INTEL ACCELERATOR

GROUNDRULES/ASSUMPTIONS

1. ROM Cost Estimate
2. Commercial temperature ranges.
3. Screened parts not required.
4. Design for production of 500 units.
5. No burn in of modules.
6. ~~Six~~ 4 layer PWB for production.
7. Standard Eurocard configuration (9" high 10" deep)
   (See assembly sketch)
8. Build 140 units (deliver 128; 12 are spares inventory or demo units).
9. Build two prototype units using multiwire.
10. Minimum size tower will be "loaned" to MSD from INTEL for system integration .
11. Units are warranted for 90 days; assume 0.05 failures per unit during warranty period.
12. Unit acceptance test (FAT) will be done on Teradyne.
13. Weekly meetings will be bid under A05, not in individual WAs.
14. Assume 1 May 1988 program start.
15. Assume three ECOs during 140 unit build:
    2 component changes, 1 board wiring change.
16. This is a commercial program -- parts, environment, documentation, practices, quality.
17. Assume multiwire board with thru-hole components. Board vendor NRE will provide boards for two proto units.
18. The board is a single-sided 9" x 10" Eurocard.

The iPSC-VX Vector Concurrent Supercomputer

The iPSC™-VX (vector extension) is a vector concurrent computer system, an enhanced member of the iPSC family of "personal supercomputers." The iPSC-VX offers true supercomputer performance for the individual scientist or research team, and provides the tools for developing efficient, high-performance applications in a large-scale parallel computing environment.

The cost of owning and maintaining a conventional supercomputer has made high-end scientific computing inaccessible to all but a privileged few. The iPSC-VX changes this. By using low cost VLSI to harness the complementary technologies of concurrent and vector processing, the iPSC-VX sets a new price-performance standard. The result is supercomputer performance at the price of a supermini.

The iPSC-VX family builds upon the basic architecture of the iPSC concurrent computer. By coupling a high-performance vector processor to each of the iPSC processing nodes, the vector enhancement results in dramatically improved computational performance for both vector and scalar operations. Optimized to meet the mathematical requirements of scientific computation, the iPSC-VX family is ideal for such applications as electronic device simulation, computational chemistry, fluid dynamics, and oil reservoir modeling.

An iPSC-VX system (Figure 1) consists of 16, 32, or 64 computational nodes. Each node consists of an independent microcomputer with its own CPU, communications control, local memory, and dedicated vector processor. In keeping with the basic iPSC architecture, the processing nodes in a system are interconnected using a hypercube topology where connected nodes are supported with reliable point-to-point message delivery service. The multiple nodes which comprise a system are collectively known as the "Cube" and may be housed in one, two, or four computational units. The Cube Manager, an Intel System 310 supermicrocomputer, provides a gateway to the system and serves as a convenient software development station.

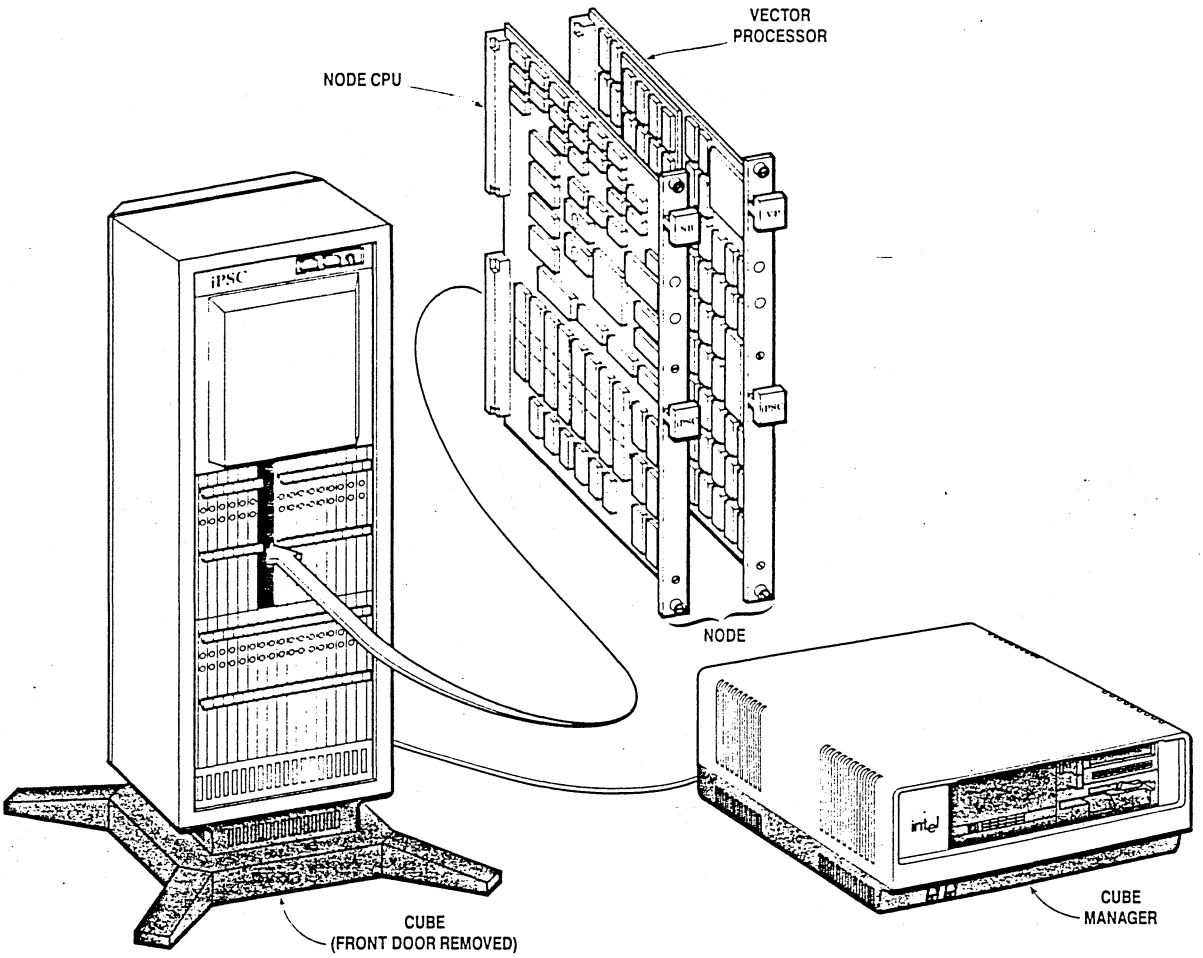

Figure 1 — iPSC-VX System

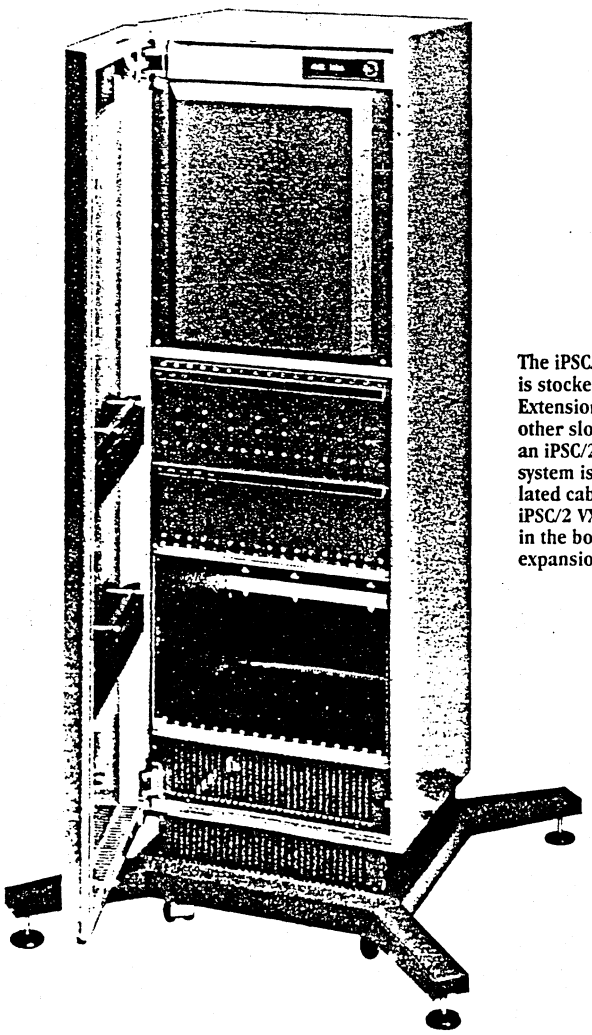

The iPSC/2 system shown is stocked with VX Vector Extension boards in every other slot, each paired with an iPSC/2 node board. The system is a partially-populated cabinet, with only 8 iPSC/2 VX nodes, and room in the bottom half for later expansion.

