Solutions

2.1

- Selection σ (a WHERE filter on rows) ✓.

2.2

- Projection π (keeping the named columns) ✓.

2.3

- WHERE filters rows before grouping; HAVING filters groups after ✓.

3.1

- Course query result: one row – `Databases | 3` ✓.
- LEFT JOIN: Algorithms 2, Calculus 5, Compilers 0, Databases 3, Linear Algebra 4, Topology 0 – 6 rows ✓.
- `SELECT student_id FROM Enrollment`: 14 rows ✓.
- DISTINCT: 8 rows – the 6 repeats collapse ✓.

3.2

- FROM/JOIN builds Department x Course rows following the foreign key ✓.
- WHERE keeps courses with credits ≥ 12 ✓.
- GROUP BY partitions by department name; SELECT projects name and the count; ORDER BY sorts descending ✓.
- The change is a HAVING clause: `HAVING COUNT(*) >= 2` ✓.
- Before: 2 result rows (one per department); after: still 2 rows – each department has exactly 2 such courses, so both survive the filter ✓.