CPEG 222 Spring 2014 HW2 Solution

Bo Lu

April 29, 2014

Solution of Question 1:

(a) $B[I][0]$ and J exhibit temporal locality since they are keeping reused in the inner loop. $A[I][J]$ exhibits spatial locality since the contiguous elements in memory are used one by one.

(b) $B[I][0]$ and I exhibit temporal locality since they are keeping reused in the inner loop. J also exhibits temporal locality since the inner loop is small and J is keeping reused in the outer loop. $A[J][I]$ now exhibits spatial locality since the contiguous elements in memory are used one by one.

Solution of Question 2:

	Flash drive	Hard disk
Latency	Almost instantaneous	Several seconds
Throughput	High on read Slow on write	Slow on read High on write
Energy for non-idle	Low	High
Standby energy	Low	High
Capacity	2 TB	6 TB
Cost	\$0.45/GB	\$0.05/GB

Solution of Question 3:

Amount of time CPU takes to process each 64KB block
$=$ (number of cycles to process each block)/(clock rate)$= (2 \times 10^6)/(3 \times 10^9) \times 10^3 = 0.67ms$
Amount of time spent in memory transfer for 64KB block
$=$ (size of the block)/(rate of memory bus)$= (64 \times 2^{10})/(640 \times 2^{20}) \times 10^3 = 0.097ms$
Amount of time spent in I/O transfer for 64KB block
$=$ seek time + rotational delay + transfer time + controller overhead
$= 9 + (64 \times 2^{10})/(64 \times 2^{20}) \times 10^3 + (10^6)/(3 \times 10^9) \times 10^3 = 10.31ms$
Therefore, the main bottleneck is I/O in the above system.

Solution of Question 4:

(a) Hit time $= 1$ cycle. So the clock rate for P1 is $1/(10^{-9}) = 1GHz$.
The clock rate for P2 is $1/(1.8 \times 10^{-9}) = 0.6GHz$.

(b) We first compute the average memory access cycle (AMAC), and then AMAT.
$AMAC_1 = 1.0 + 0.036 \times \lceil 200/1 \rceil = 8.2$ cycles, $AMAT_1 = AMAC_1 \times 1.0 = 8.2ns$
$AMAC_2 = 1.0 + 0.028 \times \lceil 200/1.8 \rceil = 4.14$ cycles, $AMAT_2 = AMAC_2 \times 1.8 = 7.44ns$

(c) Only 30% of instructions need to access the data cache. For the other 70% instructions, their CPI is 1. Therefore we have:
$CPI_1 = 0.7 \times 1 + 0.3 \times 8.2 = 3.16$ cycles, therefore $TPI_1 = 3.16 \times 1.0 = 3.16ns$
$CPI_2 = 0.7 \times 1 + 0.3 \times 4.14 = 1.94$ cycles, therefore $TPI_2 = 1.94 \times 1.8 = 3.49ns$
Since $TPI_1 < TPI_2$, processor 1 is faster.

Solution of Question 5:
Most memory accesses are sequential in their addresses. This implies that the higher bits of these addresses are identical. If the higher bits are used as the index, then contiguous cache blocks will be placed in the same set. Therefore, for a program with high spatial locality, a lot of cache conflicts and hence cache misses will be generated. This is why the middle bits are used as the index and higher bits are used as the tag.

Solution of Question 6:

(a) The number of bits used for byte offset $= log_2 16 = 4$.

(b) The number of blocks $= 128 \times 1024/16 = 8192$, so the number of bits used for index $= log_2 8192 = 13$.

(c) The number of bits used for tag $= 32 - 4 - 13 = 15$.