





## ROBUSTNESS AGAINST READ COMMITTED FOR TRANSACTION TEMPLATES WITH FUNCTIONAL CONSTRAINTS

BRECHT VANDEVOORT <sup>a</sup>, BAS KETSMAN <sup>b</sup>, CHRISTOPH KOCH <sup>c</sup>, AND FRANK NEVEN <sup>a</sup>

<sup>a</sup>UHasselt, Data Science Institute, ACSL, Belgium  
*e-mail address:* brecht.vandevooort@uhasselt.be, frank.neven@uhasselt.be

<sup>b</sup>Vrije Universiteit Brussel, Belgium  
*e-mail address:* bas.ketsman@vub.be

<sup>c</sup>École Polytechnique Fédérale de Lausanne, Switzerland  
*e-mail address:* christoph.koch@epfl.ch

**ABSTRACT.** The popular isolation level Multiversion Read Committed (RC) trades some of the strong guarantees of serializability for increased transaction throughput. Sometimes, transaction workloads can be safely executed under RC obtaining serializability at the lower cost of RC. Such workloads are said to be robust against RC. Previous work has yielded a tractable procedure for deciding robustness against RC for workloads generated by transaction programs modeled as transaction templates. An important insight of that work is that, by more accurately modeling transaction programs, we are able to recognize larger sets of workloads as robust. In this work, we increase the modeling power of transaction templates by extending them with functional constraints, which are useful for capturing data dependencies like foreign keys. We show that the incorporation of functional constraints can identify more workloads as robust than otherwise would not be. Even though we establish that the robustness problem becomes undecidable in its most general form, we show that various restrictions on functional constraints lead to decidable and even tractable fragments that can be used to model and test for robustness against RC for realistic scenarios.

### 1. INTRODUCTION

Many database systems implement several isolation levels, allowing users to trade isolation guarantees for improved performance. The highest, serializability, projects the appearance of a complete absence of concurrency, and thus perfect isolation. Executing transactions concurrently under weaker isolation levels can introduce certain anomalies. Sometimes, a transactional workload can be executed at an isolation level lower than serializability without introducing any anomalies. This is a desirable scenario: a lower isolation level, usually implementable with a cheaper concurrency control algorithm, yields the stronger isolation guarantees of serializability for free. This formal property is called robustness [Fek05, BG16]: a set of transactions  $\mathcal{T}$  is called *robust against a given isolation level* if every

*Key words and phrases:* concurrency control, robustness, complexity.

\* The present paper is the full version of [VKKN22] and supplies all proofs.

possible interleaving of the transactions in  $\mathcal{T}$  that is allowed under the specified isolation level is serializable.

Robustness received quite a bit of attention in the literature. Most existing work focuses on Snapshot Isolation (SI) [ACFR08, BBE19a, Fek05, FLO<sup>+</sup>05] or higher isolation levels [BBE19b, BG16, CBG15, CGY17]. It is particularly interesting to consider robustness against lower level isolation levels like multi-version Read Committed (referred to as RC from now on). Indeed, RC is widely available, often the default in database systems (see, e.g., [BDF<sup>+</sup>13]), and is generally expected to have better throughput than stronger isolation levels.

In previous work [VKKN21], we provided a tractable decision procedure for robustness against RC for workloads generated by transaction programs modeled as transaction templates. The approach is centered on a novel characterization of robustness against RC in the spirit of [Fek05, KKNV20] that improves over the sufficient condition presented in [AF15], and on a formalization of transaction programs, called *transaction templates*, facilitating fine-grained reasoning for robustness against RC. Conceptually, transaction templates as introduced in [VKKN21] are functions with parameters, and can, for instance, be derived from stored procedures inside a database system (c.f. Figure 1 for an example). The abstraction generalizes transactions as usually studied in concurrency control research – sequences of read and write operations – by making the objects worked on variable, determined by input parameters. Such parameters are *typed* to add additional power to the analysis. They support *atomic updates* (that is, a read followed by a write of the same database object, to make a relative change to its value). Furthermore, database objects read and written are considered at the granularity of fields, rather than just entire tuples, decoupling conflicts further and allowing to recognize additional cases that would not be recognizable as robust on the tuple level.

