



Montogue

Quiz EL306 Digital Memory

Lucas Monteiro Nogueira

► PROBLEMS

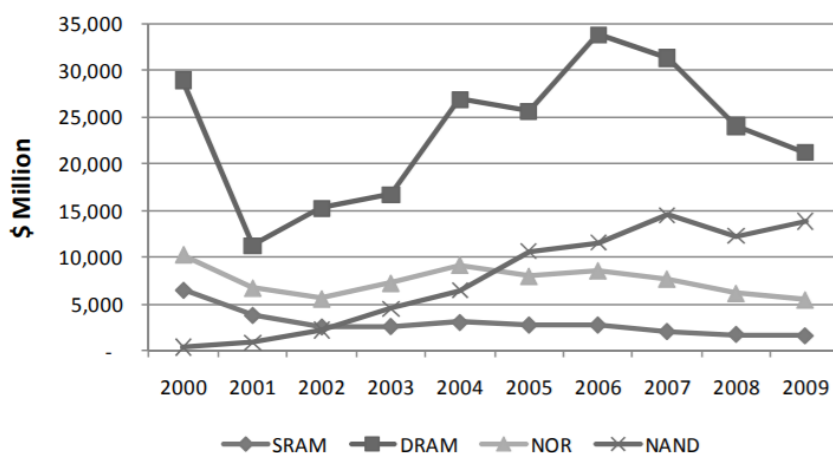
► Problem 1

The following statements refer to digital memory theory. True or false?

1. () Random access memories are non-volatile in nature.
2. () First-in, first-out (FIFO) memories are examples of random access memories.
3. () Static random access memories (SRAMs) are usually simpler in design than dynamic random access memories (DRAMs).
4. () Semiconductor random access memories can be built out of MOS or BJT transistors, but most modern commercial RAMs are based on MOSFET technology.
5. () The 7489 IC, a rather outdated digital component, is an example of RAM chip.
6. () The access time of a sequentially accessed memory is the same regardless of the position of the desired data in the storage device.
7. () EPROM memories are nonvolatile and can be erased by exposure to ultraviolet light. Importantly, a EPROM chip cannot be programmed a second time once its data is erased by the UV method.
8. () If a single read-only memory is used to design a combinational circuit described by a certain truth table, we can surmise that the number of address lines in the ROM is given by the number of input variables in the truth table. The number of output variables has no bearing on the number of address lines.
9. () Two-dimensional decoding, a celebrated technique of digital memory design, tends to lead to chip structures shaped like squares, which facilitates fabrication and packaging.

Statements 10 to 18 are somewhat more difficult to evaluate than the previous ones. Emphasis is placed on NAND Flash memory technology.

10. () NAND Flash memories are organized serially, while NOR Flash memories are organized in a parallel array architecture.
11. () The following chart shows the estimated global revenue from sales of four types of semiconductor memories in the 2000s. As can be seen, NOR memories have shown a steady growth in this period, whereas NAND memories have maintained a consistent revenue of \$ 5,000 to \$ 10,000 million dollars.



12.() Flash memory cells are mostly built upon one of two types of electron-based technologies: floating-gate cells or charge-trapping silicon oxide nitride oxide silicon (SONOS) cells. Floating-gate technology has been a prevalent non-volatile memory technology for some time, as it offers greater robustness against physical defects, less challenging production requirements, and easier scalability than SONOS cells.


13.() Decades after Gordon Moore of Intel posited that the number of transistors on a chip is to double every two years, Hwang Chang-gyu of Samsung electronics outlined an aggressive roadmap for his company that some members of the flash memory community came to call “Hwang’s law.” Specifically, Hwang established that the bit density of NAND flash memories has to double every year. ■ (A black square indicates the end of a statement.)

Unfortunately, scaled flash memory technology, be it NAND or NOR, is usually accompanied by an increased bit-error-rate (BER). One possible workaround to reduce the BER – and hence to contribute to the achievement of cutting-edge scaled technologies - is the massive introduction of Error-Correction Codes (ECCs). Two popular examples of ECCs are the Reed-Solomon and BCH codes.

