

Extending Practical Pre-Aggregation in On-Line Analytical Processing

Torben Bach Pedersen
Kommunedata
DK-8200 Århus N, Denmark
tbp@kmd.dk

Christian S. Jensen Curtis E. Dyreson
Department of Computer Science
Aalborg University, DK-9220 Aalborg Øst, Denmark
{csj,curtis}@cs.auc.dk

Abstract

On-Line Analytical Processing (OLAP) based on a dimensional view of data is being used increasingly for the purpose of analyzing very large amounts of data. To improve query performance, modern OLAP systems use a technique known as *practical pre-aggregation*, where *select* combinations of aggregate queries are materialized and re-used to compute other aggregates; full pre-aggregation, where all combinations of aggregates are materialized, is infeasible. However, this re-use of aggregates is contingent on the dimension hierarchies and the relationships between facts and dimensions satisfying stringent constraints, which severely limits the scope of practical pre-aggregation. This paper significantly extends the scope of practical pre-aggregation to cover a much wider range of realistic situations. Specifically, algorithms are given that transform “irregular” dimension hierarchies and fact-dimension relationships, which often occur in real-world OLAP applications, into well-behaved structures that, when used by existing OLAP systems, enable practical pre-aggregation. The algorithms have low computational complexity and may be applied incrementally to reduce the cost of updating OLAP structures.

1 Introduction

On-Line Analytical Processing (OLAP) systems aim to ease the process of extracting useful information from large

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or special permission from the Endowment.

**Proceedings of the 25th VLDB Conference,
Edinburgh, Scotland, 1999.**

amounts of detailed transactional data and have gained widespread acceptance in traditional business applications as well as in new applications such as health care. These systems generally offer a dimensional view of data, in which measured values, termed facts, are characterized by descriptive values, drawn from a number of dimensions; and the values of a dimension are typically organized in a containment-type hierarchy. A prototypical query applies an aggregate function, such as average, to the facts characterized by specific values from the dimensions.

Fast response times are required from these systems, even for queries that aggregate large amounts of data. The perhaps most central technique used for meeting this requirement is termed *pre-aggregation*, where the results of aggregate queries are pre-computed and stored, i.e., materialized, for later use during query processing. Pre-aggregation has attracted substantial attention in the research community, where it has been investigated how to optimally use pre-aggregated data for query optimization [6, 3] and how to maintain the pre-aggregated data when base data is updated [18, 23]. Further, the latest versions of commercial RDBMS products offer both query optimization based on pre-computed aggregates and automatic maintenance of the stored aggregates when base data is updated [28].

The fastest response times may be achieved when materializing aggregate results corresponding to all combinations of dimension values across all dimensions, termed *full* (or *eager*) pre-aggregation. However, the required storage space grows rapidly, to quickly become prohibitive, as the complexity of the application increases. This phenomenon is called *data explosion* [25, 20] and occurs because the number of possible aggregation combinations grows rapidly when the number of dimensions increase, while the sparseness of the multidimensional space decreases in higher dimension levels, meaning that aggregates at higher levels take up nearly as much space as lower-level aggregates. In some commercial applications, full pre-aggregation takes up as much as 200 times the space of the raw data [20]. Another problem with full pre-aggregation is that it takes too long to update the materialized aggregates

when base data changes.

With the goal of avoiding data explosion, research has focused on how to select the best subset of aggregation levels given space constraints [10, 8, 30, 1, 26] or maintenance time constraints [9], or the best combination of aggregate data and indices [7]. This approach is commonly referred to as *practical* (or partial or semi-eager [4, 10, 27]) pre-aggregation. Commercial OLAP systems now also exist that employ practical pre-aggregation, e.g., Informix MetaCube [12] and Microsoft Decision Support Services (Plato) [17].

The premise underlying the applicability of practical pre-aggregation is that lower-level aggregates can be *re-used* to compute higher-level aggregates, known as summarizability [15]. Summarizability occurs when the mappings in the dimension hierarchies are *onto* (all paths from the root to a leaf in the hierarchy have equal lengths), *covering* (only immediate parent and child values can be related), and *strict* (each child in a hierarchy has only one parent); and when also the relationships between facts and dimensions are many-to-one and facts are always mapped to the lowest levels in the dimensions [15]. However, the data encountered in many real-world applications fail to comply with this rigid regime. This motivates the search for techniques that allow practical pre-aggregation to be used for a wider range of applications, the focus of this paper.

Specifically, this paper leverages research such as that cited above. It does so by showing how to transform dimension hierarchies to obtain summarizability, and by showing how to integrate the transformed hierarchies into current systems, transparently to the user, so that standard OLAP technology is re-used. Specifically, algorithms are presented that automatically transform dimension hierarchies to achieve summarizability for hierarchies that are non-onto, non-covering, and non-strict. The algorithms have low computational complexity and are thus applicable to even very large databases. The algorithms can also be used to contend with non-summarizable relationships between facts and dimensions and may be modified to accommodate incremental computation, thus minimizing the maintenance cost associated with base-data updates [22].

To our knowledge, this work is the first to present algorithms to automatically achieve summarizability for non-covering and non-onto hierarchies. The research reported here is also the first to demonstrate techniques and algorithms for achieving summarizability in non-strict hierarchies. The integration of the techniques into current systems, transparently to the user, we believe is a novel feature. The only past research on the topic has been on how to manually, and not transparently to the user, achieve summarizability for non-covering hierarchies [24].

The next section presents a real-world clinical case study that exemplifies the non-summarizable properties of real-world applications. Section 3 proceeds to define the aspects of a multidimensional data model necessary for describing the new techniques, and defines also important properties related to summarizability. Section 4 presents algorithms

that transform dimension hierarchies to achieve summarizability and discusses how the algorithms may be applied to non-summarizable relationships between facts and dimensions. Section 5 demonstrates how the techniques may be integrated into current systems, and Section 6 summarizes and points to topics for future research.

2 Motivation—A Case Study

The case study concerns patients in a hospital, their associated diagnoses, and their places of residence. The data analysis goal is to investigate whether some diagnoses occur more often in some areas than in others, in which case environmental or lifestyle factors might be contributing to the disease pattern. An ER diagram illustrating the underlying data is seen in Figure 1.

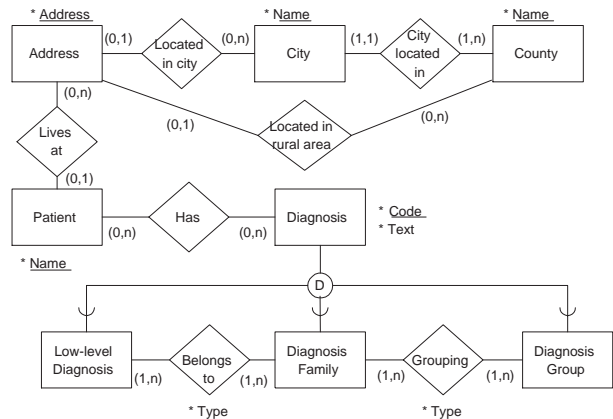


Figure 1: ER Schema of Case Study

The most important entities are the *patients*, for which we record the name. We always want to count the number of patients, grouped by some properties of the patients. Thus, in multidimensional terms, the patients are the *facts*, and the other, describing, entities constitute the *dimensions*.

Each patient has a number of *diagnoses*, yielding to a *many-to-many* relationship between facts and the diagnosis dimension. When registering diagnoses of patients, physicians use different levels of granularity, ranging from very precise diagnoses, e.g., “Insulin dependent diabetes during pregnancy,” to more imprecise diagnoses, e.g., “Diabetes,” which cover wider ranges of patient conditions. To model this, the relationship from Patient to diagnoses is to the supertype “Diagnosis,” which then has three subtypes, corresponding to different levels of granularity, the *low-level diagnosis*, the *diagnosis family*, and the *diagnosis group*. Examples of these are “Insulin dependent diabetes during pregnancy,” “Insulin dependent diabetes,” and “Diabetes,” respectively. The higher-level diagnoses are both (imprecise) diagnoses in their own right, but also serve as groups of lower-level diagnoses, i.e., a diagnosis family consists of 5–20 related low-level diagnoses and a diagnosis group consists of 5–20 related diagnosis families.

