for some hardware models (Brand's technical report [B78]).

A gap remains which must be filled to create a unified theory for the criterion for consistency of programs. Section 7.9 presents an extended criterion which may be checked by symbolic execution. However, we have proven hierarchical, partitioned verification only for other definitions. Two approaches for closing this gap are the following:

- Prove hierarchical, partitioned verification for the extended criterion.
- Prove partitioned verification for the simpler criterion from Section 7.8.

Then prove a correspondence between this definition and the extended definition. The results for hierarchical, partitioned verification with the simple criterion may also hold for the extended criterion, according to this correspondence.

Section 7.10 shows a correspondence between the two definitions, but only for deterministic program models; Brand demonstrated partitioned verification only for nondeterministic models.

## 7.2. Notation

This chapter will use graph models to represent programs. The discussion will use terms from graph theory. A "vertex" represents a program state, while an edge represents an action of the program in changing the program state.

When used as subscripts, the letters "s", "i", and "j" will denote particular programs. In the first parts of the chapter, we use only "s" and "i", denoting **specification and implementation**. Later we add "j", and the three letters denote **specification, intermediate, and low level** in a hierarchical design.

When used as subscripts, "a" and "b" will denote unspecified programs. That is, program *a* might be program *s*, *i*, or *j*.

Appendix A describes notational conventions used here for relations.

## 7.3. Simulation between Deterministic Programs

This section reviews Milner's discussion of algebraic simulation between programs [Mil71]. We will omit certain details not of concern here.

Milner uses an abstract graph-like representation for a program. The set  $D_a$  (a vertex set) is the set of all possible states in the program. A state is an assignment of values to the program variables, including the contents of the stack and the location

counter. Relation  $F_a \subseteq D_a \times D_a$  (an edge set) specifies how the program maps any state into a successor state. Milner requires that the program be deterministic; that is, that  $F_a$  be a function. A program  $P_a$  (a graph) is then the ordered pair  $\langle D_a, F_a \rangle$ .

Figure 7-1 illustrates the verification problem. Programs  $P_s = \langle D_s, F_s \rangle$  and  $P_i = \langle D_i, F_i \rangle$  are proposed specification and implementation programs. A relation  $R_{ab} \subseteq D_a \times D_b$  specifies how vertices in  $D_b$  should correspond to vertices in  $D_a$ . In particular,  $R_{si} \subseteq D_s \times D_i$ , and  $R_{is} = R_{si}^{-1}$ . Relation  $R_{si}$  is usually called the "simulation relation." By definition,  $P_i$  is consistent with  $P_s$  under relation  $R_{si}$  if and only if

$$\forall \langle s, u \rangle \in D_s \times D_i \mid s R_{si} u : F_s(s) R_{si} F_i(u)$$

Milner expresses this, " $R_{si}$  is a simulation of  $P_s$  by  $P_i$ " emphasizing the importance of relation  $R_{si}$  in the definition. Since  $F_s$  and  $F_i$  are functions, this criterion is equivalent to both of the following:

$$R_{si} F_i \subseteq F_s R_{si} \quad (7.1)$$

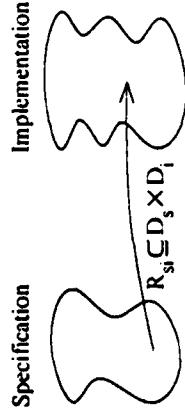
$$R_{is} F_s \subseteq F_i R_{is} \quad (7.2)$$

Note that  $F_s$  is required to be defined for all vertices in  $D_s$ . This disallows don't-care conditions that could lead to more optimal designs.

## 7.4. Simulation between Nondeterministic Programs

Brand modifies Milner's development in three ways [Br78]. First, Brand allows a program  $P_a$  to be nondeterministic; relation  $F_a$  need not be a function.

Second, given  $P_a$ ,  $P_b$ , and  $R_{ab}$ , he defines new program  $P_{ab}$ . Note that  $P_{ab}$  is a version of  $P_a$  modified using  $R_{ba}$  ( $= R_{ab}^{-1}$ ):



**Figure 7-1:** Functional verification problem

$$\begin{aligned} D_{a/b} &= R_{ba}(D_b) \\ F_{a/b} &= \{(s,i) \in D_{a/b} \times D_{a/b} \mid (\exists x_0, x_1, \dots, x_n \in D_a \text{ } (n \geq 1)) \mid \end{aligned}$$

$s = x_0 F_a x_1 F_a \dots F_a x_n = t$ , and

of the several  $x_i$ , only  $x_0$  and  $x_n$  are in  $D_{a/b}\}$

$$P_{a/b} = \langle D_{a/b}, F_{a/b} \rangle \quad (7.3)$$

This allows a coarse definition of relation  $R_{ab}$ . Third, he adds closure to Eq. (7.1) to obtain a new criterion for consistency using  $P_{si}$  and  $P_{is}$ :

$$R_{si} F_{i/s} \subseteq F_{si} * R_{si} \quad (7.4)$$

This development leading to Eq. (7.4) is slightly different from Bland's, but the effective result is the same.

Eq. (7.4) has the unacceptable property of "verifying" the correctness of incorrect implementations, regardless of the choice of  $R_{si}$ . In fact, we now show that if program  $P_s$  is strongly connected, then an arbitrary implementation  $P_i$  satisfies the above criterion.

**PROOF:** Consider an arbitrary pair  $\langle s, v \rangle \in R_{si} F_{i/s}$ . It will suffice to show that this pair is in  $F_{si} * R_{si}$ . From the definition of  $D_{si}$ ,  $D_{i/s}$ , and  $F_{i/s}$  in Eq. (7.3),

must be in  $D_{si}$  and  $v$  in  $D_{i/s}$ . Therefore there exists a vertex  $t$  in  $D_{si}$  such that  $t R_{si} v$ . If  $P_s$  is strongly connected, then by the construction of  $D_{si}$  and  $F_{si}$  using Eq. (7.3), program  $P_{si} = \langle D_{si}, F_{si} \rangle$  is also strongly connected. Hence  $s F_{si} * t$ . Now  $s F_{si} * t R_{si} v$ , or  $\langle s, v \rangle \in F_{si} * R_{si}$ . QED.

In view of the comments on the modeling of changing inputs in Section 7.5 (later), it is unreasonable to expect that  $P_s$  will be strongly connected. However, the above proof is instructive. For every path in  $P_i$ , Eq. (7.4) ensures only that a path connecting corresponding points exists in  $P_s$ . It proves nothing about desirable properties of the path itself, that the path passes through appropriate intermediate states.

It is not surprising that some incorrect implementations can be shown correct for some choices of  $R_{si}$ . By both Eqs. (7.1) and (7.4), every implementation is consistent with the specification under the empty relation. However, the above result says that every (incorrect) design may satisfy the criterion for consistency for every choice of  $R_{si}$ .

This problem arises because of the presence of closure in Eq. (7.4). The justification for closure is as follows: Consider states  $s_n$  and  $t$  in  $D_{si}$ , and  $u$  and  $v$  in  $D_{i/s}$  such that (Figure 7-2)

- $s_n R_{si} u F_{i/s} v$  and
- $s_n F_{si} / R_{si} v$ .

Suppose state  $s_n$  is inside a tight loop not including state  $t$ , while  $u$  is not in any such loop. This might occur because the implementation  $P_i$  uses a more suitable operator so that the loop is not necessary. Then all the corresponding states  $s_j$  for the same

location in every pass through the loop are also related to  $u$  by  $R_{s_i}$ . For each such state we have  $s_i R_{s_i} u F_{s_i} v$ . Figure 7-2 illustrates the situation. The relation  $s_i F_{s_i} R_{s_i} v$  does not hold, although  $s_i F_{s_i}^* R_{s_i} v$  is true. The removal of closure from the right hand side of Eq. (7.4) prevents the verification of a possibly correct implementation in this case. The argument for closure in the left hand side is similar.

121

state we have  $s_i R_{s_i} u F_{s_i} v$ . Figure 7-2 illustrates the situation. The relation  $s_i F_{s_i} R_{s_i} v$  does not hold, although  $s_i F_{s_i}^* R_{s_i} v$  is true. The removal of closure from the right hand side of Eq. (7.4) prevents the verification of a possibly correct implementation in this case. The argument for closure in the left hand side is similar.

122

### 7.5. Program Models

Milner and Brand used different program models and proposed different criteria for consistency. The two approaches have different shortcomings. The next few sections will take a more general look at possible types of program models and consistency criteria.

This section identifies several types of program models which might be used to describe hardware. A later section will consider definitions of consistency for each of these models.

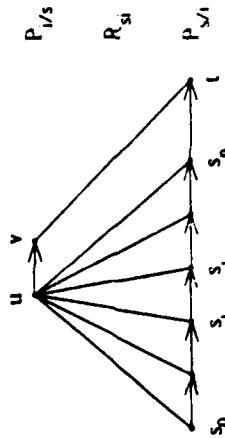


Figure 7-2: Example motivating use of closure

This argument assumes that states in the relation  $R_{s_i}$  are identified solely by their location counters, so that all states sharing the same location counter are related to the same states by  $R_{s_i}$ . This assumption is not necessary. The definition of states in the relation  $R_{s_i}$  may include values of variables (including the stack) as well as the location counter. Closure is not required for this purpose.

This refutation is technically correct, but ignores a practical reason for desiring that states in relation  $R_{s_i}$  be identified solely by their location counters. The practical consideration involves the use of symbolic simulation to compare program models. Section 7.9 will discuss this question in detail.

A hardware component generally reads its inputs many times during its normal operation. These inputs usually are allowed to change many times. Therefore, if a program model is to be used to represent hardware, then it is not reasonable to require that all inputs be presented before the start of execution.

On first inspection, the use of a deterministic model appears to prohibit the modeling of a program with changing inputs. If inputs may change, then from the program's point of view, these changes are nondeterministic. It seems that the model for this program must be nondeterministic. However, this is true only if the model explicitly includes the random changing of inputs. A model might include only the actions taken by the program itself, not the random changes from the

program's environment. In fact, we will always assume that the graph does not model changes to inputs by the environment. Then both deterministic and nondeterministic models can be used for modeling hardware.

Milner did not allow don't-cares in the deterministic program model; for every state, a next state had to be defined. Brand did not explicitly support don't-cares, although in a nondeterministic model, one can indicate a don't-care condition for state  $s$  by specifying every state as a possible successor to  $s$ . Depending on one's design philosophy, the use of don't-cares may or may not be desirable.

In a program graph model  $P_a$ , a vertex is a don't-care vertex if it has no successors. In that case, we assume that any vertex is allowed to be the successor. Any vertex which has one or more successors explicitly specified in the model is not a don't-care vertex; its next vertex must be one of the defined successor vertices.  $D_a^*$  is the set of don't-care vertices in  $D_a$ :

$$D_a^* = \{ \text{vertices } s \in D_a \mid \neg(\exists t, s F_a t) \} \quad (7.5)$$

This definition breaks down for the program  $P_{a/b}$  constructed from Eq. (7.3). It is possible for a vertex  $s$  in  $D_{a/b}$  to have no successor in  $P_{a/b}$  although it does not lead to a don't-care vertex in  $P_a$ . For example, the only path from vertex  $s$  in  $P_a$  could be  $s F_a t F_a t$ , where  $t$  is not in  $D_{a/b}$ . In this case,  $P_{a/b}$  fails to properly characterize the behavior of  $P_a$  at  $s$ . We can eliminate this anomaly by the following restriction.

**Requirement 1:** Every path in  $P_a$  leading from any vertex in  $D_{a/b}$  must eventually reach either a vertex in  $D_{a/b}$  or a don't-care vertex.  $\square$

We also assume the state space  $D_a$  is finite. Then Requirement 1 is satisfied if and

only if every loop in  $P_a$  that is reachable from a vertex in  $D_{a/b}$  contains a vertex in  $D_{a/b}$ .

Even with this restriction, a possible anomaly still exists for nondeterministic models. Given vertex  $s$  in  $D_{a/b}$ , if a path in  $P_a$  from  $s$  reaches a don't-care vertex without passing through intervening vertices in  $D_{a/b}$ , then  $s$  should be considered a don't-care vertex in  $P_{a/b}$ . However, if other paths from  $s$  reach vertices in  $D_{a/b}$ , then  $s$  will have explicitly specified successors in  $P_{a/b}$  (by Eq. (7.3)), and so will not be recognized as a don't-care vertex. We can handle this problem by modifying Eq. (7.3):

$$\begin{aligned} D_{a/b} &= R_{ba}(D_b) \\ F_{a/b} &= \{ \langle s, t \rangle \in D_{a/b} \times D_{a/b} \mid \\ &\quad (\exists x_0, x_1, \dots, x_n \in D_a \mid n \geq 1) \mid \\ &\quad s = x_0 F_a x_1 F_a \dots F_a x_n = t, \text{ and} \\ &\quad \text{of the several } x_i, \text{ only } x_0 \text{ and } x_n \text{ are in } D_{a/b} \text{ and} \\ &\quad (\text{there do not exist } y_0, y_1, \dots, y_n \in D_a \mid n \geq 1) \mid \\ &\quad s = y_0 F_a y_1 F_a \dots F_a y_n, \\ &\quad y_n \text{ is a don't-care state, and} \\ &\quad \text{of the several } y_i, \text{ only } y_0 \text{ is in } D_{a/b} \} \end{aligned} \quad (7.6)$$

Given Eq. (7.6) and Requirement 1 above, then vertex  $s$  is a don't-care vertex in  $P_{a/b}$  if and only if  $s F_a t$ , where  $t$  is a don't-care vertex (possibly  $s$ ) in  $P_a$  and some path  $s F_a t$  includes no vertices in  $D_{a/b}$  except  $s$ . The definition of don't-care vertices is now the same for both  $P_a$  and  $P_{a/b}$ . A don't-care vertex is a vertex with no explicitly declared successors.

We will consider consistency criteria for four types of program models:

deterministic with don't-cares, deterministic without don't-cares, nondeterministic with don't-cares, and nondeterministic without don't-cares.

## 7.6. Consistency Criteria

Eqs. (7.1) and (7.2) may be applied to programs  $P_{s,i}$  and  $P_{i,s}$  defined by Eq. (7.6):

$$R_s F_{i,s} \subseteq F_{s,i} R_s \quad (7.7)$$

$$R_s F_{s,i} \subseteq F_{i,s} R_s \quad (7.8)$$

Eq. (7.7) specifies that the existence of an edge in the implementation  $P_{i,s}$  implies the existence of a corresponding edge in the specification  $P_{s,i}$ . We will call criteria of this type "implied specifications." Eq. (7.8) specifies that the existence of an edge in the specification implies a corresponding edge in the implementation. Such criteria are "implied implementations."

The next section will consider the choice between an implied specification and an implied implementation. This section considers the selection of one implied specification from several simple alternatives. An analogous discussion applies to implied implementations.

Four simple implied specifications are as follows:

$$R_s F_{i,s} R_{i,s} \subseteq F_{s,i} \quad (7.9)$$

$$R_s F_{i,s} \subseteq F_{s,i} R_{i,s} \quad (7.7)$$

$$F_{i,s} R_{i,s} \subseteq R_s F_{s,i} \quad (7.10)$$

$$F_{i,s} \subseteq R_s F_{s,i} R_{i,s} \quad (7.11)$$

Consider an edge  $\langle u, v \rangle$  in  $P_{i,s}$ . Eq. (7.9) requires that any vertex related to  $v$  by  $R_{i,s}$

must be a successor of every vertex related to  $u$  by  $R_s$ . This is clearly unreasonable; Eq. (7.9) will not be considered further.

Given the edge  $\langle u, v \rangle$  in  $P_{i,s}$ : Eq. (7.7) says that there is an edge  $\langle s, i \rangle$  in  $F_{s,i}$  corresponding to  $\langle u, v \rangle$  for every vertex  $s$  in  $D_{s,i}$  related to  $u$  by  $R_{s,i}$ . Eqs. (7.10) and (7.11) guarantee only that there is an edge  $\langle s, i \rangle$  corresponding to  $\langle u, v \rangle$  for some vertex  $s$  in  $D_{s,i}$  related to  $u$  by  $R_{s,i}$ . The former requirement is preferable. It guarantees that the action of the implementation at state  $u$  is correct regardless of how it happens to be related to the specification during a particular execution of programs  $P_s$  and  $P_i$ .

One might argue that a corresponding edge  $\langle s, j \rangle$  need not exist for a particular vertex  $s_j$  related to  $u$  by  $R_{s,i}$ , because we happen to know that  $P_s$  will never be at  $s_j$  when  $P_i$  is at  $u$ . But if this is the case, then  $\langle s, u \rangle$  should not be in  $R_{s,i}$ . Alternatively, we might happen to know that when  $P_s$  is at  $s_j$ , then we do not care what the next state is. But in this case, the problem is that the specification need not be implied by the implementation at this point. This consideration argues for the choice between Eqs. (7.7) and (7.8), not the choice between (7.7), (7.10), and (7.11).

Eq. (7.10) guarantees that there is an edge  $\langle s, i \rangle$  in  $F_{s,i}$  corresponding to  $\langle u, v \rangle$  for every vertex  $i$  in  $D_{s,i}$  related to  $v$  by  $R_{s,i}$ . This test is of little value. Given two corresponding states  $i$   $R_{s,i}$   $v$ , we are concerned not with how the programs  $P_s$  and  $P_i$  came to be in those states, but rather with the consistency of the future actions of the programs.

The above considerations motivate the use of Eq. (7.7) rather than (7.9).

(7.10), or (7.11). This confirms the general form of the criteria for consistency used by both Milner and Brand.

### 7.7. Implied Implementation vs. Implied Specification

This section considers four program models, determining whether Eq. (7.7) or (7.8) is preferable for each case.

#### 7.7.1. Deterministic Programs without Don't-Cares

If  $P_s$  and  $P_i$  are deterministic with no don't-cares, then Eq. (7.6) and Requirement 1 from Section 7.5 guarantee that  $P_{s,i}$  and  $P_{i,s}$  will also be deterministic with no don't-cares. Milner noted that Eqs. (7.7) and (7.8) are exactly equivalent for such programs. The choice of one or the other is arbitrary.

#### 7.7.2. Deterministic Programs with Don't-Cares

Three issues arise in defining consistency for deterministic programs with don't-cares.

1. The specification should be completely implemented. This is ensured by Eq. (7.8) but not by Eq. (7.7). Brand claimed that one could overcome this problem with Eq. (7.7) by providing a sufficiently detailed relation  $R_s$  [Br78]. However, his discussion assumed that there were no don't-cares in the implementation. If this is not the case, then the detailed relation does not solve the problem.

2. Since don't-cares are allowed in the specification, we should not require that every action of the implementation be explicitly given in the specification. Eq. (7.7) makes this requirement, but Eq. (7.8) does not. The following example illustrates this point.

Figure 7-3 shows a graph representation of a specification program  $P_{s,i}$ . The letter inside each vertex is a name for the program state. The notation "i/xx" gives the system inputs and outputs associated with each state; "i" is the Boolean input while "xx" are the Boolean outputs. For example, the values 0/11 on state  $i$  indicates that the system enters state  $i$  from states  $e$  or  $h$  when the input is 0. (This is different from typical finite state machine graphs where the I/O designations are attached to edges, not vertices.)

Asterisks separate pairs of states which are identical except for the value of the input. Since the graph only models the behavior of the program itself, the edges ( $F_{s,i}$ ) connect only states with identical inputs. However, in an actual execution, the environment may cause transfers between states separated by asterisks, for example from  $g$  to  $h$ . Hence, the model permits changing inputs, but  $F_{s,i}$  does not explicitly represent the effects of changing inputs.

The designer may combine states  $j$  and  $i$  in the implementation, yielding program  $P_{i,s}$  shown in Figure 7-4. The relations  $F_{s,i}$  and  $F_{i,s}$  appear in Figure 7-5, along with relation  $R_s \subseteq D_{s,i} \times D_{i,s}$ .

The four relations  $(R_s, F_{i,s})$ ,  $(F_{s,i}, R_s)$ ,  $(R_s, F_{s,i})$ , and  $(F_{i,s}, R_s)$  appear in Figure 7-6. The relations satisfy Eq. (7.8), but not Eq. (7.7). In general, Eq. (7.7) does not allow incomplete specifications as in Figure 7-3. This makes optimizations awkward, requiring the modification of  $P_{s,i}$  to include the optimizations performed.

