Box 29 ISSCC presentation – 1992 3 of 5 102749955
folder 3

ISSCC PRESENTATION
FEBRUARY, 1992

100 MHz Macropipelined CISC CMOS Microprocessor

Roy W. Badeau, R. Iris Bahar, Debra Bernstein, Larry L. Biro, William J. Bowhill, John F. Brown, Michael A. Case, Ruben W. Castelino, Elizabeth M. Cooper, Maureen A. Delaney, David R. Deverell, John H. Edmondson, John J. Ellis, Timothy C. Fischer, Thomas F. Fox, Mary K. Gowan, Paul E. Gronowski, William V. Herrick, Anil K. Jain, Jeanne E. Meyer, Daniel G. Miner, Hamid Partovi, Victor Peng, Ronald P. Preston, Chandrasekhara Somanathan, Rebecca L. Stamm, Stephen C. Thierauf, G. Michael Uhler, Nicholas D. Wade, William R. Wheeler

Presenter - William V. Herrick

Digital Equipment Corporation, Hudson, MA

Outline

- Overview
- Selected Box Details
- Clock Generation and Distribution
- CAD and Verification
- Summary

Vital Statistics

- 0.75 µm N-well CMOS, 3.3 V technology, allowing 5 V input signals
- Tox = 105 Å, Leff = 0.5 µm
- 3 layers Al, SRAM LI, self-aligned silicide
- 1.3 M transistors, 1.62 X 1.46 cm^2
- 261 signal pads, 121 supply pads
- 339 pin THPGA, 5.5 mil SWB, thermal slug
- WC power dissipation = 16 W (avg), 18 W (peak)

RISC-Like Design

The macropipelined CISC chip is a RISC-like design with an instruction fetch and decode front end:

- Multiple semi-autonomous fully pipelined boxes
- Queues used at critical box interfaces to normalize operating rates
- Scoreboards used to resolve data dependencies
- Branch prediction used to keep pipeline full
- L/S μarchitecture with read-and-run, write-and-run

Functional Overview

- ***Ibox*:** macroinstruction fetch, parse, and dispatch
- ***Ebox*:** integer and logic instruction execution, and interrupt and exception processing
- ***Fbox:*** integer multiply and floating point instruction execution
- ***Mbox:*** memory management and reference processing, first-level on-chip Pcache
- ***Cbox:*** interface to second-level off-chip Bcache and memory subsystem

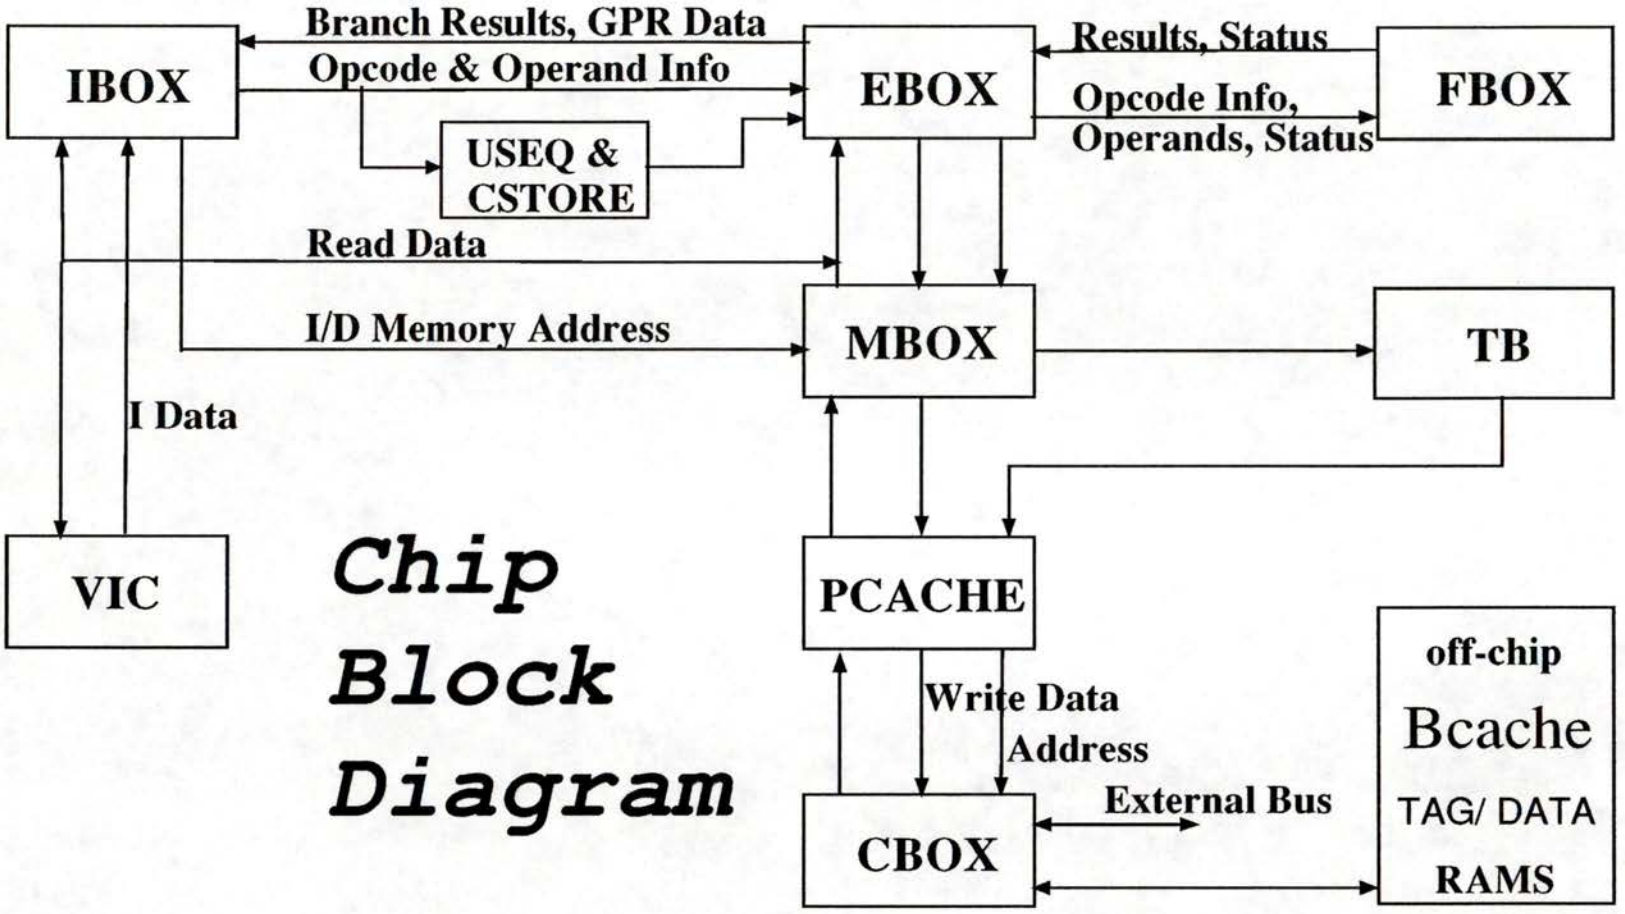
IBOX
Branch Results, GPR Data
Opcode & Operand Info
EBOX
Results, Status
Opcode Info,
Operands, Status
FBOX
USEQ &
CSTORE
Read Data
I/D Memory Address
MBOX
TB
I Data
VIC
Chip
Block
Diagram
PCACHE
off-chip
Bcache
TAG/ DATA
RAMS
Write Data
Address
External Bus
CBOX

Color Chip Photo

Chip Photo

Ibox

- Instruction stream prefetching, instruction parsing, operand specifier processing, and branch prediction
- 2 KB direct-mapped Virtual Instruction Cache
- 16 byte Prefetch Queue
- Instruction Decode and Issue Unit
- Complex Specifier Unit
- 512 x 4 bit conditional Branch Prediction array
- Inter-box Queues and a GPR Scoreboard

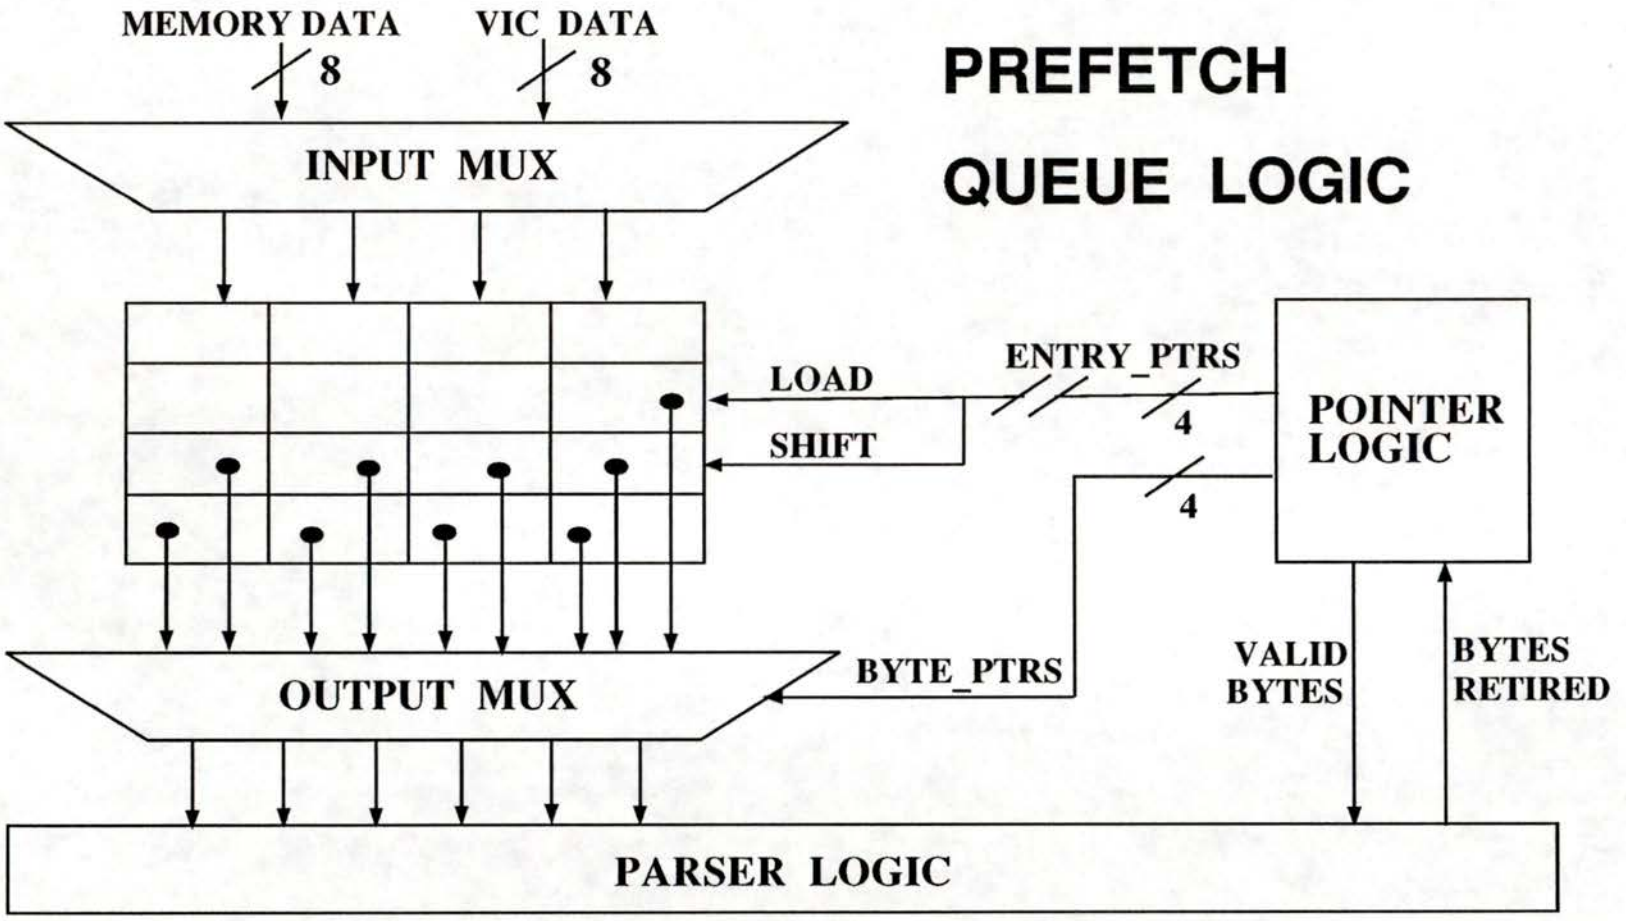
MEMORY DATA
8
VIC DATA
8
INPUT MUX
PREFETCH
QUEUE LOGIC
LOAD
SHIFT
ENTRY_PTRS
4
4
POINTER
LOGIC
BYTE_PTRS
OUTPUT MUX
VALID
BYTES
BYTES
RETIRED
PARSER LOGIC