An important insight obtained from [VKKN21] is that more accurate modeling of the workload allows to recognize larger sets of transaction programs as robust. Processing workloads under RC increases the throughput of the transactional database system compared to when executing the workload under SI or serializable SI, so larger robust sets mean better performance of the database system. In this work, we increase the modeling power of transaction templates by extending them with *functional constraints*, which are useful for capturing data dependencies like foreign keys (inclusion dependencies). This appears to be a sweet spot for strengthening modelling power – as we show in this paper, it allows us to remain with abstractions that have been well established within database theory, without having to move to general program analysis, and it pushes the robustness frontier on popular transaction processing benchmarks. Generally speaking, workloads can profit more from richer modelling the larger and more complex they get, so the fact that adding functional constraints yields larger robust sets already on these simple benchmarks suggests that these techniques are practically useful. Our contributions can be summarized as follows:

- We argue in Section 2 through the SmallBank and TPC-C benchmarks that the incorporation of functional constraints can identify more workloads as robust that otherwise would not be, and that they reduce the extent to which changes need to be made to workloads to make them robust against RC.
- In Section 4, we establish that robustness in its most general form becomes undecidable. The proof is a reduction from PCP and relies on cyclic dependencies between functions allowing to connect data values through an unbounded application of functions.

- We consider a fragment in Section 5 that only allows a very limited form of cyclic dependencies between functions and assumes additional constraints on templates that, together, imply that functions behave as bijections. Robustness against RC can be decided in NLOGSPACE and this fragment is general enough to model the SmallBank benchmark.
- In Section 6, we obtain an EXPSPACE decision procedure when the schema graph is acyclic (so, no cyclic dependencies between functions). Even for small input sizes, such a result is not practical. We provide various restrictions that lower the complexity to PSPACE and EXPTIME, and which allow to model the TPC-C benchmark as discussed. Notice that, for robustness testing, an exponential time decision procedure is considered to be practical as the size of the input is small and robustness is a static property that can be tested offline.

These contributions should be contrasted with our earlier work [VKKN21], where we focused on a characterization for robustness against RC for basic transaction templates without functional constraints and performed an experimental study to show how the robustness property can improve transaction throughput.

## 2. BENCHMARKS

We present a small extension of the SmallBank benchmark [ACFR08] to exemplify the modeling power of transaction templates and discuss how the addition of functional constraints can detect larger sets of transaction templates to be robust. Finally, we discuss in the context of the TPC-C benchmark how the incorporation of functional constraints requires less changes to templates in making them robust.

The SmallBank schema consists of three tables: *Account*(Name, *CustomerID*, *IsPremium*), *Savings*(CustomerID, *Balance*, *InterestRate*), and *Checking*(CustomerID, *Balance*). Underlined attributes are primary keys. The *Account* table associates customer names with IDs and keeps track of the premium status (Boolean); *CustomerID* is a UNIQUE attribute. The other tables contain the balance (numeric value) of the savings and checking accounts of customers identified by their ID. *Account* (*CustomerID*) is a foreign key referencing both the columns *Savings* (*CustomerID*) and *Checking* (*CustomerID*). The interest rate on a savings account is based on a number of parameters, including the account status (premium or not). The application code can interact with the database through a fixed number of transaction programs:

- *Balance*( $N$ ): returns the total balance (savings & checking) for a customer with name  $N$ .
- *DepositChecking*( $N, V$ ): makes a deposit of amount  $V$  on the checking account of the customer with name  $N$ .
- *TransactSavings*( $N, V$ ): makes a deposit or withdrawal  $V$  on the savings account of the customer with name  $N$ .
- *Amalgamate*( $N_1, N_2$ ): transfers all the funds from  $N_1$  to  $N_2$ .
- *WriteCheck*( $N, V$ ): writes a check  $V$  against the account of the customer with name  $N$ , penalizing if overdrawing.
- *GoPremium*( $N$ ): converts the account of the customer with name  $N$  to a premium account and updates the interest rate of the corresponding savings account. This transaction program is an extension w.r.t. [ACFR08].

