

ANALYSIS OF SEQUENCING IN COMPUTER
PROGRAMS AND SYSTEMS

by

William Preston Alexander, III

August 1974

TR-35

ANALYSIS OF SEQUENCING IN COMPUTER
PROGRAMS AND SYSTEMS*

by

William Preston Alexander, III

*This paper is supported in part by the National Science Foundation grants GJ36424 and GJ1084.

ANALYSIS OF SEQUENCING IN COMPUTER PROGRAMS AND SYSTEMS

by

WILLIAM PRESTON ALEXANDER, III, B.A., M.A.

TR-35

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

August 1974

ACKNOWLEDGMENTS

I would like to thank the members of my committee, each of whom helped in some unique way, and especially Dr. J. C. Browne for his encouragement and Dr. J. H. Howard for his time and effort.

I also acknowledge my debt to my parents, who are largely responsible.

This work was supported in part by NSF grants GJ36424 and GJ1084.

May 1974

ABSTRACT

Methods are developed for determining the orders in which a computer program performs its operations and for verifying from its source code whether a program adheres to any given sequencing rule. These methods are based on analysis of the program's state graph, and techniques are presented for compressing the graph to eliminate unnecessary detail and for expanding it to add necessary distinctions. A program is described which automates most of these methods and which has been used to analyze parts of a real operating system. The methods are extended to apply to systems of parallel processes. Systems are modeled as state graphs and analyzed for the order in which they perform operations as well as for blockage properties such as deadlock.

TABLE OF CONTENTS

CHAPTER	PAGE
I. INTRODUCTION	1
1.1 The Problem	1
1.2 Summary of the Research	2
1.3 The Context	3
1.4 Summary of Chapters	5
II. ANALYZING INDIVIDUAL PROCESSES	6
2.1 Introduction	6
2.2 State Graphs and Traces	7
2.3 Prototype Graphs	11
2.4 Comparison of Graphs	13
2.5 State Graph Construction	20
2.6 Splitting	23
2.7 Folding	28
2.8 Summary	32
III. PRACTICAL EXPERIENCE	35
3.1 Introduction	35
3.2 Building the State Graph	36
3.3 Splitting	38
3.4 Folding	39
3.5 Actual Uses	41
3.6 Conclusions	44
IV. ANALYZING SYSTEMS OF PROCESSES	47
4.1 Introduction	47

CHAPTER	PAGE
4.2 System State Graphs	48
4.3 Analyzing Sequencing Properties of Systems	54
4.4 Deadlock	60
4.5 Analyzing a System for Deadlock	64
4.6 Permanent Blocking	66
4.7 An Example	67
V. SUMMARY AND CONCLUSIONS	79
5.1 Summary of Previous Chapters	79
5.2 Suggestions for Further Research	80
5.3 Graphs As Models	80
5.4 Conclusion	83
REFERENCES	84

LIST OF TABLES

TABLE	PAGE
2.1 Partial Execution of the Comparison Algorithm	19
3.1 Some Executions of Program TRACE	43
4.1 Example of Use of the System Comparison Algorithm	59
4.2 Analysis of the System of Figure 4.6	72

LIST OF FIGURES

FIGURE	PAGE
2.1 Critical Section Prototype Graph	12
2.2 Graphs of Two Erroneous Programs	17
2.3 Example Graphs	18
2.4 Segment of Source Code	24
2.5 Segment of Initial State Graph (Incorrect)	22
2.6 Segment of Initial State Graph (Corrected)	24
2.7 Graph Segment During Splitting	26
2.8 Graph Segment After Splitting	26
2.9 Illustration of the "Only Successor" Folding Rule	31
2.10 Graph Segment After Folding	31
2.11 Prototype for One-use-only	32
4.1 A System of Parallel Processes	48
4.2 Prototype Graph for Mutual Exclusion of Two Processes with Critical Section	55
4.3 Two Parallel Processes with Interlock	57
4.4 A Restless System	63
4.5 Inaccurate Representation of Two Processes	69
4.6 Accurate Representation of Two Processes	70
5.1 Use of the Graph Conversion Algorithm	82

