



L7: Memory Basics and Timing



Acknowledgement: Nathan Ickes, Rex Min

J. Rabaey, A. Chandrakasan, B. Nikolic, "Digital Integrated Circuits: A Design Perspective" Prentice Hall, 2003 (Chapter 10)





Read-Wr	ite Memory	Non-Volatile Read-Write Memory	Read-Only Memory (ROM)			
Random Access	Non-Random Access	EPROM E ² PROM	Mask-Programmed			
SRAM DRAM	FIFO LIFO	FLASH				

Key Design Metrics:

- **1. Memory Density (number of bits/μm²) and Size**
- 2. Access Time (time to read or write) and Throughput
- 3. Power Dissipation













 Works fine for small memory blocks (e.g., small register files)
 Inefficient in area for large memories – density is the key metric in large memory circuits

How do we minimize cell size?

Static RAM (SRAM) Cell (The 6-T Cell)





- State held by cross-coupled inverters (M1-M4)
- Retains state as long as power supply turned on
- Feedback must be overdriven to write into the memory



Interacting with a Memory Device





- Address pins drive row and column decoders
- Data pins are bidirectional and shared by reads and writes
- Output Enable gates the chip's tristate driver
- Write Enable sets the memory's read/write mode
- Chip Enable/Chip Select acts as a "master switch"



MCM6264C 8k x 8 Static RAM



Same (bidirectional) data bus On the outside: used for reading and writing Chip Enables (E1 and E2) Address -□ E1 must be low and E2 must be high to enable the chip Chip Enables E1 E2 **MCM6264C** Write Enable (W) Data DQ[7:0] When low (and chip is enabled), Write Enable W the values on the data bus are Output Enable \overline{G} written to the location selected by the address bus **Output Enable (G)** When low (and chip is enabled On the inside: with W=0), the data bus is driven with the value of the selected A2 memory location DQ[7:0] A3 ecode $\overline{E1}$ 28 VCC NС 🛛 A4 Memory matrix E2 A12 🛛 2 27 W A5 256 rows $\overline{}$ Α7 Π E2 26 A7 ROW 32 Column A6 25 A8 A8 A5 A9 5 24 A9 A4 🛙 A11 W 23 П A11 Pinout G G A3 🛛 7 22 🛛 . . . 21 T A10 A2 [] 8 Sense Amps/Drivers A1 E1 Пg 20 🛛 Column Decoder A0 10 19 П DQ7 DQ0 E 11 18 🛛 DQ6 DQ1 12 DQ5 17 A0 A1 A6 A10 DQ2 13 DQ4 16 V_{SS} [] 14 15 DQ3



Reading an Asynchronous SRAM





E2 assumed high (enabled), W = 1 (read mode)

- Read cycle begins when all enable signals (E1, E2, G) are active
- Data is valid after read access time
 □ Access time is indicated by full part number: MCM6264CP-12 → 12ns
- Data bus is tristated shortly after G or E1 goes high







E2 assumed high (enabled), $\overline{W} = 1$ (read mode)

- Can perform multiple reads without disabling chip
- Data bus follows address bus, after some delay







E2 and \overline{G} are held high

Data latched when W or E1 goes high (or E2 goes low)

- Data must be stable at this time
- □ Address must be stable before W goes low

Write waveforms are more important than read waveforms Glitches to address can cause writes to random addresses!