The transaction templates for these programs are presented in Figure 1. The corresponding SQL code is given in Appendix A.

Balance: $R[X : \text{Account}\{N, C\}]$ $R[Y : \text{Savings}\{C, B\}]$ $R[Z : \text{Checking}\{C, B\}]$ $Y = f_{A \rightarrow S}(X), X = f_{S \rightarrow A}(Y)$ $Z = f_{A \rightarrow C}(X), X = f_{C \rightarrow A}(Z)$	DepositChecking: $R[X : \text{Account}\{N, C\}]$ $U[Z : \text{Checking}\{C, B\}\{B\}]$ $Z = f_{A \rightarrow C}(X), X = f_{C \rightarrow A}(Z)$	TransactSavings: $R[X : \text{Account}\{N, C\}]$ $U[Y : \text{Savings}\{C, B\}\{B\}]$ $Y = f_{A \rightarrow S}(X), X = f_{S \rightarrow A}(Y)$
Amalgamate: $R[X_1 : \text{Account}\{N, C\}]$ $R[X_2 : \text{Account}\{N, C\}]$ $U[Y_1 : \text{Savings}\{C, B\}\{B\}]$ $U[Z_1 : \text{Checking}\{C, B\}\{B\}]$ $U[Z_2 : \text{Checking}\{C, B\}\{B\}]$ $X_1 \neq X_2,$ $Y_1 = f_{A \rightarrow S}(X_1), X_1 = f_{S \rightarrow A}(Y_1)$ $Y_2 = f_{A \rightarrow S}(X_2), X_2 = f_{S \rightarrow A}(Y_2)$ $Z_1 = f_{A \rightarrow C}(X_1), X_1 = f_{C \rightarrow A}(Z_1)$ $Z_2 = f_{A \rightarrow C}(X_2), X_2 = f_{C \rightarrow A}(Z_2)$	WriteCheck: $R[X : \text{Account}\{N, C\}]$ $R[Y : \text{Savings}\{C, B\}]$ $R[Z : \text{Checking}\{C, B\}]$ $U[Z : \text{Checking}\{C, B\}\{B\}]$ $Y = f_{A \rightarrow S}(X), X = f_{S \rightarrow A}(Y)$ $Z = f_{A \rightarrow C}(X), X = f_{C \rightarrow A}(Z)$	GoPremium: $U[X : \text{Account}\{N, C\}\{I\}]$ $R[Y : \text{Savings}\{C, I\}]$ $U[Y : \text{Savings}\{C\}\{I\}]$ $Y = f_{A \rightarrow S}(X), X = f_{S \rightarrow A}(Y)$

FIGURE 1. Transaction templates for SmallBank.

Based on this benchmark, we give an informal description of transaction templates and functional constraints to illustrate their modeling power. More formal definitions can be found in Section 3. In short, a transaction template is a sequence of read (R), write (W) and update (U) statements over typed variables ( $X, Y, \dots$ ) with additional equality and disequality constraints. For instance,  $R[Y : \text{Savings}\{C, I\}]$  in GoPremium indicates that a read operation is performed to a tuple in relation *Savings* on the attributes *CustomerID* and *InterestRate*. We abbreviate the names of attributes by their first letter to save space. The set  $\{C, I\}$  is the read set. Write operations have an associated write set while update operations contain a read set followed by a write set: e.g.,  $U[X : \text{Account}\{N, C\}\{I\}]$  in GoPremium first reads the *Name* and *CustomerID* of tuple  $X$  and then writes to the attribute *InterestRate*. To capture the dependencies between tuples induced by the foreign keys, we use two unary functions:  $f_{A \rightarrow S}$  maps a tuple of type *Account* to a tuple of type *Savings*, while  $f_{A \rightarrow C}$  maps a tuple of type *Account* to a tuple of type *Checking*. As *Account(CustomerID)* is **UNIQUE**, every savings and checking account is associated to a unique *Account* tuple. This is modelled through the functions  $f_{C \rightarrow A}$  and  $f_{S \rightarrow A}$  with an analogous interpretation. Notice that the equality constraints for each template in Figure 1 imply that these functions are bijections and each others inverses.