14.() The BCH and Reed-Solomon codes have a very similar structure, but Reed-Solomon codes require fewer parity bits; this is one of the reasons why they are preferred for ECCs embedded in NAND memories. ■

One way to circumvent the issues encountered with planar NAND scaling without necessarily resorting to error-correction codes is to use three-dimensional architectures. Although several approaches have been used to fabricate a 3D NAND chip, the most effective solution employs a vertical channel with a horizontal gate, as illustrated below. Several architectures with vertical channels and horizontal gates have been proposed, such as BiCS, VRAT, and TCAT.

15.() Regardless of the architecture, much effort is required to develop the suitable integration modules in order to have at least a minimum of, say, 32 layers along the vertical dimension. In view of this and other reasons, as of 2021 no 3D-NAND-based SSD had made it to the mass market.



16.() The hard drive latency period of a hard disk drive (HDD) is defined as the time it takes a hard drive to spin up from 0 rpm to 3600 rpm.

17.() There are three main interface protocols used to connect solid-state drives (SSDs) into server or storage infrastructure: SAS, SATA, and PCI Express. By the mid-2010s, most commercial computers had hard drives based on PCI Express, but there is a growing trend towards devices designed with SATA; one reason is the inherently greater transfer speed of the latter. ■

One unorthodox random access memory concept is the ferroelectric RAM, or FeRAM. FeRAM memories are constituted of ferroelectric materials such as the perovskite PZT, which can be polarized spontaneously by an electric field and exhibit a residual permanent polarization that can be manipulated with external voltage to read and write data. Unfortunately, FeRAM concepts face a number of limitations that hinder their practical implementation, such as the inherently complicated integration with the standard CMOS process and the need to prevent hydrogen contamination of the ferroelectric material, which may reduce the permanent-polarization capability of the device.

18.() In spite of the many technical challenges, FeRAM is a potential contender for niche markets, especially very low-voltage embedded applications.

► **Problem 2** (Mano and Ciletti, 2004)

Word number 723 in a 1024×16 memory contains the binary equivalent of 3,451. List the 10-bit address and the 16-bit memory content of the word. Also write the address and memory content in hexadecimal form.

► **Problem 3** (Modified from Mano and Ciletti, 2004)

The memory units that follow are specified by the number of words times the number of bits per word. How many address lines are needed in each case? What is the number of bytes stored in each memory?*

Problem 3.1: $8K \times 16$

Problem 3.2: $2G \times 8$

Problem 3.3: $16M \times 32$

Problem 3.4: $256K \times 64$

Problem 3.5: What is the total **byte** capacity of a ROM that has 14 address lines and 4 data outputs?

*Recall that symbols K, M and G equal $2^{10}, 2^{20}$ and 2^{30} , respectively.

► **Problem 4** (Modified from Floyd, 2015, w/ permission)


Draw a basic logic diagram for a $2K \times 4$ -bit static RAM, showing all the inputs and outputs.

► **Problem 5** (Modified from Mano and Ciletti, 2004)

Show the memory cycle timing waveforms for the read and write operations. Assume a CPU clock of 100 MHz and a memory cycle time of 25 ns.


► **Problem 6** (Modified from Floyd, 2015, w/ permission)

For the ROM array illustrated below, determine the outputs for all possible input combinations, and summarize them in tabular form. A_i are inputs, and O_i are outputs. A blue cell is a 1, a gray cell is a 0.



► **Problem 7** (Modified from Floyd, 2015, w/ permission)

Determine the truth table for the ROM illustrated below. As in the previous problem, A_i are inputs and O_i are outputs. A blue cell is a 1, a gray cell is a 0.



▶ **Problem 8** (Mandal, 2010, w/ permission)

Design a 4-input diode ROM to convert BCD codes to excess-3 codes. The BCD-to-XS-3 conversions are provided below.