Highly configurable systems give the application developer the flexibility to meet user needs exactly, and offer the user a hardware investment that is preserved over time. The investment is preserved in two ways, in software that is unchanging as the number of nodes increases, and in hardware that can be modified—instead of replaced—as new technology becomes available.

iPSC/2 systems can be upgraded easily because of careful design for modularity and standardization. Intel's strategy is to apply modularity to elements that change the most rapidly. For the iPSC/2 system, evolving performance and capability is expected of numeric accelerators, memory modules, and the Direct-Connect communication modules.

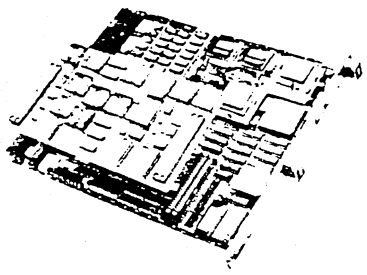

DIRECT-CONNECT MODULE
The Direct-Connect Module on each iPSC/2 node board is mounted on two connectors, and can be easily replaced or updated in the field. Future enhancements in communication technology can provide performance improvement without replacement of entire nodes.

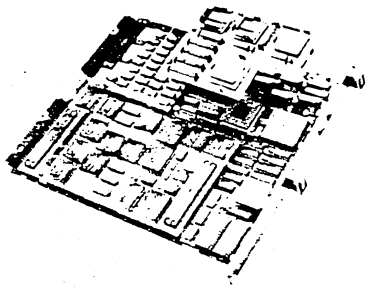

MEMORY MODULE
Memory technology has historically been the area of most rapid development: memory capacity has doubled every two years for the past decade and a half. Today, using surface mount manufacturing techniques and readily available components, it is possible to configure an 8-Mbyte memory in the area of a 3x5 inch card. Within the lifetime of the iPSC/2 system, this capacity can be expected to continue to increase dramatically.

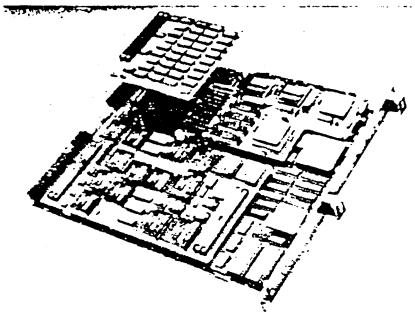

SX MODULE
Arithmetic capabilities are supported in a modular fashion using an 80387 arithmetic coprocessor or the SX Scalar Extension module on each iPSC/2 node. The iLBX-II interconnect is maintained for the VX Vector Extension, and for future enhancements.

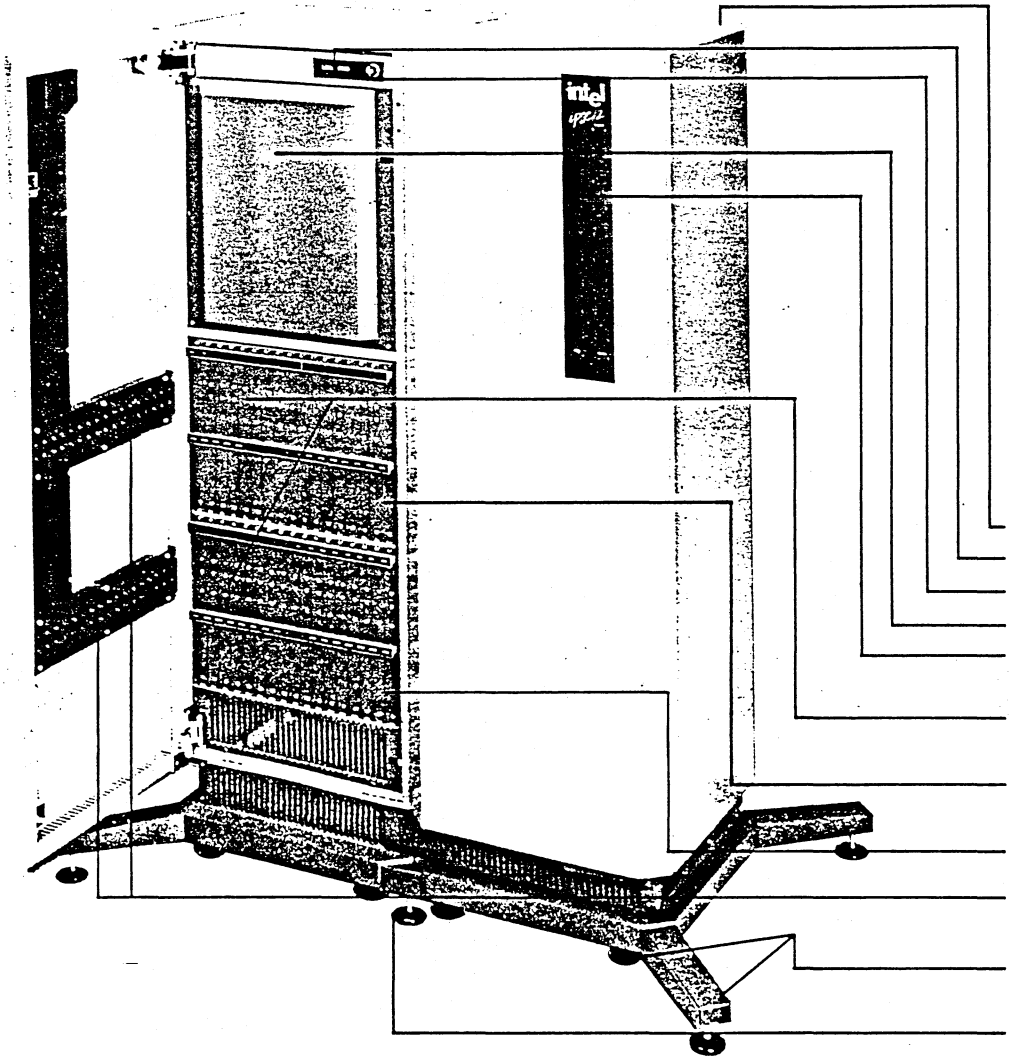

iPSC/2 nodes fit into modular cabinet units; each unit provides 32 slots for node boards or expansion boards. Two additional slots in each cabinet unit contain the Unit Service Module (USM) and a spare node board. Up to four cabinet units can be configured as one system, with up to 128 basic nodes or 64 nodes with expansion boards. The USM provides initialization and diagnostic services for each cabinet unit.

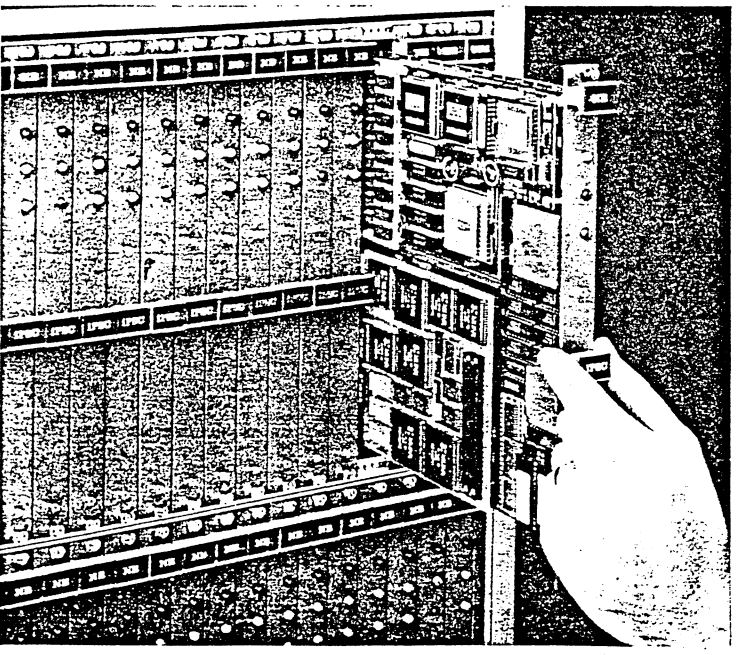

A panel on the front door of each cabinet unit displays the red and green status LED indicators of each board within. Cabinet units are mounted on sturdy casters, and can be moved about a floor easily. The cabinet units are secured in place with stability arms splayed outward at each corner.

