

Parallel Computing –Lab 3


Parallel Architectures

In this Lab you are going to explore **parallel architectures**.

A. Cache coherency and false sharing (5 marks)

The simplest cache coherency protocol is **MESI**. Although cache coherency is important for making sure a program executes correctly, **it leads to false sharing**, which can cause performance problems. In order to remove false sharing, we have to either **move variables in memory** so that **they fit in different cache lines** or **Align shared global data to cache line boundaries**. For example, in the situation shown in this code:

```
struct foo {  
    int x;  
    int y;  
};  
int numThreads = 2;  
  
struct foo z;  
if (threadID ==0) {  
    int sum1 = 0;  
    int i;  
    for (i = 0; i < 100000; ++i)  
        sum1 += z.x;}  
  
if (threadID ==1) {  
    int sum2 = 0;  
    int j;  
    for (j = 0; j < 100000; ++j)  
        sum2 += z.y;}
```



cache line size is 64 bytes = $64/4=16$ integer space

- x is in the first integer space**


The minimum number of integer elements the padding array is 15

- The minimum number of integer elements the padding array is 9

- ZERO THEY ARE IN DIFFERENT CACHE LINE

- Consider the following organization of a NUMA machine:**

Machine (512GB)



Host: cantor

Indexes: physical

Date: Wed 12 Jun 2013 12:48:31 BST

a- Can you specify the following: [5 Marks]

1- How many NUMA nodes?

8

2- How many NUMA regions?

4

3- How many cores per node?

6

4- How many cores per region?

12

b- Within each NUMA node, can you specify:

1- How many cache levels?

4

2. Which levels of cache are shared? and between which cores?

L1i , L2 is shared between two adjacent cores

L3 is shared between all six cores in the same node

3. Which levels of cache are private?

L1d is private