Decimal	BCD				XS-3			
0	0	0	0	0	0	0	1	1
1	0	0	0	1	0	1	0	0
2	0	0	1	0	0	1	0	1
3	0	0	1	1	0	1	1	0
4	0	1	0	0	0	1	1	1
5	0	1	0	1	1	0	0	0
6	0	1	1	0	1	0	0	1
7	0	1	1	1	1	0	1	0
8	1	0	0	0	1	0	1	1
9	1	0	0	1	1	1	0	0

▶ **Problem 9** (Mandal, 2010, w/ permission)

Design a 4-input diode ROM to convert binary codes to Gray codes. The BCD-to-Gray conversions are provided below.

Decimal	Binary				Gray			
0	0	0	0	0	0	0	0	0
1	0	0	0	1	0	0	0	1
2	0	0	1	0	0	0	1	1
3	0	0	1	1	0	0	1	0
4	0	1	0	0	0	1	1	0
5	0	1	0	1	0	1	1	1
6	0	1	1	0	0	1	0	1
7	0	1	1	1	0	1	0	0
8	1	0	0	0	1	1	0	0
9	1	0	0	1	1	1	0	1
10	1	0	1	0	1	1	1	1
11	1	0	1	1	1	1	1	0
12	1	1	0	0	1	0	1	0
13	1	1	0	1	1	0	1	1
14	1	1	1	0	1	0	0	1
15	1	1	1	1	1	0	0	0

▶ **Problem 10** (Modified from Mano and Ciletti, 2004)

A $16K \times 4$ memory uses coincident decoding by splitting the internal decoder into X-selection and Y-selection.

Problem 10.1: Name one advantage of using coincident decoding in lieu of memory decoding with a single decoder.

Problem 10.2: What is the size of each decoder, and how many AND gates are required for decoding the address?

Problem 10.3: Determine the X and Y selection lines that are enabled when the input address is the binary equivalent of 6,000.

▶ **Problem 11**

A dynamic RAM (DRAM) chip uses two-dimensional address multiplexing. It has 13 common address pins, with the row address having one bit more than the column address. What is the capacity of the memory?

▶ **SOLUTIONS**

P.1 → **Solution**

1. False. Random access memories are volatile. Read-only memories are examples of non-volatile memory.

2. False. FIFO memories are actually sequential memories.

3. True. Mano and Ciletti (see reference below) observe that a typical SRAM cell consists of six transistors, whereas a DRAM cell is constructed with a single MOS transistor and a capacitor.

Reference: Mano and Ciletti (2004).

4. True. Modern SRAMs are generally manufactured with MOSFET-based flip-flops, while modern DRAMs are typically structured with MOS capacitors.

5. True. The 7489 IC is a 64-bit RAM chip organized as a 16-word by 4-bit array.

6. False. Independence of position is a feature of random access memories, not sequentially accessed memories.

7. False. Erasable programmable read-only memory (EPROM) resists memory deletion by electrical means but may have its content erased with UV light. Usually, EPROM chips are reprogrammable after their data is deleted by this method.

8. False. In actuality, the number of address lines in the ROM is given by the sum of the number of input variables and the number of output variables in the truth table.

9. True. This one is borrowed from Wakerly (see reference below).
Reference: Wakerly (2006).

10. True. NOR memories are organized in a parallel array architecture in which each cell may be accessed via a contact; this provides a superior random performance relatively to NAND Flash chips, which are organized serially and offer rather slow random performance because there are no contacts directly accessing the memory cells. On the other hand, NAND Flash memories have much lower effective chip size and lower cost per bit because their architecture involves only two contacts for every 32 memory cells.

11. False. This simple observational exercise requires only inspection of the chart: the curve for NAND memories shows a steady increase from 2000 to 2009, whereas the one for NOR memories is kept circumscribed by the \$5,000 - \$10,000 interval. Importantly, the chart also shows that dynamic RAMs comfortably dominate semiconductor memory sales, but began to decline at around 2006.

