

# High Performance Computer Architecture

## Homework 5

Total Points = 50

Assigned April 22, 2002, Due April 27, 2005

**Question 1.** (20 points) Assume that the following program that adds two 64-element vectors A and B is running on a shared memory multiprocessor that uses an interconnection network to connect 64 CPUs with 64 Memory units

```
{Initially Sum = 0}
for (j=0; j<64; j++)
    Sum = Sum + A(j) * B(j)
```

How will you map the different elements into the memory modules and rewrite the program so that it leads to maximum speedup. The major issue here is to avoid various types of contentions and data dependencies as much as possible.

**Question 2.** (20 points) Let B be a (1024 x 1024) matrix. Consider the execution of the following program on a superscalar processor that fetches four instructions per cycle:

```
for (j=0; j<1024, j=j+1) {
    S[j] = 0;
    for (k=0, k <1024, k=k+1) {
        S[j] = S[j] + B[j,k]
    }
}
```

How will a compiler unroll the loop to maximize ILP? Do not convert it to assembly language

**Question 3.** (10 points) Consider the following specifications of a disk. If the average queuing delay is 20 ms, then estimate how much time will it take a to read a file of size 100 MB (There is no RAID here). Ignore controller delay.

	<b>Seagate Cheetah Ultra160 SCSI</b>
Diameter	3.5"
Formatted capacity	73.4 GB
Cylinders	14,100
Disks	12
Heads	24
Bytes / sector	512-4096
Sectors / track	~424
RPM	10,033
Seek time	4.0 ms
Data transfer rate	80 MB/s