Each diagnosis has an alphanumeric code and a descriptive text, which are specified by some standard, here the World Health Organization’s International Classification of

Diseases (ICD-10) [29], or by the physicians themselves. In the hierarchy determined by the WHO, one lower-level item belongs to exactly one higher-level item, leading to a *strict, covering* hierarchy. In the user-defined hierarchy, one lower-level item may belong to zero or more higher-level items, e.g., the family “Diabetes during pregnancy” may belong to both the “Diabetes” and “Other pregnancy related diseases” groups. Thus, the user-defined hierarchy is *non-strict* and *non-covering*.

We also record the addresses of the patients. If the address is located in a city, we record the *city*; otherwise, if the address is in a rural area, we record the *county* in which the address is located. A city is located in exactly one county. As not all addresses are in cities, we cannot find all addresses in a county by going through the “City located in” relationship. Thus, the mapping from addresses to cities is *non-covering* w.r.t. addresses. For cities and counties, we record the name. Not all counties have cities in them, so the mapping from cities to counties is *into* rather than *onto*.

In order to exemplify the data, we assume a standard mapping of the ER diagram to relational tables and the use of surrogate keys, named *ID*, with globally unique values. The three subtypes of the Diagnosis type are mapped to a common Diagnosis table, and because of this, the “Belongs to” and “Grouping” relationships are mapped to a common “Grouping” table. The resulting tables with sample data are shown in Table 1 and will be used in examples throughout the paper.

If we apply pre-aggregation to the data from the case study, several problems occur. For example, if the counts of patients by City are pre-computed and we use these for computing the numbers of patients by county, an incorrect result will occur. In the data, the addresses “123 Rural Road” and “1 Sandy Dunes” (one of them is the address of a patient) are not in any city, making the mapping from City to County not *covering* w.r.t. addresses.

Next, if the counts of patients by Low-Level Diagnosis are pre-computed and we use these for computing the total count of patients, an incorrect result again ensues. First, patients only with lung cancer are not counted, as lung cancer is not present at the level of Low-Level Diagnosis; the mapping from Low-Level Diagnosis to Diagnosis Family is *into*. Second, patients such as “Jim Doe” only have higher-level diagnoses and will not be counted; the fact-to-dimension mapping has *varying granularity*. Third, patients such as “Jane Doe” have several diagnoses and will be counted several times; the relationship between facts and dimensions is *many-to-many*. Fourth, Low-Level diagnoses such as “Insulin dependent diabetes during pregnancy” are part of several diagnosis families, which may also lead to “double” counting when computing higher-level counts; the dimension hierarchy is *non-strict*.

These problems yield “non-summarizable” dimension hierarchies that severely limit the applicability of practical pre-aggregation, leaving only full pre-aggregation, which requires huge amounts of storage, or no pre-aggregation, which results in long response times for queries.

ID	Name
1	John Doe
2	Jane Doe
3	Jim Doe

Patient

PatientID	DiagID	Type
1	9	Primary
2	5	Secondary
2	9	Primary
3	11	Primary

Has

PatientID	AddressID
1	50
2	51
3	52

LivesAt

ID	Address
50	21 Central Street
51	34 Main Street
52	123 Rural Road
53	1 Sandy Dunes

Address

ID	Code	Text	Type
4	O24	Diabetes during pregnancy	DF
5	O24.0	Ins. dep. diab. during pregn.	LLD
6	O24.1	Non ins. dep. diab. during pregn.	LLD
9	E10	Insulin dependent diabetes	DF
10	E11	Non insulin dependent diabetes	DF
11	E1	Diabetes	DG
12	O2	Other pregnancy related diseases	DG
13	A1	Cancer	DG
14	A11	Lung cancer	DF

Diagnosis

ParentID	ChildID	Type
4	5	WHO
4	6	WHO
9	5	User
10	6	User
11	9	WHO
11	10	WHO
12	4	WHO
13	14	WHO

Grouping

CityID	CountyID
20	30
21	31

CityLocatedIn

ID	Name
20	Sydney
21	Melbourne

City

AddressID	CityID
50	20
51	21

LocatedInCity

ID	Name
30	Sydney
31	Melbourne
32	Outback

County

ID	Name
52	31
53	32

LocatedInRuralArea

Table 1: Tables for the Case Study

The properties described above are found in many other real-world applications. Many-to-many relationships between facts and dimensions occur between bank customers and accounts, between companies and Standard Industry Classifications (SICs), and between students and departments [14, 15]. Non-strict dimension hierarchies occur from cities to states in a Geography dimension [24] and from weeks to months in a Time dimension. In addition, hierarchies where the change over time is captured are generally non-strict. The mapping from holidays to weeks as well as organization hierarchies of varying depth [11] offer examples of “into” mappings. Non-covering relationships exist for days-holidays-weeks and for counties-cities-states, as well as in organization hierarchies [11].

Even though many real-world cases possess the properties described above, current techniques for practical pre-aggregation require that facts are in a many-to-one relationships to dimensions and that all hierarchies are strict, onto,

and covering.

3 Data Model Context and Concepts

This section describes the aspects of a multidimensional data model that extend practical pre-aggregation. The full model is described elsewhere [21]. Next, the data model context is exploited for defining properties of hierarchies relevant to the techniques.

The particular data model has been chosen over other multidimensional data models because it quite naturally captures the data described in the case study and because it includes explicit concepts of dimensions and dimension hierarchies, which is very important for clearly presenting the techniques. However, the techniques are also applicable to other multidimensional or statistical data models, as will be discussed in Section 5.

3.1 A Concrete Data Model Context

For each part of the model, we define the *intension* and the *extension*, and we give an illustrating example.

An *n-dimensional fact schema* is a two-tuple $\mathcal{S} = (\mathcal{F}, \mathcal{D})$, where \mathcal{F} is a *fact type* and $\mathcal{D} = \{\mathcal{T}_i, i = 1, \dots, n\}$ is its corresponding *dimension types*.

Example 1 In the case study from Section 2, *Patient* is the fact type, and *Diagnosis*, *Residence*, and *Name* are the dimension types. The intuition is that *everything* that characterizes the fact type is considered to be *dimensional*.

A dimension type \mathcal{T} is a four-tuple $(\mathcal{C}, \leq_{\mathcal{T}}, \top_{\mathcal{T}}, \perp_{\mathcal{T}})$, where $\mathcal{C} = \{C_j, j = 1, \dots, k\}$ are the *category types* of \mathcal{T} , $\leq_{\mathcal{T}}$ is a partial order on the C_j 's, with $\top_{\mathcal{T}} \in \mathcal{C}$ and $\perp_{\mathcal{T}} \in \mathcal{C}$ being the top and bottom element of the ordering, respectively. Thus, the category types form a lattice. The intuition is that one category type is “greater than” another category type if members of the former’s extension logically contain members of the latter’s extension, i.e., they have a larger value size. The top element of the ordering corresponds to the largest possible value size, that is, there is only one value in it’s extension, logically containing all other values.

We say that C_j is a *category type of \mathcal{T}* , written $C_j \in \mathcal{T}$, if $C_j \in \mathcal{C}$.

Example 2 Low-level diagnoses are contained in diagnosis families, which are contained in diagnosis groups. Thus, the *Diagnosis* dimension type has the following order on its category types: $\perp_{Diagnosis} = Low\text{-level } Diagnosis < Diagnosis\ Family < Diagnosis\ Group < \top_{Diagnosis}$. Other examples of category types are *Address*, *City*, and *County*. Figure 2, to be discussed in detail later, illustrates the dimension types of the case study.

