

## The effect of collisions on the performance of reduction filters.

**J.M. Morrissey**

School of Computer Science  
University of Windsor  
Windsor, Ontario N9B 3P4  
joan@uwindsor.ca

**W.K. Osborn**

Department of Computer Science  
University of Manitoba  
Winnipeg, Manitoba R3T 2N2  
osborn@cs.umanitoba.ca

**Keywords:** optimization, hash-semijoin, evaluation.

**Topics:** Computer networks and communications, software engineering.

### Introduction

The optimization of general queries in a distributed database management system is an important research topic. The problem is to select the best sequence of database operations that will process the query and minimize costs. Approaches include algorithms which are join-based [1, 2], semijoin-based [3, 4, 5, 6, 7, 8, 9, 10], or a combination of both [11]. Many algorithms concentrate on special types of queries [5, 12, 13, 14, 15, 6]. Some algorithms have been developed which improve on strategies produced by other heuristics [16, 17]. However, it has been shown [18, 19] that finding an optimal solution for a given general query is NP-hard so most algorithms aim to produce an efficient but (perhaps) suboptimal evaluation strategy. A new approach is to use hash-semijoins, which employ reduction filters (also known as Bloom filters), to reduce costs. Bloom filters are first described in [20], and have mostly been used to improve the relational join operation [21, 22, 23, 24, 25]. Bloom filters have been applied to other relational operations [26, 27, 28] and file processing operations [29]. Mullin [30, 31] employs Bloom filters to improve the semijoin operation. His method is limited to the case where there are only

two sites involved in the query and it has not been expanded to cater for general queries. Tseng and Chen [32] present a new relational operator called a hash-semijoin. They propose a method which will transform a sequence of semijoins, produced by any existing semijoin-based algorithm, into a more cost effective sequence by replacing some of the semijoins by hash-semijoins. But their approach has not been experimentally evaluated.

### The algorithm

In our recent work [33, 34] we use reduction filters. Each filter is an array of bits that functions as a very compact representation of the values of a join attribute in a relation. A hash function is used to set bits in the filter. It is then used to identify tuples (in other relations) which cannot belong to the result. The relations can be reduced significantly in size, data transfers can be reduced and costs minimized.

Our algorithm, which is based on reduction filters, can process general queries consisting of an arbitrary number of relations and join attributes. Each query is represented by a graph and an adjacency list. Each relation is usually only processed once — to minimize data transfers. However, if a filter changes during use then certain relations must be processed again — a queue is used to record this information. The algorithm consists of two phases: during phase one the adjacency list is used to determine the order in which the filters are constructed

and used.

**Phase one:**

1. Select the relation with lowest in-degree for processing.
2. Scan adjacency list to see which filters must be constructed. If a filter is already available then concurrently use it to reduce the relation and produce all required filters.
3. If a filter has changed the use the following “filter rule”: if a filter for a relation changes then add that relation to the queue only if it has already been processed; it is not already on the queue; and it is not the most recently processed relation.
4. Use adjacency list to “remove edges” from query graph — that is, reduce the in-degree of each relation in the list by 1.
5. Mark relation as processed.

Repeat all steps until each relation has been processed once.

**Phase two:**

1. Remove relation from queue.
2. Reduce relation using all appropriate filters.
3. If a filter changes then use the “filter rule”.

Repeat all steps until the queue is empty.

This work has one serious limitation — we used a “perfect” hash function which means that collisions are not possible. Experiments showed that our algorithm could reduce data transfers very substantially [33]. However, in a database collisions in the reduction filter<sup>1</sup> are a reality and a problem. Such collisions mean that the hash-semijoin cannot achieve the same

---

<sup>1</sup> A collision occurs when two attribute values hash to the same address in the filter.

amount of reduction as a traditional semijoin would. In this paper we investigate the effect that collisions have on the performance of our algorithm.

## The experiments

We have an experimental framework consisting of a collection of software to generate random queries with certain characteristics. Software is also used to generate the relations needed to execute each query. We assume a distributed relational database management system connected via a point to point network. There is no fragmentation nor replication. We consider only select-project-join (SPJ) queries, each consisting of a number of relations located at different sites which must be joined and the result made available at the query site. To evaluate, we compare our method not against another algorithm but against the effects of a “full reducer” — an algorithm which reduces the relations to just those tuples which can be part of the final result. We believe that this method is more objective and gives a better assessment of the performance of the algorithm.