CHAPTER I

INTRODUCTION

1.1 The Problem

Many important properties of computer programs can be expressed in terms of the order in which they perform actions. The order in which a program produces its external effects can be important in many real time applications. Most methods for coordinating multiple parallel asynchronous processes require that each process observe some protocol dictating the order in which transactions between the processes may occur. Other properties involve the internal workings of the memory function, such as the usually desirable property that a program assign a value to a cell before using its contents (the initialization problem). For large and complex programs, traditional testing and debugging methods cannot assure that a program has such properties, and none of the formal proof procedures for programs are designed to prove ordering properties.

The first objective of this study is the development and automation of a method for stating and proving assertions about the order in which a program performs its operations. A second goal is to develop a method for stating and proving similar assertions about systems of processes running in parallel. The methods must verify rather than simulate or test; the source code of the object programs must be analyzed, not executed.

1.2 Summary of the Research

In this study a comprehensive method is developed for making and proving statements about the order in which programs perform actions. Individual programs are modeled empirically as directed graphs called state graphs which are generated from the source code, and paths through this graph indicate possible sequences of actions by the program. The statements to be proved are made in terms of prototype state graphs, small compact graphs specifying allowable sequences of actions. Careful attention is given to the problems of dealing with state graphs of large programs in limited space and time. Techniques are presented for expanding and compressing state graphs to reflect just the amount of detail needed and no more. Based on formal language theory, an algorithm is presented which compares the sequences represented by one graph with those of another.

The method presented for proving sequencing properties of programs is practical and can be almost completely automated. A program is described which implements most of the techniques and which has been used to analyze parts of a real operating system.

The method developed for analyzing individual processes is extended to apply to systems. Systems are modeled by their state graph, and this graph can then be analyzed to determine the order in which the system performs actions as well as for system blockage properties such as deadlock.

The program proof method presented here does not improve on or replace any existing program correctness technique; rather its importance is that it provides a whole new class of statements about programs which can be proved. The method for analyzing systems cannot be applied directly to real large-scale operating systems, but it does provide a unified model for expressing a wide

class of operating system properties and should prove useful in analyzing simple systems and in evaluating protocols for coordinating parallel processes.

1.3 The Context

As more and more people learn how to program computers, the seller's market once enjoyed by programmers is turning into a buyer's market in which customers can demand correct, dependably correct, and soon even provably correct programs. For this reason, and perhaps because the challenge of applying mathematical rigor to proofs about objects with effects in the real world is irresistible, there has arisen beside the traditional interest in debugging and testing programs a growing interest in proving statements about them a priori. A survey of this area was recently provided by London [20]. An exhaustive bibliography by Hetzel [12] covers program testing and debugging as well as program proving.

The object of the proofs referred to is program semantics, not syntax. Throughout this study, "program" means "syntactically correct computer program." The semantics of a program is the set of all the effects of executing the program. In the broadest sense, "correctness" means that the semantics of a program match the desire of the user or the intent of the writer. Obviously then, we do not prove programs correct; in the current state of the art we prove that a program has some property or, more specifically, that its execution has certain effects or that its effects have certain properties.

Perhaps because this area in computer science has taken models as well as personnel from mathematics, programs have most often been viewed as functions whose range and domain are the memory of the (virtual) machine upon

which they execute. One can take as argument the initial state, or the vector of initial values in all the memory cells, the value as the final state, or vector of final values, and conceive of a program as the function which maps one set of vectors to the other. There are many techniques for proving statements, usually expressed in first order predicate calculus, about the program-as-function. These include computational induction, structural induction, recursion induction, fixed-point induction, and inductive assertions. Manna, Ness, and Vuillemin [21] discuss all of these methods and prove some equivalences among them. Many interesting properties of programs can be expressed using these techniques, and London [19] has even been able to show that a simple but real compiler produces correct code in the sense that the compiler-produced machine code computes the same function as the source code.

