

Multi-source materialized views maintenance: multi-level views

Josep Silva, Jorge Belenguer, Matilde Celma

Computer Science Department
Technical University of Valencia
Camino de Vera s/n
E-46021 Valencia – Spain
{jsilva, jorbefa, mcelma}@dsic.upv.es

Abstract. In many information systems, the databases that make up the system are distributed in different modules or branch offices according to the requirements of the business enterprise. In these systems, it is often necessary to combine the information of all the organisation's databases in order to perform analysis and make decisions about the global operation. This is the case of Data Warehouse Systems. From a conceptual point of view, a Data Warehouse can be considered as a set of materialized views which are defined in terms of the tables stored in one or more databases. These materialized views store historical data that must be maintained in either real time or periodically by means of batch processes. During the maintenance process the systems must perform selections, projections, joins, etc. that can affect several databases. This is a complex problem since making a join among several tables requires (at least temporarily) having the information from these tables in the same place. This requires the Data Warehouse to store auxiliary materialized views that in many cases contain duplicated information. In this article, we study this problem, and we propose a method that minimizes the duplicated information in the auxiliary materialized views and also reduces the response time of the system.

Key words: Data Warehousing, materialized views, multi-source views, view fragmentation, multi-level views.

1 Introduction

The use of Data Warehouse systems is becoming one of the critical factors that determine the success of many companies and organisations. The information gathered in the Warehouse can be used to make decisions about the processes of the organization, and should therefore be consistent. The information should also be as up-to-date as possible. Having the information of the operational systems up-to-date makes the results of the queries carried out on the Warehouse Database to be closer to the reality of the organization.

Many previous approaches [3, 6, 7, 8, 12, 13, 14, 15] consider that the information of a Data Warehouse System consists of a set of materialized views that store information from one or more databases that the organization uses in its operational systems. The process of loading this information is usually done using daily batch processes at night in order to avoid the slowdown of the operation systems. However, in some cases, the organization needs to compare the historical information in the Warehouse with the most recent information available in the operational system and therefore the Warehouse must be maintained in real time.

Maintenance of materialized views in general, and of Data Warehouses in particular, is a very relevant problem that has been studied in many works. For instance, [14, 15] deal with the problem of updating materialized views in real time, and [5, 7, 12] outline the general maintenance problem of materialized views. Unfortunately, none of these works have studied the case in which each independent view has been defined over multiple sources of data. In this case, the problem is more complex since each single materialized view can involve several operational databases. In [13] we proposed a qualitative solution to the problem; however the solution left the quantification of the different cases for future study. In this article we propose a quantitative approach to the problem and we analyze the most efficient solution for each case.

The article is organized as follows: In section 2, the statement of the problem is presented along with the current state of the art. In section 3, the different parameters for quantifying the problem are studied. Also, different ways of measuring the time and the space required in the maintenance of the materialized views are analyzed. In section 4, the different cases are analyzed and the most efficient solution for each one is presented. Section 5 concludes.

2 State of the art

The updating of materialized views in real time is usually performed by first establishing a communication channel between the Warehouse and the underlying operational systems. This is done so that every time a modification takes place in the tables of the operational systems, these systems inform the Warehouse of the changes by sending the updates that are necessary to maintain the consistency of the materialized views.

When a view is defined on tables from databases of different operational systems, neither the Warehouse nor any of the operational systems can make a join among the tables to solve the view. This is because all the necessary information must be (at least temporarily) in the same location. At first glance, this may appear to be a problem of small granularity; however, this is not always the case. Many times business enterprises have branch offices located in different parts of the world that share the same Data Warehouse. In this context, the volume of data needed for maintaining the views might contain millions of tuples.

The solutions that have been proposed [7, 9, 12] and used to solve the problem of the multi-source views are based on the duplication of information [5, 6]; the definition of maintenance transactions [10]; or the redefinition of the original views,

in an attempt to avoid (if it is possible) the joining of tables from different databases [2]; or establishing a hierarchy of auxiliary multi-level views [13].

As it was stated in [13], in many cases the most appropriate solution consists of extending the materialized views definition with a hierarchy of views whose top level are the original views. Although this proposal solves the problem, it has only been defined at a qualitative level, and has left the quantification for future work. An outline is presented in Figure 1, where the solution to this problem is presented.

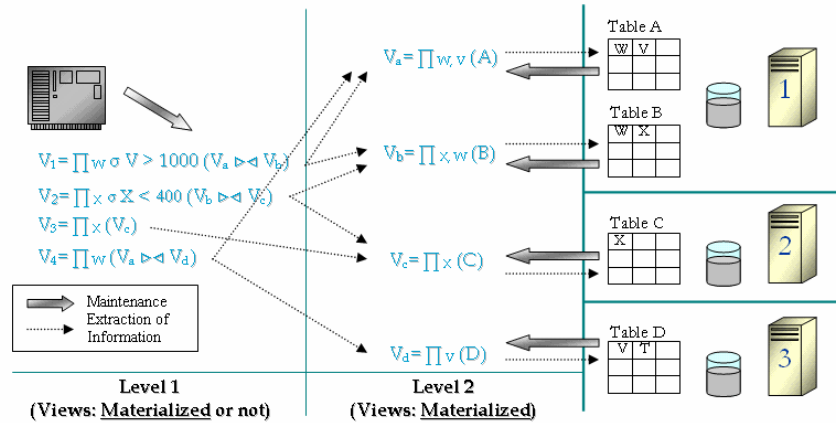


Fig. 1. Data warehouse maintenance using auxiliary multi-level views.