A transaction  $T$  over a database  $\mathbf{D}$  is an *instantiation* of a transaction template  $\tau$  if there is a variable mapping  $\mu$  from the variables in  $\tau$  to tuples in  $\mathbf{D}$  that satisfies all the constraints in  $\tau$  such that  $\mu(\tau) = T$ . For instance, consider a database  $\mathbf{D}$  with tuples  $\mathbf{a}_1, \mathbf{a}_2, \dots$  of type *Account*,  $\mathbf{s}_1, \mathbf{s}_2, \dots$  of type *Savings*, and  $\mathbf{c}_1, \mathbf{c}_2, \dots$  of type *Checking* with  $f_{A \rightarrow S}^{\mathbf{D}}(\mathbf{a}_i) = \mathbf{s}_i$ ,  $f_{A \rightarrow C}^{\mathbf{D}}(\mathbf{a}_i) = \mathbf{c}_i$ ,  $f_{S \rightarrow A}^{\mathbf{D}}(\mathbf{s}_i) = \mathbf{a}_i$ ,  $f_{C \rightarrow A}^{\mathbf{D}}(\mathbf{c}_i) = \mathbf{a}_i$  for each  $i$ . Then, for  $\mu_1 = \{X \rightarrow \mathbf{a}_1, Y \rightarrow \mathbf{s}_1\}$ ,  $\mu_1(\text{GoPremium}) = U[\mathbf{a}_1]R[\mathbf{s}_1]U[\mathbf{s}_1]$  is an instantiation of GoPremium whereas  $\mu_2(\text{GoPremium})$  with  $\mu_2 = \{X \rightarrow \mathbf{a}_1, Y \rightarrow \mathbf{s}_2\}$  is not as the functional constraint  $Y = f_{A \rightarrow S}(X)$  is not satisfied. Indeed,  $\mu_2(Y) = \mathbf{s}_2 \neq \mathbf{s}_1 = f_{A \rightarrow S}^{\mathbf{D}}(\mathbf{a}_1) = f_{A \rightarrow S}^{\mathbf{D}}(\mu_2(X))$ . We then say that a set of transactions is *consistent* with a set of templates if every transaction is an instantiation of a transaction template. More formal definitions are given in Section 3.

Functional constraints are different from the more usual data consistency constraints like key constraints, functional dependencies or denial constraints, etc. The latter are intended to verify data consistency, whereas the former are intended to verify whether a set of transactions instantiated from templates are indeed consistent with these templates. The abstraction of functional constraints provides a straightforward mechanism to capture dependencies between tuples implied by e.g. foreign key constraints. Consider for example variables  $X$  and  $Y$  in GoPremium. Rather than specifying that the value of the attribute *CustomerID* in the tuple assigned to  $X$  should agree with the value of the attribute *CustomerID* in the tuple assigned to  $Y$  and combining this information with the defined foreign key from *Account* to *Savings* to conclude that two instantiations of GoPremium that agree on the tuple assigned to  $X$  should also agree on the tuple assigned to  $Y$ , the functional constraint  $Y = f_{A \rightarrow S}(X)$  expresses this dependency more directly. An additional benefit of our abstraction is that this approach is not limited to dependencies implied by foreign keys. For the SmallBank benchmark, for example, we can infer from the fact that *Account(CustomerID)* is **UNIQUE** that each checking and savings account is associated to exactly one *Account* tuple, even though no foreign key from respectively *Checking* and *Savings* to *Account* is defined in the schema. Since functional constraints are based on unary functions, they are limited to expressing tuples being implied by a single other tuple (e.g., the *Savings* tuple being implied by the *Account* tuple in GoPremium). More complex relationships where a tuple is implied by the co-occurrence of two or more other tuples cannot be captured by our formalism.