We designed a set of experiments to test the performance with different collision rates; specifically 0% through 9% and then 10% through 60% collisions. (Our framework allows us to specify and guarantee a certain percentage of collisions in the filter. For full details see [35]). At each collision rate we generated and executed 600 random queries, where each relation has between 200 and 600 tuples; the active domains had between 150 and 250 values; the selectivity of each attribute was set between 0.5 and 0.95; and each relation had approximately 75% of the joining attributes. For convenience the queries are grouped by type, where type 3-2 means that the query has

three relations and two join-attributes. The result tables are presented below:

Type	0%	1%	2%	3%	4%
3-2	83.54	83.48	84.60	84.93	84.69
3-3	95.34	97.38	93.95	92.48	95.15
3-4	99.26	100	100	99.81	99.60
4-2	91.51	92.37	94.92	95.69	94.88
4-3	99.05	99.09	98.66	97.21	98.29
4-4	100	100	100	99.87	100
5-2	99.55	96.72	97.59	98.55	96.91
5-3	100	100	99.85	99.82	99.96
5-4	100	100	100	100	100
6-2	98.55	98.92	98.74	98.98	99.03
6-3	100	100	99.97	99.99	100
6-4	100	100	100	100	100
Avg	97.23	99.33	97.36	97.27	97.38

Table 1 Average reduction: 0% - 4%

Type	5%	6%	7%	8%	9%
3-2	83.87	84.19	81.94	81.59	83.36
3-3	93.95	95.56	93.07	96.15	92.95
3-4	100	100	100	99.51	100
4-2	94.97	93.03	90.57	91.85	91.68
4-3	98.89	97.96	99.36	98.52	99.21
4-4	100	99.72	99.77	100	100
5-2	97.87	96.40	96.75	97.68	96.86
5-3	99.99	99.84	98.70	98.65	99.98
5-4	100	100	100	100	100
6-2	99.67	99.11	99.12	98.30	98.11
6-3	100	99.76	99.74	100	99.93
6-4	100	100	100	100	100
Avg	97.43	97.13	96.59	96.85	96.84

Table 2 Average reduction: 5% - 9%

Type	10%	20%	30%	40%	50%	60%
3-2	85.26	80.11	75.21	73.31	72.74	69.56
3-3	95.75	93.15	87.27	90.53	86.10	88.58
3-4	99.55	99.02	98.95	97.92	98.58	97.11
4-2	93.54	92.13	89.29	86.25	87.43	81.23
4-3	97.45	97.83	96.41	98.34	94.53	96.62
4-4	100	100	100	100	99.63	98.31
5-2	96.01	96.39	95.39	92.22	91.21	89.37
5-3	99.64	98.38	98.98	96.56	98.72	98.10
5-4	100	100	100	100	100	99.77
6-2	99.28	97.88	96.53	98.55	97.41	95.81
6-3	100	100	99.58	99.13	99.59	99.44
6-4	100	100	100	100	100	100
Avg	97.21	96.24	94.80	94.40	93.83	92.82

Table 3 Average reduction: 10% - 60%

## Conclusions

The results show that when the collision rate is between 0% and 10% then we can achieve, on average, a 97% reduction in the relations, and in approximately 83% of all queries the relations were fully reduced<sup>2</sup>. It would appear that the query type has more of an influence on the reduction rate than the percentage of collisions. This is not an unexpected result since, as the number of joining attributes increases we use more filters on a relation and this reduces the impact of the collisions.

When we looked at collision rates in the range 10% to 60% we found that we could still achieve 95% reduction on average, and approximately 68% of all queries were fully reduced.

The conclusions of this paper are very significant. To our knowledge no one has ever examined the effect of collisions on the performance of a hash-semijoin based algorithm.

<sup>2</sup> The relations had been reduced to just those tuples which can participate in the final join.

Our main conclusion is that the collisions, even at 60%, do not significantly affect performance.

A full discussion of the results will be presented at conference time.

