Partitioned hash join example:

Bucket 0	0	4	8	12
" 1	1	5	9	13
" 2	2	6	10	14
" 3	3	7	11	15

Relation R: 10 pages   Each page can hold 2 tuples

Relation S: 8 pages

Memory: 5 pages

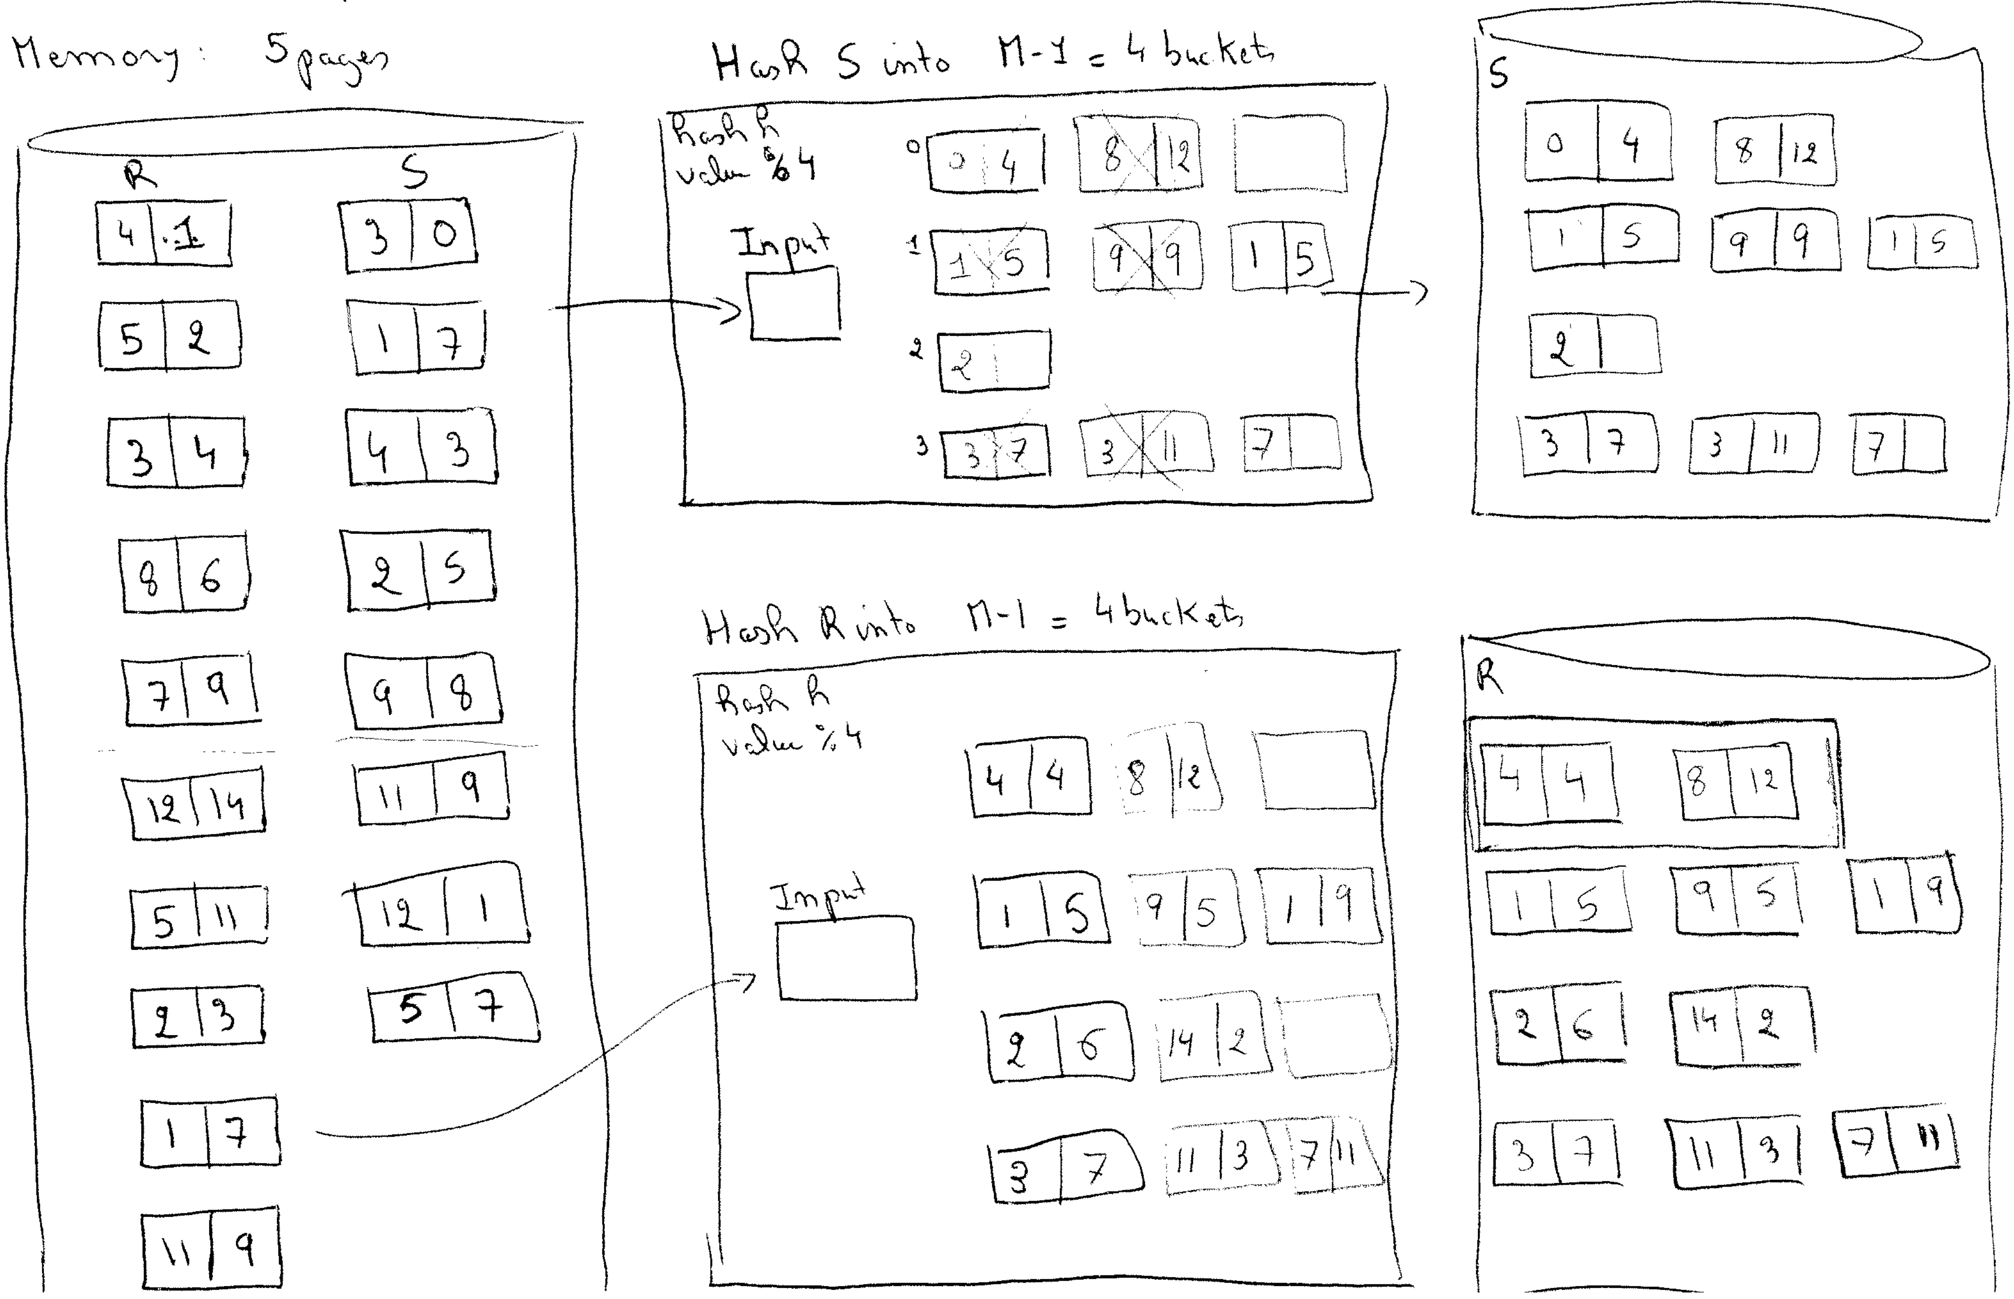