PREFETCH QUEUE SCHEMATIC

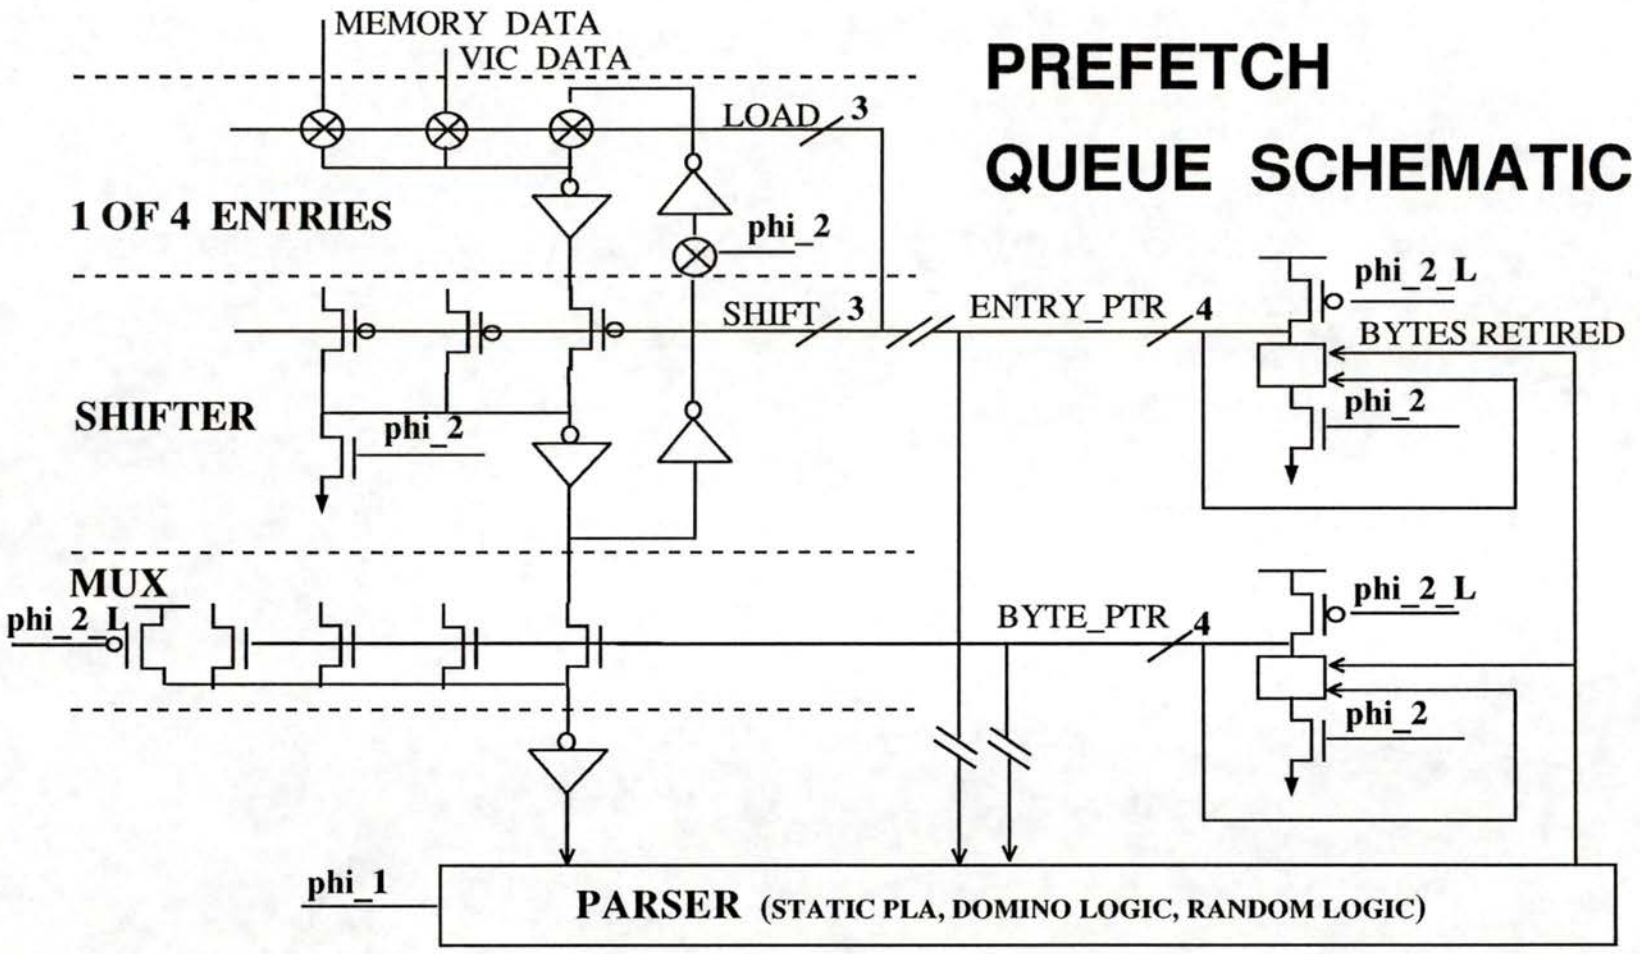

Ebox

- Instruction execution and coordination, including Fbox operand steering and instruction retire and result store orchestration (Fbox or FP scoreboard)
- Trap, fault, interrupt, and exception handling, and CPU initialization and state control
- Microsequencer and 1600 61-bit word Control Store with 20 patchable locations
- ALU, Shifter, and 5-port Register File
- Specialized HW features to speed execution

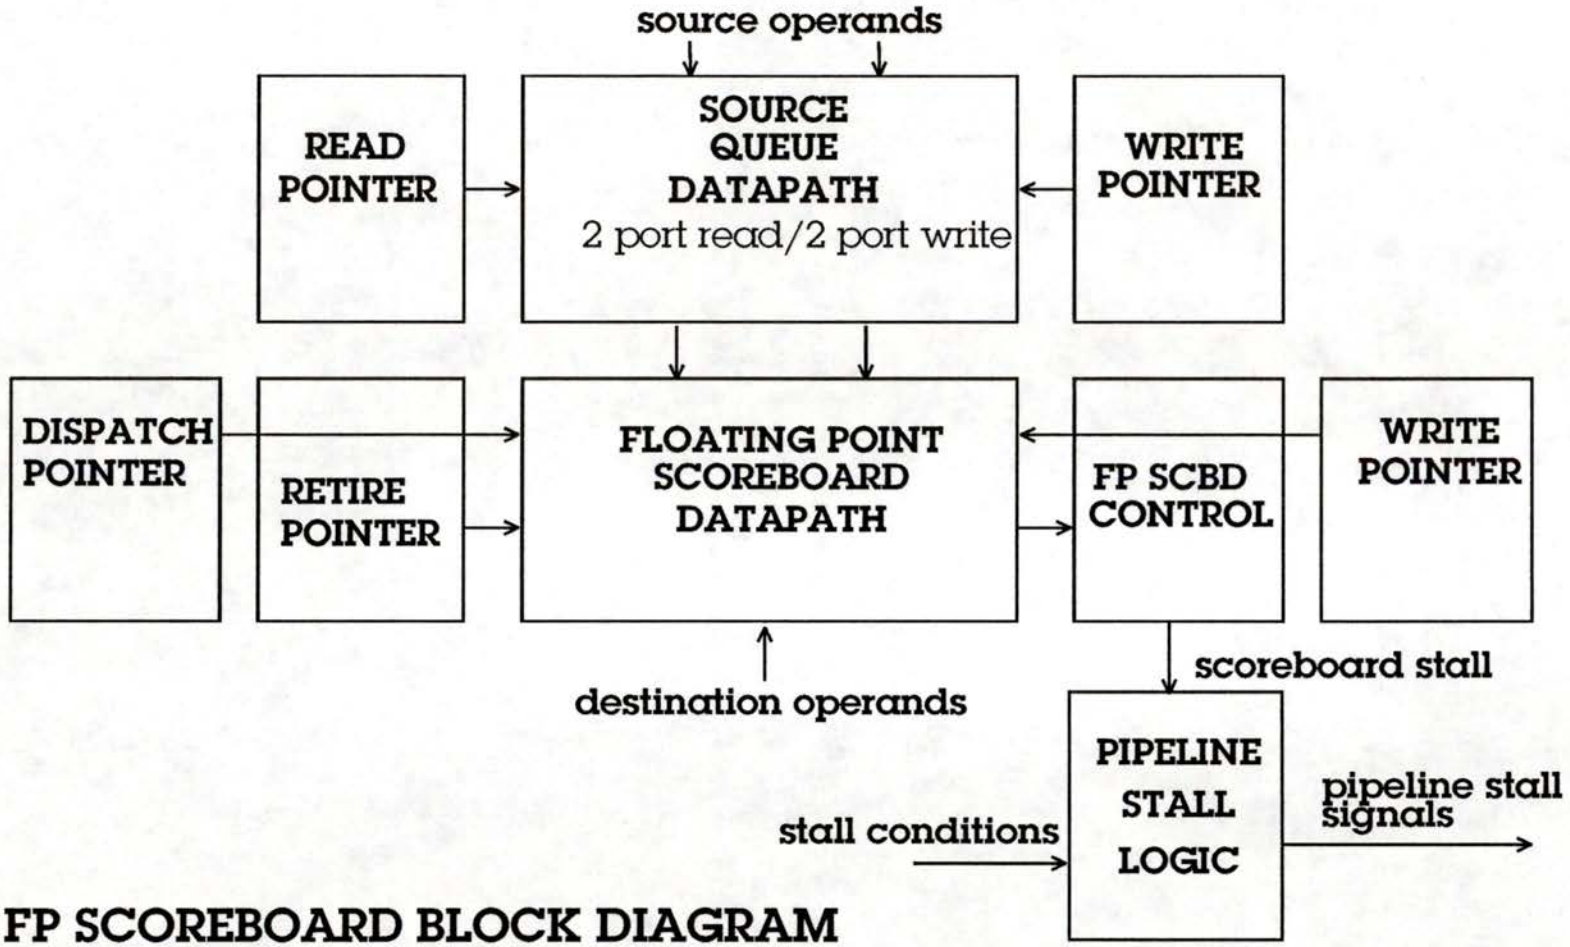

FP SCOREBOARD BLOCK DIAGRAM

FP SCOREBOARD CIRCUIT STRUCTURE

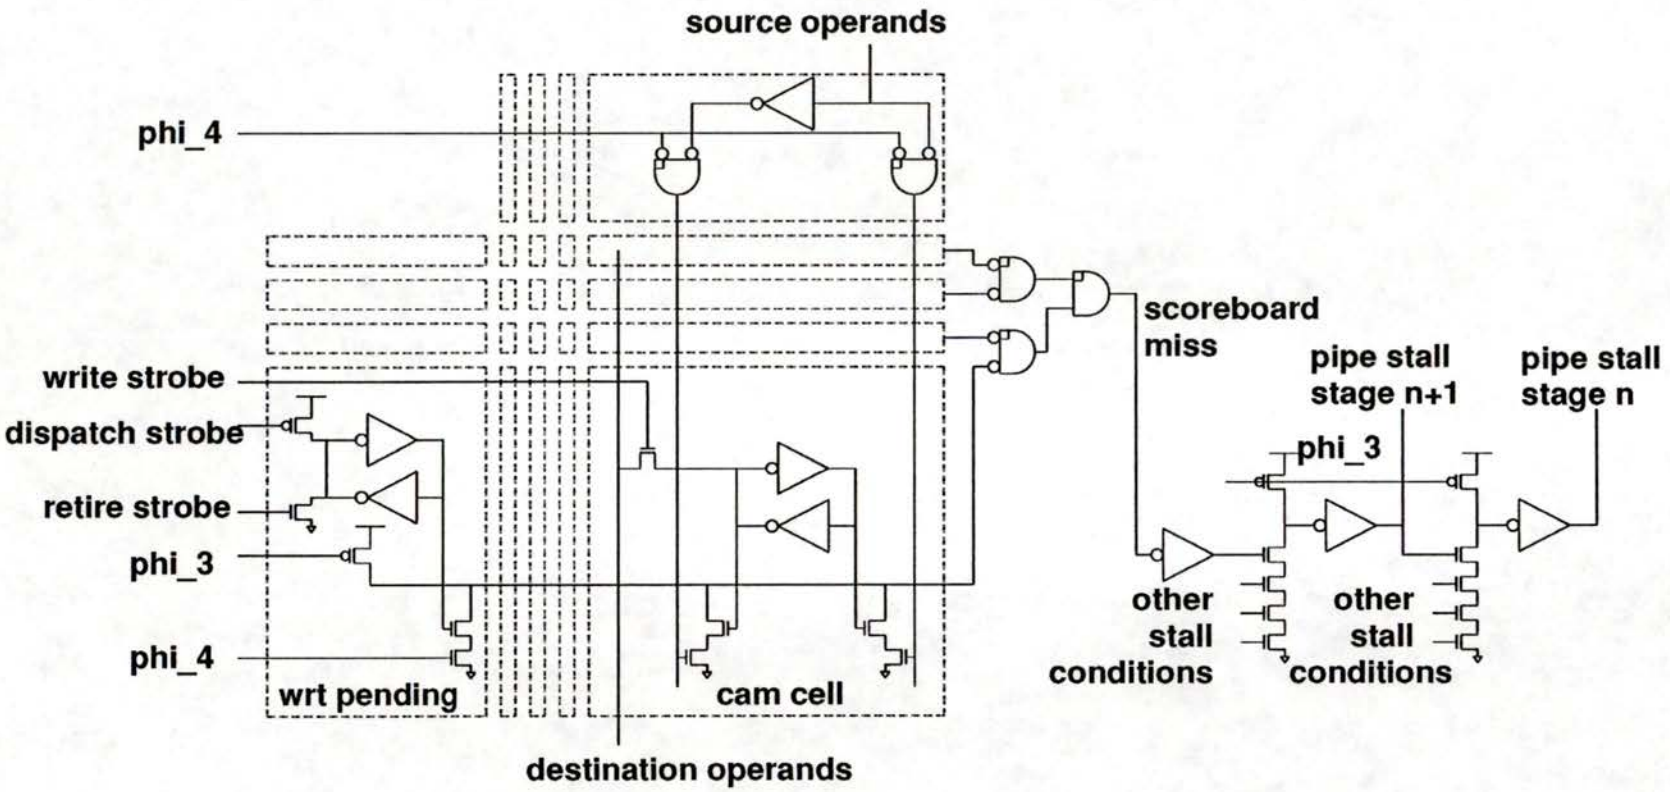

Fbox

- 4-Stage pipelined execution unit with a non-pipelined floating point divider and conditional bypass of the final stage
- Except for DIV instructions, single precision floating point instructions may be started every cycle; double precision and integer multiply, every two cycles
- Ebox supplies operands, stores Fbox results, and coordinates fault and exception handling with the Fbox

