# Discovering and merging related analytic datasets 

Rutian Liu, Eric Simon, Bernd Amann, Stéphane Gançarski

## To cite this version:

Rutian Liu, Eric Simon, Bernd Amann, Stéphane Gançarski. Discovering and merging related analytic datasets. Information Systems, inPress, 91, pp.101495. 10.1016/j.is.2020.101495 . hal-02459098

## HAL Id: hal-02459098

## https://hal.science/hal-02459098

Submitted on 21 Jul 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Discovering and Merging Related Analytic Datasets^ 

Rutian Liu ${ }^{\text {a, },}$, Eric Simon ${ }^{\text {b }}$, Bernd Amann ${ }^{\text {c }}$, Stéphane Gançarski ${ }^{\text {d }}$<br>${ }^{a}$ SAP France, LIP6, Paris, France<br>${ }^{b}$ SAP France, Paris, France<br>${ }^{c}$ LIP6, Sorbonne Université, CNRS, Paris, France<br>${ }^{d}$ LIP6, Sorbonne Université, CNRS, Paris, France


#### Abstract

The production of analytic datasets is a significant big data trend and has gone well beyond the scope of traditional IT-governed dataset development. Analytic datasets are now created by data scientists and data analysts using big data frameworks and agile data preparation tools. However, despite the profusion of available datasets, it remains quite difficult for a data analyst to start from a dataset at hand and customize it with additional attributes coming from other existing datasets. This article describes a model and algorithms that exploit automatically extracted and user-defined semantic relationships for extending analytic datasets with new atomic or aggregated attribute values. Our framework is implemented as a REST service in SAP HANA and includes a careful theoretical analysis and practical solutions for several complex data quality issues.


Keywords: schema augmentation, schema complement, data quality, SAP HANA

## 1. Introduction

### 1.1. Context and motivation

Enterprise Business Intelligence (BI) provides business users with solutions for managed data reporting, analysis, aggregation and visualization. These solutions heavily rely on trusted and well documented analytic datasets comprising multidimensional facts that hold measures and refer to one or more hierarchical dimensions [1]. Traditionally, these analytic datasets are created by the IT department in the form of data warehouses and data marts [2] or by enterprise application software vendors in the form of predefined and customizable analytic models. For example, SAP provisions thousands of predefined and customizable analytic datasets, also known as "virtual data models" for various business application domains (e.g., SCM, CRM, ERP) [3, 4]. These datasets are defined as views over the transactional data stored and managed by the SAP S4/HANA business suite and carry information, including sophisticated measures, which is easily understandable by business users, and ready for consumption by BI tools. Recently, agile data preparation tools $[5,6,7]$ emerged to empower business users and data scientists to easily create their own high-quality analytic datasets from transactional data and existing analytic datasets and to build insightful and interactive personalized data visualizations.

[^0]Data preparation tools have the ability to partially automate data cleaning and transformation operations, perform record linkage and duplicate elimination, and even extract structured information from unstructured data. Another typical requirement of business users or data scientists is to augment the schema of an existing analytic dataset with new attributes, coming from one or more semantically related datasets that may represent additional details on dimensions or new measures. This is a critical need in many scenarios such as the creation of data mashups (e.g., add demographics to sales information), or the engineering of features within datasets used to train predictive models (e.g., features engineering to learn identifying customers that are likely to sign up for a credit card in the following quarter). Despite the importance of this task, existing data preparation tools poorly support users in augmenting the schema of an analytic dataset. This lack of assistance compels business users to redefine multiple times similar analytic datasets in a possibly inconsistent manner or to depend on their IT department to create their customized datasets. This creates an important bottleneck on the IT organization and pushes forward the deadline for making the desired datasets available to business users.

Addressing the problem of agile schema augmentation in the context of building trusted analytic datasets for specific BI tasks is a major business opportunity for several reasons. First, companies generally manage large collections of analytic datasets which keep growing with the need for new business data analysis tasks. At the scale of a large company, the facility offered by agile data preparation
tools increases considerably the number of available analytic datasets with respect to those created by the IT organization or those predefined by enterprise software application vendors. Second, semantic relationships between analytic datasets, which are essential to support schema augmentation, can be accurately and automatically extracted from the dataset definitions. Indeed, analytic datasets are often defined over other datasets using queries, scripts, or views. For example, in BI applications supported by SAP, it is common to find analytic datasets with five or more levels of nested view definitions. By parsing these view definitions it is possible to discover the dimensions shared by different analytic datasets. Third, analytic datasets are generally complemented by rich and carefully designed metadata (e.g., about which attributes represent levels of hierarchical dimensions, or about the units and currencies of measure attributes) to support complex data analysis scenarios, that can be exploited during the schema augmentation process. Finally, IT organization and business users invest time to create analytic datasets containing high quality data (i.e., clean data) for business process optimization and decision making. Reusing these datasets and their metadata for schema augmentation is therefore worthwhile. For all these reasons, analytic datasets are a "gold mine" of high-quality and interrelated data that is relevant to business users and data scientists, although under-exploited by current data preparation tools.

Table 1: Sales

| ORG_ID | CITY | COUNTRY | YEAR |
| :--- | :--- | :--- | ---: |
| SALESORGNUE |  |  |  |
| Oh_01 | Dublin | USA | 2018 |
| Ca_01 | Dublin | USA | 2018 |
| Ie_01 | Dublin | Ireland | 2018 |

### 1.2. Main challenges

The manual augmentation of an analytic dataset is often a cumbersome and error-prone process, raising multiple challenges illustrated hereafter. Imagine a simple use case with two analytic datasets represented by fact tables Sales and Dem(ographics) as shown in Tables 1 and 2. Attributes ORG_ID, CITY and COUNTRY in table Sales come from dimension table $S A L E S O R G$, attribute YEAR in tables Sales and Dem is from dimension table TIME, and attributes CITY, STATE, COUNTRY in table Dem are from

Table 2: Dem (Demographics)

| CITY | STATE | COUNTRY | YEAR | POP | UNEMP |
| :--- | :--- | :--- | :--- | ---: | ---: |
| REGION |  |  |  | TIME | $(\mathrm{K})$ |$|(\%)$.

dimension table REGION. Dimension table names are in italic font to distinguish them from fact tables. Assume that attribute ORG_ID is unique in Sales while attributes CITY, STATE, COUNTRY, and YEAR are unique in Dem. In addition, assume that SALESORG and REGION are also related through common attributes CITY and COUNTRY which means that these attributes have the same meaning in both tables. Then, tables Sales and Dem are naturally related through their common attribute YEAR which comes from the same dimension and attributes CITY and COUNTRY which come from different dimensions but are semantically equivalent.

In our example, all tables are defined as views over transactional data as shown on Figure 1. Dimension tables are represented by rounded rectangles, fact tables by square rectangles. Plain edges represent data dependencies, expressing for instance that the definition of Sales depends on SALESORG, TIME, and some transactional data. A data analyst now might want to build a new analytic table SalesDem by augmenting the schema of dataset Sales with the measure attributes POP(ulation) and UNEMP(loyment rate) of dataset Dem. This augmentation then corresponds to yet another fact table (view) SalesDem defined by a left-outer join with Dem on the common attributes CITY, COUNTRY and YEAR.


Figure 1: The construction of table SALES - DEM
The goal of this article is to propose a solution for assisting the user in the definition of this new augmentation view. In particular, we will show that by exploiting all available information on analytic datasets, possibly extended by some other user-defined metadata, it is possible to automatically generate for a given dataset, a set of useful and correct schema augmentations. This is done by solving several challenges we will describe below.

Relationship extraction. The first challenge is to discover the relationships between analytic datasets that can be used for schema augmentation. As explained before, many useful relationships can be extracted from the definitions of analytic datasets. For example, the view definition of fact table SALES in Figure 1 uses attributes from dimensions SALESORG and TIME. Similarly, fact table DEM uses attributes from dimension TIME. Thus, relationships can be extracted to identify semantically equivalent attributes of TIME in both fact tables and to generate left-outer
join queries like the query for the augmentation SalesDem of table Sales.

Row multiplication. In our example, the outer-join operation increases the number of rows in Sales (e.g., generates two rows with ORG_ID $=$ 'Oh_01') because the join attributes do not constitute a unique identifier in dataset Dem. However, for some application scenarios in feature engineering, data enrichment or data analysis, such a multiplication of rows in Sales is undesirable. There, schema augmentation should be controlled to keep the number of rows in Sales constant, leading to the notion of schema complement introduced in [8]. This raises the challenge of identifying row multiplication by computing the unique, and possibly minimal, identifiers of analytic datasets.

Incorrect and ambiguous reduction. A simple solution to avoid row multiplication is to transform table Dem by reducing partitions identified by the common attributes with Sales into single tuples. For instance, a possible reduction operation is to pre-aggregate the measures of table Dem along attribute STATE before performing a left outer join. However, this raises two new challenges. The first challenge is to choose the correct aggregation functions that can be applied to the aggregated attributes. For this we need the knowledge that POP can be summed and averaged while only a maximal or minimal value can be computed on the unemployment rate values of attribute UNEMP. Furthermore, if an averaged POP is used to augment Sales, we must determine what aggregation functions are applicable on this new attribute in the augmented Sales dataset to avoid future incorrect aggregation operations. The second challenge is the ambiguity of the aggregated measure attributes added to Sales which, for instance, contain an aggregated value for all cities 'Dublin' in country 'USA', thereby destroying the distinction between cities in different states and sharing the same name. Ambiguous measures to should be detected and controlled, e.g., by assigning a null value to POP and UNEMP for tuples with ORG_ID $=$ 'Oh_01' and ORG_ID $=$ 'Ca_01'.

Incomplete merge. Another solution to avoid row multiplication would be to add STATE in the schema of Sales (although this is not necessary since ORG_ID was supposed to be unique in Sales). Then the common join attributes would form a unique identifier in Dem and a schema complement for dataset Sales can be obtained without any preaggregation. However, even in that case, summing POP in the augmented Sales dataset along CITY will not return the same result as summing POP along CITY in Dem, even for the same values of COUNTRY and STATE in Sales (e.g., POP value of city 'San Jose' in 'California' will not be counted). This can be a problem if in the augmented Sales dataset, REVENUE should be compared with POP in each STATE of a Country. This issue arises because the merge of Sales with Dem is incomplete with respect to Dem and detecting incomplete merge is necessary to avoid such an erroneous analysis.

### 1.3. Research contributions

This paper presents a new solution to assist business users and data scientists in augmenting the schema of an analytic dataset with attributes coming from other datasets. This is achieved by automatically discovering related analytic datasets and suggesting ways of building a schema augmentations, including schema complements, that resolve the challenges presented before.

More specifically, we make the following technical contributions.

- We introduce attribute graphs as a novel concise and natural way to define literal functional dependencies over the level types of hierarchical dimensions from which we can easily infer unique identifiers in both dimension and fact tables, and minimal identifiers in the case of dimensions.
- We give the formal definitions for schema augmentation, schema complement and merge query in the context of analytic tables. We then present several reduction operations for transforming schema augmentations with row multiplication into schema complements extending each row in the source table by a single row in the augmented table. These operations extend previous contributions on schema augmentation and schema complement (e.g., [8], [9], [10]) to the case of analytic datasets.
- We also present formal quality criteria for schema augmentations, schema complements and merge queries. These criteria are used to define automatic repair operations (1) to notify the generation of ambiguous attributes, (2) to infer applicable aggregation functions on new attributes, and (3) to supplement merge results obtained by incomplete schema augmentation.
- We describe the implementation of our solution as a REST service within the SAP HANA platform and provide a detailed description of our algorithms. We separate the generic part of the algorithms from the specific implementation optimizations done by leveraging the capabilities of SAP HANA.
- We evaluate the performance of our algorithms to compute unique identifiers in dimension and fact tables, and analyze the effectiveness of our REST service using two application scenarios.


### 1.4. Paper outline

The rest of this paper is structured as follows. In Section 2, we describe the multi-dimensional data model for analytic datasets, which is used to capture all usual concepts of hierarchical dimensions, facts, measures, and cubes [1]. This makes our results easily applicable to any other database system implementing the concepts of multidimensional models. We also introduce the new concept of
attribute graph as a way to define literal functional dependencies between dimension and fact table attributes and to compute dimension and fact table identifiers. In Section 3, we describe the relationships between analytic tables that are used to support schema augmentation and that can be automatically extracted from the definition of fact and dimensions tables. We then provide the formal definitions of schema augmentation and of natural and reduction-based schema complement. Section 4 introduces formal quality criteria for schema augmentations, schema complements and merge queries. Section 5 describes the implementation of our solution as a REST application service, not in production yet, in the SAP HANA platform while Section 6 details our algorithms. Section 8 presents our experiments to evaluate the performance and effectiveness of our solution. Section 7 compares our research results with previous work on schema augmentation as well as work on the querying of single or multiple OLAP data cubes. We conclude in Section 9.

## 2. Data Model

This section presents our extension of the relational data model to model analytic tables and their semantic relationships. We assume as input to our model any relational database containing a set of non-analytic tables. Analytic tables, i.e., dimension and fact tables, are then defined bottom-up as views over these non-analytic tables and other analytic tables. We use conventional relational database notations [11]. Each table $T$ is a finite multiset of tuples over a set of domains of values $S=\left\{A_{1}, \ldots, A_{n}\right\}$, called attributes, where each domain may contain a null marker. We call S the schema of $T$.

### 2.1. Hierarchies and dimension tables

We consider a multidimensional data model with dimensions, i.e. hierarchies of values defined by hierarchy types. A hierarchy type $\mathbf{H}=(\mathrm{L}, \preccurlyeq)$ is a set of level types $\mathrm{L}=\left\{L_{1}, \cdots, L_{n}\right\}$ that is organized by a partial order $\preccurlyeq$. $L_{i}$ is called a child level type of $L_{j}$ if there exists an edge $L_{i} \preccurlyeq L_{j}$ and $L_{i} \preccurlyeq^{*} L_{j}$ denotes that $L_{i}$ is a descendant level type of $L_{j}$. We call all types $L_{i}$ where there exists no type $L_{j}$ such that $L_{j} \preccurlyeq L_{i}$ or $L_{i} \preccurlyeq L_{j}$, respectively the lower and the upper bounds of $\mathbf{H}$.

Example 1. Two examples of hierarchy types are shown in Figure 2 where an arc from $A$ to $B$ means that $A \preccurlyeq B$.
(a) CITY $\longrightarrow$ STATE $\longrightarrow$ COUNTRY $\longrightarrow$ CONTINENT
(b) DAY $\longrightarrow$ WEEK $^{\longrightarrow}$ MONTH $\longrightarrow$ YEAR

Figure 2: Hierarchy types GEOGRAPHY (a) and TIME (b)

Each level type $L_{i}$ represents a domain $N_{i}$ of values related to the values of the domains $N_{j}$ of other level types
$L_{j}$. More precisely, a hierarchy $H=(N, \leq)$ of hierarchy type $\mathbf{H}=(\mathrm{L}, \preccurlyeq)$ is a set of values $N$ and a partial order $\leq$ where $N$ contains for each level type $L_{i} \in \mathrm{~L}$ a non empty subset of values $N_{i} \subseteq N$ such that each order relation $v_{i} \leq$ $v_{j}$ preserves the ancestor/descendant relation $\preccurlyeq^{*}$ between the corresponding hierarchy types $L_{i}$ and $L_{j}$, i.e., $v_{i} \in$ $N_{i}, v_{j} \in N_{j} \Rightarrow L_{i} \preccurlyeq^{*} L_{j}$. We also assume that ( $N, \leq$ ) is transitively reduced, i.e., there is no pair of nodes that is connected by an edge and a sequence of two or more edges.


Figure 3: Hierarchy REGION of type GEOGRAPHY

Example 2. Figure 3 shows a hierarchy REGION (right part) of type GEOGRAPHY (left part).

Hierarchies can naturally be represented in relational tables, called dimension tables or, more simply, dimensions.

Definition 1 (Dimension table). Any hierarchy $H=(N, \leq$ ) of type $\mathbf{H}=(\mathrm{L}, \preccurlyeq)$ defines a dimension table $D(\mathrm{~S})$ with a one-to-one mapping $\phi: \mathrm{L} \rightarrow \mathrm{X}$ from level types $L_{i} \in \mathrm{~L}$ to $a$ subset of attributes $\mathrm{X} \subseteq \mathrm{S}$ in the schema of $T$ such that for each maximal path $v_{1} \cdot v_{2} \ldots . v_{k}$ in $N_{1} \times N_{2} \times \cdots \times N_{k}$ in $H$ there exists a tuple $t \in T$ where $t . \phi\left(L_{1}\right)=v_{1}, t . \phi\left(L_{2}\right)=$ $v_{2}, \cdots, t . \phi\left(L_{k}\right)=v_{k}$ and $t . A_{j}=$ null for all other attributes in X.

Example 3. Dimension table REGION in Table 3 represents the hierarchy REGION of hierarchy type GEOGRAPHY in Figure 3. Each level type is translated into an attribute of the same name and each tuple represents a maximal path in the hierarchy (null markers are denoted with "-").

Table 3: Dimension table REGION

| CITY | STATE | COUNTRY | CONTINENT |
| :--- | :--- | :--- | :--- |
| Dublin | Ohio | United States | North America |
| Dublin | California | United States | North America |
| Dublin | - | Ireland | Europe |
| Dublin | Ontario | Canada | North America |
| Dublin | - | Belarus | Europe |
| Palo Alto | California | United States | North America |
| - | - | - | Antarctica |

All attributes of a dimension table $D$ corresponding to level types are called the dimension attributes of $D$ and the
hierarchy type $\mathbf{H}=(\mathrm{L}, \preccurlyeq)$ can naturally be mapped into a dimension attribute hierarchy $\mathcal{A}=(\mathrm{S}, \preccurlyeq)$ where $\mathrm{A}_{i} \preccurlyeq \mathrm{~A}_{j}$ iff $L_{i} \preccurlyeq L_{j}$. The remaining attributes are called detail attributes and functionally depend on one or more dimension attributes (dimension table REGION only contains dimension attributes).

Null markers in dimension tables represent regular values yielding the same interpretation for equality as null values in SQL unique constraints [11]: two attribute values $t_{1}$. A and $t_{2}$. A are literally equal, denoted by $t_{1} \cdot \mathrm{~A} \equiv t_{2} \cdot \mathrm{~A}$, if $t_{1}$.A and $t_{2}$.A are equal or both null markers. Literal equality naturally extends to sets of attributes and leads to the notion of literal functional dependencies (LFD) [12]. A literal functional dependency LFD X $\mapsto Z$ holds between two subsets of attributes if for any two tuples $t_{1}$ and $t_{2}$ of $T$, when $t_{1} \cdot \mathrm{X} \equiv t_{2} \cdot \mathrm{X}$ then $t_{1} \cdot Z \equiv t_{2} \cdot Z$. Note that if X does not contain any nullable attribute, which is for instance enforced in SQL by declaring X as a primary key, an LFD $\mathrm{X} \mapsto Z$ is equivalent to a functional dependency with nulls.

Definition 2 (Dimension identifier). Let $\mathrm{X} \subseteq \mathrm{S}$ be the set of dimension attributes in a schema S and $\mathcal{L}$ be a set of LFDs defined on X . Then, $\mathrm{K} \subseteq \mathrm{X}$ is a dimension identifier of S if K is a minimal set such that $\mathrm{K} \mapsto \mathrm{S}$ holds for any instance $T$ over S satisfying all LFDs in $\mathcal{L}$.

The set of LFDs defined on a dimension table schema characterizes the set of all valid schema instances. However, defining this set of LFDs may be difficult for the designer of a dimension table. We introduce the notion of attribute graphs, which are defined over attribute hierarchies and which allow the schema designer to characterize the set of all valid dimension tables by additional semantic properties in a simple and natural way.

Definition 3 (Attribute graph). An attribute graph over some attribute hierarchy $\mathcal{A}=(\mathrm{S}, \preccurlyeq)$ is a directed labeled graph $\mathcal{D}=\left(\mathrm{S}, \mathrm{R}, \lambda_{\mathrm{R}}, \perp, \mathrm{T}\right)$, where S is the set of attributes in $\mathcal{A}, \perp$ and $T$ are two special attributes with empty domains (by definition, t. $\perp \equiv t . \top \equiv$ null for all tuples), $\mathrm{R} \subseteq(\mathrm{S} \cup\{\perp, \top\})^{2}$ is a set of edges such that there exists an edge:

1. $\left(A_{i}, A_{j}\right) \in \mathrm{R}$ for each edge $A_{i} \preccurlyeq A_{j}$ in $\mathcal{A}$;
2. $\left(\perp, A_{i}\right) \in \mathrm{R}$ for each lower bound in $\mathcal{A}$ and
3. $\left(A_{i}, T\right) \in \mathrm{R}$ for each upper bound in $\mathcal{A}$.

There might exist also other edges $\left(A_{i}, A_{j}\right) \in \mathrm{R}$ between any two nodes connected by a path in $\mathcal{D}$.

In the following we denote by $\mathrm{R}\left(\mathrm{A}_{i}, \mathrm{~A}_{j}\right)=l$ an edge $\left(\mathrm{A}_{i}, \mathrm{~A}_{j}\right) \in \mathrm{R}$ labeled by $l=\lambda_{\mathrm{R}}\left(\mathrm{A}_{i}, \mathrm{~A}_{j}\right)$. The edge labeling function $\lambda_{R}: R \rightarrow\{+, \mathbf{1}, \mathbf{f}\}$ assigns to each edge a unique label encoding the presence of functional and literal functional dependency constraints between the connected attributes of a dimension table.

Definition 4 (Valid dimension table). A dimension table $T$ with schema S is valid with respect to some attribute graph $\mathcal{D}=\left(\mathrm{X}, \mathrm{R}, \lambda_{\mathrm{R}}\right)$ where $\mathrm{X} \subseteq \mathrm{S}$, if all following conditions hold in $T$ :

1. If there exists a tuple $t \in T$ such that $t . A_{i}$ is not null, then there exists either an edge $\mathrm{R}\left(\perp, A_{i}\right)$ in $\mathcal{D}$ or an edge $\mathrm{R}\left(A_{k}, A_{i}\right)$ in $\mathcal{D}$ such that $t . A_{k}$ is not null.
2. For all tuples $t_{1}, t_{2}$ in $T$ and all edges $\mathrm{R}\left(A_{i}, A_{j}\right)$ in $\mathcal{D}$ the following holds:
(a) If $\mathrm{R}\left(A_{i}, A_{j}\right)=\mathbf{f}$, then $t_{1} \cdot A_{i} \equiv t_{2} \cdot A_{i}$ implies $t_{1} \cdot A_{j} \equiv t_{2} \cdot A_{j} ;$
(b) If $\mathrm{R}\left(A_{i}, A_{j}\right)=\mathbf{1}$, then $t_{1} \cdot A_{i}=t_{2} \cdot A_{i}$ implies $t_{1} \cdot A_{j} \equiv t_{2} \cdot A_{j}$.
In Definition 4, $t_{1} \cdot \mathrm{~A}=t_{2} \cdot \mathrm{~A}$ is true if both tuples $t_{1} \cdot \mathrm{~A}$ and $t_{2}$.A are equal non-null values, false if both $t_{1}$.A and $t_{2}$.A are different non-null values and unknown otherwise. From this definition, it follows naturally that $t_{1} \cdot \mathrm{~A}=t_{2} \cdot \mathrm{~A}$ implies $t_{1} . \mathrm{A} \equiv t_{2}$. . but not the opposite. Observe also that by Definition $4, \mathrm{R}\left(\mathrm{A}_{i}, \mathrm{~A}_{j}\right)=\mathbf{f}$ is equivalent to $\mathrm{A}_{i} \mapsto \mathrm{~A}_{j}$ and $\mathrm{A}_{i} \rightarrow \mathrm{~A}_{j}$ implies $\mathrm{R}\left(\mathrm{A}_{i}, \mathrm{~A}_{j}\right)=\mathbf{1}$ whereas $\mathrm{R}\left(\mathrm{A}_{i}, \mathrm{~A}_{j}\right)=\mathbf{1}$ does not imply $\mathrm{A}_{i} \rightarrow \mathrm{~A}_{j}$. Consequently, if some dimension table $T$ is valid w.r.t. some attribute graph $\mathcal{D}$, it is also valid w.r.t. to all attribute graphs obtained by replacing $\mathbf{f}$ edge labels by $\mathbf{1}$ edge labels and $\mathbf{1}$ edge labels by + edge labels. We can also see that all edges $\mathrm{R}\left(\perp, \mathrm{A}_{i}\right)$ are labeled by + or $\mathbf{f}$. Indeed, $t_{i} . \perp \equiv t_{j} . \perp$ holds for all couples $\left(t_{i}, t_{j}\right)$, and either $t_{i} \cdot \mathrm{~A}_{i} \equiv t_{j} . \mathrm{A}_{i}$ also holds for all couples $\left(t_{i}, t_{j}\right)$, i.e. $\mathrm{R}\left(\perp, \mathrm{A}_{i}\right)=\mathbf{f}$ or not, i.e. $\mathrm{R}\left(\perp, \mathrm{A}_{i}\right)=+\left(t_{i} \cdot \mathrm{~A}_{i} \equiv t_{j} \cdot \mathrm{~A}_{i}\right.$ does not hold for at least one couple). Symmetrically, all edges $\mathrm{R}\left(\mathrm{A}_{i}, \top\right)$ are labeled by $\mathbf{f}$ since $t_{i} . \top \equiv t_{j}$. T for all tuples $t_{i}$ and $t_{j}$.

Example 4. Figure 4 shows an attribute graph for hierarchy type GEOGRAPHY that validates dimension table REGION. The lower and upper bound attributes are respectively CITY (connected to node $\perp$ ) and Continent (connected to node T ). The arc labels of attribute graph have the following semantics. First, since $\mathrm{R}(S T A T E, C O U N T R Y)=\mathbf{1}$, for each non-null value of attribute STATE, we can determine a unique value of attribute COUNTRY. Second, R (COUNTRY,CONTINENT) $=\mathbf{f}$ states that the attribute COUNTRY literally determines the attribute CONTINENT. Third, the same value of CITY can have multiple values for STATE $(\mathrm{R}($ CITY, STATE $)=+)$ or multiple values for COUNTRY $(\mathrm{R}(C I T Y, C O U N T R Y)=+)$ and no value for STATE. Finally, there exists a single continent without countries, states and cities, which is represented by the arc $\mathrm{R}(\perp$, CONTINENT) with label $\mathbf{f}$.


Figure 4: Attribute graph for dimension REGION

Proposition 1. Let $\mathcal{D}=\left(\mathrm{S}, R, \lambda_{R}, \perp, \top\right)$ be an attribute graph. Then, the subset $\mathrm{X}_{D} \subseteq \mathrm{~S}$ of all attributes in S with at least one + labeled in-edge and no $\mathbf{f}$ labeled in-edge is a dimension identifier for all valid dimension tables with attributes S .

The proofs of the previous and the following propositions can be found in the Appendix A.