12. False. Richter (see reference below) notes that, while floating-gate tech has been the most successful non-volatile memory technology in the past few decades, SONOS cells offer advantages in some respects including a greater robustness against physical defects, less challenging production requirements, and easier scalability.

Reference: Richter (2014).

13. True. Indeed, Hwang believed that Samsung could double bit density in its NAND Flash memories every year, and that such developments would rapidly enable solid-state drive solutions to supersede HDD as the storage technology of choice in consumer electronics.

14. False. While it is true that BCH and Reed-Solomon codes have similar structures, it is BCH codes that require fewer parity bits and hence happen to be particularly useful in NAND memories. Pirovano (see reference below) notes that while the encoding of an ECC takes few cycles of latency, the decoding phase can require more cycles and visibly reduce read performance, as well as the memory response time at random access. Moreover, ECC requires additional cells to encode the logical data, thus using part of the memory chip just to improve the overall reliability.

Reference: Pirovano, in Gastaldi and Campardo (2017).

15. False. Micron and Intel unveiled a commercial 3D NAND all the way back in 2015. Their model was based on floating-gate tech, but charge-trap 3D-NAND SSDs are already available as well.

16. False. The hard drive latency period is actually the time that it takes for the required sector to spin under the head once the head is positioned.

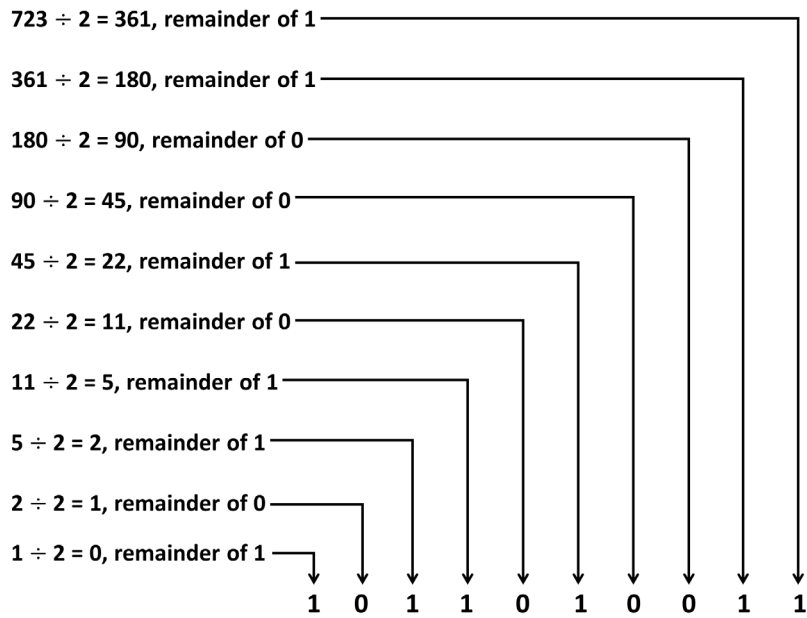
17. False. Swap “SATA” with “PCIe” in the last sentence and you’ll obtain a true statement: it is PCIe, not SATA, that offers superior transfer speed. For this and other reasons, PCIe-based devices are bound to supplant SATA-based ones in most applications.

18. True. As pointed out by Pirovano (see reference below), the voltage required to switch the permanent polarization in a suitable ferroelectric material is in the range of 1.5 – 3 V for typical deposited-layer thicknesses of 70 – 100 nm. It follows that ferroelectric memories can be a valuable solution for low-power and very low-voltage applications, like a battery-operated embedded system, smartcards, and RFID applications.

Reference: Pirovano, in Gastaldi and Campardo (2017).

P.2 → **Solution**

To compute the 10-bit address, we convert 723 to binary form using the divide-by-2 method.



