

# Heuristic Scheduling of Concurrent Data Mining Queries\*

Marek Wojciechowski, Maciej Zakrzewicz

Poznan University of Technology  
Institute of Computing Science  
ul. Piotrowo 3a, 60-965 Poznan, Poland  
{marek,mzakrz}@cs.put.poznan.pl

**Abstract.** Execution cost of batched data mining queries can be reduced by integrating their I/O steps. Due to memory limitations, not all data mining queries in a batch can be executed together. In this paper we introduce a heuristic algorithm called CCFull, which suboptimally schedules the data mining queries into a number of execution phases. The algorithm significantly outperforms the optimal approach while providing a very good accuracy.

## 1 Introduction

Multiple Query Optimization (MQO) [8] is a database research area that focuses on optimizing sets of queries together by executing their common expressions only once in order to save query execution time. Many exhaustive and heuristic algorithms have been proposed for traditional MQO [3][6][7]. A specific type of a database query is a Data Mining Query (DMQ) [5], which describes a data mining task. It defines constraints on the data to be mined and constraints on the patterns to be discovered. DMQs are submitted for execution to a Knowledge Discovery Management System KDDMS [5], which is a Database Management System (DBMS) extended with data mining functionality. Traditional KDDMSs execute DMQs serially and do not try to share any common expressions between different DMQs.

DMQs are often processed in batches of dozens of queries, executed during low user activity time. Queries in a batch may show many similarities to each other, e.g., their source data sets may overlap. If such queries were executed serially, then it would be likely that many I/O operations were wasted because the same database blocks were required by multiple DMQs. If I/O steps of different DMQs were integrated and performed once, then it would be possible to decrease the overall execution cost and time of the whole batch. One of the methods to process batches of DMQs is Apriori Common Counting (ACC) [9], focused on frequent itemset discovery queries [1]. ACC is based on Apriori algorithm [2] and it integrates the steps of candidate support counting – all candidate hash trees for multiple DMQs are loaded into memory and the database is scanned only once. Basic ACC [9] assumes

---

\* This work was partially supported by the grant no. 4T11C01923 from the State Committee for Scientific Research (KBN), Poland.

that all DMQs fit in memory, which is not a common case, at least for initial Apriori iterations. If the memory can hold only a subset of all DMQs, then it is necessary to divide/schedule the DMQs into subsets called phases [10]. The way such scheduling is done determines the overall cost of batched DMQs execution. To solve the scheduling problem, in [10] we proposed an “initial” heuristic algorithm, called *CCRecursive*. According to our experiments *CCRecursive* offers acceptable accuracy and on average outperforms the optimal scheduling algorithm. However, in particular situations its execution time could increase significantly due to its recursive nature, which is the motivation for seeking novel, more predictable solutions. In this paper we present and evaluate another heuristic algorithm for scheduling data mining queries to be executed by ACC, called *CCFull*.

## 1.1 Related Work

Multiple-query optimization has been extensively studied in the context of database systems (see e.g. [8]), however very little work has been done on optimizing sets of data mining queries. To the best of our knowledge, apart from the ACC method discussed in this paper, the only other multiple query processing scheme for data mining queries is Mine Merge, presented in one of our previous papers [11]. In contrast to ACC, Mine Merge is independent of a particular frequent itemset mining algorithm. However, it was proven very sensitive to data distribution and less predictable than ACC.

## 2 Preliminaries and Problem Statement

**Data mining query.** A *data mining query* is a tuple  $DMQ = (R, a, \Sigma, \Phi, \beta)$ , where  $R$  is a relation,  $a$  is an attribute of  $R$ ,  $\Sigma$  is a condition involving the attributes of the relation  $R$ ,  $\Phi$  is a condition involving discovered patterns, and  $\beta$  is the minimum support threshold. The result of the data mining query is a set of patterns discovered in  $\pi_a \sigma_{\Sigma} R$  and satisfying  $\Phi$ .

**Problem statement.** Given is a set of data mining queries  $DMQ = \{dmq_1, dmq_2, \dots, dmq_n\}$ , where  $dmq_i = (R, a, \Sigma_i, \Phi_i, \beta_i)$ ,  $\Sigma_i$  has the form “ $(l_{1min}^i < a < l_{1max}^i) \vee (l_{2min}^i < a < l_{2max}^i) \vee \dots \vee (l_{kmin}^i < a < l_{kmax}^i)$ ”,  $l_{*}^i \in dom(a)$  and there exist at least two data mining queries  $dmq_i = (R, a, \Sigma_i, \Phi_i, \beta_i)$  and  $dmq_j = (R, a, \Sigma_j, \Phi_j, \beta_j)$  such that  $\sigma_{\Sigma_i} R \cap \sigma_{\Sigma_j} R \neq \emptyset$ . The problem of *multiple query optimization* of  $DMQ$  consists in generating such an algorithm to execute  $DMQ$  that has the lowest I/O cost.