3. The first two considerations motivate a preference for Eq. (7.8). The third

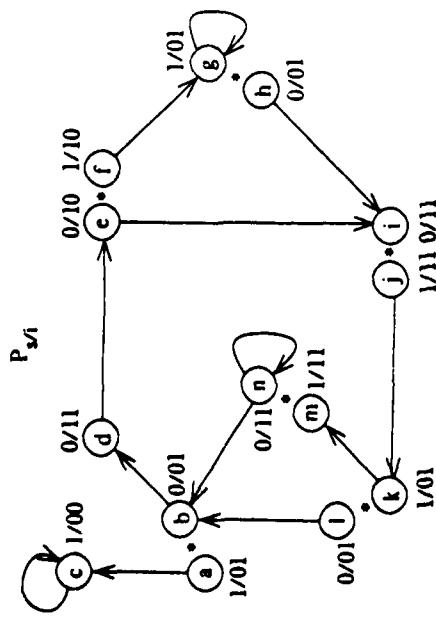


Figure 7-3: Example - Specification program

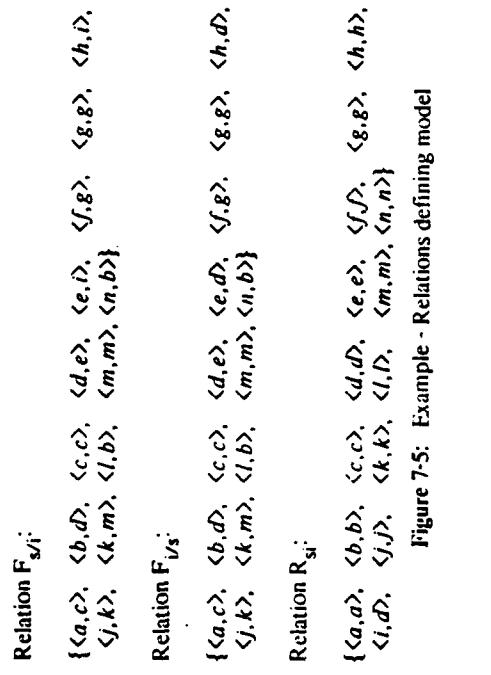


Figure 7-5: Example - Relations defining model

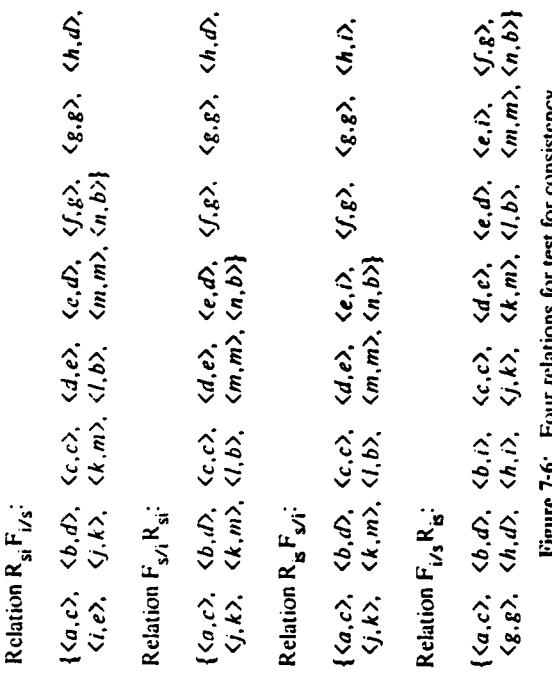


Figure 7-6: Four relations for test for consistency

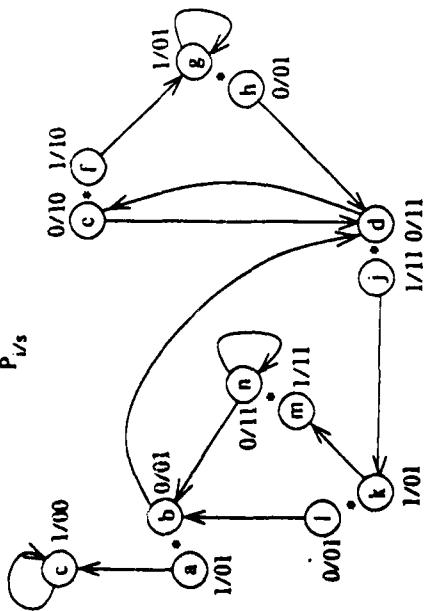


Figure 7-4: Example - Implementation program

consideration is that any action taken by the implementation should correspond either to an action declared in the specification or to a don't-care condition in the specification. Eq. (7.7) clearly satisfies this requirement, but Eq. (7.8) is also acceptable.

Assume we have deterministic programs  $P_{s,i}$  and  $P_{i,s}$ , and relation  $R_{is}$  satisfying Eq. (7.8). Suppose some action  $u F_{i,s} v$  performed by  $P_{i,s}$  corresponds to neither an edge nor a don't-care condition in the specification. Figure 7-7 shows the problem. Vertex  $u$  is related to some vertex  $s$  by  $R_{is}$ . Since  $u$  is not related to a don't-care vertex, there exists an edge  $\langle s, t \rangle$  in  $F_{s,i}$ . Since  $\langle u, v \rangle$  does not correspond to  $\langle s, t \rangle$ , we have  $v \notin R_{si}(t)$ . ( $R_{si}(t)$  is the set of states  $w$  such that  $t R_{si} w$ )

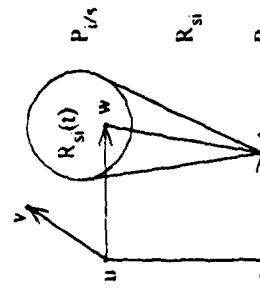


Figure 7-7: Incorrect programs satisfying Eq. (7.8)

We have  $u R_{is} s F_{si} t$ . Therefore, from Eq. (7.8) we also have  $u F_{i,s} w R_{is} t$  for some state  $w$  in  $R_{si}(t)$ . Now  $P_{i,s}$  specifies both  $u F_{i,s} w$  and  $u F_{i,s} v$ , so  $P_{i,s}$  is nondeterministic. This contradicts the condition that the programs are deterministic. Therefore, Eq. (7.8) guarantees the third desirable property for deterministic programs with don't-cares.

### 7.7.3. Nondeterministic Programs without Don't-Cares

Again, three issues arise in defining consistency for nondeterministic programs without don't-cares.

1. For deterministic models, it is desirable that the specification be completely implemented. For nondeterministic models, this is not necessary. Suppose a specification indicates three possible successor states for state  $s$ . An implementation which always chooses the same one of those three states would be correct, even though it does not completely implement the specification. We do not require that each possible successor is sometimes the actual successor. We care only that the actual successor is one of the specified possible successors.

If for some case we do wish to require that each possible successor is sometimes the actual successor, then we can introduce a dummy input variable. We then choose the successor deterministically from the input value. Brand made this point in his technical report [Br78]. Therefore, the complete implementation of the specification is not an issue for nondeterministic models.

2. Even though the complete implementation of all edges in the specification  $F_{si}$  is not necessary, we should implement at least one edge incident out of each vertex in  $D_{si}$ . Eq. (7.8) obviously guarantees this by requiring the implementation of every edge. However, Eq. (7.7) is also satisfactory.

Consider arbitrary vertex  $s$  in  $D_{si}$ . It is related to some vertex  $u$  in  $D_{is}$  by  $R_{is}$ . Since there are no don't-cares,  $u F_{i,s} v$  for some vertex  $v$ . Then  $s R_{si} F_{si} v$ . If Eq. (7.7) is satisfied, then  $s F_{si} t R_{si} v$  for some vertex  $t$ . Therefore, one edge incident out of  $s$  is implemented.

3. When don't-cares were allowed, we did not require every edge in  $F_{s/s}$  to be explicitly mirrored in  $F_{s/I}$ . However, when don't-cares are not allowed, this requirement is necessary. Eq. (7.7) enforces this requirement, while Eq. (7.8) fails to do so, as noted by Brand [B78].

The first two points above do not discriminate between Eqs. (7.7) and (7.8). The third point motivates the preference for Eq. (7.7), the implied specification, for nondeterministic program models without don't-cares.

#### 7.7.4. Nondeterministic Programs with Don't-Cares

A nondeterministic program with don't-cares can be converted to a nondeterministic program without don't-cares. This is done by specifying every vertex in  $D_{s/b}$  as a possible successor to every don't-care vertex in  $D_{s/b}$ . The discussion for nondeterministic programs without don't-cares then applies, and Eq. (7.7) is preferred.

The addition of these edges can drastically increase the size of the edge set. We can replace the test of Eq. (7.7) against the programs without don't-care vertices by an equivalent test against the program with don't-cares.

Let  $P_{s/I} = \langle D_{s/I}, F_{s/I} \rangle$  and  $P_{s/s} = \langle D_{s/s}, F_{s/s} \rangle$  be nondeterministic programs with don't-cares. Also let  $D_{s/I}^*$  and  $D_{s/s}^*$  be the don't-care vertices in  $P_{s/I}$  and  $P_{s/s}$ . Then the application of Eq. (7.7) to the programs with don't-care edges added is described by

$$R_s(F_{s/s} \cup (D_{s/s}^* \times D_{s/s})) \subseteq (F_{s/I} \cup (D_{s/I}^* \times D_{s/I})) R_s \quad (7.12)$$

We will assume that no one vertex  $s$  in  $D_{s/I}$  has immediate successors  $F_{s/I}(\{s\})$  related by  $R_s$  to all vertices in  $D_{s/s}$ . This is expressed by

$$\text{not } (\exists s \in D_{s/I} | R_s(F_{s/I}(\{s\})) = D_{s/s}^*) \quad (7.13)$$

Note that this condition applies *before* the addition of don't-care edges. Under this condition, Eq. (7.12) is equivalent to the following tests on  $P_{s/I}$ ,  $P_{s/s}$ , and  $R_s$ :

$$R_s F_{s/s} \subseteq (F_{s/I} R_s) \cup (D_{s/I}^* \times D_{s/I}) \quad (7.14)$$

$$R_s (D_{s/s}^*) \subseteq D_{s/I}^* \quad (7.15)$$

Eqs. (7.14) and (7.15) are preferable to Eq. (7.12) because the left hand side of Eq. (7.14) is generally much smaller than the left hand side of Eq. (7.12). If Eq. (7.13) does not hold, then Eqs. (7.14) and (7.15) are stronger than Eq. (7.12); under no condition are they weaker.

#### 7.8. Deterministic Model Reconsidered

Section 7.7.2 shows that Eq. (7.8) is appropriate for deterministic models with don't-cares. Sections 7.7.3 and 7.7.4 show that Eq. (7.7) is better for nondeterministic models, with or without don't-cares. This seems paradoxical, because a deterministic model is just a special case of a nondeterministic model, and the same type of test should be suitable for both.

To resolve the paradox, note that in Section 7.7.4, we had to modify either the programs, adding don't-care edges, or the consistency criterion, using Eqs. (7.14) and (7.15) instead of (7.7). We may conjecture that the application of Eqs. (7.14) and (7.15) is equivalent to the application of Eq. (7.8) for deterministic models with don't-cares. This section proves such a conjecture.

The following lemmas and theorem assume we are given programs  $P_s$  and  $P_i$  and relation  $R_s \subseteq D_s \times D_i$ . Assume that  $F_s$  and  $F_i$  are deterministic with don't-

cares allowed. That is, each vertex in  $D_s$  and  $D_t$  has exactly zero or one successor. Define  $P_{s,t}$  and  $P_{t,s}$  by Eq. (7.6). Finally,  $D_{s,t}^x \subseteq D_{s,t}$  and  $D_{t,s}^x \subseteq D_{t,s}$  are the sets of vertices that have no successors in  $P_{s,t}$  and  $P_{t,s}$ ; the don't-care vertices.

This section will show that Eqs. (7.14) and (7.15) are equivalent to Eq. (7.8) for these deterministic models. The equations are repeated here:

$$R_s F_{t,s} \subseteq (F_{s,t}, R_s) \cup (D_{s,t}^x \times D_{t,s}) \quad (7.14)$$

$$R_s (D_{t,s}^x) \subseteq D_{s,t}^x \quad (7.15)$$

$$R_s F_{s,t} \subseteq F_{t,s} R_s \quad (7.8)$$

**Lemma 7.1:** If Eqs. (7.14) and (7.15) hold for deterministic  $P_{s,t}$ ,  $P_{t,s}$ , and  $R_s$ , then Eq. (7.8) also holds.  $\square$

**PROOF:** Consider arbitrary vertices  $u, s$ , and  $t$  such that  $u R_s s F_{s,t} t$ . It will suffice to show that  $u F_{t,s} R_s t$ . Suppose that  $v \in D_{t,s}^x$ . Then by Eq. (7.15),  $s \in D_{s,t}^x$ . This contradicts the fact that  $s$  has a successor ( $t$ ), so we must have instead  $u \notin D_{t,s}^x$ .

Then  $v = F_{t,s}(u)$  is defined.

Now  $s R_s u F_{t,s} v$ , so from Eq. (7.14),  $s F_{s,t} R_s v$ , since  $s \notin D_{s,t}^x$ . Since  $F_{s,t}$  is deterministic, then we necessarily have  $s F_{s,t} t R_s v$ . Then  $u F_{t,s} v R_s t$ . QED.

**Lemma 7.2:** If Eq. (7.8) holds for deterministic  $P_{s,t}$ ,  $P_{t,s}$ , and  $R_s$ , then Eq. (7.14) also holds.  $\square$

**PROOF:** Consider arbitrary vertices  $s, u$ , and  $v$  such that  $s R_s u F_{t,s} v$ . It will suffice to show that if  $s \notin D_{s,t}^x$ , then  $s F_{s,t} R_s v$ . Assume  $s \notin D_{s,t}^x$ . Then  $t = F_{s,t}(s)$  is defined.

Now  $u R_{t,s} s F_{s,t} t$ , so from Eq. (7.8),  $u F_{t,s} R_{t,s} t$ . Since  $F_{t,s}$  is deterministic, then we necessarily have  $u F_{t,s} v R_{t,s} v$ . QED.

**Lemma 7.3:** If Eq. (7.8) holds for deterministic  $P_{s,t}$ ,  $P_{t,s}$ , and  $R_s$ , then Eq. (7.15) also holds.  $\square$

**PROOF:** Consider arbitrary vertices  $u$  and  $s$  such that  $u R_{t,s} s$  and  $u \in D_{t,s}^x$ . We must show that  $s \in D_{s,t}^x$ . Suppose this is not the case. Then  $t = F_{s,t}(s)$  is defined, and we have  $u R_{t,s} s F_{s,t} t$ . Eq. (7.8) then implies that  $u F_{t,s} R_{t,s} t$ . This contradicts the original condition that  $u$  has no successors. QED.

**Theorem 7.4:** Eq. (7.8) holds for deterministic  $P_{s,t}$ ,  $P_{t,s}$ , and  $R_s$  if and only if Eqs. (7.14) and (7.15) hold.  $\square$

**PROOF:** This follows directly from Lemmas 7.1, 7.2, and 7.3. QED.

This theorem unifies the consistency criterion for the several program models considered. Note that if  $P_{s,t}$  and  $P_{t,s}$  have no don't-cares, then Eqs. (7.14) and (7.15) reduce to Eq. (7.7). Then Eqs. (7.14) and (7.15) constitute a criterion for consistency that is suitable for all four program models discussed in Section 7.7.

- deterministic models without don't-cares (reduce to Eq. (7.7) in Section 7.7.1))
- deterministic models with don't-cares (this section).
- nondeterministic models without don't-cares (reduce to Eq. (7.7) in Section 7.7.3), and
- nondeterministic models with don't-cares (Section 7.7.4).

We will call Eqs. (7.14) and (7.15) the "simple" criterion.

## 7.9. Closure Again

Section 7.4 indicated that Brand had a practical reason for identifying states in the mapping  $R_{si}$  solely by their location counters. Here, we discuss this reason and suggest a more practical criterion for consistency that allows the use of mappings defined in this way. This definition will be related to the simpler definition from the preceding development.

### 7.9.1. Stopping Points in Loops

Consider the programs  $P_s = \langle D_s, F_s \rangle$  and  $P_i = \langle D_i, F_i \rangle$  and relation  $R_s \subseteq D_s \times D_i$  (which defines new programs  $P_{si}$  and  $P_{is}$ ). Also define sets  $D_{si}^* \subseteq D_{si}$  and  $D_{is}^* \subseteq D_{is}$  containing vertices with no successors in  $P_{si}$  and  $P_{is}$ . The criterion for consistency of  $P_i$  with  $P_s$  is that  $P_{si} P_i$  and  $R_{si}$  satisfy Eqs. (7.14) and (7.15).

The program models  $P_s$  and  $P_i$  are so complex that this operation carried out for individual state pairs is intractable. However, symbolic simulation allows the analysis of many state pairs simultaneously. The states in  $D_{a/b}$  are sometimes called "stopping points" [B78], because in a symbolic simulation carrying out the above test, the simulation stops when it reaches a set of states in  $D_{a/b}$ .

A symbolic simulation including a loop that does not contain any stopping points can generate many cases, one case for each possible number of passes through the loop. For this reason, it is advantageous to include in  $D_{a/b}$  a state from each pass through a loop in  $P_a$ . An easy way to do this is to identify stopping points solely by their location counters, and then to specify at least one such stopping point

inside each loop. This is the practical consideration behind Brand's argument in Section 7.4.

### 7.9.2. New Criterion for Consistency

Figure 7-8a repeats Figure 7-2, in which a simple operation in  $P_i$  implements a loop in  $P_s$ . This figure does not satisfy Eq. (7.14), because  $s_i R_{si} u F_{is} v$ , but  $s_i \notin D_{si}^*$  and  $\langle s_i, v \rangle \notin F_{si} R_{is}$ . We argue that nonetheless,  $P_i$  and  $P_s$  are still consistent:  $s_i F_{si} * i R_{si} v$  through a sequence of states  $s_j$  related to  $u$ . A similar remark holds if  $s_i$  is a don't-care vertex, as shown in Figure 7-8b.

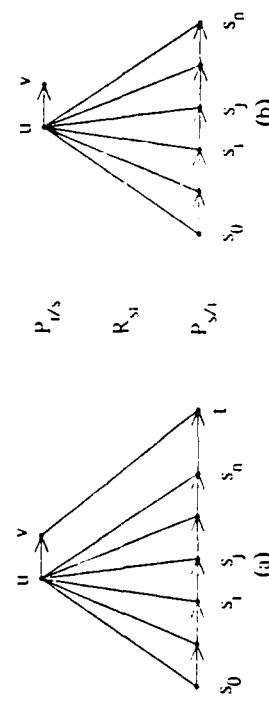


Figure 7-8a: Original example motivating use of closure

Now assume that the following condition is satisfied (see Figure 7-9):

$$\forall \text{ states } s, i \in D_{si} \text{ and } u, v \in D_{is}: (u R_{is} s F_{si} / R_{si} v R_{is} s) \supseteq (u R_{is} i) \quad (7.16)$$

Given  $u R_{is} s_i F_{si} s_j$  with Eq. (7.16) holding, it follows that  $s_i R_{si} u R_{is} s_j$  if and only if  $s_i R_{si} R_{is} s_j$ , as proved here:

**PROOF:** Only if:  $s_i R_{si} R_{is} s_j$  trivially follows from  $s_i R_{si} u R_{is} s_j$ . If: Since

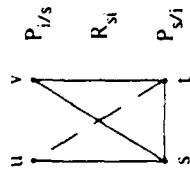


Figure 7-9: Condition for simplified definition

$s_i R_{si} R_{is} s_j$ , there exists a state  $w \in D_{is}$  such that  $s_i R_{si} w R_{is} s_j$ . Then Eq. (7.16) applies with  $\langle s, t, u, v \rangle = \langle s_i, s_j, u, w \rangle$ . QED.

This extension to the notion of consistency then can be expressed by

$$R_{si} F_{is} \subseteq (R_{si} R_{is} \cap F_{si})^* ((F_{si}, R_{si}) \cup (D_{si}^X \times D_{is})) \quad (7.17)$$

Now consider the converse situation in Figure 7-10. This figure also does not satisfy Eq. (7.14):  $s R_{si} u$ ,  $F_{is} u$ , but  $s \notin D_{si}^X$  and  $\langle s, u \rangle \notin F_{si}, R_{si}$ . Again, we argue that  $P_i$  and  $P_s$  are consistent. State  $u_j$  is actually the same point (same location counter) as  $u_i$ , and therefore rightly corresponds not to  $t$ , but to  $s$ . Eq. (7.18) adds this extension to Eq. (7.14). Relation  $I$  is the identity relation.

$$R_{si} F_{is} \subseteq ((F_{si}, I) R_{si}) \cup (D_{si}^X \times D_{is}) \quad (7.18)$$

Eq. (7.18) allows verification in Figure 7-10, but it also incorrectly verifies the example in Figure 7-11. There,  $P_i$  specifies that the program stays in a loop, while  $P_s$  does not allow staying in a loop. If state  $u$  in  $D_{is}$  is in a loop of states all related to  $s$  in  $D_{si}$ , then the successor of  $s$  ( $t$  in Figure 7-11) should be related by  $R_{si}$  to some vertex in the loop. This requirement and Eq. (7.16) imply that the successor of  $s$  must be related to every vertex in the loop. Eq. (7.19) tests this condition:

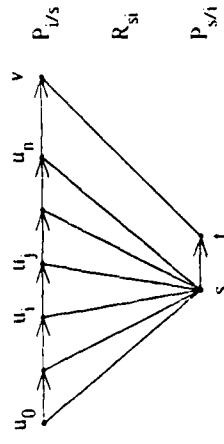


Figure 7-10: Second example motivating use of closure

$$R_{si}((R_{is} R_{si} \cap F_{is})^* \cap I) \subseteq (F_{si}, R_{si}) \cup (D_{si}^X \times D_{is}) \quad (7.19)$$

The plus sign denotes  $\vdash$ -closure; if  $Q$  is a relation, then  $Q \vdash = QQ^*$ .

