CSE 444: Database Internals

Lecture 8

Operator Algorithms (part 2)

WHITE BOARD EXAMPLE

Relation R: 10 pages
Relation S: 8 pages
Memory: 5 pages

Each page holds 2 tuples
We only show value of join attribute for each tuple

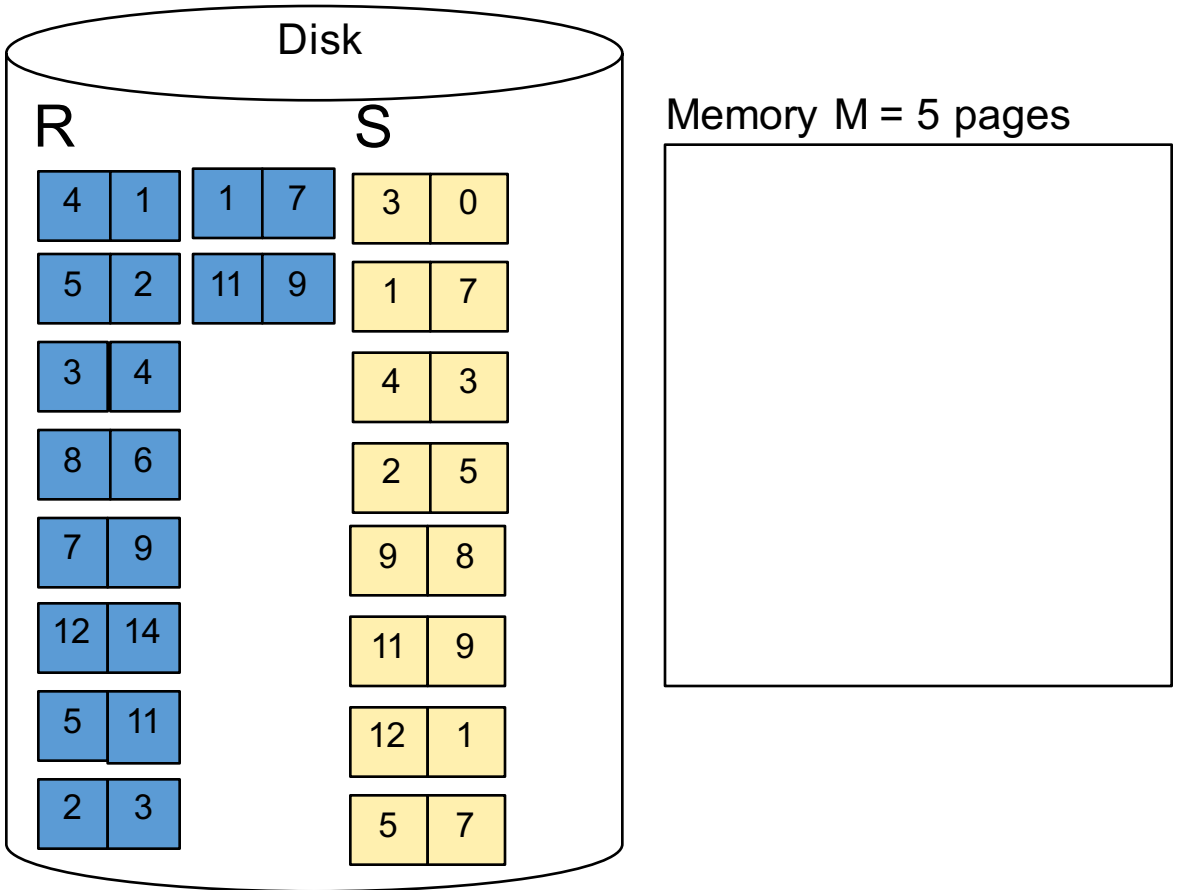

Grace Join

Also called

Partitioned Hash-Join Example

Step 1: Read relation S one page at a time and hash into M-1 (=4 buckets)