FBOX

4 - Stage

Pipeline

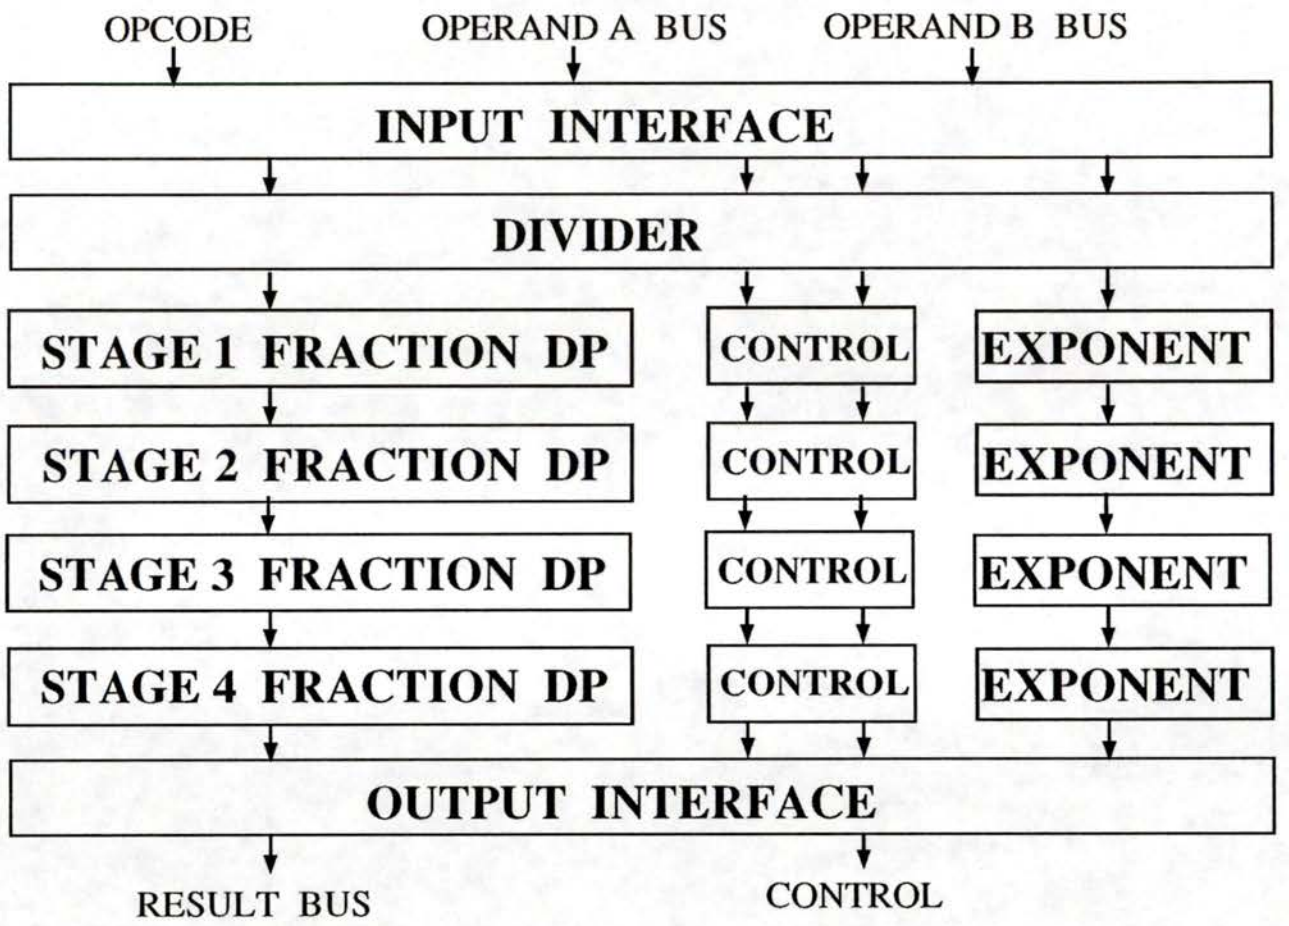

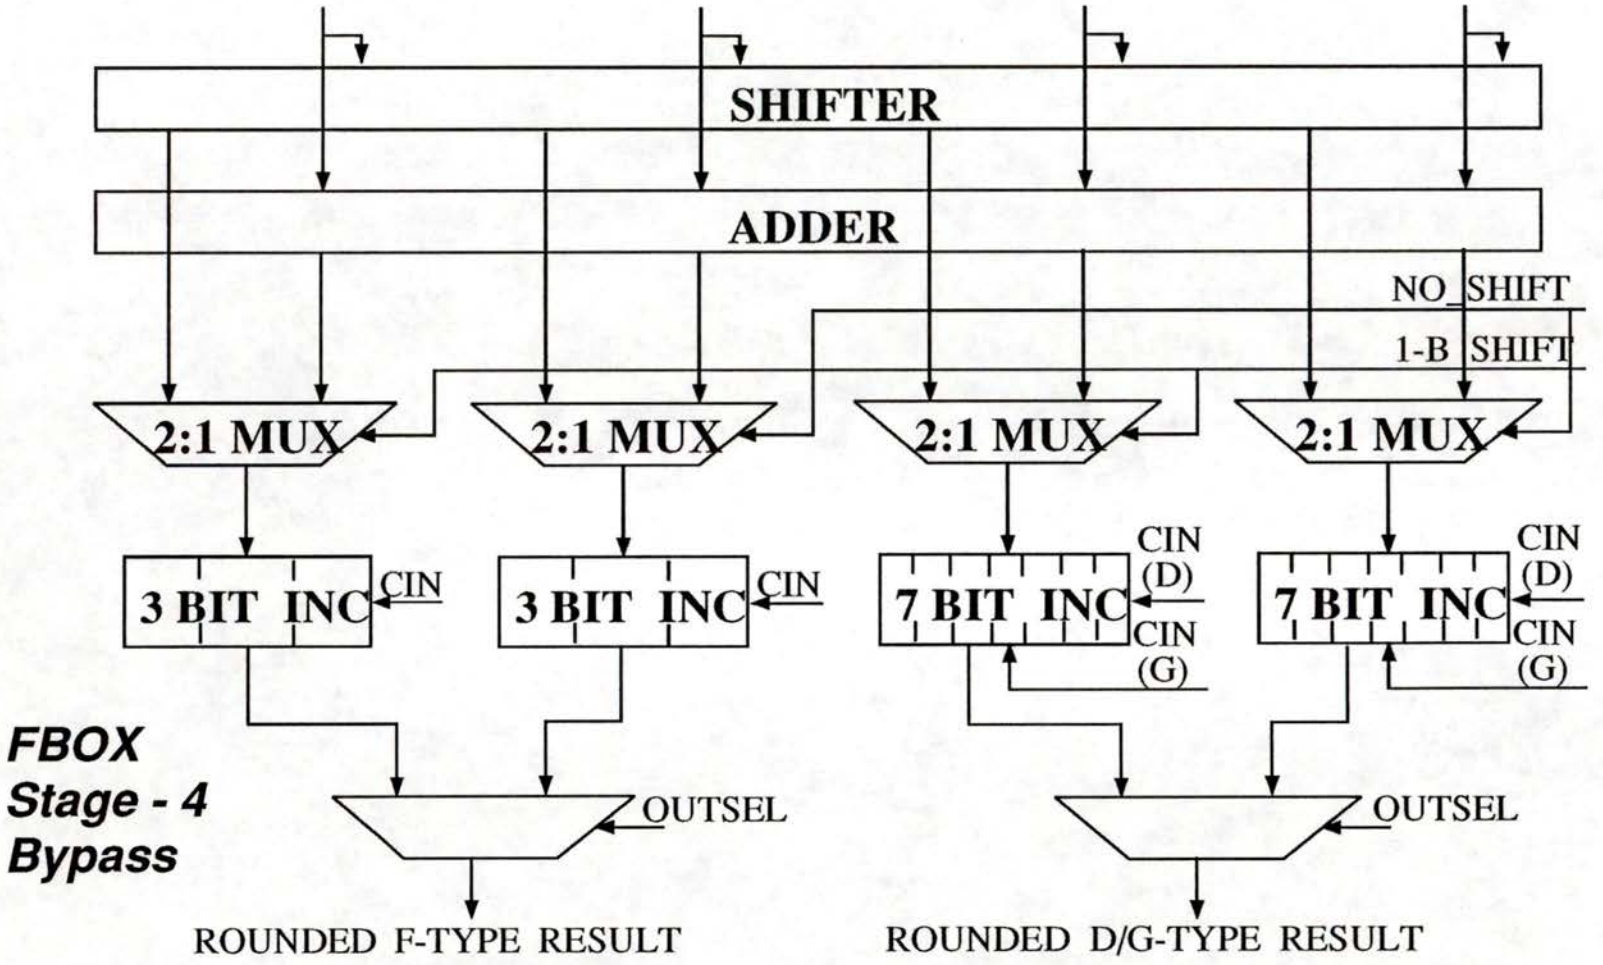

FBOX
Stage - 4
Bypass

Mbox

- Memory management (address translation, access checks), memory reference processing (prioritizing, sequencing), and Pcache control
- 96-Entry fully-associative TB
- 8 KB 2-way set-associative write-through physical address I-stream/D-stream Pcache
- Flow-through circuit techniques (domino structures, transparent latches, and Pcache TAG comparison)