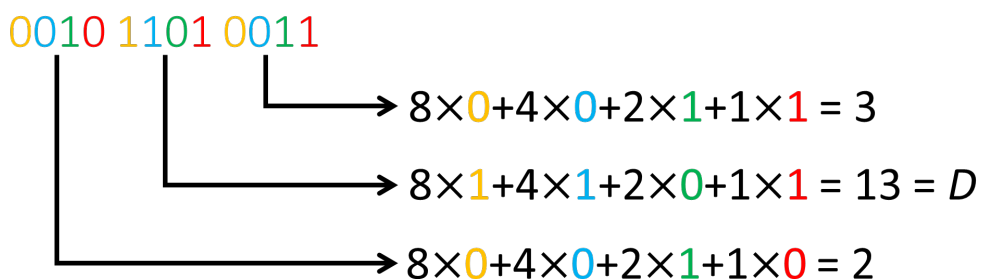
Thus, the 10-bit address of the word is 1011010011. To find the 16-bit memory content of the word, we first convert 3,451 to binary form, giving 11010111011; these are 12 digits, so we add 4 zero bits to obtain the correct form of the word in question. The results are summarized below.

Address: 10 1101 0011
Memory content: 0000 1101 0111 1011


Most readers should have already mastered conversion of binary numbers to hexadecimal form at this point in their digital principles course, but we nonetheless provide a quick review of the process. First, we write the binary number in groups of 4 digits:

Address: 0010 1101 0011

Next, we take each group and multiply the rightmost digit by 2^0 (or 1), the second-from-the-right digit by 2^1 (or 2), the third-from-the-right digit by 2^2 (or 4), and the leftmost digit by 2^3 (or 8). Then, we add the results and repeat the procedure for all groups of four, as shown. Note that the conversion in the middle resulted in decimal 13, which in hexadecimal notation is usually denoted by the letter *D*.



Lastly, we order the results such that the result obtained from conversion of the rightmost group of 4 digits will be the rightmost digit in the hexadecimal representation, while the result obtained from the leftmost group will be the leftmost digit. The number thus obtained is the hexadecimal form of the address.



Proceed similarly with the binary number that represents memory content and you'll find $(3,451)_{10} = (D7B)_{16}$. The results are summarized below.

	Decimal	Binary	Hex.
Address	723	10 1101 0011	2D3
Memory content	3,451	0000 1101 0111 1011	D7B

P.3 → Solution

Problem 3.1: This memory has 8K words, which amounts to $8 \times 2^{10} = 2^{13}$ words; recalling that a memory with 2^k words requires k address lines, the memory in question requires 13 address lines. Further, each word has 16 bits, so the memory contains $2^{13} \times 16$ bits. Since one byte corresponds to 8 bits, the memory stores a total of $(2^{13} \times 16)/8 = 16.4$ kilobytes.

Problem 3.2: This memory has 2G words, or $2 \times 2^{30} = 2^{31}$ words; accordingly, 31 address lines are needed. Each word has 8 bits, so the memory contains $2^{31} \times 8$ bits. Knowing that one byte corresponds to 8 bits, the memory stores a total of $(2^{31} \times 8)/8 = 2.15$ gigabytes.

Problem 3.3: This memory has 16M words, or $16 \times 2^{20} = 2^{24}$ words; accordingly, 24 address lines are needed. Each word has 32 bits, so the memory contains $2^{24} \times 32$ bits. Knowing that one byte corresponds to 8 bits, the memory stores a total of $(2^{24} \times 32)/8 = 67.1$ megabytes.

Problem 3.4: This memory has 256K words, or $256 \times 2^{10} = 2^{18}$ words; accordingly, 18 address lines are needed. Each word has 64 bits, so the memory stores a total of $(2^{18} \times 64)/8 = 2.10$ megabytes.