However, a memory function is not a completely satisfactory model for a computer program. A program operates over a length of time, while a function is usually thought of as acting instantaneously. If a function produces as its value a vector, the order in which the values in the vector are produced does not usually matter. But because a program operates in time, the order in which it performs its actions may be important.

This is not to say that the inductive proof methods do not take any account of the sequential nature of programs. They all can allow composition of functions and provide computation rules for order of evaluation. Even more explicitly, the inductive assertion method involves explicitly marking paths through the program and making assertions which apply before, and others after, certain path segments have been executed. Perhaps the most explicit recognition of time is Good's [9] addition of a variable, t , representing the time at which a given point in the program is reached, to the domain of the assertions. The problem with using the augmented inductive assertion method to

prove statements about the order in which programs perform operations is that unless one already knows the order it is very difficult to assign values to t at each point in the program or to make statements about it, particularly in any standard predicate calculus. This is because when a program has loops one must assign a different t for each time the program reaches a given point in the loop. Otherwise, if two different points on a loop were simply assigned an unsubscripted t and t' , respectively, it might be the case that both $t < t'$ and $t > t'$ were true.

In fact this extension to the inductive assertion method was proposed to allow one to prove additional statements about the memory function. It can also be said of the path notation in the inductive assertion method and of the composition of functions that they all allow the prover to more accurately take into account the sequential nature of programs in order to make or prove statements about the program as memory function, but they are not well suited to stating or proving sequencing properties of programs. For this important class of properties, a new method is needed.

1.4 Summary of Chapters

The structure of the rest of this study is as follows: Chapter II introduces state graphs and develops the method for proving a wide class of sequencing assertions about individual computer programs. Chapter III describes an actual program written to implement the method of Chapter II and details the uses to which this program has been put. In Chapter IV the method of Chapter II is extended to apply to systems of parallel processes. The final chapter contains some general observations and suggestions for further work.

CHAPTER II

ANALYZING INDIVIDUAL PROCESSES

2.1 Introduction

In this chapter a method is developed for analyzing the orders in which individual computer programs perform their actions. The method consists of three steps: First the program is modeled by its state graph. State graphs are defined in the next section, and the process of building a program's initial state graph from its source code is discussed in Section 2.5. In the second step of the method the state graph is manipulated to eliminate unnecessary detail and to make additional distinctions as needed. The two main graph manipulatory techniques, folding and splitting, are explained in Sections 2.6 and 2.7. The statements to be proved about programs are embodied in small graphs called prototypes which specify allowable sequences. The third and final step in the method presented in this chapter is to compare the program's state graph with the prototype graph.

This chapter is organized so as to present as quickly as possible the broad outline of the method, leaving until later certain details. Thus the next three sections explain what a state graph is, what a prototype is, and how to compare them. The later sections explain how to get the state graph and how to assure that it compactly and accurately reflects the program's sequencing properties.

2.2 State Graphs and Traces

In the course of its execution, a program carries the machine which is interpreting it through many distinct states. By a state is meant a combination of values of all the memory cells and registers, including the program counter, used by the program.

A state graph of a program is a finite directed graph with nodes corresponding to states of the program/machine and arcs labeled with the program actions which carry the machine from one state to the next. A state graph is somewhat like a very detailed flow-chart of a program except that the roles of nodes and arcs are reversed; arcs represent actions while nodes represent states of the machine between actions.

A node may have several incident and exiting arcs, possibly with the same labels, and two nodes may have more than one arc connecting them in either direction. We use the notation $s \rightarrow s'$ to mean that there is an arc from node s to node s' and $s \xrightarrow{x} s'$ to indicate that this arc is labeled "x". Node s' is a successor of s , and s is a predecessor of s' . An arc labeled x may be called an "x-arc". An arc from a node to itself is a unit arc.

$s_0 \Rightarrow s_n$ means that there is a sequence of nodes, s_0, s_1, \dots, s_n , $n \geq 0$, such that $s_{i-1} \rightarrow s_i$ for $1 \leq i \leq n$. In this case we say that there is a path from s_0 to s_n and that s_n is reachable from s_0 .