Attribute graphs concisely describe well-know properties of hierarchies, as described for instance in [1]. Edges $\mathrm{R}(\perp, \mathrm{A})$ can model "unbalanced" hierarchies, edges $\mathrm{R}\left(\mathrm{A}_{i}, \mathrm{~A}_{j}\right)$ such that $\mathrm{A}_{i} \npreceq \mathrm{~A}_{j}$ in $\mathcal{A}$ model "non-covering hierarchies", and edges $\mathrm{R}\left(\mathrm{A}_{j}, \mathrm{~A}_{i}\right)$ with label 1 represent "nonstrict" hierarchies.

Example 5. By Proposition 1, the dimension identifier for dimension REGION is (CITY, STATE, COUNTRY). Figure 5 shows an attribute graph for the schema of a dimension table WAREHOUSE. It expresses that attribute Wh_ID literally determines CITY and STATE but not COUNTRY. By Proposition 1, the identifier of W AREHOUSE is (WH_ID, COUNTRY).


Figure 5: Attribute graph for dimension WAREHOUSE

### 2.2. Fact tables

Fact tables associate measures with dimension values. A fact table over a set of dimensions $D_{1}, \cdots, D_{n}$ is a table $T$ whose schema S contains a non-empty subset $\mathrm{X}_{i}$ of dimension attributes from each dimension $D_{i}$, and a nonempty set of attributes Z representing one or more measures where $\mathrm{X}_{1} \cup \cdots \cup \mathrm{X}_{n} \mapsto \mathrm{~S}$. The active domain of each dimension attribute $T . A \in \mathrm{X}_{i}$ is a subset of the active domain of $D_{i}$.A. Each measure is represented by a group of attributes where one attribute has the role of Value and other attributes have the role Detail. Value attributes carry the actual measure values while Detail attributes provide auxiliary information on the measure. In particular, each measure attribute can be associated with a Unit or Currency detail attribute to control the application of aggregate functions. For simplicity, the term measure attribute will refer in the following to a Value attribute.

Definition 5 (Fact identifier). Let $\mathrm{K} \subseteq \mathrm{S}$ be a set of dimension attributes in the schema of a fact table $T(\mathrm{~S})$ and $\mathcal{L}$ be a set of LFDs defined on S . Then, K is a fact identifier of S if $\mathrm{K} \mapsto \mathrm{S}$ holds for any instance T over S satisfying all LFDs in $\mathcal{L}$.

Each attribute graph of some dimension represents a set of LFDs and can be used to compute fact identifiers for some schema $S$ as stated in the following proposition.

Proposition 2. Let $T(\mathrm{~S})$ be a fact table defined over a set of dimensions $D_{1}, \cdots, D_{n}$ and $\mathrm{K}_{1}, \cdots, \mathrm{~K}_{n}$ be the dimension identifiers of $D_{1} \cap \mathrm{~S}, \cdots, D_{n} \cap \mathrm{~S}$ respectively. Then $\mathrm{K}=$ $\mathrm{K}_{1} \cup \cdots \cup \mathrm{~K}_{n}$ is a fact identifier of $T$, and K is a minimal identifier if all dimensions in $T$ are mutually independent.

Example 6. Consider fact tables SALES and INVENTORY below. SALES contains the sales of products, is defined over dimensions STORE, TIME, and PROD. INVENTORY is defined over dimensions WAREHOUSE, TIME, TAX and PROD. Measure attributes are in italics. Figure 6 shows the attribute graphs for dimensions STORE, PROD and TAX with identifiers STORE_ID, PROD_SKU and TAX_NO respectively.

SALES ( PROD_SKU, BRAND, MONTH, YEAR, CITY, STATE, COUNTRY, AMOUNT, CURRENCY)

INVENTORY ( $\underline{\text { PROD_SKU }}, B R A N D, \underline{Y E A R}, \underline{W H-I D}$, CITY, COUNTRY, TAX_NO, RATE, TAX_DESC, QTY_ON_HAND)


Figure 6: Attribute graphs of dimensions $S T O R E, P R O D$ and $T A X$
The fact identifiers of fact tables SALES and INVENTORY are defined by the union of the corresponding dimension identifier attributes. If all dimensions are mutually independent, they form a minimal set.

### 2.3. Aggregable attributes

In analytic tables, attributes are aggregable, but an attribute does not necessarily aggregate with any aggregate function along any dimension, as shown by extensive previous work in statistical and OLAP databases [13].

Example 7. In the schema of fact table SALES, measure AMOUNT is aggregable using SUM along all dimension attributes. However, measure QTY_ON_HAND in fact table INVENTORY is not aggregable along the TIME dimension, i.e. summing the quantity of products QTY_ON_HAND over all months produces is considered senseless. So, a SUM over QTY_ON_HAND must at least be grouped by YEAR.

Our model enables the designer of an analytic table to declare for each attribute the aggregate function that is applicable and the maximal set of dimension attributes along which this aggregation can be computed. Clearly, if some attribute A is aggregable along a set of dimension attributes Z , then it is also aggregable along any subsets
of Z . In the following we denote by $\operatorname{agg}_{\mathrm{A}}(\mathrm{F}, \mathrm{Z})$ the $a g$ gregable property of A and state that property $\operatorname{agg}_{\mathrm{A}}(\mathrm{F}, \mathrm{Z})$ holds in $T$ if Z is the maximal set of attributes along which A is aggregable using F in $T$. More precisely, we give the following definition for aggregable properties $\operatorname{agg}_{\mathrm{A}}(\mathrm{F}, \mathrm{Z})$ :
Definition 6. Let X be the set of dimension attributes in an analytic table $T$ over schema $\mathrm{S}, A$ be an aggregable attribute in S and F be an aggregate function on A . Aggregable property $\operatorname{agg}_{A}(\mathrm{~F}, \mathrm{Z})$ holds in $T$ if

1. all aggregations along any subsets of Z are considered as meaningful by the user;
2. Z only contains dimension attributes which determine $A$, i.e. for all dimension attributes $B \in \mathrm{Z} \cap \mathrm{X}$ there exists a minimal subset of dimension attributes $\mathrm{U} \subseteq \mathrm{X}$ such that $\mathrm{U} \mapsto A$ and $B \in \mathrm{U}$;
3. F is applicable to $A$.

Example 8. Considering Example 7, attribute QTY_ON_HAND depends on all dimension attributes and can be aggregated along these attributes except attribute YEAR of dimension TIME which is considered meaningless to the user (item 1 in Definition 6). We then state that $\operatorname{agg}_{Q T Y-O N-H A N D}(\mathrm{SUM}, Z)$ holds in INVENTORY where Z contains all dimension attributes of INVENTORY except YEAR.

Aggregable properties generalize the approach of [14] and [2], which only focused on SUM aggregations and categorized measures as being additive, fully-additive, nonadditive, or semi-additive.

For determining valid aggregate function for new attributes, we provide a default categorization of attributes. Two categories, NUM and DESC, are obtained from the (SQL) data type of attributes. A third category STAT captures numerical values resulting from the use of some aggregate functions. Table 4 describes the five common SQL aggregate functions applicable to each category.

Table 4: Categories of aggregable attributes

| Category | Properties |
| :--- | :--- |
| NUM | $\bullet \quad$ Numerical values |
|  | $\bullet \quad$ Applicable functions: SUM, AVG, COUNT, |
|  | MIN, MAX |
| DESC | $\bullet$ |
|  | Descriptive or categorical values |
| STAT | • Applicable functions: COUNT |
|  | $\bullet$ |

Example 9. Attribute STORE_ID of data type string in dimension table STORE is of category DESC and can only be counted. Attributes amount and QtY_on_hand in tables SALES and INVENTORY respectively are of category NUM. Thus, by default, both attributes can be summed up unless a user-defined aggregable property is specified, as in Example 8 for attribute QTY_ON_HAND.

If A is an attribute of $T$ for which no aggregable property was defined by a user, we first use the default category of A to define which aggregate function $F$ is applicable to A and then state that a default aggregable property $\operatorname{agg}_{\mathrm{A}}(\mathrm{F}, \mathrm{Z})$ holds in $T$ where Z is the set of all dimension attributes of $T$ (by default, A literally depends on all dimension attributes).

### 2.4. Summary of metadata acquisition

To conclude this section, we summarize the data and related metadata introduced by our data model in Table 5 and indicate how they are obtained. We distinguish between data and metadata that can be user-defined, automatically computed from other data or metadata and derived by default rules. Most human effort is required on the definition of dimension and fact tables (analytic schemas) with a careful description of their attribute graphs. We shall see in Section 6.1 that attribute graphs can also be automatically computed from dimension tables or samples thereof. This essentially reduces the human-effort for enabling schema augmentations and complements to defining analytic schemas and aggregable properties.

Table 5: Summary of data model concepts

| Metadata | Source |
| :--- | :--- |
| Dimension and fact tables | User-defined |
| Attribute graphs | User-defined or computed |
| Dimension identifiers | Computed |
| Fact identifiers | Computed |
| Aggregable properties | User-defined or default rule |

## 3. Schema Augmentation and Complement

In this section we first describe the relationships that can be automatically extracted from the definitions of dimension and fact tables and formally define the notion of schema augmentation between tables connected by such relationships. We then introduce natural and reductionbased schema complements as a special case of schema augmentation where the tuples of a table are augmented with new attribute values obtained from another table without generating new tuples in the result.

### 3.1. Relationships

We distinguish between join relationships and attribute-mapping relationships. Both kinds of relationships link semantically equivalent attributes and can be extracted from the view definitions of analytic tables. New join relationships between tables that are not related through views can also be added by the business user if necessary.
Definition 7 (Join relationship). $A$ join relationship $\mathrm{R}\left(T_{1}, T_{2}\right)$ between two tables $T_{1}$ and $T_{2}$ is defined by a nonempty set of equality atoms $\left\{P_{1}, \cdots, P_{k}\right\}$ where each $P_{i}$ is
of the form $T_{1} . A=T_{2} . B$. It relates all tuples of $T_{1}$ and $T_{2}$ satisfying the join predicate $P_{1} \wedge \cdots \wedge P_{k}$, i.e., the set of all related tuples is defined by $T_{1} \bowtie_{P_{1} \wedge \cdots \wedge P_{k}} T_{2}$.


Figure 7: The construction of INVENTORY

Example 10. Fact table INVENTORY of Example 6 is a view over dimensions PROD, TIME, TAX and WAREHOUSE, the view is defined by a star join between a non-analytical table ct_INVENTORY and these dimensions as shown in Figure 7. The analysis of the view definition yields four join relationships between ct_INVENTORY and the dimension tables.

Join relationships might also exist in the database schema between non-analytic tables (i.e., tables that are neither dimension nor fact tables) through foreign key constraints which can be easily extracted from the schema definition.

Definition 8 (Attribute mapping relationship). Let $T_{1}$ and $T_{2}$ be two tables such that $T_{1}$ is derived from $T_{2}$ using a query $Q$ and auxiliary tables: $T_{1}=Q\left(T_{2}, T_{i}, \cdots, T_{n}\right)$ with $n \geq 0$. Query $Q$ defines an attribute mapping from $T_{2} \cdot B$ to $T_{1} \cdot A$, denoted by $T_{2} \cdot B \rightarrow T_{1} . A$, if for all possible values $x \in \operatorname{dom}(A)$ and all possible instances of $T_{2}, T_{i}, \cdots, T_{n}$ :

$$
\sigma_{A=x}\left(\mathcal{Q}\left(T_{2}, T_{i}, \cdots, T_{n}\right)\right)=\sigma_{A=x}\left(Q\left(\sigma_{B=x}\left(T_{2}\right), T_{i}, \cdots, T_{n}\right)\right)
$$

An attribute mapping relationship $\mathrm{R}\left(T_{1}, T_{2}\right)$ between $T_{1}$ and $T_{2}$ is defined by a non-empty set of attribute mappings from $T_{1}$ to $T_{2}$.
Example 11. Dimension PROD from Example 6 is defined as a view over three non-analytic tables, where X is the set of attributes in $P R O D$, and $P_{1}, P_{2}$ are join predicates: $P R O D=\pi_{\mathrm{X}}\left(\mathrm{PRODUCT} \bowtie_{P_{1}}\right.$ SUBCATEGORY $\bowtie_{P_{2}}$ CATEGORY). Predicate $P_{1}$ defines a join relationship between PRODUCT and SUBCATEGORY on SUBCATEGORY_ID whereas predicate $P_{2}$ defines a join relationship between SUBCATEGORY and CATEGORY on CATEGORY_ID. An attribute mapping relationship between PROD and each non-analytic table providing an attribute in X is also extracted from the view definition.

Example 12. In the view definition of Example 10, since INVENTORY is a projection of the star join result, there exists an attribute mapping relationship between table INVENTORY and the other five tables.

Each relationship $\mathrm{R}\left(T_{1}, T_{2}\right)$ between two tables $T_{1}\left(\mathrm{~S}_{1}\right)$ and $T_{2}\left(\mathrm{~S}_{2}\right)$ defines a partial mapping $\mu_{\mathrm{R}}$ from attributes A of $T_{1}$ to attributes B of $T_{2}$ where $\mu_{\mathrm{R}}\left(T_{1} \cdot \mathrm{~A}\right)=T_{2} \cdot \mathrm{~B}$ if (i) R is a join relationship containing an atom $T_{1} \cdot \mathrm{~A}=T_{2} \cdot \mathrm{~B}$ or (ii) R contains an attribute mapping $T_{1} \cdot \mathrm{~A} \rightarrow T_{2}$. or $T_{2}$. $\rightarrow T_{1}$. A. Relationship R is well-formed if $\mu_{\mathrm{R}}$ is a one-to-one mapping. For the sake of simplicity, unless specified differently, we assume in the following that all relationships R define a natural mapping $\mu_{\mathrm{R}}\left(T_{1} \cdot \mathrm{~A}\right)=T_{2}$. A for some attributes A of $\mathrm{S}_{1} \cap \mathrm{~S}_{2}$ shared by table $T_{1}$ and $T_{2}$ and the set of attributes A is called the common attributes of $T_{1}$ and $T_{2}$.

Example 13. In Example 11, the common attributes for join relationships R (CATEGORY, SUBCATEGORY) and R (PRODUCT, SUBCATEGORY) are respectively \{CATEGORY_ID\} and \{SUBCATEGORY_ID\}. Similarly, the attribute mapping relationship $\mathrm{R}(P R O D, S U B C A T E G O R Y)$ between dimension table $P R O D$ and non-analytical table SUBCATEGORY has common attributes \{PROD_SKU, BRAND $\}$.

Consider now a fact table ORDER recording detail descriptions for orders from different stores and customers. ORDER has two attributes ORDER_DAY, SHIP_DAY, each one referring to attribute DAY of dimension TIME. A single attribute mapping relationship R (ORDER, TIME) mapping the two ORDER attributes to the same attribute TIME.DAY, $\mu_{\mathrm{R}}$ (ORDER.ORDER_DAY) $=$ TIME.DAY, $\mu_{\mathrm{R}}$ (ORDER.SHIP_DAY) $=$ TIME.DAY, is not a one-to-one mapping and therefore not well-formed. Thus, two distinct relationships $\mathrm{R}_{1}$ (ORDER,TIME) and $\mathrm{R}_{2}$ (ORDER,TIME) are needed, one for ORDER_DAY and the other for SHIP_DAY.

Figure 8 shows a graph of relationships issued from previous examples. Dimension tables are represented by rounded rectangles, fact tables by bold square rectangles, and non-analytic tables by square rectangles. Each node contains its identifier (ID), and edges indicate relationships labelled with their common attributes. Edges with solid lines are extracted relationships while edges with bold dashed lines are user-defined relationships. For instance, a join relationship is explicitly defined by a designer between dimensions STORE and WAREHOUSE, using equality atoms on attributes CITY, STATE and COUNTRY.

### 3.2. Schema augmentation

We first introduce the notion of schema augmentation using well-formed relationships and the notion of merge query.

Definition 9 (Schema augmentation). Let $T_{0}\left(\mathrm{~S}_{0}\right)$ and $T(\mathrm{~S})$ be two tables related by a relationship R with a set of common attributes $\mathrm{Y}=\mathrm{S}_{0} \cap \mathrm{~S}$. Then table $T$ is a schema augmentation to $T_{0}$ with respect to R if $\mathrm{Y} \subset \mathrm{S}$.

Observe that, by definition, common attributes Y is not empty and $T$ must bring new attributes to $T_{0}$.


Figure 8: Examples of relationships

Definition 10 (Merge query). Let $T_{0}\left(\mathrm{~S}_{0}\right)$ and $T(\mathrm{~S})$ be two tables related by a relationship R with a set of common attributes Y . Then the merge of $T_{0}$ and $T$ is a left-outer join query $Q=\Pi_{\mathrm{X}}\left(T_{0} \searrow_{P_{1} \wedge \cdots \wedge P_{k}} T\right)$, where $\mathrm{K}=|\mathrm{Y}|$ is the number of common attributes of $\mathrm{R}, \Pi$ is a duplicate elimination projection. $Q$ is defined as follows:

1. If all common attributes are dimension attributes, and there exists a pair of common attributes $A_{1}, A_{2} \in$ Y in $T_{0}$ such that $T_{0} . A_{1} \preccurlyeq T_{0} . A_{2} \wedge T . A_{1} \npreceq T . A_{2}$, then $\mathrm{X}=\mathrm{S}_{0} \uplus \mathrm{~S}$ else $\mathrm{X}=\mathrm{S}_{0} \cup \mathrm{~S}$.
2. For each $A_{i} \in \mathrm{Y}, \exists P_{i} \in P$ such that $P_{i}$ is of the form: $\left(T_{0} \cdot A_{i}=T \cdot A_{i}\right) O R\left(T_{0} \cdot A_{i}=\right.$ null $A N D T \cdot A_{i}=$ null $)$.

In the following, we will abbreviate $Q=\Pi_{\mathrm{X}}\left(T_{0} \searrow_{P_{1} \wedge \cdots \wedge P_{k}}\right.$ T) to $Q=T_{0} \searrow_{Y} T$.

Item 1 checks if the attribute hierarchy is preserved by relationship R and keeps the common attributes separately (disjoint union $\uplus$ ) for each table if this is not the case. Item 2 manages the SQL semantics of joins in presence of nulls by adapting the join predicates.

Example 14. Table SALES is a schema augmentation to INVENTORY with respect to common attributes: $\mathrm{Y}=$ \{PROD_SKU, BRAND, YEAR, CITY, COUNTRY\} and adds MONTH, STATE and AMOUNT to INVENTORY. The merge query is INVENTORY $\perp_{\mathrm{Y}}$ SALES and the schema of the resulting merge does not duplicate the common attributes.

After merging a table with a schema augmentation, the identifier of the new augmented table might change. The following proposition describes how to compute the identifiers of augmented tables.

Proposition 3. Let $T_{0}^{\prime}\left(\mathrm{S}_{0}^{\prime}\right)$ be a merge of table $T_{0}\left(\mathrm{~S}_{0}\right)$ with a target schema augmentation $T(\mathrm{~S})$. Let K be the identifier of $T_{0}$ and $\mathrm{S}_{\text {new }} \subseteq \mathrm{S}_{0}^{\prime}-\mathrm{S}_{0}$ be the set of dimension attributes added to $T_{0}$ in the merge. Then, for all minimal subsets $\mathrm{K}_{\text {new }} \subseteq \mathrm{S}_{\text {new }}$ where $\mathrm{K}_{\text {new }} \mapsto \mathrm{S}_{\text {new }}, \mathrm{K} \cup \mathrm{K}_{\text {new }}$ is an identifier of $T_{0}^{\prime}$.

Example 15. Consider the merged result INVENTORY' in Example 14. Then, $\mathrm{K}=\{$ PROD_SKU, YEAR, WH_ID, COUNTRY\} and $\mathrm{S}_{\text {new }}=\{$ MONTH, STATE $\}$. The set $\mathrm{S}_{\text {new }}$ is also the minimal subset where $\mathrm{S}_{\text {new }} \mapsto \mathrm{S}_{\text {new }}$. Thus, the identifier of INVENTORY' ${ }^{\prime}$ is $\cup \mathrm{S}_{\text {new }}=\left\{\right.$ PROD_SKU, YEAR, $^{2}$, WH_ID, COUNTRY, MONTH, STATE\}.

Schema augmentations define no restriction on the common attributes in underlying table relationship R . The consequence of this unrestricted setting is that many augmentations lead to tuple duplication and unexpected results. For example, in the merge results of Example 14, tuples in INVENTORY will be duplicated because of the newly added dimension attributes MONTH and STATE. Then an aggregation with functions SUM or AVG of QTY_ON_HAND could easily produce wrong results. To better control the schema augmentation process, we introduce new restrictions that can be applied for filtering and transforming schema augmentations before merging them with the source tables.

### 3.3. Natural schema complement

We extend the original definition of natural schema complement [8] by considering LFD and relationships including attribute mappings.

Definition 11 (Natural schema complement). Let $T(\mathrm{~S})$ be a schema augmentation of $T_{0}\left(\mathrm{~S}_{0}\right)$ with a set of common attributes $\mathrm{Y}=\mathrm{S}_{0} \cap \mathrm{~S}$. Table $T$ is a natural schema complement to a table $T_{0}$ with respect to R if $\mathrm{Y} \mapsto \mathrm{S}$.

Condition $\mathrm{Y} \mapsto \mathrm{S}$ guarantees that by joining $T_{0}$ and $T$ on their common attributes Y , each tuple of $T_{0}$ will match at most one tuple of $T$. The merge of $T_{0}$ and $T$ through natural schema complement follows Definition 10 for merging schema augmentations and, by Definition 11 and Proposition 3, the identifier of the merge result is equal to the identifier of the source table $T_{0}$. We refer to the merge with a schema augmentation and natural schema complement by augmented merge and natural merge respectively.
Example 16. The non-analytical table PRODUCT (PROD_SKU, BRAND, ...) is a natural schema complement to dimension table PROD since the common attribute PROD_SKU is a primary key in PRODUCT. This constraint guarantees that the corresponding merge query $P R O D \bowtie^{P R O D}{ }^{\prime} S K U$ PRODUCT generates exactly one tuple for each tuple in PROD. The schema of the natural merge only has one occurrence of the common attribute PROD_SKU (natural outer join).

### 3.4. Reduction queries

When the common dimension attributes of $T_{0}$ and $T$ do not literally determine all attributes of $T$, a natural schema complement can still be obtained by transforming $T$ before applying the outer join merge. These transformations can be obtained by queries which reduce each partition of $T$ identified by the shared attributes to a single tuple. We refer to those queries as reduction queries.

Definition 12 (Reduction query). Let $T(\mathrm{~S})$ be a table with identifier K . A query $Q(T)$ producing the table $T^{\prime}\left(\mathrm{S}^{\prime}\right)$ is called a reduction of $T$ on a subset of attributes $\mathrm{K}^{\prime} \subseteq \mathrm{S}^{\prime}$ if $\mathrm{K}^{\prime} \subset \mathrm{K}$ is a proper subset of identifier K and $\mathrm{K}^{\prime} \mapsto \mathrm{S}^{\prime}$ holds in $T^{\prime}$ ( $\mathrm{K}^{\prime}$ is an identifier of $T^{\prime}$ ). The attributes in $\mathrm{K}-\mathrm{K}^{\prime}$ are called the reduced attributes.
Example 17. Consider the fact table SALES with identifier $\mathrm{K}=\{$ PROD_SKU, MONTH, YEAR, CITY, STATE, COUNTRY\}. A query SALES $^{\prime}=Q$ (SALES) that filters COUNTRY $=$ 'USA' and projects on all other attributes is a reduction of SALES on attributes $\mathrm{K}^{\prime}=\mathrm{K}-\{$ COUNTRY $\}$. $\mathrm{K}^{\prime}$ is the key of SALES ${ }^{\prime}$ and COUNTRY is the reduced attribute.

Each reduction query $\mathcal{Q}(T)$ of table $T$ with key K on attributes $\mathrm{K}^{\prime} \subset \mathrm{K}$ reduces each partition of $T$ with the same values for attributes $\mathrm{K}^{\prime}$ and different values for the reduced attributes $\mathrm{K}-\mathrm{K}^{\prime}$ into a single tuple ( $\mathrm{K}^{\prime}$ is the key of the result).

For a given table and set of attributes, there might exist a great variety of reduction queries. In our work, we focus on three types of reduction queries: aggregate, filter and pivot.

Aggregate reduction. Let $T(\mathrm{~S})$ be an analytic table, A be an aggregable attribute in $S$, and $F$ be an aggregate function such that aggregable property $\operatorname{agg}_{\mathrm{A}}(\mathrm{F}, \mathrm{Z})$ holds in $T$. We denote by $Q(T)=\mathcal{A g g}_{T}(\mathrm{~F}(\mathrm{~A}) \mid \mathrm{X}), \mathrm{S}-\mathrm{Z} \subseteq \mathrm{X} \subseteq \mathrm{S}$, an aggregate query on table $T$ that aggregates A using aggregate function F with group-by attributes X. When $\mathrm{K}^{\prime} \subset \mathrm{K}$, where K is the identifier of $T$, we call query $\mathcal{Q}(T)=\mathcal{A g g}_{T}\left(\mathrm{~F}(\mathrm{~A}) \mid \mathrm{K}^{\prime}\right)$ an aggregate reduction of $T$ on attributes $K^{\prime}$, and $K-K^{\prime}$ are the reduced attributes. Note that the SQL group-by operator implements literal equality semantics for null values (null values are not distinct).

Filter reduction. Let $T(\mathrm{~S})$ be a table, we denote by $Q(T)=\operatorname{Filter}_{T}(P \mid \mathrm{X}), P=\left\{P_{1} \wedge \cdots \wedge P_{i}\right\}$, a filter query that filters $T$ by a set of predicates $P$ on attributes X of $T$. When $\mathrm{K}^{\prime} \subset \mathrm{K}$, where K is the identifier of table $T$, and for each attribute $\mathrm{A} \in \mathrm{K}-\mathrm{K}^{\prime}$, there exists a predicate $P_{i} \in P$ of the form $\mathrm{A}=v_{i}, v_{i} \in \operatorname{dom}(\mathrm{~A})$, we call query $Q(T)=\mathcal{F i l t e r}_{T}\left(P \mid \mathrm{K}-\mathrm{K}^{\prime}\right)$ a filter reduction of $T$ on attributes $\mathrm{K}^{\prime}$, and $\mathrm{K}-\mathrm{K}^{\prime}$ are the reduced attributes.

Pivot reduction. Let $T(\mathrm{~S})$ be a table, and A be an attribute in S. We denote by $Q(T)=\mathcal{P i v o t}_{T}(\mathrm{~A} \mid \mathrm{X}), \mathrm{X} \subset \mathrm{S}$, a pivot query with a set of attributes X whose values are transformed into columns, and the values of attribute A are pivoted into the new columns. The query replaces attributes A and X by converting every distinct tuple $t \in$ $\pi_{\mathrm{X}}(T)$ in the projection of $T$ to a new attribute T.t storing the corresponding values of A . When $\mathrm{K}^{\prime} \subset \mathrm{K}$, where K is the identifier of $T$, we call query $\operatorname{Pivot}_{T}\left(\mathrm{~A} \mid \mathrm{K}-\mathrm{K}^{\prime}\right)$ a pivot reduction of $T$ on attributes $\mathrm{K}^{\prime}$, and $\mathrm{K}-\mathrm{K}^{\prime}$ are the reduced attributes.