The iPSC/2 system has only ordinary power and cooling requirements, needing 220 volts AC phased power at approximately 20 amperes per unit. European specification systems are available.

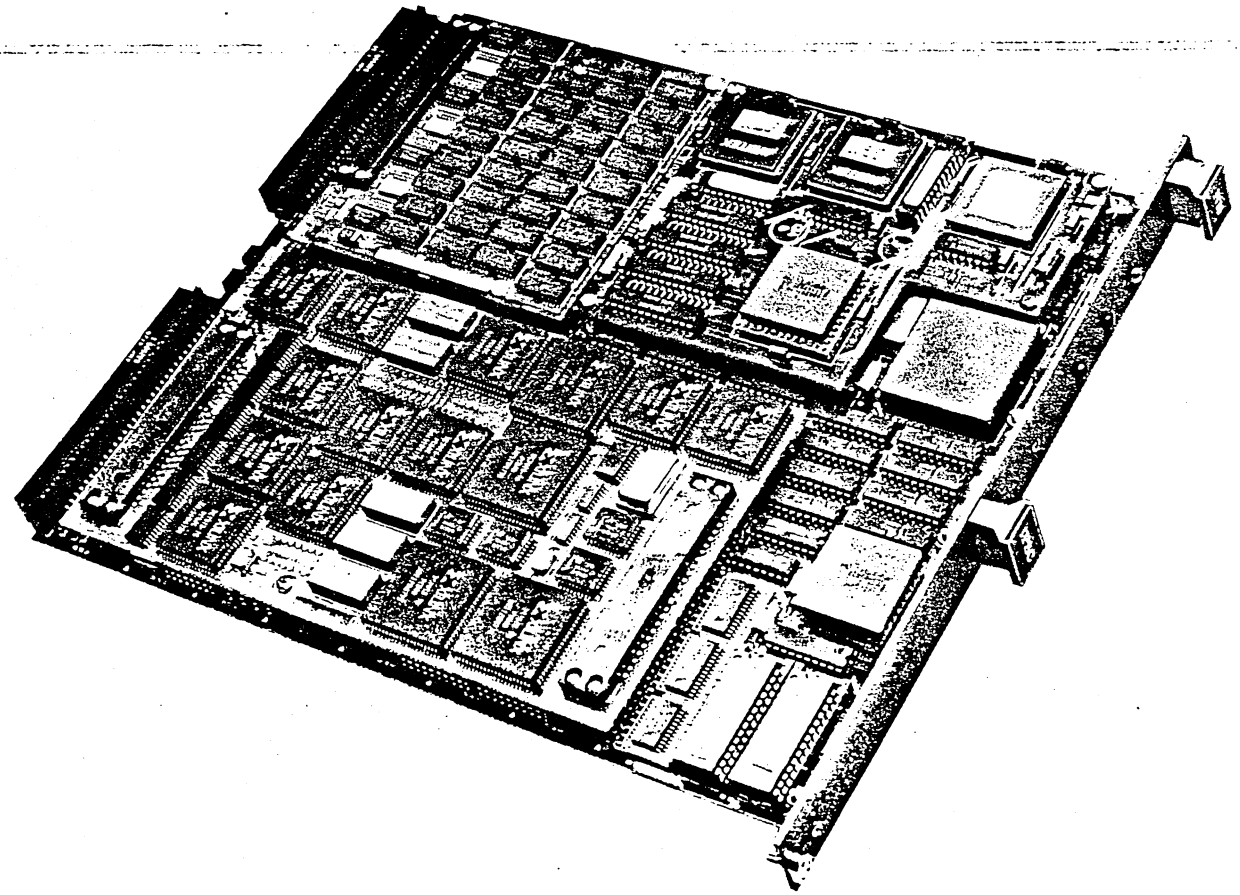

NODE PROCESSOR The heart of an iPSC/2 node is the powerful Intel 80386 32-bit microprocessor. Used only in its native 32-bit mode, the 80386 executes FORTRAN, C, and LISP programs at up to 4 MIPS and addresses of up to 16 megabytes of physical memory. The 80386 features an on-chip page-based memory management unit. This capability should be especially attractive to designers of advanced operating systems and database or knowledge base systems wishing to implement virtual memory on either a local node or global system basis.

NUMERIC COPROCESSOR Augmenting the 80386 processor is its companion 80387 numeric coprocessor, providing IEEE-compatible scalar floating-point performance of up to 300 kiloflops at 64-bit precision. The 80387 instruction set also includes a number of transcendental and elementary arithmetic functions such as sine, arc tangent, and logarithm.

MEMORY Memory accesses from the node processor and floating-point units are buffered by a 64-Kbyte fast static cache memory. Node memory is modular and may be configured to be 1, 4, or 8 megabytes. In those applications that require very large node memories, two 8-megabyte memory modules may be stacked on a single node. The height of this 16-megabyte module requires that the adjacent slot in the cabinet be left empty.

DIRECT MEMORY ACCESS CONTROLLER For high-speed transfer of large blocks of data within a node and between nodes via the hypercube network, each node contains a multichannel direct memory access (DMA) controller. One channel is dedicated to inbound messages and another to outbound. Two other channels are available to software, with a total DMA bandwidth of 10.7 megabytes per second.

