

Multiversion-Based View Maintenance Over Distributed Data Sources

SONGTING CHEN, BIN LIU, and ELKE A. RUNDENSTEINER
Worcester Polytechnic Institute

Materialized views can be maintained by submitting maintenance queries to the data sources. However, the query results may be erroneous due to concurrent source updates. State-of-the-art maintenance strategies typically apply compensations to resolve such conflicts and assume all source schemata remain stable over time. In a loosely coupled dynamic environment, the sources may autonomously change not only their data but also their schema or semantics. Consequently, either the maintenance or the compensation queries may be broken. Unlike compensation-based approaches found in the literature, we instead model the complete materialized view maintenance process as a view maintenance transaction (VM_Transaction). This way, the anomaly problem can be rephrased as the serializability of VM_Transactions. To achieve VM_Transaction serializability, we propose a multiversion concurrency control algorithm, called *TxnWrap*, which is shown to be the appropriate design for loosely coupled environments with autonomous data sources. *TxnWrap* is complementary to the maintenance algorithms proposed in the literature, since it removes concurrency issues from consideration allowing the designer to focus on the maintenance logic. We show several optimizations of *TxnWrap*, in particular, (1) space optimizations on versioned data materialization and (2) parallel maintenance scheduling. With these optimizations, *TxnWrap* even outperforms state-of-the-art view maintenance solutions in terms of refresh time. Further, several design choices of *TxnWrap* are studied each having its respective advantages for certain environmental settings. A correctness proof based on transaction theory for *TxnWrap* is also provided. Last, we have implemented *TxnWrap*. The experimental results confirm that *TxnWrap* achieves predictable performance under a varying rate of concurrency.

Categories and Subject Descriptors: H.2.4 [Database Management]: Systems—*Distributed databases; relational databases; concurrency*

General Terms: Theory, Algorithms, Experimentation

Additional Key Words and Phrases: View maintenance, transaction processing

1. INTRODUCTION

1.1 Materialized Views and Their Environment

Materialized views [Gupta and Mumick 1995; Agrawal et al. 1997] built by gathering data from possibly distributed data sources and integrating it into

This work was supported in part by several grants from NSF, namely, the NSF NY1 grant IRI 97-96264, the NSF CISE Instrumentation grant IRIS 97-29878, and the NSF grant IIS 9988776. Author's addresses: Department of Computer Science, Worcester Polytechnic Institute, Worcester, MA 01609; email: {chenst,binliu,rudenst}@cs.wpi.edu.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2004 ACM 0362-5915/04/1200-0675 \$5.00

one repository customized to users' needs are a well recognized technology for data integration, e-business, and data warehousing. One important task of a view manager is to maintain the materialized view upon source changes, since frequent data updates are common for many applications, such as stock trading or telephone call recording.

In dynamic environments like the World Wide Web (WWW), the data sources may change their schema, semantics as well as their query capabilities. A schema change could occur for numerous reasons during the software lifecycle, including design errors, schema redesign during the early stages of database deployment, the addition of new functionalities, and even new developments in the application domain, such as new tax laws or Y2K problems. Even in fairly standard business applications, rapid schema changes have been observed. In Marche [1993], significant changes (about 59% of attributes on average) were reported for seven different applications ranging from project tracking, sales management to government administration. A similar report can also be found in Sjoberg [1993]. Second, the emerging schema mapping techniques [Miller et al. 2000; Madhavan et al. 2001] are critical for information integration over heterogeneous data sources, which aim at mapping between heterogeneous sources, are semiautomatic and depend on domain knowledge. As a result, a good mapping is thus hard to find and may evolve over time [Lee et al. 2002; Velegarakis et al. 2003].

Moreover, data sources in a distributed environment are typically owned by different information providers and function independently from one another. The relationship between materialized views and such autonomous data sources hence must be loosely coupled [Zhuge et al. 1995]. That is, the source updates are committed without any concern of how and when the view manager will incorporate them into the view [Widom 1995; Lee et al. 2002]. This causes problems which we called *maintenance anomalies*. These are the issues we now address in this work.

1.2 Motivating Example of the Maintenance Anomaly Problems

We will now illustrate the anomaly problems via motivating examples. We first distinguish between three main maintenance tasks, namely, *view maintenance*, *view synchronization*, and *view adaptation*. *View maintenance* [Zhuge et al. 1995; Agrawal et al. 1997; Salem et al. 2000] maintains the materialized view extent under source data updates. In contrast, *view synchronization* [Nica et al. 1998; Lee et al. 2002] aims at rewriting the view definition when the source schema has been changed. Thereafter, *view adaptation* [Gupta et al. 1995; Nica and Rundensteiner 1999] incrementally adapts the view extent again to match the newly changed view definition.

If there is no concurrency among source updates, namely, the current view maintenance process completes before the next source update occurs, then view maintenance incorporates each source *data update* while view synchronization and view adaptation together incorporate the source *schema change* into the view. However, the data sources are autonomous and may undergo changes concurrently, thus causing the maintenance anomaly problems illustrated below.

DS_1 : Customer(Name, Address, Phone): Customer Info.
DS_2 : Tour(TourID, TourName, Age, Type, NoDays): Tour. Info.
DS_3 : Participant(Participant, TourID, StartDate, Loc): Participant. Info.
DS_4 : FlightRes(Name, Age, FlightNo, Dest): Reservation. Info.

Fig. 1. Description of data sources.

Example 1. Assume we have four data sources (DS) as shown in Figure 1. The materialized view *Asia-Customer* is defined by the SQL query in Equation (1). Assume the data update “ $\Delta C = INSERT INTO Customer VALUES ('Ben', MA, 213)$ ” has happened at DS_1 . In order to determine the delta effect on the view extent, this now requires us to send the *view maintenance query* Q [Zhuge et al. 1995] defined in Equation (2) down to the *FlightRes* relation at DS_4 .

```

CREATE VIEW Asia-Customer AS
SELECT      C.Name, F.Age, F.FlightNo, F.Dest
FROM        Customer C, FlightRes F (1)
WHERE       C.Name = F.Name
            AND F.Dest = 'Asia'
SELECT      'Ben' as Name, F.Age, F.FlightNo, F.Dest
FROM        FlightRes F (2)
WHERE       F.Name = 'Ben'
            AND F.Dest = 'Asia'

```

We distinguish between two kinds of anomaly problems that may arise:

- Data update anomaly*: If during the transfer time of the query Q to the relation *FlightRes* in the DS_4 , *FlightRes* has already committed a new *data update* “ $\Delta F = INSERT INTO FlightRes VALUES ('Ben', 18, AA43, Asia)$.” This new tuple would also be included in the join result of Q and the extra tuple $(Ben', 18, AA43, Asia)$ would be inserted into the view. However, later when the materialized view manager starts processing the ΔF , the same tuple would be inserted into the view again. A data update anomaly appears. Similarly, a delete operation could cause anomaly problems.
- Schema change anomaly*: If during the transfer time of the query Q to DS_4 , the *FlightRes* relation in DS_4 has a *schema change*, for example, the attribute *FlightRes.Age* is dropped, then the query Q faces a schema conflict. In this case, the selected attribute *Age* required by the maintenance query Q is no longer available in the source. Hence the query Q cannot be processed by DS_4 . We then say that *the query Q is broken*. That is, a schema change anomaly appears.

As illustrated by the example above, sources being autonomous may commit updates that are concurrent with and hence may conflict with the view maintenance process. A concurrent data update may result in an *incorrect query result* while a concurrent schema change may result in a *broken query* that cannot be processed by the respective data source.

1.3 Contributions of this Work

A preliminary version of this article focusing on the TxnWrap model itself has been published in the 21st International Conference on Conceptual Modeling

[Chen et al. 2002]. Our contributions there included the following:

- (1) We illustrated that the anomaly problems as described in Section 1.2 can be reformulated as a transaction problem by modeling the complete maintenance process using a transaction model, called a *VM_Transaction*. Based on our *VM_Transaction* model, we demonstrated that the view maintenance anomaly problems can be mapped to the problem of the serializability of *VM_Transactions*.
- (2) We then proposed a multiversion concurrency control algorithm, called *TxnWrap*, to achieve such serializability that is appropriate for the view maintenance in a distributed environment with autonomous data sources.
- (3) To the best of our knowledge, *TxnWrap* is the first work to integrate the existing materialized view maintenance works in the literature such as view maintenance, view synchronization, and view adaptation algorithms within the context of a sound theory, the serializability theory, to remove concurrency concerns from their maintenance logic.
- (4) We have implemented the *TxnWrap* solution and integrated it into a fully functioning view maintenance prototype system. Experiments comparing maintenance performance with and without *TxnWrap* enabled confirm that *TxnWrap* avoids the maintenance anomaly problems and achieves stable performance.

In addition, in this extended journal article, we now contain the following:

- (1) We propose a new multiversion management strategy. Instead of using a global identifier as both a version number and its corresponding *VM_Transaction* identifier, we now propose to separate identifier design into two concepts. That is, we introduce a *local identifier* within each data source's scope for version management and a *global identifier* for *VM_Transaction* management where the global identifier is composed out of a vector of local identifiers. This avoids the coordination cost for generating global identifiers in distributed environments.
- (2) We analyze the design choices for *TxnWrap* version placement. Three different wrapper design choices are discussed and their tradeoffs, such as space costs and maintenance costs, are evaluated experimentally.
- (3) We illustrate that the development of a sound transactional foundation for view maintenance offers the advantage of the ability to easily add additional services to the system and assure their correctness. In particular, we show how *TxnWrap* with minimal effort can be optimized by applying a parallel instead of sequential maintenance scheduler. This way, we achieve a better maintenance performance. An initial version of parallel view maintenance work has been published in *DaWak'2002* [Liu et al. 2002a]. We now provide an extensive performance study of the parallel feature. Furthermore, we also present a performance study that demonstrates how the existing self-maintenance approach, such as [Quass et al. 1996] can benefit from our *TxnWrap* solution.

- (4) We demonstrate that our TxnWrap solution can be easily extended to efficiently handle the scenarios when there are multiple tables in one data source, a rather important practical requirement not supported by previous work such as SWEEP [Agrawal et al. 1997].
- (5) We introduce several space optimization strategies for versioned data materialization. Static optimization strategy is to analyze the view definition and the existing constraints to reduce the wrapper space during version creation. Dynamic optimization is a runtime strategy that aims to reduce the intermediate wrapper version space on the fly. We also conduct an experimental study to measure the impact that these space optimization strategies have on maintenance performance .
- (6) Finally, a correctness proof for both sequential and parallel TxnWrap is provided utilizing concepts from transactional theory [Bernstein et al. 1987]. These correctness proofs, primarily due to taking a transactional approach toward solving this problem, are missing from most existing solutions in the literature that typically offer (hard-coded) algorithmic approaches.