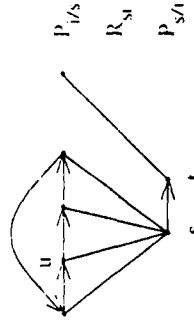


Figure 7-11: Incorrect verification by Eq. (7.18)

The combination of Eqs. (7.17) and (7.18) yields Eq. (7.20):

$$R_{si} F_{is} \subseteq (R_{si} R_{is} \cap F_{si})^* (((F_{si} \cup I) R_{si}) \cup (D_{si}^X \times D_{is})) \quad (7.20)$$

Eq. (7.20) puts closure back into Eq. (7.14), but in a limited fashion that retains intuitive notions of correctness. Eqs. (7.15), (7.19), and (7.20) comprise a criterion for consistency that supports stopping points in loops in both  $P_s$  and  $P_i$  subject to the condition in Eq. (7.16). We will call this criterion the "extended" criterion.

## 7.10. Relation between Extended and Simple Criteria

### 7.10.1. review

The object of Section 7.10 is to demonstrate that the extended criterion for consistency, designed to facilitate the verification of programs with loops, still retains intuitive notions of correctness. The approach is to show that

- if  $P_i$  is consistent with  $P_s$  under  $R_{si}$  by the extended criterion from

Section 7.9,

- then  $P_i$  is also consistent with  $P_s$  under a reasonable simulation relation  $R'_{si}$  by the simple criterion from Section 7.8.

Therefore, the extended criterion retains the intuitive notions of correctness captured by the simple definition. This will be proven only for deterministic models.

This discussion takes place in three parts.

$$\begin{aligned} R_{si}(R_{si} \cap R_{si}) + \cap 1 &\subseteq (F_{si} \cap R_{si}) \cup (D_{si}^x \times D_{si}^x). \\ R_{si}(R_{si} \cap F_{si}) \cap (F_{si} \cap R_{si})^* ((F_{si} \cup 1) R_{si}) \cup (D_{si}^x \times D_{si}^x) \end{aligned} \quad (7.19)$$

First, Section 7.10.2 defines the new relation  $R'_{si}$ . To do this, it defines an equivalence relation on  $D_{ab}$  and selects a distinguished ("terminal") vertex from each equivalence class. Relation  $R'_{si}$  is then defined on the distinguished vertices.

Second, Section 7.10.3 demonstrates that  $R'_{si}$  is a reasonable simulation relation. It does this by showing that "important" behavioral details captured by  $R_{si}$  are also captured by  $R'_{si}$ .

Third, Section 7.10.4 shows that  $P_i$  is consistent with  $P_s$  under  $R'_{si}$  by the simple criterion, assuming consistency under  $R_{si}$  by the extended criterion.

The discussion in Section 7.10 will assume that the program model is deterministic with don't-cares allowed. For each state  $s \in D_a$ , there is exactly zero or one state  $t$  such that  $s F_a t$ . Note that  $P_{a/b}$  will also be deterministic with don't-cares.

Given programs  $P_s$  and  $P_i$  and relation  $R_{si} \subseteq D_s \times D_i$ , which satisfy Eqs. (7.15), (7.19), and (7.20), we will show how to derive relation  $R'_{si}$  satisfying Eqs. (7.14) and (7.15). This discussion makes the following assumptions about  $P_s$ ,  $P_i$ ,  $R_{si}$ , programs  $P_{si}$  and  $P_{i/s}$  defined by Eq. (7.6), and sets  $D_{si}^x$  and  $D_{i/s}^x$  defined by Eq. (7.5).

- All programs  $P_a$  have a finite number of states in  $D_a$ .
  - In all programs  $P_a$ ,  $F_a$  is a function with domain a subset of  $D_a$ :
- $$\forall \text{ vertices } s, t, r \text{ in } D_a : (s F_a t \text{ and } s F_a r) \supseteq (t = r). \quad (7.21)$$
- Eqs. (7.15), (7.19), and (7.20) are satisfied:

$$\begin{aligned} R_{si}(D_{i/s}^x) &\subseteq D_{si}^x. \\ R_{si}((R_{si} R_{si} \cap F_{i/s}) + \cap 1) &\subseteq (F_{si} \cap R_{si}) \cup (D_{si}^x \times D_{si}^x). \\ R_{si} F_{si} \subseteq (R_{si} R_{is} \cap F_{si})^* ((F_{si} \cup 1) R_{si}) \cup (D_{si}^x \times D_{si}^x) \end{aligned} \quad (7.19)$$

- The following holds. This includes Eq. (7.16) and the symmetric condition.

$$\begin{aligned} \forall \text{ states } s, t \in D_{a/b} \text{ and } u, v \in D_{b/a} : \\ (s R_{ab} u F_{ba} v R_{ba} t R_{ab} u) \supseteq (s R_{ab} v) \end{aligned} \quad (7.22)$$

### 7.10.2. New Relation

Define relation  $R_{aa}$  on  $D_{a/b} \times D_{a/b}$ :

$$R_{aa} = (R_{ab} R_{ba}) \cap (F_{a/b} \cup F_{a/b}^{-1})^*$$

Relation  $R_{aa}$  is an equivalence relation since it is reflexive, symmetric, and transitive. Denote by  $R_{aa}[s]$  the equivalence class defined by  $R_{aa}$  that contains vertex  $s$ . Relations  $R_{ss}$  and  $R_{ii}$  are defined by Eq. (7.23), with resulting equivalence classes  $R_{ss}[s]$  and  $R_{ii}[u]$ .

Now define a *terminal vertex* in  $D_{a/b}$  to be any vertex  $t$  reachable from all other vertices in  $R_{aa}[t]$  via a path in  $R_{aa}[t]$ :

$t$  is a terminal vertex iff

$$\forall \text{vertices } s \in R_{aa}[t]: s (R_{ab} R_{ba} \cap F_{a/b})^* t \quad (7.24)$$

Choose a set  $D'_{a/b}$ , a set of terminal vertices with exactly one terminal vertex from each equivalence class. If an equivalence class has more than one terminal vertex, the choice of the terminal vertex is arbitrary. (Lemma 7.5 will show that every equivalence class has at least one terminal vertex.) Define  $t_{ae}[s]$  to be the terminal vertex of  $R_{aa}[s]$  that is in  $D'_{a/b}$ . (Subscript "e" denotes "equivalence class.")

Terminal vertices in  $D_{ss}$  and  $D_{ii}$  and the vertex sets  $D'_{ss}$  and  $D'_{ii}$  are defined in this way. Then define the new relation  $R'_{ss}$  by

$$R'_{ss} = R_{ss} \cap (D'_{ss} \times D'_{ss}) \quad (7.25)$$

New programs  $P'_{ss} = \langle D'_{ss}, F'_{ss} \rangle$  and  $P'_{ii} = \langle D'_{ii}, F'_{ii} \rangle$  are then defined from  $P_s$ ,  $P_i$ , and  $R'_{ss}$  by Eq. (7.6). Don't-care vertex sets  $D'_{ss}$  and  $D'_{ii}$  are defined for  $P'_{ss}$  and  $P'_{ii}$  by Eq. (7.5). The major claims of this discussion will be that  $R'_{ss}$  is a suitable relation, and that  $P_s$ ,  $P_i$ , and  $R'_{ss}$  satisfy Eqs. (7.14) and (7.15).

$$\begin{aligned} R_{ss} F'_{i/s} &\subseteq (F'_{ss} R'_{ss}) \cup (D'_{ss} \times D'_{i/s}) \\ R_{ss} (D'_{i/s}) &\subseteq D'_{ss} \end{aligned}$$

### 7.10.3. Suitability of New Relation

Note that it is trivial to define a relation  $R''_{ss}$  that satisfies Eqs. (7.14) and (7.15) given arbitrary programs  $P_s$  and  $P_i$ . In fact, the empty set is such a relation. Obviously, the empty set does not capture the desired intuitive correspondence between  $P_s$  and  $P_i$ .

However,  $R'_{ss}$  does capture this correspondence as expressed by  $R_{ss}$ , in the following sense: For all vertices  $u$  and  $v$  in  $D_{a/b}$  but not in the same equivalence class defined by  $R_{aa}$ , we have (Figure 7-12)

- if  $u F'_{a/b} v$ , then  $u F'_{a/b} t_{ae}[v]$ ; also,  $u R'_{ab} s$  for some vertex  $s$  in  $D'_{a/b}$  (this guarantees that  $u F'_{i/s} v$  is a correct edge if Eqs. (7.14) and (7.15) is satisfied with  $R'_{ss}$ ); and
- if  $u F'_{a/b} v$ , then there is a vertex  $w$  such that  $u F'_{a/b} w$  and  $v = t_{ae}[w]$ .

Figure 7-12 illustrates these assertions. In the figure, boxed vertices are in  $D'_{a/b}$  or  $D'_{b/a}$ , and dashed edges are edges whose existence is implied by the presence of other solid edges. A small oval denotes an equivalence class, and a wavy edge denotes a path within a single equivalence class.

Proofs of the above assertions follow. Each theorem directly proves a portion of the assertions. The lemmas facilitate the proof of the theorems. All lemmas and theorems below assume the definitions and assumptions from Sections 7.10.1 and 7.10.2, including Eqs. (7.21) through (7.25). The reader may wish to skip the proofs on first reading.

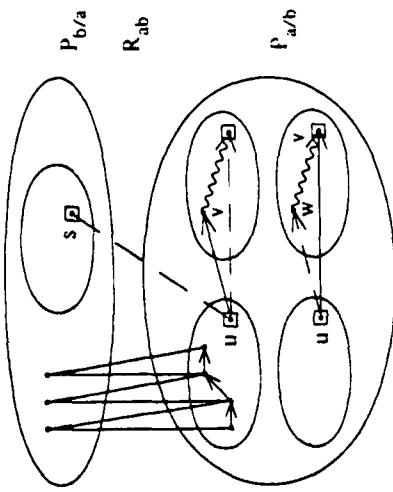


Figure 7-12: Properties of new simulation relation

**Lemma 7-5:** Every equivalence class  $R_{ab}[s]$  has at least one terminal vertex.  $\square$

**PROOF:** Select an arbitrary vertex  $s$  in  $D_{ab}$  and look at its descendants in the directed edge set  $(R_{ab} R_{ba} \cap F_{ab})$ . Since  $D_{ab}$  has a finite number of states, one of three cases in Figure 7-13 arises: (a) The path reaches a vertex that has no successor.

- (b) The path ends at a self-loop. (c) The path enters a larger loop.

Consider case (a). Let  $t$  be the vertex with no successor. For every vertex  $m$  in  $R_{ab}[t]$ , there exists a loop-free path  $m (R_{ab} R_{ba} \cap (F_{ab} \cup F_{ba}^{-1}))^*$   $t$ , as illustrated in Figure 7-14. For each pair of adjacent vertices  $j$  and  $k$  in this path, either  $j F_{ab} k$  or  $k F_{ab} j$ . In particular, in Figure 7-14, either  $p F_{ab} t$  or  $t F_{ab} p$ . But  $t$  has no successor in the edge set  $(R_{ab} R_{ba} \cap F_{ab})$ . Therefore  $p F_{ab} t$ .

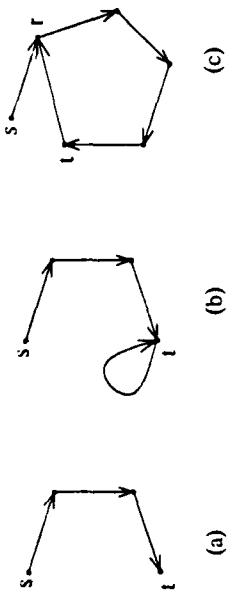


Figure 7-13: Three types of terminal vertices

Moving to the next pair, either  $p F_{ab} q$  or  $q F_{ab} p$ . But  $p$  has at most one successor in the edge set  $(R_{ab} R_{ba} \cap F_{ab})$  (Eq. (7.21)).  $p F_{ab} t$ , and  $t$  is distinct from  $q$  (since the path from  $m$  to  $t$  is loop-free). Therefore  $q F_{ab} p$ . Applying this step inductively, we find  $m (R_{ab} R_{ba} \cap F_{ab})^* t$ . This satisfies Eq. (7.24), so  $t$  is by definition a terminal vertex.

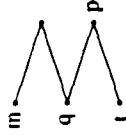


Figure 7-14: Reachability of terminal vertex for case (a)

In cases (b) and (c) of Figure 7-13, we can remove edge  $\langle t, t \rangle$  or  $\langle t, r \rangle$  without modifying the equivalence class partition. The edge removal changes these cases to case (a), since  $t$  then has no successor. In each modified case,  $t$  is a terminal vertex by the above analysis. Putting the deleted edge back into  $F_{ab}$  does not alter the reachability of  $t$ , so  $t$  is a terminal vertex in each original case as well. QED.

**Lemma 7.6:** Given arbitrary vertices  $s$  in  $D_{s,t}$  and  $u$  in  $D_{t,s}$ : If  $s R_s u$ , then all vertices reachable from  $s$  by  $(R_s R_s \cap F_{s,t})^*$  (including  $s$ ) are related by  $R_{s,t}$  to all vertices reachable from  $u$  by  $(R_t R_t \cap F_{t,s})^*$  (including  $u$ ).  $\square$

**PROOF:** By induction, it suffices to show that if  $s (R_s R_s \cap F_{s,t})^* t$  then  $t R_{s,t} u$ , and if  $u (R_t R_t \cap F_{t,s})^* s$  then  $s R_{s,t} v$ . Both of these follow from Eq. (7.22). QED.

**Lemma 7.7:** Given arbitrary vertices  $s$  in  $D_{s,t}$  and  $u$  in  $D_{t,s}$ : If  $s R_s u$ , then  $t_x[s] R_{s,t} t_x[u]$ .  $\square$

**PROOF:** By definition of terminal vertices,  $t_x[s]$  is reachable from  $s$  by  $(R_s R_s \cap F_{s,t})^*$ . Similarly,  $t_x[u]$  is reachable from  $u$  by  $(R_t R_t \cap F_{t,s})^*$ . Then Lemma 7.6 applies. QED.

**Lemma 7.8:** Given arbitrary vertices  $s$  and  $t$  in different equivalence classes  $R_{aa}[s]$  and  $R_{ab}[t]$ : If  $s F_{ab} t$ , then  $s$  is the unique terminal vertex for  $R_{aa}[s]$ .  $\square$

**PROOF:** If  $s$  and  $t$  are in different equivalence classes and  $s F_{ab} t$ , then vertex  $s$  has no successor in the edge set  $(R_{ab} R_{ba} \cap F_{ab})$ . This is an example of case (a) in the proof of Lemma 7.5, so  $s$  is a terminal vertex. We show it is the only terminal vertex by contradiction. Assume  $R_{aa}[s]$  has a second terminal vertex  $r$  distinct from  $s$ . By the definition of terminal vertices,  $r$  is reachable from  $s$  by  $(R_{ab} R_{ba} \cap F_{ab})^*$ , and  $s$  is reachable from  $r$ . Since  $F_{ab}$  is a directed edge set, this implies the existence of a loop including  $r$  and  $s$ , as shown in Figure 7-15. But then  $s$  has two successors, violating Eq. (7.21). Therefore,  $r$  cannot be a distinct terminal vertex, so  $s$  is unique. QED.

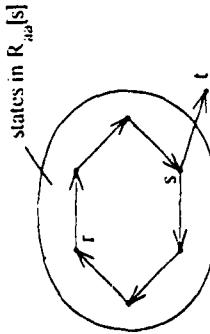


Figure 7-15: Impossibility of two distinct terminal vertices

**Lemma 7.9:** Given arbitrary vertex  $s$  in  $D_{ab}$ : If  $s$  is a don't-care vertex in  $P_{ab}$ , then  $s$  is the unique terminal vertex for  $R_{aa}[s]$ .  $\square$

**PROOF:** The proof parallels that for Lemma 7.8. Vertex  $s$  corresponds to case (a) in the proof of Lemma 7.5, so  $s$  is a terminal vertex. If  $R_{aa}[s]$  had a second terminal vertex, violating Eq. (7.21). Therefore,  $s$  is the unique terminal vertex for  $R_{aa}[s]$ .  $\square$

distinct terminal vertex  $t$ , then  $s$  would have a successor in the edge set  $(R_{ab} R_{ba} \cap F_{ab})$ , contradicting the given condition that  $s$  is a don't-care vertex in  $P_{ab}$ . QED.

The following three theorems demonstrate the assertions claimed at the beginning of this section.

**Theorem 7.10:** Given arbitrary vertices  $u$  and  $v$  in different equivalence classes  $R_{aa}[u]$  and  $R_{aa}[v]$ : If  $u F_{ab} v$ , then  $u F'_{ab} t_{ae}[v]$ .  $\square$

**PROOF:** By Lemma 7.8,  $u$  is the unique terminal vertex for  $R_{aa}[u]$ . Therefore  $u$  is in  $D'_{ab}$ . Vertex  $t_{ae}[v]$  is also in  $D'_{ab}$ . By definition of a terminal vertex, there exists path(s)  $u F_{ab} v (R_{ab} R_{ba} \cap F_{ab})^*$   $t_{ae}[v]$ . Choose such a path in which each vertex appears only once (no loops). We have  $u F_a * t_{ae}[v]$  along this path. Also, since each equivalence class has only one vertex in  $D'_{ab}$ , then only the endpoints of this path are in  $D'_{ab}$ . Therefore, by definition of  $F'_{ab}$ ,  $u F'_{ab} t_{ae}[v]$ . QED.

**Theorem 7.11:** Given arbitrary vertices  $u$  and  $v$  in different equivalence classes  $R_{aa}[u]$  and  $R_{aa}[v]$ : If  $u F_{ab} v$ , then  $u R'_{ab} s$  for some vertex  $s$  in  $D'_{ba}$ .  $\square$

**PROOF:** Since  $u$  is in  $D'_{ab}$ , there must exist a vertex  $r$  in  $D_{ba}$  such that  $u R_{ab} r$ . By Lemma 7.7,  $t_{ae}[u] R_{ab} t_{be}[r]$ . Since  $t_{ae}[u]$  and  $t_{be}[r]$  are in  $D'_{ab}$  and  $D'_{ba}$ , respectively, then by definition of  $R'_{ab}$ ,  $t_{ae}[u] R'_{ab} t_{be}[r]$ . But by Lemma 7.8,  $u$  is the unique terminal vertex for  $R_{aa}[u]$ , so that  $u = t_{ae}[u]$ . Then the theorem is true with  $s = t_{ae}[u]$ . QED.

**Theorem 7.12:** Given arbitrary vertices  $u$  and  $v$  in  $D_{ab}$ : If  $u F'_{ab} v$ , then  $u F_{ab} w$  for some vertex  $w$  such that  $v = t_{ae}[w]$ .  $\square$

**PROOF:** If  $u F'_{ab} v$ , then certainly  $u F_{ab} w$  for some vertex  $w$  in  $D_{ab}$ . Complete the proof by contradiction: Suppose that  $v \neq t_{ae}[w]$ . Then  $v \notin R_{aa}[w]$  since  $v = t_{ae}[v]$ . We have  $u F_{ab} w F_{ab} * v$  since  $F_{ab}$  is deterministic. The path  $w F_{ab} * v$  must exit  $R_{aa}[w]$  at some point, and by Lemma 7.8 it can do so only through the unique terminal vertex  $t_{ae}[w]$ . Therefore we necessarily have  $u F_{ab} w F_{ab} * t_{ae}[w] F_{ab} F_{ab} * v$ .

But then a midpoint (namely  $t_{ae}[w]$ ) of any path from  $u$  to  $v$  is in  $D'_{ab}$ , contradicting the given fact that  $u F'_{ab} v$ . Then the supposition that  $v \neq t_{ae}[w]$  is false, proving the theorem. QED.