**Apriori Common Counting (ACC).** If the set of data mining queries was executed serially, i.e. one data mining query at a time, then the total execution cost would be the sum of execution costs of data selection formulas for each data mining query separately. ACC executes a set of data mining queries by integrating their I/O operations. It is based on the traditional Apriori approach to discover frequent

itemsets. In the first step, for each data mining query we build a separate hash tree for 1-candidates. Next, for each distinct data selection formula we scan its corresponding database partition and we count candidates for all the queries that contain the formula. Such a step is performed for 2-candidates, 3-candidates, etc. Notice that if a given distinct data selection formula is shared by many queries, then its corresponding database partition is read only once. An overview of ACC is shown in Fig. 1.

```

for ( $i=1; i \leq n; i++$ ) /*  $n$  = number of data mining queries */
   $C_1^i = \{\text{all 1-itemsets from } \sigma_{s_1 \cup s_2 \cup \dots \cup s_k} R, \forall s_j \in S: (dmq_{i,s_j}) \in E\}$  /* generate 1-candidates */
  for ( $k=1; C_k^1 \cup C_k^2 \cup \dots \cup C_k^n \neq \emptyset; k++$ ) do begin
    for each  $s_j \in S$  do begin
       $CC = \bigcup C_k^i: (dmq_{i,s_j}) \in E;$  /* select the candidates to count now */
      if  $CC \neq \emptyset$  then  $\text{count}(CC, \sigma_{s_j} R);$ 
    end
    for ( $i=1; i \leq n; i++$ ) do begin
       $F_k^i = \{C \in C_k^i \mid C.\text{count} \geq \text{minsup}^i\};$  /* identify frequent itemsets */
       $C_{k+1}^i = \text{generate\_candidates}(F_k^i);$ 
    end
  end
  for ( $i=1; i \leq n; i++$ ) do
     $\text{Answer}^i = \bigcup_k F_k^i;$  /* generate responses */

```

**Fig. 1.** Apriori Common Counting

### 3 Heuristic Scheduling of Concurrent Data Mining Queries

The basic ACC assumes unlimited memory and therefore the candidate hash trees for all DMQs can completely fit in memory. If, however, the memory is limited, then ACC execution must be divided into multiple *phases*, so that in each phase only a subset of DMQs is processed. In such a case, the key question to answer is: which data mining queries from the set should be executed together in one phase and which data mining queries can be executed in different phases? We refer to the task of data mining queries partitioning as to *data mining query scheduling*.

The problem of data mining query scheduling is a combinatorial problem which can be solved by generating all possible schedules and then choosing the best one. Such approach can be used for a small number of data mining queries, however, for a realistic case it is infeasible. The number of all possible schedules is determined by the *Bell number*, e.g., for 13 queries we get over 4 million schedules. Therefore, we propose a heuristic algorithm *CCFull*, which quickly finds a suboptimal schedule.

#### 3.1 Algorithm CCFull

In the first step we generate a *gain graph* for the set of data mining queries. The gain graph is a full hypergraph, in which vertices represent the data mining queries while

edges are described with weights which represent the amount of I/O cost reduction to be achieved if data mining queries connected with the edge were executed together (in the same phase). If common execution of given data mining queries results in no reduction of I/O cost, the weight of the connecting edge is zero. A sample gain graph is shown in Fig. 2. For example, it can be noticed that common execution of the data mining queries  $dmq_0$ ,  $dmq_2$ , and  $dmq_3$  would reduce the total I/O cost by 16 units (the weight of the connecting hyperedge) compared with the sequential execution, since for  $dmq_0$  and  $dmq_2$  the cost of redundant I/O operations is 5 units, for  $dmq_2$  and  $dmq_3$  the cost of redundant I/O operations is 8 units, and for  $dmq_0$  and  $dmq_3$  the cost of redundant I/O operations is 3 units. Using the same example, it can be also noticed, that common execution of only the data mining queries  $dmq_1$  and  $dmq_2$  provides no cost reduction (the weight of the connecting hyperedge is zero).