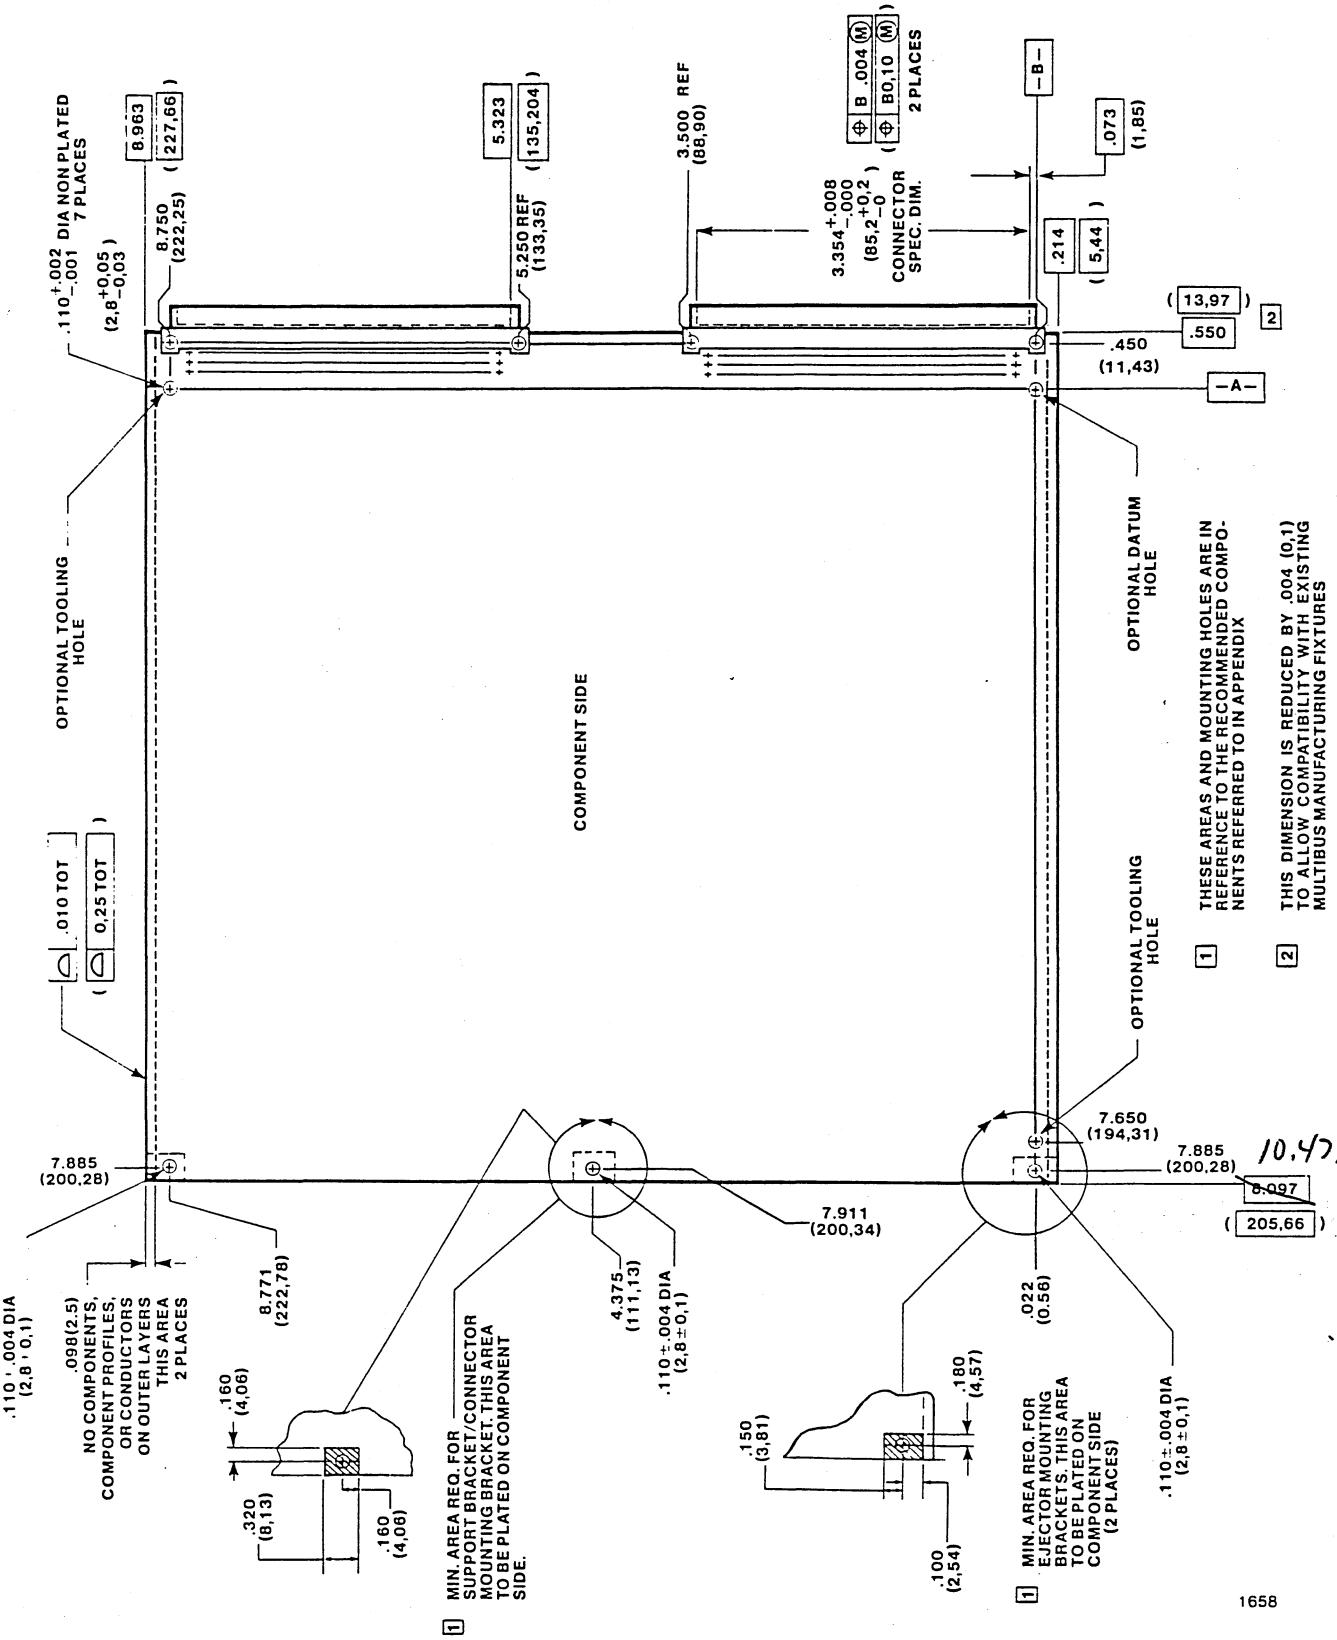

Figure 1-3. Double-High Board Size

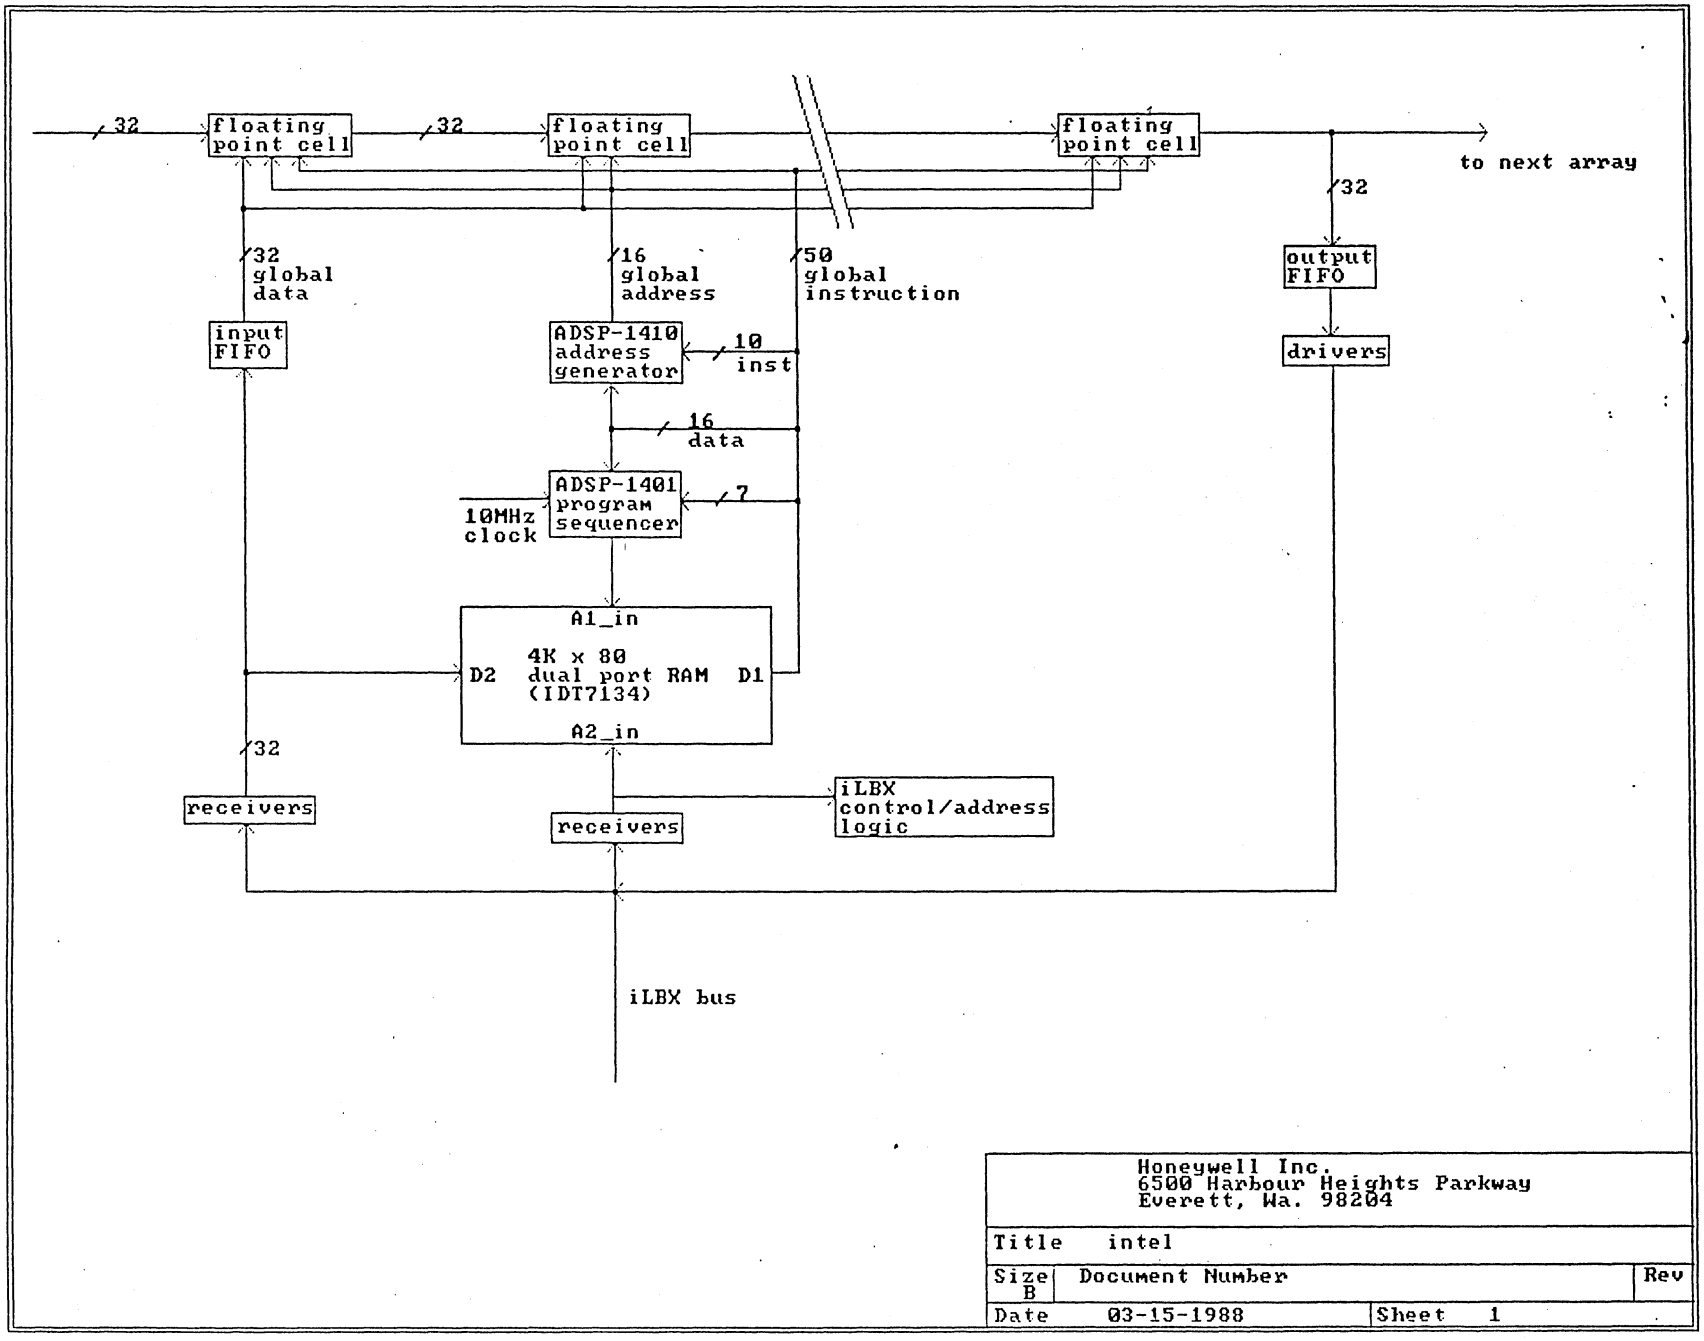

Honeywell Inc. 6500 Harbour Heights Parkway Everett, Wa. 98204		
Title intel		
Size B	Document Number	Rev
Date 03-15-1988	Sheet 1	

3-15-88

INTEL ACCELERATOR PARTS LIST (PER UNIT)

part	package	manufacturer	PINS	quantity
ADSP-1401	plastic DIP	Analog Devices	48	1
ADSP-1410	plastic DIP	Analog Devices	48	8
IDT7134	LCC	IDT	52	38
74ACT244	LCC	generic	20	10
XC-2064	PLCC	Xilinx	68	9
WTL 3332	PGA	Weitek	168	7
10MHz oscillator		-	14	1
74ACT00	PLCC	generic	14	18
IDT7201	LCC	IDT	32	8
.1 μF cap (ceramic)	-	-	2	[illegible]
10 μF [illegible]	-	-	2	1
96-pin connector	-	-	-	2
9" X 10" 4-layer PWB				1