Theorems 7.10, 7.11, and 7.12 prove the desirable properties of  $R'$  asserted at the beginning of this section. If  $s F_{ab} t$ , and  $s$  and  $t$  are in different equivalence classes, then edge  $\langle s, t \rangle$  in  $F'_{ab}$  will uniquely correspond to edge  $\langle s, t \rangle$  in  $F'_{ab}$ . If  $s F_{ab} t$ , and  $s$  and  $t$  are in the same equivalence class, then edge  $\langle s, t \rangle$  in  $F_{ab}$  is probably "lost" in  $F'_{ab}$ , not corresponding to any edge there. (The exception arises when  $s$  and  $t$  are both in a loop within  $R_{aa}[s]$ .) The reasoning is that  $s$  and  $t$  both are related to the same vertices in  $D_{ba}$  by  $R_{ab}$  only to simplify the test for consistency by symbolic simulation; the essential information in  $R_{ab}$  is retained if we keep only one stopping point from each equivalence class.

The next section will show that if  $P_s$ ,  $P_t$ , and  $R_{aa}'$  satisfy Eqs. (7.15), (7.19), and (7.20), then  $P_s$ ,  $P_t$ , and  $R_{aa}'$  will satisfy Eqs. (7.14) and (7.15).

This section will demonstrate that  $P_i$  is consistent with  $P_s$  under  $R'_{si}$  by the simple criterion, assuming consistency under  $R_{si}$  by the extended criterion. This is expressed by Theorem 7.13. The following statement of the theorem summarizes the definitions and restrictions assumed in the preceding discussion.

**Theorem 7.13:** Given the following:

- Programs  $P_s = \langle D_s, F_s \rangle$  and  $P_i = \langle D_i, F_i \rangle$  and relation  $R_s \subseteq D_s \times D_i$ ;  $D_s$  and  $D_i$  are finite.
- $F_s$  and  $F_i$  functions on subsets of  $D_s$  and  $D_i$ ; Eq. (7.21).
- $P_{si}$  and  $P_{i/s}$  defined from  $P_s$ ,  $P_i$ , and  $R_{si}$  by Eq. (7.6).
- $P_{si}$ ,  $P_{i/s}$ , and  $R_{si}$  satisfy Eqs. (7.15), (7.19), (7.20), and (7.22).
- $D'_{si}$ , subset of  $D_{si}$ , containing exactly one state from each equivalence class generated by relation  $(R_{si} \cap (F_{si} \cup F_{si}^{-1}))^*$ .
- $D'_{i/s}$ , subset of  $D_{i/s}$ , containing exactly one state from each equivalence class generated by relation  $(R_{i/s} \cap (F_{i/s} \cup F_{i/s}^{-1}))^*$ .
- $R'_{si} = R_{si} \cap (D'_{si} \times D'_{i/s})$ .
- $P'_{si}$  and  $P'_{i/s}$  defined from  $P_s$ ,  $P_i$ , and  $R'_{si}$  by Eq. (7.6).
- Don't-care vertex sets  $D'_{si}$  and  $D'_{i/s}$  defined from  $P'_{si}$  and  $P'_{i/s}$  by Eq. (7.5).

The following result holds:  $P'_{si}$ ,  $P'_{i/s}$ , and  $R'_{si}$  satisfy Eqs. (7.14) and (7.15):

$$\begin{aligned} R'_{si} F_{i/s} &\subseteq (F'_{si} R'_{si}) \cup (D'_{si} \times D'_{i/s}) \\ R'_{si} (D'_{i/s}) &\subseteq D'_{si} \quad \square \end{aligned} \quad (7.26)$$

$$(7.27)$$

First, we state and prove several lemmas. All the lemmas will assume the conditions given in Theorem 7.13. Lemmas 7.14 and 7.15 prove useful properties.

Lemmas 7.16 and 7.17 prove Eq. (7.26) for two cases, thereby partially proving

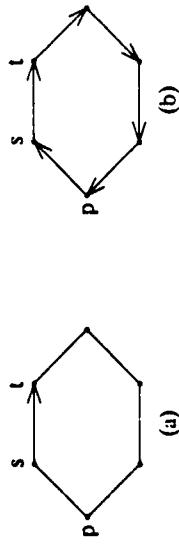
Theorem 7.13. The theorem will follow directly from the lemmas. The reader may wish to skip the proofs on first reading.

**Lemma 7.14:** Given states  $s$  and  $t$  in  $D_{ab}$ : If  $s \not\sim_{ab} t$  and by  $F_{ab}$ , then they are also related by  $R_{ab} R_{ba}$ .

**Lemma 7.14:** Given states  $s$  and  $t$  in  $D_{ab}$ : If  $s \not\sim_{ab} t$  and  $t \in R_{ab}[s]$ , then  $s (R_{ab} R_{ba} \cap F_{ab}) t$ .  $\square$

**PROOF:** Prove by contradiction (Figure 7-16). Assume  $s \not\sim_{ab} t$  and  $t \in R_{ab}[s]$ , but  $\langle s, t \rangle \notin (R_{ab} R_{ba} \cap F_{ab})$ . Then  $\langle s, t \rangle \notin R_{ab} R_{ba}$ . Since  $s$  and  $t$  are in the same equivalence class, we know  $s$  and  $t$  are joined by a path  $(R_{ab} R_{ba} \cap (F_{ab} \cup F_{ab}^{-1}))^*$ , which we choose to be loop-free. This path cannot include edge  $\langle s, t \rangle$  since  $\langle s, t \rangle \notin R_{ab} R_{ba}$ . Therefore the situation is as shown in Figure 7-16a. The directed edge is edge  $s F_{ab} t$ . The undirected edges are edges in  $(R_{ab} R_{ba} \cap (F_{ab} \cup F_{ab}^{-1}))$ .

Now  $(R_{ab} R_{ba} \cap (F_{ab} \cup F_{ab}^{-1})) \subseteq (F_{ab} \cup F_{ab}^{-1})$ , so for each undirected edge  $\langle j, k \rangle$  in the path, either  $j F_{ab} k$  or  $k F_{ab} j$ . In particular, in Figure 7-16a, either  $p F_{ab} s$  or  $s F_{ab} p$ . But  $s F_{ab} t$ ,  $F_{ab}(s)$  is single-valued, and  $t$  is distinct



**Figure 7-16:** Connectivity between states in equivalence class

from  $\rho$  (since  $s R_{ab} R_{ba} \rho$ ). Then we cannot have  $s F_{ab} p$ , so  $p F_{ab} s$ . Applying a similar argument along the path recursively, we find  $t F_{ab} + s$  as shown in Figure 7-16b. Since  $j R_{ab} R_{ba} k$  for each edge  $\langle j, k \rangle$  in the path, then  $t (R_{ab} R_{ba} \cap F_{ab}) + s$ .

There is a state  $u$  in  $D_{ba}$  such that  $t R_{ab} u$ . By Lemma 7.6, we also have  $s R_{ab} u$  since  $s$  is reachable from  $t$  by  $(R_{ab} R_{ba} \cap F_{ab})^*$ . (Note that the conditions for Lemma 7.6 are included in the conditions for Theorem 7.13.) Then  $s R_{ab} + R_{ba} t$ , which implies  $s (R_{ab} R_{ba} \cap F_{ab}) t$ . This contradicts the original assumption. QED.

Now consider an arbitrary vertex  $s$  in  $D_{ab}$ . Lemma 7.15 says that  $s$  is in  $D'_{ab}$  and has no successors in  $P'_{ab}$  if and only if  $s$  has no successors in  $P_{ab}$ .

**Lemma 7.15:**  $D'_{ab} = D_{ab}$ .  $\square$

**PROOF:** Consider an arbitrary dom<sup>c</sup> care vertex  $s$  in  $D_{ab}$  ( $s \in D_{ab}^x$ ). By Lemma 7.9, vertex  $s$  is the unique terminal vertex of  $R_{ab}[s]$ . Then  $s \in D'_{ab}$  and so  $s \in D_{ab}^x$ . This shows  $D'_{ab} \subseteq D_{ab}^x$ .

Now suppose some vertex  $r$  is in  $D_{ab}^x$  but not  $D'_{ab}$ . We have  $r F_{ab} t$  for some vertex  $t \in D_{ab}$ . Then  $r F_{ab} + t_{ab}[t]$ , so  $r \notin D_{ab}^x$ , since  $t_{ab}[t]$  is in  $D'_{ab}$ . This contradicts the supposition that  $r \in D_{ab}^x$ . This shows  $D_{ab}^x \subseteq D'_{ab}$ . QED.

**Lemmas 7.16 and 7.17 prove Eq. (7.26) in Theorem 7.13.** Each lemma assumes a vertex pair  $\langle s, v \rangle$  that is in the left hand side (lhs) of Eq. (7.26). This implies the existence of a vertex  $u$  also assumed in the lemmas. Either  $u$  and  $v$  are in the same equivalence class or they are not; the lemmas cover each case in turn.

Now if  $s$  is in  $D_{sv}^x$ , then Eq. (7.26) is immediately satisfied. The lemmas therefore assume that  $s$  is not in  $D_{sv}^x$  and show that Eq. (7.26) is still satisfied.

**Lemma 7.16 covers the case where  $u$  and  $v$  are in the same equivalence class,** illustrated in Figure 7-17. To avoid clutter, arrows are shown for only one edge in each loop in the figure.

**Lemma 7.16:** Given states  $s$  in  $D'_{sv}$  and  $u$  and  $v$  in  $D'_{sv}$ : If  $s R'_s u F'_{vs} v, s \notin D_{sv}^x$ , and  $v$  is in  $R_u[u]$ , then  $s F'_{sv} R'_v v$ .  $\square$

**PROOF:** Make the following observations. (In proofs with numbered paragraphs, the text in each paragraph justifies the initial assertion in the paragraph.)

- (1)  $s R'_s v F'_{vs} v$ . Since each equivalence class in  $D_{vs}$  has only one state in  $D'_{vs}$ , then  $u = v$ .
- (2)  $v F'_{vs} + v$  through a sequence of vertices in  $R_u[u]$ . From (1) by definition of  $F'_{vs}$ ,  $v F'_{vs} + v$  along a path that does not include any other states in  $D'_{vs}$ .

(11)  $s \in F'_{S/I} R'_{S/I} v$ , from (1) and (10). QED.

Lemma 7.17 covers the case where  $u$  and  $v$  are in different equivalence classes, illustrated in Figure 7-18. The figure shows  $s$  separate from  $t$  to facilitate following the proof, although the proof will show  $s = t$ .

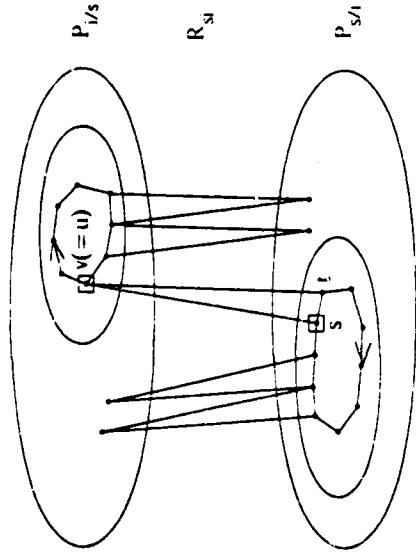


Figure 7-17: Proof of Eq. (7.26): first case

$D'_{S/I}$ . But if the path leaves  $R_u[v]$ , it must necessarily pass through another state in  $D'_{S/I}$ , when it returns to  $R_u[v]$  (see Lemma 7.8 and proof of Theorem 7.12).

(3)  $v \in (R_{S/I} R_{S/I} \cap F_{S/I})_t \setminus v$ , by recursive application of Lemma 7.14 to (2).

(4)  $s \in R_u[v]$ , from (1) since  $R'_u$  is a subset of  $R_{S/I}$ .

(5)  $s \in F_{S/I} t R_{S/I} v$  for some state  $t$  in  $D_{S/I}$ , from (3) and (4) by Eq. (7.19); note  $s \notin D'_{S/I}$  and  $D'_{S/I} = D^*_{S/I}$  by Lemma 7.15.

(6)  $s \in (R_{S/I} R_{S/I} \cap F_{S/I})_t$ , from (4) and (5).

(7)  $t \in (R_{S/I} R_{S/I} \cap F_{S/I})^* s$ , because  $s \in D'_{S/I}$  is a terminal vertex, and from (6),  $t$  is in  $R_{S/I}[s]$ .

(8)  $s \in (R_{S/I} R_{S/I} \cap F_{S/I})^* s$ , from (6) and (7).

(9)  $s \in F_{S/I} + s$  through a path in  $R_{S/I}[s]$ , from (8).

(10)  $s \in F_{S/I} s$ , from (9), since  $s$  is the only state from  $R_{S/I}[s]$  in  $D'_{S/I}$ .

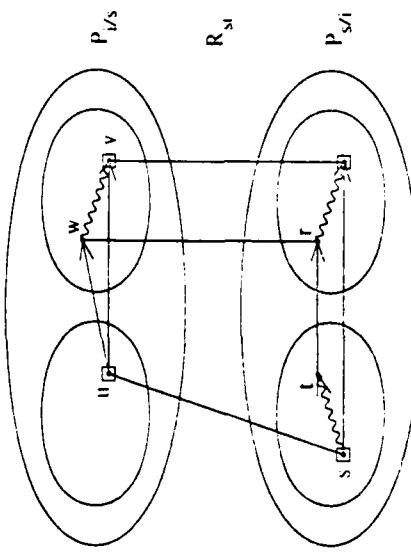


Figure 7-18: Proof of Eq. (7.26): second case

**Lemma 7.17:** Given states  $s$  in  $D'_{S/I}$  and  $u$  and  $v$  in  $D'_{S/I}$ : If  $s \in R'_{S/I} u F'_{S/I} v, s \notin D^*_{S/I}$ , and  $v$  is not in  $R_u[u]$ , then  $s \in F'_{S/I} R'_{S/I} v$ .  $\square$

**PROOF:** Make the following observations.

- (1)  $s \in R'_{S/I} u F'_{S/I} v$ , with  $v$  not in  $R_u[u]$ , by assumption.
- (2)  $u \in F'_{S/I} w (R_{S/I} \cap F_{S/I})^* v$  for some state  $w$  in  $D_{S/I}$ , from (1) and Theorem 7.12, and by definition of a terminal vertex. Note that  $w$  is in the same equivalence class as  $v$ , and therefore in a different equivalence class from  $u$ .

- (3)  $s \in R_{sI} F_{sI} w$ , from (1) and (2), since  $R_{sI} \subseteq R_{sI}$ .  
 (4)  $\langle s, w \rangle \notin R_{sI}$ ; otherwise, we would have  $u (R_{sI} R_{sI} \cap F_{sI}) w$ . This would imply  $w \in R_{sI}[u]$ , contradicting (2).

- (5)  $s (R_{sI} R_{sI} \cap F_{sI})^* t ((F_{sI} \cup D) R_{sI}) \cup (D_{sI}^x \times D_{sI}) w$  for some state  $t$ , by (3) and Eq. (7.20).

(6)  $t \in D_{sI}^x$ , by assumption. (This will lead to a contradiction.)

- (7)  $s = t$ . From (1),  $s = t_{sI}[s]$ . Step (5) gives  $t_{sI}[s] = t_{sI}[t]$ . Finally, by

Lemma 7.9, (6) implies  $t_{sI}[t] = t$ .

- (8)  $s \in D_{sI}^x$ , from (6) and (7). This contradicts the given condition that

$s \notin D_{sI}^x$ . The assumption of  $t \in D_{sI}^x$  in (6) led to a contradiction in (8).

- (9)  $t \notin D_{sI}^x$ . We have  $r \in D_{sI}^x \cap F_{sI}$  in (10) and (11).

- (10)  $s (R_{sI} R_{sI} \cap F_{sI})^* t (F_{sI} \cup D) r R_{sI} w$  for some vertex  $r$ , from (5) and (9).

(11)  $s \neq r$ , because  $r R_{sI} w$  by (10), but  $\langle s, w \rangle \notin R_{sI}$  by (4).

- (12) WLOG assume  $s (R_{sI} R_{sI} \cap F_{sI})^* t F_{sI} r R_{sI} w$ . If  $t \neq r$  in (10), then

from (11),  $s \neq t$ , so  $s (R_{sI} R_{sI} \cap F_{sI})^* t$ . Then we can redefine  $t$  to denote the predecessor of  $r$  in this path.

- (13)  $r \in R_{sI}[t]$  by assumption. (This will lead to a contradiction.)

- (14)  $s R_{sI} w$ . We have  $r R_{sI} w$  from (12), and  $s = t_{sI}[r]$  from (12) and (13), since  $s = t_{sI}[s]$  by (1). Therefore  $s$  is reachable from  $r$  by  $(R_{sI} R_{sI} \cap F_{sI})^*$ . Then Lemma 7.6 gives  $s R_{sI} w$ . This directly contradicts (4).

- (15)  $r$  and  $t$  are in different equivalence classes. The assumption of  $r \in R_{sI}[t]$  in (13) led to a contradiction in (14).

- (16)  $s = t$ . From (1),  $s = t_{sI}[s]$ . Step (12) gives  $t_{sI}[s] = t_{sI}[t]$ . Finally, by Lemma 7.8, (12) and (15) imply  $t_{sI}[t] = t$ .

- (17)  $s F_{sI} r R_{sI} w$ , by (12) and (16).

- (18)  $t_{sI}[r] R_{sI} v$ . By Lemma 7.7 and (17),  $t_{sI}[r] R_{sI} t_{sI}[w]$ . Since  $t_{sI}[r] \in D_{sI}$  and  $t_{sI}[w] \in D_{sI}$ , then  $t_{sI}[r] R_{sI} t_{sI}[w]$ . From (2),  $t_{sI}[w] = t_{sI}[v]$ , and from (1),  $t_{sI}[v] = v$ .

- (19)  $s F_{sI} t_{sI}[r]$ . From (15), (16), and (17) by Theorem 7.10.

- (20)  $s F_{sI} R_{sI} v$ , by (18) and (19). QED.

**PROOF OF THEOREM 7.13:** To prove Eq. (7.26), consider an arbitrary state pair  $\langle s, v \rangle$  such that  $s R_{sI} F'_{sI} v$ . It will suffice to show that if  $s \notin D_{sI}^x$ , then  $s F'_{sI} R_{sI} v$ . Since  $s R_{sI} F'_{sI} v$ , there exists a state  $u$  such that  $s R_{sI} u F'_{sI} v$ . There are two cases; either  $u$  and  $v$  are in the same equivalence class or they are not.

Lemmas 7.16 and 7.17 show that  $s F'_{sI} R_{sI} v$  for each of these cases.

Eq. (7.27) follows directly from Eq. (7.15) and Lemma 7.15. QED.

## 7.11 Hierarchical Verification

A desirable characteristic of a criterion for consistency is the ability to apply it hierarchically. That is, if program  $P_j$  is consistent with  $P_s$ , and a third program  $P_i$  is consistent with  $P_j$ , we would like to conclude that  $P_i$  is therefore consistent with  $P_s$ . The following theorem allows such a conclusion, using the implied specification Eqs. (7.14) and (7.15).

**Theorem 7.18:** Given the following:

- Programs  $P_s = \langle D_s, F_s \rangle$ ,  $P_i = \langle D_i, F_i \rangle$ , and  $P_j = \langle D_j, F_j \rangle$ , and relations  $R_{sI} \subseteq D_s \times D_j$  and  $R_{iI} \subseteq D_i \times D_j$ .
- Relation  $R_{sj}$  defined by

$$R_{sl} = R_s R_u \quad (7.28)$$

• Programs  $P_{sl}, P_{us}, P_{ul}, P_{vs}, P_{sj},$  and  $P_{is}$  defined by Eq. (7.6).

• Corresponding don't-care vertex sets defined by Eq. (7.5).

Programs satisfy the following conditions:

$$R_u F_{ul} \subseteq (F_{ul} R_u) \cup (D_{ul}^s \times D_{ul}^x) \quad (7.29)$$

$$R_u (D_{ul}^s) \subseteq D_{ul}^x \quad (7.30)$$

$$R_{sl} F_{us} \subseteq (F_{us} R_{sl}) \cup (D_{us}^s \times D_{us}^x) \quad (7.31)$$

$$R_{us} (D_{us}^s) \subseteq D_{us}^x \quad (7.32)$$

$$D_{us} \subseteq D_{ul} \quad (7.33)$$

$$\forall u \in D_u, \forall i \in D_s, u R_u R_k / \supset (3, \epsilon, D_s) \cup R_{is} s \quad (7.34)$$

Then the following result holds:

$$R_u F_{us} \subseteq (F_{us} R_{sl}) \cup (D_{us}^s \times D_{us}^x) \quad (7.35)$$

$$R_u (D_{us}^s) \subseteq D_{us}^x \quad (7.36)$$

Eq. (7.33) intuitively requires that no detail in  $F_{us}$  be lost by  $R_u.$  It is also required to establish the result. (A weaker form similar to Eq. (7.34) might suffice.)

