BSP Problem 1

In class we discussed a one-to-all broadcast that sent messages from one processors to all other processors in using $log(p)$ BSP supersteps. Another algorithm is the all-to-all broadcast which can be naively implemented as p successive one-to-all broadcasts such that each processor starts with one data item and ends with p data items that contains the data provided by each processor. This naive implementation would be performed in $plog(p)$ BSP supersteps. Since these p broadcasts are indepenendent, you could perform the same operation in $log(p)$ BSP supersteps. Describe such an algorithm and derive its running time.

BSP Problem 2

In class we discussed a optimal broadcast algorithm that could adjust to the relative magnitude of the system bandwidth (g) and sytem latency (l). Describe an algorithm that could perform a similar optimization for summing n numbers on p processors. Derive the running time for this algorithm.