Our previous work [VKKN21], which did not consider functional constraints, has shown that  $\{\text{Am,DC,TS}\}$ ,  $\{\text{Bal,DC}\}$ , and  $\{\text{Bal,TS}\}$  are maximal robust sets of transaction templates. This means that for any database, for any set of transactions  $\mathcal{T}$  that is consistent with one of the three mentioned sets, any possible interleaving of the transactions in  $\mathcal{T}$  that is allowed under RC is *always* serializable! Using the results from Section 5, it follows that when functional constraints are taken into account GoPremium can be added to each of these sets as well:  $\{\text{Am,DC,GP,TS}\}$ ,  $\{\text{Bal,DC,GP}\}$ ,  $\{\text{Bal,TS,GP}\}$  are maximal robust sets.

We argue that incorporating functional constraints is crucial. Indeed, without functional constraints it's easy to show that even the set  $\{\text{GoPremium}\}$  is not robust. Consider the schedule over two instantiations  $T_1$  and  $T_2$  of GoPremium, where we use the mappings  $\mu_1$  and  $\mu_2$  as defined above for respectively  $T_1$  and  $T_2$  (we show the read and write sets to facilitate the discussion):

$$\begin{array}{l} T_1 : \text{U}_1[\mathbf{a}_1\{\text{N,C}\}\{\text{I}\}] \text{R}_1[\mathbf{s}_1\{\text{C,I}\}] \qquad \qquad \qquad \text{U}_1[\mathbf{s}_1\{\text{C}\}\{\text{I}\}] \text{C}_1 \\ T_2 : \qquad \qquad \qquad \text{U}_2[\mathbf{a}_2\{\text{N,C}\}\{\text{I}\}] \text{R}_2[\mathbf{s}_1\{\text{C,I}\}] \text{U}_2[\mathbf{s}_1\{\text{C}\}\{\text{I}\}] \text{C}_2 \end{array}$$

The above schedule is allowed under RC as there is no dirty write, but it is not conflict serializable. Indeed, there is a rw-conflict between  $\text{R}_1[\mathbf{s}_1\{\text{C,I}\}]$  and  $\text{U}_2[\mathbf{s}_1\{\text{C}\}\{\text{I}\}]$  as the former reads the attribute  $I$  that is written to by the latter, which implies that  $T_1$  should occur before  $T_2$  in an equivalent serial schedule. But, there is a ww-conflict between  $\text{U}_2[\mathbf{s}_1\{\text{C}\}\{\text{I}\}]$  and  $\text{U}_1[\mathbf{s}_1\{\text{C}\}\{\text{I}\}]$  as both write to the common attribute  $I$  implying that  $T_2$  should occur before  $T_1$  in an equivalent serial schedule. Consequently, the schedule is not serializable. However, taking functional constraints into account,  $\{T_1, T_2\}$  is not consistent with  $\{\text{GoPremium}\}$  as  $\mu_2(Y) = \mathbf{s}_1 \neq \mathbf{s}_2 = f_{A \rightarrow S}(a_2) = f_{A \rightarrow S}(\mu_2(X))$  implying that the above schedule is *not* a counterexample for robustness.

The second benchmark is based on the TPC-C benchmark [TC]. We modified the schema and templates to turn all predicate reads into key-based accesses. The schema consists of six relations:

$f$	$dom(f)$	$range(f)$
$f_{D \rightarrow W}$	District	Warehouse
$f_{C \rightarrow D}$	Customer	District
$f_{O \rightarrow C}$	Order	Customer
$f_{L \rightarrow O}$	OrderLine	Order
$f_{L \rightarrow S}$	OrderLine	Stock
$f_{S \rightarrow W}$	Stock	Warehouse

TABLE 1. Function names for the TPC-C benchmark schema.