Figure 7-19 shows a simple example that satisfies all the conditions except Eq. (7.33), but that fails to satisfy Eq. (7.35). In particular, vertex  $v$  does not satisfy Eq. (7.33), and  $\langle s, v \rangle$  is in the LHS but not the RHS of Eq. (7.35).

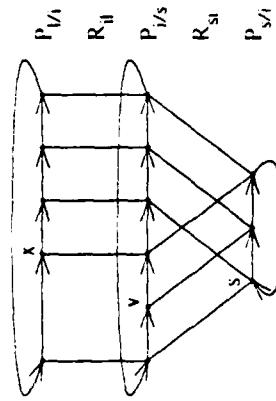


Figure 7-19: Need for Eq. (7.33)

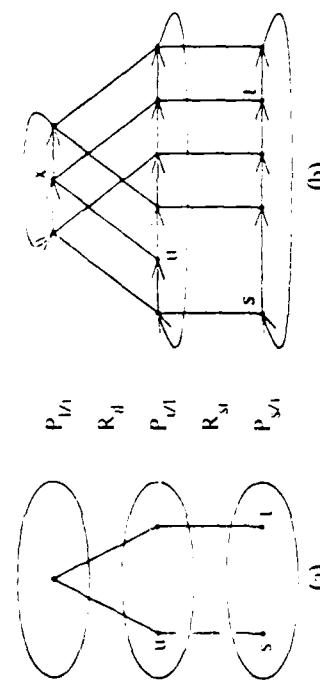


Figure 7-20: Need for Eq. (7.34)

Figure 7-20a illustrates Eq. (7.34). The existence of the solid edges implies the existence of the dashed edge for some vertex  $s.$  This condition is also required to establish the result. Figure 7-20b shows an example that satisfies all the conditions except Eq. (7.34), but that fails to satisfy Eq. (7.35). In particular, vertex  $u$  does not satisfy Eq. (7.34), and  $\langle s, v \rangle$  is in the LHS but not the RHS of Eq. (7.35).

We now present some lemmas. The lemmas assume the conditions given in Theorem 7.18. The proof of the theorem follows the lemmas.

**Lemma 7.19:**  $D_{sl} = D_{sjr} \square$

**PROOF:**  $D_{\vee I} \subseteq D_{\vee I'}$ . Given a vertex  $s$  in  $D_s$ , if  $v R_{sI} x$  for any vertex  $x$ , then by Eq. (7.28),  $s R_{sI} u$  for some vertex  $u$ . (A similar argument also shows  $D_{vI} \subseteq D_{vI'}$ .)

$D_{sI} \subseteq D_{sI'}$ : Prove by contradiction. Assume there is a vertex  $s$  in  $D_s$  such that  $s$  is in  $D_{sI}$  but not  $D_{sI'}$ . Then  $s R_{sI} u$  for some vertex  $u$  in  $D_I$ . This implies that  $u$  is in  $D_{vI}$ , but not in  $D_{vI'}$ ; if  $u$  were in  $D_{vI'}$  then  $s$  would be in  $D_{sI'}$ . This contradicts Eq. (7.33). QED.

**Lemma 7.20:**  $(D_{vI}^x \cap D_{vI'}) \subseteq D_{vI''}^x$   $\square$

**PROOF:** Show by contradiction. Suppose  $u \in (D_{vI}^x \cap D_{vI'})$ , but  $u \notin D_{vI''}^x$ . Since  $u$  is in  $D_{vI}$  but not  $D_{vI''}$ , then  $u F_{vI}^x v F_{vI''}^x w$  for some vertex  $w$ . Now  $D_{vI} \subseteq D_{vI'}$  by Eq. (7.33), so  $\{u, v\} \subseteq D_{vI'}$ . Therefore  $u F_{vI'}^x v$ . This contradicts  $u \in D_{vI}^x$ . QED.

Figure 7.21 illustrates the following lemma. This lemma states conditions under which the existence of a path in  $P_i$  implies the existence of a related path in  $P_i'$ .

**Lemma 7.21:** Given vertices  $u \in D_{vI''}$ ,  $w \in D_{vI'}$ , and  $v \in D_{vI}$  such that  $u \notin D_{vI''}^x$  and  $w R_{vI} F_{vI''}^x v$  along a path containing no vertices in  $D_{vI''}$  except  $w$  and possibly  $v$ . Then there exists a vertex  $v$  in  $D_{vI}$  such that  $u F_{vI}^x v R_{vI} x$  along a path containing no vertices in  $D_{vI}$  except  $u$  and possibly  $v$ . We have  $v \in D_{vI'}$  if and only if  $x \in D_{vI'}$ .  $\square$

**PROOF:** Assume vertices  $u$ ,  $w$ , and  $x$  as given. If  $x = w$ , then the result is immediate with  $v = u$ . Assume  $x \neq w$ . From Eq. (7.28) we know  $D_{vI''} \subseteq D_{vI'}$ , so

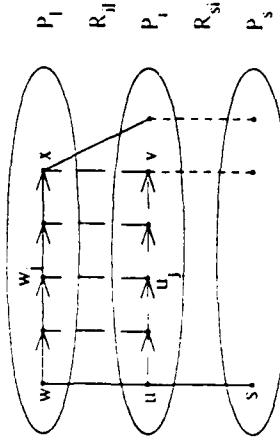


Figure 7.21: Lemma asserting existence of path in  $P_i'$

$w \in D_{vI'}$ . Then  $w F_{vI}^x v$  along a path containing no vertices in  $D_{vI}$  except  $w$  and possibly  $v$ . Let  $w_0, w_1, \dots, w_n$  be the sequence of vertices  $w_j$  in such a path, with  $n \geq 1$ ,  $w_0 = w$ ,  $w_n = x$ . We show by induction that we can find vertices  $u_j, \dots, u_n$  such that for  $0 < j \leq n$ ,  $u F_{vI}^x u_j F_{vI}^x w_j R_{vI} w_j$  along a path containing no vertices in  $D_{vI}$  except  $u$  (and  $u_n$  iff  $j = n$  and  $w_n \in D_{vI''}$ ).

First consider  $j = 1$ . Lemma 7.20 implies  $u \notin D_{vI''}$ , since  $u \in D_{vI}$ , but  $u \notin D_{vI''}$ . Then Eq. (7.29) gives  $u F_i + u_j R_{vI} w_j$  for some vertex  $u_j$ , along a path containing no vertices in  $D_{vI}$  except  $u$ . Since  $D_{vI} \subseteq D_{vI'}$  (Eq. (7.33)), only  $u$  and  $u_j$  in this path could be in  $D_{vI''}$ . If  $n = 1$  and  $w_n \in D_{vI''}$ , then Eq. (7.34) implies that  $u_j \in D_{vI''}$ . If  $n > 1$  or  $w_n \notin D_{vI''}$ , then  $u_j \notin D_{vI''}$ ; otherwise Eq. (7.28) would give  $w_j \in D_{vI''}$ , a contradiction.

Now assume the inductive assertion is true for  $j = j_0$  and consider  $j = j_0 + 1 \leq n$ . Suppose  $u_{j_0} \in D_{vI}^x$ . Then there is a don't-care vertex  $z$  in  $D_i^x$  such that  $u_{j_0} F_i^* z$  along a path not including any vertices in  $D_{vI}$  except  $u_{j_0}$ . Since

$D_{\nu s} \subseteq D_{\nu l}$ , the path also includes no vertices in  $D_{\nu s}$  except possibly  $u_{\rho}$ , but  $u_{\rho} \notin D_{\nu s}$  since  $\rho < n$ . Then from the inductive assertion,  $u F_i + z$  along a path including no vertices in  $D_{\nu s}$  except  $u$ . This implies  $u \in D_{\nu s}^x$ , a contradiction. Therefore  $u_{\rho} \notin D_{\nu l}^x$ .

Now Eq. (7.29) implies  $u_{\rho} F_i + u_{\rho+1} R_{il} w_{\rho+1},$  for some vertex  $w_{\rho+1}$ , along a path containing no vertices  $D_{\nu l}$  except  $u_{\rho}$  and  $w_{\rho+1}$ . From the argument at the end of the preceding paragraph,  $u_{\rho+1} \in D_{\nu s}$  if and only if  $n = \rho + 1$  with  $w_n \in D_{\nu s}$ .

This completes the induction. The inductive premise satisfies the lemma for  $j = n$  with  $v = u_n$ . QED.

**Lemma 7.22:**  $R_h(D_{\nu s}^x) \subseteq D_{\nu s}^x \quad \square$

**PROOF:** Consider arbitrary vertices  $w$  and  $u$  such that  $w \in D_{\nu s}^x$  and  $w R_h u$ . It will suffice to show that  $u \in D_{\nu s}^x$ . Since  $w \in D_{\nu s}^x$ , there exists vertex  $y \in D_{\nu s}^x$  such that  $w F_i^* y$  along a path containing no vertices in  $D_{\nu l}$  except  $w$ . Let  $x$  be the last vertex of this path that is in  $D_{\nu l}$ . Now  $x \in D_{\nu l}^x$ .

Theorem 7.18 says that if program  $P_j$  is consistent with  $x$ , which in turn is consistent with  $P_s$  according to Eqs. (7.14) and (7.15), then we may conclude that  $P_j$  is consistent with  $P_s$  without comparing  $P_1$  against  $P_s$ . This result holds provided that relations  $R_{s0}$  and  $R_{hl}$  satisfy the conditions in Eqs. (7.33) and (7.34).

The proof will be by contradiction. Eq. (7.34) implies  $u \in D_{\nu s}$ . Suppose  $u \notin D_{\nu s}^x$ . Then by Lemma 7.21, there exists  $v$  such that  $u F_i^* v R_{il} x$  along a path containing no vertices in  $D_{\nu s}$  except  $u$ . By Eq. (7.30),  $v \in D_{\nu l}^x$ , so there exists  $z \in D_{\nu l}^x$  such that  $v F_i^* z$  along a path containing no vertices in  $D_{\nu l}$  except  $v$ . Since  $D_{\nu s} \subseteq D_{\nu l}$  (Eq. (7.33)), we now have  $u F_i^* z$  along a path containing no vertices in  $D_{\nu s}$  except  $u$ . Therefore  $u \in D_{\nu s}^x$ . This contradicts the supposition and proves the desired result. QED.

#### PROOF OF THEOREM 7.18: To prove Eq. (7.35).

$$R_{sl} F_{\nu s} \subseteq (F_{s0} R_{s0}) \cup (D_{s0}^x \times D_{\nu s}) \quad (7.35)$$

make the following observations.

(1)  $s R_{sl} F_{\nu s} x$ , where  $s \notin D_{\nu s}^x$ , by assumption. It will suffice to show that

$$s F_{s0} R_{s0} x.$$

- (2)  $s R_{sl} u R_{il} w F_{\nu s} x$  for some vertices  $u$  and  $w$ , from (1) by Eq. (7.28).
- (3)  $u \notin D_{\nu s}^x$ . Otherwise, Eq. (7.32) would imply  $s \in D_{\nu s}^x$ . This would contradict (1), since Lemma 7.19 and Eqs. (7.6) and (7.5) imply  $D_{s0}^x = D_{s0}^x$ .
- (4)  $u F_{\nu s} v R_{il} x$  for some vertex  $v$ , from (2) and (3) by Lemma 7.21, with  $v \in D_{\nu s}$  since  $x \in D_{\nu s}$ .

- (5)  $s F_{s0} t R_{s0} v$  for some vertex  $t$ , from (2) and (4) by Eq. (7.31), since  $s \notin D_{\nu s}^x$  (see (3)).
- (6)  $s F_{s0} R_{s0} x$ , from (4), (5), and Eq. (7.28). Lemma 7.19 and Eq. (7.6) imply  $F_{s0} = F_{s0}$ .

This proves Eq. (7.35). Eq. (7.36) follows from Lemmas 7.19 and 7.22 and Eqs. (7.28) and (7.32). QED.

### 7.12. Summary

Milner and Brand proposed definitions of program consistency for two particular types of program models, deterministic and nondeterministic programs without don't-cares. This chapter considered these model types plus the corresponding models with don't-cares. Section 7.8 gave a criterion for consistency suitable for all four models and compatible with the definitions proposed by Milner and Brand.

Brand argued for the inclusion of closure in the criterion for practical reasons. We saw that this leads to the verification of clearly incorrect designs, but the criterion without closure is difficult to test. Section 7.9 added closure to the criterion in a limited way. Section 7.10 showed the correspondence between this new criterion and the simple criterion for deterministic models.

Finally, Section 7.11 described hierarchical verification. Milner defined a "strong simulation" consistency criterion that he described as a transitive reflexive relation leading to a "quasi-ordering" of programs. This notion is similar to that of hierarchical verification, but applies only to deterministic models without don't-cares with all inputs presented before execution, and execution terminating with a final output value. That model is not suitable for hardware descriptions. Section 7.11 showed that hierarchical verification is possible for more general models.

Several problems remain to be solved to achieve a unified, useful theory.

These problems are outlined in Chapter 9.

AD-A136 466

VERIFICATION OF HARDWARE DESIGN CORRECTNESS: SYMBOLIC  
EXECUTION TECHNIQUE..(U) STANFORD UNIV CA COMPUTER  
SYSTEMS LAB W E CORY JUN 83 TR-83-241 DAAG29-80-K-0046

2/1

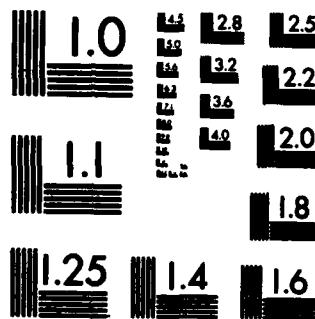
UNCLASSIFIED

F/G 9/2

N

END  
DATA  
FRAMED

2-84  
DTIC



MICROCOPY RESOLUTION TEST CHART  
NATIONAL BUREAU OF STANDARDS-1963-A

## Chapter 8

### Partitioned Verification

A hardware specification to be compared against a proposed implementation may consist of several interconnected components. In that case, it is desirable to verify the design for each component individually.

For example, consider Figure 8-1. Part (a) represents a system specification with two components, a CPU and a memory. Part (b) represents a proposed design for this system. We may verify the CPU design ignoring the memory component, and then similarly verify the memory design. The verification of each component individually is a partitioned verification. This is especially desirable if one component type is used many times, or if the design is specified hierarchically.

The implicit assumption in partitioned verification is that if each individual component design is correct, then the combination of these designs is also correct. Brand demonstrated this property under limited conditions [Br78]. This thesis will not expand on this theoretical base, although a number of improvements are needed as mentioned in Chapter 9. However, this chapter will discuss the impact of the requirements for Brand's theorem.

Section 8.1 reviews Brand's result. Section 8.2 presents an example that

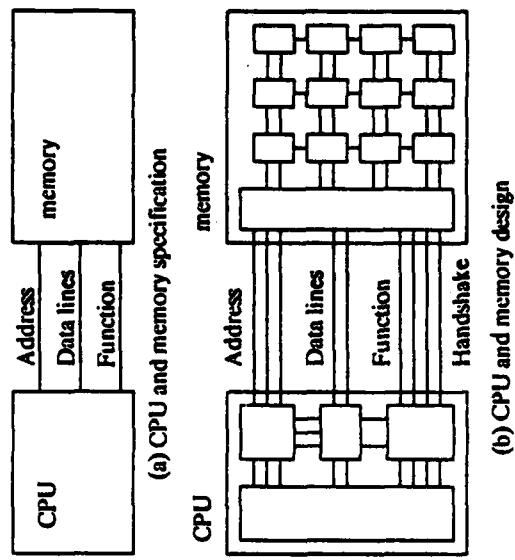


Figure 8-1: Changes in data type

demonstrates the characteristic of a design that may prevent the designer from satisfying the requirements of this theorem. Finally, Section 8.3 describes a modification of the design process that allows the requirements to be satisfied. This modification uses transients, a concept proposed by Hill [Hi80].

#### 8.1. Review of Partitioned Verification

This section repeats Brand's theorem for partitioned verification. No new proofs are presented; the purpose of this section is to explain in more detail what partitioned verification is and to highlight the restrictions under which the theorem holds.

This theorem allows partitioned verification for a design with two components. This result easily extends to an indefinite number of components, and we state the result below for the general case. For brevity, we omit the elegant notion of extending relations in the development of the program model.

In the following discussion, the definition of "referring to  $\Delta'$ " (Section 8.1.2) is stronger than Brand's. We also assume a restriction from Section 7.5 not required by Brand. These changes are required to repair a minor error in Brand's proof. We do not repeat or correct that proof here. The interested reader may consult Appendix B, which discusses the correction of the proof. The appendix assumes familiarity with Brand's technical report [Br78].

### 8.1.1. Brand's Program Model

The program (hardware description) consists of the combination of  $n$  components. Intuitively, a component is an autonomous block specified separately from the other blocks. We use graphs to represent programs, as in Chapter 7. Assume nondeterministic models with no don't-cares.

In the graph model, component  $i$  is a program  $P_i^l = \langle D_a, F_a^l \rangle$ . (Recall that subscript "l" denotes program "s" or "t".) By convention, a superscript will identify a component. Then  $P_a$  is the union of the components;  $P_a = \langle D_a, U_j^a, I_j^a, F_a^l \rangle$ . The state space  $D_a$  for program  $P_a$  is split among the  $n$  components:

$$D_a = A_a \times D_a^1 \times D_a^2 \times \dots \times D_a^n$$

Then  $D_a^j$  is the set of local variables for component  $j$ , while  $A_a$  is the set of variables shared by the components. If  $s \in D_a$ , then  $s^j$  is the element of  $s$  in  $D_a^j$ ;  $s^0$  is the element of  $s$  in  $A_a$ . Also,  $s(r, v)$  will denote static  $s$  with element  $s^j$  set to  $r$ .

169

170

$F_a^l$  describes the actions of component  $l$ , and so can modify only its own local variables in  $D_a^l$  and the shared variables in  $A_a$ . These actions are independent of the values of the local variables for any other component. The following assertion expresses these properties:

$\forall$  states  $s, t$  and integers  $i, j \models F_a^l i$ :

$$((s^j \Rightarrow r) \supset (j=i \text{ or } j=0)) \text{ and}$$

$$((j \neq i \text{ and } j \neq 0) \supset (V, v \in D_a^j; s(r, v) \models F_a^l (r, v)))$$

Finally, relation  $R_{ab}$  is defined in a piecewise fashion. We have  $R_{ab}^i \subseteq D_a^i \times D_b^i$ ; also  $R_{ab}^0 \subseteq A_a \times A_b$ . Then  $s R_{ab} t$  if and only if, for  $0 \leq i \leq n$ ,

$$R_{ab} = R_{ab}^0 \& R_{ab}^1 \& \dots \& R_{ab}^n = \bigcup_{i=0}^n R_{ab}^i$$

Now  $D_{a/b}, D_{a/b}^1, F_{a/b}$ , etc. are defined by Eq. (7.6). We assume the requirement from Section 7.5 for all the edge sets  $F_a^l$ ; namely, every path in  $P_a^l = \langle D_a, F_a^l \rangle$  leading from any vertex in  $D_{a/b}$  must eventually reach a vertex in  $D_{a/b}$ .

### 8.1.2. Definitions

A "stopping point of  $D_{a/b}$  for component  $l'$ " is a state  $s$  with  $s^l \in D_{a/b}$  and  $s^0 \in A_{a/b}$ . Note that a state in  $D_{a/b}$  is a stopping point of  $D_{a/b}$  for every component.

A transition  $s F_a^l r$  refers to  $\Delta_a$  if there exist  $r \in D_a$  and  $d_j, d'_j \in A_a$  such that  $s(d_j, v) \models F_a^l r$  and either (i)  $d_j \Rightarrow r^0$  or (ii)  $\langle s(d_j, \emptyset), r(d_j, v) \rangle \notin F_a^l$ .

We say "there is a stopping point of  $D_{a/b}$  for component  $l$  between any two references to  $\Delta_a$  by  $F_a^l$ " if whenever

$\exists_0 F_a^1 s_1, F_a^2 \dots F_a^i s_i F_a^{i+1} \dots F_a^m s_m$  /  
 $(m \geq 1)$ , where the first and last transitions refer to  $\Delta_a$ , then some  $s_k$  in the path,  
 $1 \leq k \leq m$ , is a stopping point of  $D_{a/0}$  for component  $a$ .

### 8.1.3. Theorem for Partitioned Verification

Braud's theorem, with the minor changes noted above, is as follows.