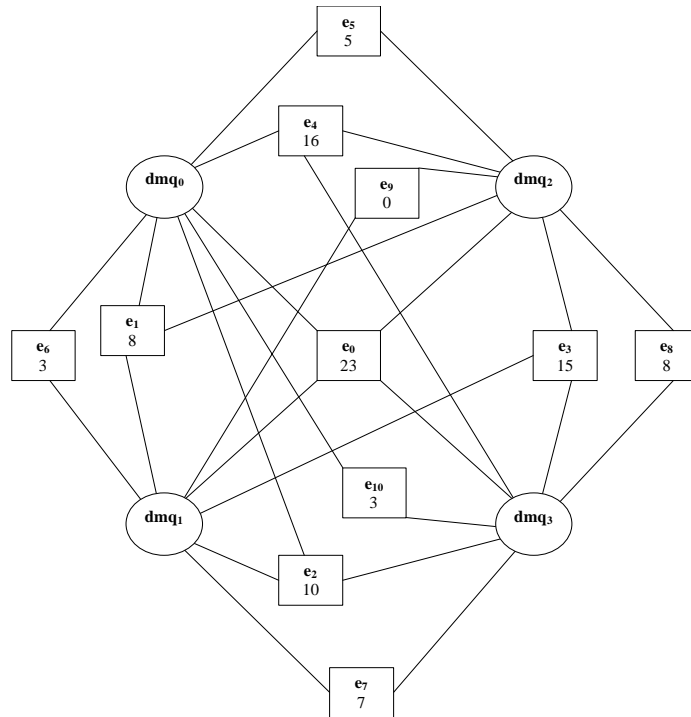


Fig. 2. Sample gain graph

The gain graph can be generated using the algorithm *GenerateGainGraph* shown in Fig. 3. The algorithm takes two arguments: the set of all distinct data selection formulas and the set of all data mining queries. First, the algorithm builds a full hypergraph whose nodes are the data mining queries (line 1). Each hyperedge receives the weight of zero, initially (line 3). Then, for each hyperedge  $e$ , we create a set  $P$  of distinct data selection formulas involved in all data mining queries connected with the hyperedge  $e$  (line 4). I/O costs for executing the distinct data selection

formulas from  $P$  are then summarized and the result is assigned to the hyperedge  $e$  weight (line 5 and 6).

```

GenerateGainGraph( $S, DMQ$ ):
begin
1.   generate a full hypergraph  $G = \{V, E\}$ ,  $V = DMQ$ 
2.   for each  $e \in E$  do begin
3.      $e.gain = 0$ ;
4.      $P = \{s_i \in S \mid \exists dmq_j \in e, dmq_j = (R, a, \Sigma_j, \Phi_j, \beta_j), s_i \subseteq \Sigma_j\}$ 
5.     for each  $s \in P$  do begin
6.        $e.gain += cost(s) * (|\{dmq_j: dmq_j \in e, dmq_j = (R, a, \Sigma_j, \Phi_j, \beta_j), s_i \subseteq \Sigma_j\}| - 1)$ 
8.     end
9.   end
10.  return  $G$ 
end

```

**Fig. 3.** Gain graph generation algorithm

After having created the gain graph, *CCFull* performs the following steps. All hyperedges are sorted in descending order according to their weights. Next, *CCFull* iterates over the hyperedges and checks if data mining queries connected with the current hyperedge have been already scheduled. If none of the data mining queries has been scheduled so far, and if their hash trees fit in memory, then a new phase is generated and the data mining queries are assigned to it. Otherwise, if only some of the data mining queries have been already scheduled to different phases, then *CCFull* tries to combine all those phases together with the unscheduled data mining queries. If such combined phase does not fit in memory, then the current hyperedge is ignored and *CCFull* continues with the next one. The algorithm ends when all hyperedges are processed. The algorithm *CCFull* is shown in Fig. 4.