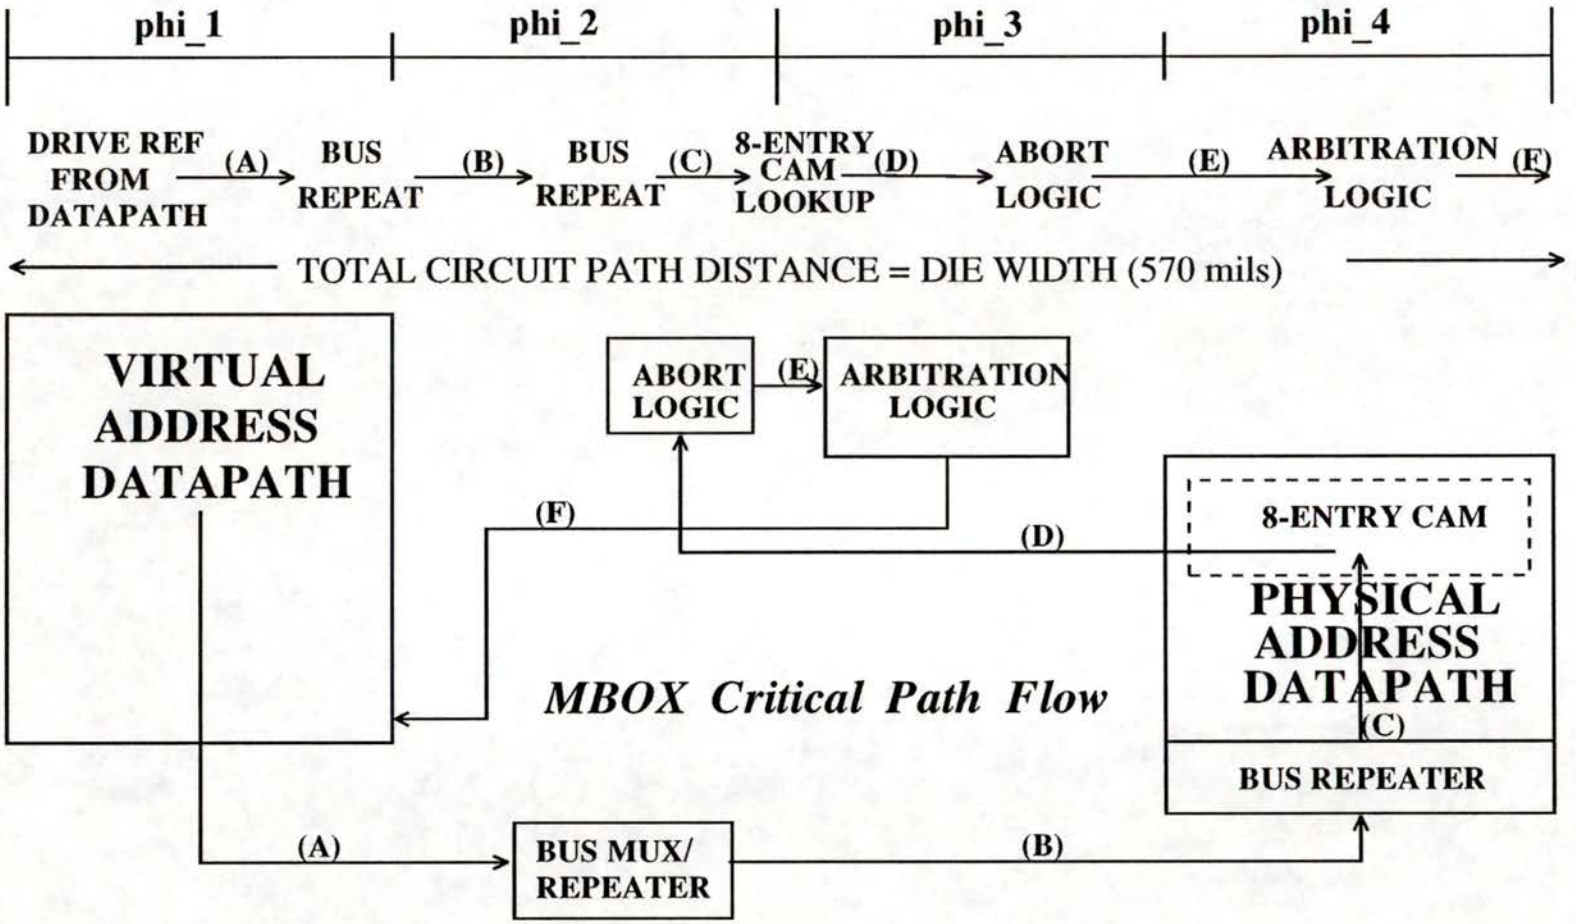
phi_1
phi_2
phi_3
phi_4
DRIVE REF FROM DATAPATH
(A)
BUS REPEAT
(B)
BUS REPEAT
(C)
8-ENTRY CAM LOOKUP
(D)
ABORT LOGIC
(E)
ARBITRATION LOGIC
(F)
TOTAL CIRCUIT PATH DISTANCE = DIE WIDTH (570 mils)
VIRTUAL ADDRESS DATAPATH
ABORT LOGIC
(E)
ARBITRATION LOGIC
(F)
(D)
8-ENTRY CAM
PHYSICAL ADDRESS DATAPATH
(C)
BUS REPEATER
MBOX Critical Path Flow
(A)
BUS MUX/ REPEATER
(B)

Chip Photo

PCACHE TAG COMPARE / BLOCK SELECT LOGIC

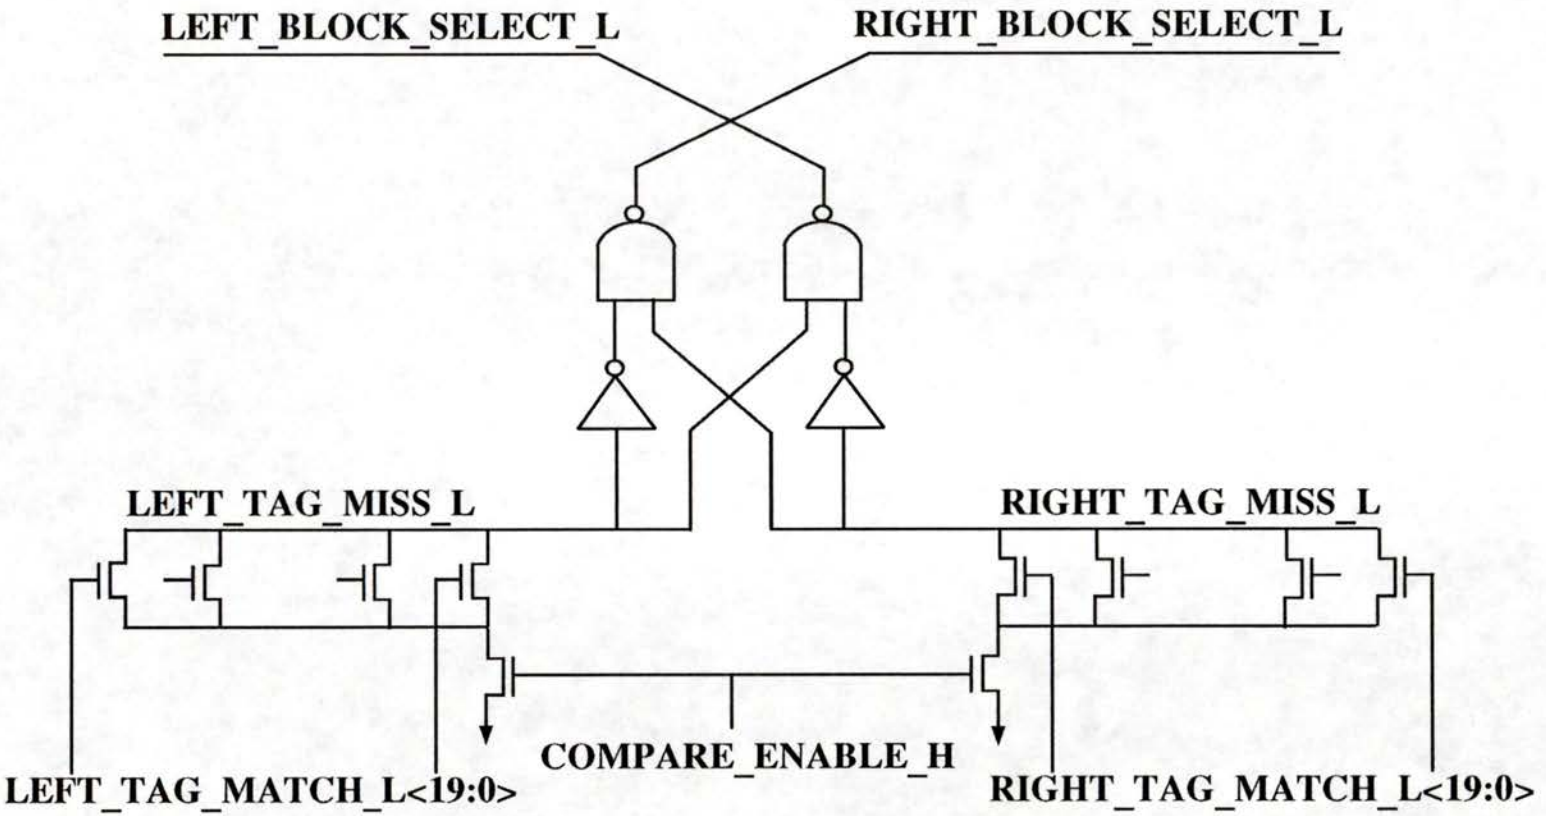

Cbox

- Controller for direct-mapped, writeback, off-chip Bcache, with support for 4 sizes (2 MB, 512 KB, 256 KB, or 128 KB) and 5 RAM speeds (2 TAG and 3 DATA)
- Write packing for sequential writes
- External bus interface with support for cache coherency protocols in multiprocessing systems
- External bus runs at Tcycle = 30 ns

Clock Generation and Distribution

- 400 MHz ECL oscillator generates 100 MHz 4-phase internal clocks and synchronized external clocks
- Distribution and buffering minimized clock skews and edge rates
- Across the chip: clock skews < 500 ps, rise and fall times < 650 ps
- RC delays: central clock routing < 30 ps, buffered clock routing < 125 ps

Chip Photo

Ebeam Traces 1

Ebeam Traces 2

CAD Tools and Methods

Full custom circuit and layout techniques, supported by structured custom methods and in-house tools:

- *Synthesis* - logic and layout
- *Design Capture* - schematic and layout
- *Simulation* - RTL, gate, switch, and circuit
- *Extraction* - wirelist, resistance, capacitance
- *Verification* - functional, timing, circuit, layout

Verification Results

- 1.1 Billion CPU cycles simulated on the RTL model
- 75 Million cycles simulated on the gate level model (derived from the transistor netlist); the operating system was booted on this model
- 600 Thousand cycles simulated at the switch level
- Pseudo-random tests improved functional coverage
- The timing verifier, static circuit checker, and electrical checker tools minimized or eliminated speed, circuit, and reliability problems in the design

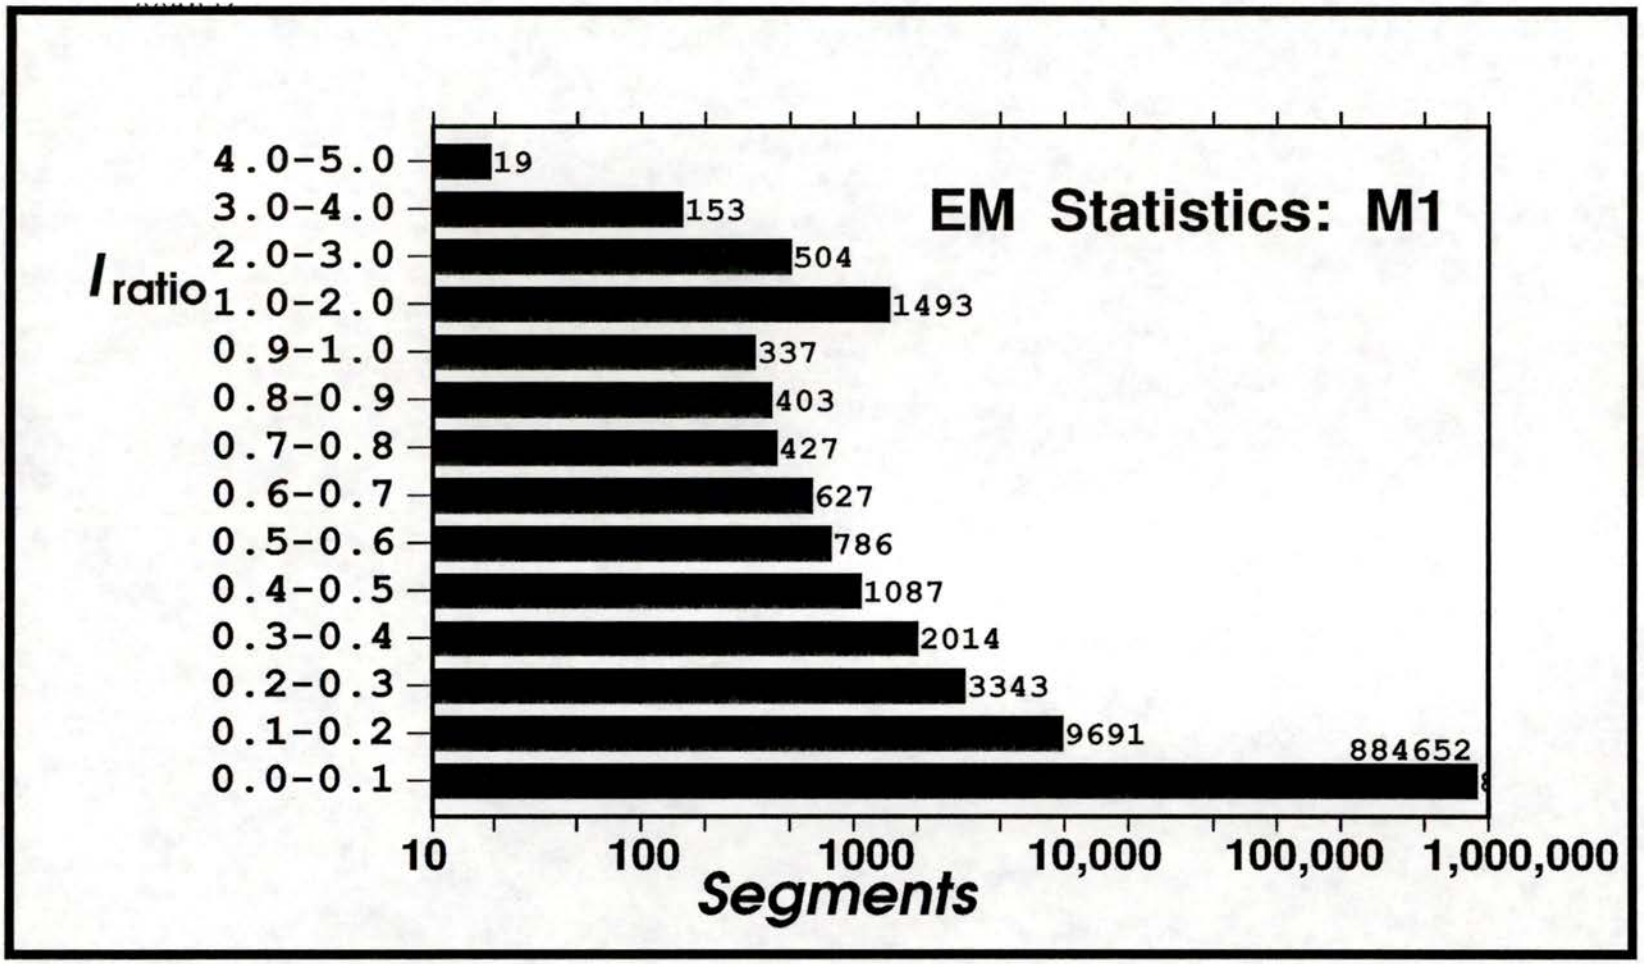
EM Statistics: M1
I_{ratio}
4.0-5.0
3.0-4.0
2.0-3.0
1.0-2.0
0.9-1.0
0.8-0.9
0.7-0.8
0.6-0.7
0.5-0.6
0.4-0.5
0.3-0.4
0.2-0.3
0.1-0.2
0.0-0.1
19
153
504
1493
337
403
427
627
786
1087
2014
3343
9691
884652
10
100
1000
10,000
100,000
1,000,000
Segments