The outline of the rest of the article is as follows. The transaction model called *VM.Transaction* is introduced in Section 2. A multiversion concurrency control algorithm to handle the conflicts between VM.Transactions is proposed in Section 3. Section 4 discusses the architectural solutions for versioned data materialization. Wrapper optimization strategies are discussed in Section 5. Extensions of TxnWrap, in particular, parallel processing of VM.Transactions is introduced in Section 6. Section 7 presents experimental results. In Section 8 we review the related work while Section 9 concludes the article. A proof of TxnWrap maintenance is given in the electronic appendix.

2. TRANSACTIONAL MODELING OF VIEW MAINTENANCE

We first analyze the view maintenance anomaly problems. Typically, the maintenance process involves reads from individual data sources to calculate and ultimately refresh the view extent. On the other hand, the data source update process may continue to commit various changes throughout the same period. Due to the autonomy of the data sources, that is, the loosely-coupled nature of the environment, concurrency between these two processes may occur. This maintenance anomaly problem is similar to the following process in a traditional centralized DBMS: assume a join is defined upon two source relations. Without any appropriate concurrency control, concurrent data updates or schema changes of other write transactions on these two relations during the computation of the join process would cause a concurrency problem. This is clearly similar to the view maintenance anomaly problem.

A traditional DBMS deals with this kind of problem (1) by applying a transaction model to encapsulate all operations that need to be executed atomically into one transaction, such as the two reads of the two relations and the full join computation, and (2) by using concurrency control strategies to guarantee the ACID properties of each transaction [Gray and Reuter 1992]. The ACID properties assure a consistent view of all data inside each transaction and hence solve the problem.

Similarly, in our view maintenance context, the maintenance process can be viewed as a series of read operations over sources that should be executed atomically, while the source updates are independent write operations. There are read and write conflicts since the source writes are autonomous. The above discussion leads us to the idea of applying a transactional model to solve the anomaly problem, as described in detail below.

2.1 Transactions in a Materialized View Maintenance Environment

In a materialized view maintenance environment, each view maintenance process is composed of the following two types of transactions:

- (1) A *source update transaction* at some data source is committed, denoted as “ $T_s = w(DS_i)c(s)$,” where $w(DS_i)$ means the write operation on DS_i , i is the index of the data source, and $c(s)$ means the commit operation.
- (2) A *view maintenance transaction* that computes the delta effects caused by source updates in order to refresh the view extent. During this period, the maintenance transaction reads the view definition to generate the maintenance queries to be sent to the different sources and to compute the maintenance result. After that, it refreshes the view extent. We denote this as “ $T_v = r(VD)r(DS_1)r(DS_2) \cdots r(DS_n)w(v)c(v)$,” where $r(VD)$ denotes the generation of the source-specific maintenance queries based on the view definition VD , $r(DS_i)$ denotes reading from DS_i , $w(v)$ denotes refreshing the view extent, and $c(v)$ denotes the commit process.¹

However, these two types of transactions are not completely independent. As suggested above, the relationship is that each source update transaction would trigger a corresponding view maintenance transaction. Furthermore, this view maintenance transaction is required to see a consistent state of all sources. That is, a view maintenance transaction cannot conflict with other source update transactions. Clearly, some coordination between these two transactions is needed, as otherwise an anomaly may occur as shown in Section 1.2.

These observations give us the hint that a high-level computational transaction model that integrates both of these two subtransactions may be able to achieve such coordination. In addition, the ACID properties of such a high-level transaction should be sufficient to resolve the conflicts between its subtransactions, namely, source update transactions and view maintenance transactions.

2.2 A Transactional Model: VM.Transaction

We model this high-level transaction that integrates both a source update transaction and its corresponding view maintenance transaction as a VM.Transaction.

Definition 1. A VM.Transaction starts after a local successfully committed source update transaction T_s and commits when the materialized view has

¹In schema change maintenance transactions, we may need to update the view definition (VD) for any schema change that affects the view schema before we can generate the maintenance queries. More discussion on this issue can be found in Section 6.2.

been successfully refreshed, that is, the commit of the *view maintenance transaction* T_v . We denote it as “ $T = T_s T_v = w(DS_i)c(s)r(VD)r(DS_1)r(DS_2) \cdots r(DS_n)w(v)c(v)$.”

Previous research has already hinted at this notion of a “global view maintenance transaction” [Zhuge et al. 1996]. However, Zhuge et al. [1996] stated that a global transactional approach is not feasible since the sources may not tolerate such a wait time before being allowed to commit their local updates. We now observe that the nested structure of a VM.Transaction as indicated in Definition 1 can be exploited to avoid this potential drawback. In fact, our work will show the benefits and feasibility of such a transactional approach toward view maintenance. Here, we adopt this idea of taking a transactional approach for view maintenance and put forth the following supporting arguments:

- (1) The VM.Transaction model is a *conceptual* rather than a real transaction model. It encapsulates two types of transactions, namely, source update transactions and their corresponding view maintenance transactions. We observe that the source update transaction can in fact commit autonomously as a sub-transaction within the complete VM.Transaction. Thus sources never need to wait.
- (2) With the help of the VM.Transaction concept, the view maintenance anomaly problem [Zhuge et al. 1995] can be rephrased as the well-known “read dirty data” transaction problem. With an appropriate concurrency control solution, the conflicts between the VM.Transactions, in other words, the view maintenance anomaly problems can naturally be solved.
- (3) Last, this formal transaction-based solution lays a solid foundation for many other maintenance functions and services including handling multiple relations per source or parallel view maintenance, as we will demonstrate later.

2.3 Start, Commit, Abort, and Rollback of VM.Transactions

Since our notion of a VM.Transaction is a conceptual transaction model, it does not have an automatic rollback or abort mechanism. A VM.Transaction is started when a source update has been committed. After that, we must have this source update incorporated into the view to keep the view consistent.² The reason is that the source update transaction is out of the control of the view manager. Thus the rollback of this update cannot be achieved. Hence whenever an error happens during the VM.Transaction, we have to redo the maintenance again until it is successfully committed.

2.4 VM.Transaction Scheduling

Two types of VM.Transaction schedulers are possible, namely, sequential and parallel ones. A sequential scheduler processes one view maintenance transaction at a time, and continues on the next one only after the current one has

²We adopt the view consistency notions from Zhuge et al. [1995] and Zhang et al. [2004]. More details will be discussed in Section 6.3.

been completed. Most view maintenance algorithms in the literature [Zhuge et al. 1995; Agrawal et al. 1997] implicitly employ such sequential scheduling for view maintenance though without utilizing this transaction terminology. Parallel scheduling is more challenging, since out-of-order view maintenance transaction processing introduces additional conflicts [Zhang et al. 2004]. In this article, we will first introduce our solution assuming a sequential scheduler to keep the description simple. Later in Section 6 we demonstrate that our transactional model can be easily extended to support such a parallel scheduler. This is in effect one of the advantages of adopting a transactional model for view maintenance. That is, many services can be easily added to this solid transactional framework.

2.5 Conflicts Between VM_Transactions

We now examine the operations encapsulated within VM_Transactions to identify potential conflicts between them. From Definition 1, we know that there are four basic read and write operations within a VM_Transaction, namely, $w(DS_i)$, $r(VD)$, $r(DS_i)$, and $w(v)$. Hence, three types of conflicts may arise between these operations.³ First, the write-write conflict can occur when some write operation $w_1(DS_i)$ in VM_Transaction T_1 conflicts with $w_2(DS_i)$ in another VM_Transaction T_2 . Second, some write operation $w_1(v)$ in VM_Transaction T_1 conflicts with $w_2(v)$ in another VM_Transaction T_2 . Obviously such write-write conflicts would naturally be taken care of by the local data manager of the sources or the view manager respectively. The third conflict is the read-write conflict, namely, one read operation $r_1(DS_i)$ of VM_Transaction T_1 may read some inconsistent query result written by $w_2(DS_i)$ of another VM_Transaction T_2 .

We can see that the view maintenance anomaly problems illustrated in Section 1.2 arise due to the read-write conflicts of VM_Transactions. That is, we rephrase the view maintenance anomaly problem in terms of the read-write conflicts of VM_Transactions. Hence, as long as we can resolve such read-write conflicts, we solve the view maintenance anomaly problem.

2.6 Serializability of VM_Transaction

Above discussions lead us to a concurrency control strategy for our VM_Transaction model to resolve the conflicts. In other words, we only have to guarantee the serializability of VM_Transactions.

Definition 2. A history of VM_Transactions S is serializable iff it is equivalent to some serial schedule S_1 of these VM_Transactions.

Note that a serial schedule of VM_Transactions S_1 assumes source update occurs after the previous maintenance process. Obviously no view maintenance anomaly would occur in this case. Consequently, no anomaly would happen in its equivalent history S either. Hence, Definition 2 describes the correctness

³A read-write view definition (VD) conflict may arise between VM_Transactions if we schedule view maintenance transactions in parallel. We investigate this type of conflict more closely in Section 6.

criterion of the execution of VM_Transactions. In the next section, we will propose a concurrency control strategy, called *TxnWrap*, to achieve such serializability for VM_Transactions.

3. MULTIVERSION CONCURRENCY CONTROL ALGORITHM

Generally, two basic kinds of concurrency control algorithms [Bernstein et al. 1987], namely, *lock-based* or *multiversion-based*, can be used to address the serializability of traditional transactions. Below we examine which of these two strategies may be appropriate to achieve the serializability of the VM_Transactions in a loosely-coupled view maintenance context.

The traditional *lock-based* concurrency control algorithms [Bernstein et al. 1987] work as follows: for every write, an exclusive write lock is obtained and released at the end of the transaction; and for every read, a read lock is obtained, after the read operation completes, the read lock can be released. This is, however, too restrictive for the schedule of VM_Transactions. By Definition 1, we know that VM_Transaction will read from all sources. Hence a read lock over all sources is required. However, the sources in a distributed environment are often autonomous and would not be willing to accept locks from external agents. Hence locking algorithms are not suitable in such environments.

In the traditional *multiversion* concurrency control algorithms [Bernstein et al. 1987], a writer writes on a version that is different from the one that the reader reads. That is, the reader will read an older version than the one that the writer is currently writing to, thus avoiding the conflicts. In particular, we further differentiate between two categories of multiversion algorithms, that is, the finite version [Mohan et al. 1992; Quass and Widom 1997] and the unrestricted version [Chan and Gray 1985; Agrawal and Sengupta 1989] algorithms. Finite-version algorithms preset a fixed number of maximal versions (often two or some constant n). There is a so called “version expiration” [Mohan et al. 1992; Quass and Widom 1997] phase when the number of versions exceeds the allowed maximum. In that case, we have to switch the reader to a newer existing version to allow for the writers to start making a new version. To achieve this, we require that either the writer synchronizes with all readers on the old version or aborts those readers. In comparison, in unrestrictive multiversion algorithms, the writers can always create a new version and the readers can find whatever the most appropriate versions that they need at any time. No blocking between the reader and writer would ever occur.