The above definitions of aggregate and pivot reductions can be easily generalized by replacing attribute A with a set of attributes.

Example 18. Consider a table $T(\mathrm{~S})$ in Table 6 a with two attributes $A_{1}, A_{2}$ forming an identifier of $T$. The result of a pivot reduction query of $T$ on attribute $A_{1}$, as $\operatorname{Pivot}_{T}(M \mid$ $\left\{A_{2}\right\}$ ), is shown in Table 6b. Domain values of $A_{2}$ are pivoted into two new attributes $M_{-} b_{1}, M_{-} b_{2}$ storing the value of $M$.

Table 6: Example of pivot reduction
(a) Table T

|  |  |  |  |
| :--- | :--- | :--- | :--- |
|  | $\mathrm{A}_{1}$ | $\mathrm{~A}_{2}$ | M |
| $t_{1}$ | $a_{1}$ | $b_{1}$ | $x_{1}$ |
| $t_{2}$ | $a_{2}$ | $b_{1}$ | $x_{2}$ |
| $t_{3}$ | $a_{1}$ | $b_{2}$ | $x_{3}$ |

(b) Table $\operatorname{Pivot}_{T}\left(\mathrm{M} \mid\left\{\mathrm{A}_{2}\right\}\right)$

| $T^{\prime}$ | $\mathrm{A}_{1}$ | $\mathrm{M}_{-} b_{1}$ | $\mathrm{M}_{-} b_{2}$ |
| :--- | :--- | :--- | :--- |
| $a_{1}$ $x_{1}$ $x_{3}$ <br> $a_{2}$ $x_{2}$ - |  |  |  |

Proposition 4. Let $T_{0}\left(\mathrm{~S}_{0}\right)$ and $T(\mathrm{~S})$ be two tables such that $T$ is a schema augmentation to $T_{0}$ with common attributes Y. Let K be an identifier of $T$. The result of a reduction query $Q(T)$ on attributes $\mathrm{K}^{\prime}=\mathrm{K} \cap \mathrm{Y}$, is a natural schema complement to $T_{0}$.

Example 19. Fact tables SALES and INVENTORY in the relationship graph of Figure 8 have common attributes $\mathrm{Y}=\{$ PROD_SKU, BRAND, YEAR, CITY, COUNTRY $\}$. INVENTORY is a schema augmentation to SALES since Y is not a fact identifier of INVENTORY (it misses \{WH_ID, TAX_NO\}). We can then define several reduction queries which transform INVENTORY into a natural schema complement of SALES. A first reduction query is $Q($ INVENTORY $)=\mathcal{A g g}_{\text {INVENTORY }}\left(\operatorname{AVG}\left(Q T Y Y_{-}\right.\right.$ON_HAND $\mid$ $\mathrm{X}), \mathrm{X}=\left\{P R O D_{-} S K U, Y E A R\right.$, COUNTRY $\}$ which aggregates measure attribute QTY_ON_HAND on X. It produces a natural schema complement to SALES. Similarly, a pivot query $\mathcal{P i v o t}_{\text {INVENTORY }}\left(Q T Y_{-} O N_{-} H A N D \mid\left\{W_{-} I D, T A X \_N O\right\}\right)$ or a filter query $\mathcal{F i l t e r}_{\text {INVENTORY }}\left(W H_{-} I D=' 1234\right.$ ' $\wedge T A X$ _NO $=$ ' $a b c d$ ') reduces attributes WH_ID, TAX_NO and transforms INVENTORY into a natural schema complement of SALES. Finally, a reduction query $Q$ (SALES) that first applies a filter SALES $_{\text {FIL }}=\mathcal{F i l t e r}_{\text {SALES }}\left(\left\{\right.\right.$ STATE $=^{\prime}$ Ohio $\left.\left.^{\prime}\right\}\right)$ followed by a pivot query $\operatorname{Pivot}_{\text {SALES }_{F I L}}(\{\operatorname{AMOUNT}\}$ | \{MONTH\}) reduces attributes MONTH and STATE and produces a natural schema complement to INVENTORY.

### 3.5. Derived relationships

We derive new relationships by composition and fusion of relationships which respectively perform the intersection and union of common attributes. Derived relationships offer more schema augmentation opportunities, especially between fact tables. For example, in Figure 8, there is no direct relationship between fact tables INVENTORY and SALES (edge with plain line). However, both tables share dimension tables $P R O D$ and TIME which can be used to derive a new relationship shown by the edge with a dashed line.

Proposition 5. Composition of relationships. Let $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ and $\mathrm{R}_{2}\left(T_{2}, T_{3}\right)$ be two well-formed relationships between tables $T_{1}, T_{2}$ and $T_{3}$ with respective common attributes $\mathrm{Y}_{1}$ and $\mathrm{Y}_{2}$. If $\mathrm{Y}_{3}=\mathrm{Y}_{1} \cap \mathrm{Y}_{2} \neq \emptyset$, then there exists a well-formed relationship $\mathrm{R}_{3}\left(T_{1}, T_{3}\right)$ that is a composition of $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ and $\mathrm{R}_{2}\left(T_{2}, T_{3}\right)$ with common attributes $\mathrm{Y}_{3}$.

Proposition 6. Fusion of relationships. Let $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ and $\mathrm{R}_{2}\left(T_{1}, T_{2}\right)$ be two well-formed relationships between two tables $T_{1}$ and $T_{2}$ with respective common attributes $\mathrm{Y}_{1}$ and $\mathrm{Y}_{2}$. If $\forall A \in \mathrm{Y}_{1} \cap \mathrm{Y}_{2}, \mu_{\mathrm{R}_{1}}(A)=\mu_{\mathrm{R}_{2}}(A)$ then there exists a well-formed relationship $\mathrm{R}_{3}\left(T_{1}, T_{2}\right)$ that is a fusion of $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ and $\mathrm{R}_{2}\left(T_{1}, T_{2}\right)$ with common attributes $\mathrm{Y}_{3}=\mathrm{Y}_{1} \cup \mathrm{Y}_{2}$.

Given a set of well-formed relationships $\mathrm{R}_{0}\left(T_{0}, T\right), \cdots, \mathrm{R}_{n}\left(T_{0}, T\right), n \geq 0$ between two tables $T_{0}, T$ with respective common attributes $\mathrm{Y}_{0}, \cdots, \mathrm{Y}_{n}$. The common attributes of $T_{0}$ and $T$ is the union of $\mathrm{Y}_{0}, \cdots, \mathrm{Y}_{n}$.

Example 20. In Figure 8, relationships $\mathrm{R}(S A L E S, P R O D)$ and $\mathrm{R}(P R O D, I N V E N T O R Y)$ are composed to yield $\mathrm{R}(S A L E S$, INVENTORY) with common attributes $\left\{P R O D \_S K U, B R A N D\right\}$. Similarly, R (SALES, TIME) and $\mathrm{R}($ TIME, INVENTORY ) are composed to generate $\mathrm{R}($ SALES, INVENTORY $)$ with common attribute YEAR. Finally, R (SALES, STORE), R (STORE, WAREHOUSE) and $\mathrm{R}($ WAREHOUSE, INVENTORY) are composed to produce $\mathrm{R}(S A L E S$, INVENTORY) with common attributes \{CITY, COUNTRY\}. Then the fusion of $\mathrm{R}(S A L E S$, INVENTORY $), \mathrm{R}(S A L E S$, INVENTORY) and $\mathrm{R}($ SALES, INVENTORY $)$ yields $\mathrm{R}^{\prime}($ SALES, INVENTORY $)$ with common attributes \{YEAR, PROD_SKU, BRAND, CITY, COUNTRY\}. However, the fusion of the two relationships $\mathrm{R}_{1}($ ORDER,TIME) and $\mathrm{R}_{2}($ ORDER,TIME) in Example 13 is not possible because $\mu_{\mathrm{R}_{1}}(D A Y) \neq \mu_{\mathrm{R}_{2}}(D A Y)$ (both relationships map two different attributes ORDER.ORDER_DAY and ORDER.SHIP_DAY to the same attribute TIME.DAY.

Consider a table $T_{0}^{\prime}\left(\mathrm{S}_{0}^{\prime}\right)$ which is the result of merging table $T_{0}\left(\mathrm{~S}_{0}\right)$ with a target schema augmentation $T(\mathrm{~S})$ through a path of relationships reaching $T$ from $T_{0}$. Then, each relationship $\mathrm{R}\left(T, T_{i}\right)$ (resp. $\mathrm{R}\left(T_{0}, T_{i}\right)$ ) with common attributes Y yields a relationship $\mathrm{R}\left(T_{0}^{\prime}, T_{i}\right)$ if $\mathrm{S}_{0}^{\prime} \cap \mathrm{Y} \neq \emptyset$.

## 4. Quality Guarantees

We now describe the quality guarantees managed by our system. Section 4.1 evaluates the correctness of aggregated attribute values produced by reduction operations and merge queries. Sections 4.2 and 4.3 deal with the generation of ambiguous and incomplete attribute values during the construction of schema augmentations (and schema complements).

### 4.1. Propagation of aggregable properties

When the schema of a table $T_{0}$ is augmented with a new attribute A after a merge with a table $T$, stating the aggregable property of A that holds in the augmented table is critical to avoid incorrect aggregations on that augmented table. This raises two problems:

- The first problem is to define the aggregable properties of all new attributes which are computed by some reduction query. These properties include the identification of the applicable aggregate functions and the set of dimension attributes along which each new attribute can be aggregated in the query result (before merging).
- The second problem is to define the aggregable properties of the new attributes in the augmented table obtained after merging the source table with the reduction query result.

To address the first problem we determine which aggregate functions are applicable to A in the result $T=\mathcal{Q}\left(T^{\prime}\right)$ of a reduction query $Q$ over $T^{\prime}$. This falls into one of the following cases:

1. A is an attribute of $T^{\prime}$, filtered or not. Then if $F$ is applicable to A in $T^{\prime}$, it is also applicable to A in $T$.
2. $\mathrm{A}=F\left(\mathrm{~A}^{\prime}\right)$ for some $\mathrm{A}^{\prime}$ in $T^{\prime}$, i.e., $\mathrm{A}^{\prime}$ has been aggregated. Then if $F$ is a distributive aggregate function (see Definition 13 below), $F$ is also applicable to A in $T$. Other default functions applicable to A are determined by the co-domain categories of function $F$ (NUM, DESC or STAT) using Table 4.
3. A holds pivoted values of an attribute A ${ }^{\prime}$ of $T^{\prime}$. Then, if $F$ is applicable to $\mathrm{A}^{\prime}$ in $T^{\prime}$, it is also applicable to A in $T$.

Definition 13 (Distributive aggregate function). Let F be an aggregate function applicable to a set of domain values V . If for any partition $\mathrm{V}_{1}, \mathrm{~V}_{2}$ of $\mathrm{V}, \mathrm{F}\left(\mathrm{V}_{1} \cup \mathrm{~V}_{2}\right)=$ $\mathrm{F}\left(\mathrm{F}\left(\mathrm{V}_{1}\right) \cup \mathrm{F}\left(\mathrm{V}_{2}\right)\right)$ then F is said to be a distributive aggregate function.

Let us illustrate the previous case analysis. Suppose that an attribute A contains values that are aggregated from attribute $\mathrm{A}^{\prime}$ using the average function AVG (case 2). First, we can easily see that AVG is not distributive, i.e. cannot be applied to A. Then, according to Table 4, statistical functions like AVG and STDEV have the domain category NUM, the co-domain category STAT and the only functions that can be applied are COUNT, MIN and MAX. We can conclude that the default aggregate functions that are applicable on A are COUNT, MIN and MAX.

Now, if $\mathbf{a g g}_{\mathrm{A}^{\prime}}\left(\mathrm{F}, \mathrm{Z}^{\prime}\right)$ holds in $T^{\prime}$ and F is applicable to A in the reduction query result $T=Q\left(T^{\prime}\right)$, the aggregable property $\operatorname{agg}_{\mathrm{A}}(\mathrm{F}, \mathrm{Z})$ holds in $T$ for some subset $\mathrm{Z} \subseteq \mathrm{Z}^{\prime}$ of attributes that are not reduced by the reduction query on
$T^{\prime}$. For determining Z, we use Definition 6 , which states that it is not possible to aggregate an attribute A along dimension attributes which do not determine A (item 2). The following example illustrates how Z can be validated using Definition 6.

Example 21. Suppose that the fact table PRICE_LIST(PROD_SKU, YEAR, PRICE) is a natural schema complement to SALES with common attributes PROD_SKU and YEAR. Let SALES' be the natural merge of SALES with PRICE_LIST adding a new attribute PRICE to SALES. PRICE is of category NUM and aggregable along some set of variables X using SUM, i.e. $\operatorname{agg}_{\text {PRICE }}(\mathrm{SUM}, \mathrm{X})$. Suppose that the user decides to define $\mathrm{X}=\{$ CITY, COUNTRY\}. Then, it is easy to see that an aggregation $\mathcal{A g g}_{\text {SALES }^{\prime}}\left(\mathrm{SUM}(\right.$ PRICE $) \quad \mid \quad\left\{P R O D_{-} S K U\right.$, YEAR, REGION $\left.\}\right)$ would wrongly sum up the PRICE values for all products sold in all cities of the same REGION. Using Definition 6 for aggregable properties, we can see that since attribute PRICE depends on attributes PROD_SKU and YEAR of table PRICE_LIST, these two attributes are the only dimension attributes that might appear in X and $\operatorname{agg}_{\text {PRICE }}(\mathrm{SUM}, \mathrm{X})$ does not hold in SALES' for $\mathrm{X}=\{$ CITY, COUNTRY $\}$.

We now consider the problem of determining the aggregable properties of the new attributes after merge. The following proposition states which aggregable properties hold for attribute A in the augmented table $T_{0}^{\prime}$, knowing the aggregable properties of A that hold in the used schema augmentation $T$.
Proposition 7. Let $T_{0}^{\prime}\left(\mathrm{S}_{0}^{\prime}\right)$ be a merge of table $T_{0}\left(\mathrm{~S}_{0}\right)$ with a target schema augmentation $T(\mathrm{~S})$. Then the following aggregable properties hold for all aggregable attributes $A \in$ $\mathrm{S}_{0}^{\prime}$ :

1. If $\operatorname{agg}_{A}\left(\mathrm{~F}, \mathrm{~V}_{0}\right)$ holds in $T_{0}$ and $A \in \mathrm{~S}_{0}^{\prime} \cap \mathrm{S}_{0}$ then $\operatorname{agg}_{A}\left(\mathrm{~F}, \mathrm{~V}_{0}\right)$ holds in $T_{0}^{\prime}$.
2. If $\operatorname{agg}_{A}(\mathrm{~F}, \mathrm{~V})$ holds in $T$ and $A \in \mathrm{~S}_{0}^{\prime}-\mathrm{S}_{0}$, then $\operatorname{agg}_{A}\left(\mathrm{~F}, \mathrm{~V} \cap \mathrm{~S}_{0}^{\prime}\right)$ holds in $T_{0}^{\prime}$.
Example 22. Suppose SALES ${ }^{\prime}$ is the complete natural merge of SALES and INVENTORY in Figure 8 reduced by the aggregate query $\mathcal{A g g}_{\text {Inventory }}(\operatorname{AVG}(A) \mid \mathrm{X})$ where $A=Q T Y_{-} O N_{-} H A N D$ and $\mathrm{X}=\left\{P R O D_{-} S K U, Y E A R, C O U N T R Y\right\}$. Function AVG is applied on A that transforms A from category NUM to attribute $A V G(A)$ of category STAT. Then, only COUNT, MIN and MAX are applicable on $\operatorname{AVG(A)}$ by default. Because $\operatorname{agg}_{A}(\mathrm{AVG}, \mathrm{V})$ holds in INVENTORY for $\mathrm{V}=\left\{P R O D_{-} S K U, Y E A R\right.$, WH_ID, CITY, COUNTRY, TAX_NO\}, then by Item 2 in Proposition 7, $\mathbf{a g g}_{\mathrm{AVG}(A)}\left(\mathrm{COUNT}, \mathrm{V}^{\prime}\right)$ holds in SALES ${ }^{\prime}$ where $\mathrm{V}^{\prime}=\{Y E A R$, COUNTRY $\}\left(V^{\prime}\right.$ only contains dimension attributes which are necessary to determine attribute $A V G(A))$.

### 4.2. Non-ambiguous aggregable attributes

As illustrated in Section 1, merge queries may produce attribute values that are ambiguous with respect to a nonstrict dimension hierarchy. The next definition states that
a schema augmentation $T(\mathrm{~S})$ may contain ambiguous attribute values with respect to a dimension $D$ if S misses some attributes of $D$ which are necessary to distinguish two tuples of $T$ that are literally equal on their $D$ attributes.

Definition 14 (Ambiguous analytic table). Let $T(\mathrm{~S})$ be an analytic table over a dimension $D\left(\mathrm{~S}_{D}\right)$. Let $\mathrm{X}=\mathrm{S} \cap$ $\mathrm{S}_{D}$ be the set of attributes in S from dimension $D$ and $\mathrm{X}^{*}=\left\{A_{j} \in \mathrm{~S}_{D} \mid \exists A_{i} \in \mathrm{X}, A_{i} \preccurlyeq^{*} A_{j}\right\}$ be the ancestors of all attributes in X in the hierarchy type of $D$. Table $T$ is said to be a non-ambiguous with respect to $D$ if the literal functional dependency $\mathrm{X} \mapsto \mathrm{X}^{*}$ holds in $T \bowtie_{\mathrm{X}} D$ and ambiguous otherwise.

In the following, $\mathrm{X}^{*}$ is called the closure of X in $D$. We also say that a table $T$ is ambiguous if it is ambiguous w.r.t. at least one dimension and non-ambiguous otherwise.

Example 23. Fact table SALES contains attributes $\mathrm{X}=$ \{CITY, STATE, COUNTRY\} from dimension STORE (Figure 8). X is closed w.r.t. the ancestor relationship, i.e. $\mathrm{X}^{*}=\mathrm{X}$. Then $\mathrm{X} \mapsto \mathrm{X}^{*}$ holds in table SALES $\bowtie_{\mathrm{X}}$ STORE and SALES is not ambiguous with respect to STORE. Now suppose that SALES' only contains two attributes $\mathrm{X}_{1}=\{$ CITY, COUNTRY $\}$ from dimension STORE. Then $\mathrm{X}_{1}^{*}=\{$ CITY, STATE, COUNTRY $\}$ and $\mathrm{X}_{1} \mapsto \mathrm{X}_{1}^{*}$ does not hold in SALES' $\bowtie_{\mathrm{X}}$ STORE (the city of Dublin exists in two different states of country USA). In this case, SALES' is ambiguous with respect to dimension STORE.

We can show that merge queries over two nonambiguous tables generate non-ambiguous results. Consider two tables $T_{1}$ and $T_{2}$ that are non-ambiguous w.r.t. to some common dimension $D$, i.e. both literal functional dependencies $\mathrm{X}_{1} \mapsto \mathrm{X}_{1}^{*}$ and $\mathrm{X}_{2} \mapsto \mathrm{X}_{2}^{*}$ hold in $D$. Then, by composition of literal functional dependencies, $\mathrm{X}_{1} \cup \mathrm{X}_{2} \mapsto \mathrm{X}_{1}^{*} \cup \mathrm{X}_{2}^{*}$ holds in $T_{1} \bowtie_{P} T_{2}$.

In the case of reduced schema augmentations, it may happen that a reduction query generates an ambiguous schema complement from a non-ambiguous input table (by removing some dimension attributes). The next proposition defines a sufficient condition for reduction queries to produce non-ambiguous schema augmentations.
Proposition 8. Let $T(\mathrm{~S})$ be a non-ambiguous analytic table w.r.t. a dimension $D$ of schema $\mathrm{S}_{D}$. Let $T^{\prime}\left(\mathrm{S}^{\prime}\right)=\mathcal{Q}(T)$ be a reduction of $T$. Let $\mathrm{X}=\mathrm{S} \cap \mathrm{S}_{D}$ and $\mathrm{X}^{\prime}=\mathrm{S}^{\prime} \cap \mathrm{S}_{D}$. If $\mathrm{X}^{\prime} \mapsto \mathrm{X}$ in $D$, then $T^{\prime}$ is non-ambiguous w.r.t. $D$.

From Proposition 8 we can directly conclude that any filter reduction over a non-ambiguous tables generates a non-ambiguous table (for filter reduction $\mathrm{X}^{\prime}=\mathrm{X}$ and $\mathrm{X}^{\prime} \mapsto$ X holds in all dimensions $D$ ). However, for pivot and aggregate reduction queries we must check if $\mathrm{X}^{\prime} \mapsto \mathrm{X}^{\prime *}$ holds in $T^{\prime} \bowtie_{\mathrm{X}^{\prime}} D$ (Definition 14) to guarantee non-ambiguity ( $\mathrm{X}^{\prime} \mapsto \mathrm{X}$ holds in $D$ is a sufficient but not necessary condition for non-ambiguity).
Example 24. Fact table INVENTORY contains attributes $\mathrm{X}=\{$ WH_ID, CITY, COUNTRY $\}$ from dimension

WAREHOUSE. $\quad \mathrm{X}^{*}=\{\text { Wh_ID, CITY, STATE }, \text { COUNTRY }\}^{\text {G }}$ is the schema of WAREHOUSE and $\mathrm{X} \mapsto \mathrm{X}^{*}$ holds in INVENTORY $\bowtie_{\mathrm{X}}$ WAREHOUSE. Therefore, table INVENTORY is not ambiguous with respect to WAREHOUSE. Now, let $T$ be the result of an aggregate reduction as $\mathcal{A g g}_{\text {INVENTORY }}\left(\mathrm{AVG}\left(Q T Y \_O N \_H A N D\right) \mid \mathrm{X}^{\prime}\right)$ where $\mathrm{X}^{\prime}=\left\{\right.$ CITY, COUNTRY\}. Using Proposition 8, $\mathrm{X}^{\prime} \mapsto$ X does not hold in $D$. Thus, $T$ might be ambiguous with respect to WAREHOUSE. We then check whether the condition $\mathrm{X}^{\prime} \mapsto \mathrm{X}^{\prime *}$ holds in $T \bowtie_{\mathrm{X}^{\prime}}$ WAREHOUSE, where $\mathrm{X}^{*}=\{$ CITY, STATE, COUNTRY\}. Based on the attribute graph in Figure 5, $\mathrm{X}^{\prime} \mapsto \mathrm{X}^{\prime *}$ does not hold, and $T$ is ambiguous with respect to WAREHOUSE

### 4.3. Complete merge results

A third quality problem we presented in the introduction concerns merge queries which might produce results which are incomplete with respect to the original schema augmentation table.

Example 25. Consider the two analytic tables $T_{0}$ and $T$ below related by a well-formed relationship R with common attributes $A_{1} \preccurlyeq A_{2}$ from dimension $D_{1}$, and $B_{1} \preccurlyeq B_{2}$ from dimension $D_{2}$. Identifiers of $T_{0}$ and $T$ are $\mathrm{K}_{0}=$ $\left\{A_{1}, A_{2}, B_{1}, B_{2}\right\}$ and $\mathrm{K}=\left\{A_{1}, A_{2}, A_{3}, B_{1}, B_{2}\right\}$ respectively. Clearly, $T$ is a schema augmentation to $T_{0}$ with respect to R. However, merging $T_{0}$ with $T$ with a left outer join yields a table $T_{0}^{\prime}$ containing tuples $t_{5}$ and $t_{6}$ augmented by two new null valued attributes $A_{3}$ and $M$ (there's no matching tuple in $T$ ). Then, by Proposition 7, if aggregable property $\operatorname{agg}_{M}\left(F,\left\{A_{1}, A_{3}, B_{1}\right\}\right)$ holds in $T$, it also holds in $T_{0}^{\prime}$. Then, some valid aggregate query $\operatorname{Agg}\left(\mathrm{F}(M) \mid\left\{A_{2}, B_{2}\right\}\right)$ will generate a null value for partition $\left(b_{1}, e_{1}\right)$ in the merge $T_{0}^{\prime}$ while the same query has value $\mathrm{F}\left(\left\{y_{1}, y_{4}\right\}\right)$ for the same partition in $T$. This possibly undesirable situation occurs because $T_{0}^{\prime}$ is not a complete merge with respect to $T$.

Table 7: Example of incomplete merge

$$
\begin{aligned}
& \begin{array}{l|l|l|l|l} 
& \\
\\
\hline \mathrm{A}_{1} & \mathrm{~A}_{2} & \mathrm{~B}_{1} & \mathrm{~B}_{2} \\
\hline t_{5} & a_{1} & b_{1} & d_{1} & e_{1} \\
t_{6} & a_{2} & b_{3} & d_{3} & e_{3} \\
\hline
\end{array} \\
& \begin{array}{l|l|l|l|l|l|l|l} 
\\
\hline
\end{array}
\end{aligned}
$$

Definition 15 (Candidate completion tuples). Let $T_{0}^{\prime}$ be the merge of $T_{0}$ and $T$ that are related by a relationship R with a set of common attributes Y. Let $\mathrm{Y}^{\text {top }} \subseteq \mathrm{Y}$ be the set of the "highest" attributes of Y in the corresponding attribute hierarchies. We can define the following two tables $T^{c t} \subseteq T$ and $T^{c a n d} \subseteq T:$

- the natural semi-join $T^{c t}=T \ltimes_{\mathrm{Y}} T_{0}^{\prime}$, where $\mathrm{Y}=$ $\wedge_{i}\left(T . A_{i}=T_{0}^{\prime} \cdot A_{i}\right)$ for each $A_{i} \in \mathrm{Y}$, is called the completion table of $T$ with respect to $T_{0}^{\prime}$.
- the natural semi-join $T^{\text {cand }}=T \ltimes_{\mathrm{Y} \text { top }} T_{0}^{\prime}$, where $\mathrm{Y}^{\text {top }}=\wedge_{i}\left(T \cdot A_{i}=T_{0}^{\prime} \cdot A_{i}\right)$ for each $A_{i} \in \mathrm{Y}^{\text {top }}$, is called the candidate completion table of $T$ with respect to $T_{0}^{\prime}$.

Example 26. In Example 25, we have $\mathrm{Y}=\left\{A_{1}, A_{2}, B_{1}, B_{2}\right\}$ and $\mathrm{Y}^{\text {top }}=\left\{A_{2}, B_{2}\right\}$. Then completion table and candidate completion of $T$ with respect to $T_{0}^{\prime}$ are $T^{c t}=T \ltimes_{\mathrm{Y}} T_{0}^{\prime}=\emptyset$ and $T^{\text {cand }}=T \ltimes_{\mathrm{Y}^{\text {top }}} T_{0}^{\prime}=\left\{t_{1}, t_{4}\right\}$.