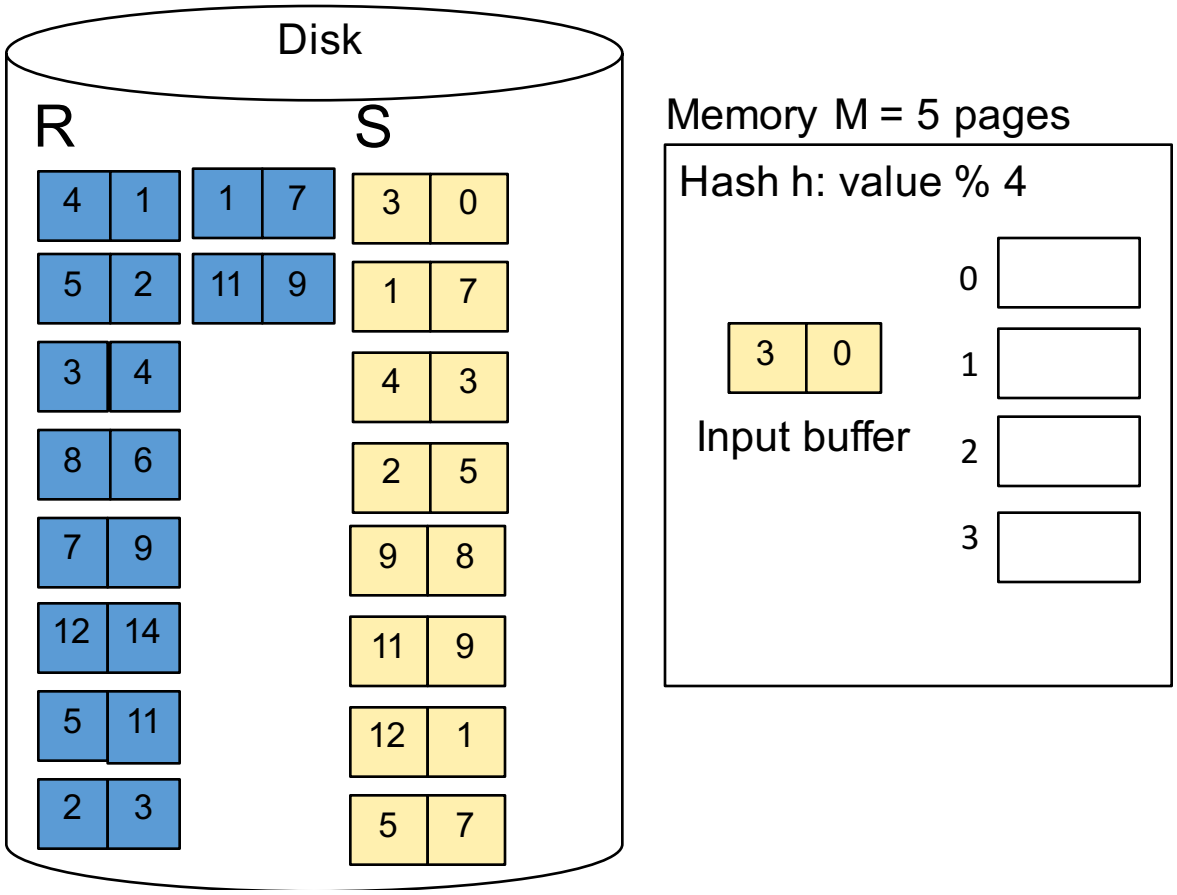

Step 1: Read relation S one page at a time and hash into the 4 buckets

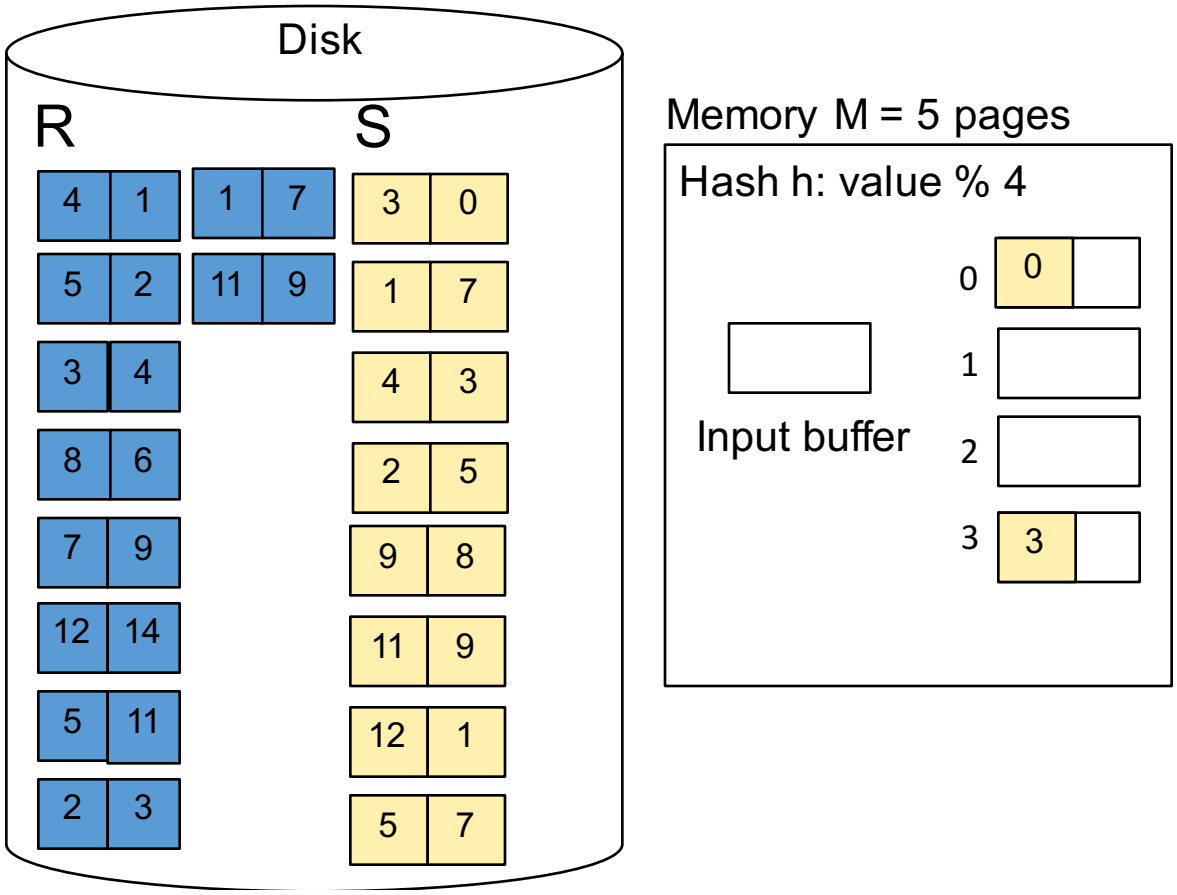

Step 1: Read relation S one page at a time and hash into the 4 buckets

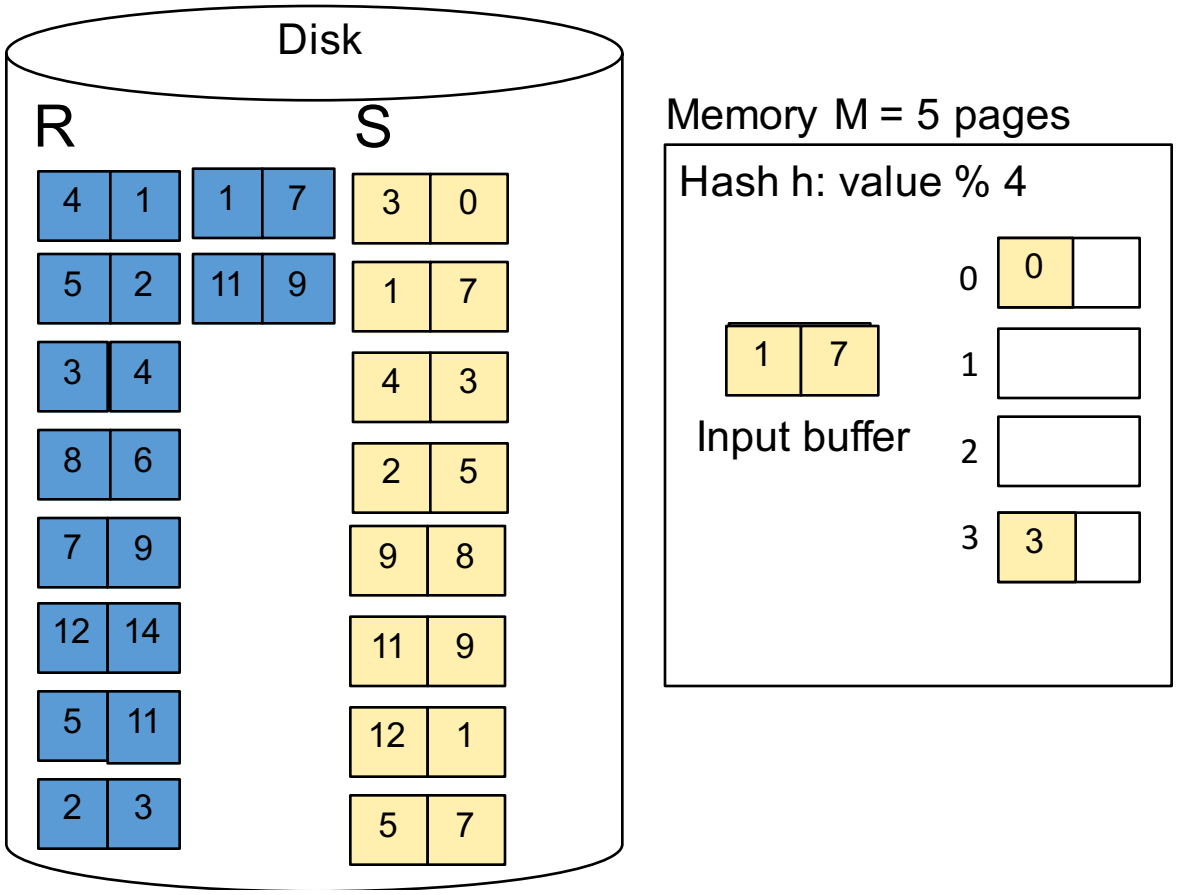

Step 1: Read relation S one page at a time and hash into the 4 buckets

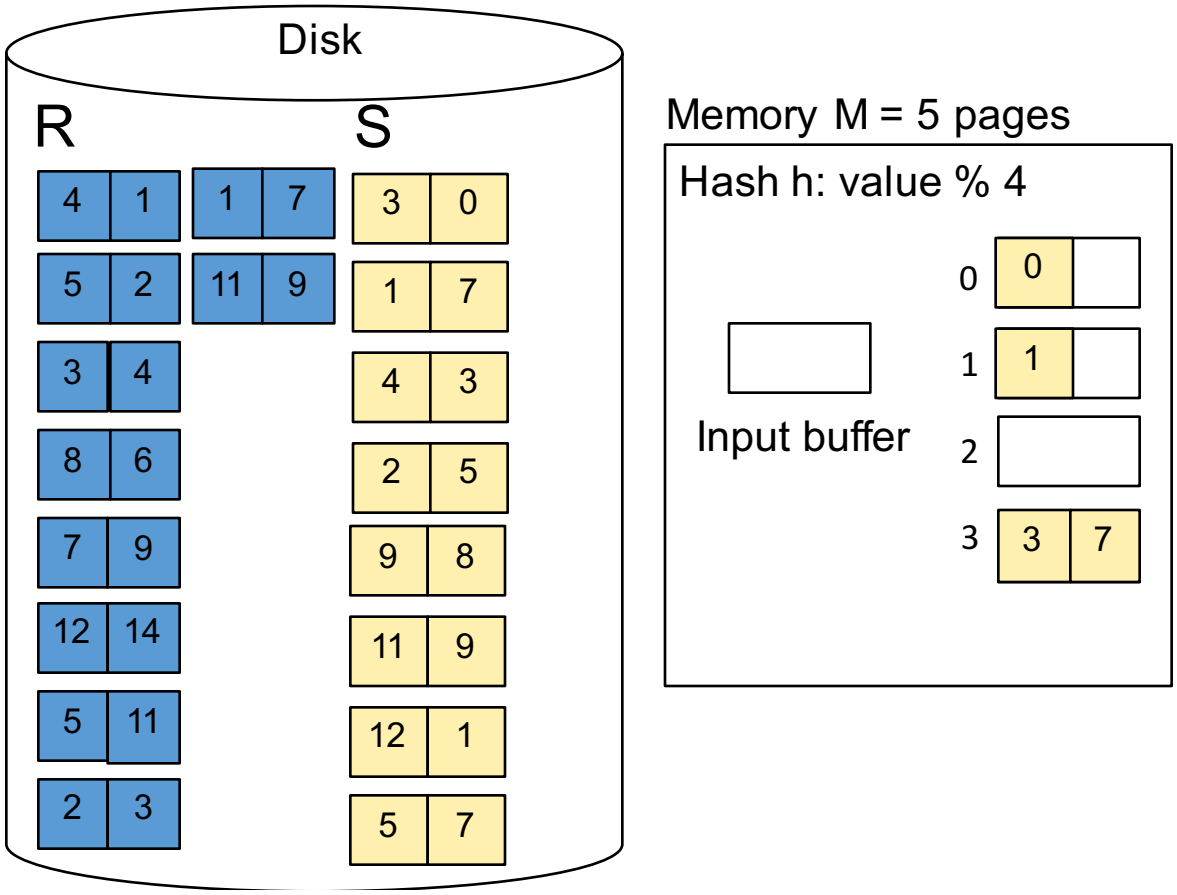

Step 1: Read relation S one page at a time and hash into the 4 buckets