A *category C_j of type \mathcal{T}* is a set of *dimension values e* . A *dimension D of type $\mathcal{T} = (\{C_j\}, \leq_{\mathcal{T}}, \top_{\mathcal{T}}, \perp_{\mathcal{T}})$* is a two-tuple $D = (C, \leq)$, where $C = \{C_j\}$ is a set of categories C_j such that $Type(C_j) = C_j$ and \leq is a partial order on $\cup_j C_j$, the union of all dimension values in the individual categories. We assume a function $Pred : C \mapsto \mathcal{P}^C$ that gives the set of immediate predecessors of a category C_j .

Similarly, we assume a function $Desc : C \mapsto \mathcal{P}^C$ that gives the set of immediate descendants of a category C_j . For both $Pred$ and $Desc$, we “count” from the category $\top_{\mathcal{T}}$ (of type $\top_{\mathcal{T}}$), so that category $\top_{\mathcal{T}}$ is the ultimate predecessor and category $\perp_{\mathcal{T}}$ (of type $\perp_{\mathcal{T}}$) is the ultimate descendant.

The definition of the partial order is: given two values e_1, e_2 then $e_1 \leq e_2$ if e_1 is logically contained in e_2 . We say that C_j is a category of D , written $C_j \in D$, if $C_j \in C$. For a dimension value e , we say that e is a dimensional value of D , written $e \in D$, if $e \in \cup_j C_j$.

The category of type $\perp_{\mathcal{T}}$ in dimension of type \mathcal{T} contains the values with the smallest value size. The category with the largest value size, of type $\top_{\mathcal{T}}$, contains exactly one value, denoted \top . For all values e of the dimension D , $e \leq \top$. Value \top is similar to the *ALL* construct of Gray et al. [5]. When clear from the context, we refer to a category of type $\top_{\mathcal{T}}$ simply as a \top category, not to be confused with the \top dimension value.

Example 3 In our *Diagnosis* dimension we have the following categories, named by their type. The numbers in parentheses are the ID values from the Diagnosis table in Table 1. *Low-level Diagnosis* = {“Insulin dependent diabetes during pregnancy” (5), “Non insulin dependent diabetes during pregnancy” (6)}, *Diagnosis Family* = {“Diabetes during pregnancy” (4), “Insulin dependent diabetes” (9), “Non insulin dependent diabetes” (10), “Lung cancer” (14)}, *Diagnosis Group* = {“Diabetes” (11), “Other pregnancy related diseases” (12), “Cancer” (13)}, and $\top_{Diagnosis} = \{\top\}$. We have that $Pred(Low\text{-level } Diagnosis) = \{Diagnosis\ Family\}$. The partial order \leq is obtained by combining the WHO and user-defined hierarchies, as given by the Grouping table in Table 1. Additionally, the top value \top is greater than, i.e., logically contains, all the other diagnosis values.

Let F be a set of facts, and $D = (C = \{C_j\}, \leq)$ a dimension. A *fact-dimension relation* between F and D is a set $R = \{(f, e)\}$, where $f \in F$ and $e \in \cup_j C_j$. Thus R links facts to dimension values. We say that fact f is *characterized by* dimension value e , written $f \rightsquigarrow e$, if $\exists e_1 \in D ((f, e_1) \in R \wedge e_1 \leq e)$. We require that $\forall f \in F (\exists e \in \cup_j C_j ((f, e) \in R))$; thus, all fact maps to at least one dimension value in every dimension. The \top value is used to represent an unknown or missing value, as \top logically contains all dimension values, and so a fact f is mapped to \top if it cannot be characterized within the particular dimension.

Example 4 The fact-dimension relation R links patient facts to diagnosis dimension values as given by the Has table from the case study, so that $R = \{(\text{“John Doe” (1), “Insulin dependent diabetes” (9)}, (\text{“Jane Doe” (2), “Insulin dependent diabetes during pregnancy” (5)}, (\text{“Jane Doe” (2), “Insulin dependent diabetes” (9)}, (\text{“Jim Doe” (3), “Diabetes” (11)}))\}$. Note that facts may be related to values in higher-level categories. We do not require that e belongs to $\perp_{Diagnosis}$. For example, the fact “John Doe” (1) is related

to the diagnosis “Insulin dependent diabetes” (5), which belongs to the *Diagnosis Family* category. This feature will be used later to explicitly capture the different granularities in the data. If no diagnosis was known for patient “John Doe” (1), we would have added the pair (“John Doe” (1), \top) to R .

A *multidimensional object* (MO) is a four-tuple $M = (\mathcal{S}, F, D, R)$, where $\mathcal{S} = (\mathcal{F}, \mathcal{D} = \{\mathcal{T}_i\})$ is the fact schema, $F = \{f\}$ is a set of *facts* f where $Type(f) = \mathcal{F}$, $D = \{D_i, i = 1, \dots, n\}$ is a set of *dimensions* where $Type(D_i) = \mathcal{T}_i$, and $R = \{R_i, i = 1, \dots, n\}$ is a set of fact-dimension relations, such that $\forall i((f, e) \in R_i \Rightarrow f \in F \wedge \exists C_j \in D_i(e \in C_j))$.

Example 5 For the case study, we get a three-dimensional MO $M = (\mathcal{S}, F, D, R)$, where $\mathcal{S} = (Patient, \{Diagnosis, Name, Residence\})$ and $F = \{“John Doe” (1), “Jane Doe” (2), “Jim Doe” (3)\}$. The definition of the diagnosis dimension and its corresponding fact-dimension relation was given in the previous examples. The Residence dimension has the categories *Address* ($= \perp_{Residence}$), *City*, *County*, and $\top_{Residence}$. The values of the categories are given by the corresponding tables in Table 1. The partial order is given by the relationship tables. Additionally, the only value in the $\top_{Residence}$ category is \top , which logically contains all the other values in the Residence dimension. The Name dimension is simple, i.e., it just has a *Name* category ($= \perp_{Name}$) and a \top category. We will refer to this MO as the “Patient” MO. A graphical illustration of its schema is seen in Figure 2. Because some addresses map directly to counties, County is an immediate predecessor of Address.

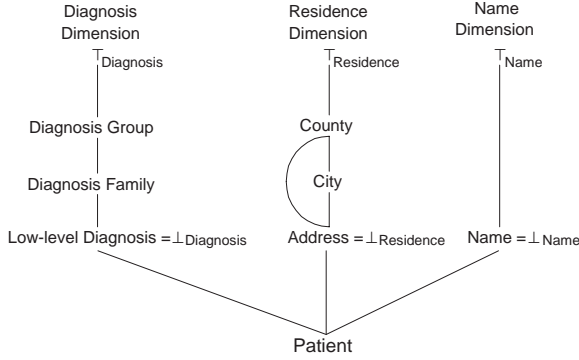


Figure 2: Schema of the Case Study

The facts in an MO are objects with *value-independent identity*. We can test facts for equality, but do not assume an ordering on the facts. The combination of dimension values that characterize the set of facts of an MO do *not* constitute a “key” for the fact set. Thus, several facts may be characterized by the same combination of dimension values. But, the facts of an MO is a set, and an MO does not have duplicate facts.

3.2 Hierarchy Properties

In this section important properties of MOs are defined, which will be used in the following sections to state precisely what problems the proposed algorithms solve. The

first important concept is *summarizability*, which intuitively means that higher-level aggregates may be obtained directly from lower-level aggregates.

Definition 1 Given a type T , a set $S = \{S_j, j = 1, \dots, k\}$, where $S_j \in 2^T$, and a function $g : 2^T \mapsto T$, we say that g is *summarizable* for S if $g(\{\{g(S_1), \dots, g(S_k)\}\}) = g(S_1 \cup \dots \cup S_k)$. The argument on the left-hand side of the equation is a multiset, i.e., the same value may occur multiple times.