Figure 1 shows three operational systems with tables (A, B, C, D) and a Data Warehouse with four materialized views (V_1 , V_2 , V_3 , V_4) defined as follows:

- $V_1 = \Pi_w \sigma V > 1000 (A \bowtie B)$
- $V_2 = \Pi_x \sigma X < 400 (B \bowtie C)$
- $V_3 = \Pi_x (C)$
- $V_4 = \Pi_w \sigma V > 1000 (A \bowtie D)$

Views V_1 and V_2 introduce a problem because they include joins on tables from different databases that are in different locations. To solve this problem, we have defined four new views (V_a , V_b , V_c , V_d) to extract the necessary information from tables A, B, C and D. The original views have been redefined from these new views (see definition in Figure 1) in such a way that none of the views include joins on tables from different databases. So the maintenance of the original views is now possible because they are defined over local views (single source); and also the maintenance of auxiliary views is possible because they are defined over a single (remote) source.

This leads to the question of how many level 2 views should be defined for each table of the operational systems in order to duplicate and maintain the smallest

quantity of information possible, and at the same time maximize the response time of the system. This question does not have an easy answer, since the answer not only depends on the definition of the views, but on other parameters such as the transmission speed and the volume of data stored in the table. The main goal of this work is to find the answer to this question based on the study of all the possible cases where this problem occurs.

3 Problem statements: study of influential parameters

If we analyze the outline of views in Figure 1, it seems quite clear that there should be at least one view for each database in the operational systems. In the simplest case, this view belongs to the last level in the view hierarchy, and it contains information from one or more tables of this database. It is also clear that it is not always possible to represent all the necessary information from one database in a single view. In general, by defining a subview for each table included in the views of the first level, we have all the necessary information for maintenance. Each subview stores the minimum necessary information of each table, and by joining them, all the Data Warehouse views can be computed. Sometimes it is more appropriate to define more than one view for each table in order to avoid storing and maintaining unnecessary information. The following example shows that by using two views, V_1 and V_2 , less information is stored than when using the single view V_3 .




Fig. 2. Removing spare data using several views.

The information stored (and maintained) in views V_1 and V_2 does not include the tuples of the blank space (condition D/field P) which is stored by V_3 . A table can be represented as a space of two dimensions where dimension Y represents the tuples of the table and dimension X represents the fields of this table. Each view defined in this table represents a subspace that includes both the fields in its projection and the tuples defined in its selection. In this sense, two different subspaces defined by two different views can be grouped in a single new view. However, information that does not belong to any of them could be included in this new view.

Figure 2 shows two opposed problems that can appear during maintenance. When using a single view V_3 , unnecessary information can be stored (spare information). When using views V_1 and V_2 the same information is stored twice (redundant information).

Data Warehouse designers must determine the number of views for each database and each table of the operational systems. This decision is crucial to ensure that the volume of data stored in the auxiliary materialized views and the maintenance time of the Data Warehouse are minimized. To address this problem, we study quantitatively all the different cases that a designer can find. For this purpose, the designer should take into account the following variables:

Name	Description
M	Number of sent messages.
C	Global temporal maintenance cost (taking into account input-output operations, bits transference, etc.).
Card V	Cardinality of view V.
Grad V	Number of attributes (fields) in view V.
M_T	Number of updates (insertions, modifications, deletions) on table T.
m_T	Percentage of M_T which corresponds with modifications on table T.
ρ	Number of updates that are relevant for several views.
k	Average of the maintenance cost (messages + index maintenance + etc) of an operation.
s-1	Number of updates cancelled before computing the view.

Tab. 1. Metrics and variables.

In [15], authors propose to compute the number of messages between the operational systems and the Data Warehouse as a metric of the maintenance cost of a materialized view. This metric corresponds with M in Table 1. For its computation, they propose the following formula whose proof can be found in annex D in [15]:

$$(1) \quad M = \left[\frac{M_T}{s} \right] \times 2$$

In this metric, M_T represents the total number of updates in the source; and for every 's' updates the operational system sends (and receives) one message to the Data Warehouse. Then, assuming that the message is sent after each update ($s=1$), two messages are generated by each update in the source.

It is clear that measuring the maintenance cost by only considering the amount of sent messages can be too rough. For this reason, we propose a different metric that takes into account other factors such as the view size, the probability that the modifications in the sources will affect the materialized views, etc. The formula that defines this metric is:

$$(2) \quad C = M_T \cdot \frac{Card(V_1) \cdot Grad(V_1) + \dots + Card(V_n) \cdot Grad(V_n)}{Card(V_T) \cdot Grad(V_T)} \cdot k + \rho \cdot k$$

where T is the table in the operational system on which the views $V_{1..n}$ are defined. The meaning of the rest of variables is explained in Table 1.

This metric takes into account the spare information as well as the redundant information. In the case that the intersection of two or more views is not the empty set, that is, there is redundant information, this redundancy is taken into account in the cost since it is computed several times in different pairs 'Card(V_i)·Grad(V_i)'. The spare information is computed in the corresponding pair in the same way.

