



**A LANGUAGE FOR MICROCODE  
DESCRIPTION AND SIMULATION IN VLSI**

**John Hennessy**

**TECHNICAL REPORT NO. 193**

**July 1980**

**This research was partially supported by the Joint Services, Electronics  
Program under contract no.  
DAAG29-79-C-0047**



# A Language for Microcode Description and Simulation in VLSI†

John Hennessy  
Computer Systems Laboratory  
Stanford University  
Stanford, Ca. 94305

Technical Report 193

July 1980

## Abstract

This paper presents a programming language based system for specifying and simulating microcode in a VLSI chip. The language is oriented towards PLA implementations of microcoded machines using either a microprogram counter or a finite state machine. The system supports simulation of the microcode and will drive a PLA layout program to automatically create the PLA.

Key Words and Phrases: microcode, VLSI, PLA, finite state machine, design automation, computer-aided design, simulation, layout.

† This research was partially supported by the Joint Services Electronics Program under contract no. DAAG29-79-C-0047.



## 1. Introduction

VLSI chip design has rapidly become an area of great importance and interest. Presently few tools exist to assist the user in designing and debugging a VLSI chip. This paper discusses a programming language useful for the design of a microcode controller which will employ PLA implementation techniques. Microcoded engines, implemented either with a finite state machine technique or a microprogram counter, frequently form the control unit on a VLSI chip and thus constitute a **significant** portion of the work involved in the design.

Design of a microcode controller is an error prone task for several reasons. First, the programming language is extremely low level. The designer must deal with a binary machine without the benefit of a human-engineered interface. Secondly, many of the microprograms are large. This leads to a relatively complex program without a great deal of structure.

Another major difficulty is the significant level of detail which must be expressed. This leads to one of two pitfalls: either the microcode description is very low level and cluttered with details, which make it impossible to understand; or the designer uses an ad hoc higher level description of the microcode. This leads to problems, since the translation to the low level microcode must be done by hand, and the ad hoc description tends to be too informal and vague. Hence, without a higher level standard representation microcode programs are difficult to correctly write and virtually impossible to understand.

For microcode machines which will be implemented with PLA's, the process of laying out the PLA is both time consuming and tedious, leading to unnecessary errors. However, given the description of the PLA terms, in an equation or state table form, the process is mechanically straightforward. Therefore, such a system should provide for PLA layout as well as microcode design.

As a solution to these problems we propose a higher level language which has the following goals:

1. Provide a symbolic higher level language suitable for designing and documenting the microprogram and oriented towards implementation with PLA technology.
2. Provide simulation tools to debug the microcode.
3. Provide for automatic layout of the PLA based on the microprogram.

## 2. Requirements

The goals for this design aid lead to language requirements. The desire to simulate the microcode leads to the requirement to be able to describe the subsystems which interact with the microcontroller. We will refer to these subsystems as the *environment*. Describing the environment

can be easily done in a conventional programming language, whose interaction with the microcode occurs in a restricted and well defined manner.

The micromachine description must be separate from the environment description. There are several reasons for this:

- it increases comprehensibility of the micromachine structure;
- a specialized language is more appropriate for the microprogram design, since in general it is impossible to determine if a given program describes a finite state machine;
- it simplifies the translation process.

Since the end product of this design is a finite state machine implemented with PLA techniques, a method of incorporating details about the implementation is required. This specification should include mappings between functional specifications in the environment and actual PLA outputs, as well as timing specifications which force outputs to occur earlier or later than they occur in the program. Including these details separately allows a more functional orientation in the microcode description.

- Lastly, details concerning the actual PLA layout are needed. These include such things as the number of PLA's and the positioning on the PLA of each signal.

### 3. Environment Definition

The environment of the finite state micromachine can be described in a conventional programming language; it must interface with the micromachine specification both for simulation and generation of the PLA. The environment will consist of data structures and variables which can be used to simulate the structure of the subsystems. The environment/controller interface is based on a set of functions and procedures. The functions, which must be type boolean, correspond to the inputs to the machine, while the procedures correspond to output signals.

Thus a conventional programming language can represent the environment as well as provide the necessary interface. We have chosen Pascal as the language for this purpose.

