Index Join

Jingren Zhou
Alibaba Group, Hangzhou, China

Synonyms

[Index join](#); [Index loop join](#); [Index nested loop join](#)

Definition

The index join is a variant of the nested loop join in database systems using index. The join predicate can be either an equality predicate or a range predicate. The algorithm starts with reading the *outer* relation R . For each tuple $R \in R$, instead of scanning the entire *inner* relation S , an index on S is used to find matching tuples and add them to the result.

Key Points

An index on S is applicable for an index join if one join attribute is the leading indexed key of the

index. If the join predicate is an equality predicate, an index lookup is performed for each outer tuple. If the join predicate is a range predicate, for each outer tuple, an index seek is performed to locate the first matching tuple, followed by an index scan for the rest matching tuples. Compared with the nested loop join, the index join saves disk I/Os for reading the entire inner relation S multiple times.

Cross-References

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

Recommended Reading

1. Mishra P, Eich MH. Join processing in relational databases. *ACM Comput Surv.* 1992;24(1):63–113.