Based on above observations, we find that an unrestrictive multiversion algorithm is the appropriate design choice in loosely coupled environments. First, the version that the reader (view maintenance transaction) needs should never expire before the maintenance process has been completed, otherwise the materialized view would be inconsistent. Second, the autonomous writer (source update transaction) is unlikely to synchronize with the reader. For this reason, we favor the choice of an unrestricted multiversion algorithm. Also note that since the VM_Transactions will read a monotonically increasing version of data instead of some random old ones, the version cleanup can be efficiently carried out for old versions. As an added benefit, we can easily extend this concurrency

control strategy to also work with a parallel scheduler as we will show in Section 6.

We now introduce an unrestricted multiversion algorithm based on the above discussion, called *TxnWrap*. The main idea is to extend the functionality of the wrapper for each data source beyond just communicating between sources and the view manager.⁴ That is, the wrapper is responsible for storing the versions generated due to the source updates and thus answering any maintenance query using this versioned information. Below, we describe our *TxnWrap* multiversion algorithm, in particular, first the wrapper materialization of sources at initialization time and then the version management under both update notifications and query requests.

3.1 Wrapper Schema Setup

We propose to extend the wrapper of each source to contain version relations of source data, called *wrapper relations*. A simple full replication of the source data for the wrapper relations would guarantee that we can answer a maintenance query at the wrapper without any further source access. However, we can further improve upon this by various optimization techniques to filter out unnecessary data, as we will discuss in detail in Section 5. Here, for the purpose of understanding the concurrency control strategy, we will first assume that each version relation corresponds to a full copy of the source relation.

We further propose to add two more attributes *V_Start* and *V_End* into the wrapper relation to denote the life time of each tuple. *V_Start* denotes the start version of the tuple (by insertion), and *V_End* denotes the version when the lifetime of the tuple ends (by deletion). We initialize the tuples in the wrapper relation by setting their *V_Start* to 0 and *V_End* to ∞ . This way we guarantee that all initial tuples are visible to all VM transactions.

We also build a *meta relation* at the wrapper for versioning of the source meta data. It contains six attributes, namely, *Rel*, *Attr*, *Rel'*, *Attr'*, *V_Start*, and *V_End*, respectively. The pair of attributes *Rel* and *Attr* describes the source metadata. *V_Start* and *V_End* denote the life span of the attributes, that is, the version number when it's been created, renamed or dropped. The pair of attributes *Rel'* and *Attr'* indicates the new name of a particular attribute or of a relation. The initial source wrapper status for the two sources from Example 1 are now shown in Figures 2 and 3, respectively.

3.2 Version and VM_Transaction Identifiers

A version number is needed in a multiversion management system to identify each version. Similarly, a VM_Transaction identifier is also necessary to track and identify each view maintenance transaction. One simplified strategy is to introduce a *global id* [Chen et al. 2002]. This identifier is issued by some global identifier server in the materialized view manager to assure its global

⁴Note that this is just one of the three architectural choices for where to keep the versioned data. See Section 4 for an elaboration on these choices. Without loss of generality, here we select one of these designs to describe the main idea.

Wrapper for DS ₁					
Relation C'					
Name	Add.	Phone	V_Start	V_End	
Tom	MA	123	0	∞	
Bob	CA	146	0	∞	
Meta Relation					
Rel	Attr	Rel'	Attr'	V_Start	V_End
C'	Name	-	-	0	∞
C'	Add.	-	-	0	∞
C'	Phone	-	-	0	∞

Fig. 2. Initial extent of wrapper of DS₁.

Wrapper for DS ₄					
Relation F'					
Name	Age	FlightNo	Dest	V_Start	V_End
Tom	17	AA4399	Asia	0	∞
Meta Relation					
Rel	Attr	Rel'	Attr'	V_Start	V_End
F'	Name	-	-	0	∞
F'	Age	-	-	0	∞
F'	FlightNo	-	-	0	∞
F'	Dest	-	-	0	∞

Fig. 3. Initial extent of wrapper of DS₄.

uniqueness. In such a scenario, a global ID would serve both as a version number and a VM.Transaction ID. Thus, whenever a source update is reported to the wrapper, the wrapper must first get a global ID from the view manager and then use that ID to create versions. This, however, introduces coordination cost between wrappers and view manager, which may be nontrivial in distributed environments. For this reason, we now introduce two levels of identifiers, a *local ID* within the scope of a data source to identify versions in the wrapper and a *TxnID* within the scope of the whole view maintenance context to manage VM.Transactions. This avoids the extra coordination cost between wrappers and view manager and is thus beneficial for distributed environments.

3.2.1 Local Identifier in the Wrapper. We define a *local ID* to be a timestamp that represents the time the update happened at each individual data source. Without loss of generality, we use an integer k ($k \geq 0$) to represent the local ID. Each local ID is assigned by the corresponding wrapper with each source update. For example, assume four source updates as shown in Figure 4 happened on relations as depicted in Figure 1, then the wrapper will attach a local ID 1 to the update “*INSERT INTO Customer VALUES ('Ben', 'CA', '213')*” because this is the first update to the relation *Customer*. While for the update “*RENAME Customer.Name TO Customer.First-Last,*” its local ID will be 2 because this is the second update in the relation *Customer*. The same rule also applies to the two updates on source relation *FlightRes*. The local IDs in the wrapper will be 1 and 2 respectively.

```

VM_Transaction1 (local id=1, TxnID=[1,0])
INSERT INTO Customer VALUES ('Ben', 'CA', '213');

VM_Transaction2 (local id=1, TxnID=[1,1])
INSERT INTO FlightRes VALUES ('Ben', '18', 'A34', 'Asia');

VM_Transaction3 (local id=2, TxnID=[2,1])
RENAME Customer.Name TO Customer.First-Last;

VM_Transaction4 (local id=2, TxnID=[2,2])
INSERT INTO FlightRes VALUES ('Sue', '20', 'A34', 'Asia');

```

Fig. 4. Four source updates.

3.2.2 VM_Transaction Identifier: TxnID. We also need identifiers to track each global VM.Transaction and help to construct correct maintenance queries that access the appropriate versions of data. A TxnID τ is defined to be a vector $\tau = [k_1, k_2, \dots, k_n]$ with $\tau[i] = k_i$ ($1 \leq i \leq n$ and n is the number of data sources). Each $\tau[i]$ records the current local ID of each DS_i at the time this TxnID is being generated, that is, the largest local ID that has been reported to the view manager thus far. We describe the rules for generating TxnIDs as follows:

- As initial value, the view manager has a vector T with each $T[i] = 0$ ($1 \leq i \leq n$). Each entry $T[i]$ will hold the local ID that has been reported from the source DS_i .
- Whenever an update from source DS_i ($1 \leq i \leq n$) with a local ID k is reported to the view manager, the view manager generates a TxnID τ for this source update using the following three steps⁵:
 - Set $T[i] = k$.
 - Copy T to τ (set $\tau[j] = T[j]$ for all $1 \leq j \leq n$).
 - Return τ as the VM.Transaction identifier.

It's easy to see that though the local IDs in different data sources may be the same, the TxnIDs are globally unique for each update. From the view point of the view manager, each entry of the TxnID vector records the current state of each data source on arrival of the source update to be maintained.

For example, the view *Asia-Customer* as shown in Example 1 is defined on two data sources, one *Customer* and the other *FlightRes*. We thus build a TxnID vector T with two entries, with each entry containing the current local ID of its corresponding data source. Without loss of generality, we use $T[1]$ to hold the local ID from relation *Customer* and $T[2]$ to hold the local ID from the relation *FlightRes*. Initially, we have $T[i] = 0$ ($1 \leq i \leq 2$). For the four updates depicted in Figure 4, if we also assume that they arrive at the view manager in the order as we have defined them, then the TxnIDs for these updates will be $[1, 0]$, $[1, 1]$, $[2, 1]$, and $[2, 2]$, respectively. That is, when the first data update (the VM.Transaction1) with a local ID 1 arrives at the view manager, we set $T[1] = 1$ and return $[1, 0]$ as its TxnID. While for update VM.Transaction2, we set $T[2] = 1$ and get $[1, 1]$ as the TxnID, and so on. As can be seen, each entry

⁵A FIFO transferring of updates from data source to view manager for updates from the same data sources is required, which is also assumed by most of the literature in the view maintenance area.

in TxnID records the right versioned data that needs to be accessed from the corresponding wrapper in its view maintenance transaction.

3.3 Version Operations

Basically, there are three classes of version operations that the wrapper must perform. First, the wrapper needs to create versions whenever an update is reported by a source. Second, the wrapper needs to read the appropriate versions to answer each maintenance query issued by the view manager. Third, any versions that are no longer required by any future VM.Transactions need to be cleaned up.

3.3.1 Version Creation in the Wrapper. The updates reported by a source can be categorized as data updates, which are *insert*, *delete*, and *update*, or schema changes, which are *add*, *drop*, and *rename* an attribute or *create*, *drop*, and *rename* a table. To create versions, a local ID i will be assigned as a version number by the wrapper to create a new version for each source update. The corresponding version creation rules are shown below:

- (1) Delete a tuple t : update $t[V_End]$ to be i .
- (2) Insert a tuple t : insert this tuple, set $t[V_Start]$ as i and set $t[V_End]$ ∞ .
- (3) Update a tuple t : treat it as a delete followed by an insert.
- (4) Drop an attribute: update the attribute's V_End to be i in the metarelation.
- (5) Add an attribute: insert it into the metarelation with V_Start to be i and V_End to be ∞ .
- (6) Rename an attribute: in metarelation, set Rel' and $Attr'$ of the old attribute to its new names, update V_End to i . Then add the renamed attribute to metatable.
- (7) Drop a table: treat it as drop of all its attributes.
- (8) Add a table: add all its attributes to the metarelation.
- (9) Rename a table: method similar to rename an attribute, but operate on all attributes.

As an example, consider the four source updates shown in Figure 4. After the corresponding version creation step, the wrapper relation and metadata relation are shown in Figures 5 and 6.

3.3.2 Version Read in the Wrapper. The read version operation happens when a VM.Transaction wants to query a source. Rather than being processed by the source, the queries are interpreted by the wrapper and executed by reading the appropriate versions from the wrapper data and metadata relations. Notice that the VM.Transaction requires to see a consistent state of all sources. Here a consistent state means both source data and metadata.

