

79WR373-2873

+-----+
! digital !
+-----+

INTEROFFICE MEMORANDUM

JAN 30 1979

1-388

TO: Jim Marshall, TW/A03
CC: Bill Demmer, TW/D19
Gorden Bell, ML12-1/A51

DATE: 26 JAN 79
FROM: Wayne Rosing
DEPT: PDP-11 Large Sys. Dev.
EXT: 247-2322
LOC/MAIL STOP: TW/C03

SUBJ: JANUARY STATUS REPORT

ICCS

The HYDRA ICCS activity is progressing fairly well in a number of areas and has suffered from neglect in a few others.

The areas that are going well have to do with the technical aspects of the bus. We're fairly comfortable at this point in time that we will be able to design the appropriate ECL front end logic to support bit rates in the 80 to 100 megabit range. Our current strategy is to complete a preliminary design of the incode-decode, sync detect, serializer, and buffer control logic, and CRC circuit by the end of this week. With this design we will then cost out a number of different alternatives, including all TTL interface, a mixed ECL TTL interface, and the fastest possible interface we can think of. We will then play these designs against the specs of a number of different logic families, including the Dolphin gate array, and try to determine an overall cost curve in the range of 25 through 100 megabits. With this information, and such other things as chip counts, power consumption, board area, we expect to be able to make a decision concerning these specific ICCS bit rate. The preliminary data we're getting from EMI/RFI testing indicates that, we will be able to get adequate common mode rejection with the transformer isolation scheme. This particular area of the whole ICCS effort is clearly the most controversial and risky and so a great deal of effort has been concentrated here, and I think to fairly good advantage.

The areas that are weak in terms of the ICCS basically boil down to our lack in adequately documenting what we're doing. I will make a specific attempt to remedy the situation by gathering the engineers together and tasking them during February to write up a complete preliminary specification on the chosen design as soon as we get the cost data.

JANUARY STATUS REPORT

The ICCS interconnect topology is still being considered in detail. We are looking at a single bidirectional cable versus two cables, and a topology that will allow a two cabinet point-to-point interconnect, a 8 to 16 node passive (resistor) star coupler, and the option to design an active multi-trunk exchange node. The last option is something that we're planning for future expansion. In very large systems we may find by the end of the '80's that an aggregate system bandwidth of 10 megabytes is insufficient. With the multi-trunk line we can essentially have 10 megabytes of communication between two nodes, multiplied by the number of active trunks at any given time. It's clear that even a relatively small exchange (four trunks) will provide all the bandwidth we're ever likely to require in ICCS configuration. The consideration of two coax cables boils down to recognizing that by having unidirectional cables, we can in fact add repeaters and operate over longer distances. Although we may not specifically need this for our local configurations, it is certainly useful in an area network between a number of different buildings, or in an ARPA configuration. The specific decisions in this regard are relatively minor and have no significant effect on the program as long as they're all ironed out and included in our published ICCS specifications which, as I mentioned above, we will try to get out by the end of February.

Bill Strecker, Bob Metcalfe, Bob Stewart, and myself met with Tony Lauck last week and a serious discussion of the VAX ICCS Port Interconnect scheme ensued. I perceived Tony didn't shoot any serious holes in what we were proposing. As a matter of fact, he was not fully familiar with the VAX Port and had not read the Rev. 2 document issued last October. Bill sent him a copy and I expect we will be meeting with him next week to review that document. Tony did make many constructive suggestions, many of which we have immediately incorporated into our work, and some of which are still under evaluation. The last ICCS documentation that is planned has to do with the VAX Port. Bill has committed to get Rev. 3 out before he leaves for Berkely the end of February.

To sum up, my expectation is that we will be able to have a fairly substantial documentation of the ICCS and VAX Port effort available no later than the first of March.

I had a meeting yesterday with people from CSS. They are interested in replacing their PCL11 with a new, improved intercomputer link. We had a long discussion about the terminal bus and the ICCS and they seem very happy with the overall scheme. They were a little put off by the 100 megabits, but I assured them that that was a design goal and we'd live with whatever bandwidth we could achieve for an acceptable cost. In looking at the overall product strategy, the CSS people essentially came to the

JANUARY STATUS REPORT

conclusion that a logical thing for them to do would be to design a Unibus port for the ICCS that would work into our ECL line interface card. I told them that I had no fundamental problem with CSS going ahead and joining our effort. I'm acting under the

assumption that CSS would fund the development of the Unibus port for all 11's and 10's. They asked if there was any problem with potentially shipping ICCS's earlier than HYDRA. I said no, since I couldn't imagine why there would be. Bob Metcalfe has pointed out that a CSS effort like this could result in our ability to give a rather timely response to Kahn's request for an improved interconnect for use in the ARPA network. I will be sure that Bob, Sam Fuller, and CSS follow up this opportunity.

TERMINAL BUS

The terminal bus activity is beginning to gel. I am working with Art Lim and we are beginning touring around the corporation talking with most of the heavy potential users of such a bus. I've been explaining the Tewksbury concept of the bus and what our overall implementation strategy is. I have talked already with Microcircuits, IPG, Chuck Stein (Distributed System Product Management), and Len Halio. None of these groups has expressed any significant pain with our proposals, and I think it's fair to say that nobody is so far apart that we can't arrive at a universally acceptable conclusion if we put our minds to it.

The plans for the next few weeks are for Art, Bob Metcalfe, and myself to continue discussing our ideas with the balance of the concerned groups in the company. Specifically, we have to have some extended conversations with Merrimack Communications and the Terminal Group.

Bob Pratt, Chuck Mathis, and Art Lim, are in the process of doing a detailed study of the cost versus bandwidth tradeoff for the terminal bus. In specific, we're looking at the various topologies (ring, loop, multidrop; contention versus polled) and the cost of various integrated circuit chips (Silog SIO, Intel SDLC chip, use of DMA micro-controller chips, and a DEC custom chip.) I expect to get a preliminary report out on the cost of interconnect in the communications frequency range by mid-February. We will not publish this report, however, until it's been reviewed by both the Communications group and the Terminals group.

JANUARY STATUS REPORT

STRATEGY DOCUMENT

A draft version of the strategy document was circulated early last week to approximately 16 people in Tewksbury. I had asked for responses and opinions to be forwarded to Strecker or myself by yesterday. As of yesterday evening, I had received essentially three responses on the document, one favorable, one completely neutral, and Steve Rothman's comments which are appended hereto! I continue to be disappointed by the lack of interest on the part of the people who aren't immediately working on this project in terms of their willingness to take 10 or 15 minutes and jot down their opinions. It's extremely frustrating when we find people whispering in other people's ears about how flakey our activities are, and then when you go and press them about it, you find they have either not read what documentation they've been sent, or simply totally ignorant of what we're doing at all.

My defense against the above syndrome is basically to try to get more published in the form of substantial documentation and targetted in at a lower level where people are more apt to have the time and enthusiasm to constructively criticize our efforts. Strecker, Metcalfe and myself will be rewriting this document and collating what comment there is into a coherent whole. It is my plan to have a revised "Tewksbury Strategy" ready for fairly wide broadcast by the end of next week.

NEBULA

In terms of the immediate breadboard effort, there is nothing exciting to report. The Nebula CPU currently runs its micro diagnostic exerciser. The memory controller is in the final stages of physical construction, and we expect to begin debugging next week. We expect to be able to run base machine instructions and operand specifiers by the end of March. This benchmark corresponds, in a sense, to Comet's recent running of the hard core diagnostic exerciser. With the results of that test we will get a preliminary indication of Nebula's processing speed and will have verified a sufficient amount of the hardware that I'm confident we can release for a prototype etch early in Q4.

Ken Okin and Stan Lackey are to be congratulated for an outstanding bit of hardware design. The number of errors found so far are remarkably low.

26 JAN 79

JANUARY STATUS REPORT

The major issue concerning Nebula right now has to do, of course, with the PALS from MMI. I've dealt with this topic in other memos (copies of which are attached for reference) and I won't belabor it here. We do not feel there is any significant threat to the Nebula program. The decision point that's coming up essentially is keyed around whether or not we should risk waiting for the PALS, decommit and go do a gate array machine (thereby delaying the FCS at least a year), or doing an off-the-shelf implementation. A couple of conversations with product lines has indicated that the product lines who really want Nebula want it now. I would imagine LDP and OEM would vote for a slightly more expensive machine made out of off-the-shelf components, but one that could be announced at the Fall DECUS and go into volume shipment by Q1 of '81. The alternative to a gate array approach is that we could probably design a quad form factor 2 board Nebula equivalent machine and essentially build what could be called a poor man's LSI VAX. The attraction to a quad Nebula, of course, is that it can then be packaged inside of standard terminals. I see the negatives as being the approximately two year time frame after Comet announcement before the next lower VAX comes into being, and the increased engineering investment. It's doubtful at the system level we would save enough in cost to make it worth our while, although at the personal computer level a quad machine would no doubt save us substantial money.

We now have a product manager for Nebula (Lou Phillipon) and Ken Okin has assumed the role of Technical Project Manager. I'm submitting a requisition for the program manager who will be the individual who will functionally run the program on a day to day basis.

My involvement with Nebula now is sufficiently low that I do not feel hampered in concentrating my efforts almost exclusively in IO. However, in order for the Nebula program not to go awry we must get that program manager identified and on board in the next two months.

It's our current expectation that we will be able to run VMS and all VAX instructions by the June time frame. We have a level of effort activity going, aimed at allowing the VMS people to do debugging of Nebula support of the core VMS routines during May. The reason for this activity is to insure that Nebula will be supported in Rev. 2 of VMS. If this activity does not happen it's certainly not the end of the world, but at least we do not have the problem of having Nebula availability occurring midway between Rev. 2 and Rev. 3 releases of VMS.

26 JAN 79

JANUARY STATUS REPORT

Nat Parke's group will be responsible for the development of any personal computer hardware that needs to surround Nebula and any interfacing with Universities and R&D that is required for us to launch some reasonable breadboard effort of a personal computer for the University environment. I expect to only throw potshots at this activity and look forward to seeing some spectacular breadboards in Nat's lab running by the end of next summer! I'm encouraging Nat's group to consider also the possibility of looking at some of the new Winchester technology disks that are emerging on the OEM market. (Shugart, amongst others).

BUDGETS

As far as I'm concerned, the budget issues concerning cost center 373 are under control and there are no surprises. Discussions of the other day have set a ceiling spending on the Nebula program of \$425,000. for fiscal '79. We are essentially comfortable with that number and will advise if at any time some opportunity presents itself for the constructive expenditure of more money. The projected fiscal year '80 Nebula funding (1.7 million dollars) looks fairly adequate. The biggest unknown, of course, is what fraction of that will ultimately have to be transferred to manufacturing to cover test equipment and other start-up costs. Again, I will advise if I see anything developing that would imply we could not meet the Q4 '80/Q1 '81 time frame that we're currently shooting for.

/bc

A TAXONOMY OF TECHNOLOGIES
IN FUTURE DEC PRODUCTS

A. HIGH-SPEED GATE ARRAY

ECL, e.g., DOLPHIN MCA

B. TTL GATE ARRAY

e.g., COMET, IBM Duchess (134 gates), IBM Purdue (704 gates)

C. HIGH DENSITY CUSTOM MOS

T-11, F-11, POSART, VLSI-VAX, JAWS/11

D. GLUE CHIPS

SSI, discretes, clock drivers; bridge LSI-11bus - S080bus; Unibus control

E. OFF-THE-SHELF IC'S

PLA, BIT-SLICES, ROM, μ P etc

- 2-chips are better than no chips

- 2 ps " " " " that

- fat yield.

- Design time "only" matters.

- Not sell chips!

~~400~~ \$400 LABY

→ Bkds in console 200 → 5%
→ 40K sp.

CM
1/24/79

~~TECH TRANSFER~~

TECHNOLOGY TRANSFER CALTECH TO DEC

Cheap color graphics CAD workstation
Jack Bunnos is installing 3 stations by June
Larry White (Caltech student) assigned to transportability problems

Summer students
Hank Walker IC design & architecture
Gary Tarolli CAD graphics software
Glen Okita IC design / architecture / CAD

VLSI architecture
2 MS Students
Larry Seiler Enter-chip partitioning
Greg Eflund TRIMOSBUS

VLSI VAX Advanced Development
Due by DEC # 1979:
1. CAD plan
2. Chip partitioning
3. I/O subsystem architecture
4. Chip specs

MEAD-CONWAY IC DESIGN COURSE IN FALL AT DEC
There are 3 DEC people (Paul Rubinoff, Steve Frank, Mudge) who have taken it.
(CMU Spring 78 (Sproll), MIT Fall 78 (Conway), Caltech)

Ray for Allen

CM
1/30/79

VLSI-VAX PROJECT

PHASE 1:

Starts ⁱⁿ June 1, 1979

By Dec 31, 1979 it produces

(1) VLSI-VAX Spec

- a. chip partitioning and architecture
- b. chip specs
- c. system product

(2) Design methodology

(based on experiments conducted here at DEC)

CAD tools

designer training plan

Small group

Craig Mudge

Jack Burness

← already on board

Dileep Bhandarkar

← Clayton's Task Force leader now

(Hank Walker)

(Gary Tansli)

↙ ↘ Caltech summer students

→ SSE person

→ CAD person

(Val Patel has agreed to supply)

Circuit designer

← looks impossible to fill this slot

PHASE 2:

Design & build VLSI-VAX

CM


1/30/79

VLSI-UAX

make

	hand	SC	GA	
CMOS	1/2 vol.	6L 5L 4L	X	no.
M	work DE	hi XYZ	com	
B	off for AZN.	Not Av	pro	
E				

Buy



d | i | g | i | t | a | l

interoffice memorandum

SUBJ: FOLLOW UP NOTICE

TO: Jim Cudmore, ML1-5/E30
Bill Green, ML1-4/E34
Roy Moffa, ML1-2/H26
✓ Joe Zeh, WZ2

Date: 11 JAN 79
From: Gordon Bell
Dept: OOD
MS: ML12-1/A51

Ext: 223-2237

Follow Up 1/19/79

I'd still appreciate an answer to the attached.

mj

RECEIVED
JAN 12 1979

Bill Green / Joe Zeh / Jim Cudmore / Roy Muffen + et al
KEA-016



cc: BT,)

INTEROFFICE MEMORANDUM

good point. What's right here? *W.B.A. FC/12/20*

TO: Gordon Bell cc: Skip Coombe
 Dick Clayton → Bob Kusik
 Jim Cudmore → Rich Perrin
 Bill Green → Herb Shanzer
 Jack Schneider → Ken Slater
 Joe Zeh → Mike Titelbaum
 Bill Johnson → Don Vonada

DATE: 30 NOV 78
 FROM: Keith B. Amundsen
 DEPT: Systems Interconnect
 EXT: 3360
 LOC/MAIL STOP: ML3/E67

Gordon

REC 08 12-56 1978

SUBJ: HMOS "STANDARD CELL LIBRARY" vs. "GATE ARRAY" DESIGN

Let's use HMOS NOW! (Pick 1 and proceed)

I have been involved in these tradeoff issues for a number of months now and, while not being directly associated with HMOS introduction, I have deduced that there is only a small design time benefit from the gate array approach.

