- Hash index on sid of Sailors
    - sid is key of Sailors so at most 1 matching tuple but sid is foreign key in Reserves so exactly 1 matching tuple in Sailors
    - Total cost is 221,000 I/Os
- Hash index on sid of Reserves
    - Sid not key of Reserves so 1 or more matching tuples
    - Cost of scanning Sailors is 500
    - Cost of using hash index for every tuple is (500*80)*1.2 = 48,000 I/Os
    - Assuming uniform distribution of reservations to sailors then (100,000/40,000) = 2.5 reservations
    - If index on Reserves is clustered then all reservations would likely be on 1 page so accessing matching tuples requires 40,000 I/Os
    - If index on Reserves is unclustered then 1 access per match so 40,000 * 2.5 = 100,000 I/Os
    - So in case of clustered index cost of join using index nested loops is 500 + 48,000 + 40,000 = 88,500 I/Os
        - Assuming 10ms per I/O then time is about 15 minutes
    - So in case of unclustered index cost is 500 + 48,000 + 100,000 = 148,500
        - Assuming 10 ms per I./O, time is about 25 minutes