It is easy to see that $T^{c t} \subseteq T^{\text {cand }}$ for all schema augmentations tables $T$ of some table $T_{0}$. The completeness of merge queries can now be defined by comparing the candidate completion table with the completion table.

Definition 16 (Complete merge). Let $T$ be a schema augmentation to $T_{0}$ with respect to a relationship R having a set of common attributes Y . The merge result $T_{0}^{\prime}$ of $T_{0}$ and $T$ is said to be a complete merge with respect to $T$ if $T^{c t}=T^{c a n d}$.

Example 27. Continuing with Example 26, since $T^{\text {cand }} \neq$ $T^{c t}$, we can conclude that $T_{0}^{\prime}$ is not a complete merge with respect to $T$.

The following proposition states how it is possible to complete the result of an incomplete merge query.

Proposition 9. Let $T_{0}\left(\mathrm{~S}_{0}\right)$ and $T(\mathrm{~S})$ be two analytic tables with a relationship R and common dimension attributes Y. Let $T_{0}^{\prime}\left(\mathrm{S}_{0}^{\prime}\right)$ be the merge of $T_{0}$ and $T$ and $T^{\text {miss }}=$ $T^{\text {cand }}-T^{c t}$ be the set of all tuples in $T$ lost in the merge with $T_{0}$. We can define a completion table $T^{\text {com }}\left(S_{0}^{\prime}\right)$ for all dimensions $D_{i}$ of schema $\mathrm{S}_{D_{i}}$ in $T_{0}, 0 \leq i \leq n$, where $\left(\mathrm{S}_{D_{i}} \cap \mathrm{~S}\right) \subset\left(\mathrm{S}_{D_{i}} \cap \mathrm{~S}_{0}\right):$

$$
\begin{equation*}
T^{c o m}=\Pi_{\mathrm{S}_{0}^{\prime}}\left(T^{\text {miss }} \rtimes_{\mathrm{S}_{D_{0}} \cap \mathrm{Y}} D_{0} \searrow \rtimes \cdots \searrow_{\mathrm{S}_{D_{n}} \cap \mathrm{Y}} D_{n}\right) \tag{1}
\end{equation*}
$$

If $\mathrm{Y} \mapsto \mathrm{S}_{0}$ (non-ambiguous merge condition), then

$$
\begin{equation*}
Q_{m}\left(T_{0}, T\right)=T_{0}^{\prime} \cup T^{c o m} \tag{2}
\end{equation*}
$$

is a complete merge of $T_{0}$ and $T$ with respect to $T$.
In the above proposition, Table $T^{\text {miss }}$ identifies the tuples of $T^{\text {cand }}$ that are missing in $T_{0}^{\prime}$ to get a complete merge. Expression 1 augments $T^{\text {miss }}$ with all dimension attributes and values that exist in $T_{0}$ but not in $T$. This step avoids generating null values for these attributes and can be skipped if Y already contains all the dimension attributes in $T_{0}$. The non-ambiguous merge condition for Expression 2 checks if Y contains an identifier of $T_{0}$. If this condition is true and $T_{0}$ is not ambiguous, the merge result is also not ambiguous.

We can apply Proposition 7 to compute the aggregable property of the completed merge $Q_{m}\left(T_{0}, T\right)$. However, when the merge $T_{0}^{\prime}$ of $T_{0}$ with $T$ is not complete, for all aggregable attributes A where $\mathbf{a g g}_{\mathrm{A}}(F, V)$ holds in $T$ (A is an attribute obtained from $T$ ) the aggregable property $\operatorname{agg}_{\mathrm{A}}(F, \emptyset)$ holds in $T_{0}^{\prime}$ (A is not aggregable in $T_{0}^{\prime}$ ).

Example 28. Resuming Example 25, we obtain T ${ }^{\text {miss }}=$ $T^{\text {cand }}-T^{c t}=\left\{t_{1}, t_{4}\right\}$. Thus, using the notations of Definition 16, $T^{\text {com }}=\left\{t_{1}, t_{4}\right\}$ and $Q_{m}$ returns tuple $t_{5}, t_{6}$ augmented with $M=$ null and $A_{3}=$ null completed by tuples $t_{1}$ and $t_{4}$ from $T . Q_{m}$ is a complete merge with respect to $T$.

## 5. Architecture Overview in SAP HANA

The contributions presented in this paper have been implemented and published within the SAP HANA platform [15] as a REST application service. The extended HANA architecture is depicted in Figure 9. White boxes are new components that implement the algorithms described in Section 6 whereas grey boxes are existing HANA components that have been extended.


Figure 9: Architecture overview
Analytic tables are defined in SAP HANA as nonrecursive information views [16] over non-analytic tables and other information views using a set of relational operators (e.g., union, join, projection, aggregate). For example, all dimension and fact tables in Section 3.1 can be expressed as information views. Dimension tables can be either defined as arbitrary views over non-analytic tables or as project-union views over other dimension tables. Dimension attribute hierarchies are part of the dimension table definitions. The schemas of these information views, including the role played by dimension attributes or measure attributes as defined in Sections 2.1 and 2.2, are stored in the Metadata Catalog. We extended the metadata of information views by enabling a user to declare the aggregable properties of measure attributes and the attribute graphs of dimension tables. Default values are used for
missing aggregable properties (see Section 2.3). All these metadata are also stored in the Metadata Catalog (Agg Measures and Attribute Graph).

The SAP HANA platform provides a framework of schedulable metadata crawlers that asynchronously extract metadata from the definition of tables and information views [17]. Extracted metadata are periodically stored by the Metadata Loader in the Metadata Catalog. We extended the Metadata Loader as follows. When a new or updated dimension table is crawled, the Metadata Loader calls the Identifier Builder which uses its attribute graph to compute and store its dimension identifier. Similarly, the Identifier Builder is called to compute and store the fact identifier of new fact tables from the attribute graphs of their dimensions. In both cases, the Metadata Loader stores the attribute mapping and join relationships extracted from the information views. When new non-analytic tables are crawled, primary and foreign keys (PK/FK) are extracted from the Table schemas and new relationships are created by the Metadata Loader. The Metadata Loader also recursively composes the relationships and adds them to the set of relationships. If new relationships are derived for a pair of tables, they are recursively fused, removed from the set of relationships, and kept separately for maintenance. New relationships are passed to the SC Graph Builder that updates a Schema Complement (SC) Graph.

Our new service offers four main REST APIs. The Declare Metadata API enables designers to declare relationships between tables which are passed to the SC Graph Builder. It also supports the declaration of all metadata (attribute graph, aggregable attributes, and relationships) generated during the merge of an analytic table with a schema complement, which are then passed to the Identifier Builder and the SC Graph Builder accordingly. The Get Schema Augmentation API takes as input a start table name $T_{0}$ and returns a list of schema complements and augmentations $T^{\prime}$ with their relationships from $T_{0}$ to $T^{\prime}$. The Schema Complement Finder component processes a Get Schema Augmentation request using the SC Graph and the Table schemas (see Section 6.2). The Reduction Query Generator API takes an input a target schema augmentation table, and a set of user actions to reduce the identifier of the target table. It returns a reduction query and a description of aggregated attributes in the query. Heterogeneous units and currencies of aggregate reduction queries are managed using existing core functionalities of SAP HANA. Finally, the Merge Schema Augmentation API takes as input a start table and a target schema augmentation table, or a reduction query, and returns as output a final merge query with a list of computed metadata properties (see Section 6.4). The Schema Complement (SC) Manager uses a Reduction Query sub-manager to create a reduction query. In case of aggregate reduction, it guarantees that heterogeneous units and currencies are not aggregated together or alternatively converts all units and currencies to a common one before performing an aggrega-

Table 8: A tuple from hierarchy table HT

| Attribute | Value |
| :--- | :--- |
| NODE | [North America].[United States].[Ohio].[Dublin] |
| PRED_NODE | [North America].[United States].[Ohio] |
| NODE_VALUE | Dublin |
| ATT_NAME | CITY |
| IS_LEAF | 0 |
| LEVEL_NUM | 4 |

tion. The SC Manager uses a Natural Query sub-manager to compute a query for a path of natural schema complements, detect ambiguous results and create a complement query if needed. The Merge Query Generator is responsible for producing the final merge query. It calls the Natural Query component and the Metadata Propagator to compute the list of metadata properties associated with the augmented table using metadata propagation techniques presented earlier in Section 6.4.

## 6. Algorithms and REST Service Implementation

### 6.1. Dimension and fact identifiers

In this section we describe how to compute and maintain dimension and fact identifiers which play a central role in our work. We encode attribute graphs into two tables ATtribute Graph Node (ATGN) and ATtribute Graph Edge (ATGE). Their schema are respectively:

ATGN (DT_ID, ATT_NAME, LEVEL_NUM, OPTIONAL)
ATGE (DT_ID, ATT_NAME, PARENT_ATT_NAME, LABEL)
Each hierarchy in a dimension table is identified by a distinct DT_ID value. Table ATGN stores attribute nodes, with their level in the hierarchy and whether they can take null values (OPTIONAL $=1$ ). Table ATGE stores all edges in attribute graphs where LABEL can take values: ' + ', ' 1 ', ' f '. In both tables, attribute names with value '_-_bot' or '_-top' represent the two special attributes $\perp$ or $\top$ respectively.

Example 29. Table 9 shows tuples that encode the attribute graph of dimension WAREHOUSE in ATGN and ATGE.

When an attribute graph has not been defined by a user over the hierarchy of a dimension table, Algorithm 1 provides the option to efficiently compute it from a sample of the dimension table using the indexing scheme of SAP HANA, called Hierarchy Table (HT) [18], which encodes the nodes of a hierarchy instance represented in a dimension or a fact table.

Example 30. Node 'Dublin' of 'Ohio' in the hierarchy of REGION in Figure 3 is encoded as one tuple in HT (attribute names are on the left) shown in Table 8. It contains information about its path from the top-level node, its parent node, its attribute name, whether it is a leaf node, and its level number in the hierarchy.

```
Algorithm 1 Attribute Graph Generation (ATG)
    input: HT hierarchy table
    output: attribute graph
```

1. Initialization. All attributes in ATGN are initialized with OPTIONAL $=0$; edges having PARENT_ATT_NAME $=\top$ are initialized with LABEL $=$ $f$; edges having ATT_NAME $=\perp$ are initialized with LABEL $=+$.
2. Find all attributes (except bottom level attributes $\perp$ ) in HT that have 'IS_LEAF $=0$ ' and add an additional edge from $\perp$.
3. Update ATGN by marking all attributes that are nullable with OPTIONAL $=1$.
4. For each (ATt_NAME, PARENT_ATt_NAME) pair in ATGE, if for a NODE_VALUE value for the same attribute ATT_NAME, there are more than one PRED_NODE values for PARENT_ATT_NAME, then mark the edge with label +. If there is only one single value in PARENT_ATT_NAME then mark the edge with label f.
5. For each + labeled edge (ATt_NAME, PARENT_ATT_NAME) in ATGE where ATT_NAME is optional in ATGN, if for all non-null NODE_VALUE values for the same attribute ATT_NAME, there is only one PRED_NODE value for PARENT_ATT_NAME then update the edge with label 1.

Given an attribute graph over some attribute hierarchy, a dimension identifier for all valid tables with respect to that attribute graph, is computed using the CDI algorithm. The worst-case time complexity of this algorithm is linear in the size (number of nodes and edges) of attribute graph $\mathcal{D}$.

Example 31. We apply function $C D I$ on the attribute graph of REGION in Figure 4. dimId is initialized with all attributes of REGION. Only CONTINENT is removed, because it has a label f in-edge. In Line 16, the result dimension identifier dimid is \{CITY, STATE, COUNTRY\}.

By Proposition 2, the union of the identifiers of each group of dimension attributes occurring in the fact table defines a fact identifier of a fact table. However, this identifier might not be minimal. For example, when the fact table is defined as a view from a non-analytic table for which we know a key, if all attributes of this key have an attribute mapping relationship into a (strict) subset of the fact identifier attribute, this subset is also a fact identifier.

Example 32. Recall the fact table INVENTORY of Example 6. Suppose that this table is defined as a view over non-analytic table INVENTORY_ct with primary key $\mathrm{K}=\{$ PROD_SKU, YEAR, WH_ID, COUNTRY $\}$.
(a) ATGN

| DT_ID | ATT_NAME | LEVEL_NUM | OPTIONAL |
| :--- | :--- | :--- | :--- |
| 1 | _-bot | 0 | 0 |
| 1 | WH_ID | 1 | 0 |
| 1 | CITY | 2 | 0 |
| 1 | STATE | 3 | 1 |
| 1 | COUNTRY | 4 | 0 |
| 1 | -_top | 5 | 0 |

```
Algorithm 2 Compute Dimension Identifier (CDI)
    input:
        \(\mathcal{D}\) : Attribute graph of dimension table with
            schema S
        X: A subset of S
    output:
        dimId: Dimension identifier for S
    begin
        Initialize: \(\operatorname{dimId} \leftarrow \mathrm{X}\)
        for Attribute A in \(\operatorname{dimId}\) do
            if (A has one label \(\mathbf{f}\) in-edge from X in \(\mathcal{D}\) ) or
                (A only has label \(\mathbf{1}\) in-edges from X in \(\mathcal{D}\) )
                then
                \(\operatorname{dim} I d \leftarrow \operatorname{dim} I d-\{\mathrm{A}\}\)
            end if
        end for
        return dimId
    end
```


## INVENTORY_ct ( $\underline{P R O D_{-} S K U}, \underline{Y E A R}, \underline{W H_{-} I D}, \underline{C O U N T R Y}$, $\left.T A X_{-} N O, Q T Y_{-} O N_{-} H A N D, T A X_{-} A M T\right)$

Each attribute of the primary key has a corresponding attribute mapping relationship with one attribute of the initial fact identifier $\mathrm{K}^{\prime}=\left\{P R O D_{-} S K U, Y E A R\right.$, WH_ID, COUNTRY, TAX_NO \} computed by function CDI as given in Example 6. K is strictly included in $\mathrm{K}^{\prime}$ and can be used instead of K . It implicitly follows that dimension TAX depends on the other dimensions of that fact table.

An update of a dimension table may violate the attribute graph constraints defined for that table, and hence the dimension identifier generated from that attribute graph. Thus, when a non-analytic table is updated, we must check if the attribute graph is still satisfied and the dimension identifier is still valid in $T$. If the uniqueness test fails, the update must be rejected. Since dimension tables are not frequently updated, the uniqueness test entails a small overhead in the transaction processing workload. In addition, the uniqueness test can be efficiently implemented using Hierarchy Tables.
(b) ATGE

| DT_ID | ATT_NAME | PARENT_ATT_NAME | LABEL |
| :--- | :--- | :--- | :--- |
| 1 | _-bot | WH_ID | $'+'$ |
| 1 | WH_ID | CITY | $' \mathrm{f} '$ |
| 1 | CITY | STATE | $'+$, |
| 1 | WH_ID | STATE | $' \mathrm{f} '$ |
| 1 | STATE | COUNTRY | $' 1$ |
| 1 | CITY | COUNTRY | $'+$, |
| 1 | COUNTRY | --top | $' \mathrm{f} '$ |

### 6.2. Computing schema augmentations and complements

The Metadata Catalog contains the native relationships that are extracted from the definition of analytic tables or explicitly declared by the user, and the derived relationships obtained using composition and fusion. From these relationships, schema augmentations and complements are discovered by exploring a Schema Complement graph defined as follows.

Definition 17 (Schema Complement graph). Let $\mathcal{T}=$ $\left\{T_{0}, T, \cdots, T_{n}\right\}$ be a set of tables and $\mathcal{R}=\left\{\mathrm{R}_{j k}, \cdots, \mathrm{R}_{n m}\right\}$ be a set of relationships between tables in $\mathcal{T}$, where $\mathrm{R}_{j k} \in$ $\mathcal{R}$ is a relationship between $T_{j}$ and $T_{k}$. A Schema Complement (SC) graph for $\mathcal{T}$ and $\mathcal{R}$ is a directed property graph $\mathcal{S C}=(\mathcal{T}, \mathcal{E})$ connecting the tables in $\mathcal{T}$ by a set of labeled edges $\mathcal{E}$ where:

1. for each $\mathrm{R}_{j k} \in \mathcal{R}$ there exist two edges $\mathcal{E}\left(T_{j}, T_{k}\right)$ and $\mathcal{E}\left(T_{k}, T_{j}\right)$.
2. each edge $\mathcal{E}\left(T_{j}, T_{k}\right)$ is labeled by a complement type E.CT: if the common attributes of $\mathrm{R}_{j k}$ contain the identifier of $T_{k}$ then, $\mathcal{E} . C T={ }^{\text {NAT }}{ }^{\prime}$ (natural edge) else $\mathcal{E} . C T=$ 'AUG' (augmentation edge).


Figure 10: SC graph example for Figure 8

Example 33. Figure 10 shows a partial SC graph for the relationships of Figure 8. Each edge is labelled with a reference to a relationship and its CT label. Nodes represent tables with their identifiers (ID).

The Get Schema Augmentation API is implemented by the function CSA in Algorithm 3. Given a table $T_{0}$,

```
Algorithm 3 Compute Schema Augmentation (CSA)
    input:
        \(T_{0}\) : start node (table)
        SC: schema complement graph
    output:
        result: \([(T, C T\), newAttrsrs, rid \()]:\)
        \(-T\) : destination table
        - \(C T\) : complement type
        - newAttrs: new attributes
        - rid: relationship with destination \(T\)
        altEdges: \([(T, C T\), rid \()]\) : list of optional paths
    begin
        result and altEdges are empty
        mark \(T_{0}\) as visited
        for all edges \(e\left(T_{0}, T_{i}\right)\) in \(\mathcal{S C}\) do
            if \(T_{i}\) is not visited then
                    mark \(T_{i}\) as visited
                newAttrs \(\leftarrow\) new attributes in \(T_{i}\) w.r.t. \(T_{0}\)
                add ( \(T_{i}\), e.CT, newAttrs, e.rid) to result
                if \(e . C T=\) 'NAT' then
                    add \(\operatorname{CSA}\left(T_{i}, S C\right)\) to (result, altEdges)
                end if
            else
                add ( \(T_{i}, e . C T\), e.rid) to altEdges
            end if
        end for
        return (result,altEdges)
    end
```

the function CSA explores a possibly cyclic SC graph and returns an array of schema augmentations for $T_{0}$. Lines 13 and 16 avoid infinite loops during traversal by marking visited nodes. Also, if multiple paths exist between $T_{0}$ and a schema complement $T$, a single edge from $T_{0}$ to $T$ is added to result and alternative edges are added to altEdges. Thus, a single path of relationships from $T_{0}$ to $T$ is built in result, which depends on the edge $\mathcal{E}\left(T_{0}, T\right)$ selected at the first iteration. Line 20 recursively continues the exploration for ' $N A T$ ' edges, possibly adding new arcs to result. Line 17 compares the attributes of $T_{0}$ and $T$ as follows. If $T_{0}$ has never been augmented with schema augmentations before, any attribute of $T$ that is not in common with $T_{0}$ is considered to be new. Otherwise, the attributes of $T_{0}$ that come from previous schema augmentations with $T$ have metadata that indicate their provenance from $T$ and newAttrs is empty. In both cases, a new entry is added to result to construct the full path from $T_{0}$ to a subsequent schema augmentation $T$.

Our implementation of the CSA algorithm leverages the SAP HANA graph engine [19] which supports various graph algorithms over graph data stored in a columnar table storage (using a table of nodes and a table of edges). First, we use the GET_NEIGHBORHOOD algorithm of the graph engine to compute the set Reachable of all tables that are reachable from $T_{0}$ through a sequence
of natural ' $N A T$ ' edges followed by at most one augmentation ' $A U G$ ' edge. We use the term connected for nodes which are connected by any path without taking account of this restriction. The GET_NEIGHBORHOOD performs a breadth-first search and returns all reachable tables. Second, we select all edges whose start node and end node are in: $T_{0} \cup$ Reachable and order them by respecting the ascending depth with respect to $T_{0}$. These edges are then annotated with their set of new attributes which is retrieved using SQL queries over the table schemas in the Metadata Catalog. Finally, the resulting $\mathcal{S C}$ graph is built and traversed in the main-memory of the application server engine. Given the edge ordering in the $\mathcal{S C}$ graph according to their depth with respect to $T_{0}$, algorithm CSA always returns the shortest path to a given table in Result. The use of a clever exploration strategy is left open for future work. However, in Result, schema augmentations and complements are returned sorted according to the depthfirst search exploration of the $\mathcal{S C}$ graph.

Example 34. We illustrate our iterative implementation of algorithm CSA with input $T_{1}$ on Figure 10. First, CSA computes all tables which are reachable from $T_{1}$ in SC by a sequence of NAT edges and followed by at most one AUG edge. We obtain, Reachable $=\left\{T_{2}, T_{4}, T_{3}\right\}$. The next step retrieves the list of edges connecting nodes in $T_{0} \cup$ Reachable, ordered by their distance from $T_{1}$ : Edges $=\left[\mathcal{E}\left(T_{1}, T_{2}\right), \mathcal{E}\left(T_{1}, T_{4}\right), \mathcal{E}\left(T_{1}, T_{3}\right), \mathcal{E}\left(T_{3}, T_{2}\right)\right]$. The algorithm then processes each edge $\mathcal{E}\left(T_{k}, T_{i}\right)$ in Edges as follows. If the target table $T_{i}$ has not been visited yet, the set of new attributes newAttrs ${ }_{i}$ is computed for $T_{i}$ and the edge is added to Result. All other edges are added to altEdges. In our example, this produces $\left(T_{2}, A U G\right.$, new Attrs $\left._{2}, \mathrm{R}_{12}\right)$ for edge $\mathcal{E}\left(T_{1}, T_{2}\right),\left(T_{4}, A U G\right.$, newAttrs $\left.s_{4}, \mathrm{R}_{14}\right)$ for edge $\mathcal{E}\left(T_{1}, T_{4}\right)$ and $\left(T_{3}\right.$, NAT, newAttrs $\left._{3}, \mathrm{R}_{13}\right)$ for edge $\mathcal{E}\left(T_{1}, T_{3}\right)$. For edge $\mathcal{E}\left(T_{3}, T_{2}\right)$, table $T_{2}$ is marked as visited and ( $T_{2}, A U G, \mathrm{R}_{23}$ ) is added to altEdges.

Suppose now that ( $T_{2}, A U G$, newAttrs, $\mathrm{R}_{12}$ ) is selected for augmenting the schema of table $T_{1}$, resulting in a new table $T_{1}^{\prime}$. Then there might appear new tables that are reachable from $T_{1}^{\prime}$ in the graph of Figure 10, which were not reachable from the original table $T_{1}$. For example $T_{6}$ is only accessible through two AUG edges from $T_{1}$, but by a single AUG edge from $T_{1}^{\prime}$. However, we can show that $T_{1}^{\prime}$ will never connect to more tables than to those that are already connected to $T_{1}$ and $T_{2}$ in $\mathcal{S C}$. In other terms, all schema augmentations reachable from a schema merge table $T_{1}^{\prime}$ in $S C$ are also connected to the original table $T_{1}$ in $S C$.

### 6.3. Creating a reduction query

The Compute Schema Augmentation (CSA) API returns a set of schema augmentations for some source table $T_{0}$. Assume that a user selects a schema augmentation table $T$ with complement type 'AUG' and relationship path $\mathrm{R}_{1}\left(T_{0}, T_{1}\right), \cdots, \mathrm{R}_{n}\left(T_{n-1}, T\right), n>=1$ as a schema augmentation for $T_{0}$. Let $\mathrm{Y}_{n}$ be the set of common attributes of
$T_{n-1}$ and $T$, and K be the identifier of $T$. Then, a natural schema complement to $T_{0}$ can be obtained by reducing all attributes in $\mathrm{K}-\mathrm{Y}_{n}$ through one or several reduction queries. Table 10 shows the reduction actions that a user can express on each attribute A of $\mathrm{K}-\mathrm{Y}_{n}$ and their impact in terms of reduction queries. Without loss of generality, we consider here that each user action reduces a single attribute of $\mathrm{K}-\mathrm{Y}_{n}$.

Table 10: User actions to create a reduction query

| User action | Impact on reduction query |
| ---: | :--- |
| Filter | defines a set FilP of filter predicates $\mathrm{A}_{i}$ <br>  <br> $v_{i}, v_{i} \in \operatorname{dom}\left(\mathrm{~A}_{i}\right) ;$ yields a filter reduction; |
| Pivot | defines a set of attributes PivA that are pivoted <br> as columns; yields a pivot reduction; |
| Remove | defines a set of attributes RemA which are re- <br> moved from the result of the reduction query; |
|  | yields an aggregate reduction; |
| Aggregate | defines a set AggA of aggregated attributes <br>  <br> $F_{i}\left(\mathrm{~A}_{i}\right) ;$ yields an aggregate reduction |

We now introduce the Reduction Query Generation (RQG) API which takes five inputs: ( $T$, FilP, PivA, $\operatorname{Rem} A, \operatorname{Agg} A)$, and returns two outputs as described below:

```
Algorithm 4 Reduction Query Generation (RQG)
    input:
            T: destination schema augmentation table
        FilP: a set of equality predicates of the form A}=
            for filtered attributes A \in K
        PivA: a set of pivoted attributes
    RemA: a set of attributes removed from K
        AggA: a set of aggregated attributes of the form
            F(A) where A }\in\textrm{S
    output:
        Q: reduction query
    Agg\mp@subsup{A}{}{\prime}:\quad aggregated attributes of the form F(A) in Q
```

Verify inputs and order reduction operations. The reduction process is separated into three optional reduction steps whose execution depends on the user input. Filter and Pivot are executed, respectively, when FilP and $\operatorname{Piv} A$ are not empty. Aggregate is executed when RemA or $\operatorname{Agg} A$ is not empty. The default ordering of reduction operations is: 1) Filter, 2) Pivot, 3) Aggregate.

Step 1: Apply filter reductions. A filter reduction query $Q_{\text {Fil }}=\mathcal{F i l t e r}_{T}($ FilP $)$ on the target schema augmentation table $T$ is created and executed.

Step 2: Apply pivot reductions. A pivot reduction query $Q_{\text {Piv }}$ is generated over the result of query $Q_{\text {Fil }}$. It is of the form $\operatorname{Pivot}_{Q_{F i l}}\left(\mathrm{X}_{\text {val }} \mid \operatorname{PivA}\right)$, where $\mathrm{X}_{\text {val }}$ is the set of aggregable attributes in $T$.