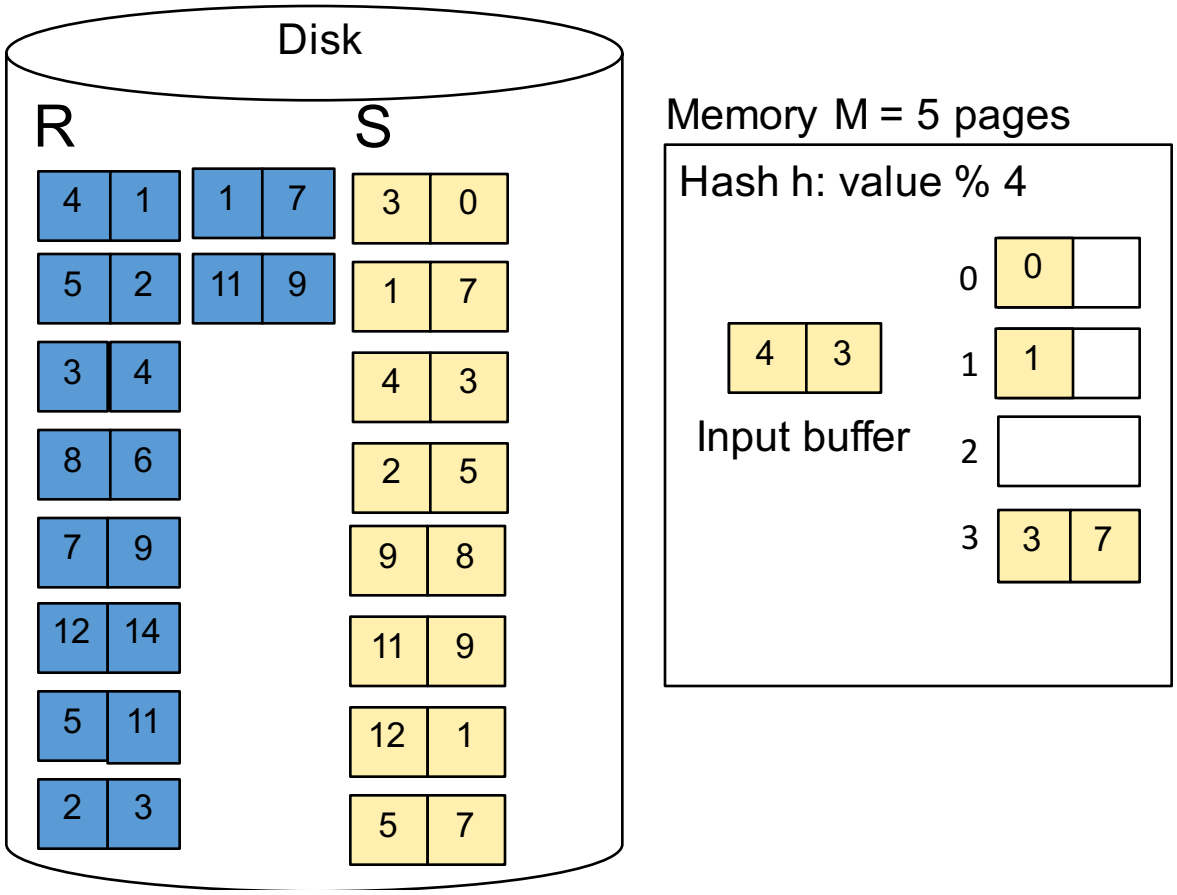

Step 1: Read relation S one page at a time and hash into the 4 buckets
When a bucket fills up, flush it to disk

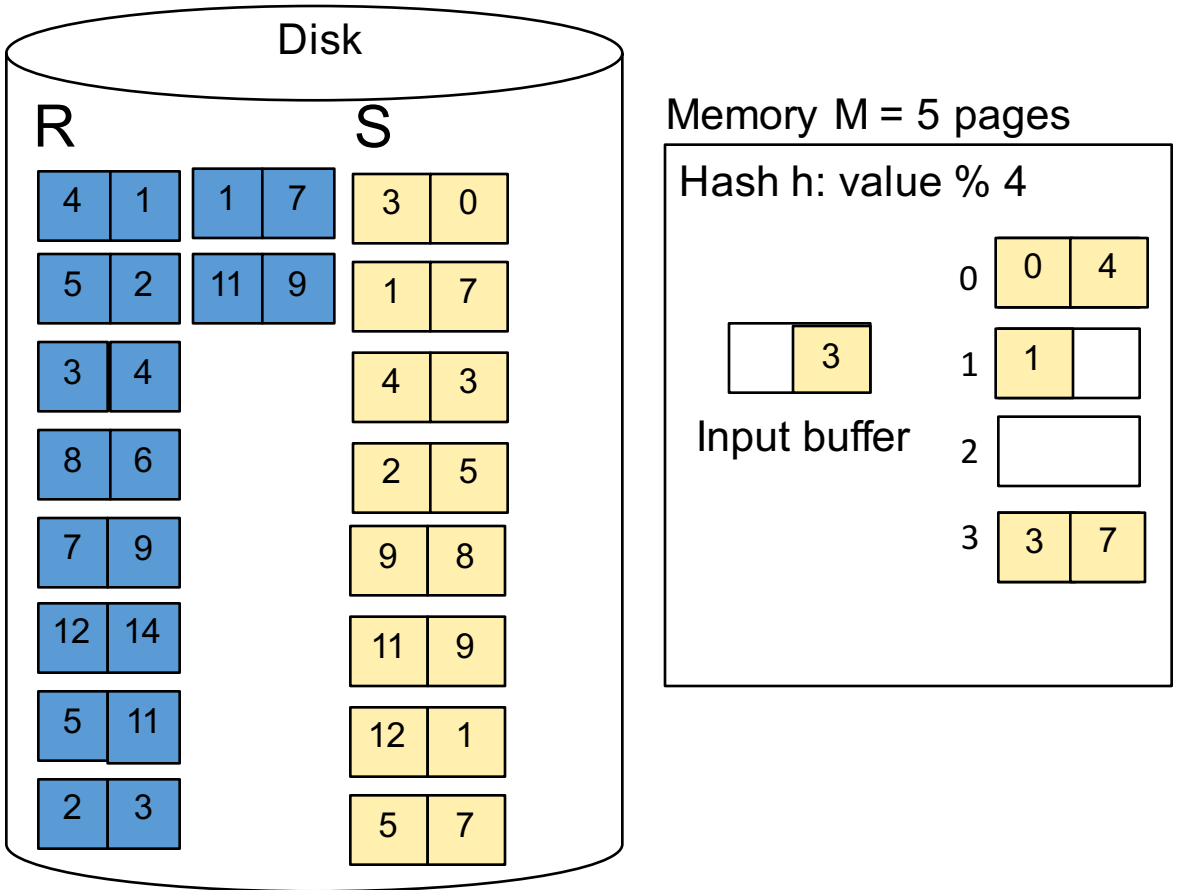

Step 1: Read relation S one page at a time and hash into the 4 buckets
When a bucket fills up, flush it to disk

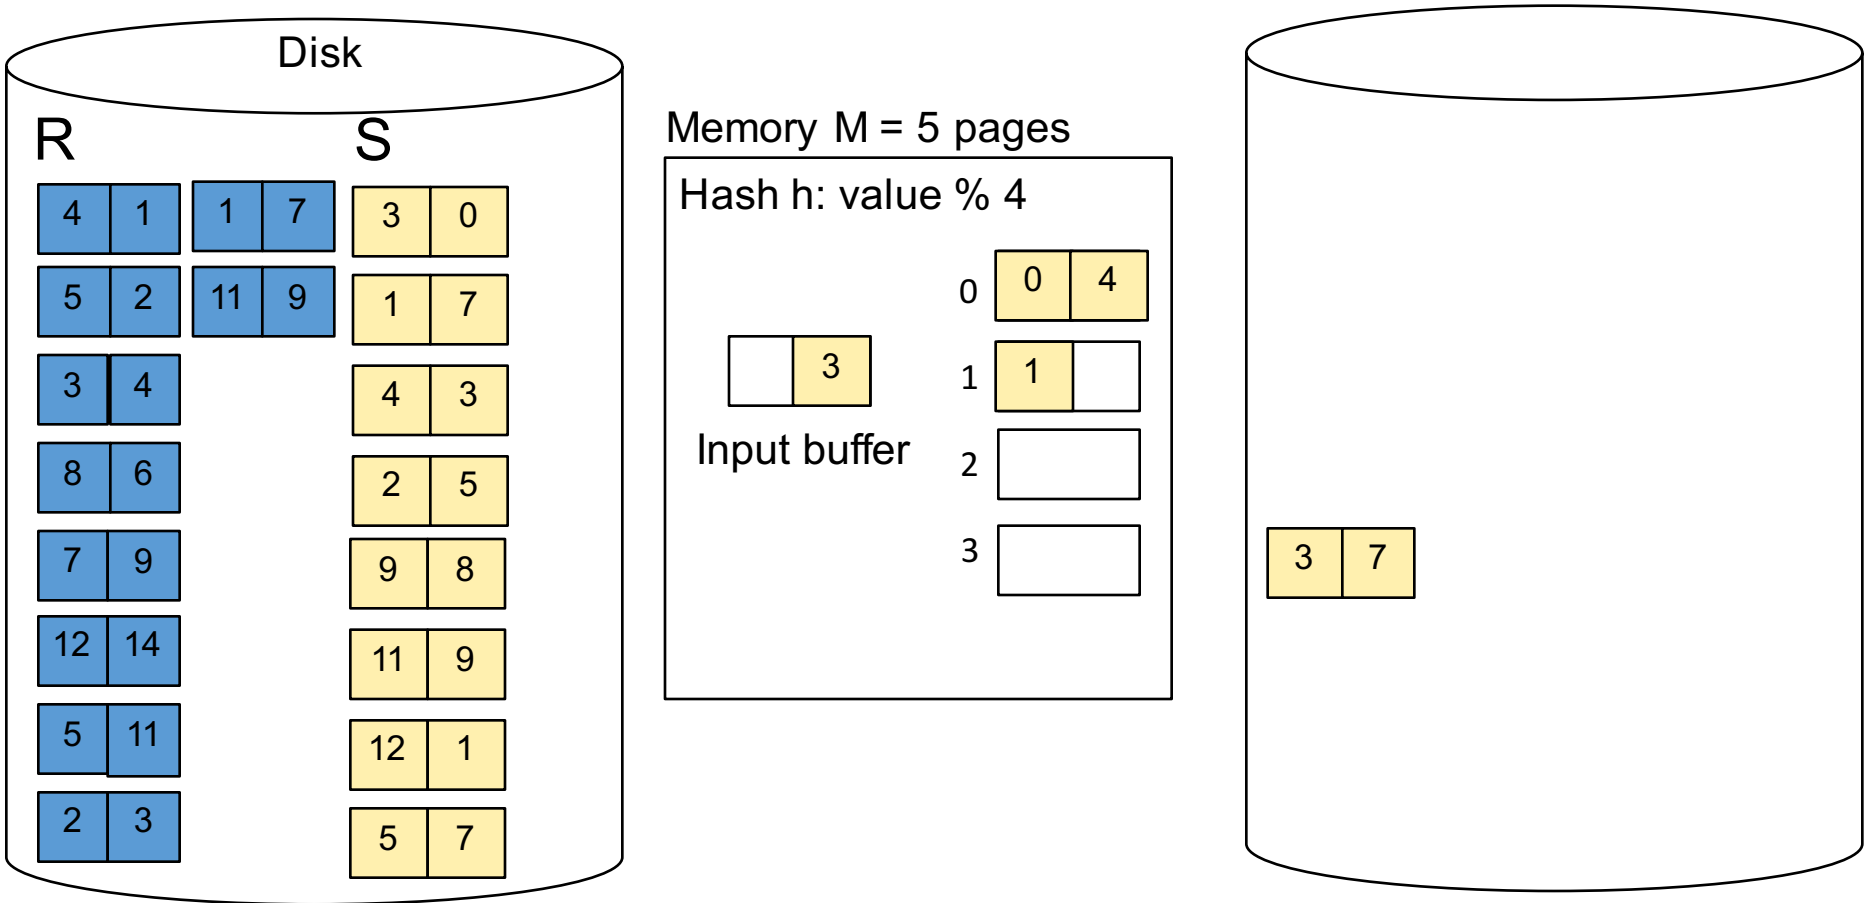

Step 1: Read relation S one page at a time and hash into the 4 buckets
When a bucket fills up, flush it to disk