- $Warehouse(WarehouseID, Info, YTD)$ ,
- $District(WarehouseID, DistrictID, Info, YTD, NextOrderID)$ ,
- $Customer(WarehouseID, DistrictID, CustID, Info, Balance)$ ,
- $Order(WarehouseID, DistrictID, OrderID, CustID, Status)$ ,
- $OrderLine(WarehouseID, DistrictID, OrderID, OrderLineID, ItemID, DeliveryInfo, Quantity)$ , and
- $Stock(WarehouseID, ItemID, Quantity)$ .

The function names belonging to this schema are given in Table 1.

We focus on five different transaction templates:

- $NewOrder(W, D, C, I_1, Q_1, I_2, Q_2, \dots)$ : creates a new order for the customer identified by  $(W, D, C)$ . The id for this order is obtained by increasing the  $NextOrderID$  attribute of the  $District$  tuple identified by  $(W, D)$  by one. Each order consists of a number of items  $I_1, I_2, \dots$  with respectively quantities  $Q_1, Q_2, \dots$ . For each of these items, a new  $OrderLine$  tuple is created and the related stock quantity is decreased.
- $Payment(W, D, C, A)$ : represents a customer identified by  $(W, D, C)$  paying an amount  $A$ . This payment is reflected in the database by increasing the balance of this customer by  $A$ . This amount is furthermore added to the YearToDate ( $YTD$ ) income of both the related  $Warehouse$  and  $District$  tuples.
- $OrderStatus(W, D, C, O)$ : requests information about the current status of the order identified by  $(W, D, O)$ . This transaction template collects information of the customer identified by  $(W, D, C)$  who created the order, the  $Order$  tuple itself, and the different  $OrderLine$  tuples related to this order.
- $Delivery(W, D, C, O)$ : delivers the order represented by  $(W, D, O)$ . The status of the order is updated, as well as the  $DeliveryInfo$  attribute of each  $OrderLine$  tuple related to this order. The total price of the order is deduced from the balance of the customer who made this order, identified by  $(W, D, C)$ .
- $StockLevel(W, I)$ : returns the current stock level of item  $I$  in warehouse  $W$ .

A detailed abstraction of each transaction template is given in Figure 2. To shorten the presentation, we only show two orderlines per order.

Incorporating functional constraints for TPC-C can not identify larger sets of templates to be robust. However, when a set of transaction templates  $\mathcal{P}$  is not robust against RC, an equivalent set of templates  $\mathcal{P}'$  can be constructed from  $\mathcal{P}$  by *promoting* certain R-operations to U-operations [VKKN21]. By incorporating functional constraints it can be shown that fewer R-operations need to be promoted leading to an increase in throughput as R-operations do not take locks whereas U-operations do. Consider for example the subset