The redundant maintenance is computed with the second expression of the formula; $\rho \cdot k$ represents the cost associated to all the operations caused by the change of a tuple from one view to another during maintenance. It is of special importance to realize that, in the case of disjoint views, this change of view cannot be caused by an *insert* or *delete* operation, but rather can only be caused by a *modify* operation. Therefore, when views do not have redundant tuples we use the parameter ρ_m which only refers to modifications. Assuming that every field in the views has the same probability to be modified, ρ_m can be computed by the formula:

$$(3) \quad \rho_m = M_T \cdot m_T \cdot \sum_{i=1}^n \left(\frac{\text{Card}(V_i) \cdot \text{Grad}(V_i) \cdot \text{Card}(\bigcup_{j \neq i}^{j=1..n} V_j)}{\text{Card}(T) \cdot (\text{Grad}(T))^2} \right)$$

Parameter ρ represents the number of times that one update on a materialized view produces another update on any other view. In order to validate the correctness of figure (3), for simplicity, we can assume that every field in table T has the same probability to be updated by one of the M_T updates (in practice, we use a statistical distribution or a weigh). Then, the probability that the view 'i' were affected by an update is:

$$(4) \quad P_{V_i} = \frac{\text{Card}(V_i) \cdot \text{Grad}(V_i)}{\text{Card}(T) \cdot \text{Grad}(T)}$$

The probability that the update affects to the set of views $V_1 \dots V_n$:

$$(5) \quad P_{V_1..V_n} = \frac{\text{Card}(\bigcup_n^{i=1} V_i) \cdot \text{Grad}(\bigcup_n^{i=1} V_i)}{\text{Card}(T) \cdot \text{Grad}(T)}$$

Using (4), the number of updates which affect to view 'i' is:

$$(6) \quad N_{V_i} = M_T \cdot P_{V_i} = M_T \cdot \frac{\text{Card}(V_i) \cdot \text{Grad}(V_i)}{\text{Card}(T) \cdot \text{Grad}(T)}$$

The number of modifications which affect the view 'i' is:

$$(7) \quad N'_{V_i} = M_T \cdot m_T \cdot P_{V_i} \cdot \frac{\text{Card}(\bigcup_{j \neq i}^{j=1..n} V_j)}{\text{Card}(T)}$$

The generalization of (7):

$$(8) \quad \rho_m = M_T \cdot m_T \cdot \sum_{i=1}^n \left(\frac{\text{Card}(V_i) \cdot \text{Grad}(V_i) \cdot \text{Card}(\bigcup_{j \neq i}^{j=1..n} V_j)}{\text{Card}(T) \cdot (\text{Grad}(T))^2} \right)$$

For the case of two views:

$$(9) \quad \rho_{1,2} = M_T \cdot \frac{\text{Card}(\bigcap V_1, V_2)}{\text{Card}(T)}$$




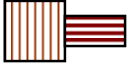

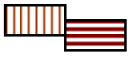

4 Case-based reasoning

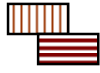

There is not only one solution to the problem described in the previous section. The solutions depend on different factors that are related to the definition of the views (for example the number of fields and tuples of a table used in several views); and they also depend on architectural factors such as the average speed of the information transmission, the average number of transmissions which are necessary to perform maintenance, the volume of information to maintain, etc. In this section, we study all the possible combinations that can take place in the definition of diverse views on a table. For each case we analyze which solution is the most appropriate according to the rest of the parameters.

Conditions		Field W	Field X	Field Y	Field Z
Condition A	Condition C	Condition F			
	Condition D	Condition E			
Condition B	Condition F	Condition D			
	Condition E				

Fig. 3. Table T with 4 fields and 6 conditions defined over the tuples space.

Figure 3 shows a hypothetical table T with 4 fields (W, X, Y, Z) and 6 conditions dividing the tuples space in 4 disjoint sets. Two views defined over this table necessarily hold in one of the following cases:

CASE 1: (Same projection, No disjoint conditions)	
$V_1 = \Pi W, X \sigma A (T)$ $V_2 = \Pi W, X (T)$	
CASE 2: (Same projection, Disjoint conditions)	
$V_1 = \Pi W, X \sigma A (T)$ $V_2 = \Pi W, X \sigma B (T)$	
CASE 3: (Same projection, Partially disjoint conditions)	
$V_1 = \Pi W, X \sigma A (T)$ $V_2 = \Pi W, X \sigma D (T)$	
CASE 4: (Distinct projection, No disjoint conditions)	
$V_1 = \Pi W, X \sigma B (T)$ $V_2 = \Pi Y, Z \sigma F \wedge D (T)$	
CASE 5: (Distinct projection, Disjoint conditions)	
$V_1 = \Pi W, X \sigma E (T)$ $V_2 = \Pi Y, Z \sigma C (T)$	
CASE 6: (Distinct projection, Partially disjoint conditions)	
$V_1 = \Pi W, X \sigma E (T)$ $V_2 = \Pi Y, Z \sigma B (T)$	
CASE 7: (Partially disjoint projections, Same conditions)	
$V_1 = \Pi W, X \sigma B (T)$ $V_2 = \Pi X, Y \sigma B (T)$	

CASE 8: (Partially disjoint projections, Disjoint conditions)	
$V_1 = \Pi W, X \sigma F(T)$ $V_2 = \Pi X, Y \sigma E(T)$	
CASE 9: (Partially disjoint projections, Partially disjoint conditions)	
$V_1 = \Pi W, X, Y \sigma F(T)$ $V_2 = \Pi X, Y, Z \sigma B(T)$	

For each case, we have computed the cost taking into account all the possibilities (using different number of views). The whole analysis can be found in Appendix 1.