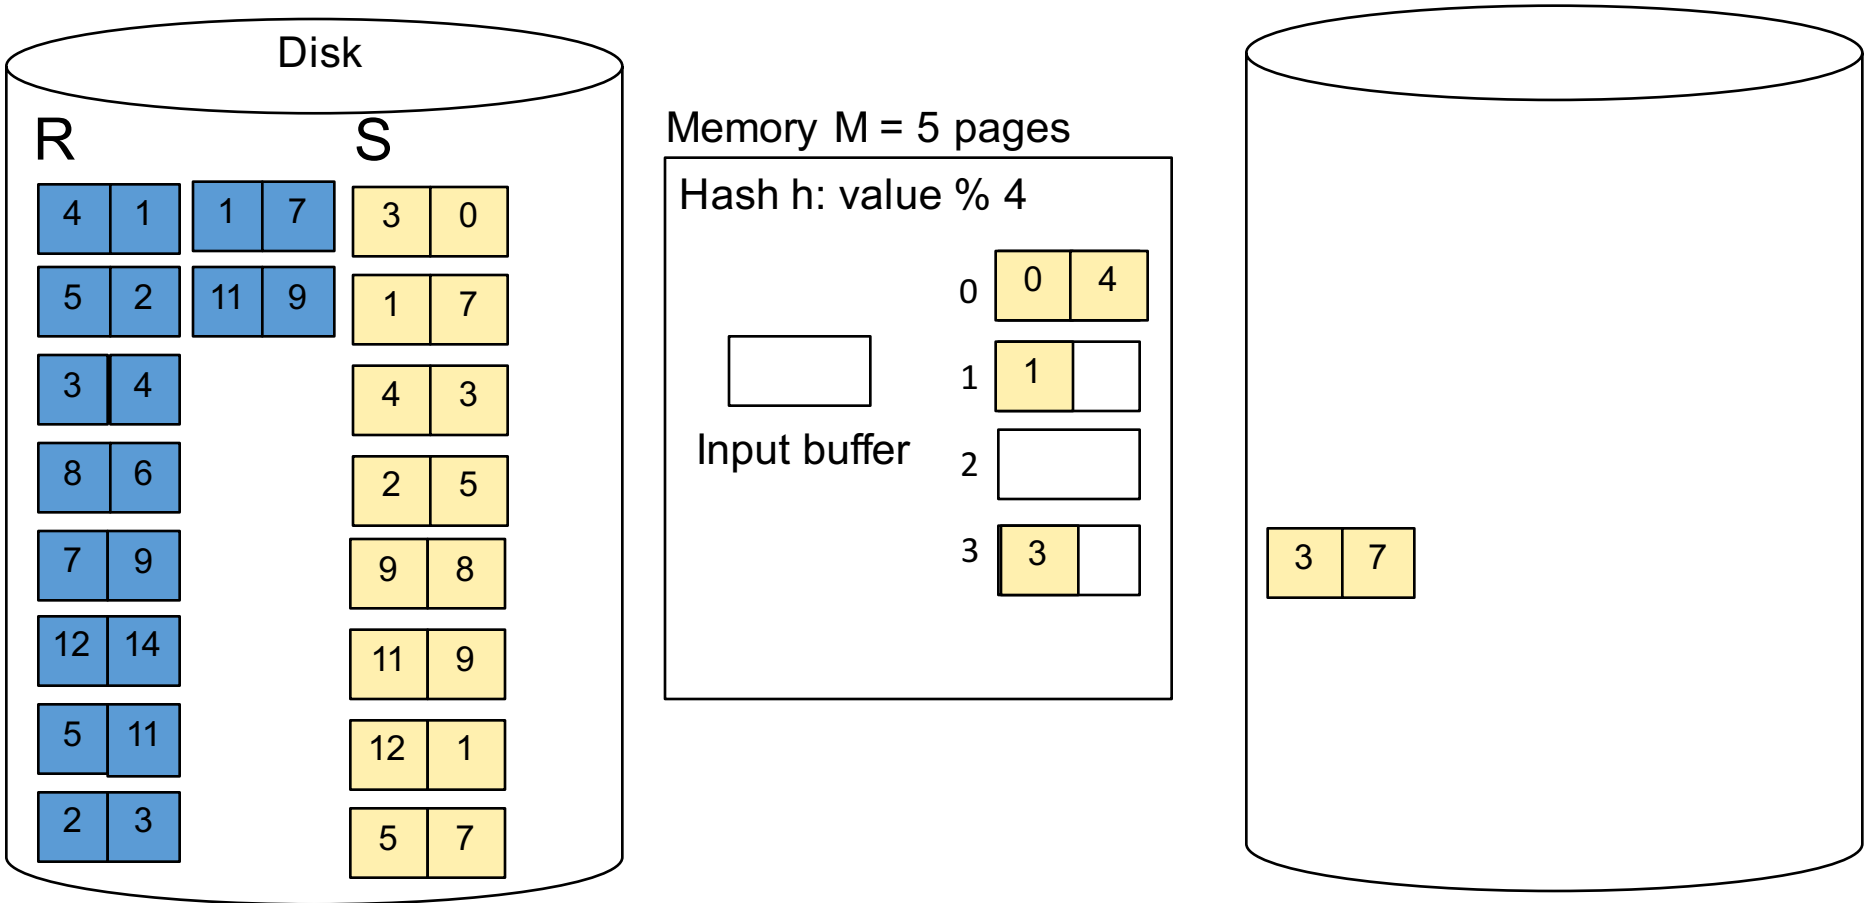

Step 1: Read relation S one page at a time and hash into the 4 buckets
At the end, we get relation S back on disk split into 4 buckets

