



# Home work 5b

CS 594 Spring 2006 Clusters Part 1  
Workloads made easy


## What to do




## What to do




## What to do




## What to do




## What to do



## What to do



## What to do



## What you will get

You will be supplied with the clients

The basic servers

2 basic sequential brokers, one for HT and one for HA workloads)

All will be MPI codes

## What you will get

You will take the broker codes and make them run better for *both* the HA and HT workloads

You will then measure their performance using the built in Stats functions (I have already done the hard work)  
(Stats are printed for the clients and servers)

You can do whatever you want to the broker codes  
(broker.h, broker\_ha\_yours.c and broker\_ht\_yours.c)  
**You should not change any other codes.**

Test with `mpirun -np 9 main {0,1,2,3} 100 1`  
and `mpirun -np 9 main {0,1,2,3} 100 10`  
(Note you need to use the head node as rank 0, +your 8 other nodes)

Email in (md5sum and dist tarball (inc writeup.txt)) by 1<sup>st</sup> March.