Step 3: Apply aggregate reductions and detect if the result is ambiguous. First, each aggregate $F_{i}\left(\mathrm{~A}_{i}\right)$ in $\operatorname{Agg} A$ is checked for correctness with respect to the aggregable properties of the attribute $\mathrm{A}_{i}$ and attributes $\operatorname{Rem} A$. Then, an aggregate query $Q_{\text {Agg }}$ is generated over the result of query $Q$ generated in the previous Filter and Pivot reduction steps. It is of the form: $\mathcal{A g g}_{Q}\left(F_{i}\left(\mathrm{~A}_{i}\right), \ldots \mid \mathrm{X}\right)$, where each $F_{i}\left(\mathrm{~A}_{i}\right)$ belongs to $\operatorname{Agg} A$ or to the new attributes whose values were pivoted in $Q$, and $\mathrm{X}=\mathrm{K}-\operatorname{Rem} A-\operatorname{Piv} A$.

Generate reduction query. The output reduction query $Q$ is generated by the previous Filter, Pivot and Aggregate reduction steps. $\operatorname{Agg} A^{\prime}$ is the set of aggregable attributes containing attributes from input $\operatorname{Agg} A$ and attributes whose values were pivoted in the pivot reduction query step (i.e., the $\mathrm{X}_{\text {val }}$ attributes). This set is needed for the propagation of aggregable properties (see Section 4.1). The identifier $\mathrm{K}^{\prime}$ of the result of $Q(T)$ is $\mathrm{K}^{\prime}=\mathrm{K}-$ FilA $-\operatorname{Rem} A-\operatorname{PivA}$ where K is the identifier of $T, \operatorname{Rem} A$ and $\operatorname{Agg} A$ are defined in Table 10 and FilA corresponds to the attributes in the filter predicates FilP.

Example 35. Consider $\mathrm{T}_{1}=$ SALES is a schema augmentation to $\mathrm{T}_{2}=$ INVENTORY in the SC graph of Figure 10. When augmenting the schema of $\mathrm{T}_{1}$ with $\mathrm{T}_{2}$, a reduction query can be expressed as an $R Q G$ call: $R Q G\left(T_{1}, \emptyset, \emptyset, \operatorname{Rem} A,\{S U M(A M O U N T)\}\right.$, where Rem $=\{$ MONTH, STATE $\}$. Because PivA and FilP are empty, a single aggregate reduction is executed: $\quad Q_{\text {Agg }}=\mathcal{A g g}_{T_{1}}(S U M(A M O U N T) \mid \mathrm{X}), \quad \mathrm{X}=$ \{PROD_SKU, YEAR, CITY, COUNTRY\}. Indeed, X is a subset of the identifier of $\mathrm{T}_{1}$ and attributes MONTH and STATE have been reduced. Thus, the RQG call finally returns a reduction query $\mathcal{Q}_{\text {Agg }}$ and $\{S U M($ AMOUNT $)\}$ as a single aggregated attribute.

Example 36. An alternative way to reduce SALES when augmenting the schema of INVENTORY is to use an $R Q G$ call: $\operatorname{RQG}\left(T_{1},\{\right.$ STATE $=$ 'Ohio' $\},\{$ MONTH $\left.\}, \emptyset, \emptyset\right)$ on the SC graph of Figure 10. Because AggA and RemA are empty, only filter and pivot reductions are applied. A filter reduction is created in step 1 as $Q_{\text {Fil }}=$ $\mathcal{F i l t e r}_{T_{1}}(\{$ STATE $=$ 'Ohio' $\})$. In step 2, the pivot reduction is applied on the result of $Q_{\text {Fil }}$ as $Q_{\text {Piv }}=$ $\mathcal{P i v o t}_{Q_{F i l}}(\{A M O U N T\} \mid\{M O N T H\})$. Thus, the $R Q G$ call returns a query $Q_{\text {Piv }}$, and an empty set of aggregated attributes.

### 6.4. Computing an augmented schema

Suppose that a user selects a table $T_{\text {dest }}$ returned by the Compute Schema Augmentation API call as a schema augmentation to a table $T_{0}$ with respect to a path of relationships $\mathrm{R}_{1}\left(T_{0}, T_{1}\right), \cdots, \mathrm{R}_{n}\left(T_{n-1}, T\right), n>=1$. We now introduce the Merge Schema Augmentation (MSA) API, which takes five inputs ( $T_{0}, T$, path, noamb, comp) and returns one merge query $Q$ as described below. The input table $T$ is defined as follows. When the chosen schema augmentation $T_{\text {dest }}$ has not been reduced, the input $T$ is equal to
$T_{\text {dest }}$ with a set of new attributes $\mathrm{X}_{\text {new }}$ and $\mathrm{R}_{n}\left(T_{n-1}, T\right)=$ $\mathrm{R}_{n}\left(T_{n-1}, T_{\text {dest }}\right)$ maps to an edge in the schema complement graph $\mathcal{S C}$. Otherwise, $T=Q_{\text {red }}\left(T_{\text {dest }}\right)$ represents the reduction query $Q_{\text {red }}$ over $T_{\text {dest }}$ obtained as an output of the Reduction Query Generation (RQG) API call, and $\mathrm{R}_{n}\left(T_{n-1}, T\right)$ corresponds to a natural schema complement edge ( $C T=$ ' $N A T$ '). The Boolean parameter noamb is true when no ambiguous value must appear in a merged result and the Boolean parameter comp is true when the merge must be complete.

```
Algorithm 5 Merge Schema Augmentation (MSA)
    Inputs:
            \(T_{0}\) : starting table
            \(T\) : destination table \(T_{\text {dest }}\) or a reduction query
                        \(Q_{\text {red }}\)
        \(\mathrm{X}_{\text {new }}\) : attributes in \(T\) that will be merged to \(T_{0}\)
        path: \(\quad\left\{\mathrm{R}_{1}\left(T_{0}, T_{1}\right), \cdots, \mathrm{R}_{n}\left(T_{n-1}, T\right)\right\}\)
    noamb: true when result must be non-ambiguous
        comp: true when merge must be complete
    Output:
            \(Q: \quad\) final merge query
```

An $M S A$ call is processed in three steps detailed thereafter.

Step 1: Create query $Q_{a}$ to update ambiguous tuples in $T$. When flag noamb $=$ false, this step does not check for ambiguous results and returns $Q_{a}=T$. When flag noamb $=$ true, we first check if $T$ is non-ambiguous. When $T$ is not ambiguous for each dimension $D$ in $T$ w.r.t. Proposition 8, then we return $Q_{a}=T$. Otherwise, by Definition 14, detecting if $T$ contains ambiguous values requires computing for each dimension $D$ used in $T$, the identifier of the ancestors $\mathrm{X}^{*}(D)$ of all key attributes in $T$ from dimension $D$. This identifier is computed over the attribute graph of dimension $D$ restricted to the attributes in $\mathrm{X}^{*}(D)$ and using Proposition 1. When $T$ is detected as possibly ambiguous with respect to a dimension $D$, we build a query $Q_{a}$ which replaces the measure values of all ambiguous tuples in $T$ by null values. A tuple $t$ in $T$ is detected as ambiguous if there exists a dimension $D$ and two tuples $t_{1}, t_{2} \in D$ such that $t_{1} \cdot \mathrm{X}_{D}^{*} \not \equiv t_{2} \cdot \mathrm{X}_{D}^{*}$ and $t_{1} \cdot \mathrm{X}_{D} \equiv t_{2} \cdot \mathrm{X}_{D} \equiv t \cdot \mathrm{X}_{D}$. For this, query $Q_{a}$ joins the result of $T$ with each dimension table $D$ on their common attributes $\mathrm{X}_{D}$ to obtain all $\mathrm{X}_{D}^{*}$ attribute values, and nullifies (invalidates) all measure attributes of all tuples $t$ where the size of the partition corresponding to $t_{\mathrm{X}_{D}}$ is greater than 1.

Step 2 (optional): Create a query $Q_{c}$ to compute a completion table for the merge of $T_{0}$ with $T$. When flag comp $=$ true, the we apply the steps of Proposition 9 to build a query $Q_{c}$ which computes a completion table $T^{c o m}$. Let Y be the set of common attributes between $T_{0}$ and $T$ (common attributes are obtained by composition of the relationships of path). If the LFD $\mathrm{Y} \mapsto \mathrm{S}_{0}$ holds in $T_{0}\left(\mathrm{~S}_{0}\right)$,
then the non-ambiguous merge condition of Proposition 9 is satisfied and a query $Q_{c}$ is generated to create the completion table $T^{\text {com }}$ that completes the merge of $T_{0}$ with $Q_{a}$ of Step 1.

Step 3: Create the final query merging $T_{0}$ with the results of Steps 1 and 2. First, a query $Q_{\text {path }}$ is created to merge $T_{0}$ with the sequence of natural schema complements $T_{1}, \cdots, T_{n-1}$ represented by $\mathrm{R}_{1}\left(T_{0}, T_{1}\right), \cdots$, $\mathrm{R}_{n-1}\left(T_{n-2}, T_{n-1}\right)$ in parameter path. Let $\mathrm{Y}_{i}$ be the common attributes in relationship $\mathrm{R}_{i}$. Then, $Q_{\text {path }}=T_{0} \triangle_{\mathrm{Y}_{1}}$ $T_{1} \cdots \searrow_{\mathrm{Y}_{n-1}} T_{n-1}$. Next, $Q_{\text {path }}$ is merged with $Q_{a}$ and we obtain $Q^{\prime}=\pi_{S_{0}^{\prime}}\left(Q_{\text {path }} \triangle \rrbracket_{\mathrm{Y}_{n}} Q_{a}\right)$, where $\mathrm{S}_{0}^{\prime}$ is the schema of $T_{0}$ augmented with the new attributes $\mathrm{X}_{\text {new }}$ coming from $Q_{a}$. Finally, if comp $=$ true, then $Q_{c}$ (obtained from Step 2) will be added to the previous result $Q^{\prime}$ and the final result is $Q=Q^{\prime} \cup Q_{c}$. Otherwise the final result is $Q=Q^{\prime}$.
Example 37. Consider a merge schema augmentation $\operatorname{MSA}\left(\mathrm{T}_{2}, \mathcal{Q}\left(\mathrm{~T}_{1}\right),\{\operatorname{SUM}(\right.$ AMOUNT $)\},\left\{\mathrm{R}_{12}\right\}$, true, true) on the SC graph of Figure 10 ( $\mathrm{T}_{1}=$ SALES is a schema augmentation to $\mathrm{T}_{2}=$ INVENTORY $)$, and $Q\left(\mathrm{~T}_{1}\right)$ is the output of the $R Q G$ call in Example 35. Since $Q\left(\mathrm{~T}_{1}\right)$ is a reduction query, we add a natural complement edge $R\left(\mathrm{~T}_{2}, \mathcal{Q}\left(\mathrm{~T}_{1}\right)\right.$ (with $C T=$ ' $N A T$ ') to the schema complement graph.

- Step 1: The destination table $Q\left(\mathrm{~T}_{1}\right)$ is a reduction query and may contain ambiguous values. We must compute all ambiguities over attribute PROD_SKU from dimension PROD, attributes CITY and COUNTRY from dimension STORE and attribute YEAR from dimension TIME. This step generates a query $Q_{a}$ which joins $Q$ with dimension table STORE and sets all aggregated measure attributes of ambiguous tuples $t \in Q\left(\mathrm{~T}_{1}\right)$ to null. For example, suppose that STORE contains two tuples $t_{1}$ and $t_{2}$ with the same values for attributes CITY and STATE and different REGION values, e.g. (Dublin, Ohio, UnitedStates) and (Dublin, California, UnitedStates). Query Q aggregates these tuples into a single ambiguous tuple $t$, and $Q_{a}$ sets the aggregated $\operatorname{SUM}(A M O U N T)$ value in $t$ to null
- Step 2: The set of common attributes between $T_{2}$ and $\mathcal{Q}\left(T_{1}\right)$ is $\mathrm{Y}=\left\{P R O D_{-} S K U\right.$, YEAR, CITY, COUNTRY $\}$. The non-ambiguous merge condition $\mathrm{Y} \mapsto A$ in Proposition 9 holds for all attributes of $T_{2}$ and a completion query $Q_{c}=$ $T^{c o m}$, as defined in Proposition 9, is created to generate tuples that complete the merge of $T_{2}$ with $Q\left(\mathrm{~T}_{1}\right)$.
- Step 3: Table $T_{2}$ is connected to $Q_{a}$ by a natural schema complement edge $R\left(\mathrm{~T}_{2}, Q_{a}\right)$ and $Q_{\text {path }}=T_{2}$. Then $\left.Q=\pi_{\mathrm{S}^{\prime}}\left(T_{2} \searrow_{\mathrm{Y}} Q_{a}\right)\right)$ ) is a natural merge, where $\mathrm{S}^{\prime}$ contains the schema of $T_{2}$ augmented with attribute $\operatorname{SUM}(A M O U N T)$, and the final result is $Q \cup Q_{c}$.

The following proposition states that the RQG and MSA API compute a correct merge result.

Proposition 10. Let query $Q$ be the result of a $M S A\left(T_{0}\right.$, $T$, path, true, true) call with a start table $T_{0}\left(\mathrm{~S}_{0}\right)$, a target schema augmentation table $T(\mathrm{~S})$ and a path of relationships: path $=\mathrm{R}_{1}\left(T_{0}, T_{1}\right), \cdots, \mathrm{R}_{n}\left(T_{n-1}, Q_{a}\right), n>=1$. If
$\mathrm{R}_{n}\left(T_{n-1}, Q_{a}\right)$ maps to an augmentation schema complement edge ( $C T=$ ' $A U G$ '), then $Q$ computes a augmented merge of $T_{0}$ with $Q_{a}$ (without ambiguous values). Otherwise, $\mathrm{R}_{n}\left(T_{n-1}, Q_{a}\right)$ maps to a natural schema complement edge ( $C T=$ ' $N A T$ ') and $Q$ computes a natural merge $T_{0}$ with $Q_{a}$.

## 7. Related Work

The notion of natural schema complement was first proposed by [8] for web tables (tables extracted from web page) where join relationships between tables are discovered by computing schema matching similarities [20] using a combination of similarity in attribute names, data types, and values (through a variant of Jaccard similarity). In the context of web tables, [9] presents direct and holistic schema matching techniques to discover natural schema complements for a given web table and user-specified attribute names. Similarly, Data civilizer [10] uses schema matching techniques to build a linkage graph between arbitrary datasets for suggesting natural schema complements. Our solution differs firstly by leveraging more "trusted" relationships which are extracted from user defined analytic fact and dimension tables (views). Join relationships are enriched by attribute mapping relationships to express mappings extracted from view definitions relating view attributes and their corresponding source table attributes. However, regardless of these particular features, our framework does not exclude the use of other schema matching techniques, including recent ones like [21], to discover new join relationships between heterogeneous dimensions (e.g., like between STORE and WAREHOUSE in the examples of this paper). Secondly, we extend previous schema complement models with the notion of reduction-based schema complements which come with new data quality issues. We formally define these issues and show how to check and solve them.

Another notion similar to schema complements is "drilling-across" fact tables related through "conformed dimensions" $[1,2]$, i.e., through dimensions which are either identical or where one is a subset of the other. Drillacross queries perform left or full outer joins operations between fact tables on their common conformed attributes. Our solution generalizes this notion by handling the more general case of merge queries between dimension tables and between fact and dimension tables. To support "drillacross" queries through two related but non-conformed dimensions, [22] proposes to build the intersection or union of these dimensions provided that they are "compatible", i.e., they can be related by a perfect matching that is coherent, sound, and consistent. Our model is more general because it relaxes each of these constraints on dimensions. Indeed, relaxation of coherence is managed by our definition of a merge query which accepts the duplication of common attributes in a merge. Soundness and consistency constraints are respectively handled by our definition of complete merge, and the detection of ambiguous
tuples in the merge results. Thus, our solution can find a schema augmentation over incompatible dimensions to a given fact table and compute a merge corresponding to a drill-across query that would not be accepted by [22]. Finally, although the idea of using arbitrary relationships between tables to express general drill-across queries has early been introduced in [23], no well-defined and implemented method to support such queries was published before. To the best of our knowledge, our solution is the first implementation that addresses the use cases presented in [23].

The quality of aggregate queries over fact tables has extensively been studied in the literature through the notion of summarizability defined as the "correct computation of aggregate values with a coarser level of detail from values with a finer level of detail" [13]. Summarizability was introduced in [24] and various conditions to ensure summarizability have been proposed [25, 26, 27, 28]. We adapt these results to the context of schema complements and the use of SQL semantics with nulls to characterize, for example, ambiguous aggregable attributes in the presence of non-strict hierarchies. We also leverage and extend previous work on additive measures [14, 2] to define applicable aggregate functions before a merge, and propose techniques to infer which aggregate functions are applicable after a merge. This enables us to propagate metadata on analytic tables across schema augmentations. Our results go beyond summarizability because, as with the approach suggested in $[29,25]$, we tolerate inaccurate measure values in a merge result, but clearly identify theses values and control their usage through their aggregable properties or the generation of null values. In addition, our method of computing complete merges is novel and was not considered in previous works on summarizability.

Dimension identifiers are essential for checking the coherence of analytic operations, and especially for verifying the semantic correctness of joins between analytic tables [1, 2, 22]. The work presented in [30] introduced the constraint that the bottom-level dimension attributes determine all other attributes in the dimension hierarchy. A similar kind of user-defined dimension identifiers are surrogate keys [2], which correspond to "artificial" bottom-level key attributes like CUSTOMER_ID or TIME_ID. An analogous approach has been introduced in [28] by enforcing linear dimension hierarchies, where each attribute has at most one parent attribute. On the one hand, the previous constraints and definitions simplify the problem of determining the dimension identifiers and to avoid the definition and analysis of LFD dependencies between dimension attributes. On the other hand, the trade-off of this simplification is the incapacity to infer new dimension identifiers for the result of operations which reduce the number of dimension attributes. This also restricts the set of composable operations to those which preserve these surrogate keys.
Example 38. Table SALES in Figure 8 contains attributes

MONTH, YEAR from dimension TIME. Suppose that the bottom-level attribute DAY is the surrogate key (dimension identifier) of dimension TIME. Then, without additional information, we cannot determine the new fact identifier of SALES, as we do in Example 6 by using the corresponding attribute graphs.

## 8. Experimental Evaluation

### 8.1. Attribute Graph and Dimension identifier Computation

All performance experiments for attribute graph generation (ATG) and dimension identifier computation (CDI) are conducted on an SAP HANA instance with 250 GB of main memory and 300 GB of disk space.

### 8.1.1. Computation of an attribute graph

Our first experiment consists in testing the performance of Algorithm 1 (ATG) for the generation of attribute graphs. We generate several hierarchies (dimension tables) and compare the performance of building attribute graphs according to four parameters as shown in Figures 11 and 12: (a) the size of the input dimension table (rows); (b) the number of hierarchy nodes / dimension attributes (without $\perp$ and $T$ ); (c) the number of +-edges in the resulting attribute graph; (d) the total number of edges in the resulting attribute graph.

The first experiments concern the generation of attribute graphs for linear strict hierarchies. The resulting attribute graphs for $n$ level types then have exactly $n$ nodes and $n-1$ f-edges (no + and 1 -edges). To achieve this, we generate for each dimension attribute a random number in the interval [ 0,9999 ], and concatenate this value with its parent attribute value. For example, some dimension table $T_{1}$ with attributes L1 $\preccurlyeq$ L2 $\preccurlyeq$ L3 $\preccurlyeq$ L4 $\preccurlyeq$ L5 might contain a tuple (L1:‘63.789.266.7426.2629', L2:‘63.789.266.7426', L3:‘63.789.266', L4:'63.789', L5:'63'). The goal of this process is to produce only $\mathbf{f}$-edges by ensuring that L1 $\mapsto \mathrm{L} 2 \mapsto \mathrm{~L} 3 \mapsto \mathrm{~L} 4 \mapsto \mathrm{~L} 5$. Figures 11a and 11b show that computation time increases linearly with the input table size ( 5 nodes/attributes) and with the number of nodes / attributes ( 10 k and 50 k rows) for these linear strict hierarchies.

The goal of the following experiments is to show the performance evolution w.r.t. to the number of additional + and 1-edges. These edges are created by ignoring the parent attribute value and allowing null values during the generation of an attribute value (multiple parents). The dimension tables used in the experiments contain 10 K rows. Figure 12a shows the performance evolution in a hierarchy over 10 attributes for additional +-edges generated by nonstrict hierarchies without null values. Figure 12b shows the performance evolution for hierarchies with other additional edges produced by allowing null values during the hierarchy generation process. We can see that the computation time slightly decreases in both cases with the number of
additional edges. For hierarchies of less than 10 nodes and samples of less than 50 K rows, which corresponds to a large number of real use cases, the computation of an attribute graph takes less than 20 seconds.

Finally, for a dimension table with a very large sample size of 1.1 M rows and 5 nodes, the computation of the attribute graph takes 67 seconds, which is still acceptable for an occasional system background call.

### 8.1.2. Computation of dimension identifiers

Algorithm 2 (CDI) takes an attribute graph $\mathcal{D}$ and a set of attributes X as input for computing the dimension identifier $Y$ of $X$, such that $Y \mapsto X$. If $X$ is the set of all attributes of a dimension table $D, \mathrm{Y}$ is the dimension identifier of $D$. We tested 20 different attribute graphs with 5 to 20 nodes and 4 to 28 edges. All dimension identifiers of these attribute graphs are computed in less than 9 milliseconds with a maximal variation of 2 milliseconds.

### 8.1.3. Comparison with related work

Inferring functional dependencies (FD) from relations is an old problem and many efficient algorithms have been proposed in the literature (see [31] for a comparison of 7 representative FD detection algorithms). These algorithms traditionally adopt NULL-NOT-EQ or NULLEQ [32] semantics for null values in the dependent attribute $A$ of a functional dependency $\mathrm{X} \rightarrow \mathrm{A}$. A third approach is to apply probabilistic semantics to generate approximate FDs [32]. In our data model, we consider NULL-EQ (literal equality) semantics and we also support null values in the determinant attribute set X. Under NULL-EQ semantics, we then could apply an appropriate skolemization function which transforms all null values a dimension table into skolems by fully respecting the existing hierarchy. Then, any existing FD detection algorithm could be used to detect all LFD dependencies with NULLEQ semantics.

However, there exists a second important difference between our setting and the standard relational setting. As it can be seen in the definition of attribute graphs, we only consider functional dependencies $\mathrm{X} \rightarrow \mathrm{A}$ where X and A appear in the same dimension table and all determinant attributes in X are descendants of the depending attribute A in the corresponding hierarchy type. This reduces the solution space that has to be explored and enables high pruning by leveraging the attribute hierarchy during the detection process. As our experiments show, the cost increases linearly with the number of attributes opposed to exponential increase for general FD algorithms [31, 33]. Thus, our solution is more efficient than general purpose algorithms for the discovery of identifiers in the limited case of dimension tables. Finally, the representation of LFD dependencies in the form of attribute graphs expresses known properties of hierarchies and is more natural to understand by a domain expert.


Figure 11: Attribute graph computation time for linear strict hierarchies (f-edges)


Figure 12: Computation time for non-strict hierarchies

### 8.2. Business intelligence scenario

Our second series of experiments checks the usage of the schema augmentation REST service, including the CSA, RQG, and MSA APIs described in Section 6, in a business intelligence setting. For this, we chose a concrete application of a worldwide clothing company which performs stock analysis in retail stores and customer segmentation. The application contains 281 non-analytic tables used for defining 42 information views representing dimension tables, and 145 carefully optimized information views representing fact tables. We classified th efact tables into two main categories: (a) 22 first-level views defined by a star-join between a non-analytic table storing facts (possibly completed with left-outer joins with other tables providing details) and dimension tables, and (b) 72 join views defined by a set of joins, aggregations and filters over fact, dimension and non-analytic tables. The remaining fact tables are defined using union and simple aggregations over other fact tables.

### 8.2.1. Experiments

In the first experiment, our goal is to select a join view, also called Hand-Crafted View, denoted by HCV, as a target and verify if we can generate an equivalent view, denoted by $G V$ for Generated View), by iteratively extending a first-level start view (table) through possibly several
schema augmentation steps, each of which consisting of a sequence of CSA - RQG - MSA API calls.

At each iteration, we simulate a user who selects a suggested target schema augmentation, defines optional filter conditions on target table attributes and adds calculated attributes after the merge.

Before running our experiment, we first create the attribute graphs of all dimension tables. Some $H C V$ views directly access non-analytic tables. For being conform to our model, which only extends analytic tables with other analytic table, we then build all missing first-level views over the non-analytic tables, which includes the definition of the aggregable properties of all measure attributes. Finally, we crawl all first-level views and their underlying non-analytic tables to extract all direct and derived relationships (including PK-FK relationships) and to compute the dimension and fact identifiers, as explained in Section 5. An extract of the resulting SC graph is shown in Figure 13.

We illustrate our protocol with the example of an existing hand-crafted view $H C V$. The definition of this view starts from a non-analytic table ct_LINEITEM, performs a sequence of left-outer joins and aggregates with five other non-analytic tables storing different facts, followed by a star-join with five dimension tables. The definition of $H C V$ is shown in Figure 14 which represents dimension tables by bold rounded rectangles, fact tables by bold square rectan-


Figure 13: Extract of the SC graph for analytic tables
gles, non-analytic tables by regular square rectangles, and intermediate tables by dashed rectangles. Intermediate tables are created to store the intermediate results during the view constructions (e.g. TEMP_1 represents the result of the join between ct_LINEITEM and ct_TRANSACTION).


Figure 14: Construction of HCV
Figure 15 gives a partial view of the SC graph for the analytic tables joined in the $H C V$ example. The table cardinalities range from $1.4 M$ rows (WEB_ORDER) to $23 M$ rows (ITEM_TAX and LINEITEM). All tables $D 1$ to $D 10$ are dimension tables and all other tables are firstlevel views defining fact tables. We use the notation WEB_ORDER(D5, D8) to state that the fact identifier of table WEB_ORDER only contains dimension attributes of dimension $D 5$ and $D 8$. For clarity, the figure does not show the edges corresponding to derived relationships. In particular, all five fact tables are pairwise connected by derived SC edges. The primary keys of non-analytic tables
are propagated to the corresponding fact tables (views) to account for the dependencies between dimensions, e.g. in WEB_ORDER, $\{D 5, D 8\} \mapsto D 7$.


Figure 15: Partial SC graph for analytic tables
Using our REST service, we start from the first level fact table LINEITEM built from table t_LINEITEM and construct a view $G V$ that is equivalent to $H C V$ by following the eight successive schema augmentation steps illustrated in Figure 15. Each target table is connected to the starting table LINEITEM by a path in the SC graph and a bold arc labeled by the number of the augmentation step. In step 1 , the merge query adds attributes of TRANSACTION from dimension D7, which did not exist in LINEITEM. We correctly infer that the SC edge from LINEITEM to TRANSACTION has label "NAT" by Definition 17. In step 2, we merge with fact table ITEM_DISCOUNT through an aggregate reduction using aggregate function SUM. In step 4, we merge WEB_ORDER with LINEITEM. These two tables share attributes from dimensions D3 and D5 and there also exists a user-specified relationship (dotted line in Figure 15) between dimension D4 (used in table LINEITEM) and dimension D8 (used in table WEB_ORDER). This allows us to merge WEB_ORDER with LINEITEM through schema augmentation with the common attributes between D4 and D8. All other steps use similar natural and reduction-based schema complements. After the assisted construction of $G V$, the equivalence both views, $G V$ and $H C V$, has been verified by comparing the total number of rows and their values.