Notice that this model of a program leaves many aspects of its behavior uninterpreted. $s \xrightarrow{x} s'$ simply says that an action labeled "x" has occurred, but says nothing about what x is or what effects it may have. The use of directed graphs other than state graphs as models of various aspects of program structure or behavior is fairly common; a survey of some of these models is provided by Baer [2].

Typically we shall not be concerned with trying to prove an assertion about all the operations which the program may perform but about a relatively small subset of them. For example, in the case of trying to determine whether or not a program observes mutual exclusion on a given resource, the assertion to be verified will only involve the four operations "reserve the resource," "use it," "release it," and "halt." The order in which the program performs any other operations is irrelevant. Thus for the purposes of this analysis most or all other operations are completely equivalent and can be denoted by the same arc label, ϵ , which will be used to mean "uninteresting operation." Further, whole sequences of such uninteresting operations can be lumped together as a single operation represented by a single ϵ -arc to the state resulting from the execution of the last one of the sequence.

A typical feature of program execution is branching on the value of some logical expression involving memory cells or registers. If the complete state of the machine is known at all times then the results of such tests can be determined, and there will be no uncertainty about the actual paths through the state graph. However, the retention of complete information is impractical, and thus it is sometimes convenient to mark the arcs leaving a given node with the logical condition under which that transition may be made. These conditions on the arcs will be in addition to the labels on the arcs identifying actions.

Since the representation of programs by their state graphs is the basis of all analysis in this study, the terms "node" and "state" will be used interchangeably when no confusion can result.

It is convenient to assume that there is a single unique starting state of the program from which all states in its graph are reachable. Real programs, however, may have several starting points, and even if they have only one it may be the case that some of their memory cells or registers are initialized

externally so that there are many possible starting states. However, since we are only interested in the order in which the program performs certain interesting actions, it can do no harm to the analysis if we simply create an extra node with no incoming arcs and an ϵ -arc from it to each of the actual possible starting states and call this new node the unique starting state of the program. Throughout the rest of this study it will be assumed that this has been done to all state graphs, and this node will be called s_0 . (The problem which arises when a program executes correctly with one set of initial values for some variables but may perform actions in an incorrect order for other initial values is discussed in Section 2.7.)

It is also convenient to assume that each state graph has identifiable final states representing the states of the machine after it has finished executing all of the instructions of interest. Certainly the states immediately following "halt" actions are always final states, and any states with no successors can always be automatically identified as final states. In practice it may be the case that certain programs never terminate and in fact do not even have a "halt" operation on any arc of their graph. An example might be the program which accepts input for a system. For these programs, the concept of a resting state [15] is useful. Intuitively, a resting state is the state of a process when it has "just completed" whatever it does and is about to "start over." More formally, a resting state is a state from which all other states in the graph can be reached. Notice that the assumption that all states are reachable from s_0 makes it a resting state. This formal definition of resting states may include more states than the intuitive one; in fact, in a strongly connected graph all the states would be resting states. Presumably only one or a few of these would be designated as final states. In any case, it is assumed throughout the rest of this study that

all state graphs have a known set of final states, whether halt states or resting states or some of each, and that these states have been explicitly so designated.

We are now ready to observe the way in which a state graph models a program's sequencing behavior. Each path through the state graph from the starting state to a final state represents a possible execution sequence of the program, and the sequence of arc labels along this path denotes a sequence of actions which the program may perform. The set of all such sequences defines all the orders in which the program may perform its actions. Such a sequence of arc labels is called a trace. The trace, α , associated with each path is the sequence of arc labels along that path. Thus the notation $s_0 \xRightarrow{\alpha} s_n$ indicates that there is a sequence of nodes $s_0, s_1, \dots, s_n, n \geq 0$, such that $s_{i-1} \xrightarrow{x_i} s_i$ for $1 \leq i \leq n$, and that $\alpha = x_1 \dots x_n$. The length of a path is defined to be the number of arcs, not necessarily distinct, traversed along it, and the trace of a path of length 0 will be denoted by λ . In a string of symbols such as a trace, x^* denotes an arbitrary number of occurrences of the symbol x , including none, while x^+ means at least one x .