From this analysis, we conclude that in cases 1, 2, 3 and 7 the best solution is to consider one single view. In cases 4, 5, 6 and 8 the best solution depends on the parameter ρ and on the cardinality of the views:

- If $\rho > M_T(\text{Card } V_3 \cdot \text{Grad } V_3 - (\text{Card } V_1 \cdot \text{Grad } V_1 + \text{Card } V_2 \cdot \text{Grad } V_2)) / \text{Card } T \cdot \text{Grad } T$ then the cost of using one view is smaller than the cost of using two views.
- If $\rho = M_T(\text{Card } V_3 \cdot \text{Grad } V_3 - (\text{Card } V_1 \cdot \text{Grad } V_1 + \text{Card } V_2 \cdot \text{Grad } V_2)) / \text{Card } T \cdot \text{Grad } T$ then the cost of using one view or two views is the same. In this case only one view is used.
- If $\rho < M_T(\text{Card } V_3 \cdot \text{Grad } V_3 - (\text{Card } V_1 \cdot \text{Grad } V_1 + \text{Card } V_2 \cdot \text{Grad } V_2)) / \text{Card } T \cdot \text{Grad } T$ then the cost of using two views is smaller than the cost of using one view.

The case number 9 is more complex since it needs 3 views for maintaining all the information without redundancy. The computation of the cost for this case is also shown in Appendix 1.

5. Conclusions and future work

There are many works [3, 5, 6, 8, 10, 14, 15] which deal with the use of materialized views in Data Warehouse systems or with the problem of maintenance of these materialized views. However, surprisingly, not much effort [7, 9, 12] has been spent to address the problem of multi-source materialized views. This is an important problem since, for example, it is not possible—or it is really very expensive due to the necessity of duplicate the tables—to maintain in real-time a view that joins tables from different sources. In [13], we proposed to extend the original views with a set of intermediate views that extract information from just one determined source. But in that work we did not consider how to find out which is the optimal configuration of views taking into account the maintenance costs.

In this paper we have carried out this study analyzing each possible case from the point of view of required space and time consumed for maintenance. To achieve this

goal ECV (Eager Compensating Algorithm) [15] have been used in order to maintain the auxiliary single-source (top level) views.

6 Bibliography

1. Ashish N., Knoblock C.A., Shahabi C.: “*Selectively materializing data in mediators by analyzing user queries*”. In Proceedings of the 4th IFCIS International Conference on Cooperative Information Systems (CoopIS), Edinburgh, Scotland, September 1999.
2. Colby L.S., Griffin T.G., Libkin L., Mumick I.S., Trickey H.: “*Algorithms for Deferred View Maintenance*”. In Proceedings of the SIGMOD: pp. 469-479, June 1996.
3. Ding L., Zhang X., Rundensteiner E.A.: “*Enhancing Existing Incremental View Maintenance Algorithms Using the Multi-Relation Encapsulation Wrapper*”. Technical Report WPI-CS-TR-99-23: pp. 9-11, Worcester Polytechnic Institute, Dept. of Computer Science, August 1999.
4. Gupta H.: “*Selection of Views to Materialize in a Data Warehouse*”. In Proceedings of the International Conference on Database Theory (ICDT 1997): 98-112, Delphi, Greece; January 1997.
5. Gupta A., Mumick I.S.: “*Maintenance of Materialized Views: Problems, Techniques, and Applications*”. In the IEEE Data Engineering Bulletin Vol. 18 N^o. 2, June 1995.
6. Kuchenhoff V.: “*On the efficient computation of the difference between consecutive database states*”. In Proceedings of the 2nd International Conference on Deductive and Object-Oriented Databases (DOOD), December 1991.
7. Moro G., Sartorio C.: “*Incremental Maintenance of Multi-Source Views*”. In Proceedings of the 12th Australasian of Electrical and Electronics Engineers (ADC 2001), 2001.
8. Quass D., Gupta A., Mumick I.S., Widom J.: “*Making Views Self-Maintainable for Data Warehousing*”. In Proceedings of the Conference on Parallel and Distributed Information Systems (PDIS '96), 1996.
9. Samtani S., Mohania M., Kumar V., Kambayashi Y.: “*Recent Advances and Research Problems in Data Warehousing*”. ER Workshops, 1998.
10. Stanoi I.R., Agrawal D., El Abbadi A.: “*Modeling and Maintaining Multi-View Data Warehouses*”. In Proceedings of the 18th International Conference on Conceptual Modelling (ER '99), November 1999.
11. Valluri S.R., Vadapalli S., Karlapalem K.: “*View Relevance Driven Selection of Materialized Views in Data Warehousing Environment*”. In Proceedings of the 13th Australasian Database Conference (ADC 2002), Melbourne, Australia 2002.
12. Widom J.: “*Research Problems in Data Warehousing*”. Computer Science Department, Stanford University. In Proceedings of the 4th International Conference on Information and Knowledge Management (CIKM); November 1995.
13. Silva J., Belenguer J., Celma M.: “*Materialización de Vistas Multi-Origen: Vistas Multinivel*”. In Proceedings of the 7th Jornadas de Ingeniería del Software y Bases de Datos (JISBD'02), El Escorial, Madrid, Spain; 2002.
14. Zhuge Y., Garcia-Molina H., Wiener J.L.: “*The Strobe Algorithms for Multi-Source Warehouse Consistency*”. In Proceedings of the 4th International Conferences on Parallel and Distributed Information Systems (PDIS'96), Miami Beach, Florida, USA; 1996.
15. Zhuge Y., Garcia-Molina H., Hammer J., Widom J.: “*View Maintenance in a Warehousing Environment*”. CSD - Stanford University; SIGMOD Conference 1995: pp. 316-327; 1995.

