Solution:

1. For P, clock period = 2.5 units. For U, clock period = 6 units.

   (a) The best-case speedup occurs when the pipeline is constantly fed with data. In this case, P produces one output every 2.5 units, while U produces a new output every 6 units. The speedup is therefore 6/2.5 = 2.4.

   (b) For the "fast" operations, the pipeline requires five cycles, and for the "slow" operations, four cycles are necessary. In both cases, the cycle time is dominated by stage 2 and is set to 1.5 units. Each "slow" operation inserts a single bubble into the pipeline. Therefore, the average number of bubbles inserted into the pipeline per cycle is 0.2.

   If a sufficiently large volume of data fed to the pipeline (>> the number of pipeline stages), a new output is produced every 1.2 cycles = 1.8 time units. The speedup over U is 6/1.8 = 3.33

   (c) Following the reasoning in (a), the speedup is now 6/1.5 =4, which is theoretically the best possible since all stages are perfectly balanced.

2. The best speedup we can get is slightly below 4, which is the case where 75% of the n operations are parallelized on m processors, and there is no overhead to doing so, so that their execution time is

Lim_{m} ▮▮▮$(0.75 \times n)/m = 0$, and the 25% of the serial operations take time 0.25 x n.

Therefore, a speedup of 6 is not possible, as a consequence of Amdahl's law.

3. 1 word = 32 bits = 4 bytes
Main memory address: $\log(1G/4) = 28$ bits
Cache address: $\log(512K/4) = 17$ bits
Number of blocks in the cache = $512K/(4*64) = 2048$

   i) *Fully associative mapping*
       The block of 64 words is mapped to one of the 2048 available blocks in the cache. The MM address is mapped as follows

Tag = 22 bits	Word # in block = 6 bits

   The tag is matched with that of each cache block to check for a hit.

ii) Direct mapping

The block # in the cache is taken from main memory address as follows

Tag = 11 bits	Cache blk # = 11 bits	Word # in block = 6 bits

The cache block number is determined from the corresponding bits, and the tag is matched with that of the corresponding cache block to check for a hit.

iii) *4-way set-associative mapping*

Number of cache sets = # cache blocks/4 = 512

Tag = 13 bits	Cache set # = 9 bits	Word # in block = 6 bits

The cache set number is determined from the corresponding bits, and the tag is matched with each block in the corresponding cache set to check for a hit.