**Theorem 8.1:** Assume partitioned, nondeterministic programs  $P_a$  and  $P_1$  without don't-care, each with  $n$  components, and relation  $R_a \subseteq D_a \times D_1$  all as described in Section 8.1.1. Suppose the following conditions hold for  $j$ ,  $1 \leq j \leq n$ :

- $R_a F_{a/1} \subseteq F_{1/a}^j \circ R_a$
- there is a stopping point of  $D_{a/0}$  for component  $j$  between any two references to  $\Delta_a$  by  $F_a^j$ .

Then the following result holds:

$$R_a F_{a/1} \subseteq F_{1/a}^j \circ R_a \quad . \square$$

One restriction above is that any two successive references to  $\Delta_a$  must be separated by a stopping point. The other restriction is that every path  $F_a^j +$  starting from a stopping point of  $D_{a/0}$  for component  $j$  must eventually reach another such (or the same) stopping point.

It is likely that any theory allowing partitioned verification will include these restrictions in some form. The requirement that every path from a stopping point reaches a stopping point ensures that  $P_{a/0}$  adequately represents the behavior of  $P_a$ .

The requirement that a stopping point separates two references to  $\Delta_a$  ensures that protocols for communication between components are strongly enough defined to permit the independent verification of components. Each of these properties is important for partitioned verification.

The example in the next section illustrates the problem these restrictions cause in a hierarchical design.

### 8.2. Partitioned Verification Problem

In a hierarchical design, different descriptions may use different levels of abstraction. This means that two descriptions to be compared may use different representations of the same data. If the change in the level of abstraction affects only the static encoding of data, then a simulation relation may readily satisfy the restrictions for Brand's theorem. However, the introduction of new timing details may hinder the fulfillment of these requirements.

Consider the example in Figure 8-1. Part (a) represents a specification of a CPU, memory, and bus at a high level of abstraction. Details of the bus protocol may be omitted or simplified. In contrast, the refined description in Figure 8-1b may specify the bus structure and protocol in greater detail.

Figure 8-2 shows the timing of a write to memory for these two descriptions. The CPU  $C/r$  is an internal control variable in the CPU.  $Mem\ Data$  is the contents of the memory being written. The *Address* and *Data* lines are controlled by the CPU. *Write Ctrl*, *Sinbe*, and *Act* are parts of the *Funcion* lines in the two descriptions. Note the labeling of states by numbers in part (a) and letters in part (b).

Figure 8-2. The state labels are taken from Figure 8-2. It is desirable to verify Figure 8-4a against Figure 8-3a and Figure 8-4b against Figure 8-3b, and so verify the design. Of course, we need a simulation relation.

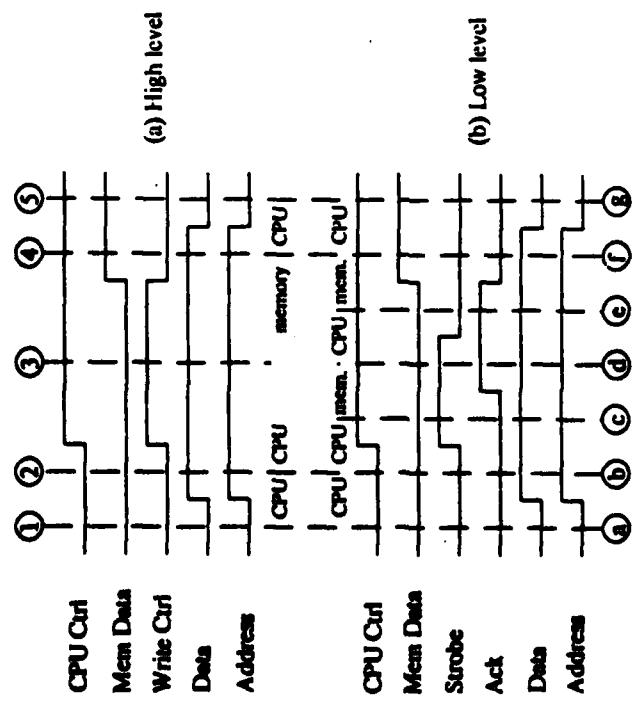


Figure 8-2: Timing diagrams for write cycle

The specification (part (a)) uses simple timing. The *Write Cntl* does not behave like a bus line in the usual sense; the CPU asserts it and the memory negates it, so it behaves more like a latch. This does not necessarily dictate how the write cycle will be implemented. In the design (part (b)), the CPU controls *Strobe* and the memory controls *Ack*.

Now Figures 8-3 and 8-4 show the graph models for the write cycles in

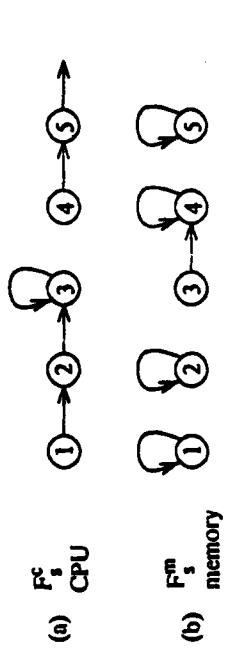


Figure 8-3: Graph models for high-level write cycle

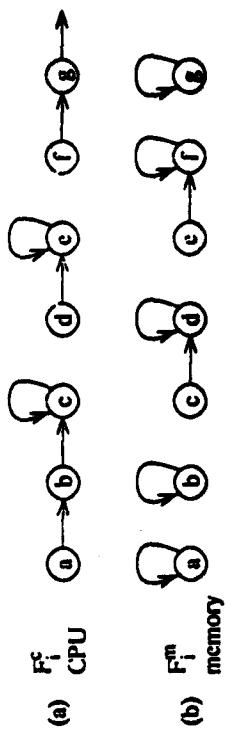


Figure 8-4: Graph models for low-level write cycle

States (1), (2), (4), and (5) naturally correspond to states (a), (b), (1), and (2), respectively. This leaves state (3) to correspond to (c), (d), or (e). But each of these choices is unsatisfactory, resulting in incomplete tests for correctness. For example, if we let (3)  $R_{sg}(d)$ , then the verification process does not check the behavior of  $F_i^c$  at (b), (c), (d), or (e), and it does not test  $F_s^m$  at (c) or (e).

Letting (3) correspond to combinations of states (c), (d), and (e) does not improve matters. For example, assume  $R_{\text{g}}(\{(3)\}) = \{(c), (d), (e)\}$ . Now suppose that through a design error, we have (d)  $F_i^c$  (d) instead of (d)  $F_i^c$  (e). The resulting design would not operate correctly, but the application of either Eq. (7.4) or (7.7) to  $F_{S/I}^c$ ,  $F_{S/I}^d$ , and  $R_{\text{g}}$  would flag no error.

The problem here is that program  $P_i$  uses finer timing detail than  $P_g$  does. For such cases, a simulation relation is likely either to not satisfy the restrictions for Theorem 8.1 or to not distinguish between correct and incorrect actions in the implementation. In the above example, the former occurred when state (3) corresponded to just one of (c), (d), and (e); the latter occurred when (3) corresponded to all three states.

The use of the same timing detail in two descriptions being compared facilitates the fulfillment of the restrictions for Theorem 8.1 by an appropriate simulation relation. A modification to the top-down design methodology can ensure this correspondence. In this new method, the designer will still have the freedom to use varying levels of abstraction. However, the design will be refined in such a way that any two descriptions to be compared will always use the same timing. Partitioned verification may then be possible.

### 8.3. Modification of the Design Methodology

In a top-down design methodology, each refinement may introduce finer structural and behavioral detail. As noted above, the addition of new timing detail at some stage of the design may hinder the fulfillment of the restrictions for

partitioned verification. The design methodology may be expanded to handle this problem.

The expanded methodology is presented below. Block diagrams are used for illustration. Blocks represent components. Lines interconnecting the blocks represent variables shared by the components.

The first step in the expanded methodology is to write a behavioral specification for a system to be designed. This is represented in Figure 8-5. This system consists of several components, and it is desirable to verify the design for this system in a partitioned fashion. Due to the difficulties discussed in the previous section, this may not be possible if the next refinement introduces new timing details.

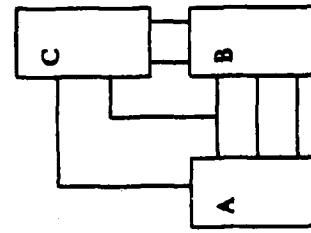


Figure 8-5: Initial system specification

To overcome this problem, if the designer wishes to use more detailed timing

177      in the next refinement, then he must first alter the specification by the addition of translators. A translator is similar to a component description, but its significance is different. While a component description represents the behavior of some hardware, a translator relates timing behavior at one level of abstraction to that at another level. Figure 8-6 shows the original specification with translators added. The translators are represented by triangles.

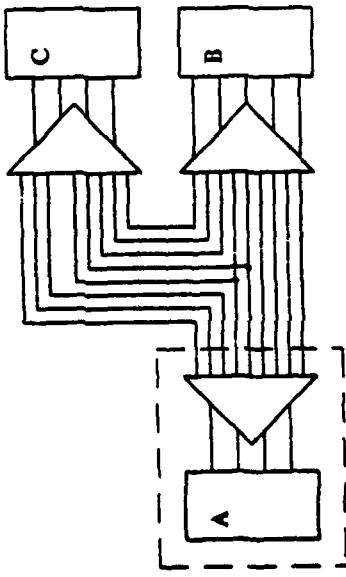


Figure 8-6: Specification with translators

The next step is to verify the correctness of the translators. The criterion for correctness is that the presence or absence of translators should be transparent in the operation of the system. The verification step consists of showing that the description in Figure 8-6 ( $P_i$  for this step) is consistent with the specification ( $\langle P_s \rangle$ ) in Figure 8-5.

The components in Figure 8-5 use high-level timing. After defining and verifying the translators, new component specifications using new low-level timing

178

may be taken from Figure 8-6. For example, the circuit inside the dotted box is a new specification for component A. It uses low-level timing for its accesses to variables shared with other components.

The designer now makes the next refinement in the design. Figure 8-7 represents the design for component A. Since this design and the specification from Figure 8-6 use the same timing for accesses to variables shared with other components, the designer should now be able to define a simulation relation satisfying the restrictions for partitioned verification. The design in Figure 8-7 may be verified separately from the designs for the other components.

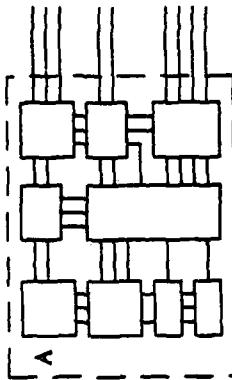


Figure 8-7: Interface design refinement

#### 8.4. Summary

Given a behavioral specification  $P$  for a system to be designed, the expanded top-down design methodology may be summarized as follows.

If the next refinement in the design will use more detailed timing, then define translators that convert the old high-level timing into the new timing. Insert these

translators into  $P$ , yielding description  $Q$  in which components communicate with each other using the new timing. Verify  $Q$  against  $P$ .

Now generate design  $R_q$  for each component  $q$  in  $Q$ . The combination of designs  $R_q$  yields design  $R$ . Define a simulation relation between descriptions  $R$  and  $Q$  that satisfies the requirements for partitioned verification. Since the descriptions  $R$  and  $Q$  use the same timing, this should be possible. Verify  $R$  against  $Q$ . This can be done in a piecemeal fashion; verify  $R_q$  for each component  $q$ .

Repeat the above cycle for each design  $R_q$  until the lowest-level design uses available components or components to be designed using other design automation tools.

This is quite similar to the original top-down design procedure. The only change is the definition and insertion of translators when needed. If the designer does not wish to introduce new timing details in some refinement of the design, then descriptions  $P$  and  $Q$  are identical. In that case, the new methodology is identical with the old. Hence, the new design style preserves the flexibility of the old style.

The new methodology is similar to a common approach in the design of systems with components that communicate over a bus. The definition of translators corresponds to the definition of a bus protocol. The design for each component is then validated against its specification and the bus protocol. The designer assumes that components which observe the bus protocol will then communicate with each other properly.

The new design methodology does require some extra work on the part of the designer; namely, the specification of translators which define the refinements in timing. However, the new methodology does not make new work where no work existed before. At some point in the old methodology, the designer had to define the refinements in timing before using them. The new methodology simply requires that these refinements be expressed formally in translator definitions.

Still defined translators in ADL/B to facilitate multi-level simulation [H80]. The use of translators to facilitate partitioned verification represents a new application of this concept.

## Chapter 9

### Summary

#### 9.1. Contributions

This thesis investigates two problems: reduction of case splitting in the symbolic execution of event-driven simulations, and the criterion for the consistency of two hardware descriptions in functional verification.

##### 9.1.1. Non-selective Traces

An event-driven simulation consists of a collection of independent components with some means of communication between them. The simulator commonly will not simulate a component unless the value of a particular variable(s) changes. This is called "selective trace" [U65]. When the variable in question receives a new value, the simulator must test whether the new value is different from the old to determine whether to simulate the component. If a symbolic execution cannot uniquely resolve this test to *TRUE* or *FALSE*, then the single execution will split into several executions (cases), one for each branch followed [HK76, Ki76a].

The presence of distinct, independent components aggravates this problem. If each of  $I$  components always splits the simulation into  $N$  cases in every time step, then after  $J$  time steps, the total number of cases will be  $N^I J$ .

Chapter 4 suggests the substitution of non-selective trace for selective trace to eliminate these case splits. The non-selective trace simulates a component whenever the proper variable is written, whether or not the operation changes the value. Of course, this substitution is correct only if the non-selective version is functionally equivalent to the selective version.

In the model for event-driven simulation given in Chapter 3, the substitution of non-selective trace for selective trace is correct for the construction typically used to describe combinational elements. More generally, consider a component described by an infinite loop that waits at one point for an event. Under certain conditions, the use of non-selective trace is appropriate for such loops. The major condition is the following: If the main body of the loop were executed even though no event occurred, then it would not modify any variables.

The annotation of a loop with assertions can facilitate a proof that the loop satisfies the conditions for use of non-selective trace. An inductive assertion analysis in turn proves the assertions. However, the assertions are often cumbersome and redundant. Section 4.2.6 gives a sufficient test that checks the loop without the formulation of the redundant assertions. This test symbolically executes the body of the loop twice in succession under certain assumptions, then checks that the second execution did not alter any variables.

Previous work did not apply symbolic execution to event-driven simulation. In particular, it did not investigate the substitution of non-selective trace for selective trace to reduce case splitting.

### 9.1.2. Case Merging

Conditional expressions can symbolically express the computational results of several parallel paths [CH79]. For example, the expression

$C_1$  and  $C_2$  or not  $C_1$  and  $C_3$

gives the value of the Boolean variable  $X$  after

```
If  $C_1$  then  $X := C_2$ 
else  $X := C_3$ ;
```

regardless of the path selected in the if statement. A flow graph similar to that used by Abu and Ullman for dataflow analysis [AU77] directs the generation of the conditional expressions.

In the simulation model of Chapter 3, each component is represented by a coroutine. The following steps facilitate the use of conditional expressions to combine the results of multiple paths in coroutines.

1. Allow the initial and final values of a coroutine's program counter to be symbolic (not constant).
2. Do a topological sort of a coroutine's flow graph to determine the order in which the coroutine's basic blocks should be processed.

Using these extensions, Section 5.2.2 describes an analysis of a hardware description yielding a new equivalent description. The new description contains no conditional branches that would introduce case splits in a symbolic execution. This allows a symbolic execution that combines all parallel paths into a single execution; hence the term "case merging."

This construction works for components in which every loop is broken by a

statement that waits for an event to occur. A modification extends the construction to apply to components with unbroken loops if the loops unfold.

The major problem with case merging is the potentially excessive complexity of the symbolic expressions generated in a symbolic execution. Section 5.4 identifies a number of optimizations that will simplify the symbolic expressions. Chapter 6 considers the tradeoff between the reduction in the number of cases and the increase in symbolic expression complexity. Case merging is most likely to be profitable for functional verification when the simulation relation is simple and the expression complexity grows less than linearly as the number of cases in a case splitting execution.

Cheatham et al. first proposed the use of conditional expressions to symbolically express the results of parallel paths. They applied this technique in a static analysis of programs including procedures but not coroutines. The contributions of Chapters 5 and 6 are (i) enhancements of case merging for coroutines, (ii) alternative treatment of unfolding loops, (iii) optimizations, and (iv) evaluation of usefulness of case merging for functional verification.

### 9.1.3. Consistency of Hardware Descriptions

The functional verification problem is that of comparing two hardware descriptions to determine their consistency. One description is a specification while the other is a proposed design. Milner and Brand considered similar questions, representing hardware descriptions (or programs in Milner's case) by graph models [Mi71, Br78]. This thesis extends their work.

Milner assumed a deterministic model with no don't-cares. Brand assumed a nondeterministic model, also without don't-cares. Section 7.8 presents a simple criterion for consistency that applies to both deterministic and nondeterministic models, with or without don't-cares. ("Simple" distinguishes this criterion from the extended criterion below.)

In some cases, a description may specify an action by a single operation while the other does not, or uses a loop to specify the same action. Brand's criterion included an extension to facilitate the verification of such cases by symbolic execution. However, its extension allows the verification of incorrect designs. Section 7.9 presents an alternate extension that also facilitates verification of descriptions with loops.

For deterministic models, Section 7.10 demonstrates the compatibility of this new extension with the new simple criterion from Section 7.8. This supports a claim that the new extension retains intuitive notions of correctness.

Milner demonstrated that under certain conditions, his criterion of consistency is transitive, a desirable property for hierarchical hardware verification. However, his model was tailored for software, not hardware. He partitioned the vertex (data) space and specified restrictions on the simulation relation relating the two graph models. He also required all inputs to be presented before execution started. These restrictions are not appropriate for hardware descriptions. Section 7.11 proves the transitivity of the new simple criterion under less restrictive conditions.

Finally, Brand proved a theorem allowing verification of a design in a

piecewise fashion. One of the conditions for this theorem prevents piecewise verification of a design if it contains timing details not included in the specification. Chapter 8 modifies the hierarchical design methodology to permit piecewise verification of a refinement that introduces new timing details. The new methodology specifies that before making such a refinement, we modify the specification so that it uses the same timing as the proposed refinement.

## 9.2. Suggestions for Further Research

### 9.2.1. Case Merging

It is desirable to obtain more definitive results on the merits of case merging vs. case splitting. The effects of the choice of canonical form, simplification, or efficient expression comparison heuristics make such an analysis difficult.

### 9.2.2. Criteria for Consistency

Ideally, we would have one criterion for consistency which achieves all of the following:

1. applies to deterministic and nondeterministic models with and without don't-cares,
2. facilitates verification of descriptions with loops,
3. supports hierarchical and partitioned verification.

However, we now have three criteria: the new simple test, new extended test, and Brand's test. These tests have the following properties:

- The new extended test satisfies (1) and (2) above, but (3) has not been demonstrated.

- The new simple test supports hierarchical verification.
- Brand's test supports partitioned verification, but only for nondeterministic models.
- For deterministic models, the new simple and extended tests are in some sense compatible.

It is desirable to close the gap between the achieved theory and the ideal, unified theory. Some possible approaches follow:

1. Prove partitioned verification for the new simple test with deterministic models. Use compatibility of simple and extended tests to prove hierarchical and partitioned verification for the extended test. This unifies the consistency theory for deterministic models.

2. Prove partitioned verification for the new simple test with nondeterministic models. Prove compatibility of simple and extended tests for nondeterministic models. Use this compatibility to prove hierarchical and partitioned verification for the extended test. This unifies the theory for nondeterministic models.

3. Prove hierarchical and partitioned verification for the new extended test with a particular model. This unifies the theory for the selected model.
4. Find conditions under which Brand's test is equivalent to the extended test. Prove hierarchical verification for Brand's test with nondeterministic models. This unifies the theory for nondeterministic models, with conditions that determine whether Brand's test is appropriate.

- Brand's result for partitioned verification allows only one resource to be shared between the several components, with only one reference to the resource between stopping points in the simulation relation. This is likely to be impractical in hierarchical hardware designs. Relaxation of this restriction is desirable.

### 9.2.3. Symbolic Execution for Functional Verification

This thesis considered symbolic execution separately from the definition of consistency for functional verification. Development is required on the interpretation of symbolic execution results to verify consistency. One problem is to guarantee that the symbolic executions adequately cover the possible input sequences.

### 9.2.4. Incorporation of Other Techniques for Verification

In some cases, symbolic execution may not be a suitable technique for comparing two hardware descriptions. The correctness of a particular component in a design may rely on significant mathematical theory not easily incorporated into symbolic execution theorem provers. For example, the comparison by symbolic execution of an arithmetic description of a 32-bit adder with a Boolean description of the same adder is difficult, due to the complex Boolean expressions generated. This problem is compounded for floating point hardware. Nonetheless, designs known to be correct are available, and the designer should be free to use them. In such cases, it is desirable to be able to establish the correctness of these components by independent means, and then to incorporate these results for the remainder of the proof.