603-2-IEC-C096-M1

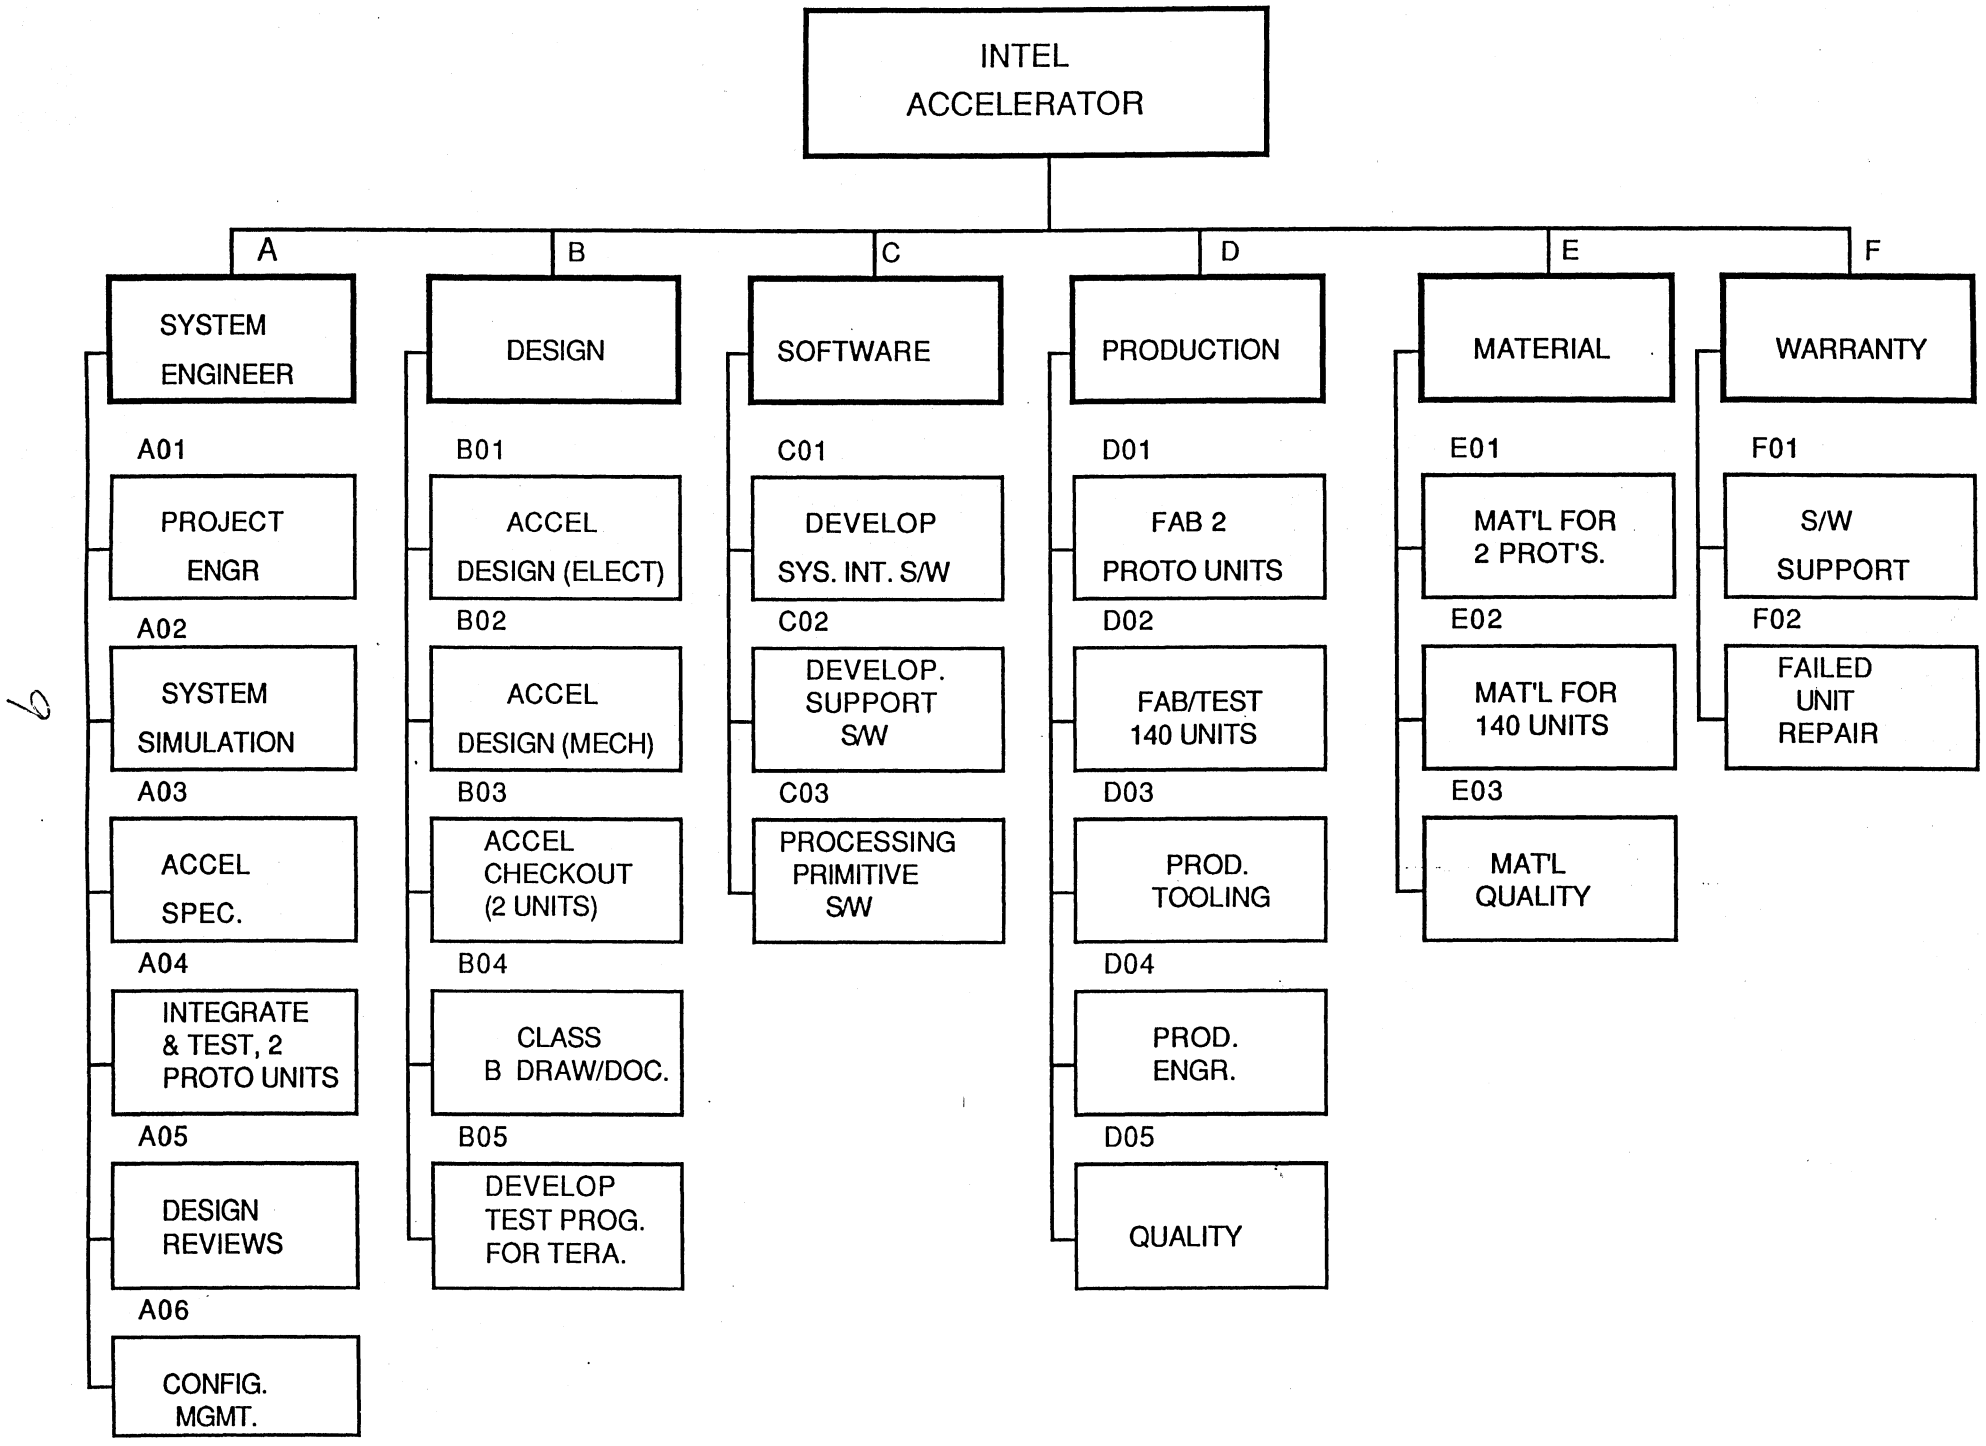
INTEL
ACCELERATOR
A
SYSTEM ENGINEER
A01
PROJECT ENGR
A02
SYSTEM SIMULATION
A03
ACCEL SPEC.
A04
INTEGRATE & TEST, 2 PROTO UNITS
A05
DESIGN REVIEWS
A06
CONFIG. MGMT.
B
DESIGN
B01
ACCEL DESIGN (ELECT)
B02
ACCEL DESIGN (MECH)
B03
ACCEL CHECKOUT (2 UNITS)
B04
CLASS B DRAW/DOC.
B05
DEVELOP TEST PROG. FOR TERA.
C
SOFTWARE
C01
DEVELOP SYS. INT. S/W
C02
DEVELOP. SUPPORT S/W
C03
PROCESSING PRIMITIVE S/W
D
PRODUCTION
D01
FAB 2 PROTO UNITS
D02
FAB/TEST 140 UNITS
D03
PROD. TOOLING
D04
PROD. ENGR.
D05
QUALITY
E
MATERIAL
E01
MAT'L FOR 2 PROT'S.
E02
MAT'L FOR 140 UNITS
E03
MAT'L QUALITY
F
WARRANTY
F01
S/W SUPPORT
F02
FAILED UNIT REPAIR

WBS ELEMENT ESTIMATOR RESPONSIBILITY MATRIX

WBS ELEMENT	ESTIMATE RESPONSIBILITY
A01	STEINER
A02	"
A03	"
A04	JOHNSON
A05	STEINER
A06	HUTCHINGS
B01	JOHNSON
B02	CORWIN
B03	JOHNSON
B05	STEINER
B05	CALLENDER
C01	KAYLOR
C02	KAYLOR
C03	KAYLOR
D01	MALPHRUS
D02	"
D03	"
D04	"
D05	"
E01	WARD
E02	"
E03	"
F01	KAYLOR
F02	MALPHRUS

INTEL ACCELERATOR
WBS DICTIONARY
3/16/88

Axx System Engineering

A01 Project Engineering

Direct all design activities in the development and transition-to-production of the accelerator circuit board and microprograms. Support four presentations to Intel or their customers. Manage all design WAs. Execute one EAC mid-year. Develop network program schedule to track status; update at weekly meetings is in A05.

A02 Simulation

Assess architecture efficacy using ADAS (Assume SRC runs ADAS) Include SRC subcontract.