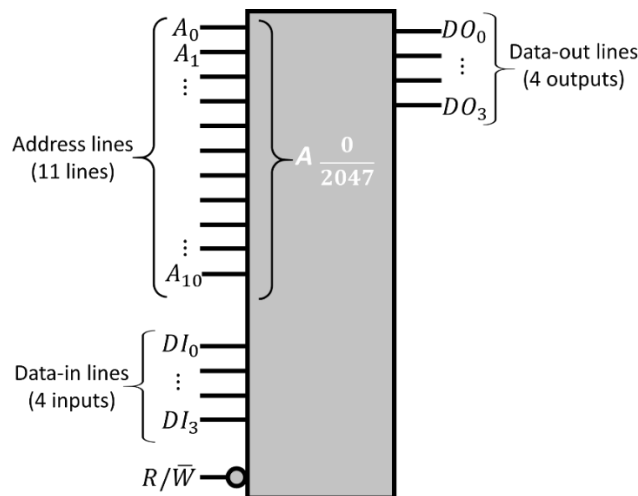
The results are summarized below. For convenience, we've rounded the storage to the nearest power of 2.

Memory	No. Address Lines	Storage
3.1	13	16 KB
3.2	31	2 GB
3.3	34	64 MB
3.4	18	2 MB

Problem 3.5: If the ROM has 14 address lines, it can hold up $2^{14} = 16,384$ addresses. Since there are 4 data outputs, the memory will store $16,384 \times 4 = 524,288$ bits or, equivalently, 65,536 bytes.

P.4 → Solution

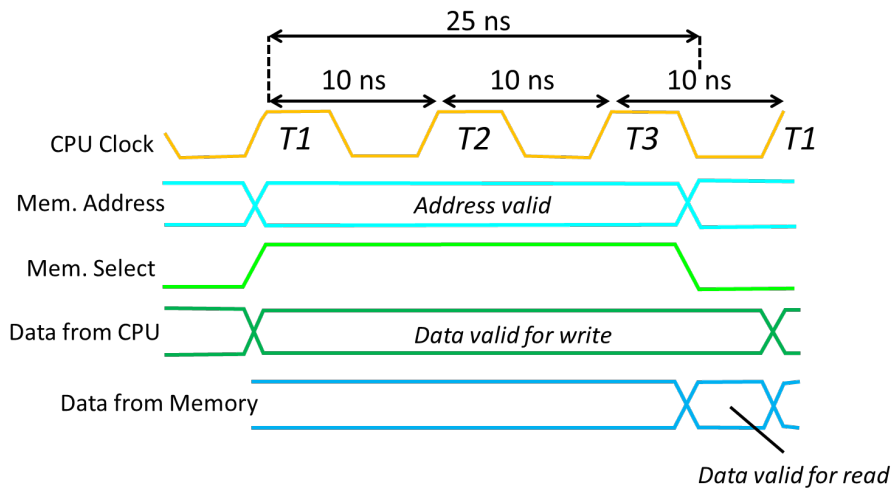
Noting that $1\text{ K} = 2^{10}$, a 2K static RAM can store 2^{11} words and has 11 address lines. Because this is a 4-bit memory, there are 4 data inputs and 4 data outputs. Needless to say, there may be other inputs, such as a chip select signal, an output enable signal, and a toggle between READ and WRITE modes. Further, data input and output channels may share the same lines. We keep things simple and draw a SRAM diagram with address lines, data-in, data-out, and READ/WRITE wires only, as shown.



P.5 → Solution

The solution is started by noting that a CPU clocked at 100 MHz operates at periods of $1/(100 \times 10^6) = 10$ ns. Since the period of the memory cycle time is 25 ns, the device will devote at least two-and-a-half, or possibly three, clock cycles for each memory request. The timing waveforms shown below shows three 10-ns cycles labeled T_1, T_2 , and T_3 . For a write operation, the CPU must provide the access and input data to the memory. This is done at the beginning of T_1 . (The two lines that cross each other in the address and data waveforms designate a possible change in value of the multiple lines.) The memory enable and the read/write signals must be activated after the signals in the address lines are stable in order to avoid destroying data in other memory words. The memory enable signal switches to the high level and the read/write signal switches to the low level to indicate a write operation. The two control signals must stay active for at least 25 ns. The

address and data signals must remain stable for a short time after the control signals are deactivated. At the completion of the third cycle, the memory write operation is completed and the CPU can access the memory again with the next $T1$ cycle. Sketching waveforms for the read cycle follows similar reasoning.