Since state graphs as defined here are finite, they may not be able to accurately model the sequencing properties of programs with loops of arbitrarily large index or with recursion to arbitrary depth. This is a limitation on all of the methods for analyzing both programs and systems presented in this study.

As a final note, observe that the complete state graph of a program actually contains all possible information about its behavior and semantics. Any assertion about the effects of a program could in theory be decided by an exhaustive examination of its state graph if such an examination were practical. Its effects as a function on memory, for example, could be discovered

by examining the values of memory cells at each of the final states or at the states immediately preceding output actions. Such a proof procedure would, of course, not only be inferior to the inductive methods mentioned in the previous chapter but would be impractical for all but the most trivial programs due to the great size of most state graphs. However, state graphs do provide an adequate and practical basis for analyzing sequencing behavior in many programs.

2.3 Prototype Graphs

In any program verification application, one must decide how to form assertions about programs and how to characterize the kinds of properties which can be analyzed. State graphs, along with the idea of traces, provide a natural way of expressing sequencing properties. If the trace xyz exists in the state graph of a program, then the program may perform an x followed by a y and then a z. Thus any sequencing property of a program may be expressed in terms of the existence and/or non-existence of traces through its state graph. As a simple example, suppose we wished to verify that a certain program never performed x, y, and z consecutively in that order. It would be a relatively easy matter to devise an algorithm which searched through a state graph looking for a given trace. However, such an algorithm would be useful only if there were one or a few specific traces whose existence one wanted to deny or confirm. Typically there are an almost unlimited number of ways in which a program can perform incorrectly or correctly. A much more general way of characterizing sequencing properties is to express them as an extremely compact state graph serving as a prototype of correct action. This can be illustrated by the following example.

Suppose we have a resource, such as a table in memory, which is to be used by several asynchronous parallel processes, and that we wish to insure that no two processes ever access the table simultaneously by using Dijkstra's semaphore protocol governing entry into critical sections [6]. This protocol is defined in terms of three primitive operations:

- P - Set an interlock which prevents any other process from performing the P operation until this process has performed a V operation.
- V - Release the interlock.
- A - Access the shared table protected by the P and V operations.

A program is defined to be correct in its usage of these operations if and only if all of the traces in its state graph from the starting state to a final state are also present from the starting state back to that state in the prototype graph of Figure 2.1.




Figure 2.1. Critical Section Prototype Graph

The node on the left is both the starting and final state of the graph and represents the noncritical sections of the program. The node on the right represents the critical section of the program when it is accessing the shared table and when other programs must be prevented from doing so. This prototype graph defines all of the desired correctness criteria for using the three primitive operations: A program may not perform an A without first having done a P because there are no traces from the starting state of the prototype

beginning with A; having done a P, a program may not halt without doing a V because the V-arc begins the only path back to the final state, etc.

Notice that each node of the prototype has a unit ϵ -arc indicating that other, uninteresting actions may be freely interspersed between the execution of the P, A, and V operations. All prototype graphs will have this feature by definition, since if some action were forbidden between, say, a P and a V, it would be interesting and should have a unique arc label other than ϵ to represent it.

All program verification techniques ultimately depend on a user-supplied specification of what constitutes "correctness." The prototype graph plays that role here; the methods of this chapter will aid the user in verifying that a given program adheres to the sequencing rules defined by a given prototype graph but say nothing about the correctness or appropriateness of that prototype. In Chapter IV a method will be presented which can be used to verify some properties of prototypes, but this method depends on a higher-level prototype graph which in turn must be supplied by the user as his definition of correctness.

2.4 Comparison of Graphs

Suppose we are given the state graph, G, of some program, and a prototype graph, P, which expresses the correctness properties we wish G to have. How do we verify that all of the traces in G are also in P? Is it possible to have a single decision procedure which will answer this question for any two graphs G and P?