To achieve this, the first step is to rewrite the original maintenance query by augmenting it with the version conditions "*AND $V_Start \leq \tau[i]$ AND $\tau[i] < V_End$* ." This rewritten maintenance query will be evaluated over the wrapper data relation. For example, Figure 7 depicts the rewritten versioned maintenance query for the three source data updates from Figure 4. The third rename

Relation C'		Wrapper for DS_1			
Name	Add.	Phone	V_Start	V_End	
Tom	MA	123	0	∞	
Bob	CA	146	0	∞	
Ben	CA	213	1	∞	

Meta Relation					
Rel	Attr	Rel'	Attr'	V_Start	V_End
C'	Name	C'	First-Last	0	2
C'	Add.	-	-	0	∞
C'	Phone	-	-	0	∞
C'	First-Last	-	-	2	∞

Fig. 5. DS_1 wrapper content after updates.

Relation F'		Wrapper for DS_4			
Name	Age	FlightNo	Dest	V_Start	V_End
Tom	17	AA43	Asia	0	∞
Ben	18	A34	Asia	1	∞
Sue	20	A34	Asia	2	∞

Meta Relation					
Rel	Attr	Rel'	Attr'	V_Start	V_End
F'	Name	-	-	0	∞
F'	Age	-	-	0	∞
F'	FlightNo	-	-	0	∞
F'	Dest	-	-	0	∞

Fig. 6. DS_4 wrapper content after updates.

VM_Transaction with TxnID=[1,0]
SELECT Name, Age, FlightNo, Dest FROM F'
WHERE Name='Ben' AND $V_Start \leq 0$ AND $V_End > 0$

VM_Transaction with TxnID=[1,1]
SELECT Name FROM C' WHERE
Name='Ben' AND $V_Start \leq 1$ AND $V_End > 1$

VM_Transaction with TxnID=[2,2]
SELECT First-Last FROM C' WHERE
Name='Sue' AND $V_Start \leq 2$ AND $V_End > 2$

Fig. 7. Version queries at wrapper.

schema change in Figure 4 results in the rewriting of view definition without any maintenance queries.

The second step is to make sure the metadata specified in the maintenance query is consistent with the wrapper data relation. If not, we need to evolve the schema of the wrapper data relation. More precisely, if the version number of the maintenance query “ $\tau[i]$ ” is larger or equivalent than any “ V_End ” in the wrapper metadata relation, then we know that the corresponding wrapper data relation’s schema is outdated. For example, to maintain the fourth update in Figure 4, the maintenance query in Figure 7 needs to access the “First-Last” column, which is not available yet in the wrapper data relation. Hence,

the corresponding wrapper data relation has to be evolved by (1) deleting the corresponding tuples in the metadata relation and (2) renaming/dropping the corresponding columns/tables in the wrapper data relations. The above has to be done as one atomic action. In this example, we delete the (C' , Name, . . .) tuple from metadata relation and rename the column “Name” to “First-Last.” After this metadata consistency checking, we can process the maintenance query. Note that we can keep those metadata tuples with $V_End <> \infty$ (there are typically only few such tuples) in memory to speed up this step.

3.3.3 Version Cleanup in the Wrapper. Finally, we clean up versions that are no longer needed by any maintenance transaction. The primary purpose of this operation is to reduce the size of the wrapper relations in order to speed up the maintenance queries and to evolve the wrapper schema. For a sequential scheduler, the cleanup is straightforward. That is, after the completion of a VM.Transaction with TxnID τ , for all the local ids $\tau[i]$ ($1 \leq i \leq n$) recorded in TxnID τ , the data with V_End less than $\tau[i]$ are no longer required by any later maintenance transaction. Hence, they can be deleted from the corresponding wrapper relation specific to DS_i . For the wrapper data relation, we simply delete the tuples with $V_End \leq \tau[i]$. For wrapper metadata relation, we atomically execute the update as “*DELETE * FROM meta-table WHERE V_End $\leq \tau[i]$* ” and perform the corresponding schema change operations on the wrapper data relation, for example, drop or rename attributes. For a parallel scheduler (Section 6), intuitively, we have to assure that all VM.Transactions that contain a local ID less than $\tau[i]$ are completed before that respective version removal can proceed.

3.4 TxnWrap: A Multiversion Algorithm

We are now ready to introduce our TxnWrap concurrency control algorithm (TxnWrap) for materialized view maintenance.

- (1) First, the TxnWrap initializes the wrapper relations as defined in Section 3.1.
- (2) When a source commits a local update transaction and reports its updates to the wrapper (or extracted by the wrapper), the TxnWrap first assigns a local ID as a version ID and uses it to create versions in the wrapper. After that, the wrapper notifies the view manager about these updates. The view manager builds a TxnID by collecting and concatenating all the current *local IDs* of each data source for this VM.Transaction.
- (3) When the wrapper receives a maintenance query from the view manager, then the TxnWrap rewrites the maintenance query according to the corresponding local ID recorded in the VM.Transaction identifier (TxnID) and executes it upon the wrapper relations.
- (4) When a VM.Transaction is committed, version cleanup starts based on the type of scheduler, that is, sequential or parallel. The cleanup action can be done by a background daemon process.

The correctness criterion of TxnWrap algorithm is characterized by Theorem 1.

THEOREM 1. *A history of VM_Transactions scheduled by the TxnWrap is serializable.*

We prove by contradiction in the electronic appendix that each TxnWrap multiversion (MV) history is serializable by showing its multiversion serial graph (MVSG) to be acyclic based on transaction theory [Bernstein et al. 1987]. By Theorem 1 and Definition 2, we know that TxnWrap maintains the views correctly. This means that the extent of the view after maintenance reflects the right states of the data sources.

3.5 Handle Multiple Relations in One Data Source

The discussions on the VM_Transaction model and its multiversion concurrency control algorithm (TxnWrap) so far assume each data source only has one relation. Such, an assumption, which has also been made by the other approaches in the literature thus far [Agrawal et al. 1997; Zhang and Rundensteiner 2000], can be easily lifted within our model as shown below. The VM_Transaction model naturally fits if we define read and write operations on relations instead of data sources. For example, assume a view is defined upon two data sources involving five relations, for example, $DS_1(R_1, R_2, R_3)$ and $DS_2(R_4, R_5)$.⁶ A VM_Transaction that models a source update to relation R_2 at DS_1 can be denoted as “ $T = w(R_2)c(R_2)r(VD)r(R_1)r(R_3)r(R_4)r(R_5)w(v)c(v)$.” Correspondingly, the TxnID will also be generated based on the local ids of source relations instead of data sources. Hence, in TxnWrap, we have the same number of wrapper data relations in the source wrapper as the relations in the data source.

For generating maintenance queries when one data source contains multiple relations, the view manager could send *one* combined maintenance query that involves all the relations in that data source instead of multiple maintenance queries to each relation. Correspondingly, the wrapper should also append all the version conditions of these relations.

For example, a data update ΔR_2 at relation R_2 with TxnID τ can be maintained using the following two queries: (1) A query “ $R_1 \bowtie \Delta R_2 \bowtie R_3$ ” to DS_1 ’s wrapper. Here we assume the result of this query is ΔDS_1^1 , which denotes the effects of ΔR_2 on data source DS_1 . (2) A query “ $\Delta DS_1^1 \bowtie R_4 \bowtie R_5$ ” to DS_2 ’s wrapper. Note that the source wrapper would append all related version conditions when receiving a combined maintenance query. For example, the version conditions for the first query will be “ $AND R_1_Wrapper.V_Start \leq \tau[1] AND \tau[1] < R_1_Wrapper.V_End AND R_3_Wrapper.V_Start \leq \tau[3] AND \tau[3] < R_3_Wrapper.V_End$.” Here $R_1_Wrapper$ and $R_3_Wrapper$ are the corresponding wrapper relations for R_1 and R_3 .

Hence, the TxnWrap algorithm can be easily extended to handle multiple source relations. This also demonstrates an additional benefit of adopting the

⁶Relations in different data sources could have same name as long as they are globally unique, that is, the concatenation of data source name with the relation name. For simplicity, we use different names in this example to illustrate the technique.

TxnWrap approach. Prior work such as SWEEP [Agrawal et al. 1997] has to send maintenance queries to each source relation even if they exist in the same database. While in a separate work, Varde and Rundensteiner [2002] tried to address this using a schema of two-level compensations.

4. TXNWRAP VERSION DESIGN STRATEGIES

4.1 Version Placement Design Choices

In the previous section, our concurrency control strategy, TxnWrap, will materialize versioned data in a special-purpose source wrapper. Clearly this is not the only design choice for the versioned data. Alternative choices include to either (1) add extra version information directly into the source data store or (2) to materialize the versioned data on the view server itself. We now discuss the tradeoffs between these version placement design options.

First, we consider the case of adding version information directly into each source. This would imply that two version attributes as described in Section 3.1 are directly added to each of the source relations. Clearly, this option reduces the space cost compared to storing versioned source copies at wrappers. However, the source update transactions would now be affected since they also have to deal with the version information. Such version information serves the sole purpose of external view maintenance support. This design choice thus may not be acceptable in a loosely coupled environment since the sources may not want to degrade their transaction performance.

The second design choice would be to materialize the versioned relations inside the view manager. One potential benefit is that instead of decomposing the view maintenance query into individual remote source queries to each source, we now could use one single combined maintenance query. The reason is that all the relations are within the same view site database. This design choice is similar to the view self-maintenance approach, such as in Quass et al. [1996], where the source data is selectively replicated in the view site for maintenance.

However, in this case, the view manager would have to maintain both the materialized view and the versioned copied data as well as answer user queries. This may become a heavy burden for the view manager. Furthermore, since all the maintenance queries are against the data on the same view site, we may not be able to improve the parallel maintenance performance (an additional functionality gained by taking a multiversion-based approach as demonstrated in Section 6). This is because a single view server may not have enough system resources to handle a large number of maintenance queries in parallel. Another potential drawback is that if we have multiple warehouses, such as data marts, that share the same source tables, then the source data has to be copied into each view site, resulting in a significant overall space cost for the entire environment.

In comparison, placing the versioned data into a dedicated wrapper achieves the independence and simplicity of both the sources and the view manager by moving the complex version management responsibility into a third dedicated place, as we have illustrated in Section 3. First, we can reduce the overall space

cost for applications such as data marts since one wrapper can support multiple warehouses. Second, the performance gain for parallel processing can be effectively explored since the maintenance query workload is distributed to each wrapper. Our experimental studies in Sections 7.6 and 7.7 show indeed significant parallel maintenance performance improvements gained by this design choice.

4.2 Comparison with Self-Maintenance Solutions

Next, we compare our solution to the self-maintenance approach such as in Quass et al. [1996]. In Quass et al. [1996], the authors proposed to selectively materialize the source data at the warehouse so that the view maintenance can be achieved without accessing the sources. Note that this approach also has to make multiple copies for data mart applications due to its data materialization at the warehouse. Since the self-maintenance approach [Quass et al. 1996] does not explicitly assume any concurrency control strategies, we now investigate how the concurrency issues could be tackled in the context of this approach.