Sample Memory Interface Logic













L7: 6.111 Spring 2005



Simulation from Previous Slide



📾 MAX+plus II - c:\documents and settings\anantha\my documents\6.111\verilog\lecture7\memtest - [memtest.scf - Waveform Editor]													
5	MAX+plus II File Edit Vie	ew Node A	ssign Utilities Option	s Window Hel	Þ								- 8 X
Þ.	Ref: 1.9us	+ +	Time: 444.0ns	In	nterval: -1.456us	6							^
Α													1.9us 💻
Æ	Name:	Value:	100.0ns 200.0n	s 300.0ns 400	0.0ns 500.0ns 60	0.0ns 700.0ns	s 800.0ns 900.0ns	: 1.Ous 1.1us	1.2us 1.3us	1.4us 1.5u	s 1.6us 1.7u	3 1.8us	1.9us 2.0
	📭 - clk	1											
	neset	1											
_	write 0 write completes												
<u>Q</u>	🗩 read	O				\sim							
	- W_b	1				X				read, a	ddress i	s stab	le
<u> </u>	– 🗊 G_b	1										\geq	
<u>ご</u> 1	📭 address	H 0009	0000	0001	0002	0003	0004	0005	(0006	(0008		0009
X	ext_address	H 0006	0000)(1	0001)		0006		
>>>> Z	📭 write_data[70]	H F9	FO	F1	F 2	F 3	(F4	F 5	(F6	(F7	(F8		F9
NY.	o∰ int_data	H F6	00	00 address/data stable Data					Data	latched into FPGA			
) <u>(</u>	🖃 data_oen	O				\mathbf{N}							
) <u>(</u>	💿 ext_data[70]	HZZ	ZZ)		F1)	ZZ		¢3) (ZZ	
<u>) (</u>	📭 ext_data[70]	H ZZ				ZZ) (d3) (ZZ	
XS	💿 read_data	н сз								X		<u>C3</u>	
	🖘 state	H6	0	X	1)	2	3	0	X	4)	5	6	X O
	🖃 data_sample	0		· · · · · · · · · · · · · · · · · · ·	write	states	1-3			rea	d states	1-3	~
	🖃 address_load	O											
	<												>



Verilog for Simple Multi-Cycle Access (conservative)



module memtest (clk, reset, G_b, W_b, address, ext_address, write_data, read_data, ext_data, read, write, state, data_oen, address_load, data_sample); input clk, reset, read, write; output G b, W b; output [12:0] ext address; reg [12:0] ext address; input [12:0] address; input [7:0] write_data; output [7:0] read_data; reg [7:0] read data; inout [7:0] ext data; reg [7:0] int_data; output [2:0] state; reg [2:0] state, next; output data oen, address load, data sample; reg G_b, W_b, G_b_int, W_b_int, address_load, data_oen, data_oen_int, data_sample; wire [7:0] ext_data; parameter IDLE = 0; parameter write1 = 1; parameter write2 = 2; parameter write3 = 3; parameter read1 = 4; parameter read2 = 5; parameter read3 = 6;

assign ext_data = data_oen ? int_data : 8'hz;

// Sequential always block for state assignment

always @ (posedge clk) begin if (!reset) state <= IDLE; else state <= next;

G_b <= G_b_int; W_b <= W_b_int; data_oen <= data_oen_int; if (address_load) ext_address <= address; if (data_sample) read_data <= ext_data; if (address_load) int_data <= write_data; end

// note that address_load and data_sample are not
// registered signals

1/4

2/4



Verilog for Simple Multi-Cycle Access



// Combinational always block for next-state // computation always @ (state or read or write) begin W b int = 1; G b int = 1; Setup the address load = 0; **Default values** data_oen_int = 0; data_sample = 0; case (state) IDLE: if (write) begin next = write1; address_load = 1; data_oen_int = 1; end else if (read) begin next = read1;address load = 1; **G_b_int = 0**; end else next = IDLE; write1: begin next = write2: W b int = 0; data_oen_int =1; end

```
write2: begin
         next = write3;
         data_oen_int =1;
         end
  write3: begin
         next = IDLE;
         data oen int = 0;
         end
  read1: begin
         next = read2;
         G_b_int = 0;
         data sample = 1;
         end
  read2: begin
         next = read3;
         end
  read3: begin
         next = IDLE;
         end
   default: next = IDLE:
  endcase
 end
endmodule
                                                   4/4
```

3/4





Common device problems

- Bad locations: rare for individual locations to be bad
- □ Slow (out-of-spec) timing(s): access incorrect data or violates setup/hold
- □ Catastrophic device failure: e.g., ESD
- □ Missing wire-bonds/devices (!): possible with automated assembly
- □ Transient Failures: Alpha particles, power supply glitch

Common board problems

- □ Stuck-at-Faults: a pin shorted to V_{DD} or GND
- Open Circuit Fault: connections unintentionally left out
- Open or shorted address wires: causes data to be written to incorrect locations
- Open or shorted control wires: generally renders memory completely inoperable

Approach

- Device problems generally affect the entire chip, almost any test will detect them
- Writing (and reading back) many different data patterns can detect data bus problems
- Writing unique data to every location and then reading it back can detect address bus problems





An idea that almost works

- 1. Write 0 to location 0
- 2. Read location 0, compare value read with 0
- 3. Write 1 to location 1
- 4. Read location 1, compare value read with 1
- 5. ...