Partitioned hash join example continued: Second pass

Read one partition of R, hash it using $h2 \neq h$

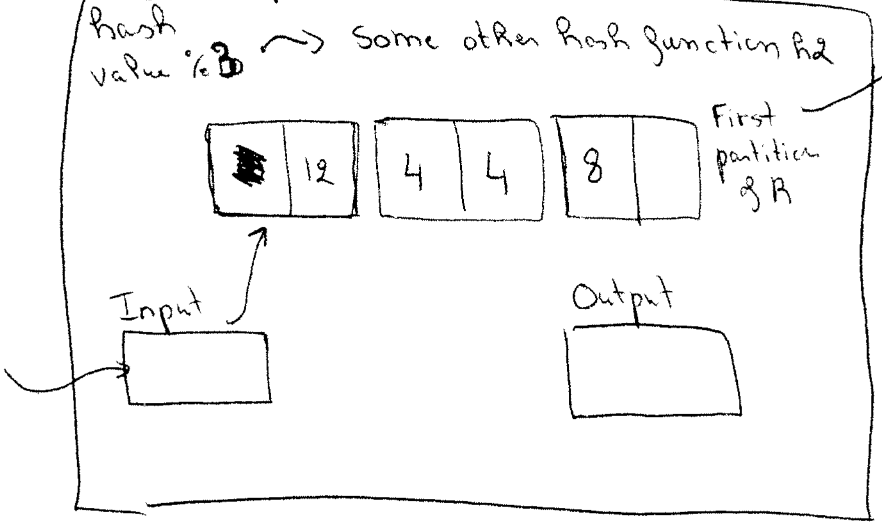

Normally a hash table requires a bit more space than raw data but we ignore this.