Since the source copies are kept at the warehouse, the updating of source copies and the updating of views can be combined into a single maintenance transaction. However, in this case, deadlocks may occur. Here is a simple example. Assume the view is $R_1 \bowtie R_2$ and there are two source updates ΔR_1 and ΔR_2 . R'_1 and R'_2 are the corresponding source copies materialized at the data warehouse. $\Delta R'_1$ and $\Delta R'_2$ are the changes to the source copies based on ΔR_1 and ΔR_2 . The maintenance transactions for these two updates are $w(R'_1)$ followed by $\Delta R'_1 \bowtie R'_2$ and $w(R'_2)$ followed by $\Delta R'_2 \bowtie R'_1$, respectively. If the write operations acquire table-level locks, then the deadlocks may occur. The reason is that the first transaction reads R'_2 table and writes R'_1 table, while the second transaction reads R'_1 and writes R'_2 table. There are cyclic read-write conflicts on R'_1 and R'_2 tables between these two transactions, resulting in deadlocks. Even if fine-granularity locking mechanisms, such as low-level locks are used, the deadlocks still may occur. For example, when $\Delta R'_1$ and $\Delta R'_2$ contain any tuples that are joinable to each other, the deadlock occurs on those tuples. Such deadlocks may introduce significant extra cost. In comparison, our TxnWrap approach does not have such deadlock problems.

An alternative transaction processing strategy for self-maintenance approaches is to separate the action of updating source copies and the updating of views as two independent transactions. In this case, although these two transactions can be executed in parallel, the same maintenance anomalies may occur as in Section 1.2. We then must apply either compensation-based or multiversion-based solutions to solve the anomalies as discussed in this article. A multiversion-based solution outperforms the compensation-based one in that it achieves a more steady performance and allows parallel maintenance as shown via experiments in Sections 7.2, 7.3, and 7.5. Hence, based on the above analysis, the existing self-maintenance approaches can be enhanced by our TxnWrap solution to achieve better maintenance performance. In short, self-maintenance approaches themselves could benefit from the TxnWrap solution.

5. WRAPPER OPTIMIZATION STRATEGIES

5.1 Static Reduction Optimization

Version management assuming a full duplication of a source relation requires significant storage space. For this reason, several optimization strategies can be applied to reduce the storage. First, we can optimize the storage by analyzing the view definition. *Selection optimization* extracts the local condition in a view definition and then pushes it down to the wrapper relation definition to reduce the wrapper size. For example, in Example 1, by pushing the local condition of the view definition, that is, “ $F.Dest = Asia$ ” into the wrapper relation definition, we can reduce the wrapper space for source relation *FlightRes*. This results in the additional advantage that some data updates with “ $F.Dest \neq Asia$ ” can now be filtered early on as being irrelevant. They thus would not even be reported to the view manager. *Projection optimization* projects any attribute of the source not involved in the view definition from the wrapper. Thus, for source relation *Customer* of DS_1 in Example 1, the attributes *Address* and *Phone* can be safely dropped from the DS_1 wrapper.

Second, we can explore existing constraints to reduce the space cost. For example, the key and foreign key constraints can also be used to reduce the versions. Assume two tables *A* and *B* with column *A.a* referencing primary key *B.b*. Assume the materialized view defined as $A \bowtie B$ with join condition $A.a = B.b$. In this case, we can avoid to make any version copy of table *A* since the maintenance does not need to access table *A* at all. Similar ideas are proposed for the self-maintenance approach such as in Quass et al. [1996]. In fact, all space optimizations strategies developed for self-maintenance approaches should also be applicable here.

5.2 Dynamic Version Optimization

While the above strategies reduce the wrapper space statically, we now consider how to optimize the wrapper space dynamically. Here we measure the versioning cost by comparing the number of tuples of the wrapper relations to that of the source relations. We find that eventually after all the updates have been maintained and all the versions have been cleaned up, the wrapper relation contains the same number of tuples as the sources with just two extra version attributes. We thus consider this to be the minimum version cost. However, the intermediate state of the wrapper relation may also contain more tuples than the sources.

The first case is when we have source updates that result in both inserts and deletes of versions. For example, the update of all “*Price*” columns to “*Price*”*0.9 will double the size of the wrapper data relation. Although such tuple explosion is temporary and will be cleaned up, it may still significantly degrade the performance.

To solve this problem, if the wrapper identifies some update that potentially affects many of the current tuples (by analyzing both statistics and the update statement itself), then we propose to hold updating the wrapper data relation and just send the view manager a special update message. When the view

manager starts to maintain this update message, the wrapper will be notified to directly update the wrapper data relation without any tuple duplication. Note that by this strategy we need to update the modified tuples' V_Start using the new local ID. This is a different versioning strategy from what we initially would do based on Section 3.3.1. Also any subsequent changes to this relation have to be recorded in a special manner at the wrapper due to semantic issues. They can be either recorded as statements (such as update and delete statements) or deltas (such as inserts). This strategy avoids the tuple explosion problems with minimal extra cost, i.e., logging an update statement is a lot cheaper than duplicating a huge chunk of the relation.

The second case is when we delete a large number of tuples from one relation. Instead of updating these tuples in the wrapper, we can apply similar techniques as above to delete them directly thus avoiding the large intermediate version space.

6. PARALLEL MAINTENANCE SCHEDULER

6.1 Parallel Scheduling in a Data Update Only Environment

One advantage of using such a principled sound approach toward view maintenance, namely, a transactional approach, is that other transactional-based features can be gotten for “free.” As one example of this, we discuss in this section how TxnWrap can be extended to support *parallel view maintenance*. First, we note that a direct extension of the sequential TxnWrap scheduler would achieve provably correct parallel maintenance for view maintenance transactions in a data update only environment. As described in Section 2.1, one *data update view maintenance transaction* is represented as $r(VD)r(DS_1)r(DS_2) \dots r(DS_n)w(v)$. Thus, there will be no read block between view maintenance transactions in TxnWrap if we apply a multiversion concurrency control strategy.⁷ Borrowing from traditional concurrency control concepts [Bernstein et al. 1987], an *aggressive scheduler* can thus straightforwardly be applied. That is, we can start maintaining data update view maintenance transactions concurrently as long as sufficient computational resources are available in the view manager because there are no read/write conflicts in these transactions due to the multiversion concurrency control strategy employed by TxnWrap.

6.2 Parallel Scheduling in a Data and Schema Change Environment

However, more issues must be dealt with if we take source schema changes into consideration. To illustrate this, let us first briefly review how schema changes are maintained [Lee et al. 2002; Chen et al. 2001]. There are three steps for maintaining a source schema change:

- Determine which view is affected by the source schema change: $[r(VD)]$.
- Find the suitable replacement for schema elements removed from the view definition and rewrite the view definition if needed: $[w(VD)]$.

⁷We consider the $w(v)$ operation in view maintenance transactions separately in the commit control discussion (Section 6.3).

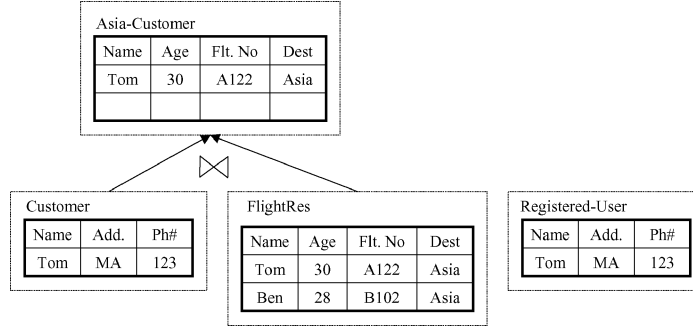


Fig. 8. View Asia-Customer before updates.

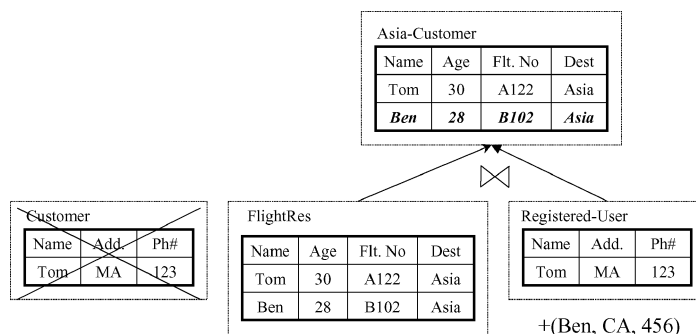
— Calculate the delta changes in term of tuples to be added or to be deleted due to the replacement between the old and the new view definition and adapt the view extent by committing these delta changes: $[r(\mathbf{VD})r(\mathbf{DS}_1)r(\mathbf{DS}_2) \cdots r(\mathbf{DS}_n)w(v)]$.

Thus, one *schema change view maintenance transaction* can be represented as $\mathbf{r}(\mathbf{VD})\mathbf{w}(\mathbf{VD})r(\mathbf{VD})r(\mathbf{DS}_1)r(\mathbf{DS}_2) \cdots r(\mathbf{DS}_n)w(v)$. Clearly, the view definition (VD) represents a critical resource among the different view maintenance transactions. Hence, when scheduling the maintenance for mixed data updates and schema changes, we have to consider the potential $r(\mathbf{VD})/w(\mathbf{VD})$ conflicts besides the anomaly problems (the read-write conflicts between the *source update transactions* and the *view maintenance transactions*).

In a view maintenance environment, we keep the assumption of a FIFO network for updates being transmitted from the sources to the view manager which come from the same data source, otherwise certain updates wouldn't be correctly maintained. For example, assume two updates “ U_1 :INSERT INTO Customer VALUES ('Ben', 'CA', '213'),” “ U_2 :DELETE FROM Customer WHERE Name = 'Ben'” happened on the same source relation Customer as shown in Example 1 in this given order. Then we should maintain U_1 before U_2 in the view manager. If not, it is possible that the maintenance result of U_2 couldn't be refreshed or may be incorrect because the corresponding tuple isn't in the view extent yet. Thus, we can't reorder view maintenance transactions randomly.

In addition, more ordering restrictions need to be imposed once we assign the corresponding TxnID to each update. To clarify this, we first define the order of TxnIDs as follows: given two TxnIDs τ_i and τ_j , $\tau_i < \tau_j \Leftrightarrow (1) \forall h, 1 \leq h \leq n$ (n is the size of TxnID vector), $\tau_i[h] \leq \tau_j[h]$. (2) $\exists k, 1 \leq k \leq n, \tau_i[k] < \tau_j[k]$. In situations that there is a mixture of data and schema changes to be maintained, we can't randomly reorder these view maintenance transactions once their TxnIDs are assigned even if these updates come from different data sources. Otherwise the maintenance result may also be inconsistent with the data source state. We use the following example to illustrate this ordering restriction.