With the correct CAD tools, an HMOS standard cell library design effort will take about one month longer than a gate array effort. On the other hand, there are many advantages in doing a standard cell design with respect to a gate array design.

There are 8 times as many standard cell types in the standard cell library as there are gate array cell configurations. Also, the cells themselves can be more densely packed. As a result, devices will be faster because functions can be synthesized with fewer levels of cells and less interconnect metal. (Faster, dynamic logic is also available only with standard cells.) Normalized functional yield will be higher, normalized functional unit cost will be lower, and parts count will be reduced, saving board real estate.

The advantage to going the gate array route is that it has the potential for getting sample first-pass parts out a few weeks sooner by starting with partially processed semiconductor wafers. Also, CAD tools are simpler for gate arrays but they are fairly complex for either method. A problem with gate array cells is that they can never all be used because of the fixed array size and topology. This causes the die to be bigger than it has to be and the yield to go down.

In my opinion, assuming the CAD tools are in place, the Standard Cell approach has the potential to yield lower cost and higher performance parts than the HMOS gate array, at a minimal cost in design time. Therefore, LSI CAD and/or CAD should research the ability of DEC to either develop or buy the appropriate auto-placement and routing CAD tools for the Standard Cell technique.

digital

INTEROFFICE MEMORANDUM

TO: Ken Slater
cc: Distribution List

DATE: 15 January 79
FROM: Joe Zeh
DEPT: Microproducts Development
EXT: 8-238-2408
LOC/MAIL STOP: WZ-2

SUBJ: HMOS "STANDARD CELL LIBRARY" VS "GATE ARRAY" DESIGN

Ken, Val Patel and I may have screwed up by temporarily pulling Ivan Dobes off the Standard Cell CAD package evaluation in favor of Dolphin. I believe that is still the right decision. However, I also believe that what Ivan was about to do is the right approach and we should make every effort to not let it slip by more than three months.

In any case, I think we can still answer Gordon's question (see attached) and Keith Amundsen's memo is a super start. As part of your Quick Turn Around effort in preparation for our February 15th OOD presentation would you be sure that we have this answer to present at that time.

JZ/ds
Attachments

Distribution List:

Ken Slater	WZ-2
Gordon Bell	ML12-1/A51
Dick Clayton	ML12-2/E71
Jim Cudmore	ML1-5/E30
Bill Green	ML1-4/B34
Jack Schneider	WZ-2
Bill Johnson	ML21-3/E87
Don Vonada	ML3-3/E67
Keith Amundsen	ML3/E67
Roy Moffa	ML5-2/E93
Mike Titelbaum	ML1-2/E65
Bob Kusik	ML21-3/E87
Peggy Wesley	WZ-2

Kusik for your work ^{and} Best Bruce, Bob Kusik, Bob Armstrong, ^{FCU}
Carnegie-Mellon University

INTER-OFFICE CORRESPONDENCE

To: Gordon Bell
From: Dan Siewiorek
Date: September 26, 1978
Subject: Semiconductor Process Control and Design Aids at Sandia

Dick Clayton, Roy Moffa, Mike Teitelbaum.
Bill Green, Joe Zeh

This is when we should be headed?
Can you get Dan or Stew McConnell
from CMU in to
talk about this system? good

CAD
LST

I have recently returned from a visit to Sandia. I was very impressed with their fabrication facilities and design aids.

Process Control

Could we run it on our KL?

The fabrication facility was completely monitored and controlled by a computer. At each step of the fabrication process there is a CRT terminal with a menu outlining the steps in fabrication for each technology. The menu is designed so that even a novice, with minimal training, can successfully fabricate wafers. The CRT terminal also serves as a data entry port so that a detailed history of the wafer batch is kept on line. The computer can be programmed to automatically turn off diffusion furnaces (for long diffusions after everyone has gone home), purge furnaces before people arrive in the morning, and automatic insertion/withdrawal of wafers from furnaces to minimize the possibility of thermal cracking.

Design Aids

Sandia has design aids that handle from gate level descriptions in standard cells (with parameterized macro expansion capability) down through automatic layout and masks. The programs are written in Fortran and are being updated and converted to a PDP-10. The software is being used to fabricate a voter chip and went from standard cell design to mask check plots in less than two weeks using a card batch version of the software and done by a CMU student unfamiliar with the system. Some of the major programs are:

SALOGS-IV - Logic simulator. The SALOGS input format is the standard entry format to all their programs. Hence this format is the natural interface into their system. I have the SALOGS III manual and a rough draft of the SALOGS IV manual. CMU will get a copy of the PDP-10 code by November.

SICLOPS - Automatic mask layout program using standard cells and macro cells (ROM, RAM, etc.). The goal is to get within 10% of hand layout. CMU should get a copy of the PDP-10 code by February 1979.

Testing - A very important, often overlooked, aspect of VLSI is chip testing. Sandia is producing automatic test generation software.

Sandia hopes to support about six different technologies and the software is accordingly parameterized.

CMU is attempting to marry the ISP to gate level design aids with the Sandia gate level to mask level software. Ultimately one might envision going from ISP

to finished chips in less than a month, automatically. CMU feels that such chips will be within a factor of two initially, and perhaps within 20-30% eventually, of careful hand design.

Sandia is definitely on the right track and CAD is essential for VLSI (e.g., Intel layout people do 12-15 devices per day or 40 man years for a 100,000 device chip).

Case Study

SSI version running. →
12 Sept. Convert to Std. Cells (300 gates) ⇔ 60 Std. Cells)

MESSAGE 7 15

DATE: 20 SEP 1978 1852-EDT
FROM: STEPHEN MCCONNELL AT CMU-10A (X361SM30)
SUBJECT: HELLO AT LAST
TO: SIENIOWEK AT CMU-10A
MESSAGE-ID: <20SEP78 185212 SM30@CMU-10A>
DAN:

JUST GOT ON THE TIP SUCCESSFULLY TODAY. HAVE SUCCESSFULLY SIMULATED THE PIRE VERSION OF THE VOTER. SOME MINOR MODIFICATIONS (BUFFERING, ETC.) ARE AWAITING THE KEYPUNCH; BUT SHOULD GET THROUGH THE SIMULATOR TOMORROW. AFTERWARDS, THE LAYOUT! NOT MUCH ELSE NEWS. GOT A NOTE FROM MARIO ABOUT THE ISPS; HE HAD APPARENTLY SEND SANDIA A KL VERSION.

IDAHO

DATE: 22 SEP 1978 1546-EDT
FROM: STEPHEN MCCONNELL AT CMU-10A (X361SM30)
SUBJECT: PROGRESS REPORT FROM SANDIA
TO: SIENIOWEK AT CMU-10A
MESSAGE-ID: <22SEP78 154621 SM30@CMU-10A>
DAN:

HAVE FINISHED SIMULATION AND FAULT ANALYSIS; AND STARTED ON LAYOUT (USING SPARC). THE FIRST RUNTHROUGH; USING COMPLETELY AUTOMATIC PLACEMENT YIELDED A CHIP 227 BY 245 MILS. BY ITSELF; THIS WON'T QUITE MAKE IT; AS THE 48-PIN DIP PACKAGE HAS A WELL ONLY 240 MILS SQUARE; REQUIRING A MAXIMUM CHIP SIZE NO GREATER THAN 220 MILS SQUARE; BUT THERE'S A LOT OF ROOM FOR GETTING THE SIZE DOWN. PLAN TO GET A FIRST PASS AT A HAND LAYOUT DONE BY THIS EVENING.

UNLESS OVERRIDING CONSIDERATIONS PREVENT IT; AM PLANNING TO CHANGE THE FLIGHT BACK TO PITTSBURGH FROM SAN FRANCISCO TO ROUTE BY BOISE; AND LAYOVER AT HOME FOR THE WEEKEND; PROBABLY RETURNING FINALLY THE FOLLOWING MONDAY. WILL OF COURSE PAY WHATEVER EXTRA CHARGES ACCRUE AT THE TIME I MAKE THE CHANGE. THAT'S ABOUT IT FOR NOW. WILL KEEP YOU POSTED.

IDAHO

MESSAGE 8 16

DATE: 25 SEP 1978 1004-EDT
FROM: STEPHEN MCCONNELL AT CMU-10A (X361SM30)
SUBJECT: CHIP SIZE DOWN OKAY
TO: SIENIOWEK AT CMU-10A
MESSAGE-ID: <25SEP78 100440 SM30@CMU-10A>
DAN:

THE HAND LAYOUT RESULTED IN A CHIP 195 BY 212 MILS SQUARE. NOW WAITING TO GET CHECK PLOTS GENERATED. WILL HAVE TO DO SOME HAND-MODS; AS THE VCC PAD HAS TO BE MANUALLY INSERTED AND CONNECTED TO THE TTL OUTPUT PADS. THAT'S ALL THE NEWS FOR NOW.

IDAHO

SEG/CAD ORGANIZATION

ARNY GOLDFEIN IS OUR NEW GROUP MANAGER.

50 CAD ENGINEERS, MIXED HARDWARE/SOFTWARE BACKGROUNDS.

5 TECHNICAL WRITERS.

V-11 IS OUR MAJOR CUSTOMER THIS YEAR.

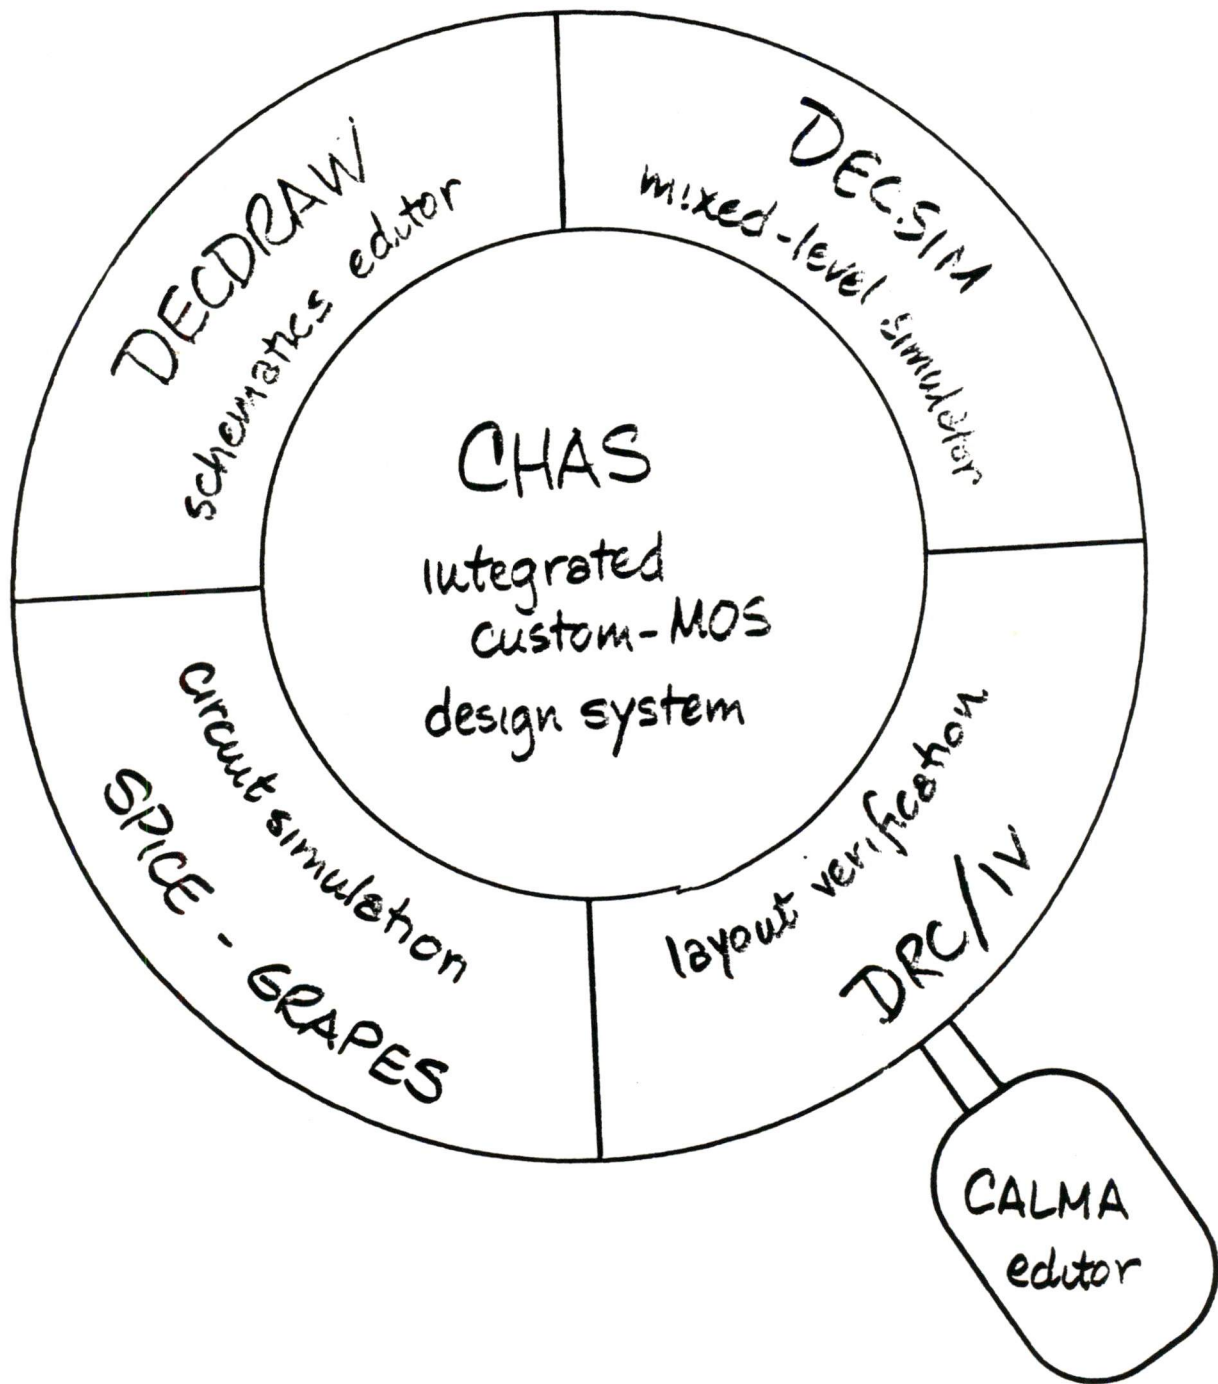
WE OWN CAD DEVELOPMENT AND SUPPORT FOR ALL TOOLS CRITICAL TO V-11.