What is the total cost ?

Scan matching partition of S search for matches

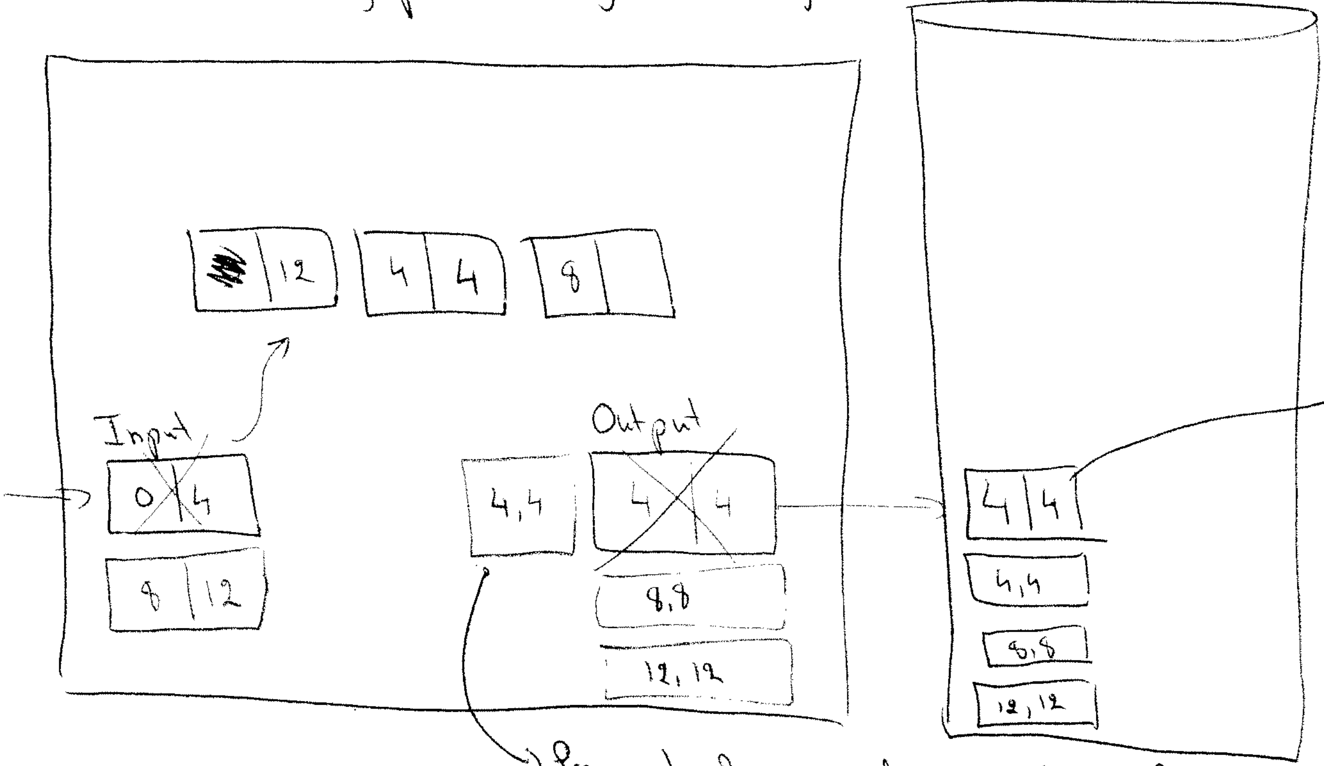

Remember that these tuples normally have more fields

Repeat for each partition.

Longer tuples [illegible] of R and S [illegible]