Appendix 1: Cost computation

CASE 1: (Same projection, No disjoint conditions)

$$V_1 = \Pi_{W,X} \sigma_A(T)$$

$$V_2 = \Pi_{W,X}(T)$$



Temporal Cost:

- **1 view (V_2)**

- Messages: $2 M_{V_2}$

- Global maintenance: $C = M_T(\text{Card } V_2 \cdot \text{Grad } V_2 / \text{Card } T \cdot \text{Grad } T) k$

- **2 views ($V_1, V_3 = \Pi_{W,X} \sigma_{\neg A}(T)$)**

- Messages: $2 M_{V_1} + 2 M_{V_3} = 2 M_{V_2} (M_{V_1} + M_{V_3} = M_{V_2})$

- Global maintenance: $C' = M_T((\text{Card } V_1 \cdot \text{Grad } V_1 + \text{Card } V_3 \cdot \text{Grad } V_3) / \text{Card } T \cdot \text{Grad } T) k + (\rho) k$

assuming $\text{Card } V_1 \cdot \text{Grad } V_1 + \text{Card } V_3 \cdot \text{Grad } V_3 = \text{Card } V_2 \cdot \text{Grad } V_2$

then $C' = M_T((\text{Card } V_2 \cdot \text{Grad } V_2) / \text{Card } T \cdot \text{Grad } T) k + (\rho) k = C + (\rho) k \cdot C$

- **2 views (V_1, V_2)**

- Messages: $2 M_{V_1} + 2 M_{V_2} > 2 M_{V_2} (M_{V_1} + M_{V_2} > M_{V_2})$

- Global maintenance: $C'' = M_T((\text{Card } V_1 \cdot \text{Grad } V_1 + \text{Card } V_2 \cdot \text{Grad } V_2) / \text{Card } T \cdot \text{Grad } T) k + (\rho') k = M_T((\text{Card } V_1 \cdot \text{Grad } V_1) / \text{Card } T \cdot \text{Grad } T) k + M_T((\text{Card } V_2 \cdot \text{Grad } V_2) / \text{Card } T \cdot \text{Grad } T) k + (\rho') k$

assuming $C = M_T(\text{Card } V_2 \cdot \text{Grad } V_2 / \text{Card } T \cdot \text{Grad } T) k$

then $C'' = M_T((\text{Card } V_1 \cdot \text{Grad } V_1) / \text{Card } T \cdot \text{Grad } T) k + C + (\rho') k > C$

If $\rho < \rho' + M_T((\text{Card } V_1 \cdot \text{Grad } V_1) / \text{Card } T \cdot \text{Grad } T)$ then $C'' > C'$.

If $\rho = \rho' + M_T((\text{Card } V_1 \cdot \text{Grad } V_1) / \text{Card } T \cdot \text{Grad } T)$ then $C'' = C'$.

If $\rho > \rho' + M_T((\text{Card } V_1 \cdot \text{Grad } V_1) / \text{Card } T \cdot \text{Grad } T)$ then $C'' < C'$.

The best solution is to use only one view.

CASE 2: (Same projection, Disjoint conditions)

$$V_1 = \Pi_{W,X} \sigma_A(T)$$

$$V_2 = \Pi_{W,X} \sigma_B(T)$$



Temporal Cost:

- **1 view ($V_3 = \Pi W, X \sigma A \vee B (T)$)**
 - Messages: $2 M_{V_3}$
 - Global Maintenance: $C = M_T (Card V_3 \cdot Grad V_3 / Card T \cdot Grad T) k$
- **2 views (V_1, V_2)**
 - Messages: $2 M_{V_1} + 2 M_{V_2} = 2 M_{V_3} (M_{V_1} + M_{V_2} = M_{V_3})$
 - Global maintenance: $C' = M_T ((Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2) / Card T \cdot Grad T) k + (\rho) k$

assuming $Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2 = Card V_3 \cdot Grad V_3$
then $C' = M_T (Card V_3 \cdot Grad V_3 / Card T \cdot Grad T) k + (\rho) k$
and assuming $C = M_T (Card V_2 \cdot Grad V_2 / Card T \cdot Grad T) k$
then $C' = C + (\rho) k \cdot C$

The best solution is to use only one view.