Summarizability is important as it is a condition for the flexible use of pre-computed aggregates. Without summarizability, lower-level results generally cannot be directly combined into higher-level results. This means that we cannot choose to pre-compute only a relevant selection of the possible aggregates and then use these to (efficiently) compute higher-level aggregates on-the-fly. Instead, we have to pre-compute the all the aggregate results of queries that we need fast answers to, while other aggregates must be computed from the base data. Space and time constraints can be prohibitive for pre-computing all results, while computing aggregates from base data is often inefficient.

It has been shown that summarizability is equivalent to the aggregate function (g) being *distributive*, all paths being *strict*, and the mappings between dimension values in the hierarchies being *covering* and *onto* [15]. These concepts are formally defined below. The definitions assume a dimension $D = (C, \leq)$ and an MO $M = (\mathcal{S}, F, D, R)$.

Definition 2 Given two categories, C_1, C_2 such that $C_2 \in Pred(C_1)$, we say that the mapping from C_1 to C_2 is *onto* iff $\forall e_2 \in C_2(\exists e_1 \in C_1(e_1 \leq e_2))$. Otherwise, it is *into*. If all mappings in a dimension are onto, we say that the dimension hierarchy is *onto*.

Mappings that are into typically occur when the dimension hierarchy has varying height. In the case study, there is no low-level cancer diagnosis, meaning that some parts of the hierarchy have height 2, while most have height 3. It is thus not possible to use aggregates at the Low-level Diagnosis level for computing aggregates at the two higher levels. Mappings that are into also occur often in organization hierarchies.

Definition 3 Given three categories, C_1, C_2 , and C_3 such that $Type(C_1) < Type(C_2) < Type(C_3)$, we say that the mapping from C_2 to C_3 is *covering with respect to C_1* iff $\forall e_1 \in C_1(\forall e_3 \in C_3(e_1 \leq e_3 \Rightarrow \exists e_2 \in C_2(e_1 \leq e_2 \wedge e_2 \leq e_3)))$. Otherwise, it is *non-covering with respect to C_1* . If all mappings in a dimension are covering w.r.t. any category, we say that the dimension hierarchy is *covering*.

Non-covering mappings occur when some of the links between dimension values skip one or more levels and map directly to a value located higher up in the hierarchy. In the case study, this happens for the “1 Sandy Dunes” address, which maps directly to “Outback County” (there are no cities in Outback County). Thus, we cannot use aggregates at the City level for computing aggregates at the County level.

Definition 4 Given an MO $M = (S, F, D, R)$, and two categories C_1 and C_2 that belong to the same dimension $D_i \in D$ such that $Type(C_1) < Type(C_2)$, we say that the mapping from C_1 to C_2 is *covering with respect to F*, the set of facts, iff $\forall f \in F (\forall e_2 \in C_2 (f \rightsquigarrow_i e_2 \Rightarrow \exists e_1 \in C_1 (f \rightsquigarrow_i e_1 \wedge e_1 \leq_i e_2)))$.

This case is similar to the one above, but now it is the mappings between facts and dimension values that may skip one or more levels and map facts directly to dimension values in categories above the bottom level. In the case study, the patients can map to diagnoses anywhere in the Diagnosis dimension, not just to Low-level Diagnoses. This means that we cannot use aggregates at the Low-level Diagnosis Level for computing aggregates higher up in the hierarchy.

Definition 5 Given two categories, C_1 and C_2 such that $C_2 \in Pred(C_1)$, we say that the mapping from C_1 to C_2 is *strict* iff $\forall e_1 \in C_1 (\forall e_2, e_3 \in C_2 (e_1 \leq e_2 \wedge e_1 \leq e_3 \Rightarrow e_2 = e_3))$. Otherwise, it is *non-strict*. The hierarchy in dimension D is *strict* if all mappings in it are strict; otherwise, it is *non-strict*. Given an MO $M = (S, F, D, R)$ and a category C_j in some dimension $D_i \in D$, we say that there is a *strict path* from the set of facts F to C_j iff $\forall f \in F (f \rightsquigarrow_i e_1 \wedge f \rightsquigarrow_i e_2 \wedge e_1 \in C_j \wedge e_2 \in C_j \Rightarrow e_1 = e_2)$. (Note that the paths to the $\top_{\mathcal{T}}$ categories are always strict.)

Non-strict hierarchies occur when a dimension value has multiple parents. This occurs in the Diagnosis dimension in the case study where the “Insulin dependent diabetes during pregnancy” low-level diagnosis is part of both the “Insulin Dependent Diabetes” and the “Diabetes during pregnancy” diagnosis families, which in turn both are part of the “Diabetes” diagnosis group. This means that we cannot use aggregates at the Diagnosis Family level to compute aggregates at the Diagnosis Group level, since data for “Insulin dependent diabetes during pregnancy” would then be counted twice.

Definition 6 If the dimension hierarchy for a dimension D is *onto*, *covering*, and *strict*, we say that D is *normalized*. Otherwise, it is *un-normalized*. For an MO $M = (S, D, F, R)$, if all dimensions $D_i \in D$ are normalized and $\forall R_i \in R ((f, e) \in R_i \Rightarrow e \in \perp_D)$ (, i.e., all facts map to dimension values in the bottom category), we say that M is *normalized*. Otherwise, it is *un-normalized*.

For normalized hierarchies and MOs, all mappings are summarizable, meaning that we can pre-aggregate values at any combination of dimension levels and safely re-use the pre-aggregated values to compute higher-level aggregate results. Thus, we want to normalize the dimension hierarchies and MOs for which we want to apply practical pre-aggregation.

4 Transformation Techniques

This section describes how dimensions can be transformed to achieve summarizability. Transforming dimensions on their own, separately from the facts, results in well-behaved

dimensions that can be applied in a number of different systems or sold to third-party users. In addition, it is discussed how to apply the presented algorithms to non-summarizable fact-dimension relations. The transformation of the dimension hierarchies is a three-step operation. First, all mappings are transformed to be *covering*, by introducing extra “intermediate” values. Second, all mappings are transformed to be *onto*, by introducing “placeholder” values at lower levels for values without any children. Third, mappings are made *strict*, by “fusing” values together. The three steps are treated in separate sections. None of the algorithms introduce any non-summarizable properties, so applying each once is sufficient.

In general, the algorithms take as input a set of tables R_{C_1, C_2} that specifies the mapping from dimension values in category C_1 to values in category C_2 . The input needs not contain all pairs of ancestors and descendants—only direct parent-child relationships are required. If there are non-covering mappings in the hierarchy, we have categories C, P, H such that $\{P, H\} \subseteq Pred(C)$ and $Type(P) < Type(H)$. In this case, the input must also contain $R_{P, H}$ tables that map P values to H values.

4.1 Non-Covering Hierarchies

The first algorithm renders all mappings in a dimension hierarchy covering w.r.t. any category. When a dimension value is mapped *directly* to another value in a category higher than the one immediately above it in the hierarchy, a new intermediate value is inserted into the category immediately above, and the two original dimension values are linked to this new value, rather than to each other.

Example 6 In the hierarchy for the Residence dimension, two links go from Address directly to County. The address “123 Rural Road” (52) is in “Melbourne County” (31), but not in a city, and the address “1 Sandy Dunes” (53) is in “Outback County” (32), which does *not* have any cities at all. The algorithm inserts two new dimension values in the City category, **C31** and **C32**, which represent Melbourne and Outback county, respectively, and links them to their respective counties. The addresses “123 Rural Road” and “1 Sandy Dunes” are then linked to **C31** and **C32**, respectively. This occurs in the first call of procedure MakeCovering (on the Address category; the procedure is given below). When MakeCovering is called recursively on the City, County, and \top categories, nothing happens, as all mappings are already covering. The transformation is illustrated graphically in Figure 3. The dotted lines show the “problematic” links, and the bold-face values and thick lines show the new dimension values and links.