The answer to the second question can be found in automata and formal language theory. Notice that state graphs as they have been defined here may

also be thought of as finite state automata (FSA). The set of all traces through the state graph is just the regular language accepted by the FSA. The problem of determining whether or not all the traces of G are also in P is exactly equivalent to determining whether $\mathcal{L}(G) \subseteq \mathcal{L}(P)$ (where $\mathcal{L}(G)$ denotes the language accepted by the FSA G). From formal language theory we know that this question is solvable for any two FSA. (This and all following assertions about the decidability or undecidability of questions in formal language theory can be found in Hopcroft and Ullman [14], Chapter 14.)

There are many possible procedures for deciding whether or not $\mathcal{L}(G) \subseteq \mathcal{L}(P)$. One, specifically adapted to the purposes and notation of this study, is presented below followed by a proof of its correctness. It must be kept in mind that the purpose of this algorithm is not merely to decide whether or not the graph conforms to the prototype, but, in case it does not, to make it as easy as possible to find the error in the original program. This consideration precludes changing the graph G in any important way.

It is assumed that the prototype graph P is deterministic, by which is meant that no node has two outgoing arcs with the same label. In addition, the algorithm requires that P be complete, by which is meant that if anywhere in the graph there is an arc with a given label then every node must have an outgoing arc with that label. If P is not already complete it can be made so by adding to it one node called a "sink node" and adding to every node which does not have an outgoing arc with a given label an arc so labeled from it to the sink state. The nodes of P are named p_0, p, p' , etc., and P has unit ϵ -arcs at every node. The nodes of G , the state graph of the program being analyzed, are named s_0, s, s' , etc. G is not necessarily either deterministic or complete.

Comparison Algorithm

1. Let L be a set of pairs (s,p) , where s is a node of G and p is a node of P . Initialize L to $\{(s_0,p_0)\}$ where s_0 and p_0 are the starting nodes of G and P respectively.
2. If there exist $s, p, x, s',$ and p' such that (s,p) is in L , $s \xrightarrow{x} s'$ is in G , and $p \xrightarrow{x} p'$ is in P , but (s',p') is not in L , then add (s',p') to L and repeat this step.
3. If there is a member (s,p) of L such that s is final but p is not, report an error, otherwise report success.

LEMMA 1. The algorithm is finite.

Proof: All the existential searches are of finite sets. The loop in step 2 must terminate eventually, because it continues only as long as it adds new elements to L , and L is a subset of the (finite) Cartesian product of the nodes of G and the nodes of P .

LEMMA 2. At the end of step 2, (s,p) is a member of L if and only if there exists an α such that $s_0 \xRightarrow{\alpha} s$ and $p_0 \xRightarrow{\alpha} p$.

Proof: To show that (s,p) in L implies the existence of α , induct on the number n of iterations of step 2 before (s,p) is added to L . For $n = 0$, we have $s = s_0$ and $p = p_0$. Clearly $\alpha = \lambda$ (the empty trace) suffices, since $s_0 \xrightarrow{\lambda} s_0$ and $p_0 \xrightarrow{\lambda} p_0$. For $n > 0$, at step $n-1$ we had (s,p) not in L , $s' \xrightarrow{x} s$ in G , $p' \xrightarrow{x} p$ in P , and (s',p') in L . By the induction hypothesis, there exists a β such that $s_0 \xRightarrow{\beta} s'$ and $p_0 \xRightarrow{\beta} p'$ because (s',p') is in L . Now $\alpha = \beta x$.

To show that $s_0 \stackrel{\alpha}{\Rightarrow} s$ and $p_0 \stackrel{\alpha}{\Rightarrow} p$ imply that (s,p) is in L , induct on the length of α . For $\alpha = \lambda$, we have $s = s_0$ and $p = p_0$, and we know that (s_0,p_0) was put in L at step 1. If $\alpha = \beta x$, then $s_0 \stackrel{\beta}{\Rightarrow} s' \stackrel{x}{\rightarrow} s$ and $p_0 \stackrel{\beta}{\Rightarrow} p' \stackrel{x}{\rightarrow} p$. So (s',p') is in L by the induction hypothesis. By step 2, (s',p') in L , $s' \stackrel{x}{\rightarrow} s$ in G and $p' \stackrel{x}{\rightarrow} p$ in P imply that (s,p) is put in L .