### 9.2.5. Strengthening the Weak Links

In a top-down design, the designer generates a chain of hardware descriptions, starting with the first formal specification of the system behavior, and ending with a formal description of the physical implementation. Each verification of two successive descriptions forms a link in the verification chain ensuring the functional correctness of the design. However, two special cases exist in which one of the two descriptions being compared is not a formal description. These special cases are the two ends of the chain.

To have confidence in the proof of correctness, the designer must be able to show that (1) the first formal specification accurately captures the intent of the informal English specification, and that (2) the last formal specification accurately models the behavior of the physical hardware. The addition of assertions to hardware descriptions, perhaps by the application of temporal logic [Bo82, MP81], may allow the proof of desirable properties about the descriptions by performance verification. This in turn would increase the designer's confidence in the design.

## Appendix A

### Notation

#### A.1. Formulae

The expression " $a(x|y)$ " will denote  $a$  with the value of  $x$  and  $y$  (if  $a$  is a record) or  $a[y]$  (if  $a$  is an array) set to  $x$ . Vertical bar " $|$ " in formulae may be read as "such that".

When a colon ":" appears in formulae, it will generally follow a specification of a range or set of values. The expression " $a: b$ " then means "for each item in  $a$ , the expression or predicate  $b$ ".

#### A.2. Relations

Given sets  $W$  and  $Y$ , a subset  $S$  of  $W \times Y$  (the Cartesian product of  $W$  and  $Y$ ) is called a binary relation on  $W \times Y$ . We write " $w S y$ " as an alternate notation for " $\langle w, y \rangle \in S$ ".

If  $W$ ,  $Y$ , and  $Z$  are sets with relations  $S$  on  $W \times Y$  and  $T$  on  $Y \times Z$ , then  $ST$  is the "composition" of the relations:

$$ST = \{ \langle w, z \rangle \in W \times Z \mid (\exists y \in Y) (w S y \wedge (y T z)) \}$$

The following statements are then synonymous:

$$\begin{aligned} & \exists y \in Y \mid ((w S y) \wedge (y T z)) \\ & \exists y \in Y \mid w S y \wedge T z \\ & w S T z \end{aligned}$$

We will also use a closure notation, given relation  $R$  on  $Z \times Z$ :

$$\begin{aligned} R^* &= \{ \langle z_1, z_2 \rangle \in Z \times Z \mid \\ &\quad (\exists \text{ a sequence of one or more elements } z_1, z_2, \dots, z_n \in Z) \\ &\quad z_1 = z_1 R z_2 R \dots R z_{n-1} R z_n = z_2 \} \\ R^+ &= R R^* \end{aligned}$$

Finally, given relation  $S$  on  $W \times Y$ , define the inverse relation:

$$S^{-1} = \{ \langle y, w \rangle \in Y \times W \mid w S y \}$$

The image of  $W$  under  $S$  is denoted by  $S(W)$ :

$$S(W) = \{ y \in Y \mid (\exists w \in W) (w S y) \}$$

For an element  $w$  in  $W$ , we may write " $S(w)$ " to denote  $S(\{w\})$ .

These conventions for relations can be used with edge sets and functions, since they may be regarded as relations.

## Appendix B

### Notes on Proof of Theorem 8.1

This appendix describes a correction of an error in the proof of Lemma 2 in Brand's technical report [BR78]. It is assumed the reader has access to the report. The lemma, its proof, and introductory material will not be repeated here. This discussion uses Brand's notation. The new restrictions from Section 8.1 are repeated here using Brand's notation:

**Definition 1:** Let  $\langle D \times \Delta, C, S \times \Sigma \rangle$  be an annotated program. A transition  $\langle x_i, d_j \rangle \in C \langle x_j, d_j \rangle$  "refers to  $\Delta'$ " if and only if there exist  $x_j \in D$  and  $d_j, d_j' \in \Delta$  such that  $\langle x_i, d_j \rangle \in C \langle x_j, d_j \rangle$  and either (i)  $d_j \neq d_j'$  or (ii)  $\langle x_j, d_j' \rangle \in C$ .

□

This definition is stronger than Brand's. It says that if some edge  $\langle x_i, d_j \rangle \in C \langle x_j, d_j \rangle$  refers to  $\Delta$  by Brand's definition, then all edges incident out of  $\langle x_i, d_j \rangle$  also refer to  $\Delta$ , as do all edges incident out of corresponding vertices  $\langle x_j, d_j \rangle$ . This makes the notion of "referring to  $\Delta'$ " a property more of the vertex element  $x_j$  than of individual edges, an intuitively appealing change which facilitates the proof of Lemma 2.

**Requirement 1:** Consider annotated program  $\langle D, C, S \rangle$ . We require that every

path  $C^*$  leading from any vertex in  $S$  must eventually reach a vertex in  $S$ . This requirement applies both to  $C_1$  and to  $C_2$ . □

#### B.1. Counterexamples

First we show two counterexamples that disprove Lemma 2 and show the need for the definition 1 and requirement 1 above.

Figure B-1 shows edge sets  $C_1$  and  $C_2$  for the first counterexample. State set  $D_1$  is  $\{0,1,2,3\}$ ,  $D_2$  is  $\{0\}$ , and  $\Delta$  is  $\{0,1\}$ . The stopping point sets are  $S_1 = \{0,1,3\}$ ,  $S_2 = \{0\}$ , and  $\Sigma = \{0,1\}$ . (Brand's treatment of stopping points is different from that in Section 8.1, due to his notion of extending relations.) In the graph of  $C_1$  in Figure B-1, states with double circles are stopping points for  $C_1$ . Edges marked by a dot refer to  $\Delta$  according to Brand's definition. The state name is  $x y d$ , where  $x \in D_1$ ,  $y \in D_2$ , and  $d \in \Delta$ .

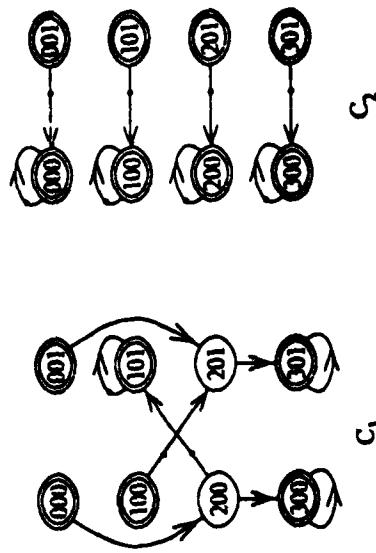


Figure B-1: First counterexample to Lemma 2

This example satisfies the conditions for Lemma 2, using Brand's definition of referring to  $\Delta$ . It also satisfies requirement 1 above for both  $C_1$  and  $C_2$ . In fact, it satisfies a stronger condition: Every state that is not a stopping has an ancestor and a descendant that are stopping points. But Lemma 2 fails for the sequence  
 $100\ C_1\ 201\ C_2\ 200\ C_1\ 101$ .

This example does not satisfy the conditions for Lemma 2 using our stronger definition of referring to  $\Delta$ , definition 1. Edges  $\langle 200, 300 \rangle$  and  $\langle 201, 301 \rangle$  in  $C_1$  refer to  $\Delta$  by this definition.

Figure B-2 shows edge sets  $C_1$  and  $C_2$  for the second counterexample. State sets  $D_1$ ,  $D_2$ , and  $\Delta$  are  $\{0,1\}$ . The stopping point sets are  $S_1 = \{0,1\}$ ,  $S_2 = \{0\}$ , and  $Z = \{0\}$ . Edges marked by a dot refer to  $\Delta$  according to definition 1 (equivalent to Brand's definition for this example). The state name is  $x\#d$ , where  $x \in D_1$ ,  $y \in D_2$ , and  $d \in \Delta$ .

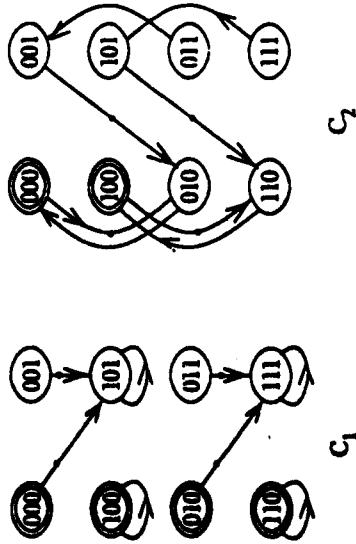


Figure B-2: Second counterexample to Lemma 2

Again, this example satisfies the conditions for Lemma 2. But Lemma 2 fails for the sequence  
 $000\ C_1\ 101\ C_2\ 110\ C_2\ 100$ .

This example does not satisfy requirement 1. The paths  $C_1^*$  from 000 and 010 do not reach stopping points.

## B.2. Comments on Proof of Lemma 2

The proof is by induction. The proof for "no references to  $\Delta$ " is trivially extended for "at most one reference to  $\Delta$ ". The inductive step is then for  $(n+1) \geq 2$  references. After assuming the first reference to  $\Delta$  is done by  $C_1$ , reorder the executions of  $C_1$  and  $C_2$ , moving as many  $C_1$  transitions as possible, as described in the proof. Then we have one of the three cases in Figure B-3.

Case I:

$$x = \langle a_0, b_0, d_0 \rangle C_1^* \langle a_1, b_0, d_0 \rangle \Sigma_1 \langle a_2, b_0, d_2 \rangle C_1^* \langle a_3, b_0, d_2 \rangle \\ \Sigma_1 \langle a_4, b_0, d_4 \rangle$$

Case II:

$$x = \langle a_0, b_0, d_0 \rangle C_1^* \langle a_1, b_0, d_0 \rangle \Sigma_1 \langle a_2, b_0, d_2 \rangle C_1^* \langle a_3, b_0, d_2 \rangle \\ \Sigma_2 \pm \langle a_4, b_0, d_4 \rangle \Sigma_1 \langle a_5, b_0, d_5 \rangle$$

Case III:

$$x = \langle a_0, b_0, d_0 \rangle C_1^* \langle a_1, b_0, d_0 \rangle \Sigma_1 \langle a_2, b_0, d_2 \rangle C_1^* \langle a_3, b_0, d_2 \rangle \\ \Sigma_2 \pm \langle a_4, b_0, d_4 \rangle = y$$

Figure B-3: Three cases in proof of Lemma 2

Figure B-3 shows the several possible types of static transition sequences.

where the state space is  $D_1 \times D_2 \times \Delta$ . The transitions that refer to  $\Delta$  are exactly those transitions for which the relation name is underlined. In the case of " $C_2 \perp$ ", more than one transition in the path may refer to  $\Delta$ . States  $x$  and  $y$  are the stopping points of interest in the proof.

For each case, the proof asserts that one of the states in the path

$$\langle a_2, b_0, d_j \rangle \circ C_1 \circ \langle a_3, b_0, d_j \rangle$$

is a stopping point. As shown by the counterexamples, this assertion fails without definition 1 and requirement 1. Now consider each case with the stronger requirements.

**Case I:** The presence of a stopping point in the path follows immediately from the requirement that two references to  $\Delta$  in  $P_1$  are separated by a stopping point.

**Case II:** Since program  $P_1$  has no don't-cares, some edge  $h$  in  $C_1$  is incident out of  $\langle a_2, b_0, d_j \rangle$ . Definition 1 implies that  $h$  refers to  $\Delta$ , since edge

$$\langle a_3, b_0, d_j \rangle \circ C_1 \circ \langle a_3, b_0, d_j \rangle$$

refers to  $\Delta$ . The addition of  $h$  to the path then makes this case like case I.

**Case III:** If  $\langle a_2, b_0, d_j \rangle$  is a stopping point, we are done. Assume  $\langle a_2, b_0, d_j \rangle$  is not a stopping point. Now  $\langle a_2, b_0, d_j \rangle$  is reachable from stopping point  $x$  by  $C_1^o$ . Then, by requirement 1, every path  $C_1^o$  from  $\langle a_2, b_0, d_j \rangle$  reaches a stopping point. Let  $z$  be the first stopping point on one such path. This path  $\langle a_2, b_0, d_j \rangle \circ C_1 \circ z$  does not refer to  $\Delta$ ; otherwise two successive references to  $\Delta$  by  $C_1$  would not be separated by a stopping point, a violation of the condition assumed in the lemma. Therefore,  $z^0 = d_j$ . Since  $z$  is a stopping point, we have  $d_j \in \Sigma$ .

Since  $x$  and  $y$  are stopping points, we have  $a_j \in S_1$  and  $b_0 \in S_2$ . Therefore  $\langle a_j, b_0, d_j \rangle$  is a stopping point.

Then, for each of the three cases, a stopping point is found on the path

$$\langle a_2, b_0, d_j \rangle \circ C_1 \circ \langle a_3, b_0, d_j \rangle$$

as asserted in the proof of the lemma. This completes the proof of Lemma 2.

In the Blackjack example, we will always have

- initial  $P_c$  equals  $TRUE$ ,
- initial  $Hil_3$ ,  $Broke_3$ , and  $Stand_3$  have constant values, and
- initial  $FF_3$ ,  $Score_3$ ,  $CardBuf_3$ ,  $Value_3$ , and  $YCnd_3$  equal symbolic variables  $f_0$ ,  $s_0$ ,  $c_0$ ,  $v_0$ , and  $y_0$ , respectively.

The only exception is when  $Lc_3$  equals  $Xf$ , in which case  $R_{3i}$  implies  $FF_3 = 0$ . This appendix presents the results of the symbolic execution of the blackjack descriptions from Chapter 6.

## Appendix C

### Symbolic Execution of Blackjack Example

This section presents the symbolic results, using case splitting, for the descriptions in Figures 6-1 and 6-2 as directed by the simulation relation  $R_{3i}$  defined by Figure 6-3. The results are obtained by the following procedure.

1. Select the initial location counters  $Lc_i$  for  $P_i$  (Figure 6-1) and  $Lc_j$  for  $P_j$  (Figure 6-2) corresponding to a pair of stopping points in  $R_{3i}$ . In the Blackjack descriptions,  $Lc_i$  necessarily equals  $top$ , while  $Lc_j$  might be  $XB$ , for example.
2. Initialize the path condition  $P_c$  and the variables of  $P_i$  as directed by  $R_{3i}$  for states with the given location counters. For example, in the Blackjack machine, if we selected  $Lc_3 = XB$ , then  $R_{3i}$  implies  $Hil_3 = 1$ ,  $Broke_3 = 0$ , and  $Stand_3 = 0$ . (For Boolean variables, we will use  $1 \equiv TRUE$  and  $0 \equiv FALSE$ .) In general, the initial path condition reflects requirements from  $R_{3i}$  that cannot be inferred from the initial values given to variables. There are no such requirements in the Blackjack example.
3. Execute  $P_i$  symbolically until  $Lc_i$  again equals  $top$ . In general, the execution of  $P_i$  could introduce more case splits or terminate with a different value of  $Lc_i$ , but that does not occur in the Blackjack example.
4. Initialize the variables of  $P_i$  as directed by  $R_{3i}$ . In the Blackjack example, we have

- $Hil_i$ ,  $Broke_i$ ,  $Stand_i$ ,  $FF_i$ ,  $Score_i$ ,  $CardBuf_i$ ,  $Value_i$ , and  $YCnd_i$  take the same initial values as their counterparts in  $P_s$ ,
- initial  $Q4_i$ ,  $Q5_i$ , and  $Q6_i$  have constant values, and
- all remaining variables initially are set to distinct symbolic variables.

The final  $P_c$  from the execution of  $P_i$  becomes the initial  $P_c$  for this execution of  $P_j$ .

5. Execute  $P_j$  symbolically until  $Lc_j$  again equals  $top$ . In general, the execution of  $P_j$  could introduce more case splits or terminate with a different value of  $Lc_j$ , but that does not occur in the Blackjack example.
6. This completes the parallel execution of  $P_s$  (step (3)) and  $P_i$  (step (5)) for a single path in each description. Record the results.

201

202

7. If step (5) introduced case splits (this never occurs for the Blackjack example), then repeat steps (5) and (6) for each separate path in  $P_j$ .

8. If step (1) introduced case splits, then repeat steps (3) through (7) for each separate path in  $P_g$ .

9. Repeat steps (1) through (8) for each possible pair of starting location counters.

The results of the above procedure are presented below for all cases. To save space, we use the following conventions:

$State_s$  and  $State_f$  are abbreviations for collections of variables. Initial values are specified only for  $State_s$  and  $State_f$ . An initial value for  $FF$  also appears only if that value is not  $s_0$ . The initial values of  $Score$ ,  $CardBuf$ ,  $Value$ , and  $YCard$  are always as given in step (2) above. The initial value of  $LC_f$  is always top. The initial and final values of all remaining variables in  $P_i$  are always irrelevant and so are never listed.

Results of symbolic executions of Figures 6-1 and 6-2 with case splitting

Figure 6-1 is description  $P_s$

( $\oplus$  “@” denotes concatenation “)

$State_s \equiv LC_s : Hit_s @ Broke_s @ Stand_s$

Figure 6-2 is description  $P_i$

$State_f \equiv Hit_f @ Broke_f @ Stand_f @ Q4_f @ Q5_f @ Q6_f$

$Pc: \text{ TRUE}$

Initial:	$State_s = XA:000$			
Final:	$State_s = XB:100$	$FF_s = 0$	$Score_s = 0$	$CardBuf_s = c_0$

Initial:	$State_f = 000000$			
Final:	$State_f = 100000$	$FF_f = 0$	$Score_f = 0$	$CardBuf_f = c_0$

$Pc: \text{ not } y_0$

Initial:	$State_s = XB:100$			
Final:	$State_s = XB:100$	$FF_s = s_0$	$Score_s = s_0$	$CardBuf_s = v_0$

Initial:	$State_f = 100000$			
Final:	$State_f = 100000$	$FF_f = s_0$	$Score_f = s_0$	$CardBuf_f = v_0$

$P_C: \gamma_0$ 

Initial:  $State_s = XC:100$   
Final:  $State_s = XC:000$     $FF_s = f_0$     $Score_s = s_0$     $CardBuf_s = v_0$

Initial:  $State_i = 100000$

Final:  $State_i = 000001$     $FF_i = f_0$     $Score_i = s_0$     $CardBuf_i = v_0$

 $P_C: \gamma_0$ 

Initial:  $State_s = XC:000$   
Final:  $State_s = XC:000$     $FF_s = f_0$     $Score_s = s_0$     $CardBuf_s = c_0$

Initial:  $State_i = 000001$

Final:  $State_i = 000001$     $FF_i = f_0$     $Score_i = s_0$     $CardBuf_i = c_0$

 $P_C: \text{and } \gamma_0$ 

Initial:  $State_s = XC:000$   
Final:  $State_s = D:000$     $FF_s = f_0$     $Score_s = s_0$     $CardBuf_s = c_0$

Initial:  $State_i = 000001$

Final:  $State_i = 000011$     $FF_i = f_0$     $Score_i = s_0$     $CardBuf_i = c_0$

$P_C: (c_0=1) \text{ and not } f_0$   
Initial:  $State_s = D:000$   
Final:  $State_s = E:000$     $FF_s = 0$     $Score_s = s_0 + 1$     $CardBuf_s = 1$   
values of  $FF_s$  and  $CardBuf_s$  unchanged; implied by  $P_C$

=====

 $P_C: \gamma_0$ 

Initial:  $State_i = 000001$   
Final:  $State_i = 000011$     $FF_i = 0$     $Score_i = s_0 + 1$     $CardBuf_i = 1$   
values of  $FF_i$  and  $CardBuf_i$  unchanged; implied by  $P_C$

Initial:  $State_i = 000001$

Final:  $State_i = 000011$     $FF_i = 0$     $Score_i = s_0 + 1$     $CardBuf_i = 1$   
values of  $FF_i$  and  $CardBuf_i$  unchanged; implied by  $P_C$

Initial:  $State_i = 000001$

Final:  $State_i = 000011$     $FF_i = 0$     $Score_i = s_0 + 1$     $CardBuf_i = 1$   
values of  $FF_i$  and  $CardBuf_i$  unchanged; implied by  $P_C$

Initial:  $State_i = 000001$

Final:  $State_i = 000011$     $FF_i = f_0$     $Score_i = s_0 + c_0$     $CardBuf_i = c_0$   
values of  $FF_i$  and  $CardBuf_i$  unchanged; implied by  $P_C$

Initial:  $State_i = 000001$

Final:  $State_i = 000010$     $FF_i = f_0$     $Score_i = s_0 + c_0$     $CardBuf_i = c_0$   
values of  $FF_i$  and  $CardBuf_i$  unchanged; implied by  $P_C$

Initial:  $State_i = 000001$

Final:  $State_i = 000011$     $FF_i = f_0$     $Score_i = s_0 + c_0$     $CardBuf_i = c_0$   
values of  $FF_i$  and  $CardBuf_i$  unchanged; implied by  $P_C$

$P_C: (c_0 \neq 1) \text{ or } f_0$   
Initial:  $State_s = D:000$   
Final:  $State_s = F:000$     $FF_s = f_0$     $Score_s = s_0 + c_0$     $CardBuf_s = c_0$   
values of  $FF_s$  and  $CardBuf_s$  unchanged; implied by  $P_C$

=====

Initial:  $State_i = 000001$

Final:  $State_i = 000011$     $FF_i = 1$     $Score_i = s_0$     $CardBuf_i = 10$   
must simplify ( $f_0$  or not  $f_0$ ) to get  $FF_i = 1$

=====

205      206

$Pc: s_0 \geq 17$	$State_s = F:000$	$FF_s = f_0$	$Score_s = s_0$	$CardBuf_s = c_0$
Initial:	$State_i = F:000$			
Final:	$State_f = XJ:100$	$FF_f = f_0$	$Score_f = s_0$	$CardBuf_f = -10$
Init:	$State_i = 000110$			
Final:	$State_f = 100000$	$FF_f = f_0$	$Score_f = s_0$	$CardBuf_f = c_0$
=====				

$Pc: s_0 \geq 22$ want optimize $Pc$ from $(s_0 \geq 17)$ and $(s_0 \geq 22)$	$State_s = F:000$	$FF_s = f_0$	$Score_s = s_0$	$CardBuf_s = c_0$
Initial:	$State_i = F:000$			
Final:	$State_f = XJ:010$	$FF_f = f_0$	$Score_f = s_0$	$CardBuf_f = -10$
Init:	$State_i = 000110$			
Final:	$State_f = 000100$	$FF_f = f_0$	$Score_f = s_0$	$CardBuf_f = c_0$
=====				

$Pc: (s_0 \geq 17) \text{ and } (s_0 \geq 22)$	$State_s = F:000$	$FF_s = f_0$	$Score_s = s_0$	$CardBuf_s = c_0$
Initial:	$State_i = F:000$			
Final:	$State_f = XJ:001$	$FF_f = f_0$	$Score_f = s_0$	$CardBuf_f = c_0$
Init:	$State_i = 000110$			
Final:	$State_f = 001000$	$FF_f = f_0$	$Score_f = s_0$	$CardBuf_f = c_0$
=====				

$Pc: f_0$	$State_s = H:000$	$FF_s = D:000$	$Score_s = s_0$	$CardBuf_s = -10$
Initial:	$State_i = 000100$			
Final:	$State_f = 000011$	$FF_f = 0$	$Score_f = s_0$	$CardBuf_f = -10$
Init:	$State_i = 000100$			
Final:	$State_f = 000000$	$FF_f = 0$	$Score_f = s_0$	$CardBuf_f = c_0$
=====				



Figure 6-2 contains no conditional branching, so the above procedure simplifies to the following:

1. Initialize the variables as given below for description  $P_i$ . We use many of the same symbolic variables used to initialize  $P_s$ , in agreement with  $R_{\text{sd}}$ .
2. Symbolically execute the loop once.

