


P6. Addresses are broken up this way when accessing the L1 cache:



Give me two possible configurations of this cache.

P7 A particular out-of-order superscalar processor has a base CPI (ideal memory) of .8 cycles/instruction. That processor successfully hides half of the latency of cache-miss load instructions (through a non-blocking cache and out-of-order execution) and hides all the latency of store instructions (through the write buffer). What is the CPI of that processor with an instruction cache miss rate of 3%, a load miss rate of 8%, a miss penalty of 30 cycles to the L2 cache (both instructions and data), a miss penalty of 50 cycles from L2 to memory, and a local miss rate of 30% for the L2. Assume 20% of instructions are loads.

P8 It is observed in the following code that each iteration requires two loads ($A[j]$, then $A[j+4096]$) and a store ($A[j]$), and the cache hit rate on the store is 0% due to conflict misses. Tell me everything you can conclude about the cache. Assume that elements of $A[]$ each require 4 bytes.

```
for (j=0; j<N; j++) {  
    A[j] = A[j] + A[j+4096];  
}
```