THEOREM. The algorithm reports an error at step 3 if and only if G generates a trace not generated by P .

Proof: If step 3 reports an error, then there exists a (s,p) in L with s final and p not. By Lemma 2 there exists an α such that $s_0 \stackrel{\alpha}{\Rightarrow} s$ and $p_0 \stackrel{\alpha}{\Rightarrow} p$. Thus G generates α . But since P is deterministic and p is not final, P does not generate α .

If there exists an α which is generated by G but not by P , then there exists a final s such that $s_0 \stackrel{\alpha}{\Rightarrow} s$. Since P is complete, there exists a p such that $p_0 \stackrel{\alpha}{\Rightarrow} p$. Since P is complete, there exists a p such that $p_0 \stackrel{\alpha}{\Rightarrow} p$. By Lemma 2 (s,p) is in L . But since P does not generate α , p is not final, and step 3 will report an error. Q.E.D.

One would probably not implement the comparison algorithm as it is written. For one thing, rather than having many pairs scattered throughout L of which s is the first member, one might prefer to associate one set of nodes of P with each node of G in order to make clearer which sections of the object program correspond to which states of the prototype. The searches can be made deterministic. More importantly, it is not actually necessary to make P complete; one can modify the test at step 2 so that if (s,p) is in L and $s \stackrel{x}{\rightarrow} s'$ is in G but there is no outgoing arc from p labeled x , then the algorithm

reports an error at once. This is because the x -arc from s to s' must be part of a trace in G which is not in P because P is deterministic. It will be assumed in the examples which follow that this second modification has been made to the algorithm.

This modification to the comparison algorithm raises a point which must be discussed in more detail. Consider the two graphs in Figure 2.2. In both graphs the node labeled 0 is the starting and only final state.




Figure 2.2. Graphs of Two Erroneous Programs

For both graphs the intended prototype is the graph P of Figure 2.1. The first graph, G_1 , is clearly wrong; it contains the illegal sequences $\epsilon(AV)^*$. The original algorithm will not report these illegal sequences, however, and technically that would be correct since they do not end at a final state and hence are not elements of $\mathcal{L}(G_1)$. Obviously, however, we would like to be told about such illegal traces in programs, and modifying the algorithm so that failure is reported when arcs in G have no corresponding arcs in P accomplishes this end. The second graph, G_2 , has no illegal sequences, but like G_1 it has a region from which there is no path to a final state. This presumably reflects an error in the program, but not one which the comparison algorithm will detect.


A very simple algorithm will detect them, however, First work from

s_0 forward along arcs, marking those states which can be reached from s_0 ; and then work from the final states backward along the arcs, marking states from which some final state is reachable. Next consider states which do not have both marks: States without the first mark are not reachable at all and may be deleted from the graph. States with the first mark but not the second have no path to a final state.

In using the comparison algorithm on the state graph G of a program, one has three choices: the above marking algorithm can be applied to G before giving it to the comparison algorithm, or one can simply assume certain properties of G without checking, or one can apply the comparison algorithm to all graphs with the understanding that technically it detects sequences rather than elements of $\mathcal{L}(G)$.

As an example of the use of the comparison algorithm, suppose we wish to verify that $\mathcal{L}(G) \subseteq \mathcal{L}(P)$ where G and P are the graphs in Figure 2.3.

G:



P:




Figure 2.3. Example Graphs

Node 3 is the final state of G ; 0 of P . Initially L contains only $(0,0)$. The two A-arcs from node 0 cause $(1,1)$ and $(2,1)$ to be added to L . Then the C-arc from 1 to 3 adds $(3,0)$ to L , and finally the B-arc from 2 to 1 causes $(1,2)$ to be added. Table 2.1 summarizes the results of the algorithm so far.