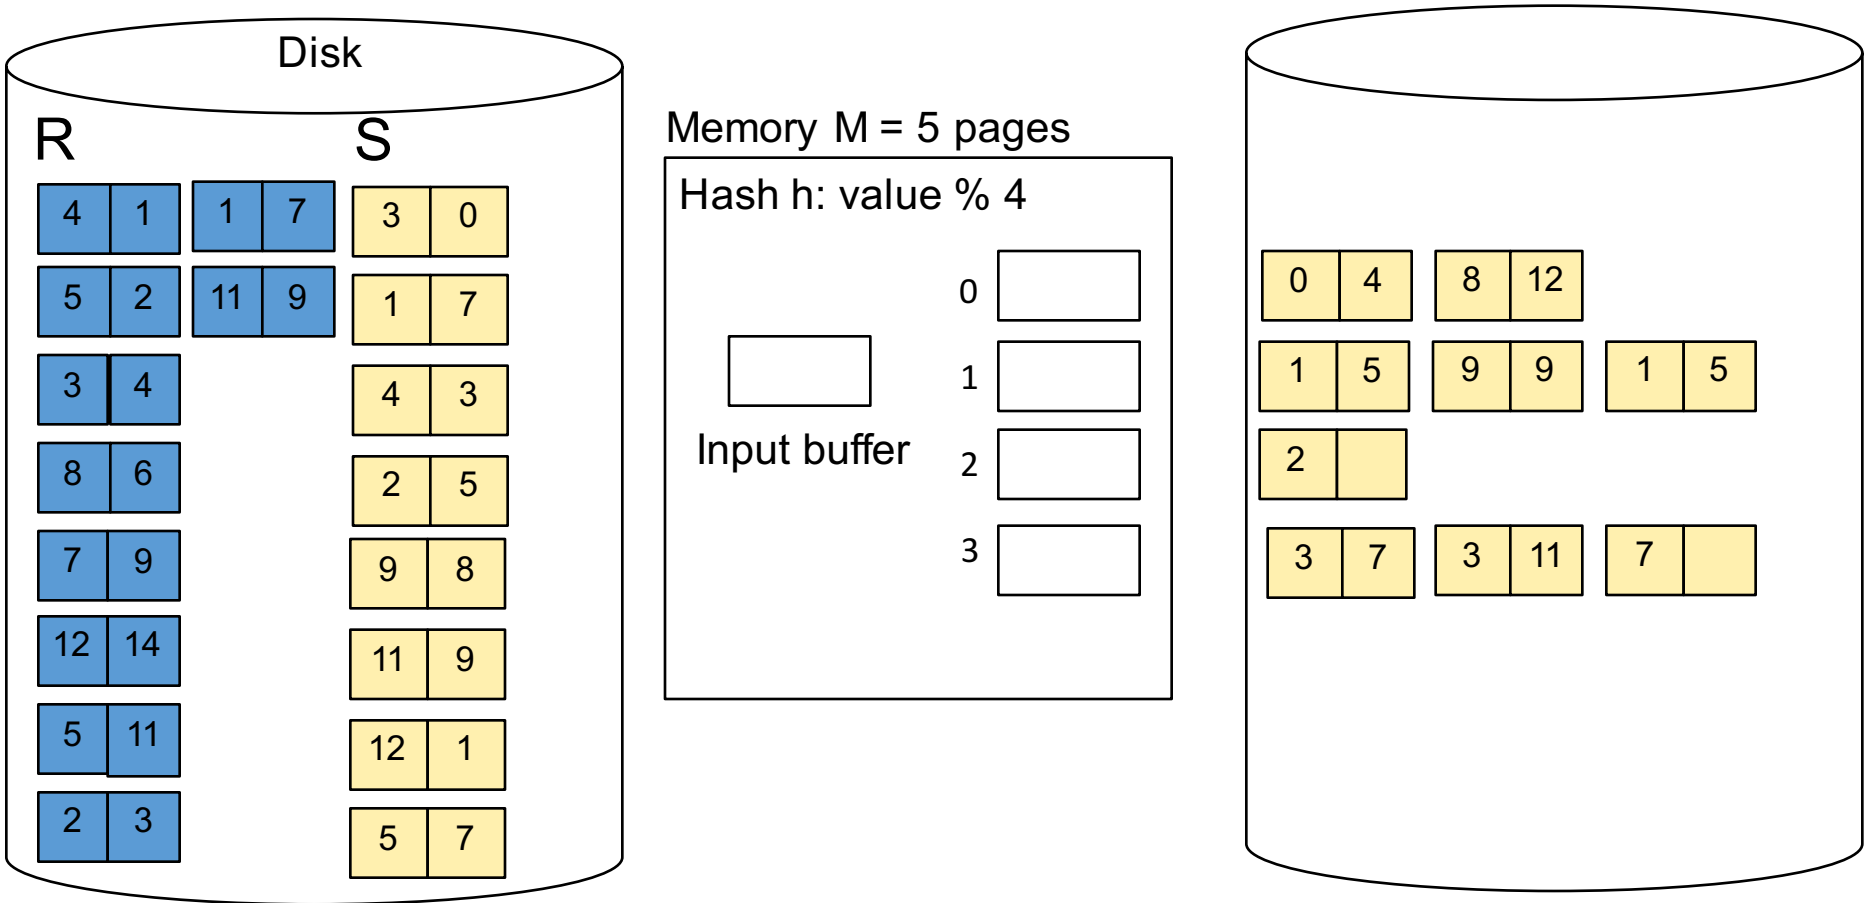

Step 2: Read relation R one page at a time and hash into same 4 buckets

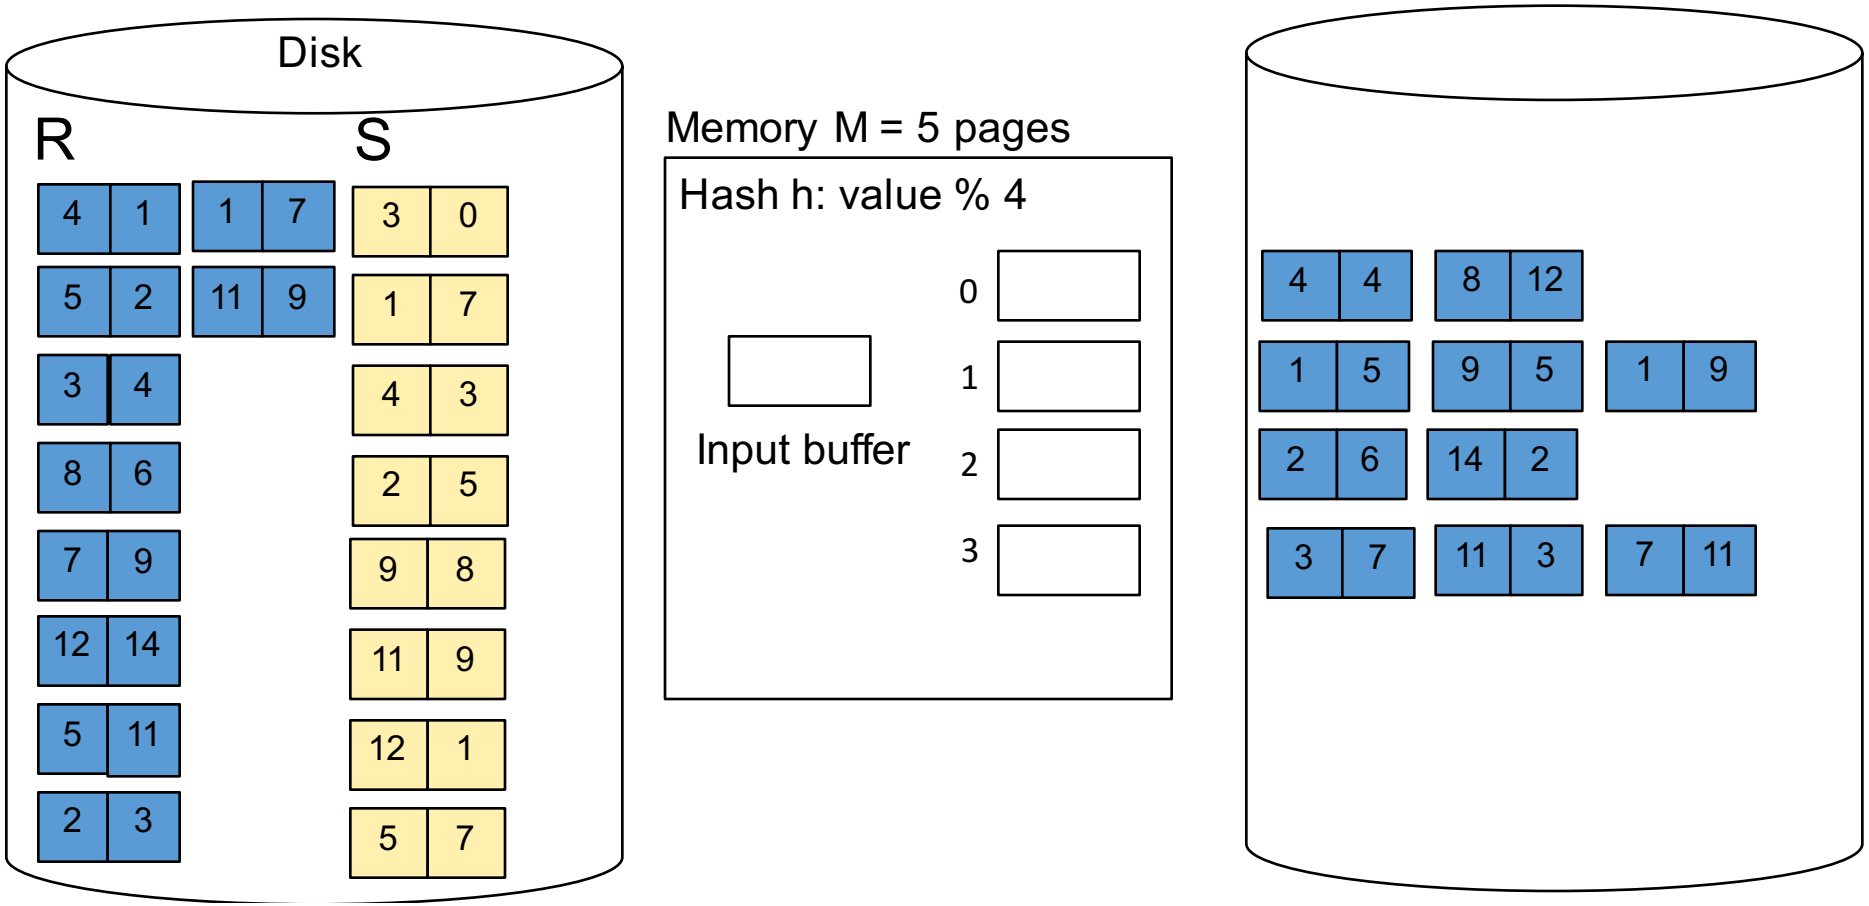

Step 3: Read one partition of R and create hash table in memory using a different hash function