The results for the two descriptions follow.

Results of symbolic executions of Figures 6-1 and 6-2 with case merging

$State_i \oplus Hit_j \otimes Broke_j \otimes Stand_j \otimes Q1_j \otimes Q2_j$

Initial:

$\forall i \in \{XA, XB, XC, D, E, F, H, XJ, XK\}: LPc_j[i] = FALSE$

$Hit_j = b_0$	$Broke_j = b_0$	$Stand_j = l_0$
$FF_j = f_0$	$Score_j = s_0$	$CardBuf_j = c_0$
$YCard_j = y_0$	$Value_j = v_0$	

Final:

$\forall i \in \{XA, XB, XC, D, E, F, H, XJ, XK\}: LPc_j[i] = FALSE$

$LPc_j[XA] = ((p_0 = XJ) \text{ and } y_0) \text{ or } ((p_0 = XK) \text{ and } y_0)$   
 $LPc_j[XB] = (p_0 = XA) \text{ or } ((p_0 = XB) \text{ and not } y_0) \text{ or }$   
 $((p_0 = F) \text{ and } (s_0 < 17))$   
 $LPc_j[XC] = ((p_0 = XB) \text{ and } y_0) \text{ or } ((p_0 = XC) \text{ and } y_0)$   
 $LPc_j[D] = ((p_0 = XC) \text{ and not } y_0) \text{ or } ((p_0 = E) \text{ or } ((p_0 = H) \text{ and } f_0))$   
 $LPc_j[E] = (p_0 = D) \text{ and } (c_0 = 1) \text{ and not } f_0$   
 $LPc_j[F] = (p_0 = D) \text{ and } ((c_0 \neq 1) \text{ or } f_0)$   
 $LPc_j[H] = (p_0 = F) \text{ and } (s_0 \geq 17) \text{ and } (s_0 \geq 22)$   
 $LPc_j[XJ] = ((p_0 = H) \text{ and not } f_0) \text{ or } ((p_0 = XJ) \text{ and not } y_0)$   
 $LPc_j[XK] = ((p_0 = F) \text{ and } (s_0 \geq 17) \text{ and } (s_0 \geq 22)) \text{ or }$   
 $((p_0 = XK) \text{ and not } y_0)$

$Hit_j = (p_0 = XA) \text{ or } ((p_0 = XB) \text{ and not } y_0) \text{ or } ((p_0 = F) \text{ and } (s_0 < 17)) \text{ or }$   
 $(h_0 \text{ and not } (y_0 \text{ and } p_0 \in \{XB, XC, XJ, XK\}))$

$Broke_j = (((p_0 = H) \text{ and not } f_0) \text{ or } ((p_0 = XJ) \text{ and not } y_0)) \text{ or }$   
 $(b_0 \text{ and not } (y_0 \text{ and } p_0 \in \{XJ, XK\}))$

$Stand_j = (((p_0 = F) \text{ and } (s_0 \geq 17) \text{ and } (s_0 \geq 22)) \text{ or } ((p_0 = XK) \text{ and not } y_0)) \text{ or }$   
 $(l_0 \text{ and not } (y_0 \text{ and } p_0 \in \{XJ, XK\}))$

211

212

$$FF_3 = (p_0 = E) \text{ or } (f_0 \text{ and } p_0 \in \{XB, XC, D, F, XJ, XK\})$$

$$Score_3 = \begin{cases} 0 & \text{if } p_0 = XA \text{ then 0 else} \\ & \text{if } p_0 = D \text{ then } s_0 + c_0 \text{ else} \\ & \quad \text{if } p_0 \in \{XA, D\} \text{ then } s_0 \\ & \quad \text{CardBuf}_3 = \begin{cases} 0 & \text{if } p_0 = XB \text{ then } v_0 \text{ else} \\ & \quad \text{if } p_0 = E \text{ then 10 else} \\ & \quad \text{if } p_0 = H \text{ then -10 else} \end{cases} \end{cases}$$

$$FF_i = \begin{cases} 0 & \text{if } p_0 \in \{XB, E, H\} \text{ then } c_0 \\ & \dots \end{cases}$$

Initial:

$$\begin{aligned} State_i &= h_0 b_0 f_0 u_0 w_0 x_0 & Score_i &= s_0 & CardBuf_i &= c_0 \\ FF_i &= f_0 & Value_i &= v_0 \\ YCard_i &= y_0 & & & & \end{aligned}$$

Final:

$$\begin{aligned} HH_i &= ((not(b_0 \text{ or } f_0 \text{ or } u_0 \text{ or } x_0)) \text{ or} \\ &\quad (v_0 \text{ and not } x_0 \text{ and } (z_0 < 17)) \text{ and not } h_0) \text{ or} \\ &\quad not(y_0 \text{ and } h_0) \\ Broke_i &= (u_0 \text{ and not } w_0 \text{ and not } f_0 \text{ and not } b_0) \text{ or} \\ &\quad not(y_0 \text{ and } b_0) \\ Stand_i &= (w_0 \text{ and not } x_0 \text{ and } (17 \leq z_0 < 22) \text{ and not } f_0) \text{ or} \\ &\quad not(y_0 \text{ and } f_0) \end{aligned}$$

$$\begin{aligned} Q4_i &= w_0 \text{ and not } u_0 \text{ or} \\ &\quad not(not(w_0 \text{ or } x_0 \text{ or } (z_0 < 22)) \text{ and } u_0) \\ QS_i &= (x_0 \text{ and not } y_0 \text{ or } u_0 \text{ and } f_0) \text{ and not } w_0 \text{ or} \\ &\quad not(u_0 \text{ and not } x_0) \text{ and } w_0 \\ Q6_i &= (h_0 \text{ and } y_0 \text{ or } u_0 \text{ and not } w_0 \text{ and } f_0) \text{ and not } x_0 \text{ or} \\ &\quad not(((c_0 \neq 1) \text{ or } f_0) \text{ and not } u_0 \text{ and } w_0) \text{ and } x_0 \\ FF_i &= (u_0 \text{ and } x_0 \text{ and not } f_0) \text{ or} \\ &\quad (h_0 \text{ or } t_0 \text{ or } w_0 \text{ or } x_0) \text{ and } f_0 \\ Score_i &= \begin{cases} 0 & \text{if not}(h_0 \text{ or } b_0 \text{ or } f_0 \text{ or } u_0 \text{ or } x_0) \text{ then 0 else} \\ & \quad \text{if } (not(u_0 \text{ and } w_0) \text{ then } s_0 + c_0 \text{ else } s_0) \end{cases} \\ CardBuf_i &= \begin{cases} 0 & \text{if } h_0 \text{ then } v_0 \text{ else} \\ & \quad \text{if } (u_0 \text{ and } x_0) \text{ then 10 else} \\ & \quad \text{if } (v_0 \text{ and not } w_0) \text{ then -10 else } c_0 \end{cases} \\ & \dots \end{aligned}$$

## References

- [AH74] A.V. Aho, J.E. Hopcroft, and J.D. Ullman. *The Design and Analysis of Computer Algorithms*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1974.
- [AU77] A.V. Aho and J.D. Ullman. *Principles of Compiler Design*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1977.
- [BB79] M.R. Barbacci, G.E. Barnes, R.G. Cattell, and D.P. Siciworek. *The ISPS Computer Description Language*. Report CMU-CS-79-137. Carnegie-Mellon University, Department of Computer Science, August 1979.
- [BF76] M.A. Breuer and A.D. Friedman. *Diagnosis & Reliable Design of Digital Systems*. Computer Science Press, Potomac, Maryland, 1976.
- [Bo82] G.V. Bochmann. Hardware Specification with Temporal Logic: An Example. *IEEE Trans. Computers* C-31(3):223-231, March 1982.
- [Br78] D. Brand. *Algebraic Simulation between Parallel Programs*. Research Report RC 7206, IBM T.J. Watson Research Center, Yorktown Heights, N.Y., June 1978. 39 pages.
- [CE77] W.C. Carter, H.A. Ellozy, W.H. Joyner Jr., and G.B. Leeman Jr. *Techniques for Microprogram Validation*. Research Report RC 6361, IBM T.J. Watson Research Center, Yorktown Heights, N.Y., January 1977. 19 pages.
- [CH79] T.E. Cheatham, Jr., G.H. Holloway, and J.A. Townley. Symbolic Evaluation and the Analysis of Programs. *IEEE Trans. Software Engineering* SE-5(4):402-417, July 1979.
- [CJ79] W.C. Carter, W.H. Joyner Jr., and D. Brand. Microprogram Verification Considered Necessary. In *Proc. National Computer Conference*, pages 657-664. AFIPS, 1978.
- [CJ79] W.C. Carter, W.H. Joyner Jr., and D. Brand. Symbolic Simulation for Correct Machine Design. In *Proc. 16th Design Automation Conference*, pages 289-296. IEEE and ACM, San Diego, 1979.
- [Co81] W.E. Cory. Symbolic Simulation for Functional Verification with ADLIB and SDL. In *Proc. 18th Design Automation Conference*, pages 82-89. IEEE and ACM, Nashville, June 1981.
- [Cv82] W.E. Cory and W.M. vanCleempout. Developments in Verification of Design Correctness-A Tutorial. In Rex Ricc (editor), *Tutorial-VLSI Support Technologies: Computer-Aided Design, Testing and Packaging*, pages 169-177. IEEE Computer Society Press, Los Alamitos, CA, 1982. Reprinted from *Proc. 17th Design Automation Conference*, pages 156-164. IEEE and ACM, Minneapolis, June, 1980.
- [Da79] J.A. Daringer. The Application of Program Verification Techniques to Hardware Verification. In *Proc. 16th Design Automation Conference*, pages 375-381. IEEE and ACM, San Diego, June 1979.
- [DD75] D.L. Dietmeyer and J.R. Dulley. Register Transfer Languages and Their Translation. In M.A. Breuer (editor), *Digital System Design Automation: Languages, Simulation & Data Base*, chapter 2. Computer Science Press, Woodland Hills, CA, 1975. pages 117-218.
- [EG76] C.J. Evangelisti, G. Goertzel, and H. Ofek. *LCD-Language for Computer Design (Revised)*. Research Report RC 6244, IBM T.J. Watson Research Center, Yorktown Heights, N.Y., October 1976. 32 pages.
- [EG77] C.J. Evangelisti, G. Goertzel, and H. Ofek. Designing with LCD-Language for Computer Design. In *Proc. Fourteenth Design Automation Conference*, pages 369-376. IEEE and ACM, New Orleans, 1977.
- [Fl67] R.W. Floyd. Assigning Meanings to Programs. In *Proc. Symposia in Applied Mathematics*, pages 19-32. American Mathematical Society, 1967. Report RC 5225, IBM T.J. Watson Research Center, Yorktown Heights, N.Y., January 1975.
- [Hl79] D.D. Hill. *EFFIGY Reference Manual*. Research Report RC 5225, IBM T.J. Watson Research Center, Yorktown Heights, N.Y., January 1979.
- [Hl79] D.D. Hill. A Modular, Strongly-Typed Computer Design Language. In *Proc. Fourth Int'l. Symposium on Computer Hardware Description Languages*, pages 75-81. IEEE and ACM, Palo Alto, 1979.
- [Hi80] D.D. Hill. *Language and Environment for Multi-Level Simulation*. PhD thesis, Stanford University, Dept. of Electrical Engineering, March 1980. also Technical Report No. 185, 170 pages.
- [Cr80] S.D. Crocker, L. Marcus, and D. van Mierop. The ISI Microcode Verification System. In *Proc. IFIP Working Conference on Firmware, Microprogramming and Restructurable Hardware*, pages 89-103. International Federation for Information Processing, Linz, Austria, April 1980.

- [HK76] S.L. Hawler and J.C. King. An Introduction to Proving the Correctness of Programs. *ACM Computing Surveys* 8(3):331-353, September 1976.
- [Ho69] C.A.R. Hoare. An Axiomatic Basis for Computer Programming. *Communications of the ACM* 12(10):576-583, October 1969.
- [HR74] H. Halliwell and J.P. Roth. SCD: A System for Computer Design. *IBM Technical Disclosure Bulletin* 17(5):1517-1519, October 1974.
- [Hv79] D.D. Hill and W.M. vanCleempt. SABLE: A Tool for Generating Structured, Multi-level Simulations. In *Proc. Sixteenth Design Automation Conference*, pages 272-279. IEEE and ACM, San Diego, June 1979.
- [Jc76] W.H. Joynet, Jr., W.C. Carter, G.B. Leeman, Jr. Automated Proofs of Microprogram Correctness. In *Proc. Ninth Annual Workshop on Microprogramming*, pages 51-55. IEEE and ACM, New Orleans, September 1976.
- [JW79] K. Jensen and N. Wirth. *Pascal User Manual and Report, Second Edition*. Springer-Verlag, New York, 1979.
- [Ki75] J.C. King. A New Approach to Program Testing. *ACM SIGPLAN Notices* 10(6):228-233, June 1975. This issue is *Proc. 1975 International Conference on Reliable Software, Los Angeles*, April 1975.
- [Ki76a] J.C. King. Symbolic Execution and Program Testing. *Communications of the ACM* 19(7):385-394, July 1976.
- [Ki76b] J.C. King. On Generating Verification Conditions for Correctness Proofs. In *4. Fachtagung der GI Programmiersprachen*, pages 253-267. Gesellschaft Fur Informatik, Erlangen, Germany, March 1976. Published by Springer-Verlag: New York; paper is in English. Also available as Research Report RC 5808, January 1976, from IBM T.J. Watson Research Center, Yorktown Heights, N.Y.
- [KS79] N. Kawato, T. Saito, F. Maruyama, and T. Uehara. Design and Verification of Large-Scale Computers by Using DDI. In *Proc. 16th Design Automation Conference*, pages 360-366. IEEE and ACM, San Diego, 1979.
- [LG79] D.C. Luckham, S.M. German, F.W. v. Henke, R.A. Karp, P.W. Milne, D.C. Oppen, W. Potak, and W.L. Scherlis. *Stanford Pascal Verifier User Manual*. Report No. STAN-CS-79-731, Stanford University, Computer Science Department, March 1979. 115 pages.
- [LJ77] G.B. Leeman Jr., W.H. Joynet Jr., and W.C. Carter. *An Automated Proof of Microprogram Correctness*. Research Report RC 6587, IBM T.J. Watson Research Center, Yorktown Heights, N.Y., June 1977. 33 pages.
- [LL79] S. Leinwand and T. Lamdan. Design Verification Based on Functional Abstraction. In *Proc. 16th Design Automation Conference*, pages 353-359. IEEE and ACM, San Diego, 1979.
- [LS79] D.C. Luckham and N. Suzuki. Verification of Array, Record, and Pointer Operations in Pascal. *ACM Trans. Programming Languages and Systems* 1(2):226-244, October 1979.
- [Ma66] P. Martin-Löf. The Definition of Random Sequences. *Information and Control* 9(6):602-619, December 1966.
- [Mc62] J. McCarthy. Towards a Mathematical Science of Computation. In *Proc. IFIP Congress 62*, pages 21-28. IFIP, Munich, August 1962.
- [Mi71] R. Milner. An Algebraic Definition of Simulation between Programs. In *Proc. Second International Joint Conference on Artificial Intelligence*, pages 481-489. British Computer Society, London, September 1971.
- [MP81] Z. Manna and A. Pnueli. *Verification of Concurrent Programs, Part I: The Temporal Framework*. Report No. STAN-CS-81-836, Stanford University, Computer Science Department, June 1981. 62 pages.
- [Ol78] H. Olfet, J.D. Lesser, C.J. Evangelisti, and G. Goertzel. *Structured Design Verification of Sequential Machines*. Research Report RC 7037, IBM T.J. Watson Research Center, Yorktown Heights, N.Y., March 1978. 20 pages.
- [Pa76] D.A. Patterson. STRUM: Structured Microprogramming System for Correct Firmware. *IEEE Trans. Computers Special Issue on Microprogramming*, October 1976.
- [PS82] V. Pitchumani and E.P. Stabler. A Formal Method for Computer Design Verification. In *Proc. 19th Design Automation Conference*, pages 809-814. IEEE and ACM, Las Vegas, June 1982.
- [Ro66] J.P. Roth. Diagnosis of Automata Failures: A Calculus and a Method. *IBM Journal of Research and Development* 10:278-291, July 1966.
- [Ro73] J.P. Roth. Verify: An Algorithm to Verify a Computer Design. *IBM Technical Disclosure Bulletin* 15(8):2646-2648, January 1973.

- [Ro75] J.P. Roth. *Generation and Verification of Hardware Designs at High Level*. Research Report RC 5779, IBM T.J. Watson Research Center, Yorktown Heights, N.Y., November 1975. 16 pages.
- [Rott] J.P. Roth. Hardware Verification. *IEEE Trans. Computers* C-26(12):1292-1294, December 1977.
- [Ul65] E.G. Ulrich. Time-Sequenced Logical Simulation Based on Circuit Delay and Selective Tracing of Active Network Paths. In *Proc. 20th National Conference*, pages 437-448. ACM, Cleveland, August 1965.
- [Uh81] S.H. Unger. Double-Edge-Triggered Flip-Flops. *IEEE Trans. Computers* C-30(6):447-451, June 1981.
- [va77] W.M. vanCleempunt. *A Structural Design Language for Computer Aided Design of Digital Systems*. Technical Report 136, Stanford University, Computer Systems Laboratory, April 1977. 27 pages.
- [va79] D. van-Mierop. *An Experiment in Automatic Verification of Microcode*. Microver note #31, University of Southern California, Information Sciences Institute, June 1979. 21 pages.
- [vM78] D. van-Mierop, L. Marcus, and S.D. Crocker. Verification of the FTSC Microprogram. In *Proc. 11th Annual Microprogramming Workshop*, pages 118. IEEE and ACM, Pacific Grove, CA, November 1978.

