HW7 Solutions

Question 1

Since virtual addresses are 11 bits, and since each address is the address of a byte, the virtual address space is 2^11 bytes, or 2KB, or 2048 bytes.

Question 2

1. VA = [20 bits | 40 bits], so VPN is 20 bits.
2. PA = [12 bits | 40 bits], so PPN is 12 bits
3. offset is 40 bits (because the page size is 2^40 bytes)

Question 3

1. Sometimes. If the page is not present, then a TLB miss and a page fault will happen at the same time. If the page is not writable but it is in the TLB, then a store to memory can cause a page fault but no TLB miss. (And if the page is present and the permissions are consistent with the requested operation, yet the VA,PA mapping is not in the TLB, then there will be a TLB miss but no page fault.)
2. Sometimes. If the page is not present, or the permission check fails, then there will be a page fault on top of the TLB miss.
3. Sometimes. Similar answer to above.
4. Always. PTE_P | PTE_U means user readable.
5. Sometimes. If PTE_W is set, then it is permissible, otherwise not.

Question 4

1. Five. One for code loading. Two for data page 0x200000, and two for data page 0x300000. Why does each data page generate two TLB misses? When virtual address 0x200000 (for example) is referenced, there is a TLB miss (because the page was not present in virtual memory and so would not be in the TLB); then, when the page actually swaps into memory and the instruction is retried, there is another TLB miss. See section 21.5 in the book for more details.
2. Two page faults