## References

- [1] J. Ahn and S. Moon, "Optimizing joins between two fragmented relations on a broadcast local network.," *Info. Syst.*, vol. 16(2), pp. 185–198, 1991.
- [2] P. Legato, G. Paletta, and L. Palopoli, "Optimization of join strategies in distributed databases," *Info. Syst.*, vol. 16(4), pp. 363–374, 1991.
- [3] P. Bernstein, N. Goodman, E. Wong, C. Reeve, and J. Rothnie, "Query processing in a system for distributed databases (SDD-1)," *ACM Trans. on Database Systems*, vol. 6(4), pp. 105–128, 1981.
- [4] P. Black and W. Luk, "A new heuristic for generating semi-join programs for distributed query processing," *IEEE COMPSAC*, vol. 581–588, 1982.
- [5] P. Apers, A. Hevner, and S. Yao, "Optimization algorithms for distributed queries," *IEEE Transactions on Software Engineering*, 9(1), pp. 51–60, 1983.
- [6] L. Chen and V. Li, "Improvement algorithms for semi-join query processing programs in distributed database systems.," *IEEE Trans. on Computers*, vol. 33(11), pp. 959–967, 1984.
- [7] H. Kang and N. Roussopoulos, "Using 2-way semi-joins in distributed query processing," in *Proc. 3rd Int. Conf. on Data Engineering*, pp. 644–650, 1987.
- [8] N. Roussopoulos and H. Kang, "A pipeline n-way join algorithm based on the 2-way semi-join program," *IEEE Trans. on Knowledge and Data Engineering*, vol. 3(4), pp. 486–495, 1991.
- [9] L. Chen and V. Li, "Domain-specific semi-join: a new operation for distributed query processing," *Info. Sci.*, vol. 52, pp. 165–183, 1990.
- [10] C. Wang, V. Li, and A. Chen, "Distributed query optimization by one-shot fixed precision semi-join execution," in *Proc. 7th Int. Conf. on Data Engineering*, pp. 756–763, 1991.
- [11] M. Chen and P. S. Yu, "Combining join and semi-join operations for distributed query processing," *IEEE Transactions on Knowledge and Data Engineering*, vol. 5(3), pp. 534–542, 1993.
- [12] D. Chiu, P. Bernstein, and Y. Ho, "Optimizing chain queries in a distributed database system," *Siam Journal of Computing*, vol. 13(1), pp. 116–134, 1984.
- [13] A. Chen and V. Li, "A method for interpreting tree queries into optimal semijoin expressions," in *ACM SIGMOD*, 1980.
- [14] D. Chiu and Y. Ho, "Optimizing star queries in a distributed database system," in *VLDB*, pp. 959–967, 1984.
- [15] C. Yu, Z. Ozsoyoglu, and K. Kam, "Optimization of distributed tree queries," *J. Comp. Sys. Sci.*, vol. 29(3), pp. 409–445, 1984.
- [16] P. Boderick, J. Pyra, and J. Riordan, "Correcting execution of distributed queries," in *Proc. of 2nd Int. Symp. on Databases in Parallel and Distributed Systems*, pp. 192–201, 1990.
- [17] P. Boderick, J. Riordan, and J. Pyra, "Deciding to correct distributed query processing," *IEEE Trans. on Knowledge and Data Engineering*, vol. 4(3), pp. 253–265, 1992.

- [18] A. Hevner, *The optimization of query processing in distributed database systems*. PhD thesis, Perdue University, 1980.
- [19] C. Wang and M. Chen, "On the complexity of distributed query optimization," tech. rep., IBM Technical Report RC 18671, 1993.
- [20] B. Bloom, "Space/time tradeoffs in hash coding with allowable errors," *Comm. ACM*, vol. 13(7), pp. 422–426, 1970.
- [21] D. DeWitt, S. Ghandeharizadeh, D. Schneider, A. Bricker, H. Hsiao, and R. Rasmussen, "The GAMMA database machine project," *IEEE Transactions on Knowledge and data engineering*, pp. 44–62, 1990.
- [22] G. Z. Qadah, *Filter-based join algorithms on uniprocessor and distributed-memory multiprocessor database machines*, vol. 303 of *Lecture Notes in Computer Science*, pp. 388–413. Springer-Verlag, 1988.
- [23] G. Z. Qadah and K. Irani, "The join algorithms on a shared-memory multiprocessor database machine," *IEEE Transactions on Software Engineering*, pp. 1168–1683, 88.
- [24] P. Valduriez and G. Gardarin, "Join and semijoin algorithms for a multiprocessor database machine," *ACM Transactions on database systems*, pp. 133–161, 1984.
- [25] J. Mullin, "Estimating the size of a relational join," *Information Systems*, vol. 18(3), pp. 189–196, 1993.
- [26] G. Graefe and K. Ward, "Dynamic query evaluation plans," in *ACM SIGMOD*, pp. 358–366, 1989.
- [27] G. Graefe, "Query evaluation techniques for large databases," *ACM computing surveys*, pp. 73–170, 1993.
- [28] C. Mohan, D. Haderle, Y. Wang, and J. Cheng, *Single table access using multiple indexes: optimization, execution and concurrency control techniques.*, vol. 416 of *Lecture Notes in Computer Science*, pp. 29–43. Springer-Verlag, 1990.
- [29] D. Severance and G. Lohman, "Differential files: their application to the maintenance of large databases," *ACM Transactions on database systems*, pp. 257–267, 1976.
- [30] J. Mullin, "(1983) A second look at bloom filters," *Comm. ACM*, vol. 26(8), pp. 570–571, 1983.
- [31] J. Mullin, "Optimal semijoins for distributed database systems," *IEEE trans. on software eng.*, vol. 16(5), pp. 558–560, 1990.
- [32] J. Tseng and A. P. Chen, "Improving distributed query processing by hash-semijoins," *Journal of Information Science and Engineering*, vol. 8, pp. 525–540, 1992.
- [33] J. Morrissey and W. Osborn, "Distributed query optimization using reduction filters," in *Proceedings of the IEEE Canadian Conference on Electrical and Computer Engineering*, (University of Waterloo, May 1998), pp. 707–710.
- [34] J. Morrissey and W. Osborn, "Experiments with the use of reduction filters in distributed query optimization," in *Proceedings of the 9th International Conference on Parallel and Distributed Computing and Systems (PDCS'97)*, (Washington, D.C., October 1997), pp. 327–330.
- [35] W. Osborn, "The use of reduction filters in distributed query optimization," Master's thesis, The University of Windsor, 1998.