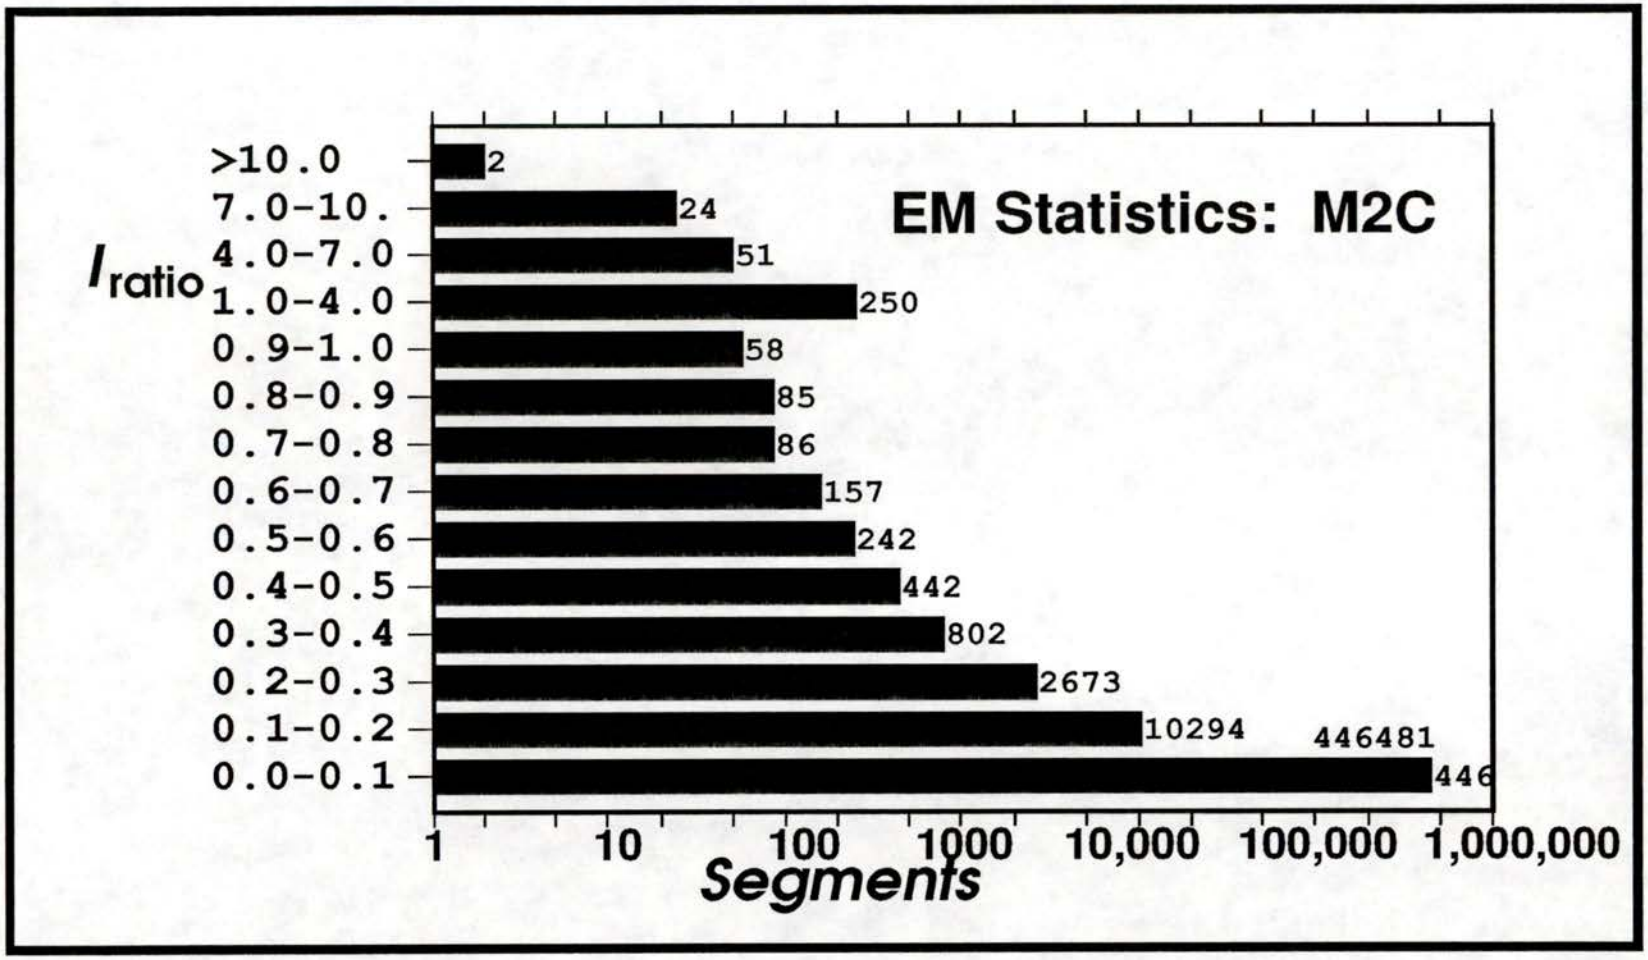
EM Statistics: M2C
I_{ratio}
>10.0
7.0-10.
4.0-7.0
1.0-4.0
0.9-1.0
0.8-0.9
0.7-0.8
0.6-0.7
0.5-0.6
0.4-0.5
0.3-0.4
0.2-0.3
0.1-0.2
0.0-0.1
2
24
51
250
58
85
86
157
242
442
802
2673
10294
446481
446
1
10
100
1000
10,000
100,000
1,000,000
Segments

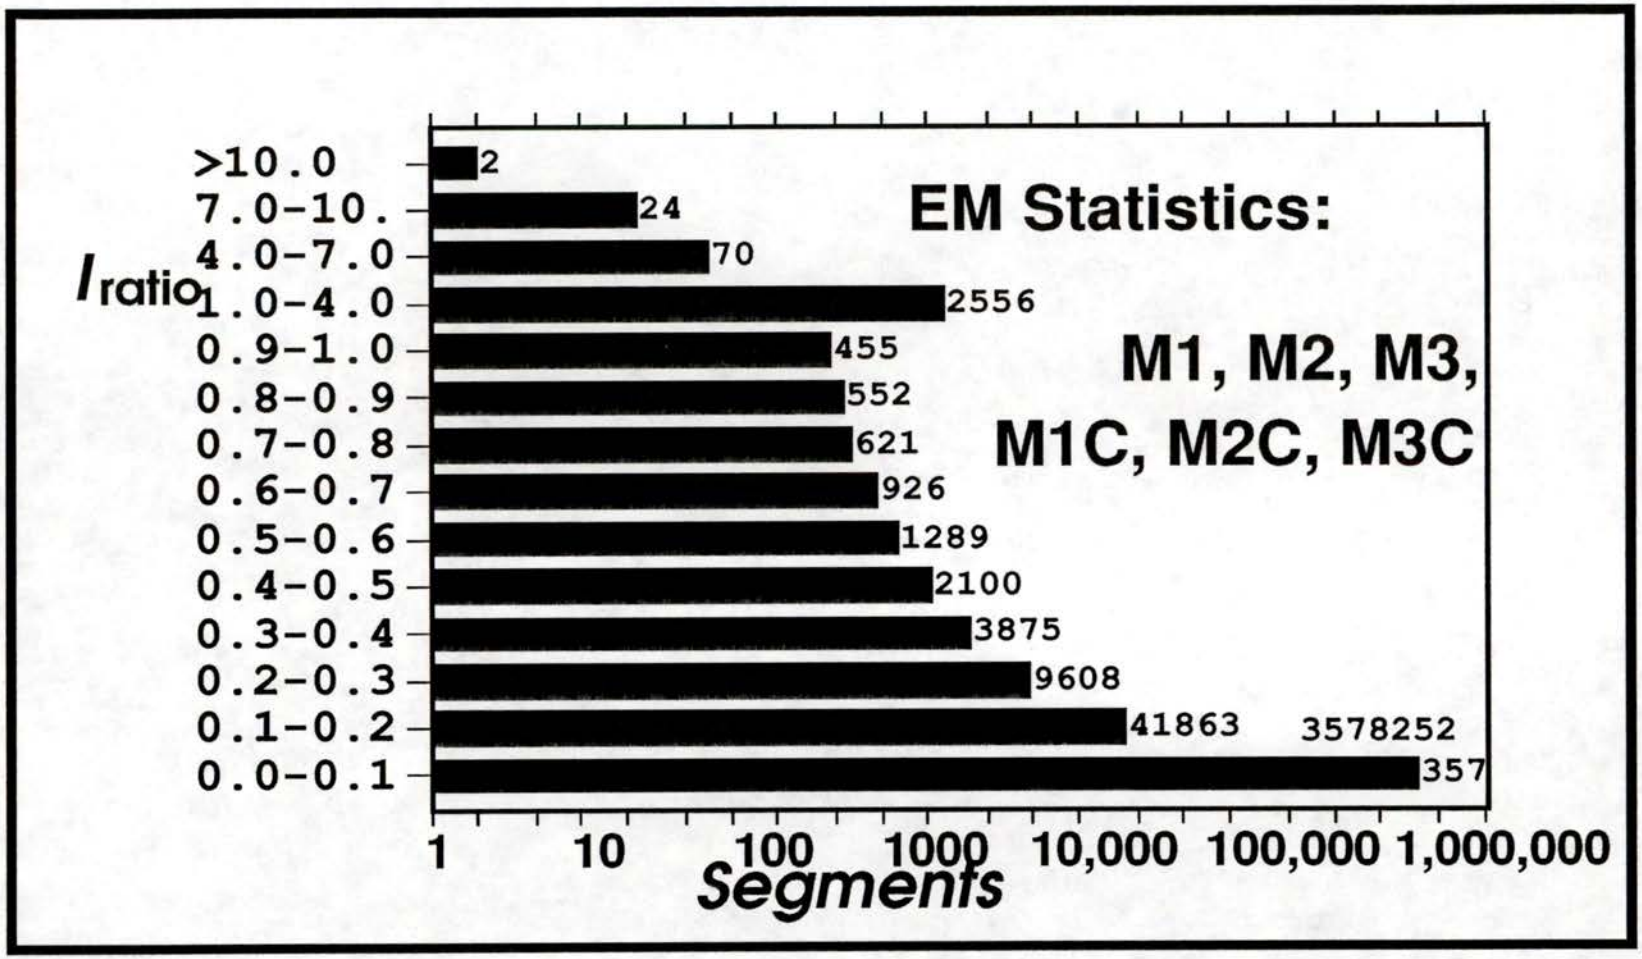
EM Statistics:
M1, M2, M3,
M1C, M2C, M3C
I ratio
>10.0
7.0-10.
4.0-7.0
1.0-4.0
0.9-1.0
0.8-0.9
0.7-0.8
0.6-0.7
0.5-0.6
0.4-0.5
0.3-0.4
0.2-0.3
0.1-0.2
0.0-0.1
2
24
70
2556
455
552
621
926
1289
2100
3875
9608
41863
3578252
357
1
10
100
1000
10,000
100,000
1,000,000
Segments

