Index Nested Loops Join

- If there is an index on the join column of one relation (say S), can make it the inner and exploit the index.
  - Cost: M + (M*p_R) *(accessing index + retrieving qualifying S tuples)
- For each R tuple, cost of probing S index is about 1.2 for hash index, 2-4 for B+ tree. Cost of then finding S tuples (assuming Alt. (2) or (3) for data entries) depends on clustering.
  - Clustered index: 1 I/O (typical), unclustered: up to 1 I/O per matching S tuple.