~40 PEOPLE ARE FOCUSED ON V-11 TOOL DEVELOPMENT.

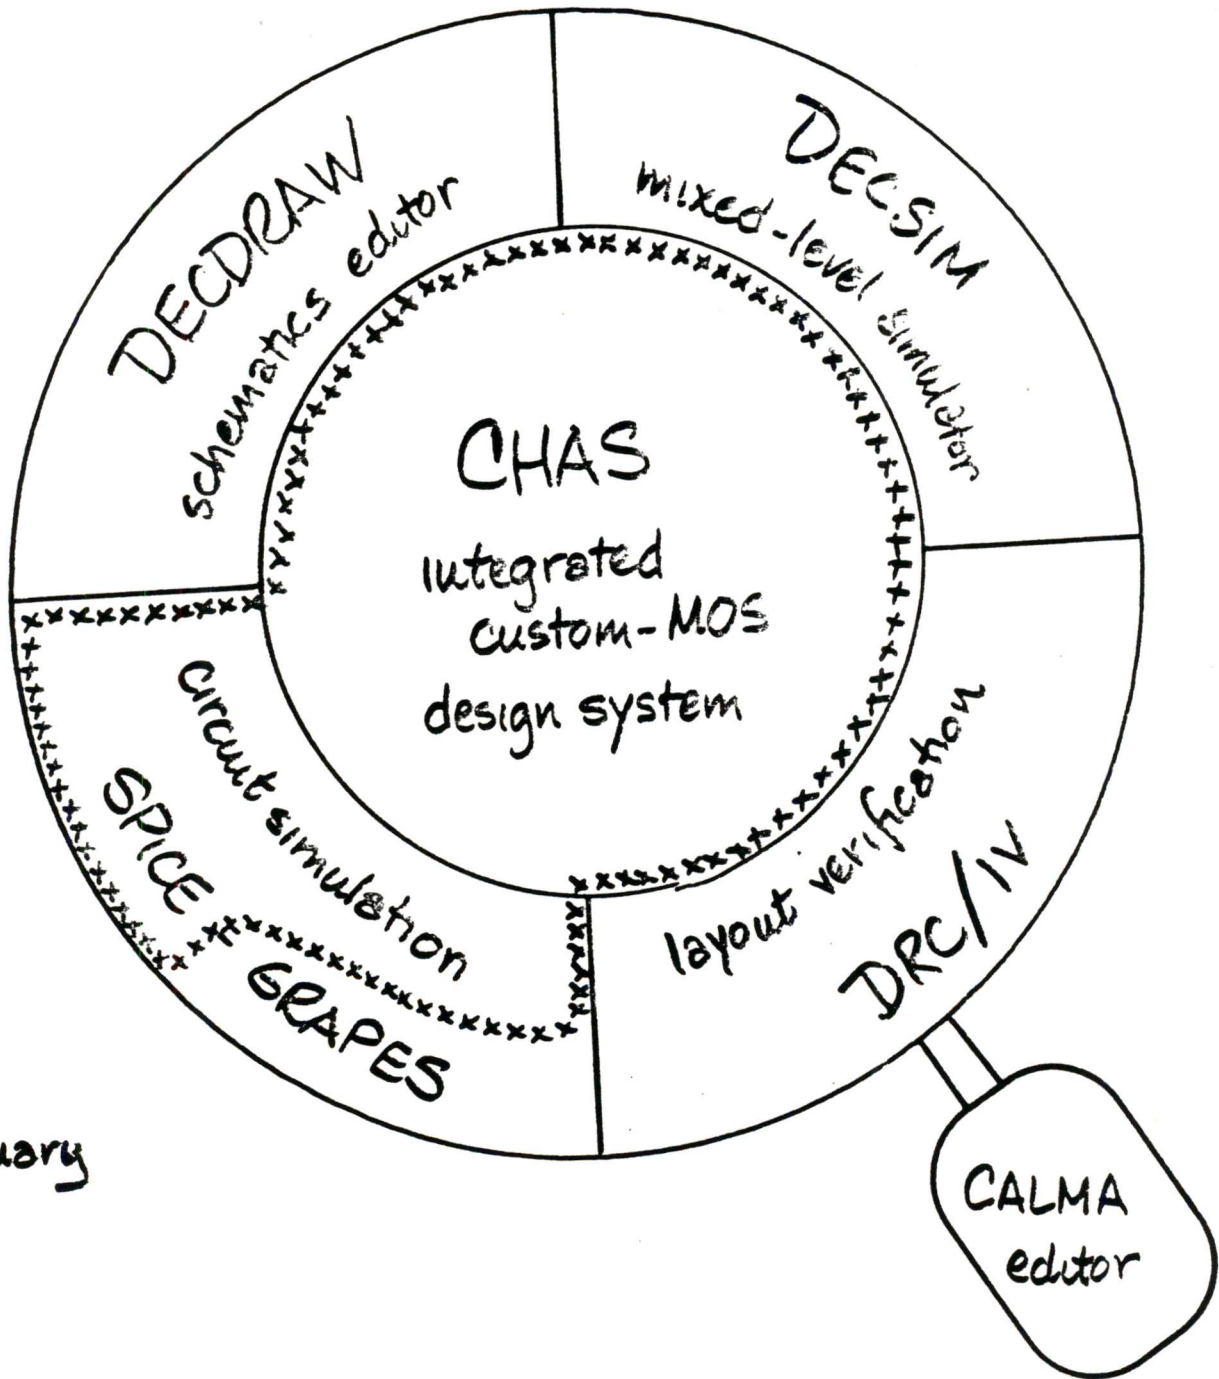
ALL TOOLS RUN ON VAX/VMS.

ALL TOOLS HAVE RESIDUAL VALUE TO THE CORPORATION.