Shmoo Plot

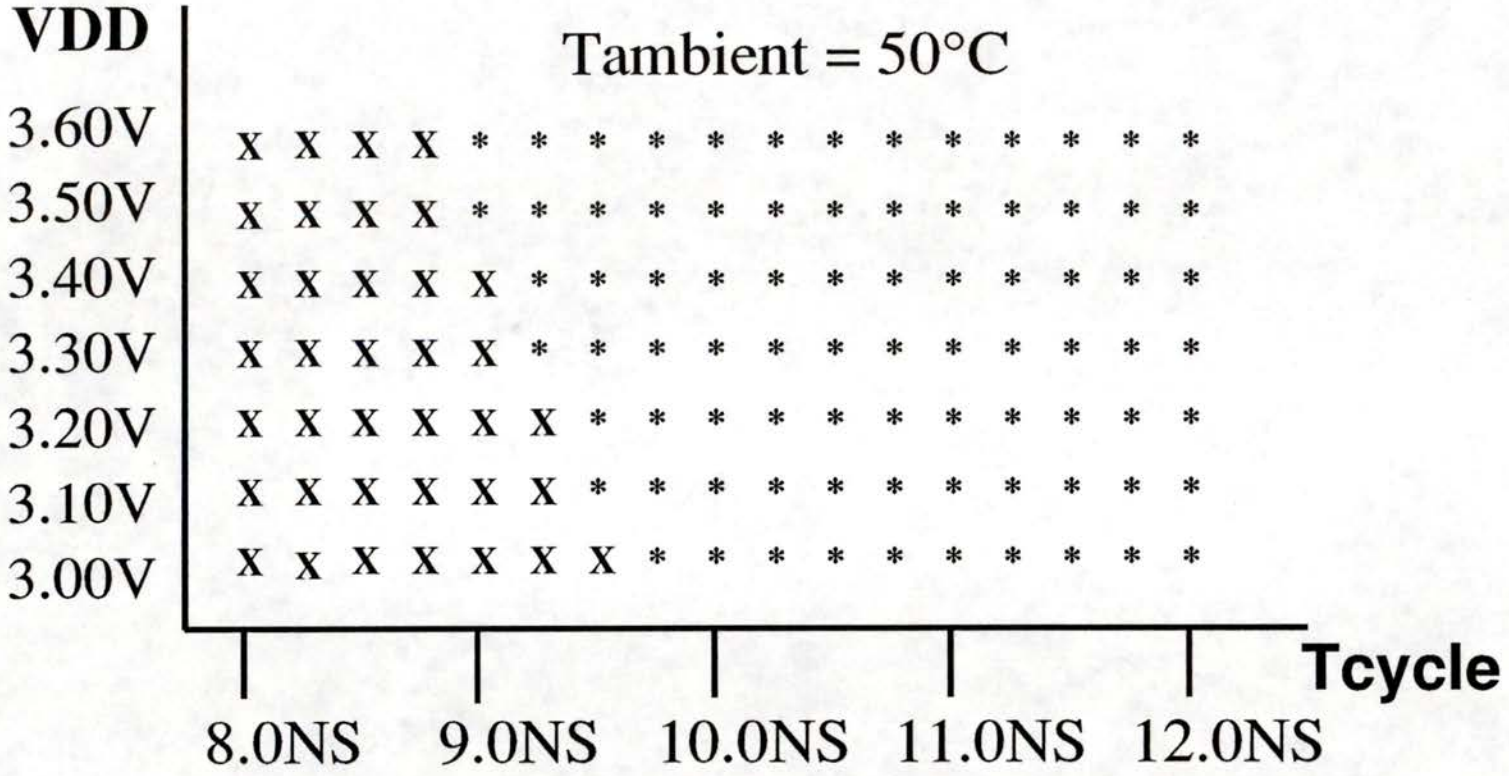
VDD
Tambient = 50°C
3.60V
3.50V
3.40V
3.30V
3.20V
3.10V
3.00V
Tcycle
8.0NS
9.0NS
10.0NS
11.0NS
12.0NS

Summary

- 100 MHz operation in a system environment has been demonstrated
- 50 SPECMarks and > 85 TPS performance levels have been measured
- Less than 2 weeks from wafers received to operating system boot
- 15 bugs found in pass 1 chips; none impeded prototype development
- Pass 2 design is shippable

100 MHz Macropipelined CISC CMOS Microprocessor

Session TA6.1

by

Roy W. Badeau, R. Iris Bahar, Debra Bernstein, Larry L. Biro,
William J. Bowhill, John F. Brown, Michael A. Case,
Ruben W. Castelino, Elizabeth M. Cooper, Maureen A. Delaney,
David R. Deverell, John H. Edmondson, John J. Ellis,
Timothy C. Fischer, Thomas F. Fox, Mary K. Gowan,
Paul E. Gronowski, William V. Herrick, Anil K. Jain,
Jeanne E. Meyer, Daniel G. Miner, Hamid Partovi, Victor Peng,
Ronald P. Preston, Chandrasekhara Somanathan, Rebecca L. Stamm,
Stephen C. Thierauf, G. Michael Uhler, Nicholas D. Wade,
William R. Wheeler

Digital Equipment Corporation
Hudson, MA

Presenter:

William V. Herrick
Digital Equipment Corporation
77 Reed Road, HL02-3/C11
Hudson, MA 01749

Phone (508) 568 4069
FAX (508) 568 4681

100 MHz Macropipelined CISC CMOS Microprocessor

Session TA6.1

by

Roy W. Badeau, R. Iris Bahar, Debra Bernstein, Larry L. Biro, William J. Bowhill, John F. Brown, Michael A. Case, Ruben W. Castelino, Elizabeth M. Cooper, Maureen A. Delaney, David R. Deverell, John H. Edmondson, John J. Ellis, Timothy C. Fischer, Thomas F. Fox, Mary K. Gowan, Paul E. Gronowski, William V. Herrick, Anil K. Jain, Jeanne E. Meyer, Daniel G. Miner, Hamid Partovi, Victor Peng, Ronald P. Preston, Chandrasekhara Somanathan, Rebecca L. Stamm, Stephen C. Thierauf, G. Michael Uhler, Nicholas D. Wade, William R. Wheeler

Digital Equipment Corporation
Hudson, MA

Presenter:

William V. Herrick
Digital Equipment Corporation
77 Reed Road, HL02-3/C11
Hudson, MA 01749

Phone (508) 568 4069
FAX (508) 568 4681

ABSTRACT

A macropipelined CISC microprocessor was implemented in a 0.75µm CMOS 3.3V technology. The 1.3M transistor custom chip measures 1.62 X 1.46 cm² and dissipates 18 watts. The TPI is 2.4X better than that of a fully micropipelined implementation of the same instruction set. 100 MHz parts were benchmarked at 50 SPECMARKS.

A macropipelined CISC microprocessor was implemented in a 0.75µm CMOS 3.3V 3 metal layer technology. The 1.3M transistor custom chip measures 1.62 X 1.46 cm^2, dissipates 18 watts (peak), and is packaged in a 339 pin PGA. The chip implements a macroinstruction pipeline to execute the instruction set of a popular CISC minicomputer. Fig.1 depicts a block diagram of the major functional units. Fig.2 shows a die photomicrograph.

The IBOX fetches and decodes macroinstructions, and contains the instruction parser/operand decoder, a 2 KByte direct-mapped virtual instruction cache, and a dynamic branch prediction unit. The microsequencer and EBOX work together as a 4 stage micropipeline to execute instructions under microprogrammed control. The control store is a 1600 61-bit microword ROM; the microcode can be patched using a CAM to substitute RAM words for ROM. The FBOX is a 4 stage pipelined floating point and integer multiply execution unit. Most floating point instructions have a latency of 4 and a repetition rate of 1 cycle. The MBOX processes memory references. The 2-way set-associative 8 KByte write-through primary cache (Pcache) is in the MBOX. Looking up a physical address from the 96-entry fully-associative TB, accessing the Pcache, and rotating and passing the Pcache data to the requester takes 1.5 cycles. Pcache operations take 1 cycle. The CBOX controls the off-chip 2nd level direct-mapped write-back cache and maintains multiprocessor cache coherency using ownership protocols. The BIU is the interface to the external bus which connects the chip to the memory subsystem.

The added parallelism of the macropipeline combined with the larger caches and more efficient compilers yields a 2.4 fold improvement in TPI over that achieved for a fully micropipelined microprocessor with the same CISC instruction set [1].

The chip has 2 sets of synchronous clocks: the 100 MHz 4 phase internal clocks, and the 4 phase external clocks that run 3 times slower and control the BIU, pad logic, and peripheral chips. The chip is functionally insensitive to overlap between adjacent phases. For performance, it was critical to keep skews small and edge rates sharp, so special attention was paid to clock generation and distribution (Fig.3). A 400 MHz ECL oscillator drives the receiver buffer at the top of the chip. The output is routed to the global clock generator (CLKGEN), a 400 MHz 12-state FSM, at the center. The internal (divide by 4) and external (divide by 12) phases are derived by decoding these states. The outputs of the phase decoders are synchronized by flip-flops (Fig.4), to eliminate generation skew. The outputs of the flip-flops are buffered by 4 inverters and distributed, using the 3rd metal layer (M3, 16 mOhms/sq.), in the central clock routing channel that extends the height of the chip. Clocks are supplied to the boxes by tapping off the central clock routing and buffering each signal with 4 inverters. This buffering minimizes the loads seen by the clocks in the routing channel, where the RC delays are held to 30 ps. The buffered clocks are then distributed east and west, again using M3 but strapped with M2. Before the clocks are used, they are inverted locally.

These final stages of buffering reduce the loading on the east-west clocks and sharpen the clock edges. Conditional clocks are generated without a delay penalty by using 2-input gates for the local buffers. The layout for the clock buffers and routing was balanced: the buffer sizes were tuned to the extracted parasitics and dummy loads were added to the more lightly-loaded phases. An on-chip decoupling capacitor (30 nF) was distributed throughout the buffer layout to minimize Vdd/Vss noise when the clocks switch rapidly.

The clock signals were measured on silicon by E-Beam probe (Fig.5). The clock skews and edge rates across this 1.62 cm chip are less than 0.5 ns and 0.65 ns, respectively.

Dynamic logic and differential cascode circuit techniques were used to achieve the 100 MHz clock rate. For example, cascode logic was used in the FBOX double precision hardware divider. The divider is a radix 2 SRT divider with overlapped partial remainder calculation (using carry save adders) and quotient bit selection (QBS). Fig.6 shows the most important QBS circuit in the divider critical path. The circuit is a cascode implementation of a 3 bit carry chain plus sum logic (the XOR stage) for one bit. Logically, the circuit outputs the sign bit (T&C) of the estimated partial remainder. The propagation delay, from the least significant G0 (generate) input to the sign outputs, was simulated to be 1.2 ns for typical parts operating under worst case conditions.

The design was extensively analyzed for logical correctness and electrical integrity before fabrication. 1 billion cycles were run on the RTL model. 75 million cycles were run on the gate-level model (abstracted from transistor netlists); the operating system was also booted on this model. 600k cycles were run on the switch-level model. A 3-D capacitance extractor was used to determine the load on every node. A trapezoid-based extractor was used to determine node resistances: clock nodes had over 100k elements. The extracted R's and C's were fed to a static timing verifier which traversed 350k signal paths and checked 42k timing constraints on a 500k transistor netlist in a single run. CAD tools were used to identify potential design and reliability problems: coupling, latchup, charge-sharing, noise, race conditions, hot carrier stresses, and EM.

The chip booted the operating system at 100 MHz on 1st silicon. The patchable control store proved invaluable: the few obscure bugs that slipped through verification were easily fixed by patching the microcode. 100 MHz parts were benchmarked at 50 SPECMARKS, out-performing all previously reported CISC microprocessors.

We would like to acknowledge the contributions of the following people: W.Anderson, E.Arnold, B.Benschneider, M.Benoit, D.Bhavsar, M.Blaskovich, P.Boucher, L.Briggs, S.Butler, R.Calcagni, S.Carroll, A.Cave, S.Chopra, M.Coiley, A.DiPace, D.Donchin, D.DuVarney, R.Dutta, H.Fair, G.Franceschi, N.Geagan, S.Goel, J.Grodstein, W.Grundmann, H.Harkness, R.Hicks, C.Holub,

R.Khanna, R.Kiser, D.Koslow, K.Ladd, S.Levitin, S.Martin, T.McDermott, K.McFadden, B.McGee, M.Minardi, T.O'Brien, J.Pan, N.Phillips, J.Pickholtz, R.Razdan, S.Samudrala, J.Siegel, K.Siegel, J.StLaurent, R.Supnik, M.Tareila, S.Watkins, Y.Yen

[1] R. Allmon, et al., "CMOS Implementation of a 32b Computer," ISSCC DIGEST OF TECHNICAL PAPERS; p. 80-81; Feb., 1989.