In the algorithm, C is a *child* category, P is a *parent* category, H is a “*higher*” category, L are the non-covering links from C to H , and N are the “higher” dimension values in L . The \bowtie operator denotes natural join. The algorithm works as follows. Given the argument category C (initially the bottom category) in line 1, the algorithm goes through all C ’s parent categories P (line 2). For each parent category P , it looks for predecessor categories H of C




Figure 3: Transformations by MakeCovering

that are “higher” in the hierarchy than P (line 4). If such an H exists, there might be links in the mapping from C to H that are not available by going through P . Line 6 finds these “non-covered” links, L , in the mapping from C to H by “subtracting” the links that *are* available by going through P from *all* the links in the mapping from C to H . Line 7 uses L to find the dimension values N in H that participate in the “non-covered” mappings. For each value in N , line 8 inserts a corresponding marked value into P ; these marked values represent the N values in P . The marked values in P are then linked to the original values in H (line 9) and C (line 10). Line 12 contains a recursive call to the algorithm on P , thus fixing mappings higher up in the hierarchy. The algorithm terminates when it reaches the \top category, which has no predecessors.

- (1) **procedure** MakeCovering(C)
- (2) **for** each $P \in \text{Pred}(C)$ **do**
- (3) **begin**
- (4) **for** each $H \in \text{Pred}(C)$
where $\text{Type}(H) > \text{Type}(P)$ **do**
- (5) **begin**
- (6) $L \leftarrow R_{C,H} \setminus \Pi_{C,H}(R_{C,P} \bowtie R_{P,H})$
- (7) $N \leftarrow \Pi_H(L)$
- (8) $P \leftarrow P \cup \{\text{Mark}(h) \mid h \in N\}$
- (9) $R_{P,H} \leftarrow R_{P,H} \cup \{(\text{Mark}(h), h) \mid h \in N\}$
- (10) $R_{C,P} \leftarrow R_{C,P} \cup \{(c, \text{Mark}(h)) \mid (c, h) \in L\}$
- (11) **end**
- (12) MakeCovering(P)
- (13) **end**
- (14) **end**

All steps in the algorithm are expressed using standard relational algebra operators. The *general* worst-case complexity of join is $\mathcal{O}(n^2)$, where n is the size of the input. However, because the input to the algorithm are hierarchy definitions, the complexity of the join in the algorithm will only be $\mathcal{O}(n \log n)$. Thus, all the operators used can be evaluated in time $\mathcal{O}(n \log n)$, where n is the size of the input. The *Mark* operation can be performed in $\mathcal{O}(1)$ time. The inner loop of the algorithm is evaluated at most once for each link between categories, i.e., at most $k^2/2$ times, where k is the number of categories (if all categories are directly linked to all others). Thus, the overall big- \mathcal{O} complexity of the algorithm is $\mathcal{O}(k^2 n \log n)$, where k is the number of categories and n is the size of the largest participating R_{C_1, C_2} relation. The worst-case complexity will

not apply very often; in most cases, the inner loop will only be evaluated at most k times.

The algorithm inserts new values into the P category to ensure that the mappings from P to higher categories are summarizable, i.e., that pre-aggregated results for P can be directly combined into higher-level aggregate results. The new values in P mean that the cost of materializing aggregate results for P is higher for the transformed hierarchy than for the original. However, if the hierarchy was not transformed to achieve summarizability, we would have to materialize aggregates for G , and perhaps also for higher level categories. At most one new value is inserted into P for every value in G , meaning that the extra cost of materializing results for P is never greater than the cost of the (otherwise necessary) materialization of results for G . This is a very unlikely worst-case scenario—in the most common cases, the extra cost for P will be much lower than the the cost of materializing results for G , and the savings will be even greater because materialization of results for higher-level categories may also be avoided.

The correctness argument for the algorithm has two aspects. First, the mappings in the hierarchy should be *covering* upon termination. Second, the algorithm should only make transformations that are semantically correct, i.e., we should get the same results when computing results with the new hierarchy as with the old. The correctness follows from Theorem 1 and 2, below. As new values are inserted in the P category, we will get aggregate values for both the new and the original values when “grouping” by P . Results for the original values will be the same as before, so the original result set is a *subset* of the result set obtained with the transformed hierarchy.

Theorem 1 Algorithm MakeCovering terminates and the hierarchy for the resulting dimension D' is covering.

Proof: By induction in the height of the lattice [22].

Theorem 2 Given dimensions D and D' such that D' is the result of running MakeCovering on D , an aggregate result obtained using D is a subset of the result obtained using D' .

Proof: Follows easily from Lemma 1 [22].

Lemma 1 For the dimension $D' = (C', \leq')$ resulting from applying algorithm MakeCovering to dimension $D = (C, \leq)$, the following holds: $\forall e_1, e_2 \in D (e_1 \leq' e_2 \Leftrightarrow e_1 \leq e_2)$ (there is a path between any two original dimension values in the new dimension hierarchy iff there was a path between them in the original hierarchy).

Proof: By induction in the height of the lattice [22].

We see that the original values in the hierarchy are still linked to exactly the same original values as before, as stated by Lemma 1, although new values might have been inserted in-between the original values. Thus, when evaluating a query using the transformed hierarchy, the results for the original values will be the same as when using the original hierarchy.

Assuming only the original result set is desired, results for the new values must be excluded, which is easy to accomplish. The new, “internal” values are marked with “mark = internal”, whereas the original values have “mark = original”. In order to exclude the new, internal values from the result set, the equivalent of an SQL HAVING clause condition of “mark = original” is introduced into the original query.

4.2 Non-Onto Hierarchies

The second algorithm renders all mappings in hierarchies onto, i.e., all dimension values in non-bottom categories have children. This is ensured by inserting placeholder values in lower categories to represent the childless values. These new values are marked with the original values, making it possible to map facts to the new placeholder values instead of to the original values. This makes it possible to only map facts to the bottom category.

Example 7 In the Diagnosis dimension, the “Lung cancer” diagnosis family (ID = 14) has no children. When the algorithm (given shortly) reaches the Diagnosis Family category, it inserts a placeholder value (**L14**) into the Low-level Diagnosis category, representing the “Lung cancer” diagnosis, and links it to the original value. Facts mapped to the “Lung cancer” value may then instead be mapped to the new placeholder value. Using this technique we can ensure that facts are mapped only to the Low-level Diagnosis Category. A graphical illustration of the transformation is given in Figure 4. The bold-faced value **L14** is the new value inserted, and the thick line between 14 and **L14** is the new link inserted.




Figure 4: Transformations by MakeOnto

In the algorithm that follows, P is a *parent* category, C is a *child* category, and N holds the parent values with *no* children. The algorithm works as follows. Given a category P (initially the \top category) in line 1, the algorithm goes through all categories C that are (immediate) descendants of P (line 2). For each C , line 4 finds the values N in P that have *no* children in C , by “subtracting” the values *with* children in C from the values in P . For each “childless” value in N , lines 5 and 6, respectively, insert into C a placeholder value marked with the parent value, and links the new value to the original. MakeOnto is then called recursively on C (line 7). The algorithm terminates when it reaches the \perp category, which has no descendants.

- (1) **procedure** MakeOnto(P)
- (2) **for each** $C \in Desc(P)$ **do**
- (3) **begin**
- (4) $N \leftarrow P \setminus \Pi_P(R_{C,P})$
- (5) $C \leftarrow C \cup \{Mark(p) \mid p \in N\}$
- (6) $R_{C,P} \leftarrow R_{C,P} \cup \{(Mark(p), p) \mid p \in N\}$
- (7) MakeOnto(C)
- (8) **end**
- (9) **end**

Following the reasoning in Section 4.1, we find that the overall big- \mathcal{O} complexity is $\mathcal{O}(k^2 n \log n)$, where k is the number of categories and n is the size of the largest participating R_{C_1, C_2} relation. However, the complexity will be $\mathcal{O}(kn \log n)$ for the most common cases.