MAJOR CAD TOOLS FOR V-11

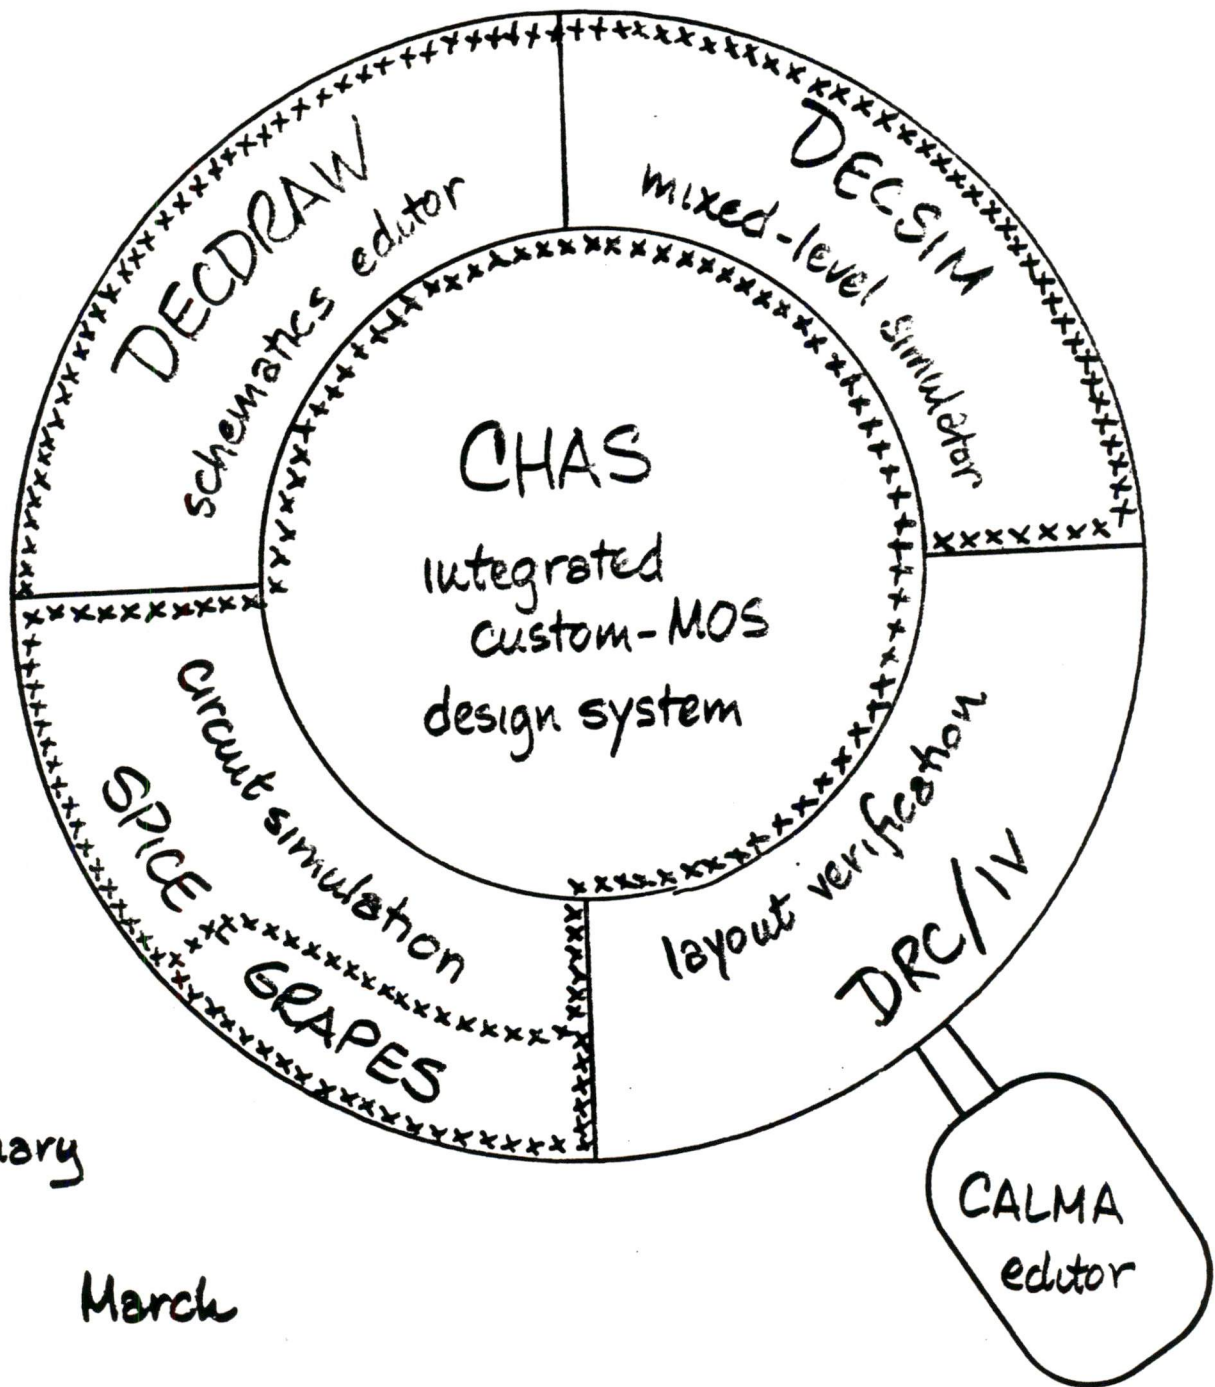


MAJOR CAD TOOLS FOR V-11

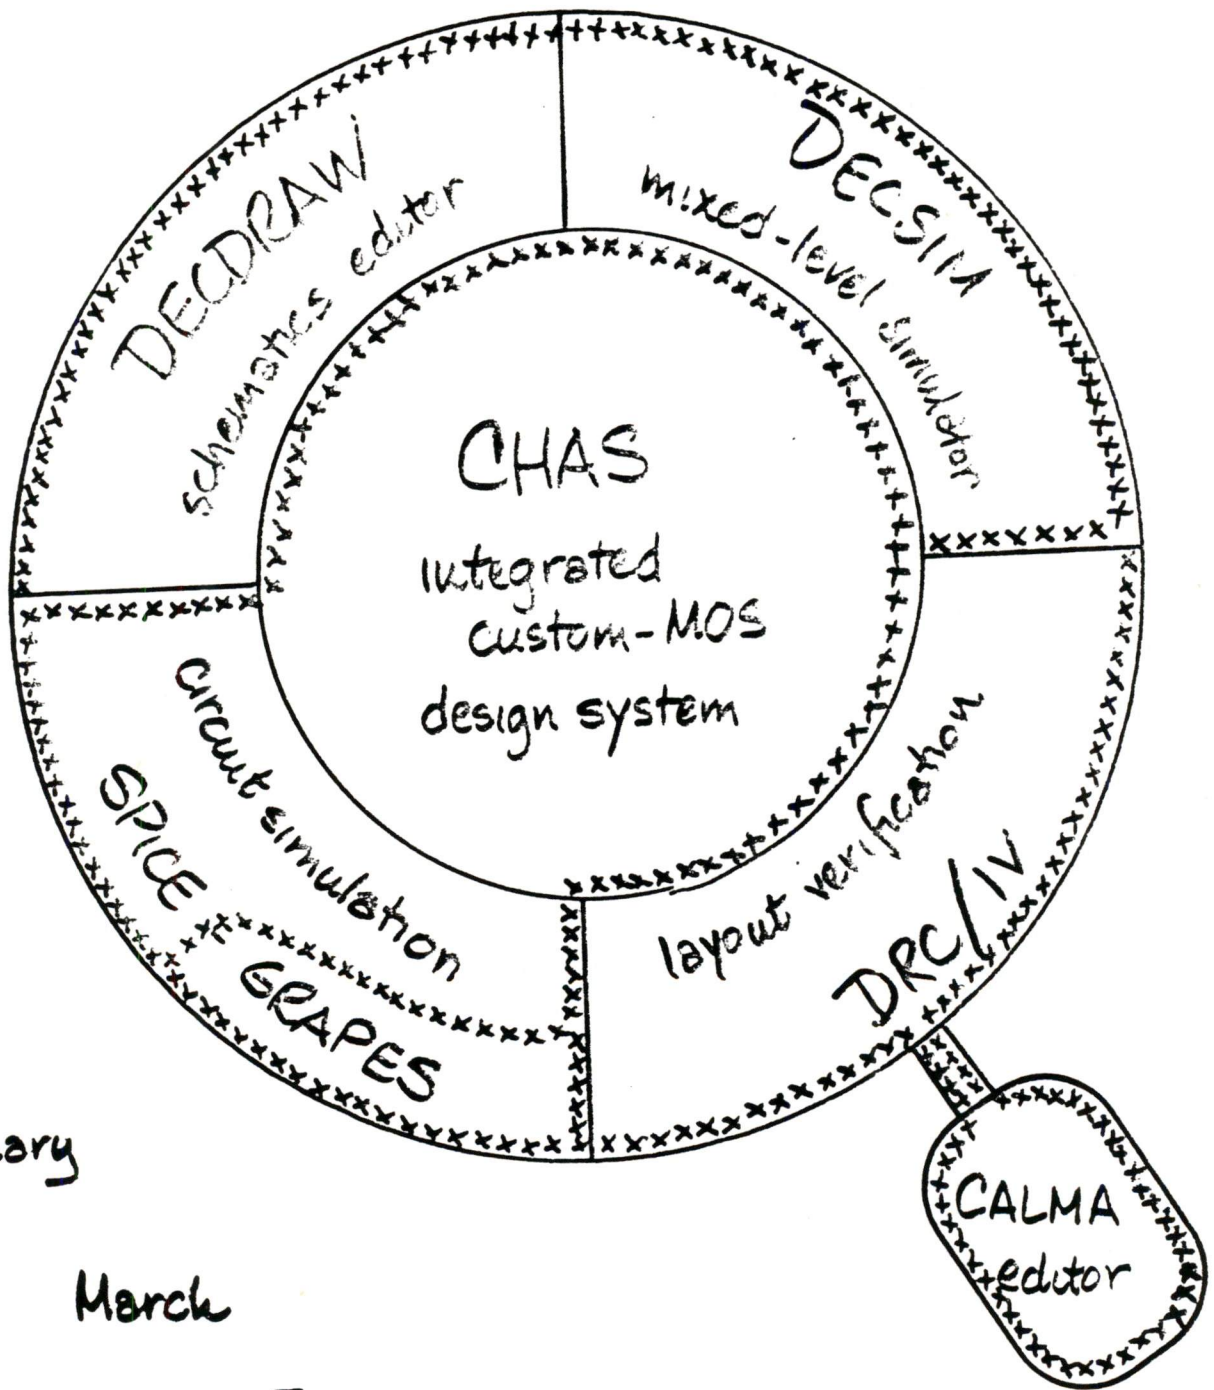


January

MAJOR CAD TOOLS FOR V-11



MAJOR CAD TOOLS FOR V-11



January

March

June

CHAS HIGHLIGHTS

TWO VERSIONS RELEASED TO DATE, FEATURING:

- MENU-DRIVEN USER INTERFACE
- HIERARCHICAL DATA BASE
- CIRCUIT SCHEMATIC EDITING ON A VT100
- SPICE WIRELISTER AND SPICE SUPPORT
- NON-GRAPHICAL FLOOR PLANNING

VERSION 3 (IN MARCH) ADDS:

- DECDRAW SCHEMATIC EDITING
- DECSIM AND SPICE WIRELISTER
- DECSIM SUPPORT
- SPICE POSTPROCESSOR SUPPORT

VERSION 4 (JUNE) WILL ADD:

- INTERACTIVE GRAPHICAL FLOOR PLANNING
AND BLOCK ASSEMBLY
- CALMA SUPPORT
- LAYOUT VERIFICATION TOOLS

CHAS WORKSTATIONS

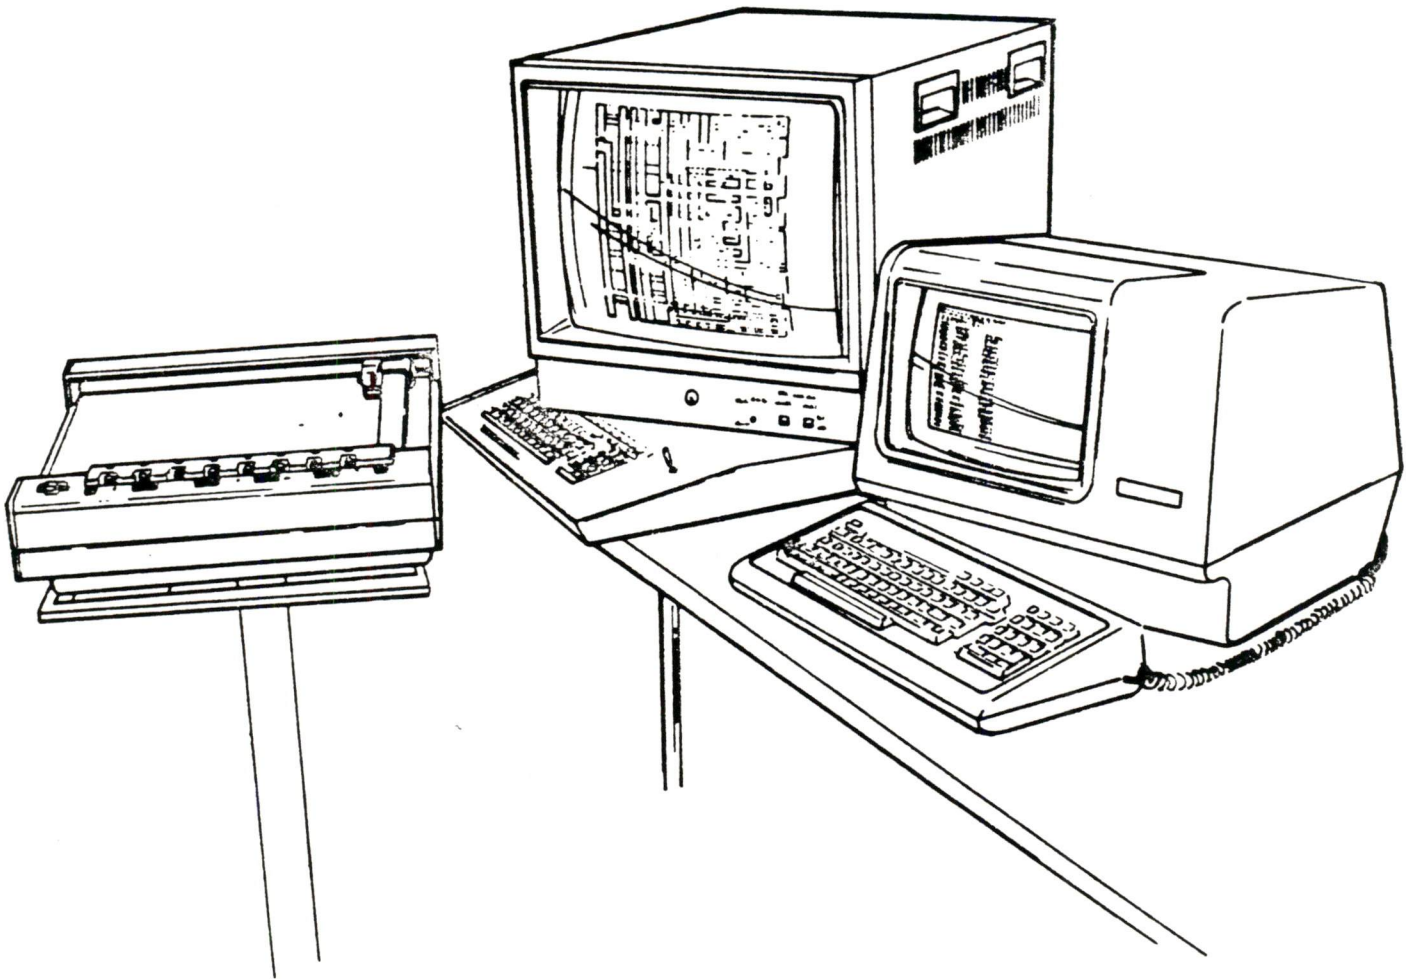


Figure B-1 CHAS Workstation

DECSIM HIGHLIGHTS

TONY HUTCHINGS IS OUR NEW DECSIM MANAGER.

TUMS AND SAGE2 DEVELOPERS HAVE BEEN TRANSFERRED TO DECSIM TEAM.

2 MILESTONES MET SINCE AUGUST:

BEHAVIORAL MODELING RELEASE IN NOVEMBER

15 DECSIM USERS IN HUDSON

V-11 TEAM HAS WRITTEN 14,000 LINES OF BEHAVIORAL MODELS

ISPS TO DECSIM TRANSLATOR AVAILABLE


REVISION OF GATE-LEVEL WIRELIST COMPILER COMPLETE IN JANUARY

NEXT RELEASES:

TIMING AND MICROCODE LOADING IN EARLY FEBRUARY

NEW NETPRO EXTERNAL RELEASE IN EARLY FEBRUARY

CONNECTABLE MODELS IN APRIL



WORKING ON A SCHEDULE FOR ADVANCED MOS MODELS

How SEG/CAD TOOLS SUPPORT THE V-11 CHIP TEAMS

CHAS PROVIDES DATA SECURITY, DESIGN EFFICIENCY, PROOF OF CORRECTNESS:

CENTRALIZED HIERARCHICAL DATA BASE

DATA INTEGRITY

DATA CONTROL

HIDING OF NASTY CAD TOOL INTERFACES

SCHEMATICS TO SPICE OR DECSIM OR IV

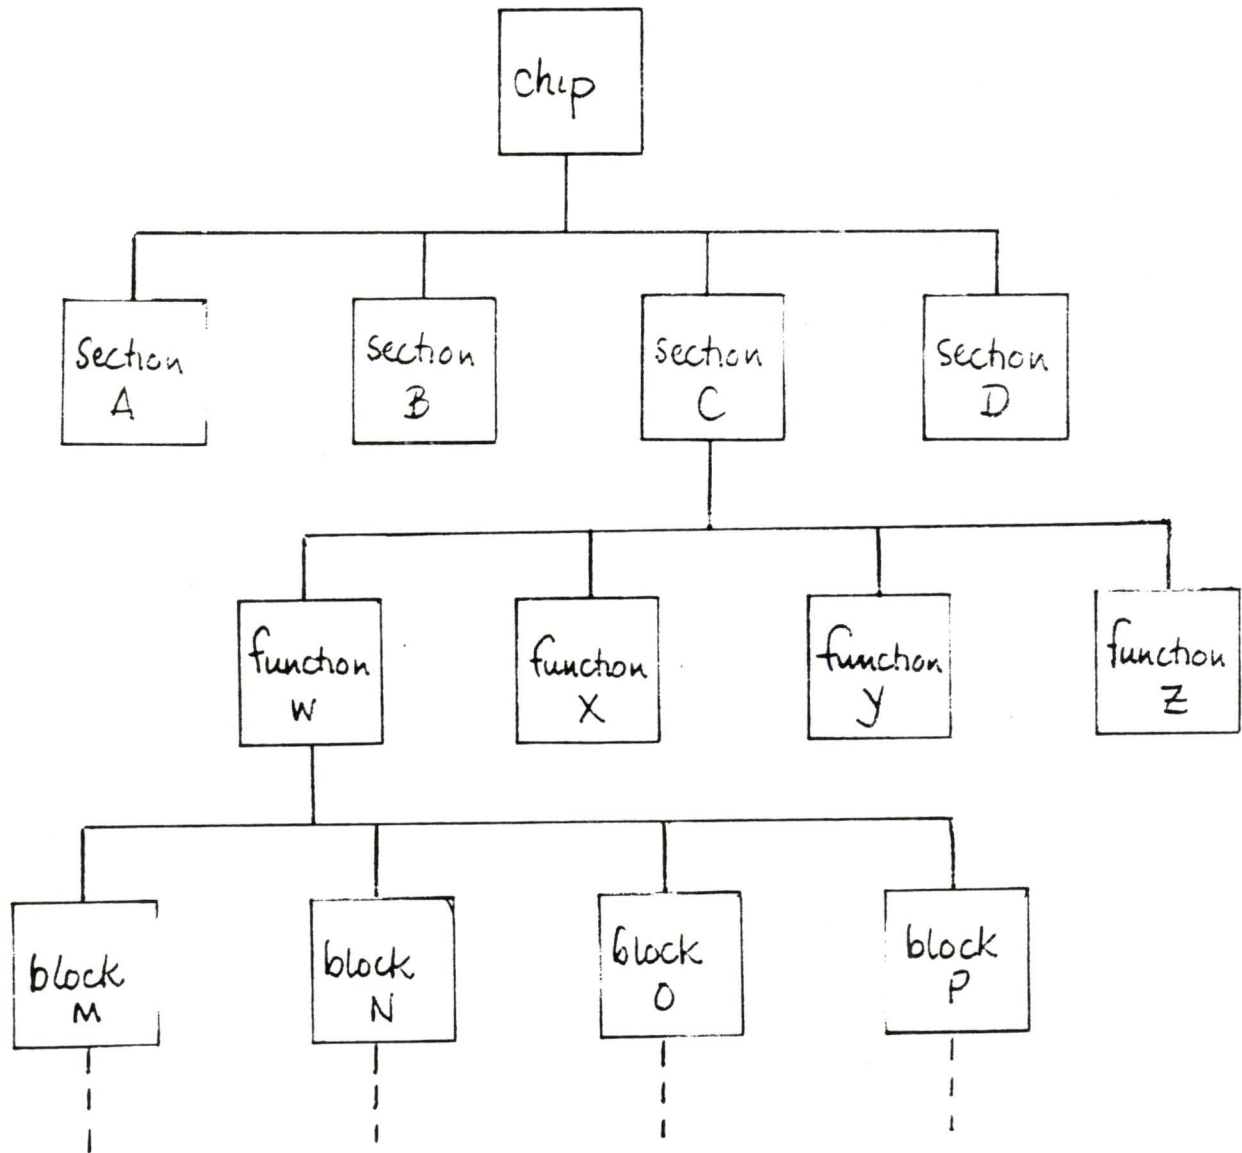
CALMA LAYOUT TO DRC AND IV

COMPUTER-AIDED CHIP VERIFICATION

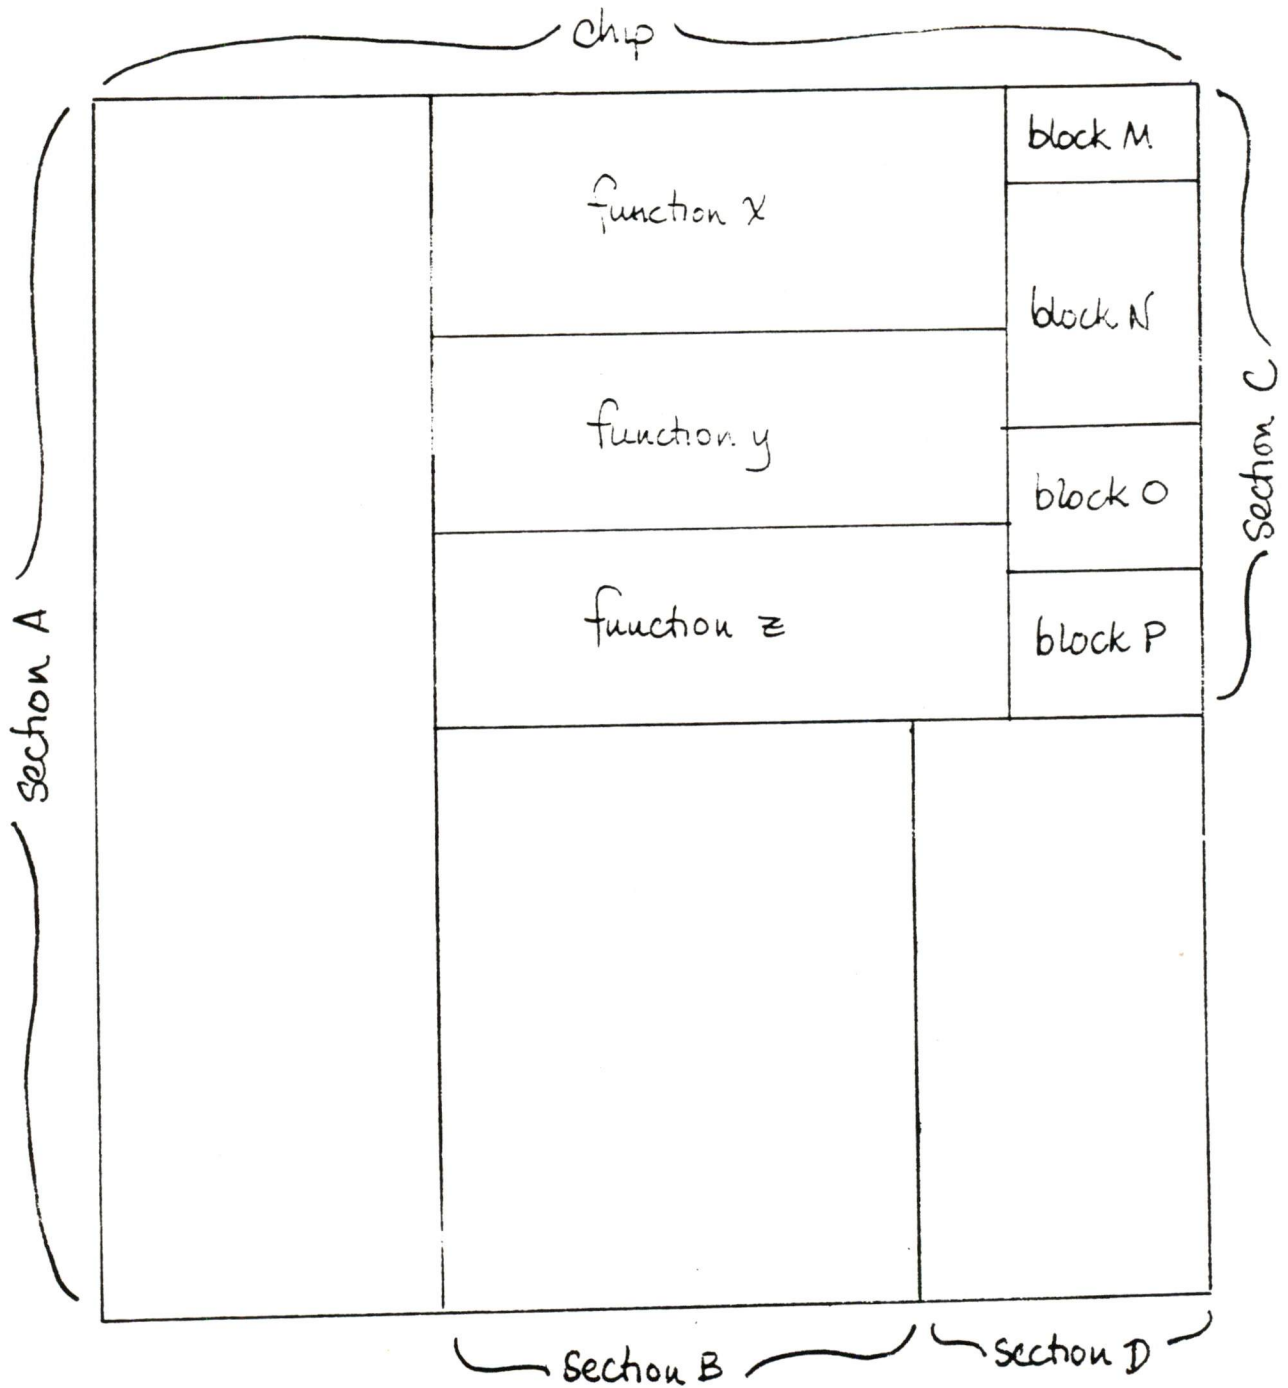
BEHAVIORAL MODELS = RTL MODELS = GATE MODELS

GATE MODELS = MASK LAYOUT

Chip Hierarchy Viewed as a Tree



Floor Plan of a Chip



How SEG/CAD TOOLS SUPPORT THE V-11 CHIP TEAMS

DECSIM PROVIDES PROOF OF LOGIC DESIGN:

BEHAVIORAL/FUNCTIONAL MODELING OF

INDIVIDUAL CHIPS AND CHIP SECTIONS

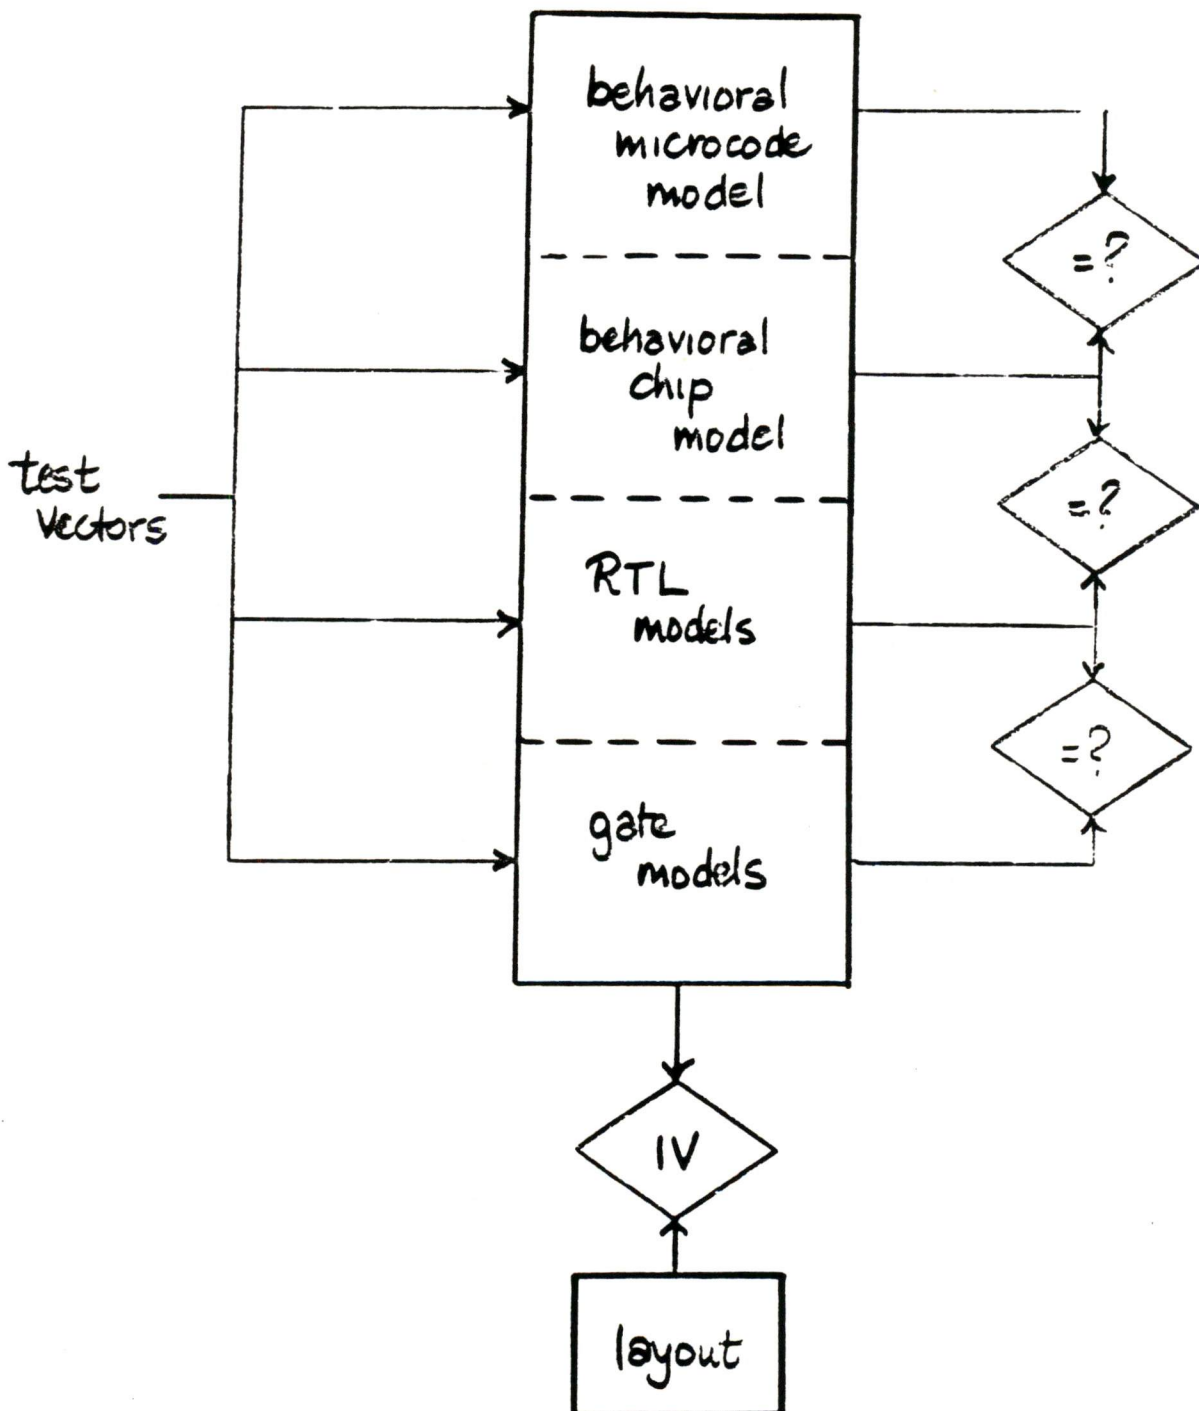
CHIP SECTIONS RUNNING TOGETHER AS A CHIP