FIGURE CAPTIONS

Figure 1 - Chip Block Diagram Showing Macropipeline Segments

Figure 2 - 100 MHz Macropipelined CISC CMOS Microprocessor

Figure 3 - Topological Block Diagram Of Clock Generation and Distribution Scheme

Figure 4 - Clock Phase Decode and Synchronizing Flip Flop

Figure 5 - On-Chip Clock Waveforms From E-Beam Probe

Figure 6 - Estimated Partial Remainder Sign Detection Logic In FBOX Divider

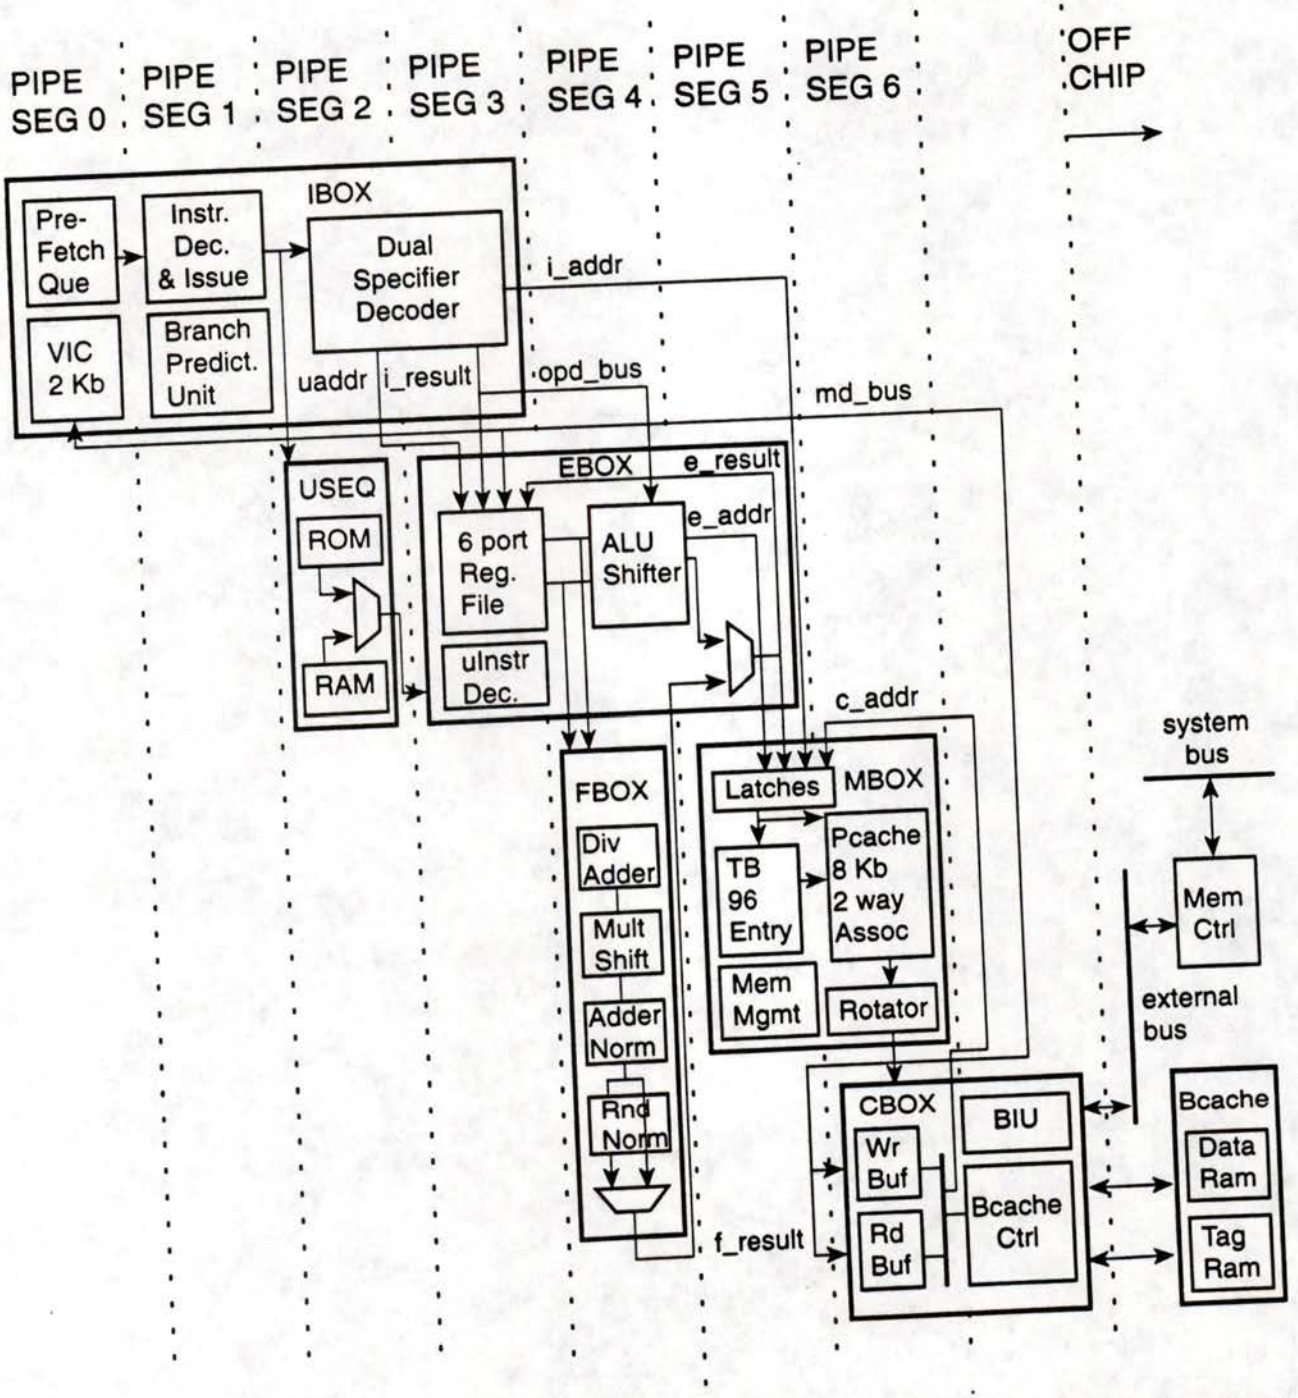

Figure 1: Chip Block Diagram Showing Macropipeline Segments.