The MakeOnto algorithm inserts new values into C to ensure that the mapping from C to P is summarizable. Again, this means that the cost of materializing results for C will be higher for the transformed hierarchy than for the original. However, if the new values were not inserted, we would have to materialize results for P , and perhaps also higher categories, as well as C . At most one value is inserted in C for every value in P , meaning that the extra cost for C is never greater than the cost of materializing results for P . As before, this is a very unrealistic scenario, as it corresponds to the case where *no* values in P have children in C . In most cases, the extra cost for C will be a small percentage of the cost of materializing results for P , and the potential savings will be even greater because pre-aggregation for higher-level categories may be avoided.

As before, the correctness argument for the algorithm has two aspects. First, the mappings in the hierarchy should be *onto* upon termination. Second, the algorithm should only make transformations that are semantically correct. The correctness follows from Theorems 3 and 4, below. Again, the result set for the original values obtained using the original hierarchy will be a subset of the result set obtained using the transformed hierarchy. The results for the new values can be excluded from the result set by adding a HAVING clause condition.

Theorem 3 Algorithm MakeOnto terminates and the hierarchy for the resulting dimension D' is onto.

Proof: By induction in the height of the lattice [22].

Theorem 4 Given dimensions D and D' such that D' is the result of applying algorithm MakeOnto to D , an aggregate result obtained using D is a subset of the result obtained using D' .

Proof: Follows easily from the observation that “childless” dimension values are linked to new placeholder values in lower categories in one-to-one relationships, meaning that data for childless values will still be counted exactly once in aggregate computations that use the new dimension.

4.3 Non-Strict Hierarchies

The third algorithm renders mappings in hierarchies strict, meaning that “double-counting” will not occur. Non-strict

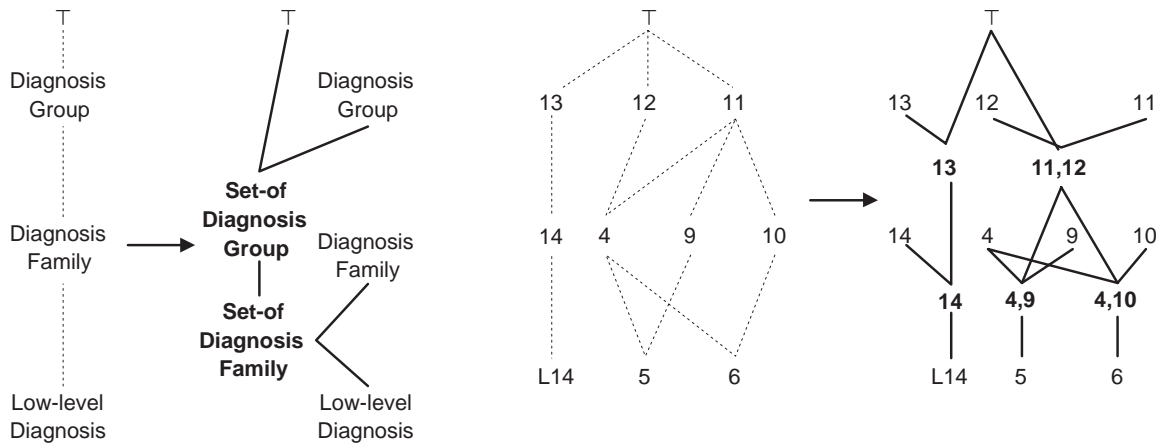


Figure 5: The Diagnosis Dimension Schema and Instance Before and After MakeStrict

hierarchies occur when one dimension value has several parent values.

The basic idea is to “fuse” a set of parent values into one “fused” value, then link the child value to this new value instead. The fused values are inserted into a new category in-between the child and parent categories. Data for the new fused category may safely be re-used for computation of higher-level aggregate results, as the hierarchy leading up to the new category is strict.

The fused value is also linked to the relevant parent values. This mapping is by nature non-strict, but this non-strictness is not a problem, as we prevent aggregate results for the parent category from being re-used higher up in the hierarchy. This is done by “unlinking” the parent category from its predecessor categories.

The categories higher up are instead reached through the fused category. This means that we can still get results for any original category, while being able to apply practical pre-aggregation throughout the hierarchy. In pre-aggregation terms, the “unlinking” of the parent categories means that we must prevent results for including this category from being materialized—only “safe” categories may be materialized. This should be given as a constraint to the pre-aggregation system that chooses which levels of aggregation to materialize.

We note that the algorithm does not introduce more *levels* in the hierarchy, only more categories, and that the number of “safe” categories in the result is the same as the number of original categories. This means that the complexity of the task of selecting the optimal aggregation levels to materialize is unaffected by the algorithm.

Example 8 The result of running the algorithm on the Diagnosis dimension is seen in Figure 5. Because of the non-strictness in the mapping from Low-level Diagnosis to Diagnosis Family, and from Diagnosis Family to Diagnosis Group, two new category types and the corresponding categories are introduced. The third picture indicates the argument to the algorithm; and, in addition, its dotted lines indicate the links deleted by the algorithm. The fourth picture

gives the result of applying the algorithm; here, the bold-face values and thick lines indicate the values and links inserted by the algorithm (note that all lines are thick, as no original links remain).

In the first call of the algorithm, the three Low-level Diagnosis values—“(low-level) Lung cancer” (L14); “Insulin dependent diabetes during pregnancy” (5); and “Non insulin dependent diabetes during pregnancy” (6)—are linked to the three new fused values—“(low-level) Lung cancer” (14); “Diabetes during pregnancy, Insulin dependent diabetes” (4, 9); and “Diabetes during pregnancy, Non insulin dependent diabetes” (4, 10)—and these are in turn linked to “Lung Cancer” (14); “Diabetes during pregnancy” (4); “Insulin dependent diabetes” (9); and “Non insulin dependent diabetes” (10). The these latter four values in the Diagnosis Family category are un-linked from their parents, as the Diagnosis Family category is “unsafe.”

When called recursively on the Set-of Diagnosis Family category, the algorithm creates the new fused values “Cancer” (13) and “Diabetes, Other pregnancy related diseases” (11, 12) in the Set-of Diagnosis Group category. These new values are linked to the values “Cancer” (13), “Diabetes” (11), and “Other pregnancy related diseases” (12) in the Diagnosis Group category, and to the \top value; and the values in the Diagnosis Group category are un-linked from their parents. Note the importance of having a \top value: the values not linked to \top are exactly the unsafe values, for which aggregate results should not be re-used.

The algorithm assumes that all paths in the dimension hierarchy have equal length, i.e., all direct links are from children to their immediate parents. This is ensured by the MakeCovering and MakeOnto algorithms. In the algorithm that follows, C is a *child* category, P is a *parent* category, G is a *grandparent* category, N is the *new* category introduced to hold the “fused” values, and \bowtie denotes natural join.

The algorithm takes a category C (initially the \perp category) as input. It then goes through the set of immediate parent categories P of C (line 2). Line 4 tests if there is non-strictness in the mapping from C to P and if P has

any parents. If this test fails, there is no problem as aggregate results for P can either be safely re-used or are guaranteed not to be re-used; and the algorithm is then invoked recursively, in line 20.

```

(1) procedure MakeStrict( $C$ )
(2) for each  $P \in \text{Pred}(C)$  do
(3) begin
(4) if  $(\exists e_1 \in C (\exists e_2, e_3 \in P (e_1 \leq e_2 \wedge e_1 \leq e_3$ 
       $\wedge e_2 \neq e_3))) \wedge \text{Pred}(P) \neq \emptyset$  then
(5) begin
(6)  $N \leftarrow \text{CreateCategory}(2^P)$ 
(7)  $R_{C,N} \leftarrow \{(e_1, \text{Fuse}(\{e_2 \mid (e_1, e_2) \in R_{C,P}\}))\}$ 
(8)  $N \leftarrow \Pi_N(R_{C,N})$ 
(9)  $R_{N,P} \leftarrow \{(e_1, e_2) \mid e_1 \in N \wedge e_2 \in \text{Unfuse}(e_1)\}$ 
(10)  $\text{Pred}(C) \leftarrow \text{Pred}(C) \cup \{N\} \setminus \{P\}$ 
(11)  $\text{Pred}(N) \leftarrow \{P\}$ 
(12) for each  $G \in \text{Pred}(P)$  do
(13) begin
(14)  $R_{N,G} \leftarrow \Pi_{N,G}(R_{N,P} \bowtie R_{P,G})$ 
(15)  $\text{Pred}(N) \leftarrow \text{Pred}(N) \cup \{G\}$ 
(16)  $\text{Pred}(P) \leftarrow \text{Pred}(P) \setminus \{G\}$ 
(17) end
(18) MakeStrict( $N$ )
(19) end
(20) else MakeStrict( $P$ )
(21) end
(22) end

```

If the test succeeds, the algorithm creates a new fused category. First, an empty category N with domain 2^P is created in line 6. The values inserted into this category represent sets of values of P . For example, the value “1, 2” represents the set consisting of precisely 1, 2. Values in C are then linked to new, fused values, representing their particular combination of parents in P (line 7). The new values are constructed using a Fuse function that creates a distinct value for each combination of P values and stores the corresponding P values along with it.