### 4. Control in the Micromachine

Control of program flow in a finite state machine uses a next-state function. Thus in the micromachine the next-state description should be included. Since the next state for most states in the machine is the successive sequential state, to have the designer specify the next state in these situations is wearisome. Therefore, support for a default next state is needed even though it may require extra work in implementation or add extra terms to the PLA.

The outputs associated with a given state should be conditional on a set of product terms only. Although arbitrary boolean expressions could be used, we choose not to because it requires a significant amount of processing to transform the expressions to a PLA oriented sum of products form. In this process the number of product terms added to the PLA may be substantial (up to  $2^n$  terms for an expression of length  $n$ ). The property that – the number of product terms in the PLA is approximately equal to the number of preconditions for the outputs – is very useful.

There are at least two major schemes for implementing a finite state machine which must be considered: a standard finite state implementation with a fixed state assignment, and an implementation based on a microcode engine with a microprogram counter. In the next section we discuss the approaches and the required control mechanisms for each.

## 5. FSM and Program Counter Implementations

The following table summarizes the advantages of the two approaches:

<u>FSM (state assignment)</u>	<u>Microcoded (program counter)</u>
Subroutines hard	Subroutines easy
Need next state terms in PLA	PC incremented by counter
Minimize states	Fixed assignment

It is clear that each approach has advantages and the language should support both approaches. In either case a method for explicitly specifying the next state is required. Additionally, in the microcoded approach a method for implementing subroutines with call and return is needed.

## 6. Defining the Micromachine

The micromachine is defined as a set of states, which are listed sequentially. Each state may optionally have a label, which denotes the state name. The state specifications are preceded by a set of specifications for outputs which are state independent. Thus the micromachine specification looks like:

```

FSM
[state-specification (for state independent outputs)]
state-name (optional) : [state-specification]
.....
```

state-name (optional) : [state-specification] .

A state specification is a list whose elements are either unconditional actions or conditional commands which are similar to guarded commands. A conditional command consists of a guard and a list of actions. A guard consists of a list of one or more product terms which are joined with or. A product term is a series of predicates joined with and. The predicate must be a call to a function in the environment; it will be associated with a PLA input. The interpretation of the command is: if all the predicates are true the actions should be executed. If there are no predicates, the guard is assumed to be true and the action is always executed in that state. Thus the form of a state specification is:

action

if  $p_1$  and . . . . . and  $p_n$  or . . . . . or  $q_1$  and . . . . . and  $q_m$   $\Rightarrow$  action

where the  $p_i$  are function invocations and the  $q_i$  are product terms, like the first term.

Each state may contain a list of such specifications; the entire state is bracketed. During simulation the specifications are evaluated and executed sequentially; however in the actual PLA implementation these operations will occur in parallel. Therefore, side effects between procedures which are outputs and functions which are inputs in the same state should be employed with great care.

### 6.1 Actions

There are two types of actions allowed: outputs and state change operations. A list of actions can be used as a single action by bracketing them. Outputs are invocations of procedures in the environment and correspond to PLA outputs. The state change directives dictate the next state: all state change directives have effect only after the current state is completed. That is, all state specifications with true guards will be executed in a state, regardless of any state change directives in that state. The state change directives are:

**next**    *state-name* - makes *state-name* the next state

**call**    *state-name* - does a procedure call to the routine at *state-name*

**return** - returns to the state sequentially following the calling state

### 6.2 A short example

```
{ Multiplication of 10 bit unsigned integers }
program multiply;
const  maxint = 1048576; (ie 220)
       one_half = 524288;
```



```

type    register = 0..maxint;
{ z ← x * y }
var     x, y, z : register;
         counter : 0..10; { count iterations }
procedure ZeroZ;
begin    Z := 0 end;
procedure GetXY;
begin    { Read in X and Y, for simulation only } read (x); read (y); end;
procedure ZeroCounter;
begin    Counter := 0      end;
procedure SHL (r: register);
begin    if r > one-half then  r := r - one-half;  r := r * 2 end;
function MSBY: boolean;
begin    MSBY := y > one-half end;
procedure AddXtoZ;
begin    Z := Z + X      end;
function LT: boolean;
begin    LT := counter < 1  end;
procedure Incr;
begin    counter := succ(counter)      end;

fsm { Definition of the machine }
  [ { No state independent actions } ; ]
Start: { Start state: Z, Counter ← 0; get simulation values for X, Y }
  [ZeroZ;      GetXY;      Zerocounter]
Loop: { Main multiplication loop }
  [SHL (z)]
  [if MSBY =>  AddXtoZ]
  [SHL (y)]
  [if LT = > Incr; next Loop { End multiplication loop}] .