CASE 3: (Same projection, Partially disjoint conditions)

$$V_1 = \Pi W, X \sigma A (T)$$

$$V_2 = \Pi W, X \sigma D (T)$$



$$V'_1 = \Pi W, X \sigma A \wedge (\neg D) (T)$$

$$V'_2 = \Pi W, X \sigma A \wedge D (T)$$

$$V'_3 = \Pi W, X \sigma D \wedge (\neg A) (T)$$

Temporal Cost:

- **1 view ($V_3 = \Pi W, X \sigma A \vee D (T)$)**
 - Messages: $2 M_{V_3}$
 - Global maintenance: $C = M_T (Card V_3 \cdot Grad V_3 / Card T \cdot Grad T) k$
- **2 views ($V_1, V_4 = \Pi W, X \sigma A \wedge (\neg D) (T)$)**
 - Messages: $2 M_{V_1} + 2 M_{V_4} = 2 M_{V_3} (M_{V_1} + M_{V_4} = M_{V_3})$
 - Global maintenance: $C' = M_T ((Card V_1 \cdot Grad V_1 + Card V_4 \cdot Grad V_4) / Card T \cdot Grad T) k + (\rho) k$

assuming $Card V_1 \cdot Grad V_1 + Card V_4 \cdot Grad V_4 = Card V_3 \cdot Grad V_3$
then $C' = M_T (Card V_3 \cdot Grad V_3 / Card T \cdot Grad T) k + (\rho) k$
and assuming $C = M_T (Card V_3 \cdot Grad V_3 / Card T \cdot Grad T) k$
then $C' = C + (\rho) k \cdot C$

• **2 views (V_1, V_2)**

- Messages: $2 M_{V_1} + 2 M_{V_2} > 2 M_{V_3}$ ($M_{V_1} + M_{V_2} > M_{V_3}$)
- Global maintenance: $C'' = M_T ((Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2) / Card T \cdot Grad T) k + (\rho) k$


assuming $Card V_2 \cdot Grad V_2 > Card V_4 \cdot Grad V_4$
 then $C'' > C' \cdot C$

The best solution is to use only one view.

CASE 4: (Distinct projection, No disjoint conditions)

$$V_1 = \Pi W, X \sigma B(T)$$

$$V_2 = \Pi Y, Z \sigma F \wedge D(T)$$



Temporal cost:

• **1 view ($V_3 = \Pi W, X, Y, Z \sigma B(T)$)**

- Messages: $2 M_{V_3}$
- Global maintenance: $C = M_T (Card V_3 \cdot Grad V_3 / Card T \cdot Grad T) k$

• **2 views (V_1, V_2)**

- Messages: $2 M_{V_1} + 2 M_{V_2} < 2 M_{V_3}$ ($M_{V_1} + M_{V_2} < M_{V_3}$)
- Global maintenance: $C' = M_T ((Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2) / Card T \cdot Grad T) k + (\rho) k$

assuming
 $Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2 < Card V_3 \cdot Grad V_3$.


The best solution depends on the parameter ρ and on the cardinality of the views.

- If $\rho > M_T (Card V_3 \cdot Grad V_3 - (Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2)) / Card T \cdot Grad T$
 then the cost for one view is smaller than for two views.
- If $\rho = M_T (Card V_3 \cdot Grad V_3 - (Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2)) / Card T \cdot Grad T$
 then the cost for one view and for two views is the same.
- If $\rho < M_T (Card V_3 \cdot Grad V_3 - (Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2)) / Card T \cdot Grad T$
 then the cost for two views is smaller than for one view.

CASE 5: (Distinct projection, Disjoint conditions)

$$V_1 = \Pi W, X \sigma E(T)$$

$$V_2 = \Pi Y, Z \sigma C(T)$$



Temporal case:

- **1 view ($V_3 = \Pi W, X, Y, Z \sigma_{C \vee E}(T)$)**
 - Messages: $2 M_{V_3}$
 - Global Maintenance: $C = M_T (Card V_3 \cdot Grad V_3 / Card T \cdot Grad T) k$
- **2 views (V_1, V_2)**
 - Messages: $2 M_{V_1} + 2 M_{V_2} < 2 M_{V_3} (M_{V_1} + M_{V_2} < M_{V_3})$
 - Global maintenance: $C' = M_T ((Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2) / Card T \cdot Grad T) k + (\rho) k$


assuming $Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2 < Card V_3 \cdot Grad V_3$.

The best solution depends on the parameter ρ and on the cardinality of the views.

- If $\rho > M_T (Card V_3 \cdot Grad V_3 - (Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2)) / Card T \cdot Grad T$ then the cost for one view is smaller than for two views.
- If $\rho = M_T (Card V_3 \cdot Grad V_3 - (Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2)) / Card T \cdot Grad T$ then the cost for one view and for two views is the same.
- If $\rho < M_T (Card V_3 \cdot Grad V_3 - (Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2)) / Card T \cdot Grad T$ then the cost for two views is smaller than for one view

CASE 6: (Distinct projection, Partially disjoint conditions)

$$V_1 = \Pi W, X \sigma E(T)$$



$$V_2 = \Pi Y, Z \sigma B(T)$$

Temporal cost:

- **1 view ($V_3 = \Pi W, X, Y, Z \sigma_{B \vee E}(T)$)**
 - Messages: $2 M_{V_3}$
 - Global maintenance: $C = M_T (Card V_3 \cdot Grad V_3 / Card T \cdot Grad T) k$
- **2 views (V_1, V_2)**
 - Messages: $2 M_{V_1} + 2 M_{V_2} < 2 M_{V_3} (M_{V_1} + M_{V_2} < M_{V_3})$
 - Global maintenance: $C' = M_T ((Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2) / Card T \cdot Grad T) k + (\rho) k$

assuming $Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2 < Card V_3 \cdot Grad V_3$.