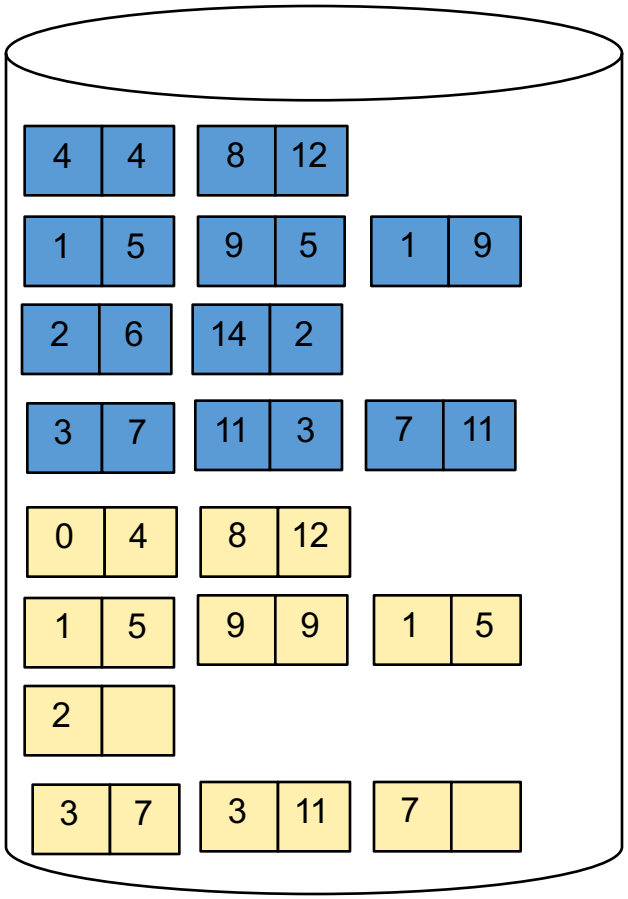

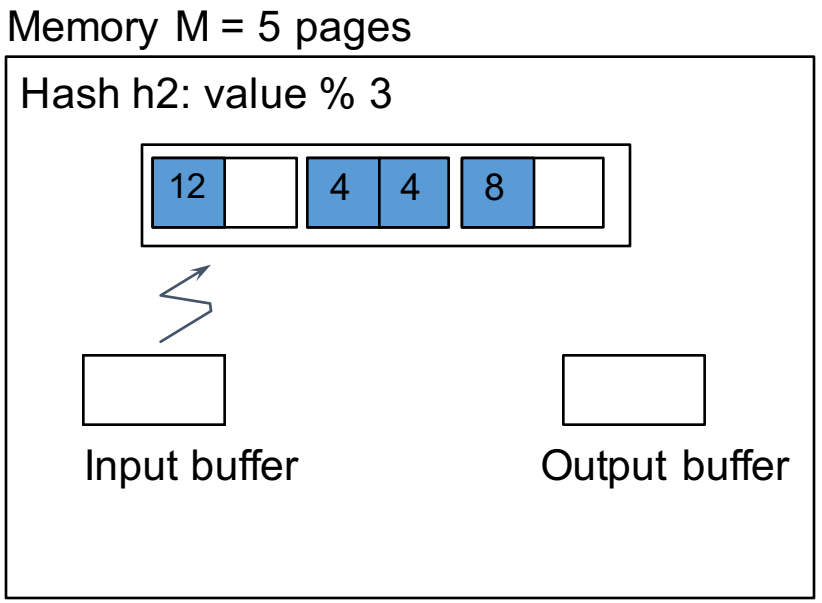

Step 4: Scan matching partition of S and probe the hash table
Step 5: Repeat for all the buckets
Total cost: 3B(R) + 3B(S)