Our second experimentation goal is to compare the performance of SQL queries using the original view $H C V$ and the equivalent generated view $G V$. The performance measures are done on a server with 2 TB of main memory and 1 TB of disk space. First, by comparing the total materialization cost and we can observe that computing and reading $G V$ is 1.5 times slower than for $H C V$. We then apply some aggregate queries with random combinations of dimension attributes and aggregable attributes on HCV and $G V$. The performance results are quite positive. For $17 \%$ of the queries, $G V$ performs from 1.15 to 3.1 faster than $H C V$, and for $74 \%$ of the queries, $H C V$ performs 1.1 to 4.5 faster than $G V$. For the remaining queries, both views have similar performance. To obtain a better under-
standing of this behaviour, we compared several couples of execution plans generated by the same query on $G V$ and $H C V$. We observed that the performance variations were caused by the optimization strategy which allocated similar query sub-plans to different execution engines. The performance variations therefore are independent of the view definitions and more related to the current HANA architecture and query optimizer.

Based on our two first experiments, we observed that for hand-crafted join views (which represents half of the fact tables of the application), an equivalent view could be semi-automatically generated using our REST service while achieving comparable execution performance. This is a very positive result because our view generation process is highly accessible to a business users who only can manipulate high-level analytic views and meaningful analytic attributes without any deep expertise about the global analytic schema and table relationships. In contrast to $G V$, the creation of $H C V$ view requires a strong expertise to manually optimize the sequence of operations in the view, and a precise knowledge of the database schema to express the join conditions, decide when a pre-aggregation is necessary, identify measure attributes and choose which aggregate functions are applicable to them. The preliminary price to pay for our approach is the creation of all the necessary first-level views. However, this is rapidly amortized with the number of join views in the application. In addition, it is possible to capitalize on the SC graph for the definition of future views.


Figure 16: Attribute graph satisfied by dimension dSTORE

As another substantial advantage, our REST service provides formal quality guarantees that are difficult to be verified and confirmed by the developer of a handcrafted $H C V$ view. To illustrate this point, suppose that a new aggregated view AGG_WEB_ORDER is created as: $\mathcal{A g g}_{\text {WEB_ORDER }}\left(S U M(\right.$ ORDER_AMOUNT $\left.) \mid \mathrm{X} \cup A_{-} S T O R E\right)$ where $\mathrm{X}=\{$ DATE, CUSTOMER_NO $\}$ and $A_{-} S T O R E$ is the set of the dimension attributes of dimension $\operatorname{dSTORE}$ whose validated attribute graph is depicted on Figure 16. Now suppose that an $H C V$ view needs to perform a join with

AGG_WEB_ORDER on attributes $\mathrm{X} \cup J A_{-} S T O R E$, where $J A_{-} S T O R E$ is a subset of $A_{-} S T O R E$, to obtain a new measure attribute $S U M$ (ORDER_AMOUNT). Table 11 shows the cases of ambiguous values for $S U M$ (ORDER_AMOUNT), depending on the attributes contained in JA_STORE, that would be detected by our RQG API if AGG_WEB_ORDER was selected as a target schema augmentation.

Table 11: Detection of ambiguous values

|  | Attributes in JA_STORE | Amb | Missing att |
| :--- | :--- | :---: | :--- |
| $Q_{1}$ | PLANT, CCOUNTRY, <br> CENTITY | N | - |
| $Q_{2}$ | CENTITY, CCLUSTER, <br> CSREGION, CREGION | N | - |
| $Q_{3}$ | CCOUNTRY, CSREGION | Y | CCLUSTER |
| $Q_{4}$ | CCLUSTER, CREGION | Y | CSREGION |

### 8.2.2. Comparison with related work

We first assess previous works with respect to their capability to generate $G V$ through the steps marked in Figure 15. The input is given by the set of analytic tables, which are essentially tables with additional metadata, so previous solutions formulated for relational databases can be applied.

The schema complement approach proposed in [8] relates tables depending on their schema similarity scores using a combination of attribute name and domain similarity measures. In our setting, we can assume that the relationships shown in Figure 15 can be found using the same schema similarity technique.

Next, the merges with natural schema complements applied in steps $1,5,6,7$ and 8 of Figure 15 can be suggested and computed by the method described in [8]. However, ITEM_TAX, ITEM_DISCOUNT and WEB_ORDER are not natural schema complements to LINEITEM, so steps 2,3 and 4 cannot be handled. Operators Extend attribute [10] and Extend using algorithm JoinTest [34] can also discover and merge natural schema complements. Although [10] proposes a technique to discover PK-FK relationships, their operator is limited to the detection of natural schema complements without reduction queries. The $A B A$ ( $A u g$ mentation By Attribute name) operator proposed in [9] augments the schema of a start table with an attribute, whose name is similar to a given name, coming from another related table. The discovery of related tables is as before based on a specific measure of schema similarity. However, due to the limitation that the key of the start table must be composed of a single attribute, the $A B A$ operator cannot be used in our example since the start table LINEITEM has an identifier composed of five attributes.

Other related work support high-level operators over OLAP cubes (similar to our analytic tables) that can express some of the merge queries that are automatically generated by our framework. Drill-across queries [2] can
merge the schemas of fact tables related through conformed dimensions, i.e., dimensions that are either identical or where one dimension is a subset of the other dimension. In our example, fact tables LINEITEM and TRANSACTION have dimensions $D 2, D 3, D 4$ and $D 5$ in common, so the merge of step 1 could be expressed as a drill-across operation. However, the remaining steps raise several issues. First, only natural schema complements can be handled due to the restriction of conformed dimensions, so the merges of steps 2,3 , and 4 cannot be expressed using drill-across operations. Second, only operations between fact tables are considered, so steps 5 to 9 would not be expressible using a drill-across operation.

Drill-across operations have been extended in [22] to fact tables with common compatible dimensions. In short, two dimensions are compatible if there exists a matching between the dimensions which preserves the hierarchy levels (coherency), level memberships (soundness) and rollup relationships between members of different levels (consistency). We can show that dimension compatibility now allows us to execute the steps 2 and 3 , since in both steps the common dimensions are compatible. However, step 4 is still not possible, since the common dimensions $D 4$ and $D 8$ violate the soundness condition [22]: the domain values of their common attributes, TRAN_NO in $D 4$ and ORDER_NO in $D 8$, are overlapping but not identical.

We now compare our detection of ambiguous values in aggregate queries with related works on summarizability. To illustrate, we use the aggregate queries of our experiment defined as: $\mathcal{A g g}_{\text {WEB_ORDER }}(S U M$ (ORDER_AMOUNT) | $\left.\mathrm{X} \cup J A_{-} S T O R E\right)$, where $\overline{\mathrm{X}}=\{$ DATE, CUSTOMER_NO $\}$ and attributes in JA_STORE is specified in Table 11. In our example, we assume that domain values of dimensional attributes in WEB_ORDER are complete. The work on summarizability addresses the following question: "is it correct to compute the above aggregate queries by summarizing the result of another aggregate query (i.e.,WEB_ORDER) or should they be computed directly from the non-analytic table ct_WEB_ORDER?"

In their pioneering work, [27] proposed three sufficient conditions to guarantee summarizability of aggregation operations: (a) category (i.e., group-by) attributes must form disjoint subsets of facts, (b) these subsets completely cover the entire set of facts, and (c) dimension attributes, measures and aggregate functions are compatible. Among the four aggregate queries in Table 11, $Q_{1}$ and $Q_{2}$ satisfy the summarizability conditions: the categorizations of the domain values are disjoint and complete, and the measure attribute ORDER_AMOUNT is compatible with function SUM and dimensions dSTORE, dTIME and dCUSTOMER. However, for query $Q_{3}$, the summarizability constraints are violated: there exists a +-edge R (CCOUNTRY, CSREGION) in the attribute graph of dSTORE, which signifies that there exists at least one CCOUNTRY value which has multiple CSREGION parent values (violation of the disjointness condition). Thus, $Q_{3}$ is not considered to be correct. For $Q_{4}$, all summarizability constraints are satisfied. Because edge
$R$ (CCLUSTER, CREGION) has label $\mathbf{f}$, values in CCSLUSTER have a unique value in CREGION, the categorization of their domain values is disjoint. However, $Q_{4}$ generates ambiguous results: since edge R (CCOUNTRY, CSREGION) is a +label edge, CSREGION should be included in the group-by clause. Otherwise, tuples with same value of CCOUNTRY and CREGION that belong to different CSREGION will be combined together by this aggregation. Thus, the summarizability conditions of [27] conclude that $Q_{4}$ is correct while we detect that it is ambiguous.

The work of [35] extended the summarizability conditions of [27] and showed that summarizability holds for aggregate queries with distributive aggregate functions over strict covering onto dimensions. In the above aggregate queries, function SUM is distributive and the hierarchy of dSTORE is onto, but not covering and strict. Therefore, the dSTORE hierarchy needs to be modified by adding intermediate artificial values for non-covering mappings and adding concatenated parent values for non-strict mappings. After these data modifications, summarizability is guaranteed for all aggregate queries $Q_{1}, Q_{2}, Q_{3}, Q_{4}$. By contrast, our solution does not require any change to the data in the dimension tables; we detect ambiguous queries and can report ambiguous values.

Finally, the work of [28] only considers in their data model linear, ordered and strict hierarchies, which does not capture the cases of dimension tables in our scenario. For example, dSTORE hierarchy is not linear, not ordered nor strict.

### 8.3. Feature Engineering Scenario

The third series of experiments checks the usage of our REST service in a predictive analytics application of a retail bank. The decision task consists in predicting whether a client without a credit card will obtain a card in the 3 months following a given reference date. The training and prediction process uses tables from a publicly available database [36] related through PK-FK relationships as shown in Figure 17. Each account is described by static characteristics (e.g. date of creation, address of the branch) given in table ACCOUNT and dynamic characteristics (e.g. payments debited or credited, balances) given in table TRANSACTION. The two tables CLIENT and CREDIT_CARD contain information about persons and credit cards respectively. Clients and accounts are linked by a many-to-many relationship stored in the associative table DISPOSITION. Tables GEOCODE and TIME provide detailed geographical and temporal information.

The development of the application includes a feature engineering step to build a training dataset. Starting from a table describing a CLIENT entity and from a reference date, the developer augments this table through a sequence of database queries with as many new features (attributes) as possible. The goal is to obtain more detailed information about the past behaviour of the clients like, for example, the monthly amount and balance of their credit cards during the past 12 months, the delivery of a new


Figure 17: Relationships for the bank retail database
credit card during the past three months, etc. The final feature dataset is denoted by FEAT. The individual augmentation queries consist of joins, projections, filters, aggregations and pivots and the whole SQL script used to build the training dataset consists of about 1,500 lines of code comprising many sub-queries involving 16 joins and 132 functional expressions to compute measure attribute values.

### 8.3.1. Experiments

The goal of our experiment is to verify if a data scientist can semi-automatically generate an view $G V$ that is equivalent to $F E A T$ by iteratively extending an initial dimension table CLIENT through several schema augmentation steps each of which consists of a sequence of CSA RQG - MSA API calls.

For running our experiment, we first created 6 dimension tables, denoted by dACCOUNT, dGEO, dCLIENT, dCARD, dTIME and dTRANSACTION, and 3 fact tables, denoted by TRANS, ACC_DISPO, and CARD_DISPO over the bank retail database tables. TRANS is defined over database table TRANSACTION and dimensions dTIME, dACCOUNT and dTRANSACTION and has two measures AMOUNT and BALANCE. ACC_DISPO is defined over database table DISPOSITION and dimensions dCLIENT and dACCOUNT. Finally, CARD_DISPO is defined as a join of tables CREDIT_CARD and DISPOSITION over dimensions dCLIENT, dTIME and dCARD. The attribute graphs are defined for all dimension tables, as well as the aggregable properties for all measure attributes of the fact tables. Finally, all tables were crawled by the HANA Crawlers to extract all direct and derived relationships (including PK-FK relationships) and compute the dimension and fact identifiers. The resulting SC graph among analytic tables is shown in Figure 18.

The view generation process starts from the dimension table dCLIENT(CLIENT_ID, GEO_ID, BIRTH_DATE, SEX, ...), which identifies a client entity, and applies a sequence of schema augmentation steps as indicated by the numbered arrows in the SC graph of Figure 18.

In Step 1, detail attributes from dimension dGEO are added to dCLIENT through a natural schema complement using common attribute GEO_ID. In step 2, dCLIENT is augmented with attribute ACCOUNTID from fact table ACC_DISPO through a schema augmentation using the


Figure 18: SC graph for the analytic tables
common attribute CLIENT_ID (thus, rows in dCLIENT are multiplied). In step 3, the resulting table is augmented with detail attributes from dimension dACCOUNT through a natural schema complement using common attribute ACCOUNTID. In step 4, table dCLIENT is augmented with measures from fact table TRANS. The common attribute between TRANS and dCLIENT is ACCOUNTID, that is, a subset of the fact identifier of TRANS which is: \{ACCOUNTID, DATE, TRANS_ID\}. The following user actions are applied to fully reduce the fact identifier. A calculated attribute MONTH(DATE) (i.e., transaction month) is pivoted as a column with values coming from measures AMOUNT and BALANCE, and attribute TRANS_ID is removed. In addition, two user actions are provided: a filter on DATE pre-selects facts within the 12 months preceding a user-given reference date, and attribute TRANS_TYPE is pivoted as a column with values coming from measures amount and balance. These two operations are injected in the query reduction generation over TRANS and the reduction query is merged with dCLIENT.

We now want to add a measure attribute storing the number of credits cards for each client. In Step 5, attributes from CARD_DISPO are added to dCLIENT through a schema augmentation using common attributes CLIENT_ID. The identifier of CARD_DISPO, consisting of \{CLIENT_ID, CARD_ID\}, is fully reduced by computing an aggregated CARD_ID using function $C O U N T$, and a reduction query is generated over CARD_DISPO. The result is then merged with dCLIENT using the remaining common attribute CLIENT_ID. We get a result table with 156 attributes where 135 of them are measures. As in the BI Scenario (Section 8.2), the correctness of $G V$ with respect to $F E A T$ is done by comparing the number of tuples and their values in both tables. We show in Figure 19 the complete definition of $G V$ where each temporary result represents a step in the view generation process.

Our semi-automatic view generation process has several advantages over the manual creation of table FEAT. First, the schema augmentations required to build a $F E A T$ table are successfully suggested by our CSA API and the order in which they are returned matches pretty well the user needs. Second, simple user actions in steps 4 and 5 yield complex reduction queries involving filter, pivot, and aggregate operations through the RQG API. Isolating


Figure 19: Construction of GV
user reduction actions in the generation process provides a great flexibility. For instance, in Step 4, filter conditions could be changed to select facts within the 2 years preceding the reference date, or pivot operations could changed to WEEK(DATE). Third, the preparatory definition of formal aggregable properties for all measures in fact table TRANS help the expert to choose the correct aggregate functions for measures amount and balance in Step 4. For instance, function SUM cannot be applied to BALANCE. Finally, our method for propagating aggregable properties controls the dimensions with respect to which measure COUNT(CARD_ID) can be aggregated after Step 5 since the measure only depends on CLIENT_ID.

Note that our REST service could be directly applied to an SC graph consisting only of the retail database tables (i.e., without creating any dimension or fact table) and would still produce table FEAT as result. However, working directly with database tables has two main drawbacks. First, the user must understand the operational data model of the database that carries many attributes that are irrelevant for business data analysis. Second, the benefits of metadata such as the separation of dimension and measure attributes and the definition of aggregable properties would be lost.

### 8.3.2. Comparison with related work

We can assess other related work with respect to their ability to generate table FEAT. As we did in Section 8.2.2, we first consider methods that find schema complements. Starting from table dCLIENT, the previous works of [8], [10] and [34] are only able to detect natural schema complements and suggest the merge of step 1 in Figure 18. The other steps require reduction queries which are not supported by these methods. Note that the $A B A$ operator of [9] can be applied on dataset dCLIENT since its identifier consists of a single attribute CLIENT.ID. However, this is again only useful for step 1 .

Concerning the use of drill-across queries [2], we can show that no step in Figure 18 can be defined by a drill-
across operation between two fact tables: steps 1 and 3 merge two dimension tables and steps 2,4 and 5 merge a dimension with a fact table.

## 9. Conclusions and Future Work

In this paper, we propose a solution to the problem of discovering and merging schema augmentations for analytic datasets. We formally define efficient algorithms for building reduction-based schema complements which generalize previously defined natural schema complements. We discuss various quality issues and their solution for generating semantically correct schema augmentations and merged tables annotated with proper metadata enabling their correct usage in future analytic queries. We also present the implementation of our solution as a REST service in SAP HANA which, to the best of our knowledge, is the first implementation of that sort. This service assists business users and data scientists to explore a large space of possible schema augmentations and, as illustrated in our experiments, to generate complex views in a controlled way with formally defined quality guarantees concerning attribute ambiguity and data completeness.

We envision several directions for future work. A first direction concerns the selection of schema complements based on user-specified criteria such as keywords and the definition of strategies for ranking the schema complements returned by our algorithm based on user profile and contextual user preferences. Another direction is to refine our relationships to include more sophisticated temporal conditions on time intervals, and the ability to deal with look-up tables. In addition, we want to suggest ways of selecting filter, pivot, and aggregate reduction operations by taking inspiration from data preparation techniques described for instance in [37]. We observed that these capabilities are required in feature engineering scenarios. Finally, use case studies with a real corpus of analytic tables defined in business verticals will be continued to assess and improve the value of our method.

## References

[1] C. S. Jensen, T. B. Pedersen, C. Thomsen, Multidimensional Databases and Data Warehousing, Synthesis lectures on data management, Morgan and Claypool Publishers, 2010.
[2] R. Kimball, M. Ross, The Data Warehouse Toolkit, John Wiley \& Sons, 2013.
[3] The virtual data model in sap s/4hana. URL https://help.sap.com/viewer/ 6b356c79dea443c4bbeeaf0865e04207/1809.000/en-US/ 8573b810511948c8a99c0672abc159aa.html
[4] A. Pattanayak, High performance analytics with sap hana virtual models, Journal of Computer and Communications 5 (07) (2017) 1-10.
[5] Data Wrangling Tools \& Software | Trifacta. URL https://www.trifacta.com/
[6] Paxata | Self-Service Data Preparation for Data Analytics. URL https://www.paxata.com/
[7] SAP Agile Data Preparation and Transformation Solution. URL https://www.sap.com/products/data-preparation.html
[8] A. Das Sarma, L. Fang, N. Gupta, A. Halevy, H. Lee, F. Wu, R. Xin, C. Yu, Finding related tables, in: Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data, ACM, 2012, pp. 817-828.
[9] M. Yakout, K. Ganjam, K. Chakrabarti, S. Chaudhuri, InfoGather: entity augmentation and attribute discovery by holistic matching with web tables, in: Proceedings of the 2012 international conference on Management of Data - SIGMOD '12, ACM Press, 2012, pp. 97-108. doi:10.1145/2213836. 2213848.
[10] D. Deng, R. C. Fernandez, Z. Abedjan, S. Wang, M. Stonebraker, A. K. Elmagarmid, I. F. Ilyas, S. Madden, M. Ouzzani, N. Tang, The Data Civilizer System, in: CIDR, 2017.
[11] J. D. Ullman, H. García-Molina, J. Widom, Database System: The Complete Book, 2nd Edition, Prentice Hall, 2009.
[12] A. Badia, D. Lemire, Functional dependencies with null markers, The Computer Journal 58 (5) (2014) 1160-1168.
[13] J.-N. Mazón, J. Lechtenbörger, J. Trujillo, A survey on summarizability issues in multidimensional modeling, Data \& Knowledge Engineering 68 (12) (2009) 1452-1469.
[14] J. Horner, I.-Y. Song, P. P. Chen, An analysis of additivity in OLAP systems, in: Proceedings of the 7th ACM international workshop on Data warehousing and OLAP, ACM, 2004, pp. 83-91.
[15] J. Lee, M. Muehle, N. May, F. Faerber, V. Sikka, H. Plattner, J. Krueger, M. Grund, High-performance transaction processing in SAP HANA, IEEE Data Eng. Bull 36 (2) (2013) 28-33.
[16] SAP HANA Modeling Guide, SAP, AG, 2019. URL https://help.sap.com/doc/ 227fc55c4fc44a43b43752d6b127bdf3/2.0.04
[17] SAP HANA smart data integration and SAP HANA smart data quality - SAP help portal.
URL https://help.sap.com/viewer/p/HANA_SMART_DATA_ INTEGRATION
[18] R. Brunel, J. Finis, G. Franz, N. May, A. Kemper, T. Neumann, F. Faerber, Supporting hierarchical data in SAP HANA, in: Data Engineering (ICDE), 2015 IEEE 31st International Conference on, IEEE, 2015, pp. 1280-1291.
[19] M. Paradies, C. Kinder, J. Bross, T. Fischer, R. Kasperovics, H. Gildhoff, GraphScript: implementing complex graph algorithms in SAP HANA, in: Proceedings of DBPL 2017, ACM Press, 2017, pp. 1-4.
[20] Z. Bellahsene, A. Bonifati, E. Rahm (Eds.), Schema Matching and Mapping, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011. doi:10.1007/978-3-642-16518-4.
[21] F. Nargesian, E. Zhu, K. Q. Pu, R. J. Miller, Table union search on open data, Proceedings of the VLDB Endowment 11 (7) (2018) 813-825.
[22] R. Torlone, Two approaches to the integration of heterogeneous data warehouses, Distributed and Parallel Databases 23 (1) (2008) 69-97.
[23] A. Abelló, J. Samos, F. Saltor, On relationships offering new drill-across possibilities, in: Proceedings of the 5th ACM international workshop on Data Warehousing and OLAP, ACM, 2002, pp. 7-13.
[24] M. Rafanelli, A. Shoshani, Storm: A statistical object representation model, in: International Conference on Scientific and Statistical Database Management, Vol. 420, Springer, 1990, pp. 14-29. doi:10.1007/3-540-52342-1\_18.
[25] J. Horner, I.-Y. Song, A taxonomy of inaccurate summaries and their management in OLAP systems, in: International Conference on Conceptual Modeling, Vol. 3716, Springer, 2005, pp. 433-448. doi:10.1007/11568322\_28.
[26] C. A. Hurtado, C. Gutierrez, A. O. Mendelzon, Capturing summarizability with integrity constraints in OLAP, ACM Transactions on Database Systems 30 (3) (2005) 854-886.
[27] H.-J. Lenz, A. Shoshani, Summarizability in OLAP and statistical data bases, in: Proceedings of the Ninth International Conference on Scientific and Statistical Database Management (SSDBM '97), 1997, pp. 132-143.
[28] T. Niemi, M. Niinimäki, P. Thanisch, J. Nummenmaa, Detecting summarizability in OLAP, Data \& Knowledge Engineering

89 (2014) 1-20.
[29] J.-P. Dittrich, D. Kossmann, A. Kreutz, Bridging the gap between OLAP and SQL, Proceedings of the 31st international conference on Very large data bases (2005) 1031-1042.
[30] W. Lehner, J. Albrecht, H. Wedekind, Normal forms for multidimensional databases, in: Scientific and Statistical Database Management, 1998. Proceedings. Tenth International Conference on, IEEE, 1998, pp. 63-72.
[31] T. Papenbrock, J. Ehrlich, J. Marten, T. Neubert, J.-P. Rudolph, M. Schönberg, J. Zwiener, F. Naumann, Functional dependency discovery: An experimental evaluation of seven algorithms, Proceedings of the VLDB Endowment 8 (10) (2015) 1082-1093.
[32] L. Berti-Équille, H. Harmouch, F. Naumann, N. Novelli, S. Thirumuruganathan, Discovery of genuine functional dependencies from relational data with missing values, Proc. VLDB Endow. 11 (8) (2018) 880-892. doi:10.14778/3204028. 3204032. URL https://doi.org/10.14778/3204028.3204032
[33] J. Liu, J. Li, C. Liu, Y. Chen, Discover dependencies from data-a review, IEEE Transactions on Knowledge and Data Engineering 24 (2) (2010) 251-264.
[34] M. J. Cafarella, A. Halevy, N. Khoussainova, Data integration for the relational web, Proceedings of the VLDB Endowment 2 (1) (2009) 1090-1101.
[35] T. B. Pedersen, C. S. Jensen, C. E. Dyreson, Extending practical pre-aggregation in on-line analytical processing (1999) 12.
[36] Home page of pkdd discovery challenge.
URL https://sorry.vse.cz/~berka/challenge/PAST/index. html
[37] A. Zheng, A. Casari, Feature Engineering for Machine Learning, O'Reilly Media, Inc., 2018.

## Appendix A. Proofs for propositions

Proposition 1. Let $\mathcal{D}=\left(\mathrm{S}, R, \lambda_{R}, \perp, \mathrm{~T}\right)$ be an attribute graph. Then, the subset $\mathrm{X}_{D} \subseteq \mathrm{~S}$ of all attributes in S with at least one + labeled in-edge and no $\mathbf{f}$ labeled in-edge is a dimension identifier for all valid dimension tables with attributes S .

Proof. Literal functional dependencies correspond to functional dependencies where null values are considered as constants $\left(t . A \equiv t^{\prime} . A\right.$ iff $t . A=t^{\prime} . A$ or $t . A$ and $t^{\prime} . A$ are both null) and, under this interpretation, the Armstrong's axioms for functional dependencies also hold for literal functional dependencies.

Let $\mathrm{W}=\mathrm{S}-\mathrm{X}_{D}$ be the subset of S such that for all $\mathrm{A} \in \mathrm{W}$, A has no + in-edges or at least one $\mathbf{f}$ in-edge. We first prove that the literal functional dependency $\mathrm{S}-\mathrm{W} \mapsto$ S holds for all valid dimension tables of $\mathcal{D}$. For this, it is sufficient to prove that (1) for all attributes $A_{i} \in W$, there exists a subset of attributes $\mathrm{S}_{i} \subseteq \mathrm{~S}-\left\{A_{i}\right\}$ where $\mathrm{S}_{i} \mapsto A_{i}$. Since attribute graphs (and the corresponding attribute dependencies) are acyclic, we can define a partial order $\leq$ over attributes in $W$ such that each attribute $A_{i} \in \mathrm{~W}$ only depends on attributes $A_{k} \leq A_{i}$. Then, from $S_{k} \subseteq S_{i}$ and the transitivity of $\mapsto$ and if for all attributes $A_{i} \in W$, $\mathrm{S}-\left\{\mathrm{A}_{i}\right\} \mapsto A_{i}$ follows $\mathrm{S}-\mathrm{W} \mapsto \mathrm{S}$. For proving (1), we distinguish two cases which separate the attributes in W into two classes. Let $\mathrm{A}_{i} \in \mathrm{~W}$ and $\mathrm{S}_{i}=\left\{\mathrm{A}_{k} \in \mathrm{~S} \mid\left(\mathrm{A}_{k}, \mathrm{~A}_{i}\right) \in\right.$ $\mathrm{R}\}$. By the definition of $W$, for each $\mathrm{A}_{i}$ there exists at least one $\mathrm{A}_{k} \in \mathrm{~S}_{i}$ where $\mathrm{R}\left(\mathrm{A}_{k}, \mathrm{~A}_{i}\right)=\mathbf{f}$ (case one) or for all $\mathrm{A}_{k} \in \mathrm{~S}_{i}: \mathrm{R}\left(\mathrm{A}_{k}, \mathrm{~A}_{i}\right)=\mathbf{1}$ (case two).