External Merge-Sort Example

Let's sort R

Step 1: Load M=5 pages of R

Step 2: Sort in memory

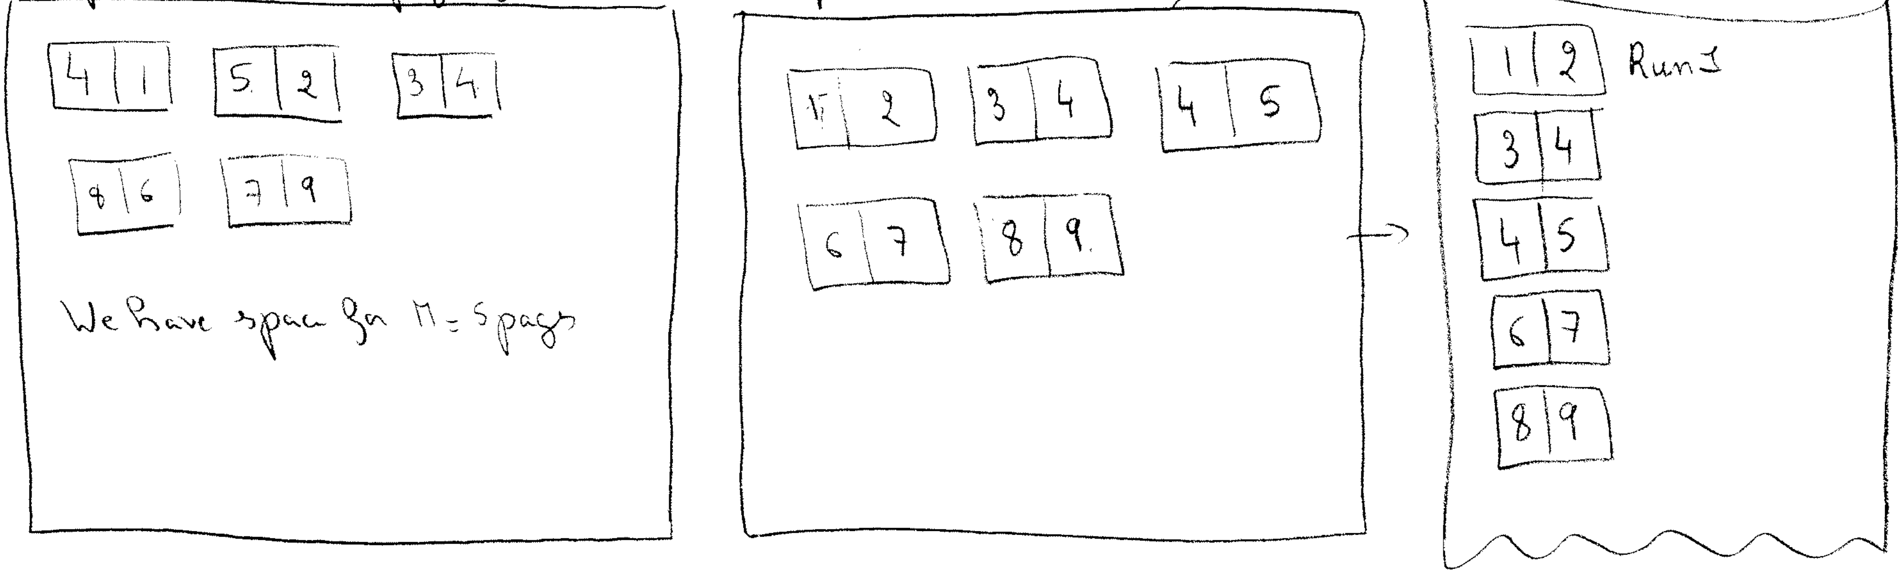

We have space for M=5 pages

Step 3: Sort next batch same as above.

