Name: _____

Directions: Work only on this sheet (on both sides, if needed); do not turn in any supplementary sheets of paper. There is actually plenty of room for your answers, as long as you organize yourself BEFORE starting writing. In order to get full credit, SHOW YOUR WORK.

1. (5) Consider Fig. 16.11, altered to be a 256Mx64 system, but still using 16Mx16 chips, and still high-order interleaved. Then address lines A_____ through A_____, inclusive, would be the inputs to a decoder which would have _____ outputs.

2. (5) Name two pairs of handshaking signals (that means four signals in all) discussed in class.

3. This question concerns designing a virtual memory system, specifically the choice of the page size.

(a) (5) The page size is set by (i) the hardware designer, (ii) the writer of the OS, (iii) the writer of the compiler, (iv) all of the above.

(b) (5) It is natural to set the page size to some multiple of the size of _____ on disk.

4. (10) The instruction **outb** has two operands. If we are using it to send a command to an 8259A chip on a PC, a likely value for the second operand is $0x_____.

5. Consider a circuit with inputs A_i, i = 0,1,2,3 and output X_j, j = 0,1, in which X_1X_0 tells us the largest index i for which A_i is 1. For example, if A_3 is 0 but A_2 is 1, then the X output codes 2, i.e. X_1 is 1 and X_0 is 0. If A_3 is 1, then the X output codes 3, i.e. X_1X_0 is 11. If the only 1 among the A_i is A_0, then X_1X_0 is 00. We are guaranteed that there always will be at least one A_i which is 1.

(a) (10) Show the K-maps for X_1 and X_0. Each of these two K-maps MUST be of the form

A_3, A_2 \ A_1,A_0	00	01	11	10
00	-	-	-	-
01	-	-	-	-
11	-	-	-	-
10	-	-	-	-

i.e. the order of the variables must be as prescribed here.

(b) (10) Use the K-maps to find minimal sum-of-product expressions for X_1 and X_0. *Each SOP must be a sum of only two terms.*

6. (10) Consider what can happen when there is a cache miss on a Pentium machine, in terms of reads from and writes to memory. To the nearest 1,000, in the worst case _____ thousand bytes will be read from memory, and _____ thousand bytes will be written to memory.

7. Consider the PowerPC pseudo-LRU cache line replacement policy, summarized in Fig. 17.21.

(a) (10) Suppose the sequence of lines accessed is L1, L0, L4, L0, L1, L5, L6, and then there is a miss. Then the line replaced will be one of _____.

(b) (10) Show the analog of Table 17.7 for the four-way associative case.

8. (10) Consider a machine with address size n, l cache lines, b bytes per block, with a direct-mapped cache. Then _____ blocks would map to the same line.

9. (10) This problem concerns the error correction example on pp.786-7. The Hamming distance for this code is _____. A pair of legal codewords which achieves this minimum value are _____ and _____.

Solutions:

1. The word size is 64 bits, i.e. 8 bytes, so A0-A2 are still used to specify byte within word.

Since we are still using chips with 16M words each, we will still have 16M system words per row. So, we still need 24 bits to specify the word within row. These bits will be A26-A3.

There will be 256M/16M = 16 rows, so we need 4 bits to specify row. These bits will be A30-A27. They will feed into the decoder, which will have 16 outputs.

2. HOLD, HLDA; DREQ, DACK

3.a i

3.b sectors (tracks OK too)

4. 20 or 21

5.a

X_1:

A_3, A_2 \ A_1,A_0	00	01	11	10
00	d	0	0	0
01	1	1	1	1
11	1	1	1	1
10	1	1	1	1

Take d = 0. The expression simplifies to $A_3 + A_2$.

5.b

X_0:

A_3, A_2 \ A_1,A_0	00	01	11	10
00	d	0	1	1
01	0	0	0	0
11	1	1	1	1
10	1	1	1	1

Again take d = 0. The expression simplifies to $A_3 + \overline{A_2}A_1$.

6. The desired block is contained in some page. (It could not straddle two pages.) We need to read that page.

The page table table (table pointing to page tables) is still resident, but the page table for our process might not be resident. In the latter case, the page table must be read from disk (a 4K write to memory), and if the previously-resident page to be replaced is dirty, it must be written to disk (a 4K read from memory).

Finally, we read the page table for our process to see where the desired block is. The page table might tell us that that page is not resident; if so, that would again result in a 4K write to memory, and a 4K read from memory if the replaced page is dirty.

So, as much as 8K would be read from memory, and as much as 8K would be written to memory.

7.a L2, L3

7.b

access	B0	B1	B2
L0	1	1	NC
L1	1	0	NC
L2	0	NC	1
L3	0	NC	0

8. Of the n bits, $\log_2 b$ bits specify the byte within block and $\log_2 l$ bits specify the cache line. So, $n - \log_2 b - \log_2 l$ bits form the tag. Thus, there are

$$2^{n-\log_2 b-\log_2 l} = \frac{2^n}{bl}$$

different blocks that map to the same line.

9. Suppose two data strings (that's data, not the whole codeword) differ only in D0. Then their codewords will differ also in P4 and P8, and thus will be a distance 3 from each other. This is the closest case. So, the Hamming distance for the code is 3.

To construct an example in which this distance is attained, use this same reasoning. Take the first data string to consist of eight 0s, and the second to consist of seven 0s and one 1, the 1 being in D0. Then the codeword for the first consists of 12 0s, while the second will have 1s in P4, P8 and D0, with 0s elsewhere.