Example 2. Figure 8 illustrates the view Asia-Customer first defined in Example 1. For simplicity, we use Customer \bowtie FlightRes to represent the view. Assume that after a certain time period, source relation Customer becomes

Fig. 9. View *Asia-Customer* after updates.

unavailable, that is, a schema change “*DROPTABLE Customer*” has happened. Based on view synchronization [Lee et al. 2002; Nica et al. 1998] techniques, the view manager will locate a replacement (here we use the relation *Registered-User* as the replacement of relation *Customer*) to update the view definition and its extent. After that, a data update “*INSERT INTO Registered-User VALUES ('Ben', 'CA', 456)*” also happened on the relation *Registered-User*. Assume these two updates are reported to the view manager for maintenance following the order as defined above. Assume they are assigned TxnIDs [1, 0, 0] and [1, 0, 1], respectively. After the correct maintenance, the final view definition will evolve to $Registered-User \bowtie FlightRes$, and it will have two tuples as depicted in Figure 9.

However, now assume we maintain the data update with TxnID [1, 0, 1] first. Then no changes to the view extent would be recorded because the source relation *Registered-User* is not in the view definition of the *Asia-Customer* yet. While maintaining the update [1, 0, 0], the view manager cannot see the data update’s effect because the TxnID tells us that it can only see the state “0” of relation *Registered-User* in this maintenance process. State “0” denotes the state of the *Registered-User* relation before the data update happened. Thus, the maintenance result of this data update will be lost. In short, we cannot change the scheduling order of these two updates after we have assigned them the corresponding TxnIDs.

As can be seen, this ordering restriction is due to the concurrency conflicts on the critical resource, namely, the view definition *Asia-Customer*. That is, the view definition could be changed by some source schema changes while other view maintenance transactions still need it. As we discussed before, this concurrency can also be resolved either using *lock-based* or *multiversion strategies*. However, in this particular context, a multiversion based strategy does not have any preference because we have to keep the TxnID order to generate the right view definition version anyway. For example, for the two updates in Example 2, the data update with TxnID [1, 0, 1] cannot be processed until the new version of the view definition has been built, that is, after the schema change with TxnID [1, 0, 0] completes at least its view synchronization (VS) portion. Thus, a *multiversion* approach doesn’t offer us any benefits rather the extra overhead

```

/* Parallel schedule the updates need to be maintained in the Txn_Queue */
1:  Parallel_Scheduler (TXN_Queue) {
2:      /* TXN_Queue contains unmaintained updates */
3:      while (TXN_Queue != empty) {
4:          VM_TXN = get first VM_TXN from TXN_Queue;
5:          /* get first VM_Transaction */
6:          if (VM_TXN.getType() = DU) { /* data update VM_Trans. */
7:              Build a new thread to maintain VM_TXN or wait until new thread is available;
8:          } else { /* a schema change VM_Trans. */
9:              Wait until all running maintenance threads commit;
10:             View synchronization (VM_TXN);
11:             Build a new thread to wrap view adaptation(VM_TXN);
12:             /* synchronize view definition update, while run view adaptation process in parallel
13:              with other processes */
14:         }
15:     } /* end of TXN_Queue equals empty */
16: } /* end of Parallel_Scheduler */

```

Fig. 10. Parallel scheduling process.




Fig. 11. Parallel scheduling example.

of versioning and version control. Based on the above analysis, we propose the following parallel maintenance scheduler using a *lock-based* control strategy to handle the conflicts in the view definition access due to some schema changes:

- Start view maintenance transactions based on their TxnID order.
- Synchronize the $w(VD)$ operations of schema change maintenance. That is, no $r(VD)$ or $w(VD)$ operations will start before the previous $w(VD)$ has been finished.

Figure 10 lists the basic logic of this parallel scheduler. A sample execution plan is depicted in Figure 11. DU and SC stand for data update maintenance transaction and schema change maintenance transaction, respectively. That is, each thread is responsible of one maintenance transaction. These threads can be run in parallel except each schema change maintenance will synchronize these processes.

6.3 Commit Control and Consistency

We adopt the notions of consistency of a view extent depending on how and when the source updates are incorporated into the materialized view in a distributed environment from Zhang et al. [2004] and Zhuge et al. [1995]. That is, five consistency levels (convergence, weak consistency, consistency, strong

consistency, and complete consistency) of the materialized view are defined based on the materialized view states according to source updates. However, in a parallel scheduling environment, even if each individual update is being maintained successfully, the final view state after committing these effects may still be inconsistent. This Variant Commit Order problem in a data update only environment has been addressed in Zhang et al. [2004]. We can apply the same commit control strategy to our mixed data update and schema change environment. The control strategy we apply in our current parallel TxnWrap work is a strict commit order control. That is, only after we commit all the effects of all previous updates, can we begin to commit the current delta changes to the view extent. As we will show in the electronic appendix, it reaches at least a strong consistency level if we apply this strict commit order control strategy.

6.4 More Discussions on the Parallel Scheduler

The limitation of the above parallel scheduler is that once we assign the TxnIDs based on the arrival order of updates at the view manager, we then must stick with this order when scheduling transactions. For example, all the subsequent data updates that arrive after a schema change have to wait for the schema change to finish its view synchronization portion. This leads us to think about whether other less strict parallel scheduling solutions may be possible within the TxnWrap context. For this, we first need to determine whether it is possible to change the scheduling order of updates in the view manager while still keeping the view consistent.

OBSERVATION 1. Any maintenance order for updates from different data sources will generate the same and correct final maintenance results.

This observation allows us to assign the corresponding TxnIDs dynamically. For example, if we assign the data update in Example 2 the TxnID [0, 0, 1] as if it comes to the view manager first, while assigning the schema change TxnID [1, 0, 1], the final maintenance result of these two updates in this newly assigned order would also be correct. This observation gives us the hint that we should be able to exchange the scheduling order of updates in the view manager that come from different data sources as long as we assign the TxnIDs correspondingly. This way, a more flexible scheduler can be designed by freely exchanging the maintenance orders for updates that come from different data sources. We are thus able to reduce the wait of the data updates that arrived after a schema change.

7. EXPERIMENTAL EVALUATION

7.1 Experimental Testbed

Since TxnWrap by itself is not a stand-alone system, we have plugged it into the existing Dynamic Data Warehouse Maintenance (DyDa) system developed at WPI [Chen et al. 2001]. The basic DyDa system uses SWEEP [Agrawal et al. 1997], a compensation-based algorithm to handle concurrent data updates. It also has some enhanced view adaption strategies to detect and correct a mixture

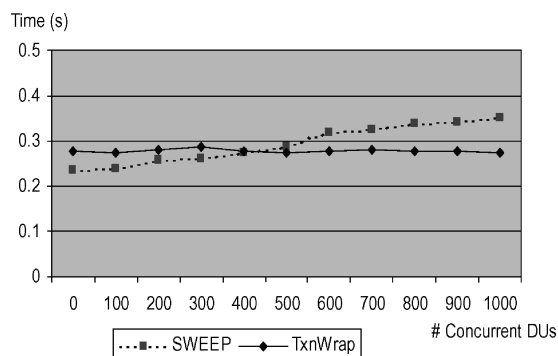


Fig. 12. Concurrent data update process.

of concurrent data updates and schema changes [Chen et al. 2004]. By disabling this concurrency logic and instead plugging TxnWrap into the DyDa system, we now have a transactional-based maintenance solution. The Java code to handle concurrencies in the DyDa system is more than 5000 lines, while in contrast the TxnWrap module contains fewer than 900 lines of code. In this section, we experimentally evaluate TxnWrap by comparing these two systems.

In our experimental study, we have conducted our experiments on four Pentium III PCs with 256M memory, running Windows NT and Oracle8i. We employ six data sources with one relation each evenly distributed over three machines. Each relation has four attributes and 100,000 tuples. There is one materialized join view defined upon these six source relations residing on the fourth machine.

Currently, we set each VM.Transaction in our experiments to contain exactly one update, as to compare the performance with other algorithms in a similar configuration. The issue of batching multiple data and schema updates in one VM.Transaction is orthogonal to this article. We thus ask readers to refer to Liu et al. [2002b] for details on this topic and corresponding experimental results.

7.2 Data Update Processing

In this experiment, we evaluate the overhead of TxnWrap for data update processing under a varying rate of concurrency compared to the SWEEP [Agrawal et al. 1997] algorithm. SWEEP applies local compensation to remove the *data update anomaly* as shown in Section 1.2. Figure 12 shows the average maintenance cost of one data update (on the *y*-axis) under different numbers of concurrent data updates (on the *x*-axis) for both systems.

From Figure 12, we can see that initially the maintenance cost of TxnWrap is larger than that of SWEEP when the number of data updates concurrent with the one currently being processed is zero (i.e., there is no concurrent data update). This cost is mainly caused by the versioned maintenance query, which requires additional version conditions to be evaluated. The maintenance cost of SWEEP increases with an increase in the number of concurrent data updates, because it has to use local compensation to deal with concurrent data updates. The more updates are found to be concurrent when handling one update, the




Fig. 13. Concurrent schema change process.

more compensation has to be undertaken. In comparison, TxnWrap exhibits an almost fixed cost for the handling of data updates independent of the number of concurrent data updates, and thus achieves a steady performance.

7.3 Schema Change Processing

In this experiment, we examine the performance of TxnWrap for schema change processing, which includes both view synchronization and view adaptation. We measure the maintenance cost for twenty schema changes while varying the time intervals between them. As described in Section 1.2, if there are some concurrent schema changes, the query would be broken and the current maintenance process may be aborted in the DyDa system. This is a significant cost. Figure 13 presents the maintenance as well as the abort cost (as part of the maintenance cost) of both systems (on the y -axis). From the figure, we see that the DyDa system has a varying abort cost under different system loads. The explanation of the curves of the DyDa system can be found in Chen et al. [2004]. Eventually when the time interval is larger than the maintenance time, the updates will not conflict with each other. In this case, both systems have the same maintenance cost. In comparison, TxnWrap, employing a multiversion algorithm, avoids any abort and hence again achieves a steady performance.

7.4 Impact of Version Management Optimizations

In this experiment, we study the performance achieved when applying the wrapper optimization strategies described in Section 5. Since various optimization strategies all have the same goal, namely, to reduce the size of wrapper relations, we study how such size reduction affects the maintenance performance.

We measure the average maintenance cost of one data update or one schema change under different size reduction factors for wrapper optimization and compare TxnWrap to the nonmultiversion solution DyDa. As shown in Figures 14 and 15, the x -axis is the size of the wrapper relations after optimization over that before optimization. Clearly, we find that the smaller the wrapper relations, the better maintenance performance.