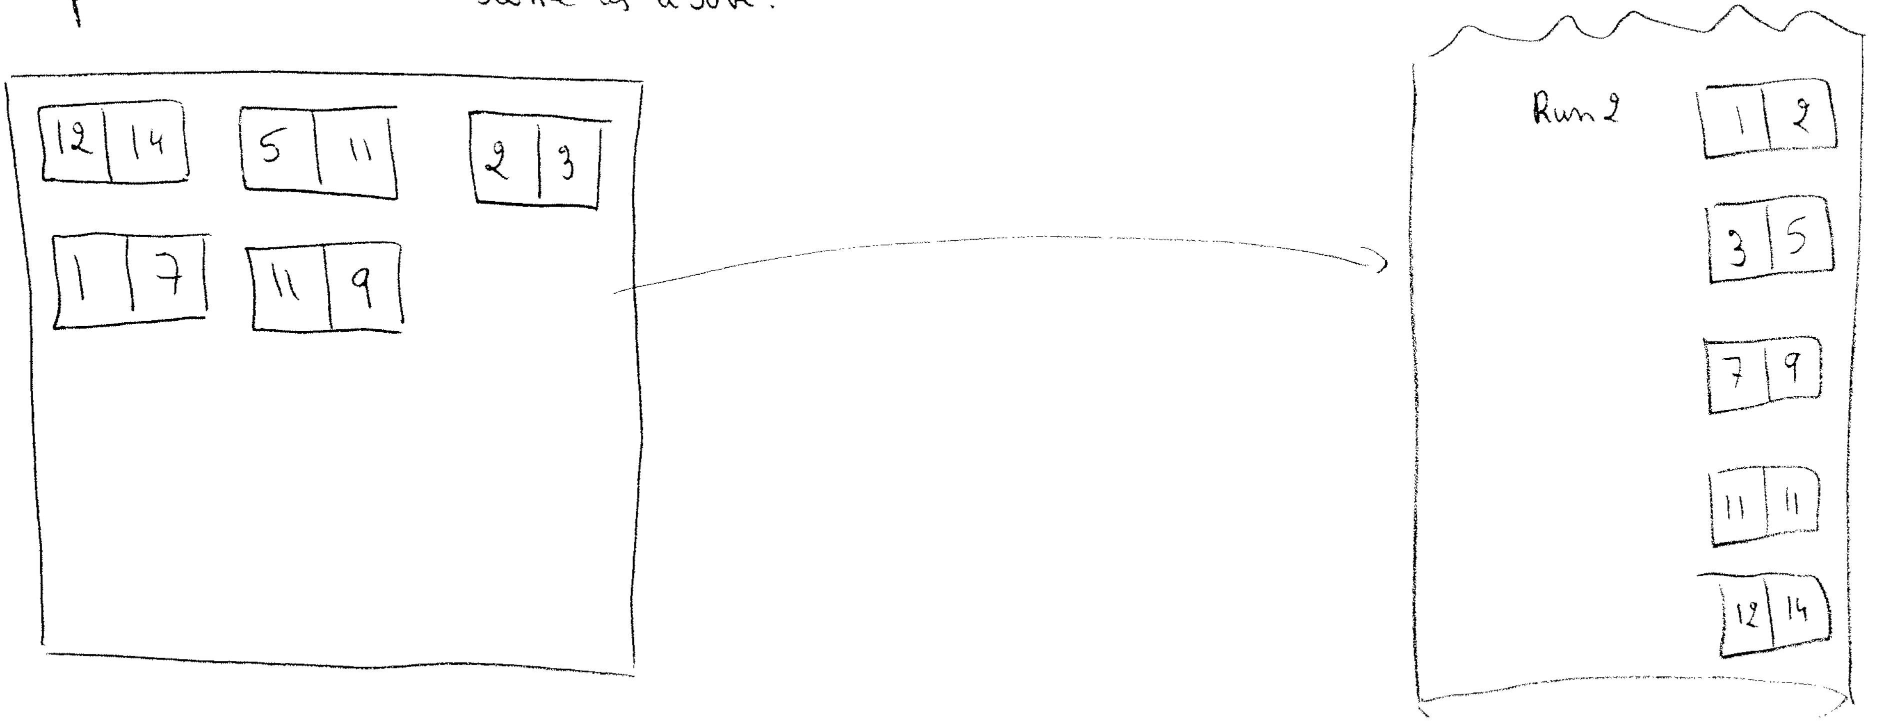

External Merge-Sort Example: continued

Merge runs together

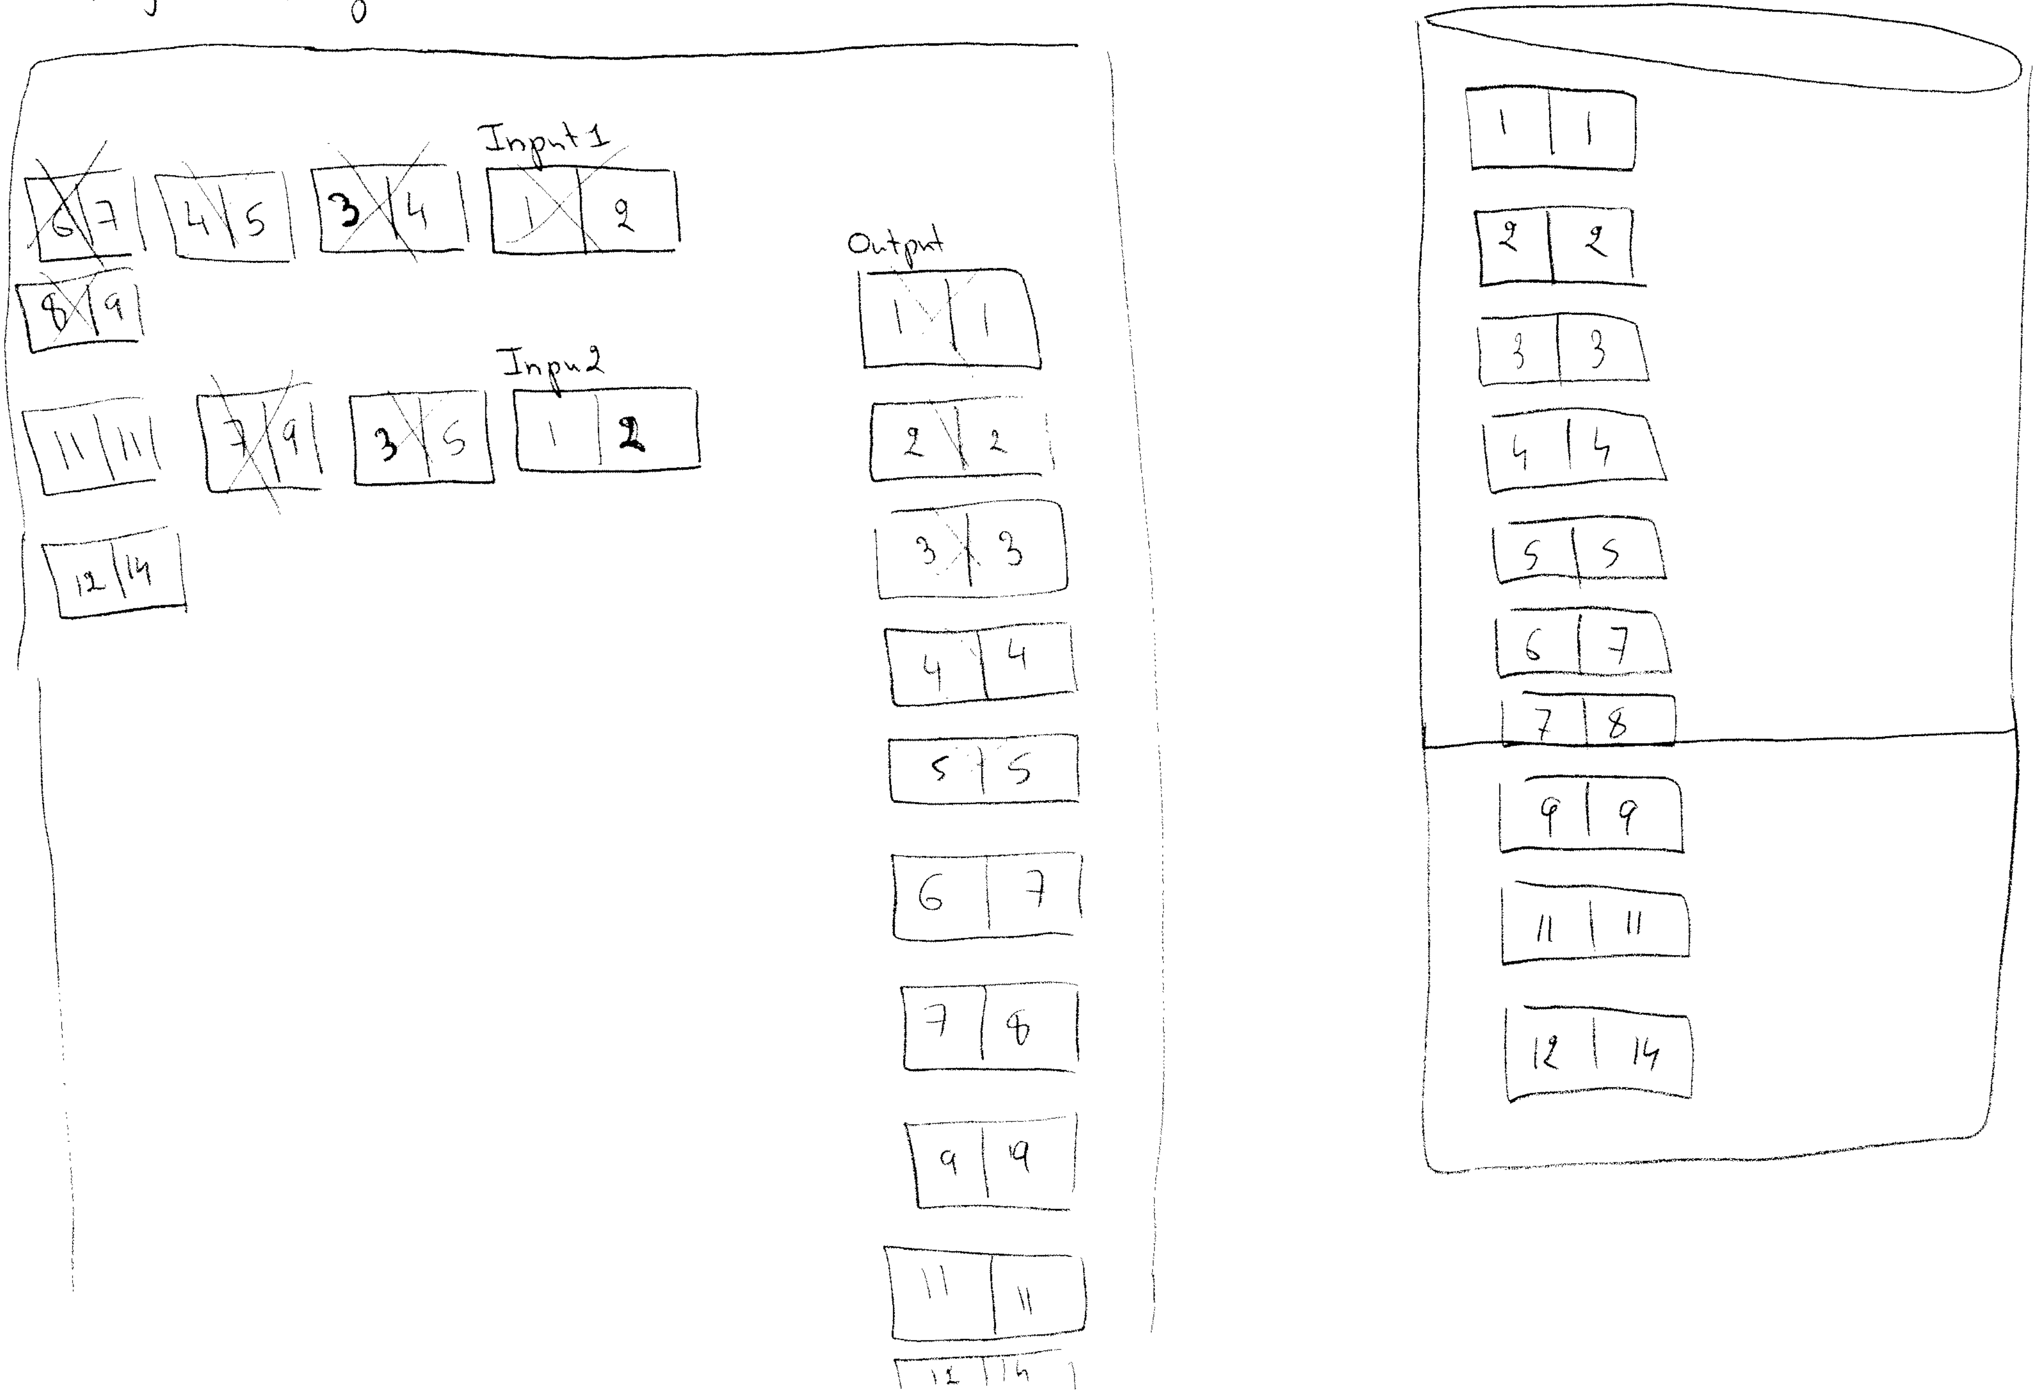