- Case one: Let $\mathrm{A}_{i} \in \mathrm{~W}$ and $\mathrm{A}_{k} \in \mathrm{~S}_{i}$ such that $\mathrm{R}\left(\mathrm{A}_{k}, \mathrm{~A}_{i}\right)=$ f. By the definition of attribute graph and label $\mathbf{f}$ edges, $\mathrm{R}\left(\mathrm{A}_{k}, \mathrm{~A}_{i}\right)=\mathbf{f}$ is equivalent with $\mathrm{A}_{k} \mapsto \mathrm{~A}_{i}$.
- Case two: Let $\mathrm{A}_{i} \in \mathrm{~W}$ and for all $\mathrm{A}_{k} \in \mathrm{~S}_{i}: \mathrm{R}\left(\mathrm{A}_{k}, \mathrm{~A}_{i}\right)=\mathbf{1}$. For any two tuples $t_{1}, t_{2} \in T$, there exists at least one attribute in $A_{k} \in \mathrm{~S}_{i}$ where $t_{1} \cdot A_{k}$ is not null. Then, if $t_{1} \cdot \mathrm{~S}_{i} \equiv t_{2} \cdot \mathrm{~S}_{i}$, there exists at least one attribute $\mathrm{A}_{k} \in \mathrm{~S}_{i}$ such that $t_{1} \cdot \mathrm{~A}_{k}=t_{2} \cdot \mathrm{~A}_{k}$ and since $\mathrm{R}\left(\mathrm{A}_{k}, \mathrm{~A}_{i}\right)=1$, we get $t_{1} \cdot \mathrm{~A}_{i} \equiv t_{2} \cdot \mathrm{~A}_{i}$.

We now prove by contradiction that $\mathrm{S}-\mathrm{W}$ is minimal, i.e. there exists no subset $\mathrm{Y} \subset \mathrm{S}-W$ where $\mathrm{Y} \mapsto \mathrm{S}$. Assume that $\mathrm{S}-\mathrm{W} \mapsto \mathrm{S}$ but $\mathrm{S}-\mathrm{W}$ is not minimal. Then there exists at least one attribute $A_{p} \in S-W$, such that ( $\mathrm{S}-$ $\left.\mathrm{W}-\mathrm{A}_{p}\right) \mapsto \mathrm{S}$. By the definition of W and $\mathrm{A}_{p} \notin \mathrm{~W}, \mathrm{~A}_{p}$ has at least one ' + ' in-edge and $\mathrm{A}_{p}$ doesn't have ' f ' inedges. Let $\mathrm{A}_{k} \in \mathrm{~S}$ such that $\mathrm{R}\left(\mathrm{A}_{k}, \mathrm{~A}_{p}\right)=+$. Then we can build a valid dimension table (hierarchy) $T$ where there exists a couple of tuples $t_{1}, t_{2} \in T$ (paths in the hierarchy $T)$ such that $t_{1} \cdot\left(\mathrm{~S}-\mathrm{A}_{p}\right) \equiv t_{2} \cdot\left(\mathrm{~S}-\mathrm{A}_{p}\right)$, and $t_{1} \cdot \mathrm{~A}_{k} \not \equiv t_{2} \cdot \mathrm{~A}_{k}$, i.e. $\mathrm{S}-\mathrm{A}_{p} \nvdash \mathrm{~A}_{k}$. Then, $\mathrm{S}-\mathrm{W}-\mathrm{A}_{p} \nvdash \mathrm{~A}_{k}$, which is in contradiction with $\left(\mathrm{S}-\mathrm{W}-\mathrm{A}_{p}\right) \mapsto \mathrm{S} . \mathrm{S}-\mathrm{W}$ is the minimal set such that $\mathrm{S}-\mathrm{W} \mapsto \mathrm{S}$.

Proposition 2. Let $T(\mathrm{~S})$ be a fact table defined over a set of dimensions $D_{1}, \cdots, D_{n}$ and $\mathrm{K}_{1}, \cdots, \mathrm{~K}_{n}$ be the dimension identifiers of $D_{1} \cap \mathrm{~S}, \cdots, D_{n} \cap \mathrm{~S}$ respectively. Then $\mathrm{K}=$
$\mathrm{K}_{1} \cup \cdots \cup \mathrm{~K}_{n}$ is a fact identifier of $T$, and K is a minimal identifier if all dimensions in $T$ are mutually independent.

Proof. Let $\mathrm{S}_{D} \subset \mathrm{~S}$ the set of dimension attributes in $T$ and $\mathrm{S}-\mathrm{S}_{D}$ be measure attributes. $K_{1}, \ldots, K_{n}$ are the dimension identifiers of $D_{1} \cap \mathrm{~S}, \ldots, D_{n} \cap \mathrm{~S}$ respectively, and we have $K_{1} \mapsto D_{1} \cap \mathrm{~S}, \ldots, K_{n} \mapsto D_{n} \cap \mathrm{~S}$. By Armstrong union axiom for LFDs, we get $\left(K_{1} \cup \ldots \cup K_{n}\right) \mapsto\left(D_{1} \cap\right.$ $S) \cup \ldots \cup\left(D_{n} \cap \mathrm{~S}\right)=\mathrm{S}_{D}$. By the definition of fact tables, all measure measures depend on the dimension attributes, i.e. $\mathrm{S}_{D} \mapsto \mathrm{~S}$ and by transitivity of LFD's, we get ( $K_{1} \cup$ $\left.\ldots \cup K_{n}\right) \mapsto \mathrm{S}$. Two dimension identifiers $K_{i}, K_{j}, i \neq j$ are independent, if there exists no attribute $A \in K_{j}$ such that $K_{i} \mapsto A$. Then, it is easy to show that if all $K_{i}$ are minimal and mutually independent, $K_{1} \cup \ldots \cup K_{n}$ is a minimal a fact identifier of $T$.

Proposition 3. Let $T_{0}^{\prime}\left(\mathrm{S}_{0}^{\prime}\right)$ be a merge of table $T_{0}\left(\mathrm{~S}_{0}\right)$ with a target schema augmentation $T(\mathrm{~S})$. Let K be the identifier of $T_{0}$ and $\mathrm{S}_{\text {new }} \subseteq \mathrm{S}_{0}^{\prime}-\mathrm{S}_{0}$ be the set of dimension attributes added to $T_{0}$ in the merge. Then, for all minimal subsets $\mathrm{K}_{\text {new }} \subseteq \mathrm{S}_{\text {new }}$ where $\mathrm{K}_{\text {new }} \mapsto \mathrm{S}_{\text {new }}, \mathrm{K} \cup \mathrm{K}_{\text {new }}$ is an identifier of $T_{0}^{\prime}$.

Proof. Let $K^{\prime}=K \cup K_{\text {new }}$ where $K_{\text {new }} \neq \emptyset$. We prove $K^{\prime} \mapsto \mathrm{S}_{0}^{\prime}$ by contradiction. Assume that $K^{\prime}$ is not the identifier of $T_{0}^{\prime}$. Then there exist two tuples $t_{1}, t_{2} \in T_{0}^{\prime}$ such that $t_{1} \cdot K^{\prime} \equiv t_{2} \cdot K^{\prime}$ but $t_{1} \cdot \mathrm{~B} \not \equiv t_{2} \cdot \mathrm{~B}, \mathrm{~B} \in \mathrm{~S}_{0}^{\prime}$. We distinguish between two cases:
$-\mathrm{B} \in \mathrm{S}_{0}$ : We know that $K \subset K^{\prime}$ and $K \mapsto \mathrm{~S}_{0}$. Then by $t_{1} \cdot K^{\prime} \equiv t_{2} \cdot K^{\prime}$ we have $t_{1} \cdot K \equiv t_{2} \cdot K$ and $t_{1} \cdot \mathrm{~B} \equiv t_{2} \cdot \mathrm{~B}$ which contradicts our assumption.
$-\mathrm{B} \in \mathrm{S}_{\text {new }}$ : We know that $K_{\text {new }} \mapsto \mathrm{S}_{\text {new }}$. Then by $t_{1} \cdot K^{\prime} \equiv$ $t_{2} \cdot K^{\prime}$ we have $t_{1} \cdot K_{\text {new }} \equiv t_{2} \cdot K_{\text {new }}$ and $t_{1} \cdot \mathrm{~S}_{\text {new }} \equiv t_{2} \cdot \mathrm{~S}_{\text {new }}$ which contradicts our assumption.

We conclude $K^{\prime} \mapsto \mathrm{S}_{0}^{\prime}$.

Proposition 4. Let $T_{0}\left(\mathrm{~S}_{0}\right)$ and $T(\mathrm{~S})$ be two tables such that $T$ is a schema augmentation to $T_{0}$ with common attributes Y . Let K be an identifier of $T$. The result of $a$ reduction query $Q(T)$ on attributes $\mathrm{K}^{\prime}=\mathrm{K} \cap \mathrm{Y}$, is a natural schema complement to $T_{0}$.

Proof. We first show that (A) if $\mathrm{K}^{\prime} \mapsto \mathrm{S}^{\prime}$, then $T^{\prime}\left(\mathrm{S}^{\prime}\right)$ is a natural schema complement of $T_{0}\left(\mathrm{~S}_{0}\right)$. By definition, $\mathrm{K}^{\prime}=K \cap \mathrm{Y}$. Then, from $\mathrm{K} \subseteq \mathrm{Y}$ and $\mathrm{K}^{\prime} \mapsto \mathrm{S}^{\prime}$ follows $\mathrm{Y} \mapsto \mathrm{S}^{\prime}$, i.e. $\mathrm{K}^{\prime} \mapsto \mathrm{S}^{\prime}$ is a sufficient condition for guaranteeing that $T^{\prime}\left(\mathrm{S}^{\prime}\right)$ is a natural schema complement of $T_{0}\left(\mathrm{~S}_{0}\right)$.

We now show that (B) if $Q(T)$ is a reduction query on $\mathrm{K}^{\prime}$, then $\mathrm{K}^{\prime} \mapsto \mathrm{S}^{\prime}$. We distinguish the case where $\mathcal{Q}(T)$ only contains one reduction operation and the case where $Q(T)$ is a sequence of at least two reduction operations. Let $T^{\prime}\left(\mathrm{S}^{\prime}\right)$ be the result of $\mathcal{Q}(T)$.

- Case 1: $Q(T)$ only contains one reduction operation. We prove for each reduction query by contradiction that $\mathrm{K}^{\prime} \mapsto \mathrm{S}^{\prime}$ in $T^{\prime}$. Suppose that $\mathrm{K}^{\prime} \nvdash \mathrm{S}^{\prime}$. Then there exist two tuples $t_{1}, t_{2} \in T^{\prime}$ such that $t_{1} \cdot K^{\prime} \equiv t_{2} \cdot K^{\prime}$ but $t_{1} \cdot \mathrm{~S}^{\prime} \not \equiv t_{2} \cdot \mathrm{~S}^{\prime}$ :
- When $Q(T)$ is a filter reduction on attributes $K-\mathrm{Y}$, by definition, $\mathrm{S}^{\prime}=\mathrm{S}$ and $t_{1} .(K-\mathrm{Y}) \equiv t_{2} \cdot(K-\mathrm{Y})$. Then, since $t_{1} \cdot K^{\prime} \equiv t_{2} \cdot K^{\prime}$ (assumption), $K=K^{\prime} \cup$ $(K-Y), \mathrm{K} \mapsto \mathrm{S}$ and $\mathrm{S}^{\prime}=\mathrm{S}$, we have $t_{1} \cdot \mathrm{~S}^{\prime} \equiv t_{2} \cdot \mathrm{~S}^{\prime}$ which contradicts our assumption. Therefore we get $K^{\prime} \mapsto S^{\prime}$.
- When $Q(T)$ is an aggregate reduction, we get $S^{\prime}=$ $\mathrm{S}-(K-\mathrm{Y})$ and $\mathrm{K}^{\prime}$ is the set of attributes in the group-by clause $Q(T)$. Then, by definition of aggregation queries, $\mathrm{K}^{\prime}$ is the identifier (key) of $T^{\prime}$, a which contradicts our assumption $K^{\prime} \not \not \nrightarrow S^{\prime}$.
- When $Q(T)$ is a pivot reduction, then attributes in $K-\mathrm{Y}$ are pivoted into new columns in $T^{\prime}$. Similar to aggregation, by the definition of pivot queries, $K^{\prime}$ is the identifier of $T^{\prime}$ and $\mathrm{K}^{\prime} \mapsto \mathrm{S}^{\prime}$ (which contradicts our assumption).
Therefore, when $Q(T)$ only contains one reduction operation, we have $\mathrm{K}^{\prime} \mapsto \mathrm{S}^{\prime}$.
-Case 2: $Q(T)$ contains multiple reduction operations. $\operatorname{Let} Q(T)=Q_{1}\left(Q_{2}\left(\ldots\left(Q_{n}(T)\right) \ldots\right)\right)$ be a sequence of of $n>1$ reduction queries on the sets of attributes $\mathrm{Y}_{1}, \mathrm{Y}_{2}, \ldots \mathrm{Y}_{n}$ where $\mathrm{Y}_{i} \subset \mathrm{Y}_{i+1}$ and $\mathrm{Y}_{1}=\mathrm{Y}$. Let $T_{i}^{\prime}\left(S_{i}^{\prime}\right)$ denote the result of reduction query $Q_{i}\left(T_{i+1}^{\prime}\right)$ on attributes $\mathrm{K}_{i}^{\prime}=\mathrm{K}_{i+1}^{\prime} \cap \mathrm{Y}_{i}^{\prime}$. Then, $T^{\prime}\left(\mathrm{S}^{\prime}\right)=T_{1}\left(\mathrm{~S}_{1}\right)$ is the result of reduction query $Q(T)=Q_{1}\left(T_{2}^{\prime}\right)$ on attributes $K_{1}^{\prime}=K_{2}^{\prime} \cap Y_{1}=K_{2}^{\prime} \cap Y$. From Case 1, we then know that $\mathrm{K}_{1}^{\prime} \rightarrow \mathrm{S}_{1}^{\prime}$.

From (A) and (B) follows that $Q(T)$ is a natural schema complement of $T_{0}$.

Proposition 5. Composition of relationships. Let $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ and $\mathrm{R}_{2}\left(T_{2}, T_{3}\right)$ be two well-formed relationships between tables $T_{1}, T_{2}$ and $T_{3}$ with respective common attributes $\mathrm{Y}_{1}$ and $\mathrm{Y}_{2}$. If $\mathrm{Y}_{3}=\mathrm{Y}_{1} \cap \mathrm{Y}_{2} \neq \emptyset$, then there exists a well-formed relationship $\mathrm{R}_{3}\left(T_{1}, T_{3}\right)$ that is a composition of $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ and $\mathrm{R}_{2}\left(T_{2}, T_{3}\right)$ with common attributes $\mathrm{Y}_{3}$.
Proof. We prove that there exists a relationship between $T_{1}$ and $T_{3}$ and the relationship is well-formed. We distinguish four cases by the types of relationships between $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ and $\mathrm{R}_{2}\left(T_{2}, T_{3}\right)$ with common attributes $\mathrm{Y}_{1}$ and $\mathrm{Y}_{2}$.

- Case 1: $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ and $\mathrm{R}_{2}\left(T_{2}, T_{3}\right)$ are both join relationships. By definition, for all attributes $\mathrm{A} \in \mathrm{Y}_{3}=\mathrm{Y}_{1} \cap \mathrm{Y}_{3}$, $T_{1} \cdot \mathrm{~A}=T_{2} . \mathrm{A}$ and $T_{2} \cdot \mathrm{~A}=T_{3} . \mathrm{A}$. Then by transitivity of equality, $T_{1} \cdot \mathrm{~A}=T_{3} . \mathrm{A}$, i.e., there $\mathrm{R}_{3}\left(T_{1}, T_{3}\right)$ is a join relationship with common attributes $\mathrm{Y}_{3}$.
- Case 2: $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ is a join relationship and $\mathrm{R}_{2}\left(T_{2}, T_{3}\right)$ is an attribute mapping relationship. We first consider the case where all attributes $\mathrm{A} \in \mathrm{Y}_{3}=\mathrm{Y}_{1} \cap \mathrm{Y}_{3}, T_{1} \cdot \mathrm{~A}=T_{2} \cdot \mathrm{~A}$ and $T_{2} . \mathrm{A} \rightarrow T_{3} . \mathrm{A}$. By definition, $T_{2} . \mathrm{A} \rightarrow T_{3}$.A indicates that there exists a query $Q_{23}$ and a set of tables such that $T_{3}=Q_{23}\left(T_{2}, T_{1}^{\prime}, \cdots, T_{n}^{\prime}\right)$ and $\forall y \in \operatorname{dom}(\mathrm{~A}):$

$$
\begin{array}{r}
\sigma_{\mathrm{A}=y}\left(Q_{23}\left(T_{2}, T_{1}^{\prime}, \cdots, T_{n}^{\prime}\right)\right) \\
=\sigma_{\mathrm{A}=y}\left(Q_{23}\left(\sigma_{\mathrm{A}=y}\left(T_{2}\right), T_{1}^{\prime}, \cdots, T_{n}^{\prime}\right)\right) \tag{A.1}
\end{array}
$$

We prove that $T_{1} . \mathrm{A} \rightarrow T_{3} . \mathrm{A}$, i.e., there exists a query $Q_{13}$ and a set of tables such that $T_{3}=$ $\left.Q_{13}\left(T_{1}, T_{1}^{\prime \prime}, \cdots, T_{m}^{\prime \prime}\right)\right)$ and $\forall y \in \operatorname{dom}(\mathrm{~A})$ :

$$
\begin{array}{r}
\sigma_{\mathrm{A}=y}\left(Q_{13}\left(T_{1}, T_{1}^{\prime \prime}, \cdots, T_{m}^{\prime \prime}\right)\right) \\
=\sigma_{\mathrm{A}=y}\left(Q_{13}\left(\sigma_{\mathrm{A}=x}\left(T_{1}\right), T_{1}^{\prime \prime}, \cdots, T_{m}^{\prime \prime}\right)\right) \tag{A.2}
\end{array}
$$

Since $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ is a join relationship, there exists a table $T^{\prime}$ such that $T_{2}=\pi_{\mathrm{S}_{2}}\left(T_{1} \bowtie T^{\prime}\right)$. Then, by replacing $T_{2}$ in Equation A. 1 and by $T_{1} \cdot \mathrm{~A}=T_{2} \cdot \mathrm{~A}$, we get:

$$
\begin{array}{r}
\sigma_{\mathrm{A}=y}\left(Q_{23}\left(\pi_{\mathrm{S}_{2}}\left(T_{1} \bowtie T^{\prime}\right), T_{1}^{\prime}, \cdots, T_{n}^{\prime}\right)\right) \\
=\sigma_{\mathrm{A}=y}\left(Q_{23}\left(\sigma_{\mathrm{A}=y}\left(\pi_{\mathrm{S}_{2}}\left(T_{1} \bowtie T^{\prime}\right)\right), T_{1}^{\prime}, \cdots, T_{n}^{\prime}\right)\right) \tag{A.3}
\end{array}
$$

Assuming $Q_{13}=Q_{23}\left(\pi_{\mathrm{S}_{2}}\left(T_{1} \bowtie T^{\prime}\right), T_{1}^{\prime}, \cdots, T_{n}^{\prime}\right)$, by replacement in Equation A.3, we get:

$$
\begin{array}{r}
\sigma_{\mathrm{A}=y}\left(Q_{13}\left(T_{1}, T^{\prime}, T_{1}^{\prime}, \cdots, T_{n}^{\prime}\right)\right) \\
=\sigma_{\mathrm{A}=y}\left(Q_{23}\left(\sigma_{\mathrm{A}=y}\left(\pi_{\mathrm{S}_{2}}\left(T_{1} \bowtie T^{\prime}\right)\right), T_{1}^{\prime}, \cdots, T_{n}^{\prime}\right)\right) \tag{A.4}
\end{array}
$$

Finally, by pushing all selections into the joins of the right side of Equation A.3, we obtain:

$$
\begin{array}{r}
\sigma_{\mathrm{A}=y}\left(Q_{23}\left(\pi_{\mathrm{S}_{2}}\left(\sigma_{\mathrm{A}=y}\left(T_{1}\right) \bowtie T^{\prime}\right), T_{1}^{\prime}, \cdots, T_{n}^{\prime}\right)\right)  \tag{A.5}\\
=\sigma_{\mathrm{A}=y}\left(Q_{13}\left(\sigma_{\mathrm{A}=y}\left(T_{1}\right), T^{\prime}, T_{1}^{\prime}, \cdots, T_{n}^{\prime}\right)\right)
\end{array}
$$

which corresponds to $T_{1} . \mathrm{A} \rightarrow T_{3}$.A.
We follow the same reasoning for proving that $T_{1} \cdot \mathrm{~A}=$ $T_{2}$.A implies $T_{3} . \mathrm{A} \rightarrow T_{2} . \mathrm{A}, T_{3} . \mathrm{A} \rightarrow T_{1}$.A.

- Case 3: $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ is an attribute mapping relationship and $\mathrm{R}_{2}\left(T_{2}, T_{3}\right)$ is a join relationship. We can follow a similar rewriting process as in Case 2 to prove for all attributes $\mathrm{A} \in \mathrm{Y}_{3}=\mathrm{Y}_{1} \cap \mathrm{Y}_{3}$ that if $T_{2} . \mathrm{A}=T_{3} . \mathrm{A}$, then $T_{1} \cdot \mathrm{~A} \rightarrow T_{2} . \mathrm{A}$ implies $T_{1}$. A $\rightarrow T_{3}$.A and $T_{2}$. A $\rightarrow T_{1}$.A implies $T_{3}$.A $\rightarrow T_{1}$.A.
- Case 4. $\mathrm{R}_{1}\left(T_{1}, T_{2}\right), \mathrm{R}_{2}\left(T_{2}, T_{3}\right)$ are both attribute mapping relationships. We consider the case where for all attributes $\mathrm{A} \in \mathrm{Y}_{3}=\mathrm{Y}_{1} \cap \mathrm{Y}_{2}, T_{1} . \mathrm{A} \rightarrow T_{2} . \mathrm{A}$ and $T_{2} . \mathrm{A} \rightarrow T_{3} . \mathrm{A}$. Then there exist two queries $Q_{12}$ and $Q_{23}$ over a set of tables such that $T_{2}=Q_{12}\left(T_{2}, \cdots, T_{m}^{\prime}\right), T_{3}=Q_{23}\left(T_{2}, T_{1}^{\prime}, \cdots, T_{n}^{\prime}\right)$ and $\forall y \in \operatorname{dom}(\mathrm{~A})$ :

$$
\begin{array}{r}
\sigma_{\mathrm{A}=y}\left(Q_{12}\left(T_{1}, \cdots, T_{n}^{\prime}\right)\right) \\
=\sigma_{\mathrm{A}=y}\left(Q_{12}\left(\sigma_{\mathrm{A}=y}\left(T_{1}\right), \cdots, T_{n}^{\prime}\right)\right) \tag{A.6}
\end{array}
$$

and

$$
\begin{array}{r}
\sigma_{\mathrm{A}=y}\left(Q_{23}\left(T_{2}, \cdots, T_{m}^{\prime}\right)\right) \\
=\sigma_{\mathrm{A}=y}\left(Q_{23}\left(\sigma_{A=y}\left(T_{2}\right), \cdots, T_{m}^{\prime}\right)\right) \tag{A.7}
\end{array}
$$

We prove that $T_{1} \cdot \mathrm{~A} \rightarrow T_{3} . \mathrm{A}$, i.e., there exists a query $Q_{13}$ and a set of tables $T_{3}=Q_{13}\left(T_{1}, \cdots, T_{k}^{\prime}\right)$ such that $\forall y \in \operatorname{dom}(\mathrm{~A})$,

$$
\begin{array}{r}
\sigma_{\mathrm{A}=y}\left(Q_{13}\left(T_{1}, \cdots, T_{k}^{\prime}\right)\right)  \tag{A.8}\\
=\sigma_{\mathrm{A}=y}\left(Q_{13}\left(\sigma_{\mathrm{A}=y}\left(T_{1}\right), \cdots, T_{k}^{\prime}\right)\right)
\end{array}
$$

By replacing $T_{2}$ by $Q_{12}$ in Equation A.7, we obtain:

$$
\begin{array}{r}
\sigma_{\mathrm{A}=y}\left(Q_{23}\left(T_{2}, \cdots, T_{m}^{\prime}\right)\right) \\
=\sigma_{\mathrm{A}=y}\left(Q_{23}\left(Q_{12}\left(T_{1}, \cdots, T_{n}^{\prime}\right), \cdots, T_{m}^{\prime}\right)\right) \tag{A.9}
\end{array}
$$

By pushing the selection $\sigma_{\mathrm{A}=y}$, we obtain:

$$
\begin{align*}
& \sigma_{\mathrm{A}=y}\left(Q_{23}\left(Q_{12}\left(T_{1}, \cdots, T_{n}^{\prime}\right), \cdots, T_{m}^{\prime}\right)\right)  \tag{A.10}\\
= & Q_{23}\left(\sigma_{\mathrm{A}=y}\left(Q_{12}\left(T_{1}, \cdots, T_{n}^{\prime}\right), \cdots, T_{m}^{\prime}\right)\right)
\end{align*}
$$

Then, by applying Equation A.6, we obtain:

$$
\begin{array}{r}
\sigma_{\mathrm{A}=y}\left(Q_{23}\left(Q_{12}\left(T_{1}, \cdots, T_{n}^{\prime}\right), \cdots, T_{m}^{\prime}\right)\right)  \tag{A.11}\\
=Q_{23}\left(\sigma_{\mathrm{A}=y}\left(Q_{12}\left(\sigma_{\mathrm{A}=y}\left(T_{1}\right), \cdots, T_{n}^{\prime}\right)\right), \cdots, T_{m}^{\prime}\right)
\end{array}
$$

By pulling the first selection $\sigma_{\mathrm{A}=y}$ out from $Q_{23}$, we obtain:

$$
\begin{array}{r}
\sigma_{\mathrm{A}=y}\left(Q_{23}\left(Q_{12}\left(T_{1}, \cdots, T_{n}^{\prime}\right), \cdots, T_{m}^{\prime}\right)\right) \\
=\sigma_{\mathrm{A}=y}\left(Q_{23}\left(Q_{12}\left(\sigma_{\mathrm{A}=y}\left(T_{1}\right), \cdots, T_{n}^{\prime}\right)\right), \cdots, T_{m}^{\prime}\right) \tag{A.12}
\end{array}
$$

Now let $\quad Q_{13}\left(T_{1}, \cdots, T_{n}^{\prime}, \cdots, T_{m}^{\prime}\right) \quad=$ $\left.Q_{23}\left(Q_{12}\left(T_{1}, \cdots, T_{n}^{\prime}\right)\right), \cdots, T_{m}^{\prime}\right)$. Then, we obtain from Equation A.12:

$$
\begin{array}{r}
\sigma_{\mathrm{A}=y}\left(Q_{13}\left(T_{1}, \cdots, T_{n}^{\prime}, \cdots, T_{m}^{\prime}\right)\right)  \tag{A.13}\\
=\sigma_{\mathrm{A}=y}\left(Q_{13}\left(\sigma_{\mathrm{A}=y}\left(T_{1}\right), \cdots, T_{n}^{\prime}, \cdots, T_{m}^{\prime}\right)\right)
\end{array}
$$

Therefore, we have for all attributes in $\mathrm{Y}_{3}=\mathrm{Y}_{1} \cap \mathrm{Y}_{2}$, $T_{1}$. A $\rightarrow T_{3}$.A, i.e. $\mathrm{R}-3\left(T_{1}, T_{3}\right)$ is an attribute mapping relationship with common attributes $\mathrm{Y}_{3}$.