```

## 7. Defining the Relationship to the PLA

The relationship between the microcode specification of the control program and the PLA is defined by the input signals, the output signals, and the mapping between environment procedures and output signals. The following sections deal with these topics.

### 7.1 Defining input and output signals

PLA signals are defined by means of input and output signal declarations; their relationship to the environment routines is established through signal definitions. Input and output signal declarations appear just before the definition of the environment procedures.

Signal declarations begin with the keyword `inputs` or `outputs`, as appropriate. The general form of each declaration is then:

```
(name [ '(' bounds ')' ] [ ':' parameters ] );
```

The list of names are the names of input or output signals being declared. The optional bounds designator indicates whether a particular signal is a single line or a number of lines. In the latter case the line is treated as an integer encoded number. The optional parameters are associated with all input/output names in the declaration.

There may be zero or more parameters; the following parameters are legal:

<u>syntax</u>	<u>meaning</u>	<u>permitted on input/output</u>
<b>pla</b> ( <i>n</i> )	Associate signal with pla # <i>n</i>	both
<b>top</b>	Position signal on top of pla	both
<b>bottom</b>	Position signal on bottom of pla	both
<b>renames</b> (signal-name)	Gives a signal another name	both
<b>earlier</b> ( <i>n</i> )	Move the signal <i>n</i> states earlier	output
<b>later</b> ( <i>n</i> )	Move signal <i>n</i> states later	output

The optional pipelined directives, i.e. `earlier` and `later`, move an output signal forward or backward in the state graph. This is very useful when a particular signal, which is logically associated with a single operation, must occur earlier. A frequently occurring example of this is precharging or enabling of alu's. Although the operation appears to occur in a single state the alu must be enabled one state earlier. The pipelined directives provide a convenient way to express such relationships without adding needless details to the microcode description. The `renames` directive gives a signal another name, without associating the other characteristics (c.g. pipelining) of the signal.

### 7.2 Describing the relationship between environment and outputs

Since a procedure of function in the environment can logically correspond to one or more signals, a method of defining the mapping between environment functions/procedures and signals is needed. Allowing this capability has two significant effects: it allows the microcode description to be functionally oriented, and it significantly decreases the amount of code needed to describe the PLA implementation of the microcode.

The mapping between an environment procedures and signals to be generated in the PLA is given in the definition section of an environment procedure. The definition section starts with the keyword `definition` and appears immediately after the function or procedure header. Procedures in the environment without a definition section are presumed to be for simulation purposes only. The definition section consists of a list of signal definitions which are separated by semicolons; the

definition section is terminated by end.

A signal definition has the form:

[ *pattern-string* : ] *signal-list*

The optional *pattern-string* is used to specify different signal combinations based on the values of the parameters to the environment procedure. The pattern string consists of a set of string patterns separated by comma and enclosed in brackets. The string patterns are used to match the substitution list on a one to one basis. If the pattern matches the list of actual parameters, the signals in the signal list are **generated** as outputs. Each string pattern can either be a alphanumeric string or a \*, the latter indicating that any actual parameter value should generate a match for the corresponding parameter.

The *signal-list* may contain the names of signals which have been declared in the output declaration section and other environment procedures. Each signal expression is a single signal, a conjunction of two signals, an assignment to a signal vector (i.e. a signal with more than one line), or a concatenation of identifiers and constants (which is then treated as a single signal). Concatenation is indicated by the symbol &. There are two forms that each single signal can have. The first is a single identifier: the second format is that of a procedure invocation. In the second case any parameters in the parameter list which correspond to environment procedure formal parameters are replaced by the actual parameter strings. If the signal identifier is an environment procedure and not a signal name, the definition section of the referenced environment procedure is used for that signal. Naturally, the procedure name can be followed by parameter strings. This facility allows multi-level procedures to produce signals by composing the definition list in each procedure.