Sort-based join example

Approach 1: Sort R, Sort S, then merge

Approach 2: Join while merging sorted runs of R & S

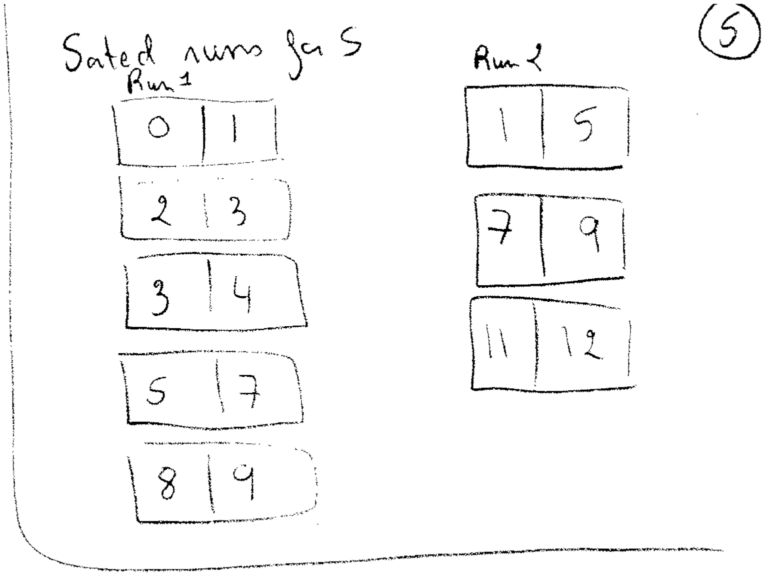
Sorted runs for S
Run 1
0 | 1
2 | 3
3 | 4
5 | 7
8 | 9
Run 2
1 | 5
7 | 9
11 | 12