Authors: R. Badeau, et.al., Session: TA6.1

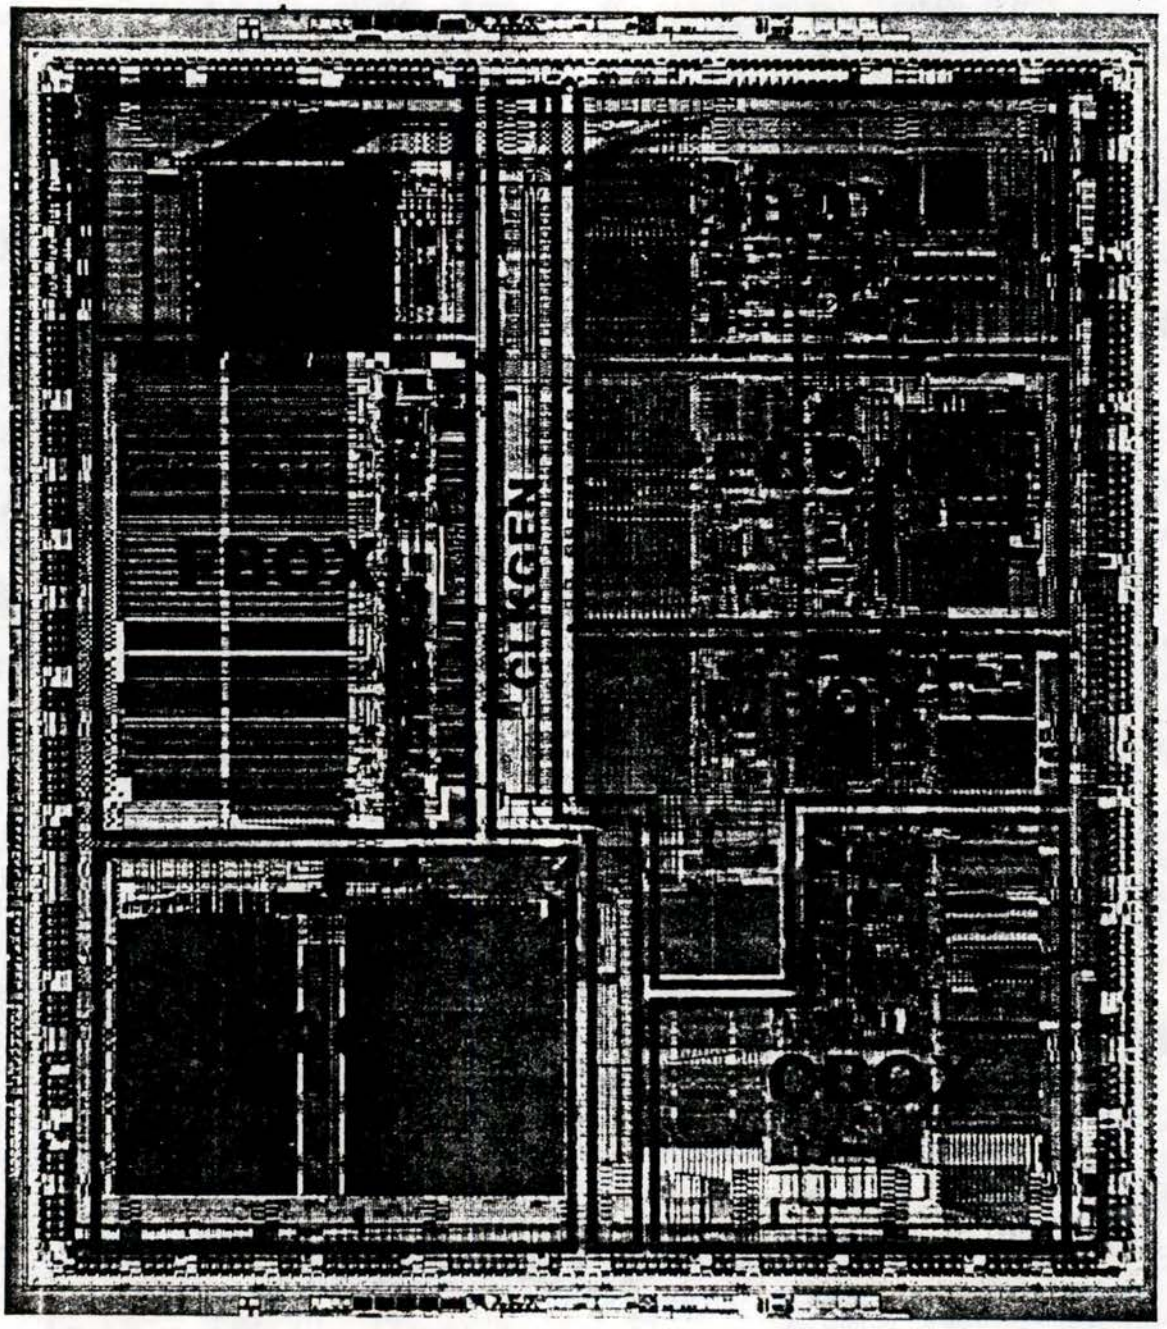

Figure 2: 100 MHz Macropipelined CISC CMOS Microprocessor

Authors: R. Badeau, et.al., Session: TA6.1

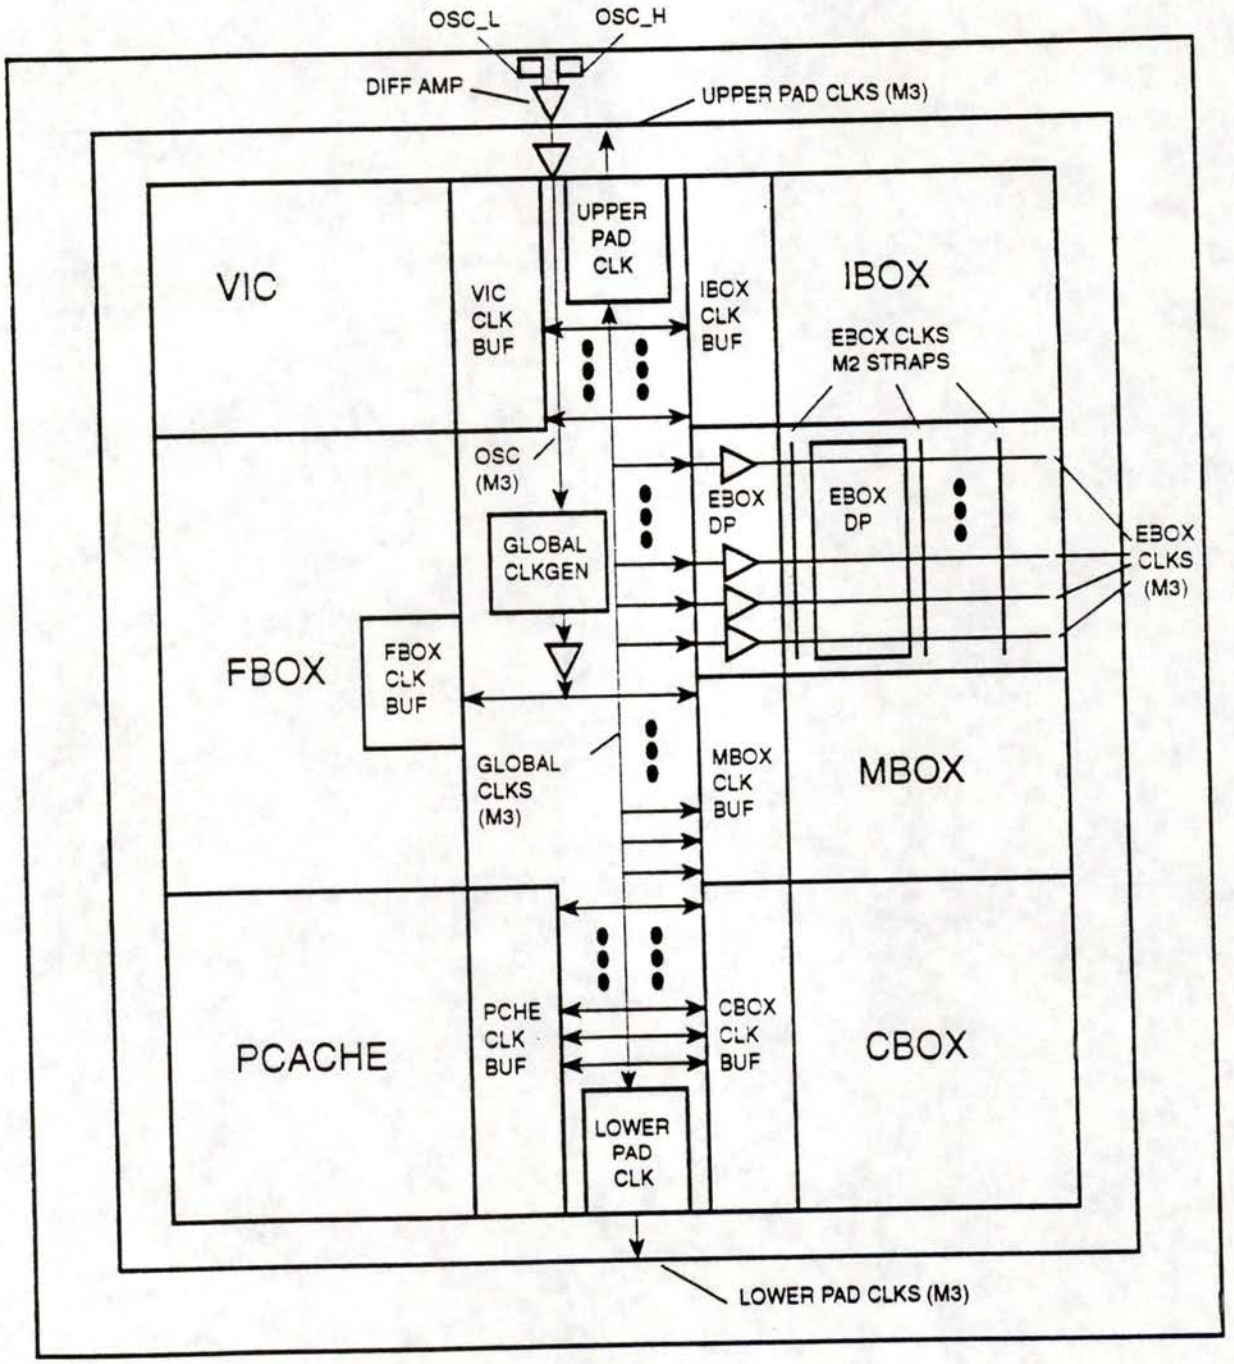

Figure 3 - Topological Block Diagram Of Clock Generation and Distribution Scheme.

Authors: R. Badeau, et.al., Session: TA6.1

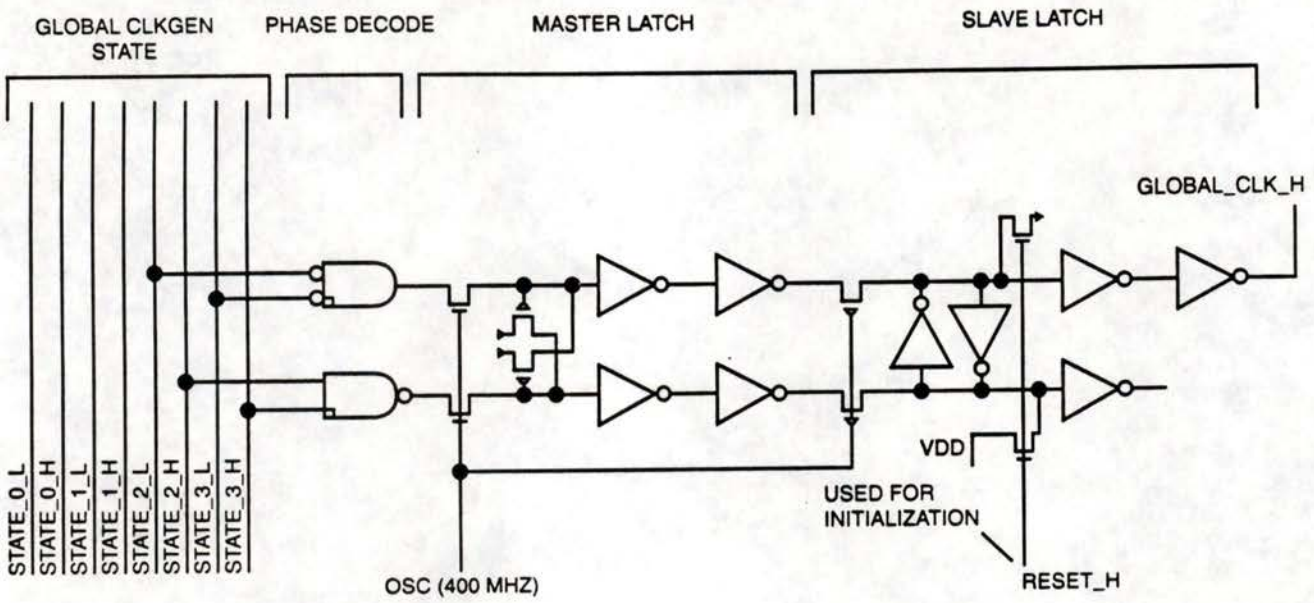

Figure 4 - Clock Phase Decode and Synchronizing Flip Flop.

Authors: R. Badeau, et.al., Session: TA6.1

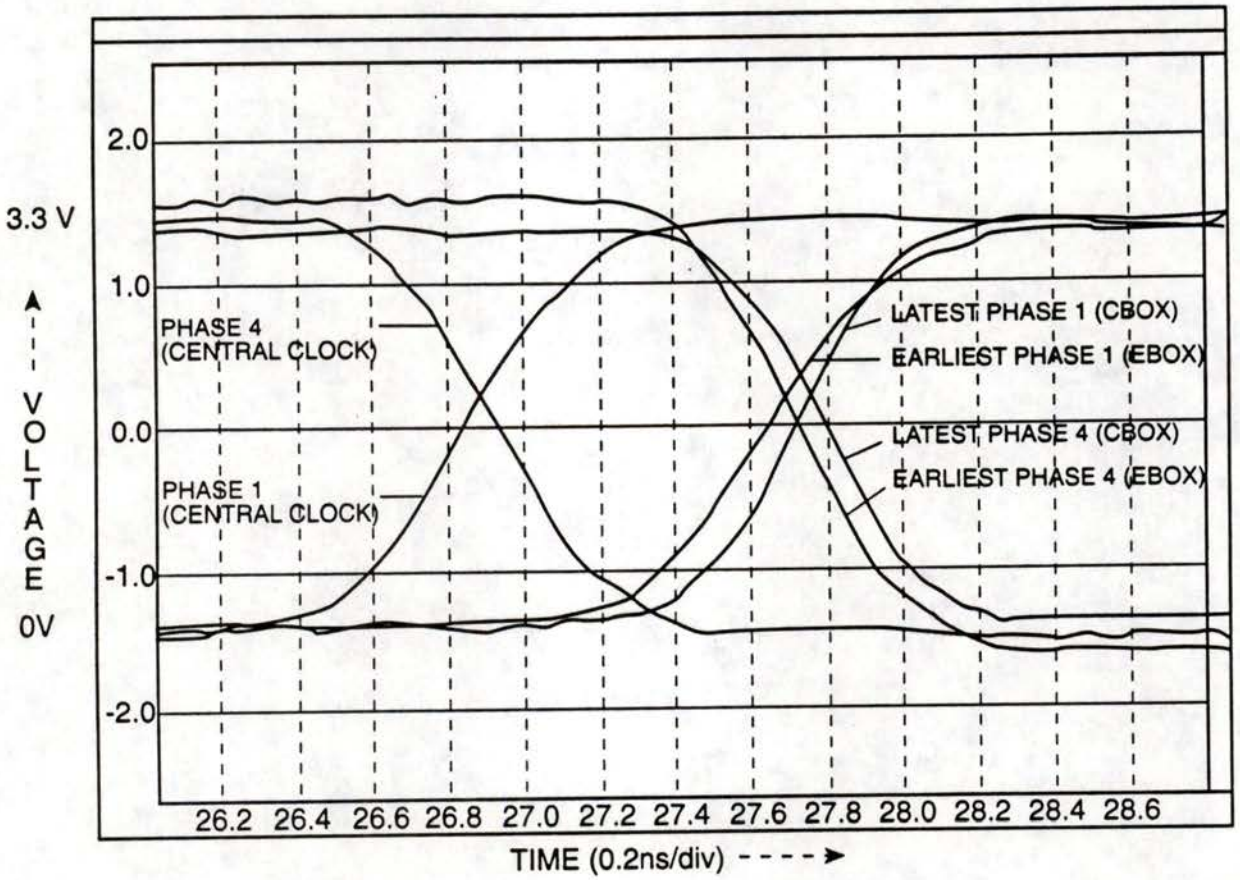

Figure 5: On-Chip Clock Waveforms From E-Beam Probe

Authors: R. Badeau, et.al., Session TA6.1

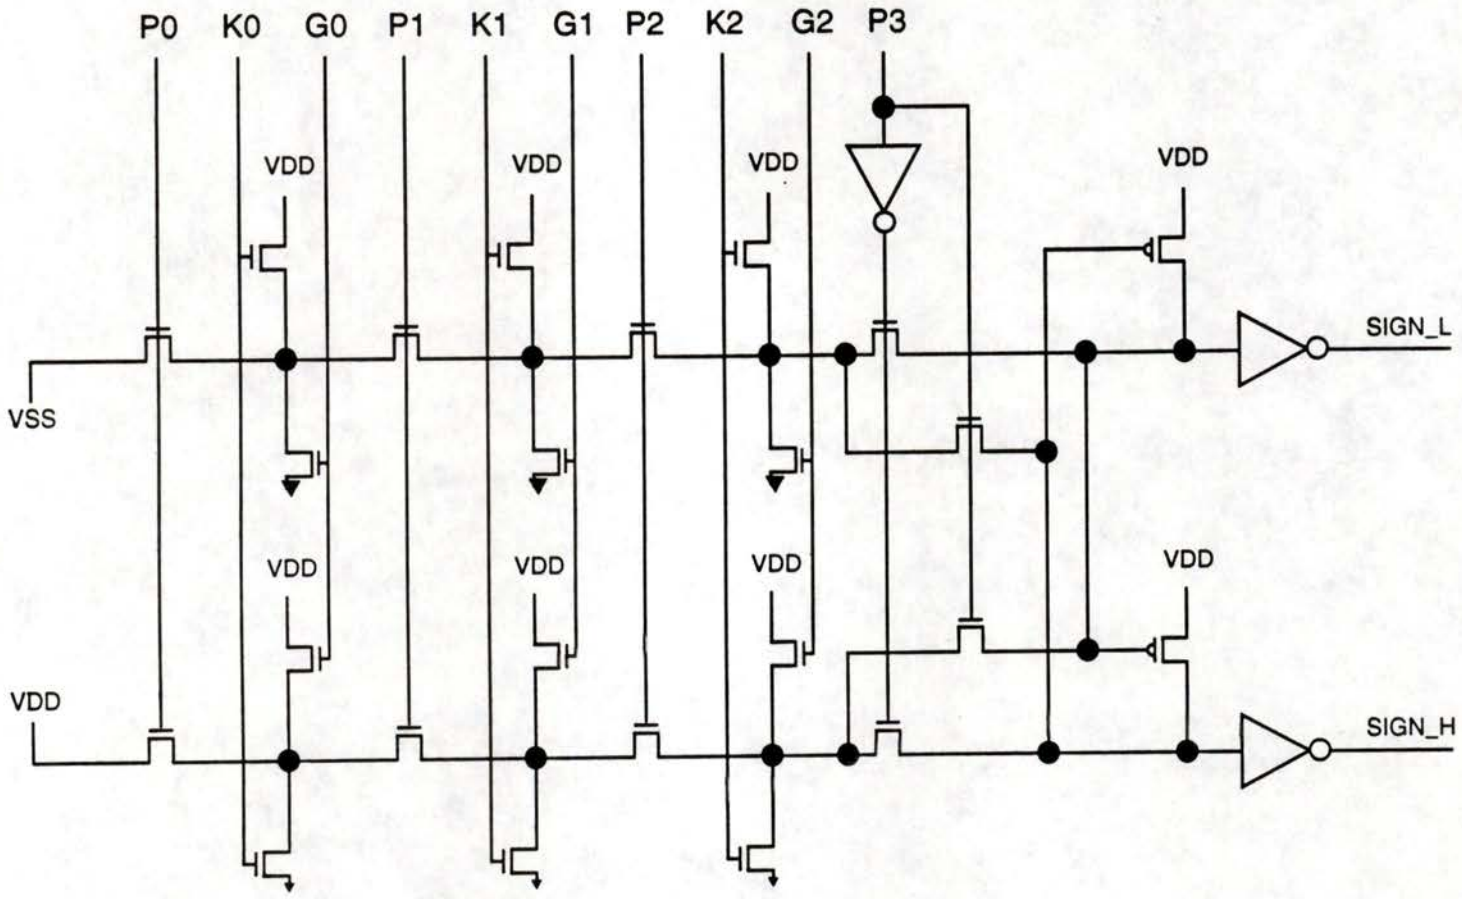

Figure 6: Estimated Partial Remainder Sign Detection Logic In FBOX Divider

Authors: R. Badeau, et.al., Session: TA6.1