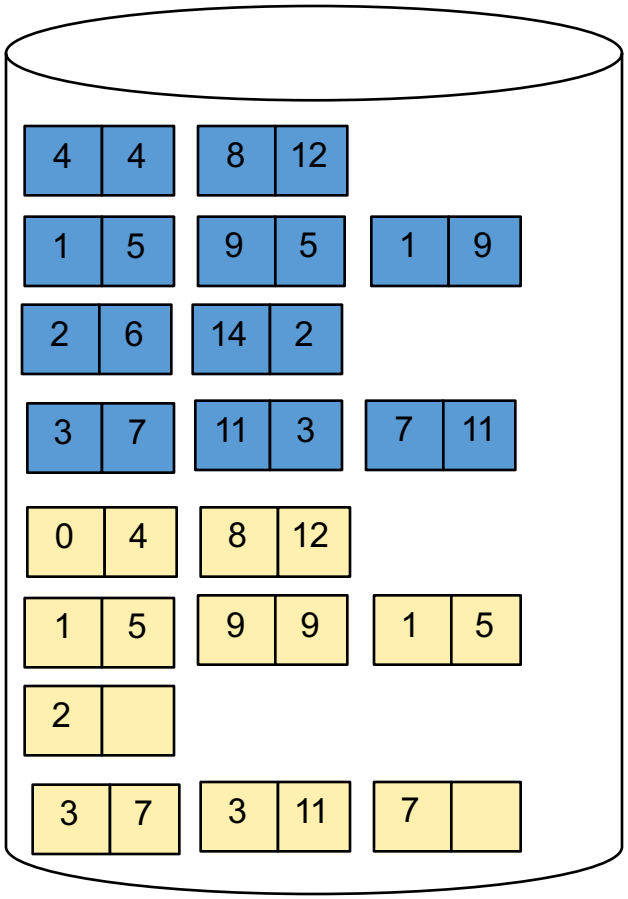

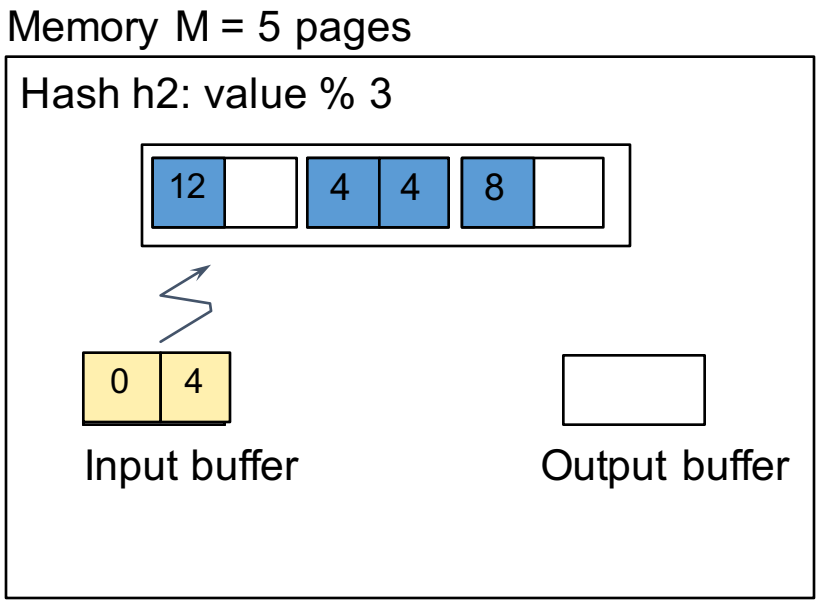

External Merge Join

Step 1: Read M pages of R and sort in memory

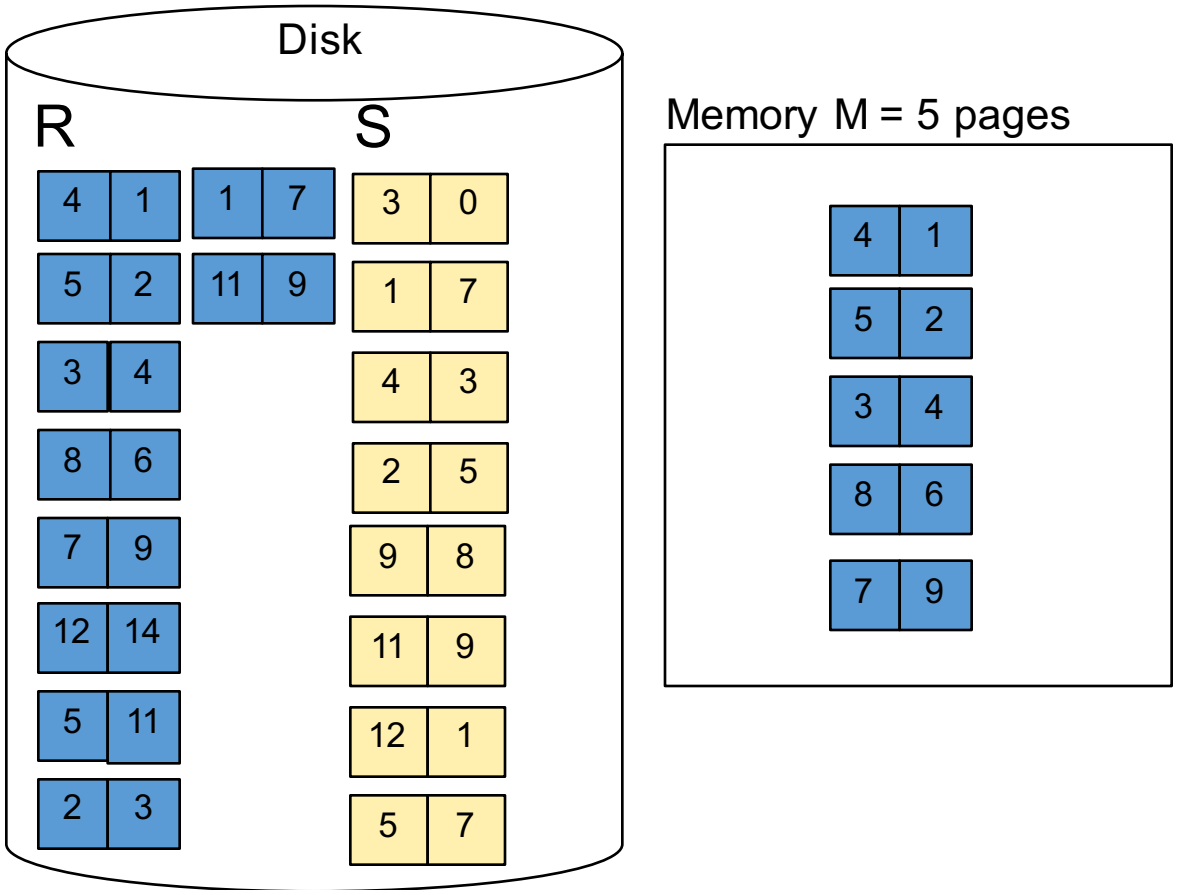

Step 1: Read M pages of R and sort in memory

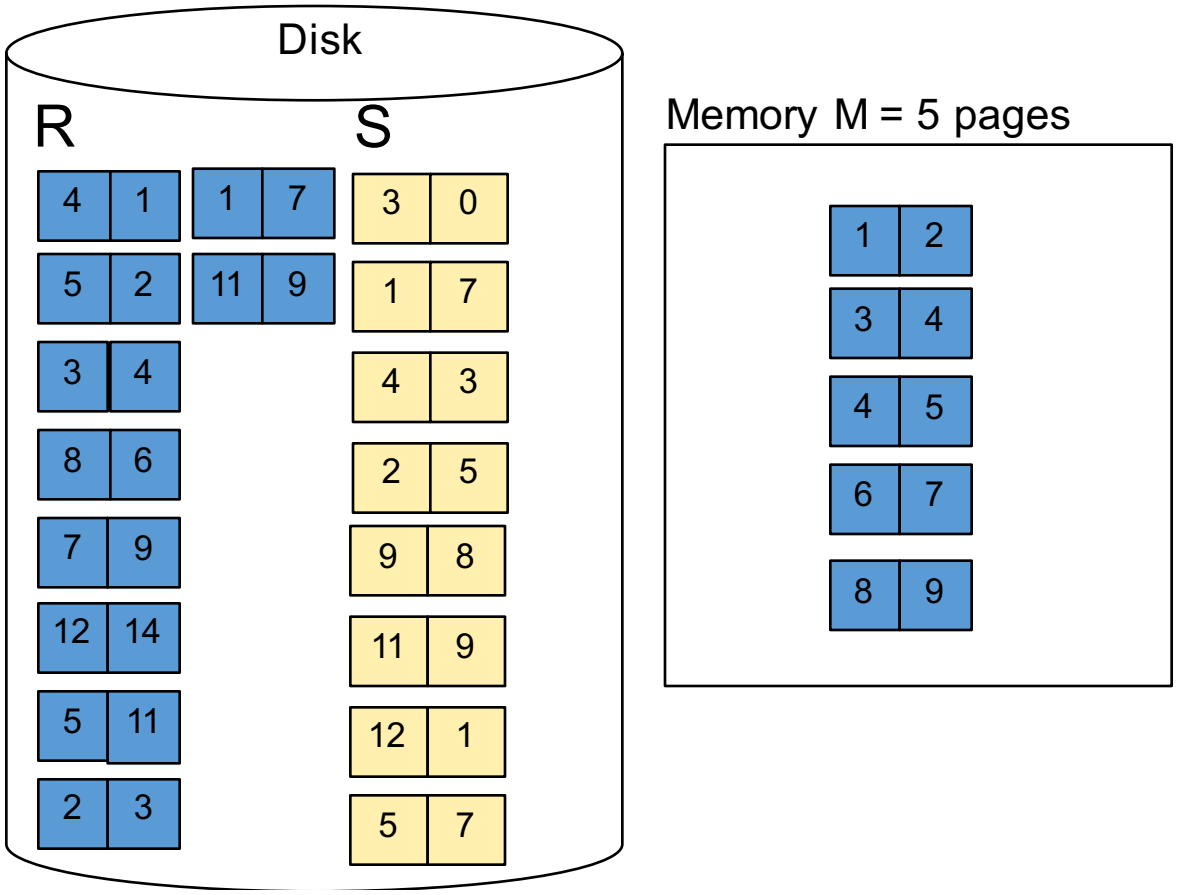