THE CHIPS RUNNING AS A VAX SYSTEM

RTL MODELING OF CHIP SECTIONS AND CHIP LOGIC FUNCTIONS

GATE MODELING OF CHIP BLOCKS

How To Make V-11 Chips Work on Pass 1



V-11 MODELING STRATEGY

PROVE THAT MULTI-LEVEL MODELS ARE IDENTICAL

BEHAVIORAL MODEL = RTL MODEL = GATE MODEL

TYPES OF SIMULATION

BEHAVIORAL MODELING OF CHIP SET FOR MICROCODE DEBUGGING

BEHAVIORAL MODELING OF CHIP SET HARDWARE

MORE DETAILED BEHAVIORAL AND RTL MODELING OF CHIP SECTIONS

RTL MODELING OF CHIP FUNCTIONS

GATE MODELING OF CHIP BLOCKS

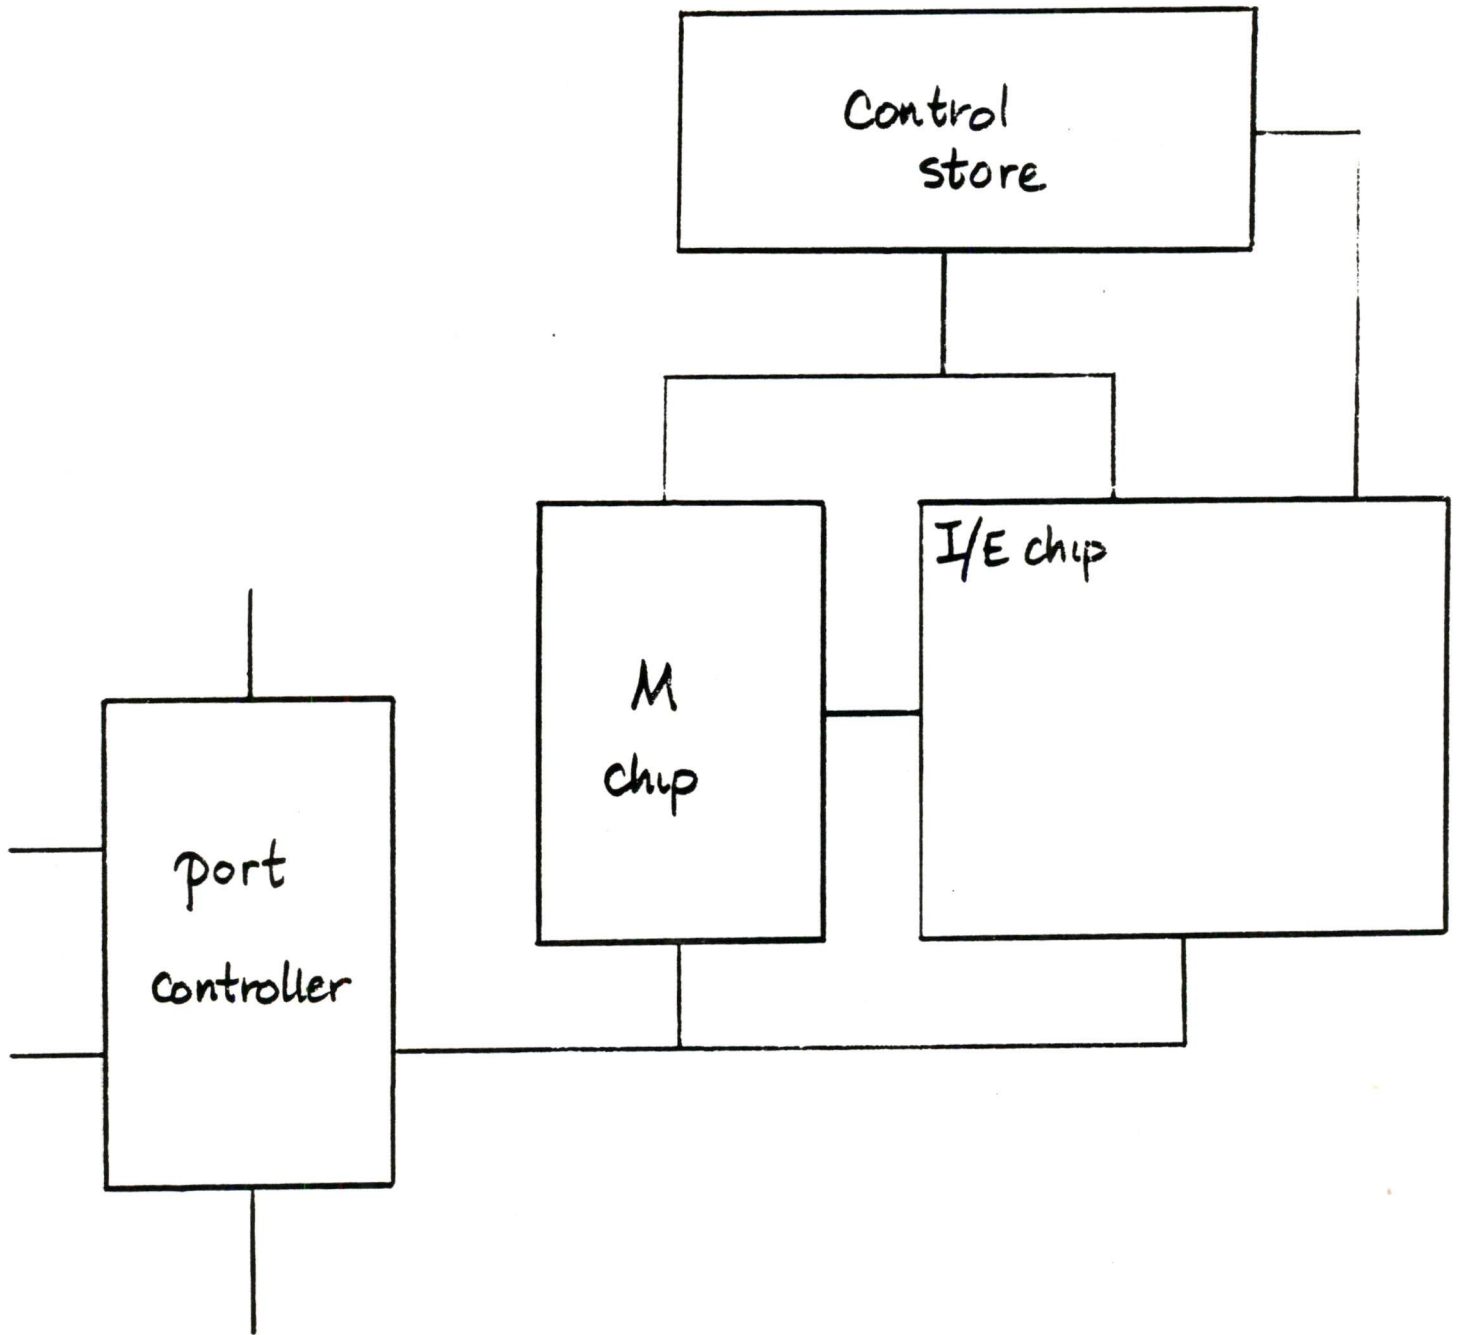
MIXED-LEVEL MODELING

REPLACE A CHIP MODEL WITH ITS EQUIVALENT SECTION MODELS

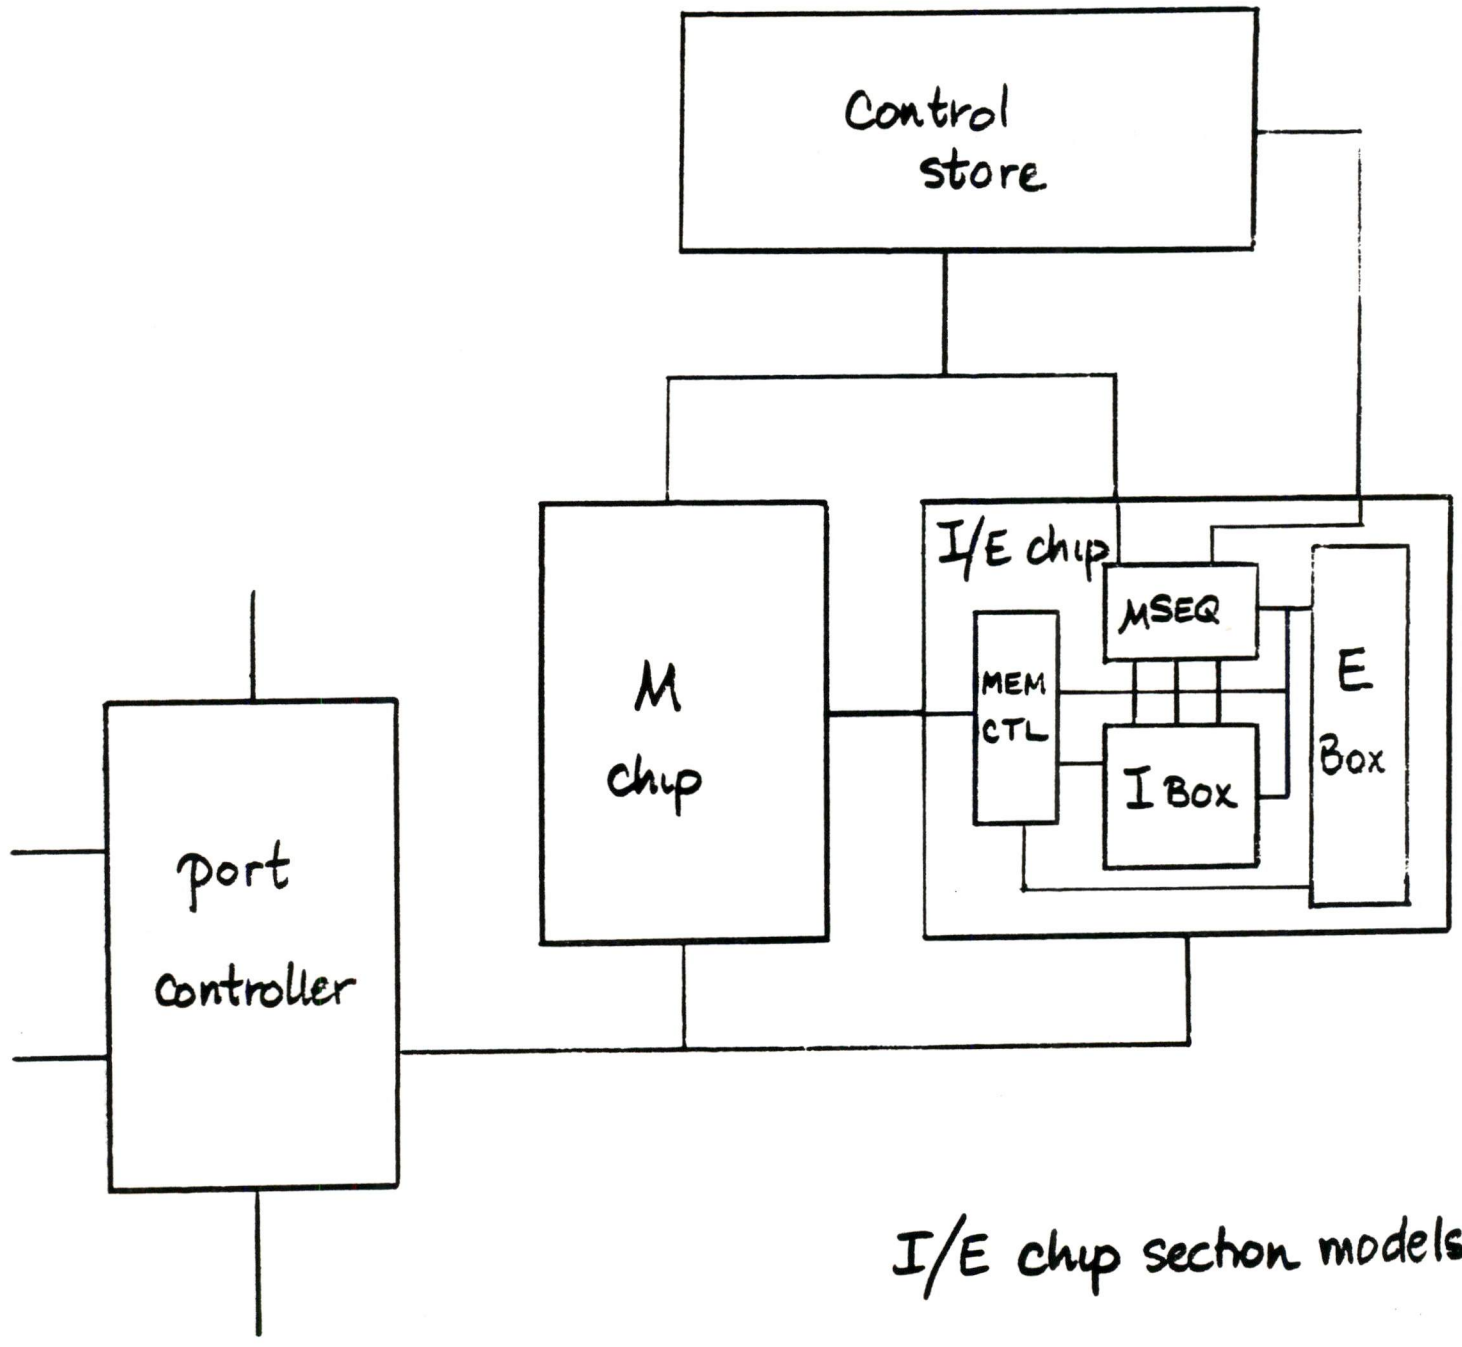
REPLACE A CHIP SECTION WITH ITS EQUIVALENT FUNCTION MODELS