### 7.3 An example

Consider the earlier example of 10 bit multiplication. We will assume that our machine will be controlled by PLA outputs which are defined using *r* as a register (which can have values 0-7) as follows:

enabler - put *r* on an alu input bus, this must occur one state before an alu operation

loadr - put the alu output bus into *r*

shiftright, shiftright - shift left and right

zero - send a zero to the alu

msbr - true if the msb of *r* is on

Itr - true if *r* is less than 1

add, incr • alu operations

Using these outputs the following declarations would be added to the example for the procedure SHL:

inputs

lt[0..7], msb [0..7]: bottom;

outputs

enable [0..7] : earlier (1) top:

load [0..7], **shiftright**, shiftright, zero, add, incr : top ;

With these output signals declared WC can specify the relationship between the environment procedures and the signals they should generate.

procedure SHL (r: register);

definition

enable <- r ; shiftright ; load <- r end;

procedure ZeroZ;

definition

zero and load <- 2; {z register}

procedure ZeroCounter;

definition

zero and load <- 3; {counter register}

function MSBY: boolean;

definition

msb [1]; { signal bit for Y (register 1) }

procedure AddXtoZ;

definition

enable <- 2; add; load <- 2; (no signal for enable x since it is 0)

function LT: boolean;

definition

Lt <- 3; (counter register}

procedure Incr;

**definition**

```
enable <- 3; Incr;
```

The appendix contains a more detailed example which requires more elaborate output definition facilities.

**8. Using the Language**

The description language can be used to drive a microcode simulation as well as generate a PLA layout. In order to do simulation the microcode description is supplied with the signal declarations and definitions. The user may supply start and end states for the finite state machine, if they exist. The simulation is presently done by creating a Pascal program which embodies the semantics of the microcode. Two types of simulation output can be produced:

1. Functional simulation with state traces and outputs of additional information as dictated by the environment procedures.
2. Functional simulation with input/output signal tracing.

PLA generation is a straightforward process, which is done in two parts. The first part analyzes the microcode structure and creates product term lists for each output. The effect of signal definition and pipelining is integrated before making these lists. The PLA layout is then done by a separate program which inputs the signal descriptions and the product term lists. This structure allows the insertion of programs which optimize the finite state machine.

**9. Other Uses of the System**

There are several other useful types of debugging and checking of microcode which can be done in the process of simulation. Most important among these are detecting potential errors which arise because the simulation does not exactly match the PLA implementation, or because the microcode does not employ the environment in a manner which the hardware is designed to support. Another class of errors arises from the fact that the simulation may fail to test all possible combinations of inputs or fail to test all states.

The major reasons why the simulation and PLA implementation might behave differently arise from the fact that the simulation treats outputs, environment procedures and the state as unique entities. In the PLA these objects are interrelated, so problems such as assigning two next states are resolved into a single well defined action in the simulation, but result in a potentially harmful action in the PLA implementation. Certain classes of these errors can be caught by predefined microcode independent methods, but others require a more general scheme, which we can also employ to resolve errors concerning the use of the hardware environment by the microcode.

Many of the errors in the use of the hardware environment by the microcode arise from situations where certain outputs are being incorrectly used, perhaps with respect to timing, or the hardware is being instructed to perform some task it is not physically able to complete. Many of the latter types of errors can be caught using a strictly type checked environment specification. For example, suppose that the register file on some microcoded processor is divided into two sections in such a way that two registers from the same section can not be gated to the alu (many hardware micromachines have this property). Microcode errors which arise because two registers from the same section are being sent to the alu can be detected by defining the machine structure with two different types of registers and specifying that the alu environment procedures have two parameters - one from each register section. This class of simple errors is detected at compile-time. Other potential types of timing errors which arise with regard to pipelining, will generate warnings.