Link together simulation models purchased from vendors to provide a micro-program development environment. Include purchase price of models.

A03 Accelerator Specification

Document the design requirements for the accelerator in internal memo format. Upgrade the spec to include schematics and mechanical details as they are developed.

Modify the specs to provide user documentation to Intel for their inclusion in data they provide to their customers. This includes principles of operation at a function call level and some marketing brochure descriptions. Detailed hardware descriptions bid in C02, of all places.

A04 Integrate/Test Prototypes

Develop a PC-base emulator of a Hypercube node. Integrate the first two units with the hypercube emulator at MSD and with a hypercube at Intel.

Execute the firmware library on the prototypes and verify proper operation (firmware will have been debugged on a software simulator as bid in C03).

A05 Design Reviews

Hold weekly progress/status meetings. Conduct a policy 6102 internal review. Include the time for a consultant chairman. Assume the review will focus on design-freeze issues rather than compliance with specs. Assume 1/2 day review plus preparation.

There wil be no customer program reviews.

A06 Configuration Management

Review drawings and maintain as-built configuration documentation.

Bxx Design

B01 Accelerator (Electrical)

Complete trade studies, including SIMD/MIMD and component selection.

Coordinating with the ADAS analysis of A02, complete the architecture design of the accelerator. Draw the schematic using OrCAD or equivalent. Perform sufficient analyses to compute operating speed.

B02 Accelerator (Mechanical)

Assume two board designs. The first is a multiwire board where we provide a netlist and pin placement drawing to the vendor.

Final board is a multilayer PWB. We will subcontract out the PWB route and design. This WA will provide technical oversight of the subcontract.

Assume drawings will be obtained from Intel or other sources for board dimensions, and mechanical parts such as connectors -- this is a standard, commercial board size.

See ~~box~~ B04 for drawings.

B03 Accelerator Checkout

(Eliminated; all effort is now included in A04)

B04 Class B Drawings

Complete the assembly drawing by modifying an Intel drawing. Put Honeywell format on schematics and parts lists that were generated by OrCAD. The only original drawing is top board layer pin placement.

B05 Teradyne Tester Programs

Run the simulation (A02) to generate raw data for tester. Modify Paul Schimpf's translator, if necessary, to translate that data to Teradyne format. Complete all software effort to prepare for board testing.

Write or buy LAZAR models for all IC types; generate a LAZAR models with ICs characterized at the boundary; and fault grade the vectors.

& Test Fixtures.

Cxx Software

C01 System Integration

Develop 80386 software to control the accelerator. The software will run on the Hypercube node as well as on our PC-based hypercube emulator.

This software will contain the accelerator drivers, function call library and basic operator intgeraction for controlling loops. Assume STEP-7 for microcode control store.

Upgrade documentation of this software to provide to Intel as draft inputs to their documentation.

C02 Development Support

Configure a commercial meta-assembler for accelerator microcode. Include purchase price of meta-assembler if necessary.

Upgrade meta-assembler documentation to provide as draft material to Intel. Include hardware description at a block level, from a micro-programmer's viewpoint.

C03 Primitives

Specify, code and debug basic primitive library. Assume scope of sky accelerator.

Upgrade documentation to provide to Intel. This will be a rough draft of primitive descriptions, listing the name, parameters and flags; describing the operation in one paragraph; and detailing exceptions, limits and formats. Assume one page or less each.

Dxx Production

D01 Fab

Fab two proto units using multiwire card for engineering to proof design.

D02 Testing

Fab and test 140 units. Unit test will be done on Teradyne. 128 units will be delivered to Intel and 12 units will be delivered to inventory at MSD. Teradyne test program will be developed, documented and verified in B05 as well as Teradyne test fixture.

D03 Development

Develop/procure any special tooling required to produce the units.

D04 Production Engineering

Production Engineering required to produce 140 units including make/buy analysis and support to design engineering to assure a producible design.

D05 Quality

Quality support to the production of 140 commercial grade units.

Exx Procurement

E01 Procure Four Units

Procure commercial grade unscreened components (IC's, connectors, etc.) for four units per attached parts lists. Two sets are used for the proto units and two sets are used for spares/production build.

E02 Procure 138 Units

Procure commercial grade unscreened components (IC's, connectors, etc.) for 138 units per attached parts lists. Include 3% average based on quantity used.

E03 Procure 140 Units

Provide material quality for commercial grade components for procurement of material for 140 units.

Fxx Follow-on Support

F01 Software Support

Provide applications S/W support for two years for development S/W and processing primitive S/W.

F02 Repairs

Provide for the repair of units failing within 90-day warranty period. Probability of unit failing in 90-day period is assumed to be 0.05. This implies six units will be returned for repair.

