


2.2 & 2.3

- 2.2) The computation performs 8 FLOPS on 2 cache lines, i.e., 8 FLOPS in 200ns. This corresponds to a computation rate of 40 MFLOPS.
- 2.3) In the best case where the vector gets cached, 8 FLOPS can be performed on 1 cache line (of the matrix). This corresponds to a peak computation rate of 80 MFLOPS. The matrix does not fit in the cache.

Alexandre David, MVP'06

1


2.10) Omega Network



Alexandre David, MVP'06

2


2.10) Moved Around



Alexandre David, MVP'06

3


2.10) Routes



Alexandre David, MVP'06

4


2.10) Switches Removed



Alexandre David, MVP'06

5

2.10) Moved Around



Alexandre David, MVP'06

6