The more complex class of errors can not be detected with a straightforward compile-time scheme. Some examples of the type of errors that may arise are: attempts to use the bus for two different quantities in the same time frame, overlapping use of environment hardware, such as an alu, and incorrect timing with respect to which state an output actually occurs during. These sorts of errors can be detected during simulation if a set of assertions which are checkable during simulation can be specified. We divide these assertions into two groups: invariant assertions and state dependent assertions. The invariant assertions specify conditions which must hold regardless of the current state, e.g. if an alu output occurs in this state, the alu was precharged in the previous state and was not doing any other operation. State dependent assertions specify properties which should hold at a particular state, e.g. a certain part of the machine should have a certain value.

### 9. / Assertions

Anywhere an action can occur, an assertion can be specified. Although the assertion generates code for simulation purposes, no PLA entries are affected or generated. Thus the usage of assertions is merely to ensure that certain properties hold. An assertion has the form `assert invocation`, where invocation must be the invocation of a boolean function. Whenever execution reaches an `assert` statement at simulation time, the simulation invokes the specified function. If the function returns false the simulation is halted with an appropriate error message.

### 10. Conclusions

This paper has proposed a language and processing system for describing microcode whose implementation orientation is PLA based. The purposes of this language are to document the microcode at a reasonable logical level, while providing a firm specification; to allow extensive simulation, debugging, and error detection; and last to automatically create the PLA layout necessary to implement the microcode description.

There are many interesting questions concerning the applicability of this tool. Although it is designed to support a wide variety of applications there may be unknown but fundamental

limitations. It would also be interesting to examine its applicability to microcode machines whose architecture is not PLA based, but whose microcontrol is extremely straightforward; a major example of this implementation strategy is the two level microcode approach used in the Motorola 68000.

#### Acknowledgements

I am indebted to Jim Clark for providing motivation and insight during this research.

## Appendix 1 Syntax of the Microcode Specification Language

This is the syntax for the non-Pascal portion of the language. Terminal symbols in the grammar are distinguished by being in quotes.

program

-> 'program' '<id>' program←parms ';' outer←block ;

outer←block

-> const←part type←def←part var←decl←part io←part proc←part  
fsm ;

proc←heading

-> 'procedure' '<id>' formal←parms ';' definition←part => proc;

func←heading

-> 'function' '<id>' formal←parms ':' '<id>' ';' definition←part => func;

io←part

-> inputs←designators outputs←designators ;

inputs←designators -> ;

-> 'inputs' designators => inputs;

outputs←designators -> ;

-> 'outputs' designators => outputs;

designators

-> spec ;

-> spec designators ;

spec

-> vector←list ';' = > io←list;

-> vector←list ':' parameters ';' => io←list ;

vector←list -> vector;

-> vector ';' vector←list;

vector -> '<id>';

-> '<id>' '[' '<int>' ':' '<int>' ']' => vector;

parameters

-> parameter ;

-> parameter parameters ;

parameter

-> 'pla' '(' '<int>' ')' => planumber;

-> 'top' => top;

-> 'bottom' => bottom;

-> 'earlier' '(' '<int>' ')' => earlier;

-> 'later' '(' '<int>' ')' => later;

-> 'renames' '(' '<id>' ')' => renames;

definition←part



```

-> 'definition' def←list = > definition;
-> ;
def←list
-> definition ;
-> definition def←list ;
definition
-> patternmatch ':' output←list ';' => defline;
-> output←list ';' => defline;
patternmatch
-> '(' pattern←list ')' = > patterned;
pattern←list
-> pattern ;
-> pattern ',' pattern←list ;
pattern
-> '*' = > anything;
-> '<id>' ;
output←list
-> output ;
-> output 'and' output←list => andnode;
output -> plain←output;
-> 'not' plain←output => notnode ;
plain←output -> invocation;
-> invocation '&' output => cat;
-> '<id>' '=' constant = > vecout;
-> '<id>' '[' '<int>' ']' => bitout;
fsm
-> 'fsm' fsm←state '.' => fsmnode;
fsm←state
-> state←ind←part states ;
state←ind←part - > '[' state←specifiers ']' = > state←indnode;
states
-> states state ;
-> state ;
state
-> label '[' state←specifiers ']' => statenode ;
state←specifiers
-> state←spec ';' state←specifiers ;
-> state←spec ;
state←spec
-> 'if or←cond '= > action = > condnode ;