Fig. 14. Data update processing with wrapper materialization optimization.

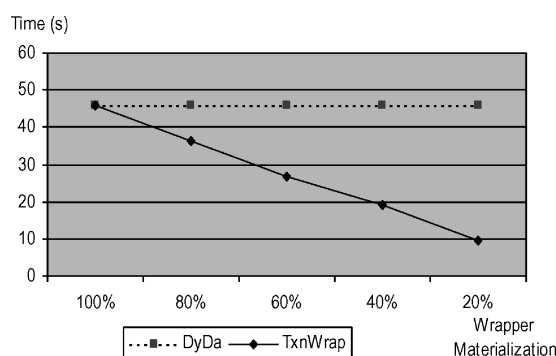


Fig. 15. Schema change processing with wrapper materialization optimization.

We note that, in Figures 14 and 15, there are still several differences between data update maintenance and schema change maintenance. There is a crossover point for data update maintenance. The reason is that the extra version attributes and conditions complicate the maintenance queries and are thus initially more costly to evaluate. In comparison, the maintenance of a schema change is much more expensive [Lee et al. 2002]. The performance largely depends on the number of tuples of the relations. The extra version attributes and conditions have a relatively small overhead. Hence, the performance improvement is more significant for schema change maintenance.

7.5 Parallel View Maintenance

7.5.1 Basic Parallel View Maintenance Performance. We have experimentally evaluated TxnWrap when extended with a parallel view maintenance scheduler. We use the same system setup as described above and use a workload of 500 data updates and vary the number of maintenance threads, that is, the maximum number of source updates that can be maintained by the system at the same time. The total maintenance time of these 500 data updates is depicted in Figure 16. The x -axis denotes the number of parallel threads in the system with S denoting the serial scheduler. The y -axis represents the total processing time.

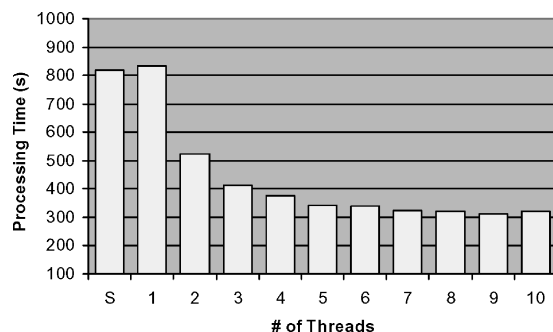


Fig. 16. Parallel data updates maintenance.




Fig. 17. Impact of source relation distribution.

If we only use one thread, then the total processing time is very similar to the serial one. There is some fairly small overhead for the parallel maintenance scheduler and thread management. Given our system setup, the total processing time reaches its minimum with five parallel threads. A maximum percentage of performance improvement of 150% is observed. When we further increase the thread number, the processing time achieves only some small extra improvement. The main reason is that the parallel processing power of wrapper databases is limited. They are not able to effectively process a large number of maintenance queries concurrently due to resource contention.

7.5.2 Impact of Source Relation Distribution. Figure 17 shows the changes of the parallel maintenance performance due to the distribution of source relations. Here we assume that each data source has a similar query processing capability. We vary the distributions of the six source relations from one data source to six data sources. The more data sources, the smaller the number of source relations located at each data source. Seen from Figure 17, if we distribute the source relations to more data sources, then a higher rate of performance improvement can be achieved. For example, if the six source relations are evenly distributed over the six data sources, then the maximal percentage of performance improvement in our setting reaches around 250% when using five parallel threads. This is as expected because more maintenance queries can be answered in parallel by all these data sources.

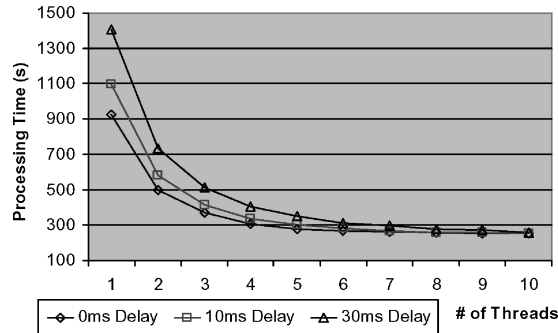


Fig. 18. Impact of network delay on parallel maintenance.




Fig. 19. Parallel maintenance of both data updates and schema changes.

7.5.3 Impact of Network Delay. In Figure 18, we measure the effect of changing the network delay in each maintenance query. A network delay is added when evaluating each maintenance query based on the average time to transfer a tuple across the network. For example, if we assume that the average time to transfer a tuple with 20 bytes is ℓ , then it takes $k \cdot 2 \cdot \ell$ to transfer k tuples with 40 bytes each. The performance changes from no network-delay to “ $\ell = 10$ ms” and then “ $\ell = 30$ m” in each maintenance query are listed in Figure 18. From the figure, we can see that while the high network delay has a negative effect on the serial TxnWrap performance, such effects are quickly offset by using our parallel TxnWrap scheme. The reason is that although the network transfer is slow, we can still run many other tasks in parallel, which improves the overall throughput. This is an especially useful feature since the wrappers and view manager are distributed over the network in our framework.

7.5.4 Impact of Schema Changes. We measure the parallel maintenance performance when schema changes and data updates are mixed. We set up a workload of 400 data updates and two schema changes. The two schema changes are evenly scattered among these 400 data updates. Figure 19 depicts changes of the total maintenance cost in terms of processing time when we vary the number of parallel threads. Similarly to what we have measured for data updates only environments (Figure 16), the total processing time reaches its minimal around thread number 5. The maximal percentage of performance improvement

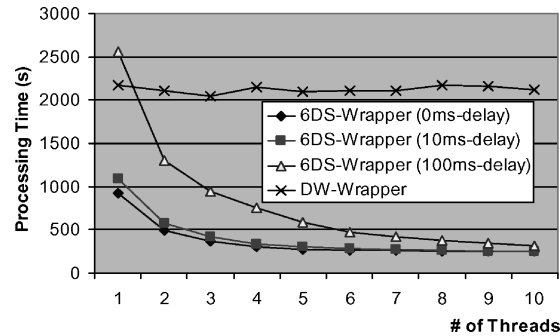


Fig. 20. Impact of wrapper placement.

we measured in this case is around 45%. This is because (1) schema change maintenance is much more time consuming than that of a data update and (2) we cannot fully maintain schema changes in parallel.

7.6 Wrapper Placement Evaluation

As we have discussed in Section 4.1, distributing the wrapper relations as done in our proposed TxnWrap architecture is likely beneficial for maximally exploiting the parallel maintenance capability. We experimentally evaluate this hypothesis in this section. Figure 20 depicts the maintenance performance under two extreme experimental settings. The “DW-Wrapper” case has six wrapper relations placed in the same database engine where the materialized view resides, while the “6DS-Wrapper” case has the wrapper relations placed in six separate database engines. We vary the number of parallel maintenance threads (x -axis). We use again the workload of 500 data updates as in Section 7.5.

From Figure 20, we can see that the maintenance performance of the “DW-Wrapper” approach could hardly be improved by adding more maintenance threads due to the high resource contention when handling a large amount of maintenance queries in parallel (the database server has a single CPU and a single disk in our experimental setting). In comparison, the “6DS-Wrapper” approach achieves better performance when adding more maintenance threads since the workload is distributed over several servers with less resource contention.

However, the “6DS-Wrapper” approach may suffer from high network delay since each maintenance query will be evaluated across the network. Seen from Figure 20, given the network delay factor “ $\ell = 100$ ms” in this experimental setting, the total maintenance time of the “6DS-Wrapper” approach is worse than that of the “DW-Wrapper” in the case of sequential maintenance. We also find that this parallel maintenance strategy makes TxnWrap less sensitive to slow network environments. This is especially useful since the wrappers and view manager are distributed over the network. The reason is that although the network transfer is slow, we can still run many other tasks in parallel. This may improve the overall throughput.

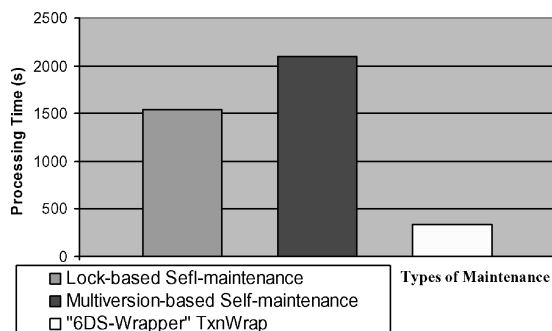


Fig. 21. TxnWrap versus self-maintenance.

7.7 TxnWrap Versus Self-Maintenance

In our system, we implement the self-maintenance approach [Quass et al. 1996] by replicating source relations in the DW server and issuing one combined maintenance query to this local DW server to maintain a source update. We compare TxnWrap to a lock-based self-maintenance solution and to a multiversion-based self-maintenance solution as described in Section 4.2.

Figure 21 compares the total maintenance time between the lock-based self-maintenance approach, the multiversion-based self-maintenance approach (i.e., the parallel “DW-Wrapper” TxnWrap approach) and the parallel “6DS-Wrapper” TxnWrap approach. The experimental setting is the same as in Section 7.6. Seen from Figure 21, the total processing time of the lock-based self-maintenance is less than the multiversion-based self-maintenance approach since the latter has extra version management overhead. Both approaches are limited in terms of parallelism performance due to the high resource contention because of the data all residing on the same machine. In comparison, distributing the replicated data under the TxnWrap approach (labeled as “6DS-Wrapper” in Figure 21) gains much more parallelism performance. While the lock-based approach may also exploit more potential parallelism when the replicated data is distributed, one problem is that the lock-based approach integrates the source updates and the maintenance process into one single transaction. Managing such a global transaction in a distributed context is non-trivial. The potential deadlocks worsen this problem. In summary, we conclude that our proposed parallel maintenance with wrapper data materialization effectively improves the maintenance performance over prior approaches while offering the extra functionality to handle concurrent schema changes.

8. RELATED WORK

Maintaining materialized views under source updates is one of the important issues of data integration [Widom 1995; Zhuge et al. 1995]. Initially, some research studied incremental view maintenance assuming no concurrency. Such algorithms for maintaining a view under source data updates are called *view maintenance algorithms* [Lu et al. 1995; Colby et al. 1996]. There has also been some work on rewriting view definitions under source schema changes [Nica

et al. 1998; Lee et al. 2002], and on adapting the view extent under source schema changes [Gupta et al. 1995; Nica and Rundensteiner 1999].