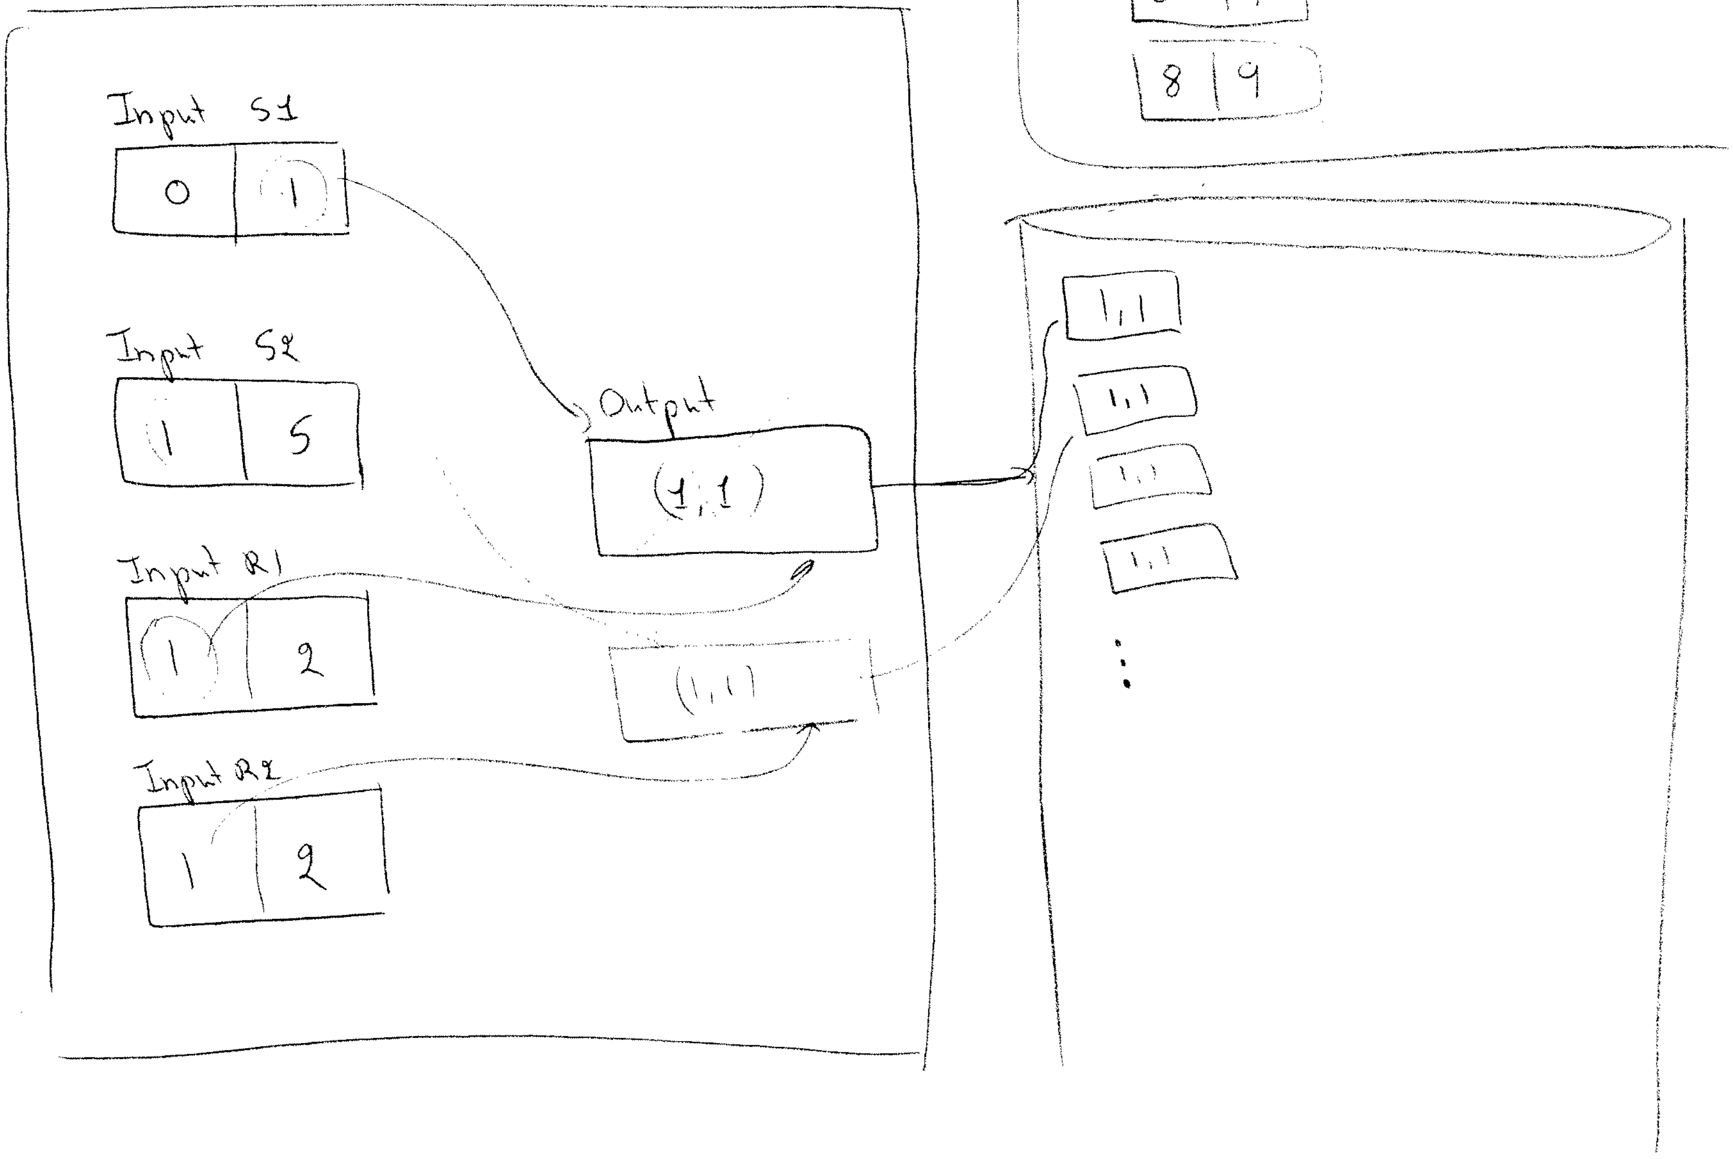
Input S1
0 | 1
Input S2
1 | 5
Input R1
1 | 2
Input R2
1 | 2
Output
(1,1)
(1,1)
1,1
1,1
1,1
1,1
...