P.6 → Solution

For two inputs, A_0 and A_1 , the decoder may receive signals 00 (equivalent to decimal 0), 01 (decimal 1), 10 (decimal 2), and 11 (decimal 3). The memory diagram provided indicates the output combinations afforded for each of these inputs; for instance, entering 00 yields $O_3 = 0, O_2 = 1, O_1 = 0$, and $O_0 = 1$. We proceed to draw up the following table.

Inputs		Outputs			
A_1	A_0	O_3	O_2	O_1	O_0
0	0	0	1	0	1
0	1	1	0	0	1
1	0	1	1	1	0
1	1	0	0	1	0


P.7 → Solution

Aside from the greater number of inputs, this ROM model is similar to the one analyzed in Problem 6. Input possibilities range from 000 (decimal 0) to 111 (decimal 7).

Inputs			Outputs			
A_2	A_1	A_0	O_3	O_2	O_1	O_0
0	0	0	0	1	0	1
0	0	1	1	1	1	1
0	1	0	1	0	1	1
0	1	1	1	0	0	1
1	0	0	1	1	1	0
1	0	1	1	0	0	0
1	1	0	0	0	1	1
1	1	1	0	1	0	1


P.8 → Solution

The ROM is illustrated below. Inputs A, B, C , and D receive a BCD number, and outputs D_0 to D_3 yield the corresponding XS-3 code. Diodes are placed in nodes set to return HIGH voltage levels.



P.9 → Solution

The ROM is illustrated below. Inputs $A, B, C,$ and D receive a binary number, and outputs D_0 to D_3 yield the corresponding Gray code. Diodes are placed in nodes set to return HIGH voltage levels.




P.10 → Solution

Problem 10.1: Coincident decoding allows for architectures with less AND gates and, consequently, more economical designs.

Problem 10.2: 16K bits amounts to $16 \times 2^{10} = 2^{14}$ bits; distributing data processing evenly, each decoder will carry $2^7 = 128$ bits. Each decoder will have 7 inputs. Each decoder has 128 AND gates, totaling 256 AND gates for two decoders.

Problem 10.3: Converting 6,000 to binary form, we obtain 1011101110000. Adding one extra zero to turn this into a 14-bit string, we write 01011101110000. This input is to be partitioned into two words, each of which is to be assigned to one of the two decoders. The address lines enabled upon entering this data are found to be 46 and 112, as shown.



P.11 → Solution

In accordance with the problem statement, the row address has 13 address lines, while the column address has 12 address lines. Accordingly, the capacity of the memory is 2^{25} words.

▶ REFERENCES

- FLOYD, T.L. (2015). *Digital Fundamentals*. 11th edition. Upper Saddle River: Pearson.
- GASTALDI, G. And CAMPARDO, G. (Eds.). (2017). *In Search for the Next Memory*. Berlin/Heidelberg: Springer.
- MANDAL, S. K. (2010). *Digital Electronics: Principles and Applications*. New Delhi: Tata-McGraw-Hill.
- MANO, M.M. and CILETTI, M.D. (2004). *Digital Design*. 4th edition. Upper Saddle River: Pearson Prentice-Hall.
- RICHTER, D. (2014). *Flash Memories*. Berlin/Heidelberg: Springer.
- WAKERLY, J.F. (2006). *Digital Design: Principles and Practices*. 4th edition. Upper Saddle River: Pearson Prentice-Hall.



Was this material helpful to you? If so, please consider donating a small amount to our project at www.montoguequiz.com/donate so we can keep posting free, high-quality materials like this one on a regular basis.