Proposition 6. Fusion of relationships. Let $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ and $\mathrm{R}_{2}\left(T_{1}, T_{2}\right)$ be two well-formed relationships between two tables $T_{1}$ and $T_{2}$ with respective common attributes $\mathrm{Y}_{1}$ and $\mathrm{Y}_{2}$. If $\forall A \in \mathrm{Y}_{1} \cap \mathrm{Y}_{2}, \mu_{\mathrm{R}_{1}}(A)=\mu_{\mathrm{R}_{2}}(A)$ then there exists a well-formed relationship $\mathrm{R}_{3}\left(T_{1}, T_{2}\right)$ that is a fusion of $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ and $\mathrm{R}_{2}\left(T_{1}, T_{2}\right)$ with common attributes $\mathrm{Y}_{3}=\mathrm{Y}_{1} \cup \mathrm{Y}_{2}$.

Proof. We can apply a similar case study as in the proof of Proposition 5. Assume that $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ and $R_{2}\left(T_{1}, T_{2}\right)$ are two natural mapping relationships with common attributes $\mathrm{Y}_{1}$ and $\mathrm{Y}_{2}$, i.e. $\forall \mathrm{A} \in \mathrm{Y}_{1} \cap \mathrm{Y}_{2}, \mathrm{~A}=\nu_{\mathrm{R}_{1}}(\mathrm{~A})=\nu_{\mathrm{R}_{2}}(\mathrm{~A})$. We prove that there exists a well-formed relationship $\mathrm{R}_{3}$ between $T_{1}$ and $T_{2}$ with common attributes $\mathrm{Y}_{3}=\mathrm{Y}_{1} \cup \mathrm{Y}_{2}$. We distinguish three cases by the types of both relation$\operatorname{ships} \mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ and $\mathrm{R}_{2}\left(T_{1}, T_{2}\right)$.

- Case 1: $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ and $\mathrm{R}_{2}\left(T_{1}, T_{2}\right)$ are both join relationships. By definition, for all attributes $\mathrm{A} \in \mathrm{Y}_{3}=\mathrm{Y}_{1} \cup \mathrm{Y}_{2}$, $T_{1} . \mathrm{A}=T_{2} . \mathrm{A}$, i.e. there exists a join relationship $\mathrm{R}_{3}\left(T_{1}, T_{3}\right)$ with common attributes $\mathrm{Y}_{3}$.
- Case 2. $\mathrm{R}_{1}\left(T_{1}, T_{2}\right), \mathrm{R}_{2}\left(T_{1}, T_{2}\right)$ are both attribute mapping relationships. By definition, for all attributes $\mathrm{A} \in \mathrm{Y}_{3}=$ $\mathrm{Y}_{1} \cup \mathrm{Y}_{2}, T_{1} . \mathrm{A} \rightarrow T_{2} . \mathrm{A}$, i.e. there exists an attribute mapping relationship $\mathrm{R}_{3}\left(T_{1}, T_{3}\right)$ with common attributes $\mathrm{Y}_{3}$.
- Case 3: $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ is a join relationship and $\mathrm{R}_{2}\left(T_{1}, T_{2}\right)$ is an attribute mapping relationship. Then, by Case 2, for showing that $\mathrm{R}_{3}\left(T_{1}, T_{2}\right)$ is an attribute mapping relationship it is sufficient to show that $\mathrm{R}_{1}\left(T_{1}, T_{2}\right)$ is an attribute mapping relationship. We show more generally that any
join relationship is also an attribute mapping relationship. By definition, $R_{1}\left(T_{1}, T_{2}\right)$ is a join relationship with common attributes $\mathrm{Y}_{1}$, i.e., for all $\mathrm{A} \in \mathrm{Y}_{1}, T_{1} \cdot \mathrm{~A}=T_{2} . \mathrm{A}$ and $T_{2}=Q_{12}\left(T_{1}, T_{2}\right)=\pi_{S_{2}}\left(T_{1} \bowtie_{Y} T_{2}\right)$ where $S_{2}$ is the schema of $T_{2}$ and $\mathrm{Y} \subseteq \mathrm{S}_{2}$. Then, $\forall y \in \operatorname{dom}(\mathrm{~A})$ :

$$
\begin{array}{r}
\sigma_{\mathrm{A}=y}\left(Q_{12}\left(T_{1}, T_{2}\right)\right) \\
=\sigma_{\mathrm{A}=y}\left(\pi_{S_{2}}\left(T_{1} \bowtie_{\mathrm{Y}_{1}} T_{2}\right)\right)  \tag{A.14}\\
=\sigma_{\mathrm{A}=y}\left(\pi_{S_{2}}\left(\sigma_{\mathrm{A}=y}\left(T_{1}\right) \bowtie_{\mathrm{Y}_{1}} T_{2}\right)\right) \\
=\sigma_{\mathrm{A}=y}\left(Q_{12}\left(\sigma_{\mathrm{A}=y}\left(T_{1}\right), T_{2}\right)\right)
\end{array}
$$

Therefore, R is an attribute mapping relationship.
Proposition 7. Let $T_{0}^{\prime}\left(\mathrm{S}_{0}^{\prime}\right)$ be a merge of table $T_{0}\left(\mathrm{~S}_{0}\right)$ with a target schema augmentation $T(\mathrm{~S})$. Then the following aggregable properties hold for all aggregable attributes $A \in$ $\mathrm{S}_{0}^{\prime}$ :

1. If $\operatorname{agg}_{A}\left(\mathrm{~F}, \mathrm{~V}_{0}\right)$ holds in $T_{0}$ and $A \in \mathrm{~S}_{0}^{\prime} \cap \mathrm{S}_{0}$ then $\operatorname{agg}_{A}\left(\mathrm{~F}, \mathrm{~V}_{0}\right)$ holds in $T_{0}^{\prime}$.
2. If $\operatorname{agg}_{A}(\mathrm{~F}, \mathrm{~V})$ holds in $T$ and $A \in \mathrm{~S}_{0}^{\prime}-\mathrm{S}_{0}$, then $\operatorname{agg}_{A}\left(\mathrm{~F}, \mathrm{~V} \cap \mathrm{~S}_{0}^{\prime}\right)$ holds in $T_{0}^{\prime}$.

Proof. Let $\operatorname{agg}_{\mathrm{A}}\left(F, V_{0}^{\prime}\right)$ be the aggregable property of A in $T_{0}^{\prime}$, i.e. $V_{0}^{\prime} \subseteq S_{0}^{\prime}$ is the maximal set of dimension attributes such that $A$ can be aggregated along $\mathrm{V}_{0}^{\prime}$ with function $F$ in $T_{0}^{\prime}$.

1. Let $\mathrm{U}_{0} \subseteq S_{0}$ be a minimal subset of attributes where $\mathrm{U}_{0} \mapsto A$ in $T_{0}$ and $\mathrm{V}_{0} \subseteq \mathrm{U}_{0}$. Since $\mathrm{V}_{0} \subseteq \mathrm{~S}_{0} \subseteq \mathrm{~S}_{0}^{\prime}$ and $\mathrm{U}_{0} \subseteq \mathrm{~S}_{0}^{\prime}$ we conclude that $\mathrm{U}_{0} \mapsto \mathrm{~A}$ in $T_{0}^{\prime}$ and $\mathrm{V}_{0} \subseteq \mathrm{~V}_{0}^{\prime}$. Since A is not an attribute in $T$, aggregation along $\mathrm{S}-\mathrm{V}$ is not meaningful and we obtain $\mathrm{V}=\mathrm{V}_{0}^{\prime}$
2. Since $T$ is a natural schema complement of $T_{0}, \mathrm{~S}_{0} \cap$ $\mathrm{S} \mapsto \mathrm{S}$. Let $\mathrm{U} \subseteq \mathrm{S}_{0} \cap \mathrm{~S}$ be a minimal subset of attributes where $\mathrm{U} \mapsto \mathrm{A}$ in $T$ and $\mathrm{V} \subseteq \mathrm{U}$. Since $\mathrm{U} \mapsto \mathrm{A}$ and $\mathrm{V} \cap S_{0}^{\prime} \subseteq \mathrm{V} \subseteq \mathrm{U}$, we conclude that $\mathrm{V} \cap \mathrm{S}_{0}^{\prime} \subseteq \mathrm{V}_{0}^{\prime}$. Since A is not an attribute in $T_{0}$, aggregation along $S_{0}-V$ is not meaningful and we obtain $\mathrm{V} \cap \mathrm{S}_{0}^{\prime}=\mathrm{V}_{0}^{\prime}$.

Proposition 8. Let $T(\mathrm{~S})$ be a non-ambiguous analytic table w.r.t. a dimension $D$ of schema $\mathrm{S}_{D}$. Let $T^{\prime}\left(\mathrm{S}^{\prime}\right)=Q(T)$ be a reduction of $T$. Let $\mathrm{X}=\mathrm{S} \cap \mathrm{S}_{D}$ and $\mathrm{X}^{\prime}=\mathrm{S}^{\prime} \cap \mathrm{S}_{D}$. If $\mathrm{X}^{\prime} \mapsto \mathrm{X}$ in $D$, then $T^{\prime}$ is non-ambiguous w.r.t. $D$.

Proof. Let $\mathrm{X}^{\prime}=\mathrm{S}^{\prime} \cap \mathrm{S}_{D}$ and $\mathrm{X}^{*}=\left\{\mathrm{A}_{j} \in \mathrm{~S}_{D} \mid \exists \mathrm{A}_{i} \in\right.$ $\left.\mathrm{X}^{\prime}, \mathrm{A}_{i} \preccurlyeq^{*} \mathrm{~A}_{j}\right\}$. To prove that $T^{\prime}$ is not ambiguous w.r.t. $D$, we must show that $\mathrm{X}^{\prime} \mapsto \mathrm{X}^{\prime *}$ holds in $T^{\prime} \bowtie \mathrm{X}^{\prime} D$ (Definition 14).

We first show that $\mathrm{X}^{\prime} \mapsto \mathrm{X}^{\prime *}$ holds in $T \bowtie \mathrm{X} D$. Let $\mathrm{X}_{D}=\mathrm{S} \cap \mathrm{S}_{D}$ and $\mathrm{X}^{*}=\left\{\mathrm{A}_{j} \in \mathrm{~S}_{D} \mid \exists \mathrm{A}_{i} \in \mathrm{X}, \mathrm{A}_{i} \preccurlyeq^{*} \mathrm{~A}_{j}\right\}$. Since $T$ is non-ambiguous, we know that (a) $\mathrm{X} \mapsto \mathrm{X}^{*}$ holds in $T \bowtie_{\mathrm{X}} D$. Since $\mathrm{X}^{\prime} \mapsto \mathrm{X}$ holds in $D, \mathrm{X}^{\prime} \subseteq \mathrm{X}$ and $\pi_{\mathrm{X}}\left(T \bowtie_{\mathrm{X}}\right.$ $D) \subseteq \pi_{\mathrm{X}}(D)$, we also obtain (b) $\mathrm{X}^{\prime} \mapsto \mathrm{X}$ holds in $T \bowtie \mathrm{X} D$. By transitivity of $\mapsto$ and (a) and (b), we obtain (c) $\mathrm{X}^{\prime} \mapsto \mathrm{X}^{*}$
holds in $T \bowtie_{\mathrm{X}} D$, and since $\mathrm{X}^{\prime *} \subseteq \mathrm{X}^{*}$, (d) $\mathrm{X}^{\prime} \mapsto \mathrm{X}^{\prime *}$ holds in $T \bowtie_{\mathrm{X}} D$.

We now show that $\pi_{\mathrm{S}_{D}}\left(T^{\prime} \bowtie_{\mathrm{X}^{\prime}} D\right) \subseteq \pi_{\mathrm{S}_{D}}\left(T \bowtie_{\mathrm{X}} D\right)$. Since $\mathrm{X}^{\prime} \mapsto \mathrm{X}$ holds in $D$ and $\mathrm{X}^{\prime} \subseteq \mathrm{X} \subseteq \mathrm{S}_{D}$, we know that $\pi_{\mathrm{S}_{D}}\left(\pi_{X^{\prime}}(T) \bowtie_{\mathrm{X}^{\prime}} D\right)=\pi_{\mathrm{S}_{D}}\left(T \bowtie_{\mathrm{X}} D\right)$ and since $\pi_{X^{\prime}}\left(T^{\prime}\right) \subseteq$ $\pi_{X^{\prime}}(T)$, we obtain (e) $\pi_{\mathrm{S}_{D}}\left(T^{\prime} \bowtie \mathrm{x}^{\prime} D\right) \subseteq \pi_{\mathrm{S}_{D}}(T \bowtie \mathrm{X} D)$.

Then, from (d) and (e), we can conclude that $\mathrm{X}^{\prime} \mapsto \mathrm{X}^{\prime *}$ holds in $T^{\prime} \bowtie_{\mathrm{X}^{\prime}} D$

Proposition 9. Let $T_{0}\left(\mathrm{~S}_{0}\right)$ and $T(\mathrm{~S})$ be two analytic tables with a relationship R and common dimension attributes Y . Let $T_{0}^{\prime}\left(\mathrm{S}_{0}^{\prime}\right)$ be the merge of $T_{0}$ and $T$ and $T^{\text {miss }}=$ $T^{\text {cand }}-T^{c t}$ be the set of all tuples in $T$ lost in the merge with $T_{0}$. We can define a completion table $T^{\text {com }}\left(S_{0}^{\prime}\right)$ for all dimensions $D_{i}$ of schema $\mathrm{S}_{D_{i}}$ in $T_{0}, 0 \leq i \leq n$, where $\left(\mathrm{S}_{D_{i}} \cap \mathrm{~S}\right) \subset\left(\mathrm{S}_{D_{i}} \cap \mathrm{~S}_{0}\right):$

$$
\begin{equation*}
T^{c o m}=\Pi_{S_{0}^{\prime}}\left(T^{m i s s} \bowtie_{\mathrm{S}_{D_{0}} \cap \mathrm{Y}} D_{0} \perp \triangleleft \cdots \bowtie_{\mathrm{S}_{D_{n}} \cap \mathrm{Y}} D_{n}\right) \tag{1}
\end{equation*}
$$

If $\mathrm{Y} \mapsto \mathrm{S}_{0}$ (non-ambiguous merge condition), then

$$
\begin{equation*}
Q_{m}\left(T_{0}, T\right)=T_{0}^{\prime} \cup T^{c o m} \tag{2}
\end{equation*}
$$

is a complete merge of $T_{0}$ and $T$ with respect to $T$.
Proof. Let $T^{c t}$ and $T^{\text {cand }}$ be the completion table and candidate completion table of $T$ and $T_{0}^{\prime}$ as defined in Definition 15. We use Definition 16 to prove that $T_{0}^{\prime}=\mathcal{Q}_{m}\left(T_{0}, T\right)$ is a complete merge w.r.t. to $T$ :

$$
T^{c t}=T \bowtie_{\mathrm{Y}} \mathcal{Q}_{m}\left(T_{0}, T\right)=T \bowtie_{\mathrm{Y} t o p} \mathcal{Q}_{m}\left(T_{0}, T\right)=T^{c a n d}
$$

We can replace $Q_{m}\left(T_{0}, T\right)$ with the right-hand side of Equation (2) and then replace $T^{c o m}$ with the right-hand side of Equation (2) to obtain:

$$
\begin{align*}
T^{c t} & =T \ltimes_{\mathrm{Y}}\left(T_{0}^{\prime} \cup T^{c o m}\right)  \tag{A.15}\\
& =T^{c t} \cup\left(T \ltimes_{\mathrm{Y}} T^{c o m}\right)
\end{align*}
$$

Equation A. 15 uses $T^{c o m}$ which extends $T_{\text {miss }}$ with all dimension attributes A in $S_{0}^{\prime}$ that don't exist in the augmentation schema $S$, i.e., $A \in S_{0}^{\prime}-S$. Then, since all these attributes are not common attributes of $T_{0}$ and $T$, i.e., $\mathrm{Y} \cap\left(\mathrm{S}_{0}^{\prime}-\mathrm{S}\right)=\emptyset$, their values do not change the result of the natural semi-join $T \ltimes_{\mathrm{Y}} T^{c o m}$. Therefore, since $\mathrm{Y} \subseteq \mathrm{S}_{0}^{\prime}$ we can safely remove the projection $\pi_{S_{0}^{\prime}}$ and replace $T^{\text {com }}$ by $T^{\text {miss }}$ in Equation A.15:

$$
\begin{equation*}
T^{c t}=T^{c t} \cup\left(T \ltimes_{\mathrm{Y}} T^{m i s s}\right) \tag{A.16}
\end{equation*}
$$

When we replace $T^{\text {miss }}$ with $T^{\text {miss }}=T^{\text {cand }}-T^{c t}$, we obtain:

$$
\begin{align*}
T^{c t} & =T^{c t} \cup\left(T \ltimes_{\mathrm{Y}}\left(T^{\text {cand }}-T^{c t}\right)\right) \\
& =T^{c t} \cup\left(\left(T \ltimes_{\mathrm{Y}} T^{c a n d}\right)-\left(T \ltimes_{\mathrm{Y}} T^{c t}\right)\right)  \tag{A.17}\\
& =\left(T^{c t} \cup\left(T \ltimes_{\mathrm{Y}} T^{\text {cand }}\right)\right)-\left(T^{c t} \cup\left(T \ltimes_{\mathrm{Y}} T^{c t}\right)\right)
\end{align*}
$$

By Definition 15 , it is easy to show that $T^{c t} \subseteq T^{\text {cand }} \subseteq$ $T$. We then obtain $T \ltimes_{\mathrm{Y}} T^{c t}=T^{c t}$ and $T \ltimes_{\mathrm{Y}} T^{c a n \bar{d}}=T^{c a n \bar{d}}$ and continue with Equation A.17:

$$
\begin{align*}
T^{c t} & =\left(T^{c t} \cup T^{c a n d}\right) \cup\left(T^{c t}-T^{c t}\right) \\
& =T^{c a n d} \cup \emptyset=T^{c a n d} \tag{A.18}
\end{align*}
$$

By $T^{c t}=T^{\text {cand }}$, we can conclude that $Q_{m}$ is a complete merge with respect to $T$.

Lemma 1 (Composition of natural schema complements). Let $T(S)$ be a natural schema complement to $T_{0}\left(S_{0}\right)$ with respect to relationship $\mathrm{R}_{0}\left(T_{0}, T\right)$ on attributes $\mathrm{Y}_{0}$, and $T_{1}\left(S_{1}\right)$ be a natural schema complement to $T$ with respect to relationship $\mathrm{R}_{1}\left(T, T_{1}\right)$ on attributes $\mathrm{Y}_{1}$. Then $T_{1}$ is also $a$ natural schema complement to the natural merge of $T_{0}$ and $T$, i.e., $T_{0} \triangle_{\mathrm{Y}_{0}} T$.

Proof. Let $T_{0}^{\prime}\left(S_{0}^{\prime}\right)$ be the natural merge of $T_{0}$ and $T, T_{0}^{\prime}=$ $T_{0} \bowtie_{\mathrm{Y}_{0}} T$.

There exists an attribute mapping relationship $\mathrm{R}^{\prime}\left(T_{0}^{\prime}, T\right)$ on common attributes $\mathrm{Y}^{\prime}=\mathrm{S}_{0}^{\prime} \cap \mathrm{S}$, and $\mathrm{Y}_{1} \subseteq \mathrm{Y}^{\prime}$. And there exists a well-formed relationship $\mathrm{R}_{1}^{\prime}\left(T_{0}^{\prime}, T_{1}\right)$ on attributes $\mathrm{Y}_{1} \cap \mathrm{~S}_{0}^{\prime}$ by the composition of relationships $\mathrm{R}^{\prime}\left(T_{0}^{\prime}, T\right)$ and $\mathrm{R}_{1}\left(T, T_{1}\right)$. Because $T_{1}$ is a natural schema complement to $T$, we have $\mathrm{Y}_{1} \mapsto \mathrm{~S}_{1}$ and $\mathrm{Y}^{\prime} \mapsto \mathrm{S}_{1}$. Therefore, $T_{1}$ is a natural schema complement to $T_{0}^{\prime}$.

Lemma 2 (Composition of schema complement and schema augmentation). Let $T(\mathrm{~S})$ be a natural schema complement to $T_{0}\left(\mathrm{~S}_{0}\right)$ with respect to relationship $\mathrm{R}_{0}\left(T_{0}, T\right)$ on attributes $\mathrm{Y}_{0}$, and $T_{1}\left(\mathrm{~S}_{1}\right)$ be a schema augmentation to $T$ with respect to relationship $\mathrm{R}_{1}\left(T, T_{1}\right)$ on attributes $\mathrm{Y}_{1}$. Then $T_{1}$ is a schema augmentation to the natural merge of $T_{0}$ and $T$, i.e., $T_{0} \bowtie_{\mathrm{Y}_{0}} T$.
Proof. Let $T_{0}^{\prime}\left(S_{0}^{\prime}\right)$ be the natural merge of $T_{0}$ and $T, T_{0}^{\prime}=$ $T_{0} \searrow_{\mathrm{Y}_{0}} T$.

There exists an attribute mapping relationship $\mathrm{R}^{\prime}\left(T_{0}^{\prime}, T\right)$ on common attributes $\mathrm{Y}^{\prime}=\mathrm{S}_{0}^{\prime} \cap \mathrm{S}$, and $\mathrm{Y}_{1} \subseteq \mathrm{Y}^{\prime}$. By the composition of relationships $\mathrm{R}^{\prime}\left(T_{0}^{\prime}, T\right)$ and $\mathrm{R}_{1}\left(T, T_{1}\right)$, there exists a well-formed relationship $\mathrm{R}_{1}^{\prime}\left(T_{0}^{\prime}, T_{1}\right)$ on attributes $\mathrm{Y}_{1} \cap \mathrm{~S}_{0}^{\prime}$. Therefore, $T_{1}$ is a schema augmentation to $T_{0}^{\prime}$.

Proposition 10. Let query $Q$ be the result of a $M S A\left(T_{0}\right.$, $T$, path, true, true) call with a start table $T_{0}\left(\mathrm{~S}_{0}\right)$, a target schema augmentation table $T(\mathrm{~S})$ and a path of relationships: path $=\mathrm{R}_{1}\left(T_{0}, T_{1}\right), \cdots, \mathrm{R}_{n}\left(T_{n-1}, Q_{a}\right), n>=1$. If $\mathrm{R}_{n}\left(T_{n-1}, \mathcal{Q}_{a}\right)$ maps to an augmentation schema complement edge ( $C T=$ ' $A U G$ '), then $Q$ computes a augmented merge of $T_{0}$ with $Q_{a}$ (without ambiguous values). Otherwise, $\mathrm{R}_{n}\left(T_{n-1}, Q_{a}\right)$ maps to a natural schema complement edge ( $C T=$ ' $N A T$ ') and $Q$ computes a natural merge $T_{0}$ with $Q_{a}$.

Proof. We can ignore steps 1 and 2 of MSA. Step 3 returns merge query $Q=\pi_{\mathrm{S}_{0}^{\prime}}\left(Q_{\text {path }} \searrow_{\mathrm{Y}_{n}} Q_{a}\right)$, where $\mathrm{S}_{0}^{\prime}$ is augmented S with attributes from $Q_{a}, Q_{\text {path }}=T_{0} \searrow_{\mathrm{Y}_{1}} T_{1}$
$\cdots \bowtie_{\mathrm{Y}_{n-1}} T_{n-1}$ and $\mathrm{Y}_{i}$ is the common attributes in relationship $\mathrm{R}_{i}$. By Algorithm 3, every relationship in $\mathrm{R}_{1}\left(T_{0}, T_{1}\right), \cdots, \mathrm{R}_{n-1}\left(T_{n-2}, T_{n-1}\right)$ maps a $S C$ edge with $C T=$ ' $N A T$ ' in $S C$. By Lemma $1, Q_{\text {path }}$ is a natural merge of $T_{0}$ with its natural schema complements $T_{1}, \cdots, T_{n-1}$. Also by Lemma 1, if $\mathrm{R}_{n}\left(T_{n-1}, T\right)$ maps a $\mathcal{S C}$ edge with $C T=$ ' $N A T$ ' in $\mathcal{S C}$, then $Q_{a}$ is a natural schema complement to the result of $Q_{\text {path }}$ and $Q$ computes the natural merge of $Q_{\text {path }}$ and $T$. Finally, by the same lemma, since $Q_{\text {path }}$ is a natural merge of $T_{0}, Q$ is a natural merge of $T_{0}$ and $\mathcal{Q}_{a}$. By Lemma 2, if $\mathrm{R}_{a}\left(T_{a-1}, T\right)$ maps to an augmentation $S C$ edge with $C T=$ ' $A U G$ ', then $Q_{a}$ is a schema augmentation to the result of $Q_{\text {path }}$ and $Q$ computes the augmented merge of $Q_{\text {path }}$ and $Q_{a}$ (and an augmented merge of $T_{0}$ and $Q_{a}$ ).


[^0]:    *This work was supported by the French Association Nationale de la Recherche Technique (ANRT) [grant CIFRE number 2016/0644];
    *Corresponding author
    Email addresses: rutian.liu@lip6.fr (Rutian Liu), eric.simon@sap.com (Eric Simon), bernd.amann@lip6.fr (Bernd Amann), stephane.gancarski@lip6.fr (Stéphane Gançarski)