The detailed steps of the algorithm from Fig. 4 are the following. In line (1) we initialize the set of scheduled phases – we start with the empty set. In line (2) we sort the list  $E$  of hyperedges from the gain graph. Hyperedges with weights equal to zero are removed from the list. In line (3) a loop starts, which iterates over the list of hyperedges. In line (4) we select all data mining queries which are connected with the current hyperedge ( $tmpV$ ). In line (5) we test if any of the selected data mining queries belongs to any of the phases scheduled so far. If not, then in line (7) we create a new candidate phase containing all the data mining queries from  $tmpV$ . Otherwise, in line (9) we create a new candidate phase containing both all the data mining queries from  $tmpV$  and data mining queries from earlier scheduled phases, to which any of the  $tmpV$  data mining queries was also scheduled. In line (10) we check if hash trees of all the data mining queries from the new candidate phase fit in memory ( $MEMSIZE$  is the available memory size). If this condition is satisfied, then in lines (11) and (12) we append the new candidate phase to the current set of scheduled phases  $Phases$ , possibly replacing some of the existing phases (when multiple phases are combined). In line (13), for each data mining query which has not been scheduled we create a new phase. In step (14) we return the generated phases.

```

CCFull( $G=(V,E)$ ):
  begin
1.  $Phases \leftarrow \{\emptyset\}$ 
2. sort  $E = \langle e_1, e_2, \dots, e_k \rangle$  in desc. order w.r. to  $e_i$ .gain, ignore edges with zero gains
3. for each  $e_i$  in  $E$  do begin
4.    $tmpV \leftarrow \{v \in V \mid v \in e_i\}$ 
5.   if ( $|\{p \in Phases \mid p \cap tmpV \neq \emptyset\}| = 0$ ) then
6.      $commonPhases \leftarrow \emptyset$ 
7.      $newPhase \leftarrow tmpV$ 
8.   else
9.      $commonPhases \leftarrow \{p \in Phases \mid p \cap tmpV \neq \emptyset\}$ 
10.     $newPhase \leftarrow tmpV \cup \bigcup_{p \in commonPhases} p$ 
11.   end if
12.   if ( $treeSize(newPhase) \leq MEMSIZE$ ) then
13.     $Phases \leftarrow Phases - commonPhases$ 
14.     $Phases \leftarrow Phases \cup newPhase$ 
15.   end if
16. end
17. add phase for each unscheduled query
18. return Phases
19. end

```

**Fig. 4.** *CCFull* algorithm

### 3.2 Example

Consider scheduling of data mining queries from Fig. 2. For the sake of simplicity, assume that hash tree sizes are 10MB for each data mining query and the available memory is 20MB.

Hyperedges of the gain graph are sorted according to their weights (skipping zero-weighted hyperedges):  $\langle e_0, e_4, e_3, e_2, e_1, e_8, e_7, e_5, e_6, e_{10} \rangle$ . In the first iteration we select the hyperedge  $e_0$ , which is connecting the data mining queries  $dmq_0$ ,  $dmq_1$ ,  $dmq_2$  and  $dmq_3$ . None of the data mining queries has been scheduled so far, and total size of their hash trees is 40MB, exceeding the available memory. Therefore, the algorithm ignores the hyperedge and starts another iteration.

In the second iteration we select the hyperedge  $e_4$ , which is connecting the data mining queries  $dmq_0$ ,  $dmq_2$  and  $dmq_3$ . None of the data mining queries has been scheduled so far, and total size of their hash trees is 30MB, exceeding the available memory again. Therefore, the algorithm ignores the hyperedge and starts another iteration. In a similar way the iterations over the hyperedges  $e_3$ ,  $e_2$  and  $e_1$  are performed – total sizes of hash trees exceed the available memory.


Yet in the sixth iteration the algorithm will behave in a different way. We select the hyperedge  $e_8$ , which is connecting the data mining queries  $dmq_2$  and  $dmq_3$ . The total size of their hash trees is 20MB, so a new phase is created:  $\{dmq_2, dmq_3\}$ . In the next iteration we select the hyperedge  $e_7$ , which is connecting the data mining queries

$dmq_1$  and  $dmq_3$ . Since  $dmq_3$  already belongs to a scheduled phase, we try to replace the existing phase  $\{dmq_2, dmq_3\}$  with a new one:  $\{dmq_1, dmq_2, dmq_3\}$ . We are unsuccessful because the total size of hash trees for the data mining queries is 30MB, what exceeds the available memory. In the next iteration we select the hyperedge  $e_5$ , and again we are unsuccessful when trying to replace the existing phase  $\{dmq_2, dmq_3\}$  with a new phase  $\{dmq_0, dmq_2, dmq_3\}$ . The next iteration operates on the hyperedge  $e_6$ , which is connecting the data mining queries  $dmq_0$  and  $dmq_1$ . These data mining queries do not belong to any of the existing phases and total size of their hash trees is 20MB. Therefore, a new phase is created:  $\{dmq_0, dmq_1\}$ .


