

J

Join Order

Jingren Zhou
Alibaba Group, Hangzhou, China

Synonyms

Join order; Join sequence

Definition

A database query typically contains multiple joins. When joins (for example, inner joins) are commutative and/or associative, there can be more than one evaluation order for joins. The join order has an enormous impact on the query cost. One of the main responsibilities of the query optimizer is to determine the optimal join order for query evaluation.

Key Points

Choosing a good join order is very important to achieve a good query performance. One important consideration for choosing a join order is to reduce the size of intermediate results as much as possible. For example, it is beneficial to first evaluate a join that returns the least result.

Other considerations include join methods, data properties, and access methods, etc.

Depending on the choice of join orders, query plans can be of different shapes.

- *Left-Deep* query plans use a base table as the inner table for each join.
- *Right-Deep* query plans use a base table as the outer table for each join.
- *Bushy* query plans use the intermediate result from other joins as join inputs.

Query optimizers typically use dynamic programming and heuristics to determine join orders.

Cross-References

- [Evaluation of Relational Operators](#)
- [Parallel Join Algorithms](#)
- [Query Optimization](#)

Recommended Reading

1. Mishra P, Eich MH. Join processing in relational databases. ACM Comput Surv. 1992;24(1):63–113.
2. Selinger PG, Astrahan MM, Chamberlin DD, Lorie RA, Price TG. Access path selection in a relational database management system. In: Proceedings of the ACM SIGMOD International Conference on Management of Data; 1979. p. 23–34.