TABLE 2.1

PARTIAL EXECUTION OF THE COMPARISON ALGORITHM

L	Reason for Being Added to L
(0,0)	Initialization
(1,1)	$0 \xrightarrow{A} 1$ in G, $0 \xrightarrow{A} 1$ in P
(2,1)	$0 \xrightarrow{A} 2$ in G, $0 \xrightarrow{A} 1$ in P
(3,0)	$1 \xrightarrow{C} 3$ in G, $1 \xrightarrow{C} 0$ in P
(1,2)	$2 \xrightarrow{B} 1$ in G, $1 \xrightarrow{B} 2$ in P

Nothing more can be added to L so we proceed to step 3. The only member of P paired with 3, the final state of G, is 0 which is the final state of P, so the algorithm reports success.

The comparison algorithm could be made considerably simpler by requiring that G be deterministic. In fact, there are algorithms for constructing from any nondeterministic FSA a deterministic one which accepts exactly the same language. However, these constructions may distort the original graph and destroy correspondences between states of the graph and sections of the program and thus, in the case when the algorithm reports failure, make it harder for the user to identify the faulty segment of code containing the illegal path. In all the methods of state graph manipulation and analysis presented in this study, one aim is the retention in the graph of as much information about the original object program as is practical.

Prototype state graphs allow one to state as correctness criteria any member of the whole class of regular languages. From an automata theory point of view, the method developed so far models the object program as an

FSA and then compares this FSA with another. With these observations in mind, it is natural to ask whether the method can be extended to other models and languages. Would it not be possible, for example, to model programs as pushdown automata and compare them to prototype pda, thus allowing ourselves to use as correctness criteria any context free language? The answer is no. From formal language theory we know that there can be no general algorithm for determining whether or not $L_1 \subseteq L_2$ when L_1 and L_2 are arbitrary context-free languages. The same is true when L_1 and L_2 are both any type of formal language except regular. The question is solvable when L_1 is regular and L_2 is an LR(k) language, but not when L_2 is context free or higher. This does not mean that the state graph of a program could not be used to determine if it adhered to some specific sequencing rule expressed as a higher type language by some ad hoc method but only that no general verification procedure is possible for such criteria. Thus one implication of the foregoing formal language theory results is that state graphs or some similar FSA-like model of programs are the most promising for development of general sequencing verification procedures.

2.5 State Graph Construction

Let us now turn our attention to the construction of state graphs representing programs. The general technique is to first read the source code of the object program line by line and construct an initial state graph which very much resembles the program's flow chart except that operations label arcs rather than nodes. This initial graph will then generally have to be extensively manipulated both to make it accurately reflect the program's sequencing properties and to compress it. In the discussion which follows, it will be

assumed that the object program is written in an assembly language; however, the basic principles are the same for programs written in some higher-level languages.

Building the initial graph is very much like assembling the object program, with the result being a graph rather than machine code. The source code is read in one line at a time and parsed. Then, depending on the operation in the line, appropriate actions are performed on the state graph.

Some lines result in no action at all, while some classes of instructions always require actions. Instructions representing actions which appear in the prototype graph cause the creation of a new arc, appropriately labeled, from the current node to a new node representing the state of the machine resulting from that action. Any instruction which affects the flow of control of the program will result in some addition to the state graph. For example, a branch statement results in two e-arcs from the current node to the two appropriate nodes. Under certain conditions these arcs may be tagged with the logical condition under which each path is taken. Any labeled instruction causes the creation of a new node corresponding to the state of the machine just before that line is executed. Program location labels may be attached to the nodes so that later the user may easily see which nodes correspond to which sections of program source code.

Techniques for building state graphs are somewhat dependent on the source language of the programs and thus may vary depending on the application, but basically are adaptations of well-known techniques in assembler and compiler writing. Chapter III contains more detail on the techniques used in one application as well as a discussion of some serious problems encountered in trying to build accurate state graphs. The present discussion will conclude with an example. Figure 2.4 shows a segment of MIX code purporting to