As first noted in Zhuge et al. [1995], if the data sources are on remote servers from the view site and an asynchronous loosely coupled environment is assumed, then the maintenance process of sending maintenance queries to the remote data sources may experience concurrency conflicts due to the autonomy of sources. Zhuge et al. [1995] introduced a compensation-based algorithm, called *ECA*, for incremental view maintenance of concurrent source *data updates* restricted to a single source. Strobe [Zhuge et al. 1996] handles multiple data sources while still assuming the schema of all sources to be static. SWEEP [Agrawal et al. 1997] applies local compensation of maintenance over distributed sources. Salem et al. [2000] proposed to materialize delta changes of both sources and views with timestamps, thus being able to asynchronously refresh the view extent. They introduced a propagation algorithm that reduces the number of compensation queries. Our earlier work on the DyDa project [Zhang and Rundensteiner 2000] was the first attempt to address mixed workloads of concurrent data updates and schema changes. However, these existing efforts including DyDa are all compensation-based. That is, they all require special algorithms to detect possible anomalies due to concurrent updates of sources and then to correct them thereafter. Furthermore, a formal basis for these algorithms is missing. This makes it more difficult to develop formal proofs of correctness or extensions in terms of services such as recovery or parallelism. Our new approach TxnWrap employs a multiversion solution removing concurrency concerns from view maintenance algorithms altogether, thus no longer requiring any compensation strategies. This is complimentary to any of the previous work mentioned above by isolating concurrency concerns from the maintenance logic.

Beyond view maintenance work, there are techniques in version management and concurrency control in general that relate to our approach. Multiversion concurrency control algorithms [Agrawal and Sengupta 1989; Chan and Gray 1985; Quass and Widom 1997; Mohan et al. 1992] can be categorized into two types, namely, finite version and unrestricted version algorithms. In this article, we demonstrate that unrestricted version algorithms are the most appropriate design choice for materialized view maintenance. In addition, our proposed algorithm can also handle versioned schema metadata. Quass and Widom [1997] proposed to utilize a two-version algorithm to resolve the concurrency between the view maintenance and its read-only sessions. Our work instead focuses on the other end of the spectrum, that is, the concurrency between frequent source updates and the view maintenance. In the former work, the read-only sessions of the materialized view can be scheduled for any version as long as it is consistent, while in our work, the view maintenance reads will ask for a particular version of source data. To keep the view consistent, the desired version should always be available to the view maintenance process.

The TxnWrap versioning approach proposed in this article is somewhat similar to the temporal database model [Zaniolo et al. 1997] in the literature. Our work shows the potential benefits for exploiting this temporal database model for view maintenance. Many existing temporal index schemes [Lomet

and Salzberg 1989] can be applied here to further improve the maintenance query performance.

9. CONCLUSIONS

In this article we take a fresh new look at the concurrency problems of materialized view maintenance. In contrast to previous work, we now present a solution based on transaction theory. While previous work [Zhuge et al. 1995; Salem et al. 2000] focused on specialized compensation algorithms for view maintenance, we instead propose a transactional model that removes concurrency considerations from the maintenance logic. In particular, we encapsulate both the source update transactions and the actual view maintenance process into what we refer to as a *VM_Transaction*. We then propose a multiversion algorithm called *TxnWrap* to achieve the serializability of *VM_Transactions*. This way *TxnWrap* addresses the maintenance anomaly problem not only under concurrent data updates but also under concurrent schema changes. Our *TxnWrap* solution can be plugged into any existing view maintenance system as a middle layer that removes concurrency issues from the view manager. As an added benefit, we find *TxnWrap* very easy to implement. The experimental results reveal that our solution achieves a rather stable maintenance performance. With even simple optimizations of the wrapper, *TxnWrap* already outperforms previous solutions from the literature [Agrawal et al. 1997; Zhang and Rundensteiner 2000] in terms of refresh rates. Another added benefit we find about the transactional approach toward view maintenance is that parallel processing can easily be added. The experimental results confirm the performance achievable by the parallel maintenance. A formal correctness proof for both *TxnWrap* and parallel *TxnWrap* is also provided based on transaction theory [Bernstein et al. 1987].

In future work, we will look more closely at supporting multiple views with interdependencies between views, since our transactional-modeling based approach removes the concurrency issues and simplifies the overall processing logic.

An electronic appendix to this article is available in the ACM Digital Library.

ACKNOWLEDGMENTS

The authors would like to thank Xin Zhang, Jun Chen, Andreas Koeller and all DSRG members at WPI for many valuable suggestions. The authors also thank the anonymous reviewers whose detailed comments helped greatly to improve the article.

REFERENCES

- AGRAWAL, D., ABBADI, A. E., SINGH, A., AND YUREK, T. 1997. Efficient view maintenance at data warehouses. In *Proceedings of SIGMOD*. 417–427.
- AGRAWAL, D. AND SENGUPTA, S. 1989. Modular synchronization in multiversion databases. In *Proceedings of SIGMOD*. 408–417.
- BERNSTEIN, P. A., HADZILACOS, V., AND GOODMAN, N. 1987. *Concurrency Control and Recovery in Database System*. Addison-Wesley, Reading, M.A.

- CHAN, A. AND GRAY, R. 1985. Implementing distributed read-only transactions. *IEEE Trans. on Softw. Eng.* 11, 205–212.
- CHEN, J., CHEN, S., AND RUNDENSTEINER, E. A. 2002. A transactional model for data warehouse maintenance. In *Proceedings of the Conference on Conceptual Modeling*. 247–262.
- CHEN, J., ZHANG, X., CHEN, S., ANDREAS, K., AND RUNDENSTEINER, E. A. 2001. DyDa: Data warehouse maintenance under fully concurrent environments. In *Proceedings of SIGMOD Demo Session*. 619.
- CHEN, S., CHEN, J., ZHANG, X., AND RUNDENSTEINER, E. A. 2004. Detection and correction of conflicting source updates for view maintenance. In *Proceedings of ICDE*. 436–448.
- COLBY, L. S., GRIFFIN, T., LIBKIN, L., MUMICK, I. S., AND TRICKEY, H. 1996. Algorithms for deferred view maintenance. In *Proceedings of SIGMOD*. 469–480.
- GRAY, J. AND REUTER, A. 1992. *Transaction Processing: Concepts and Techniques*. Morgan Kaufmann, San Francisco, CA.
- GUPTA, A. AND MUMICK, I. 1995. Maintenance of materialized views: Problems, techniques, and applications. *IEEE Data Eng. Bull.* 18, 2, 3–19.
- GUPTA, A., MUMICK, I., AND ROSS, K. 1995. Adapting materialized views after redefinition. In *Proceedings of SIGMOD*. 211–222.
- LEE, A. M., NICA, A., AND RUNDENSTEINER, E. A. 2002. The EVE approach: View synchronization in dynamic distributed environments. *IEEE Trans. Knowl. Data Eng.* 14, 5, 931–954.
- LIU, B., CHEN, S., AND RUNDENSTEINER, E. A. 2002a. A transactional approach to parallel data warehouse maintenance. In *Proceedings of DaWaK*. 307–317.
- LIU, B., CHEN, S., AND RUNDENSTEINER, E. A. 2002b. Batch data warehouse maintenance in dynamic environments. In *Proceedings of CIKM*. 68–75.
- LOMET, D. B. AND SALZBERG, B. 1989. Access methods for multiversion data. In *Proceedings of SIGMOD*. 315–324.
- LU, J. J., MOERKOTTE, G., SCHUE, J., AND SUBRAHMANIAN, V. S. 1995. Efficient maintenance of materialized mediated views. In *Proceedings of SIGMOD*. 340–351.
- MADHAVAN, J., BERNSTEIN, P. A., AND RAHM, E. 2001. Generic schema matching with cupid. In *Proceedings of VLDB*. 49–58.
- MARCHE, S. 1993. Measuring the stability of data models. *European J. Inform. Syst.* 2, 1, 37–47.
- MILLER, R. J., HAAS, L. M., AND HERNÁNDEZ, M. A. 2000. Schema mapping as query discovery. In *Proceedings of VLDB*. 77–88.
- MOHAN, C., PIRAHESH, H., AND LORIE, R. 1992. Efficient and flexible methods for transient versioning of records to avoid locking by read-only transactions. In *Proceedings of SIGMOD*. 124–133.
- NICA, A., LEE, A. J., AND RUNDENSTEINER, E. A. 1998. The CVS algorithm for view synchronization in evolvable large-scale information systems. In *Proceedings of EDBT*. 359–373.
- NICA, A. AND RUNDENSTEINER, E. A. 1999. View maintenance after view synchronization. In *Proceedings of the Conference on International Database Engineering and Applications*. 213–215.
- QUASS, D., GUPTA, A., MUMICK, I. S., AND WIDOM, J. 1996. Making view self-maintainable for data warehousing. In *Proceedings of the Conference on Parallel and Distributed Information Systems*. 158–169.
- QUASS, D. AND WIDOM, J. 1997. On-line warehouse view maintenance. In *Proceedings of SIGMOD*. 393–400.
- SALEM, K., BEYER, K. S., COCHRANE, R., AND LINDSAY, B. G. 2000. How to roll a join: Asynchronous incremental view maintenance. In *Proceedings of SIGMOD*. 129–140.
- SJOBERG, D. 1993. Quantifying schema evolution. *Inform. Softw. Tech.* 35, 1, 35–54.
- VARDE, A. S. AND RUNDENSTEINER, E. A. 2002. MEDWRAP: Consistent view maintenance over distributed multi-relation sources. In *Proceedings of DEXA*. 341–350.
- VELEGRAKIS, Y., MILLER, R. J., AND POPA, L. 2003. Mapping adaptation under evolving schemas. In *Proceedings of VLDB*. 584–595.
- WIDOM, J. 1995. Research problems in data warehousing. In *Proceedings of CIKM*. 25–30.
- ZANILOLO, C., CERI, S., FALOURSOS, C., SNODGRASS, R. T., SUBRAHMANIAN, V. S., AND ZICARI, R. 1997. *Advanced Database Systems*. Morgan Kaufmann, San Francisco, CA.
- ZHANG, X., DING, L., AND RUNDENSTEINER, E. A. 2004. Parallel multi-source view maintenance. *VLDB J.* 13, 1, 22–48.

- ZHANG, X. AND RUNDENSTEINER, E. A. 2000. DyDa: Dynamic data warehouse maintenance in a fully concurrent environment. In *Proceedings of DaWaK*. 94–103.
- ZHUGE, Y., GRACÍA-MOLINA, H., HAMMER, J., AND WIDOM, J. 1995. View maintenance in a warehousing environment. In *Proceedings of SIGMOD*. 316–327.
- ZHUGE, Y., GRACÍA-MOLINA, H., AND WIENER, J. L. 1996. The strobe algorithms for multi-source warehouse consistency. In *Proceedings of the Conference on Parallel and Distributed Information Systems*. 146–157.

Received August 2003; revised May 2004; accepted August 2004