SCHEDULE

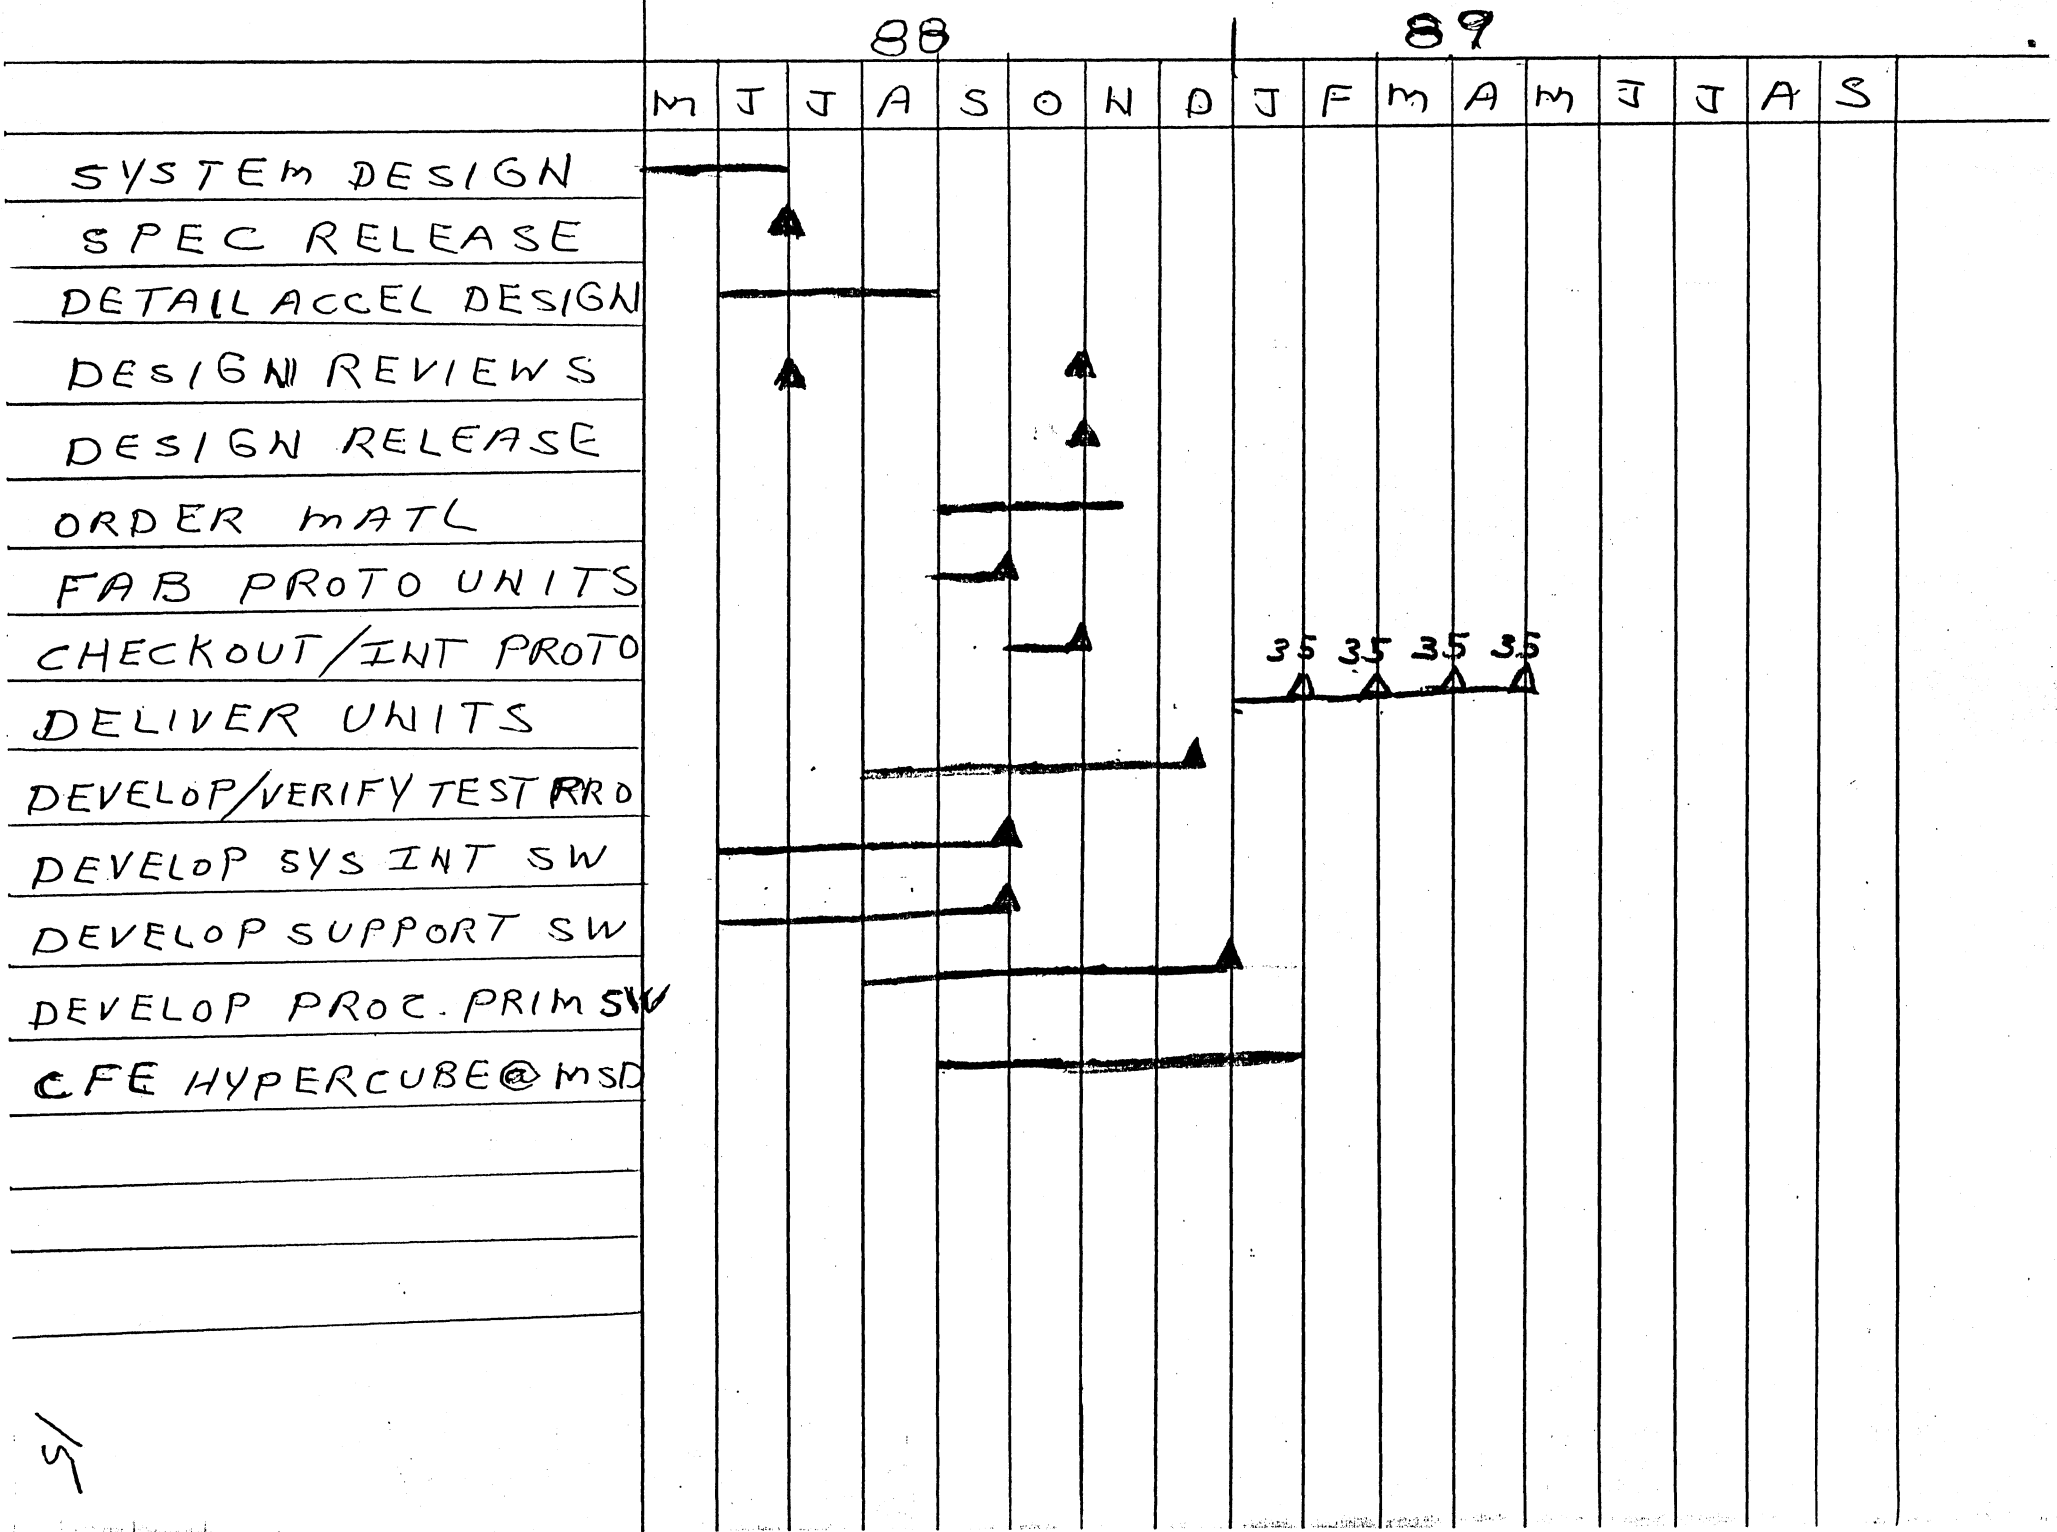
88
89
M J J A S O N D J F M A M J J A S
SYSTEM DESIGN
SPEC RELEASE
DETAIL ACCEL DESIGN
DESIGN REVIEWS
DESIGN RELEASE
ORDER MATL
FAB PROTO UNITS
CHECKOUT/INT PROTO
3.5 3.5 3.5 3.5
DELIVER UNITS
DEVELOP/VERIFY TEST PROG
DEVELOP SYS INT SW
DEVELOP SUPPORT SW
DEVELOP PROC. PRIM SW
CFE HYPERCUBE @ MSD

5/

Honeywell

SCOPE ANALYSIS

(VALIDITY OF THE BASELINE USED FOR COSTING)

PROGRAM	WBS/FUNCTION NO.
WBS/FUNCTION TITLE	

ANALYSIS/RATIONALE		GROSS ($ LBM)		PROB. %	NET ($ LBM)	
ITEM (DESCRIPTION OF RISK ELEMENT)	IMPACT (MW)	POSITIVE	NEGATIVE		POSITIVE	NEGATIVE
	TOTAL ($ LBM)			TOTAL ($ LBM)		

MSC 1246 3-79

PREPARED BY ____________________

DATE ____________________

TOTAL NET SCOPE RISK ($ LBM) [ ]

ALLOCATION OF HOURS TO PREPARE/REVIEW COST EST.

L. JOHNSON	37
B. BARTA	5
B. ENGERMAN	8
T. CALLENDER	3
R. KAYLOR	10
J. STEINER	12
QUALITY	3
PRODUCTION	12
	90

EJs COB Friday 3/18