Step 1: Read M pages of R and sort in memory, then write to disk

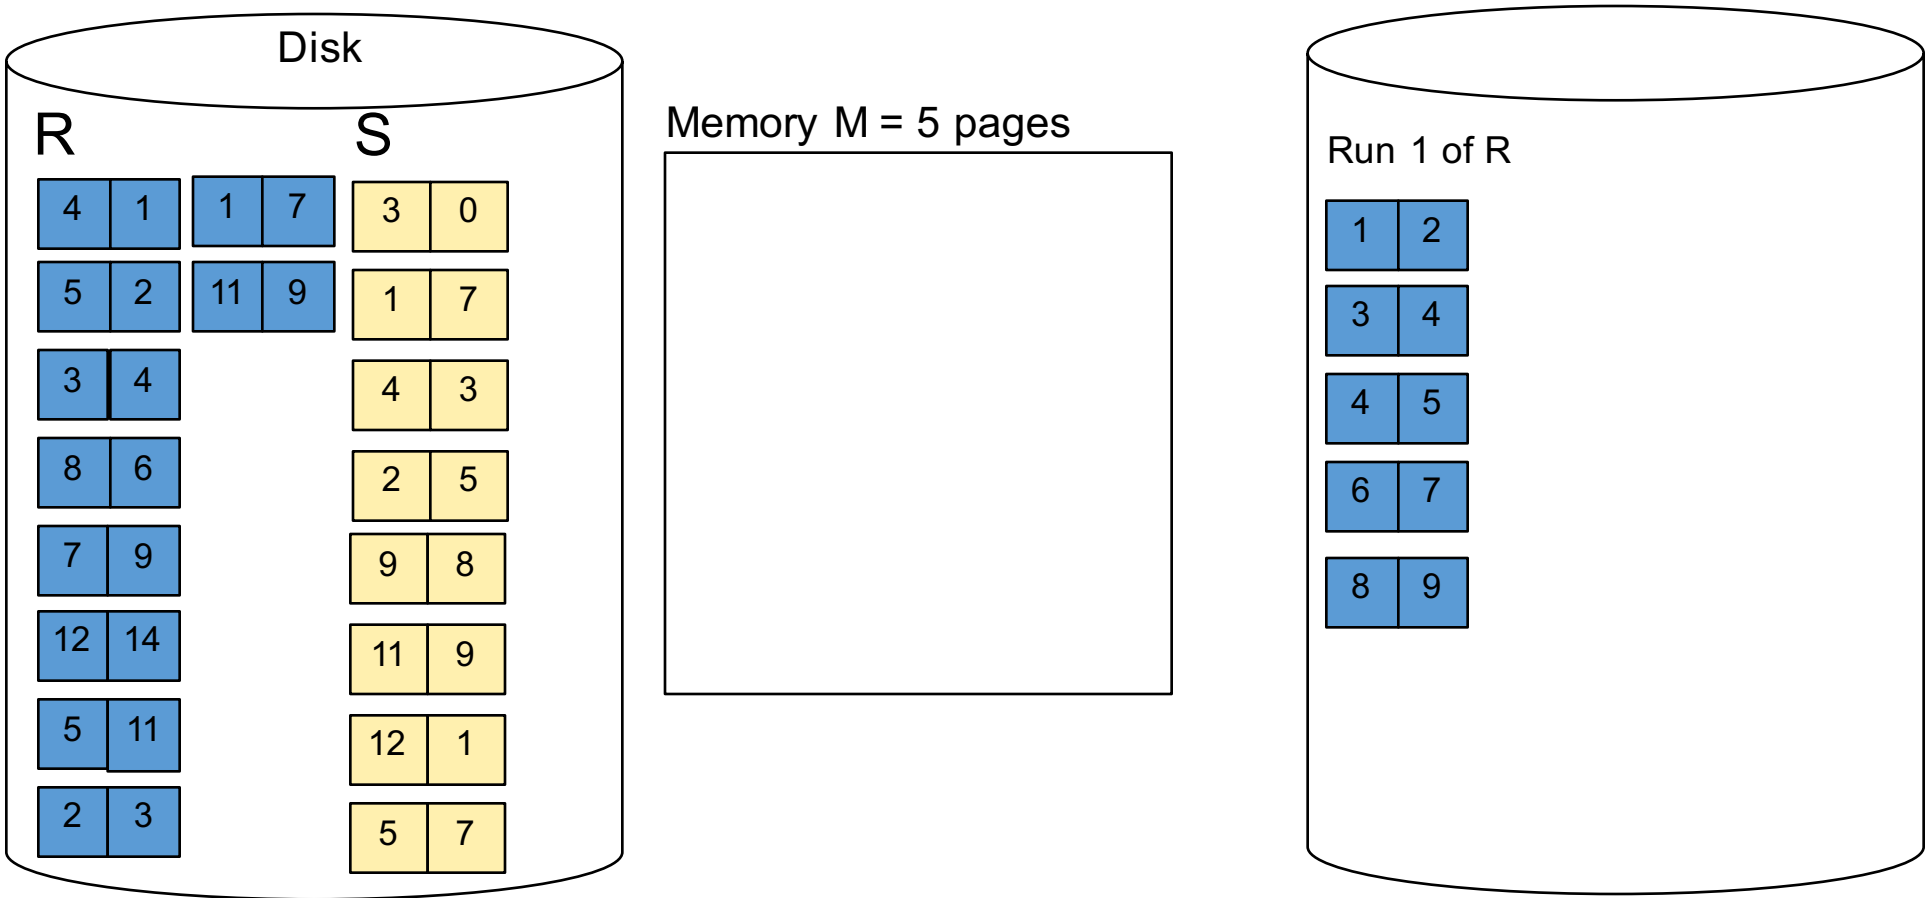

Step 2: Repeat for next M pages until all R is processed

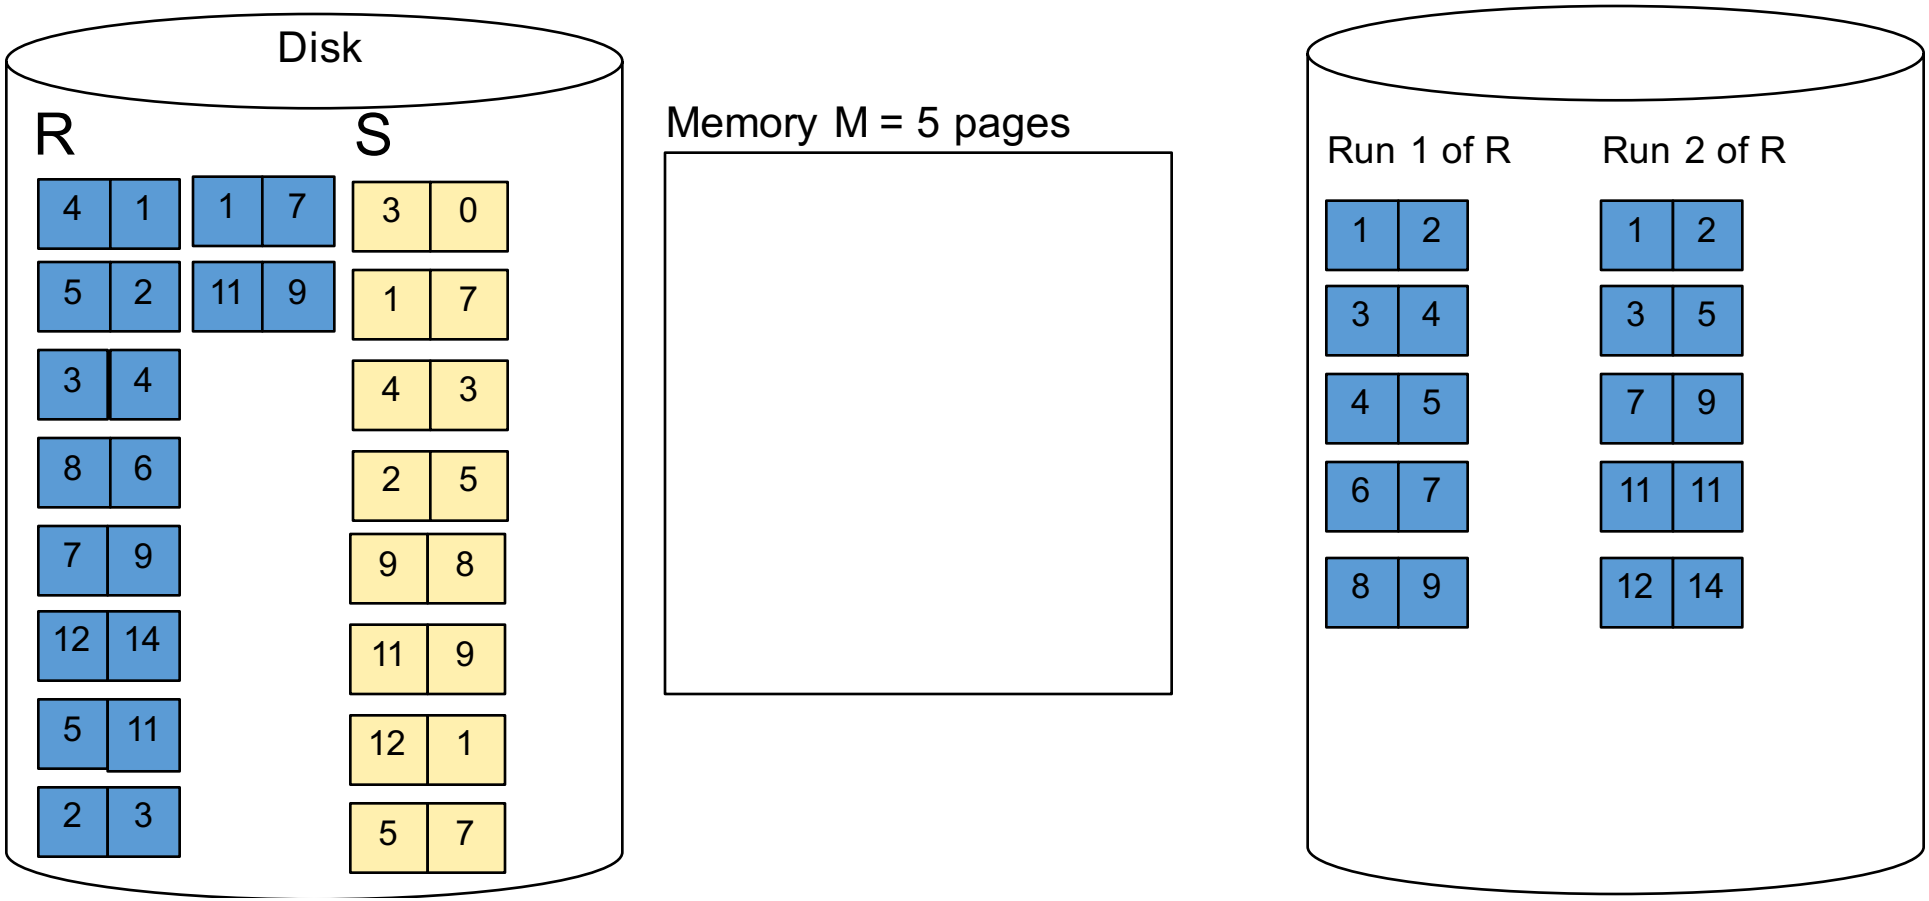

Step 3: Do the same with S

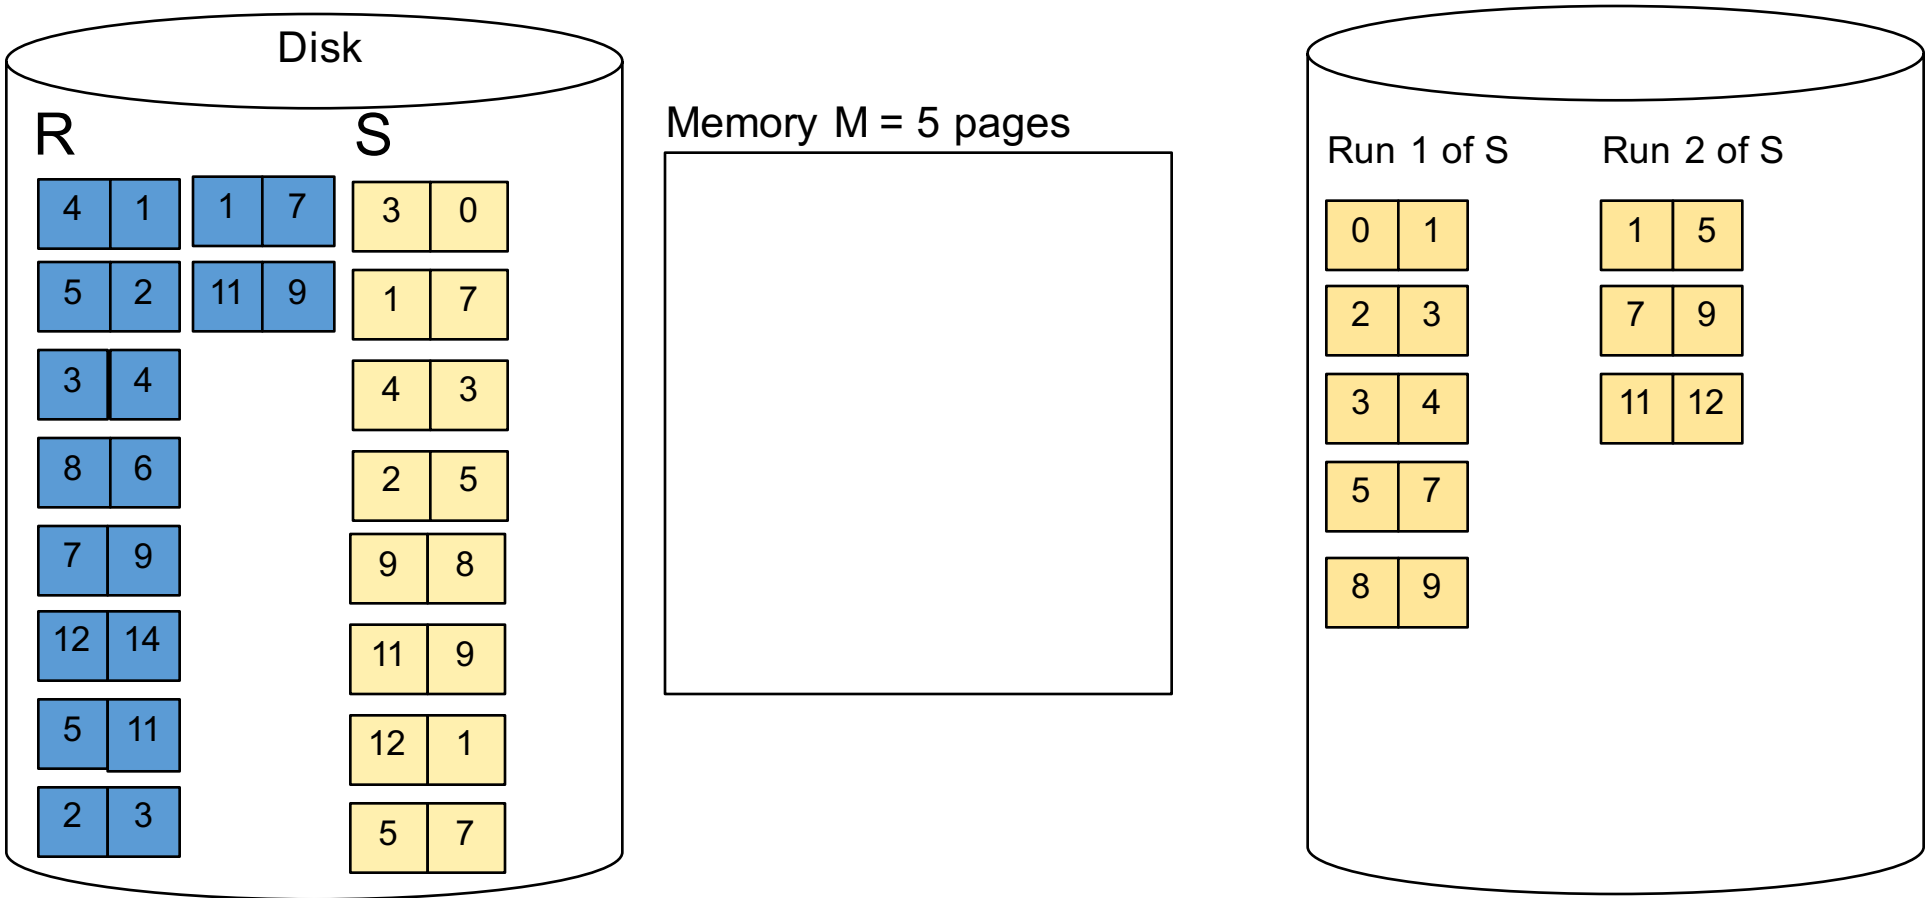

Step 4: Join while merging sorted runs

Total cost: 3B(R) + 3B(S)