```

```

-> action ;
or←cond -> cond 'or' or←cond = > ornode;
      -> cond;
cond -> cond 'and' invocation => andnode ;
      -> invocation ;
invocation
      -> '<id>' parms = > invocationnode;
parms -> ;
      -> '(' id←or←const←list ')' ;
id←or←const←list
      -> constant ;
      -> constant ',' id←or←const←list ;
id←list -> '<id>';
      -> '<id>' ',' id←list;
actions
      -> action ;
      -> action ';' actions ;
action
      -> '[' actions ']';
      -> 'assert' invocation = > assertnode;
      -> invocation => actionnode;
      -> 'next' '<id>' = > gotonode;
      -> 'call' '<id>' = > callnode;
      -> 'return' => returnnode;
      -> ;
label -> ;
      -> '<id>' ':' = > labelnode;

```

## Appendix 2. Two More Examples

### *Appendix 2. I - Computing GCD*

```

program test (input,output);
var x,y: integer;
inputs
  eq1,eq0,gtx, gty: bottom;
outputs
  aluop[1..2] : bottom ;
  enablex,enabley: top earlier (1);
procedure init;

```

```

begin
    read(x);
    read(y);
end;
procedure subt (var a,b: integer);
definition
    enable & a and enable & b and aluop = 1;
begin    a := a-b end;
Function greater (x,y:integer): boolean;
ddefinition
    gt & x ;
begin greater := x>y end;
function equal (x,y:integer): boolean;
definition eq & y;
begin eq := x = y; end;
function ne(x,y:integer): boolean;
definition not equal (x,y);
begin ne := not equal(x,y);
end;

fsm
[;]
one:  [ init ;
      assert ne(y,0); .
      if equal(x,0) => next endstate ]
      [ call two ]
      [ next one ]
two:  [ if greater(x,y) => [subt (x,y); next two ];
      if greater(y,x) => [subt (y,x); next two ] ]
three: [ assert equal(x,y);
       if equal(x,1) => [writeln(1); return];
       if ne(x,1) => [writeln(y); return]]
endstate: [ halt ].

```

*Appendix 2.2 Traffic Light Example from Mead and Conway*

```

program traffic(input,output);
const short = 2;
      long = 4;
type colortype = (green,yellow,red);
      signaltype = 0..1;

```

```

var   time: in tegcr;
      hl,fl: colortype;
      c: signaltype;
inputs
      c,tl,ts :bottom;
outputs
      st,hl[1..0],fl[1..0] : bottom;
procedure getinput; (* for simulation purposes only *)
begin
      write('cars? ');      read(c): end;
procedure timer; (* for simulation purposes only *)
begin
      if time < long then time := time + 1 end;
procedure highlight(color: colortype);
definition
      (green): hl = 0 ;
      (yellow): hl = 1 ;
      (red): hl = 2 ;
begin hl := color end;
procedure farmligh t(color: colortype);
definition
      (green): fl = 0 ;
      (yellow): fl = 1 ;
      (red): fl = 2 ;
begin fl := color end;
procedure starttimer;
definition st;
begin time := 0 end;
function cars : boolean;
definition c;
begin cars := (c = 1) end;
function notcars : boolean;
definition not c;
begin notcars := not cars end;
function timeout(length: integer):boolean;
definition
      (long): tl ;
      (short): ts ;
begin timeout := (time >= length) end;
function nottimeout(length: integer):boolean;

```

```

definition
  (long): not tl ;
  (short): not ts ;
  begin   nottimeout := not timcou t(leng th)   end;
fsm
      [ getinput; timer] (* state independent component *)
highgm: [ highlight(green); farmlight(red);
  if notcars or nottimeout(long) = > next highgrn;
  if cars and timeout(long) = > [ starttimer; next highyel]
  |
highyel: [ highlight(yellow); farmlight(red);
  if nottimeout(short) = > next highyel;
  if timeout(short) = > [ starttimer; next farmgrn ]
  |
farmgm: [ highlight(red); farmlight(green);
  if cars and nottimeout(long) = > next farmgrn;
  if notcars or timeout(long) = > [ starttimer; next farmyel]
  |
farmyel: [ highlight(red); farmlight(yellow);
  if nottimeout(short) = > next farmyel;
  if timeout(short) = > [ starttimer; next highgrn ]
  ].

```