REPLACE A CHIP FUNCTION WITH ITS EQUIVALENT GATE MODELS

V-11 System Model

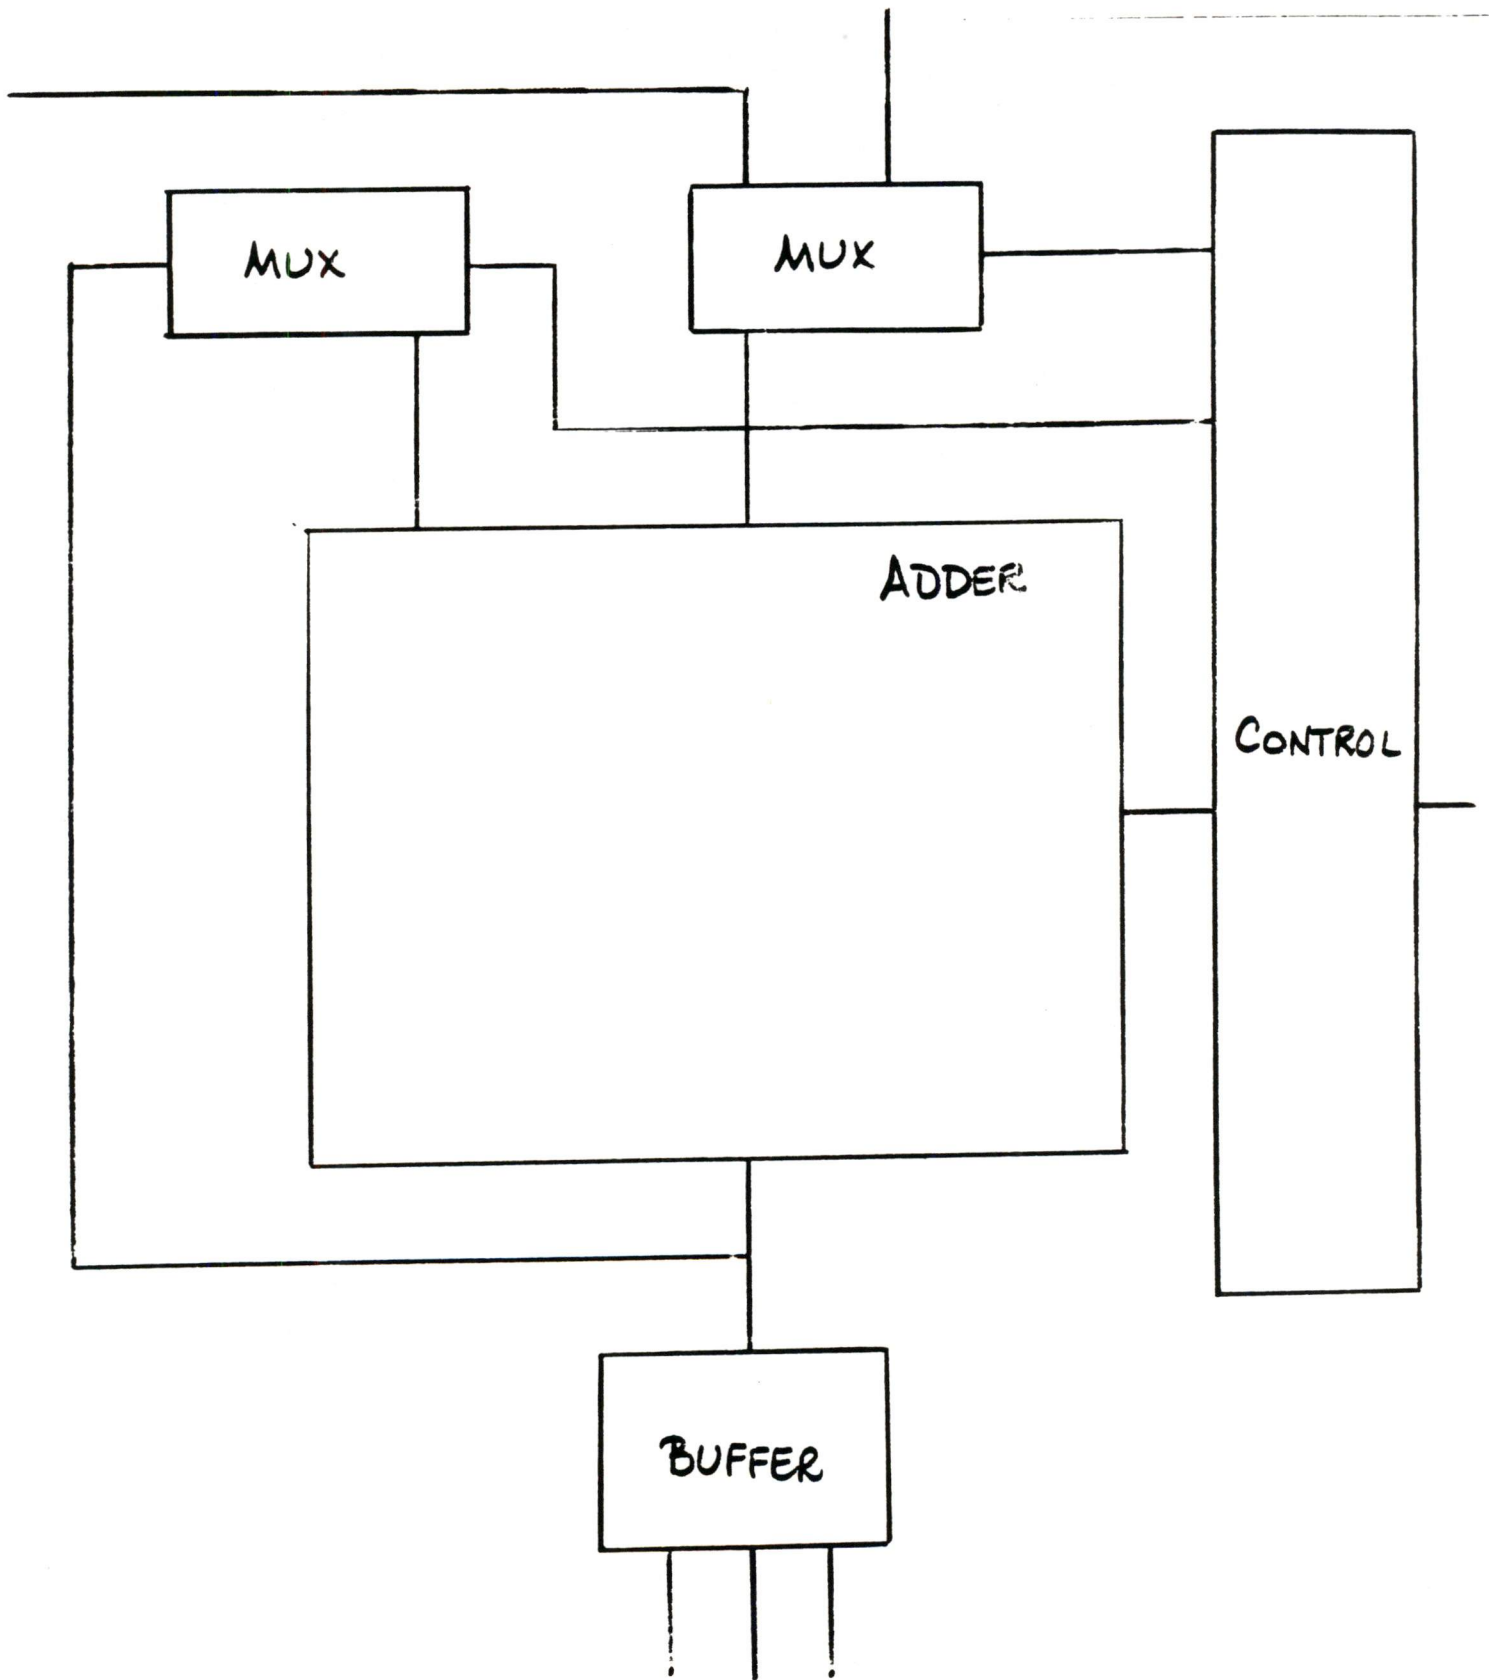


V-11 System Model

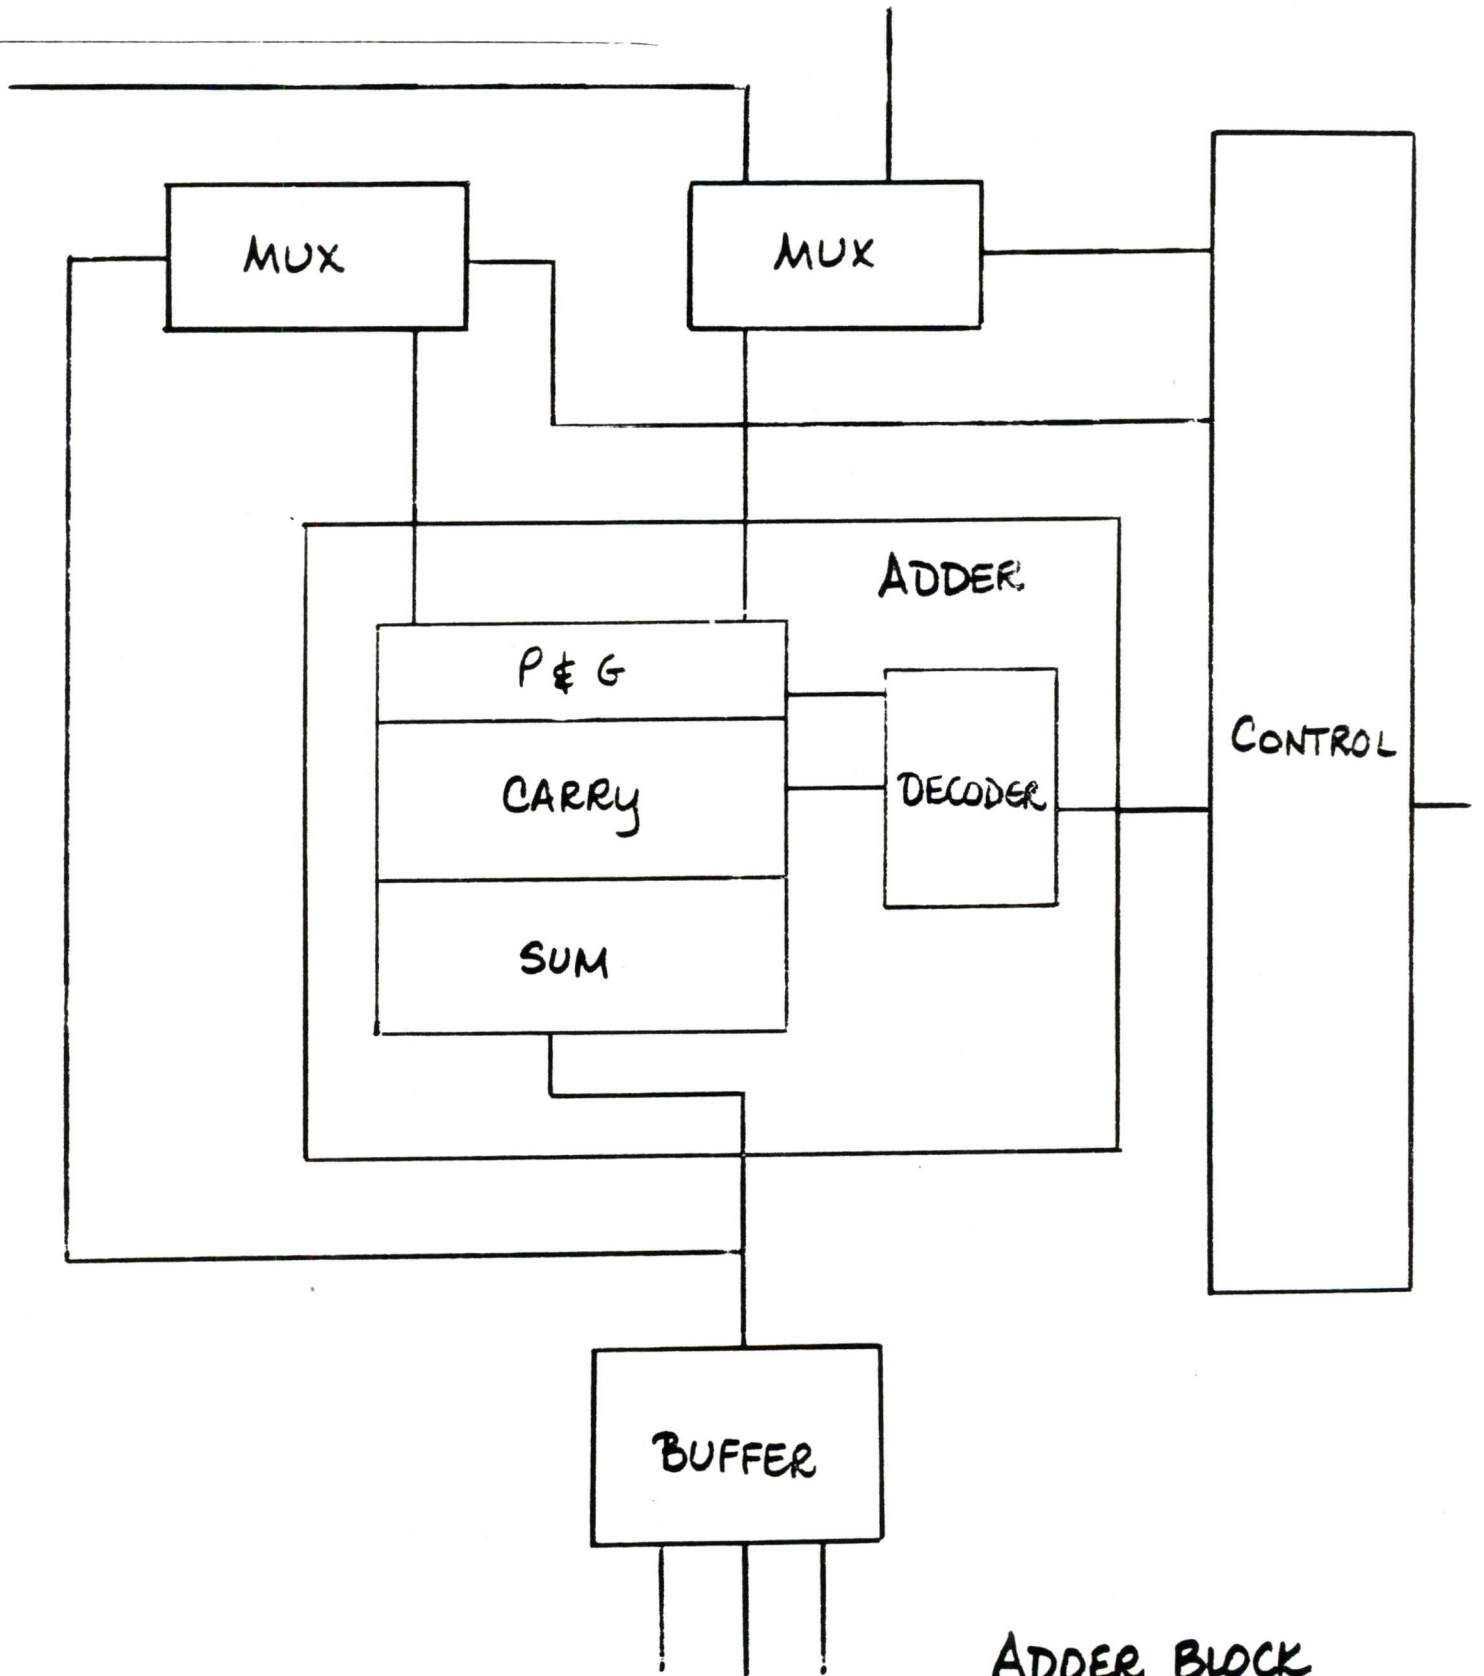


I/E chip section models

MICROSEQUENCER FUNCTION MODELS



MICROSEQUENCER FUNCTION MODELS



ADDER BLOCK
MODELS

WHAT MIXED MODELING BUYS THE V-11 TEAM

BETTER SIMULATION RATIOS

TEST PATTERN PROOF THAT MODELS AT DIFFERENT LEVELS ARE EQUIVALENT

SYSTEM SIMULATIONS PROVE THAT THE TEAMS ARE WORKING TO SAME SPECS

KEEP PROVING TO OURSELVES THAT WE'RE BUILDING A VAX

KEEP ALL TEAMS WORKING SIMULTANEOUSLY

DECSIM/TUMS COMPARISON

POINT OF COMPARISON: 170 LINE PDP8E BEHAVIORAL MODEL

COMPILE/LINK TIMES

<u>MEASURES</u>	<u>TUMS</u>	<u>DECSIM</u>
VAX CPU MIN/1000 LINES	11.7	6.7
PERFORMANCE (TUMS = 1)	1	1.7

EXECUTION TIMES (NO TRACING/EXAMINING/WATCHING DURING SIMULATION)