The resulting links are used in line 8 to insert the fused values into their category N , and an “Unfuse” function, mapping fused values from N into the corresponding P values, is used in line 9 to map the values in N to those in P . In line 10, N is included into, and P is excluded from, the sets of predecessors of C . The set of predecessors of N is set to P in line 11, meaning that the new category N resides in-between C and P in the hierarchy.

For each grandparent category G , the algorithm links values in N to values in G , in line 14; it includes G in the predecessors of N , in line 15; and it excludes G from the predecessors of P , in line 16, thereby also deleting the links from P to G from the hierarchy. The exclusion of the G categories from the predecessors of P means that aggregate results for P will not be re-used to compute results for the G categories.

In the end, the algorithm is called recursively on the new category, N . Note that the test for $\text{Pred}(P) \neq \emptyset$ in line (4) ensures that the mapping from N to P will not be altered, as P now has no predecessors.

Following the reasoning in Section 4.1, we find that the overall big- O complexity is $O(pnk \log n \log k)$, where p is the number of immediate parent and children categories in the dimension type lattice, n is the size of the largest mapping in the hierarchy, and k is the maximum number of values fused together. For most realistic scenarios, p and k are small constants, yielding a low $O(n \log n)$ complexity for the algorithm.

The MakeStrict algorithm constructs a new category N and inserts fused values into N to achieve summarizability for the mapping from N to P and from N to G . The algorithm only inserts the fused values for the combinations that are actually present in the mapping from C to P . This means that the cost of materializing results for N is never higher than the cost of materializing results for C . This is a worst-case scenario—for the most common cases, the cost of materializing results for N will be close to the cost of materializing results for P . However, without the introduction of N , we would have to materialize results not only for P , but also for G and all higher-level categories. Thus, the potential savings in materialization costs are very high.

Considering correctness, the mappings in the hierarchy should be *strict* upon termination, and the algorithm should only make transformations that are semantically correct. More specifically, it is acceptable that some mappings be non-strict, namely the ones from the new, fused categories to the unsafe parent categories that do *not* have predecessors in the resulting hierarchy, meaning that aggregate results for these categories will not be re-used.

The correctness follows from Theorems 5 and 6, below. When evaluating queries we get the same result for original values as when evaluating on the old hierarchy. The values that are deleted by the algorithm were not linked to any facts, meaning that these values did not contribute to the results in the original hierarchy. As all the new values are inserted into new categories that are unknown to the user, the aggregate result obtained will be the same for the original and transformed hierarchy. We do not need to modify the original query.

Theorem 5 Let D' be the dimension resulting from applying algorithm MakeStrict on dimension D . Then the following hold: Algorithm MakeStrict terminates and the hierarchy for the dimension D'' , obtained by removing unsafe categories from D' , is strict.

Proof: By induction in the height of the lattice [22].

Theorem 6 Given dimensions D and D' such that D' is the result of applying algorithm MakeStrict to D , an aggregate obtained using D' is the same as that obtained using D .

Proof: Follows from Lemma 2 [22].

Lemma 2 For the dimension $D' = (C', \leq')$ resulting from applying algorithm MakeStrict to dimension $D = (C, \leq)$, the following holds. $\forall e_1, e_2 \in D (e_1 \in C_1 \wedge \text{Safe}(C_1) \wedge e_1 \leq' e_2 \Leftrightarrow e_1 \leq e_2)$ (there is a path between an original dimension value in a safe category and any other original dimension value in the new dimension hierarchy iff there was a path between them in the original hierarchy).

Proof: By induction in the height of the lattice [22].

4.4 Transforming Fact-Dimension Relations

The algorithms from the previous sections may also be applied to the relationships between facts and dimensions. The basic idea is to view the facts as the bottom granularity in the lattice. The inputs to the algorithms then are the facts F , the $R_{F,C}$ tables that describe the mappings from facts to dimension values, and the C and R_{C_1,C_2} tables that describe the dimension categories and the mappings between them.

Only the covering and strictness properties are of concern for the fact-dimension relationships. An *into* mapping from facts to dimension values means that not all dimension values in the bottom category have associated facts, which does not affect summarizability. The MakeCovering and MakeStrict algorithms may be applied to render relationships summarizable (see [22] for a detailed coverage).

First, facts may be mapped *directly* to dimension values in categories higher than the \perp category. This mapping to values of *mixed* granularities means that not all facts will be accounted for when materializing aggregate results for lower categories. The MakeCovering algorithm rectifies this situation. Second, relationships between facts and dimension values may be many-to-many. This means that the hierarchy, with the facts as the bottom category, is non-strict, leading to possible double-counting of facts. The Makestrict algorithm addresses this problem.

In the case study, the mapping between patients and diagnoses is of mixed granularity and also many-to-many: a patient may have several diagnoses, each of which which may belong to any of the levels Low-Level Diagnosis, Diagnosis Family, and Diagnosis Group.

5 Architectural Context

The overall idea presented in this paper is to take unnormalized MOs and transform them into normalized MOs that are well supported by the practical pre-aggregation techniques available in current OLAP systems. Queries are then evaluated on the transformed MOs. However, we still want the users to see only the original MOs, as they reflect the users' understanding of the domain. This prompts the need for means of handling both the original and the transformed MOs. This section explores their coexistence.

A current trend in commercial OLAP technology is the separation of the front-end presentation layer from the back-end database server. Modern OLAP applications consist of an OLAP client that handles the user interface and an OLAP server that manages the data and processes queries. The client communicates with the server using a standardized application programming interface (API), e.g., Microsoft's OLE DB for OLAP [16] or the OLAP Council's MDAPI [19]. The architecture of such a system is given to the left in Figure 6.

This separation of client and server facilitates our desire to have the user see the original MO while queries are evaluated against the transformed MO. Studies have shown that queries on a data warehouse consist of 80% *navigational*

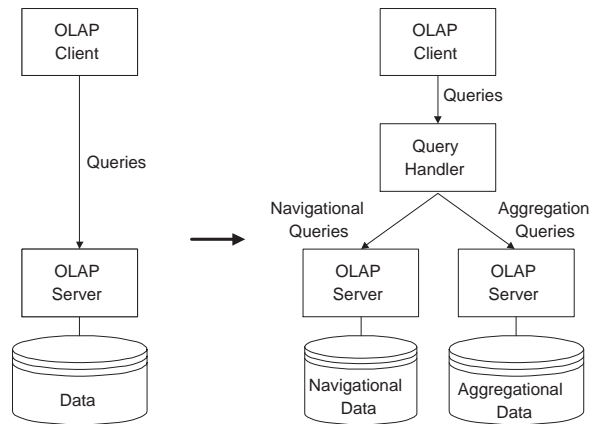


Figure 6: Architecture of Integration

queries that explore the dimension hierarchies and 20% *aggregation* queries that summarize the data at various levels of detail [13]. Examples of navigational and aggregation queries are “Show me the Low-Level Diagnoses contained in the Insulin-Dependent Diabetes Diagnosis Family” and “Show me the count of patients, grouped by Diagnosis Family,” respectively. The navigational queries must be performed on the *original* MO, while the aggregation queries must be performed on the *transformed* MO. This is achieved by introducing an extra “Query Handler” component between the client and the server. The OLAP client sends a query to the query handler, the primary task of which is to determine whether the query is a navigational query (internal to a dimension) or an aggregation query (involving the facts). Navigational queries are passed to one OLAP server that handles the original (navigational) data, while aggregation queries are passed to another OLAP server that manages the transformed (aggregation) data. This extended architecture is seen to the right in Figure 6.