What is the problem?

Suppose the memory was missing (or output enable was disconnected)



Data bus is undriven but wire capacitance briefly maintains the bus state: memory appears to be ok!



A Simple Memory Tester



- Write to all locations, then read back all locations
 - Separates read/write to the same location with reads/writes of different data to different locations
 - both data and address busses are changed between read and write to same location)

- Write 0 to address 0
- Write 1 to address 1
-
- Write (n mod 256) to address n
- Read address 0, compare with 0
- Read address 1, compare with 1
- Read address n, compare with (*n* mod 256)







 Clocking provides input synchronization and encourages more reliable operation at high speeds







- The wait state occurs because:
 - □ On a read, data is available *after* the clock edge
 - □ On a write, data is set up *before* the clock edge
- **ZBT (**"zero bus turnaround") memories change the rules for writes
 - On a write, data is set up after the clock edge (so that it is read on the following edge)
 - Result: no wait states, higher memory throughput







- Pipeline the memory by registering its output
 - □ Good: Greatly reduces CLK-Q delay, allows higher clock (more throughput)
 - □ Bad: Introduces an extra cycle before data is available (more latency)



As an example, see the CY7C147X ZBT Synchronous SRAM



EPROM Cell – The Floating Gate Transistor





Avalanche injection

Removing programming voltage leaves charge trapped

Programming results in higher V_T .

5 V

2.5 V

[Rabaey03]

5 V

D



EPROM Cell

S

Courtesy Intel

This is a non-volatile memory (retains state when supply turned off)





- Reading from flash or (E)EPROM is the same as reading from SRAM
- Vpp: input for programming voltage (12V)
 - □ EPROM: Vpp is supplied by programming machine
 - □ Modern flash/EEPROM devices generate 12V using an on-chip charge pump
- EPROM lacks a write enable
 - □ Not in-system programmable (must use a special programming machine)
- For flash and EEPROM, write sequence is controlled by an internal FSM
 - □ Writes to device are used to send signals to the FSM
 - Although the same signals are used, one can't write to flash/EEPROM in the same manner as SRAM





Dynamic RAM (DRAM) Cell





- DRAM relies on charge stored in a capacitor to hold state
- Found in all high density memories (one bit/transistor)
- Must be "refreshed" or state will be lost high overhead



Asynchronous DRAM Operation





 Clever manipulation of RAS and CAS after reads/writes provide more efficient modes: early-write, read-write, hidden-refresh, etc. (See datasheets for details)







Introductory Digital Systems Laboratory





SRAM vs. DRAM

- SRAM holds state as long as power supply is turned on. DRAM must be "refreshed" – results in more complicated control
- DRAM has much higher density, but requires special capacitor technology.
- FPGA usually implemented in a standard digital process technology and uses SRAM technology

Non-Volatile Memory

- Fast Read, but very slow write (EPROM must be removed from the system for programming!)
- □ Holds state even if the power supply is turned off

Memory Internals

Has quite a bit of analog circuits internally -- pay particular attention to noise and PCB board integration

Device details

□ Don't worry about them, wait until 6.012 or 6.374





- control signals such as Write Enable should be registered
- a multi-cycle read/write is safer from a timing perspective than the single cycle read/write approach
- it is a bad idea to enable two tri-states driving the bus at the same time
- an SRAM does not need to be "refreshed" while a DRAM does
- an EPROM/EEPROM/FLASH cell can hold its state even if the power supply is turned off
- a synchronous memory can result in higher throughput