The best solution depends on the parameter ρ and on the cardinality of the views.

- If $\rho > M_T (Card V_3 \cdot Grad V_3 - (Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2)) / Card T \cdot Grad T$ then the cost for one view is smaller than for two views.

- If $\rho = M_T(\text{Card } V_3 \cdot \text{Grad } V_3 - (\text{Card } V_1 \cdot \text{Grad } V_1 + \text{Card } V_2 \cdot \text{Grad } V_2)) / \text{Card } T \cdot \text{Grad } T$ then the cost for one view and for two views is the same.
- If $\rho < M_T(\text{Card } V_3 \cdot \text{Grad } V_3 - (\text{Card } V_1 \cdot \text{Grad } V_1 + \text{Card } V_2 \cdot \text{Grad } V_2)) / \text{Card } T \cdot \text{Grad } T$ then the cost for two views is smaller than for one view

CASE 7: (Partially disjoint projections, Same conditions)

$$V_1 = \Pi W, X \sigma B(T)$$

$$V_2 = \Pi X, Y \sigma B(T)$$

**Temporal Cost:**

- **1 view ($V_3 = \Pi W, X, Y \sigma B(T)$)**
 - Messages: $2 M_{V_3}$
 - Global maintenance: $C = M_T(\text{Card } V_3 \cdot \text{Grad } V_3 / \text{Card } T \cdot \text{Grad } T) k$
- **2 views ($V_1, V_4 = \Pi Y \sigma B(T)$)**
 - Messages: $2 M_{V_1} + 2 M_{V_4} = 2 M_{V_3} (M_{V_1} + M_{V_4} = M_{V_3})$
 - Global maintenance: $C' = M_T((\text{Card } V_1 \cdot \text{Grad } V_1 + \text{Card } V_4 \cdot \text{Grad } V_4) / \text{Card } T \cdot \text{Grad } T) k + (\rho) k$
 - assuming $\text{Card } V_1 \cdot \text{Grad } V_1 + \text{Card } V_4 \cdot \text{Grad } V_4 = \text{Card } V_3 \cdot \text{Grad } V_3$
 - then $C' = M_T(\text{Card } V_3 \cdot \text{Grad } V_3 / \text{Card } T \cdot \text{Grad } T) k + (\rho) k$
 - assuming $C = M_T(\text{Card } V_3 \cdot \text{Grad } V_3 / \text{Card } T \cdot \text{Grad } T) k$
 - then $C' = C + (\rho) k \cdot C$
- **2 views (V_1, V_2)**
 - Messages: $2 M_{V_1} + 2 M_{V_2} > 2 M_{V_3} (M_{V_1} + M_{V_2} > M_{V_3})$
 - Global maintenance: $C'' = M_T((\text{Card } V_1 \cdot \text{Grad } V_1 + \text{Card } V_2 \cdot \text{Grad } V_2) / \text{Card } T \cdot \text{Grad } T) k + (\rho) k$
 - assuming $\text{Card } V_2 \cdot \text{Grad } V_2 > \text{Card } V_4 \cdot \text{Grad } V_4$
 - then $C'' > C' \cdot C$

The best solution is to use only one view.

CASE 8: (Partially disjoint projections, Disjoint conditions)

$$V_1 = \Pi W, X \sigma F(T)$$

$$V_2 = \Pi X, Y \sigma E(T)$$

Temporal cost:

- **1 view ($V_3 = \Pi W, X, Y \sigma E \vee F(T)$)**

- Messages: $2 M_{V_3}$
- Global maintenance: $C = M_T (Card V_3 \cdot Grad V_3 / Card T \cdot Grad T) k$

• **2 views (V_1, V_2)**

- Messages: $2 M_{V_1} + 2 M_{V_2} < 2 M_{V_3} (M_{V_1} + M_{V_2} < M_{V_3})$
- Global maintenance: $C' = M_T ((Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2) / Card T \cdot Grad T) k + (\rho) k$

assuming

$$Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2 < Card V_3 \cdot Grad V_3$$

The best solution depends on the parameter ρ and on the cardinality of the views.

- If $\rho > M_T (Card V_3 \cdot Grad V_3 - (Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2)) / Card T \cdot Grad T$ then the cost for one view is smaller than for two views.
- If $\rho = M_T (Card V_3 \cdot Grad V_3 - (Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2)) / Card T \cdot Grad T$ then the cost for one view and for two views is the same.
- If $\rho < M_T (Card V_3 \cdot Grad V_3 - (Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2)) / Card T \cdot Grad T$ then the cost for two views is smaller than for one view

CASE 9: (Partially disjoint projections, Partially disjoint conditions)

$$V_1 = \prod W, X, Y \sigma F (T)$$

$$V_2 = \prod X, Y, Z \sigma B (T)$$



Temporal Cost:

• **1 view ($V_3 = \prod W, X, Y, Z \sigma B \vee F (T)$)**

- Messages: $2 M_{V_3}$
- Global maintenance: $C = M_T (Card V_3 \cdot Grad V_3 / Card T \cdot Grad T) k$

• **2 views (V_1, V_2) (*vistas originales*)**

- Messages: $2 M_{V_1} + 2 M_{V_2} < 2 M_{V_3} (M_{V_1} + M_{V_2} < M_{V_3})$
- Global maintenance: $C' = M_T ((Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2) / Card T \cdot Grad T) k + (\rho) k$

The best solution depends on the parameter ρ and on the cardinality of the views.