The OLAP server for navigation data needs to support dimension hierarchies which have non-summarizable properties, a requirement not yet supported by many commercial systems today. However, relational OLAP systems using snow-flake schemas [13] are able to support this type of hierarchies, as are some other OLAP systems, e.g., Hyperion (Arbor) Essbase [11]. If the OLAP system available does not have sufficiently flexible hierarchy support, one solution is to build a special-purpose OLAP server that conforms to the given API. This task is not as daunting as it may seem at first because only *navigational* queries need to be supported, meaning that multidimensional queries can be translated into simple SQL “lookup” queries.

We note that the only data needed to answer navigational queries is the hierarchy definitions. Thus, we only need to store the fact data (facts and fact-dimension relations, in our model) once, in the aggregational data, meaning that the overall storage requirement is only slightly larger than storing just the aggregational data. Navigational queries are evaluated on the original hierarchy definitions and do not need to be re-written by the query handler.

As described in Section 4, aggregation queries need to

be re-written slightly by adding an extra HAVING clause condition to exclude results for the new values inserted by the transformation algorithms. This can be done automatically by the query handler, giving total transparency for the user. Even though the added HAVING clause conditions are only necessary for the covering and onto transformations, they can also be applied to hierarchies transformed to achieve strictness; this has no effect, but simplifies the query rewriting.

6 Conclusion and Future Work

Motivated by the increasing use of OLAP systems in many different applications, this paper provides transformation techniques for multidimensional databases that leverage the performance-enhancing technique known as practical, or partial or semi-eager, preaggregation, making this technique relevant to a much wider range of real-world applications.

Current pre-aggregation techniques assume that the dimensional structures are *summarizable*. Specifically, the mappings in dimension hierarchies must be *onto*, *covering*, and *strict*; the relationships between facts and dimensions must be many-to-one, and facts must be mapped to the lowest categories in dimensions. The paper presents transformation techniques that render dimensions with hierarchies that do not satisfy these properties summarizable. The transformations have low, practical computational complexity, they may be implemented using standard relational database technology, and the paper also describes how to integrate the transformed hierarchies in current OLAP systems, transparently to the user. The algorithms may also be applied to the cases of non-summarizable relationships between facts and dimensions, which also occur often in real-world applications, and the algorithms can be modified to incrementally maintain the transformed hierarchies when the underlying data is modified [22].

Several directions for future research appear promising. The current techniques render the entire dimension hierarchies summarizable; extending the techniques to consider only the parts that have been selected for preaggregation appears attractive and possible. Another direction is to take into account the different types of aggregate functions to be applied, leading to local relaxation of the summarizability requirement. For example, *max* and *min* are insensitive to duplicate values, which relaxes summarizability.

References

[1] E. Baralis et al. Materialized View Selection in a Multidimensional Database. In *Proc. of VLDB*, pp. 156–165, 1997.
 [2] E. F. Codd. Providing OLAP (on-line analytical processing) to user-analysts: An IT mandate. E.F. Codd and Assoc., 1993.
 [3] S. Dar et al. Answering SQL Queries Using Views. In *Proc. of VLDB*, pp. 318–329, 1996.
 [4] C. E. Dyreson. Information Retrieval from an Incomplete Data Cube. In *Proc. of VLDB*, pp. 532–543, 1996.
 [5] J. Gray et al. Data Cube: A Relational Aggregation Operator Generalizing Group-By, Cross-Tab and Sub-Totals. *Data Mining and Knowledge Discovery*, 1(1):29–54, 1997.

[6] A. Gupta et al. Aggregate Query Processing in Data Warehousing Environments. In *Proc. of VLDB*, pp. 358–369, 1995.
 [7] H. Gupta et al. Index Selection for OLAP. In *Proc. of ICDE*, pp. 208–219, 1997.
 [8] H. Gupta. Selection of Views to Materialize in a Data Warehouse. In *Proc. of ICDT*, pp. 98–112, 1997.
 [9] H. Gupta and I.S. Mumick. Selection of Views to Materialize Under a Maintenance-Time Constraint. In *Proc. of ICDT*, pp. 453–470, 1999.
 [10] V. Harinarayan et al. Implementing Data Cubes Efficiently. In *Proc. of SIGMOD*, pp. 205–216, 1996.
 [11] Hyperion Corporation. Hyperion Essbase OLAP Server. <www.hyperion.com/downloads/essbaseolap.pdf>. Current as of May 20, 1999.
 [12] Informix Corporation. MetaCube ROLAP Option for Informix Dynamic Server. <www.informix.com/answers/english/pdf_docs/metacube/4189.pdf>. Current as of May 20, 1999.
 [13] R. Kimball. *The Data Warehouse Toolkit*. Wiley Computer Publishing, 1996.
 [14] R. Kimball. Help with Multi-Valued Dimension. *DBMS Magazine*, 11(9), 1998.
 [15] H. Lenz and A. Shoshani. Summarizability in OLAP and Statistical Data Bases. In *Proc. of SSDBM*, pp. 39–48, 1997.
 [16] Microsoft Corporation. OLE DB for OLAP Version 1.0 Specification. Microsoft Technical Document, 1998.
 [17] Microsoft Corporation. OLAP Services White Paper. <www.microsoft.com/sql/70/whpprs/olapoverview.htm>. Current as of May 20, 1999.
 [18] I. S. Mumick et al. Maintenance of data cubes and summary tables in a warehouse. In *Proc. of SIGMOD*, pp. 100–111, 1997.
 [19] The OLAP Council. *MDAPI Specification Version 2.0*. OLAP Council Technical Document, 1998.
 [20] The OLAP Report. *Database Explosion*. <www.olapreport.com/DatabaseExplosion.htm>. Current as of May 20, 1999.
 [21] T. B. Pedersen and C. S. Jensen. Multidimensional Data Modeling for Complex Data. In *Proc. of ICDE*, 1999. Extended version available as TimeCenter TR-37, <www.cs.auc.dk/TimeCenter>, 1998.
 [22] T. B. Pedersen et al. Extending Practical Pre-Aggregation in On-Line Analytical Processing. TR R-99-5004, Comp. Sci. Dept., Aalborg University, <www.cs.auc.dk/~tbp/articles/R995004.ps>, 1999.
 [23] D. Quass and J. Widom. On-Line Warehouse View Maintenance for Batch Updates. In *Proc. of SIGMOD*, pp. 393–404, 1997.
 [24] M. Rafanelli and A. Shoshani. STORM: A Statistical Object Representation Model. In *Proc. of SSDBM*, pp. 14–29, 1990.
 [25] A. Shukla et al. Storage Estimation for Multidimensional Aggregates in the Presence of Hierarchies. In *Proc. of VLDB*, pp. 522–531, 1996.
 [26] D. Theodoratos and T. Sellis. Data Warehouse Configuration. In *Proc. of VLDB*, pp. 126–135, 1997.
 [27] J. Widom. Research Problems in Data Warehousing. In *Proc. of CIKM*, pp. 25–30, 1995.
 [28] R. Winter. Databases: Back in the OLAP game. *Intelligent Enterprise Magazine*, 1(4):60–64, 1998.
 [29] World Health Organization. *International Classification of Diseases (ICD-10)*. Tenth Revision, 1992.
 [30] J. Yang et al. Algorithms for materialized view design in a data warehousing environment. In *Proc. of VLDB*, pp. 136–145, 1997.