In the last iteration we select the hyperedge  $e_{10}$ , which is connecting the data mining queries  $dmq_0$  and  $dmq_3$ . Since both data mining queries have already been scheduled to some phases, the algorithm tries to combine the existing phases  $\{dmq_2, dmq_3\}$  and  $\{dmq_0, dmq_1\}$ . However, the phases are not merged since the total size of hash trees of their data mining queries is 40MB and exceeds the available memory. The algorithm has completed. The constructed scheduling of the four data mining queries consists of 2 phases:  $\{dmq_2, dmq_3\}$  and  $\{dmq_0, dmq_1\}$ .

## 4 Experimental Evaluation

In order to evaluate performance and accuracy of the *CCFull* scheduling algorithm we performed several experiments using the MSWeb dataset from the UCI KDD Archive [4]. The experiments were conducted on a PC with AMD Duron 1.2 GHz processor and 256 MB of main memory. The datasets used in all experiments resided in flat files on a local disk. Memory was intentionally restricted to 10kB-50kB. Each experiment was repeated 100 times. The queries were randomly generated.



**Fig. 5.** Accuracy of data mining query scheduling algorithms



**Fig. 6.** Execution time of data mining query scheduling algorithms

Fig. 5 shows disk I/O costs of schedules generated by the optimal scheduling algorithm, by the *CCFull* algorithm, and by a random algorithm (which randomly builds phases from queries). For example, for the set of 10 data mining queries, the *CCFull* algorithm misses the optimal solution by only 6%.

Fig. 6 illustrates execution times for the optimal scheduling algorithm and for *CCFull*. Notice that the optimal algorithm needs ca. 1000s to schedule 12 data mining queries while *CCFull* executes in 30s.

Comparing accuracy and performance of *CCFull* with the experimental results obtained for *CCRecursive* (reported in [10]), we have to admit that on average *CCRecursive* outperforms *CCFull*, while offering a slightly better accuracy. However, in particular, rare situations *CCRecursive* took more time to complete than the optimal scheduling algorithm due to its recursive nature. We have not observed such problems with *CCFull*, which makes it a more predictable solution.

## 5 Conclusions and Future Work

In this paper we have introduced a new heuristic algorithm *CCFull* to schedule data mining queries for Apriori Common Counting. The algorithm offers a significant reduction of execution time over the optimal algorithm while providing a very good accuracy, and is predictable compared to our previous heuristics *CCRecursive*.

*CCFull* assumes that the set of data mining queries is static. However, in a real system, new queries may arrive while other queries are being executed. In the future we plan to investigate methods allowing for dynamic scheduling of the arriving queries.

## References

1. Agrawal R., Imielinski T., Swami A: Mining Association Rules Between Sets of Items in Large Databases. Proc. of the 1993 ACM SIGMOD Conf. on Management of Data (1993)
2. Agrawal R., Srikant R.: Fast Algorithms for Mining Association Rules. Proc. of the 20th Int'l Conf. on Very Large Data Bases (1994)
3. Alsabbagh J.R., Raghavan V.V.: Analysis of common subexpression exploitation models in multiple-query processing. Proc. of the 10th ICDE Conference (1994)
4. Hettich S., Bay S. D.: The UCI KDD Archive [<http://kdd.ics.uci.edu>]. Irvine, CA: University of California, Department of Information and Computer Science (1999)
5. Imielinski T., Mannila H.: A Database Perspective on Knowledge Discovery. Communications of the ACM, Vol. 39, No. 11 (1996)
6. Jarke M.: Common subexpression isolation in multiple query optimization. Query Processing in Database Systems, Kim W., Reiner D.S. (Eds.), Springer (1985)
7. Roy P., Seshadri S., Sundarshan S., Bhohe S.: Efficient and Extensible Algorithms for Multi Query Optimization. ACM SIGMOD Intl. Conference on Management of Data (2000)
8. Sellis T.: Multiple query optimization. ACM Transactions on Database Systems, Vol. 13, No. 1 (1988)
9. Wojciechowski M., Zakrzewicz M.: Evaluation of Common Counting Method for Concurrent Data Mining Queries. Proc. of the 7th ADBIS Conference (2003)
10. Wojciechowski M., Zakrzewicz M.: Data Mining Query Scheduling for Apriori Common Counting. Proc. of the 6th Int'l Baltic Conf. on Databases and Information Systems (2004)
11. Wojciechowski M., Zakrzewicz M.: Evaluation of the Mine Merge Method for Data Mining Query Processing. Proc. of the 8th ADBIS Conference (2004)