- If $\rho > M_T (Card V_3 \cdot Grad V_3 - (Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2)) / Card T \cdot Grad T$ then the cost for one view is smaller than for two views.
- If $\rho = M_T (Card V_3 \cdot Grad V_3 - (Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2)) / Card T \cdot Grad T$ then the cost for one view and for two views is the same.
- If $\rho < M_T (Card V_3 \cdot Grad V_3 - (Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2)) / Card T \cdot Grad T$ then the cost for two views is smaller than for one view

• **3 views ($V'_1 = \prod W \sigma F (T), V'_2 = \prod Y, Z \sigma B \vee F (T), V'_3 = \prod X \sigma B (T)$)**

- Messages: $2 M_{V'_1} + 2 M_{V'_2} + 2 M_{V'_3} = 2 M_{V_1} + 2 M_{V_2} < 2 M_{V_1} + 2 M_{V_2} < 2 M_{V_3} (M_{V_1} + M_{V_2} + M_{V_3} = M_{V_1} + M_{V_2} < M_{V_1} + M_{V_2} < M_{V_3})$

- Global maintenance: $C'' = M_T((Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2 + Card V'_3 \cdot Grad V'_3) / Card T \cdot Grad T) k + (\rho') k$

then

$$Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2 + Card V'_3 \cdot Grad V'_3 < Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2$$

and also

$$Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2 + Card V'_3 \cdot Grad V'_3 < Card V'_3 \cdot Grad V'_3$$

The best solution depends on the parameters ρ , ρ' and on the cardinality of the views.

- If $\rho > \rho'$ then the cost for three views is smaller than for two views.
 - If $\rho' > M_T(((Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2 + Card V'_3 \cdot Grad V'_3) - Card V'_3 \cdot Grad V'_3) / Card T \cdot Grad T)$ then the cost for one view is smaller than for three views.
 - If $\rho' = M_T(((Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2 + Card V'_3 \cdot Grad V'_3) - Card V'_3 \cdot Grad V'_3) / Card T \cdot Grad T)$ then the cost for one view and for three views is the same.
 - If $\rho' < M_T(((Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2 + Card V'_3 \cdot Grad V'_3) - Card V'_3 \cdot Grad V'_3) / Card T \cdot Grad T)$ then the cost for three views is smaller than for one view.
- If $\rho < \rho'$ then...
 - If $\rho' - \rho > M_T((Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2) - (Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2 + Card V'_3 \cdot Grad V'_3) / Card T \cdot Grad T)$ then the cost for three views is smaller than for two views.
 - If $\rho' > M_T(((Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2 + Card V'_3 \cdot Grad V'_3) - Card V'_3 \cdot Grad V'_3) / Card T \cdot Grad T)$ then the cost for one view is smaller than for three views.
 - If $\rho' = M_T(((Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2 + Card V'_3 \cdot Grad V'_3) - Card V'_3 \cdot Grad V'_3) / Card T \cdot Grad T)$ then the cost for one view and for three views is the same.
 - If $\rho' < M_T(((Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2 + Card V'_3 \cdot Grad V'_3) - Card V'_3 \cdot Grad V'_3) / Card T \cdot Grad T)$ then the cost for three views is smaller than for one view.
 - If $\rho' - \rho = M_T((Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2) - (Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2 + Card V'_3 \cdot Grad V'_3) / Card T \cdot Grad T)$ then the cost for two views and for three views is the same.
 - If $\rho' > M_T(((Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2 + Card V'_3 \cdot Grad V'_3) - Card V'_3 \cdot Grad V'_3) / Card T \cdot Grad T)$ then the cost for one view is smaller than for three views; and consequently, the cost for one view is smaller than for two views.

- If $\rho' = M_T(((Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2 + Card V'_3 \cdot Grad V'_3) - Card V_3 \cdot Grad V_3) / Card T \cdot Grad T)$ then the cost for one view and for three views is the same.
 - If $\rho' < M_T(((Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2 + Card V'_3 \cdot Grad V'_3) - Card V_3 \cdot Grad V_3) / Card T \cdot Grad T)$ then the cost for three views is smaller than for one view; and consequently, the cost for two views is smaller than for one view.
- If $\rho' - \rho < M_T((Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2) - (Card V'_1 \cdot Grad V'_1 + Card V'_2 \cdot Grad V'_2 + Card V'_3 \cdot Grad V'_3) / Card T \cdot Grad T)$ then the cost for three views is greater than for two views.
 - If $\rho > M_T((Card V_3 \cdot Grad V_3 - (Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2)) / Card T \cdot Grad T)$ then the cost for one view is smaller than for two views.
 - If $\rho = M_T((Card V_3 \cdot Grad V_3 - (Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2)) / Card T \cdot Grad T)$ then the cost for one view and for two views is the same.
 - If $\rho < M_T((Card V_3 \cdot Grad V_3 - (Card V_1 \cdot Grad V_1 + Card V_2 \cdot Grad V_2)) / Card T \cdot Grad T)$ then the cost for two views is smaller than for one view.