<u>MEASURES</u>	<u>TUMS v4.1</u>	<u>TUMS v5.0</u>	<u>DECSIM v1C</u>
VAX CPU TIME	4:29:40	3:00:06	0:11:59
PERFORMANCE (TUMS v4.1 = 1)	1	1.5	22
SIMULATION RATIO	2600:1	1700:1	110:1

EXECUTION TIMES (3 WATCH POINTS)

<u>MEASURES</u>	<u>TUMS v4.1</u>	<u>TUMS v5.0</u>	<u>DECSIM v1C</u>
VAX CPU TIME	0:12:45	0:10:38	0:03:48
PERFORMANCE (TUMS v4.1 = 1)	1	1.2	3.4

DECSIM/SAGE2 COMPARISON

POINTS OF COMPARISON: 3 NETWORKS RANGING FROM 1200 TO 12000 GATES

COMPILATION TIMES

<u>GATE COUNT</u>	<u>SAGE2 CPU SEC. ON KL10</u>	<u>SAGE2 GATES/ CPU SEC.</u>	<u>DECSIM CPU SEC. ON VAX</u>	<u>DECSIM GATES/ CPU SEC.</u>
1177	113	10.4	248	4.8
5885	480+	12.3+	293	20.1
11770	-	-	354	33.3

COMPILATION PERFORMANCE OF 5885 GATE NETWORK

	<u>SAGE2 ON KL10</u>	<u>DECSIM ON VAX</u>
PERFORMANCE (SAGE2 = 1)	1	1.6

EXECUTION TIMES OF 1177 GATE NETWORK

<u>MEASURES</u>	<u>SAGE2 ON KL10</u>	<u>DECSIM ON VAX</u>
EVENTS/CPU SECOND	3333	5000
PERFORMANCE (SAGE2 = 1)	1	1.5

How SEG/CAD TOOLS SUPPORT THE V-11 CHIP TEAMS

DECDRAW PROVIDES DESIGN EFFICIENCY:

EASY-TO-LEARN AND EASY-TO-USE SCHEMATIC DRAWING TOOL
LOGIC AND TRANSISTOR COMPONENTS
SUPPORT OF THE CHIP HIERARCHY WITHIN DRAWINGS
AUTOMATIC WIRELISTING TO SPICE OR DECSIM

SPICE AND SPICE MODELS PROVIDE PROOF OF SPEED/POWER ASSUMPTIONS:

CORRELATION BETWEEN PROCESS AND SPICE SIMULATOR
SPEED, POWER ESTIMATES OF CIRCUITS
GRAPES OFFERS INTERACTIVE GRAPHICS FOR WAVEFORM ANALYSIS

How SEG/CAD TOOLS SUPPORT THE V-11 CHIP TEAMS

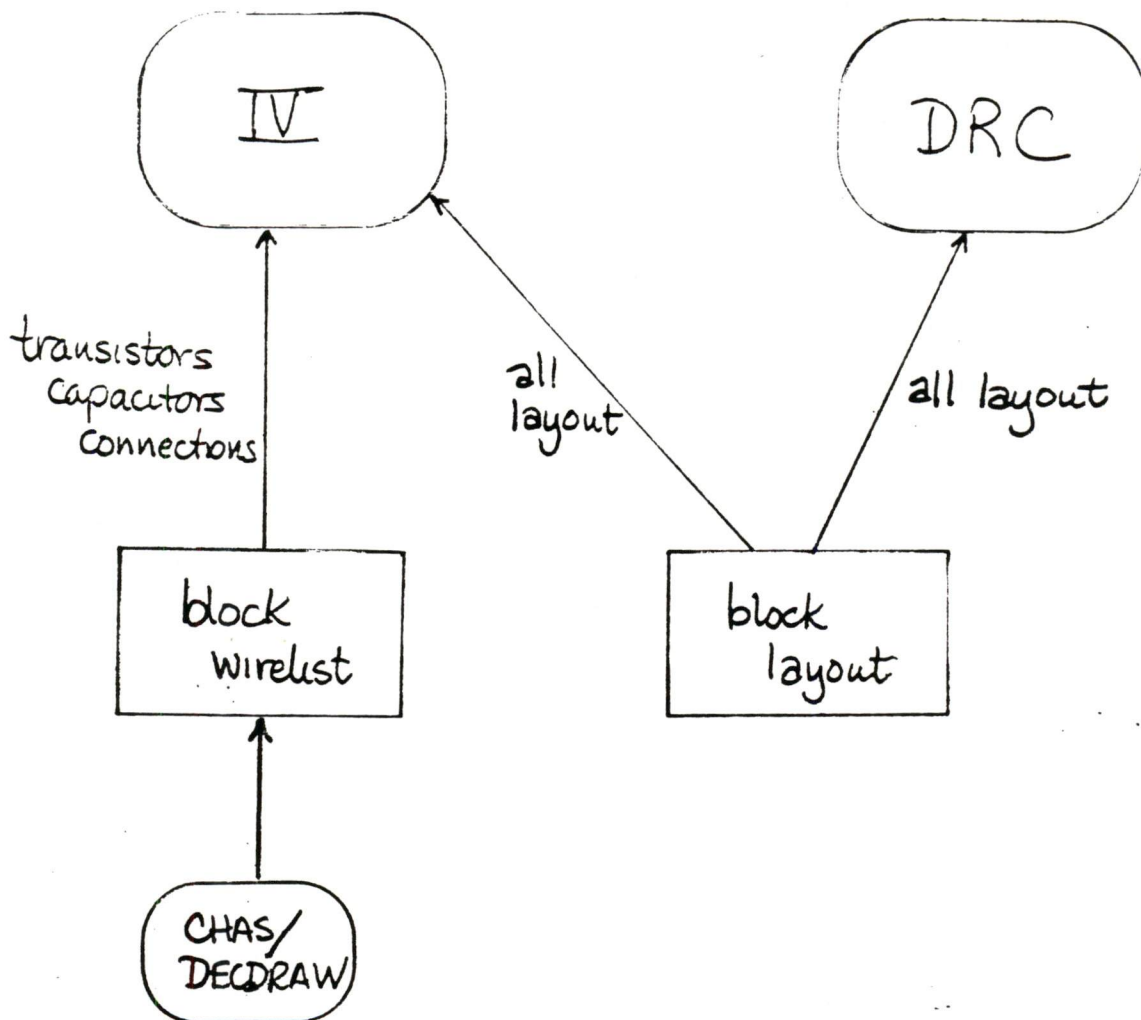
DRC AND IV PROVIDE PROOF THAT CHIP WILL WORK WHEN FABRICATED:

PROOF OF LAYOUT'S ADHERENCE TO PROCESS LAYOUT RULES

PROOF OF GATE TO LAYOUT EQUIVALENCE

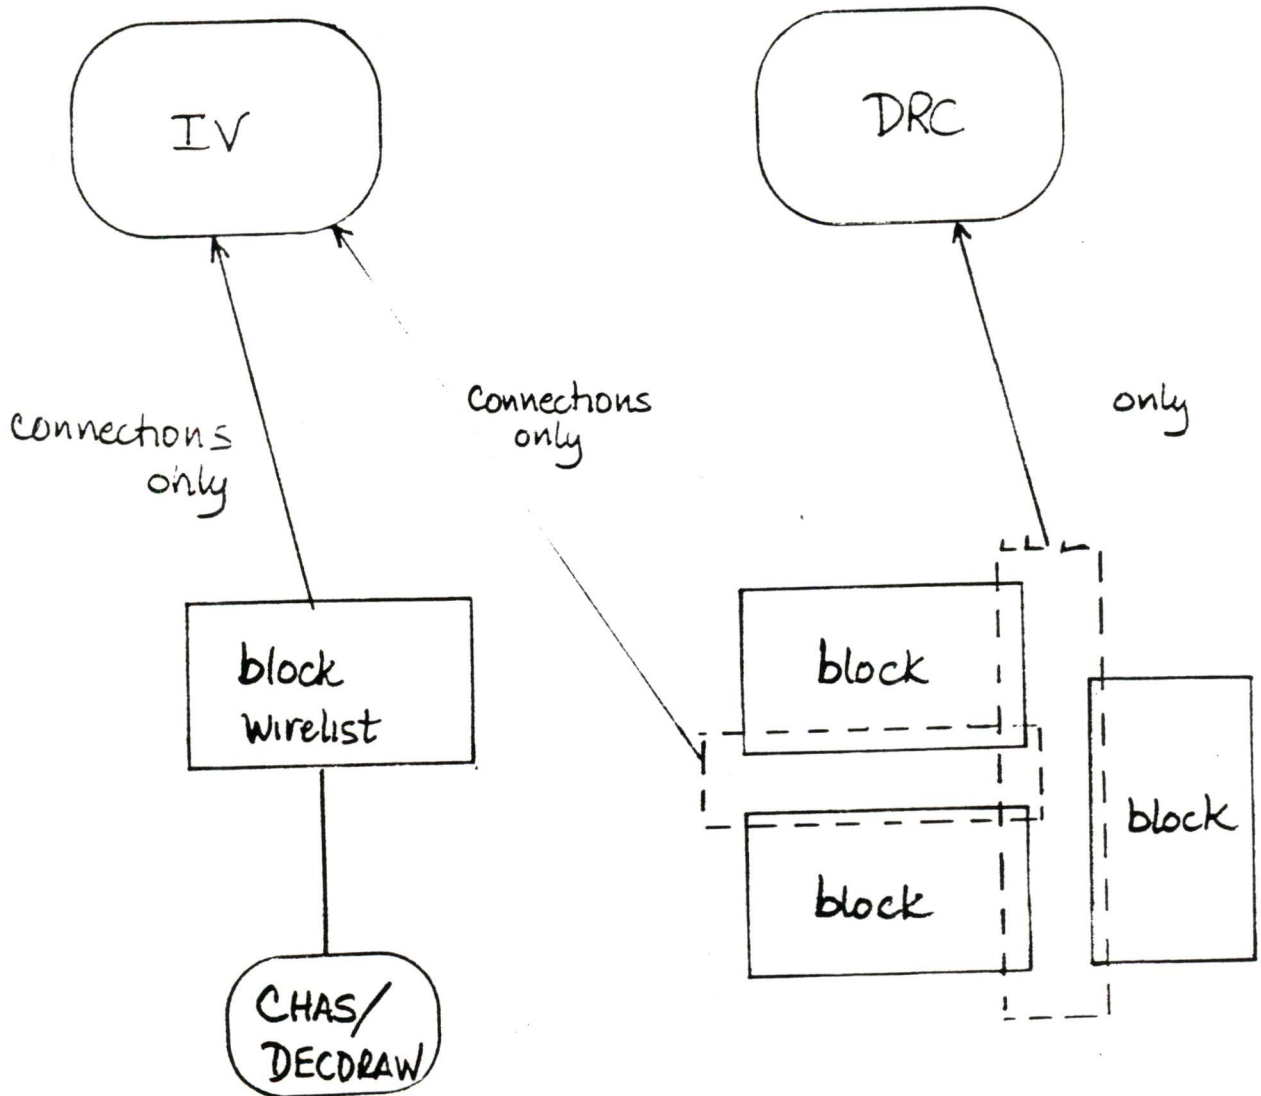
FEEDBACK OF ACTUAL PARASITICS INTO SIMULATION

Block-level DRC/IV



IV = node extractor
parasitic extractor
wirelist comparator

Hierarchical DRC/IV



V-11 CAD Risks

WHAT IF THE DECSIM SCHEDULE SLIPS?

THE SCHEDULE IS AGGRESSIVE

DECSIM TEAM HAS A GOOD TRACK RECORD FOR THE PAST FOUR MONTHS

DECSIM TEAM HAS STRONG MANAGEMENT AND IS A QUALITY TEAM

V-11 COULD USE PL/1 AS A BACKUP FOR DECSIM BEHAVIORAL MODELING

GATE MODELING EXISTS

ADVANCED MOS MODELING IS THE GREATEST "WHEN?" RIGHT NOW

2 PEOPLE WORKING ACTIVELY TOWARD A SOLUTION

PLANS, SCHEDULE ARE BEING DEVELOPED NOW

WHAT IF DECSIM SIMULATION RATIOS ARE TOO HIGH?

BENCH STRENGTH -- WE COULD DIVERT WITEK/PETERS FOR SPEED UP

MIXED-LEVEL MODELING WILL KEEP THE RATIOS DOWN

WHAT IF CHAS DOESN'T COME THROUGH FOR DESIGN VERIFICATION?

CHAS WILL COME THROUGH

EXCELLENT TEAM

BACKUP STRENGTH IN CAD A/D IN SEG

WE CAN FALL BACK TO OLDER TOOLS

VERIFY CHIPS IN SECTIONS

MANUAL CHECKING OF SECTION INTERFACES