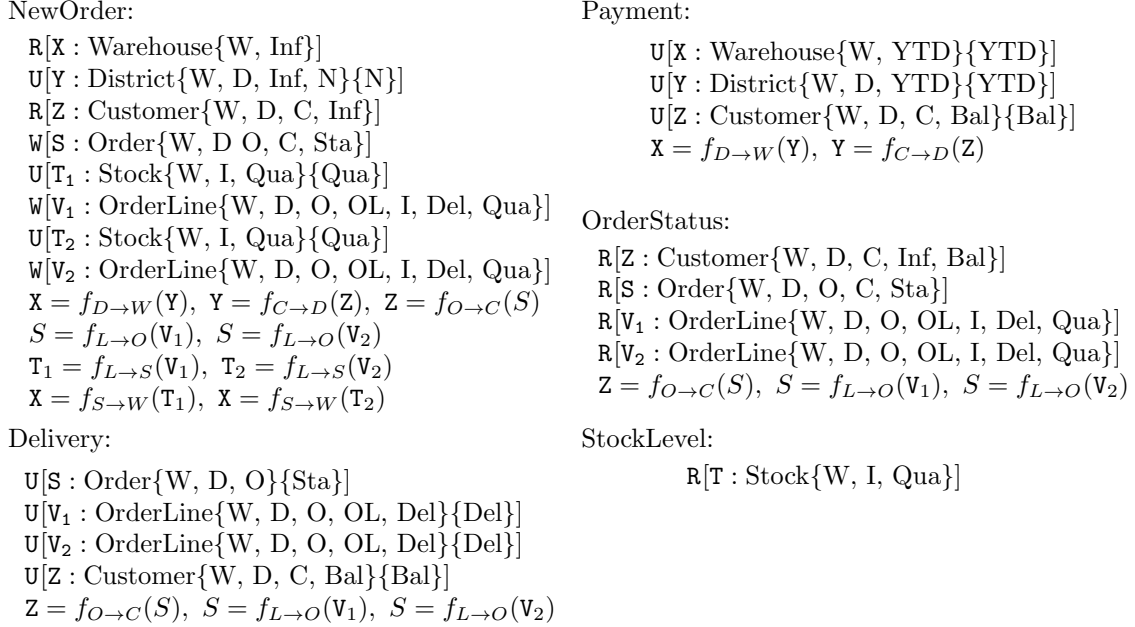


FIGURE 2. Abstraction for the TPC-C transaction templates. Attribute names are abbreviated.

$\mathcal{P} = \{\text{Delivery}, \text{OrderStatus}\}$  of the TPC-C benchmark, given in Figure 2, where functional constraints are added to express the fact that a tuple of type *OrderLine* implies the tuple of type *Order* (function  $f_{L \rightarrow O}$ ), which in turn implies the tuple of type *Customer* (function  $f_{O \rightarrow C}$ ). This set  $\mathcal{P}$  is not robust against RC, but robustness can be achieved by promoting the R-operation over *Customer* in *OrderStatus* to a U-operation. However, without functional constraints, this single promoted operation no longer guarantees robustness, as witnessed by the following schedule:

$$\begin{array}{l}
T_1(\text{Orderstatus}) : U_1[c] R_1[a] \qquad \qquad \qquad R_1[b_1] R_1[b_2] C_1 \\
T_2(\text{Delivery}) : \qquad \qquad \qquad U_2[a] U_2[b_1] U_2[b_2] U_2[c'] C_2
\end{array}$$

Notice in particular how this schedule implicitly assumes in  $T_2$  that *Order*  $a$  belongs to *Customer*  $c'$  instead of *Customer*  $c$  to avoid a dirty write on  $c$ . Without functional constraints,  $\mathcal{P}$  is only robust against RC if *all* R-operations in *OrderStatus* are promoted to U-operations.

### 3. DEFINITIONS

We recall the necessary definitions from [VKKN21] and extend them with functional constraints.

**3.1. Databases.** A *relational schema* is a pair (Rels, Funcs) where Rels is a set of relation names and Funcs is a set of function names. A finite set of attribute names  $\text{Attr}(R)$  is associated to every relation  $R \in \text{Rels}$ . Relations will be instantiated by abstract objects that serve as an abstraction of relational tuples. To this end, for every relation  $R \in \text{Rels}$ , we fix an infinite set of tuples  $\mathbf{Tuples}_R$ . Furthermore, we assume that  $\mathbf{Tuples}_R \cap \mathbf{Tuples}_S = \emptyset$  for all  $R, S \in \text{Rels}$  with  $R \neq S$ . We then denote by  $\mathbf{Tuples}$  the set  $\bigcup_{R \in \text{Rels}} \mathbf{Tuples}_R$  of all

possible tuples. Notice that, by definition, for every  $\mathbf{t} \in \mathbf{Tuples}$  there is a unique relation  $R \in \mathbf{Rels}$  such that  $\mathbf{t} \in \mathbf{Tuples}_R$ . In that case, we say that  $\mathbf{t}$  is of *type*  $R$  and denote the latter by  $\text{type}(\mathbf{t}) = R$ . Each function name  $f \in \mathbf{Funcs}$  has a domain  $\text{dom}(f) \in \mathbf{Rels}$  and a range  $\text{range}(f) \in \mathbf{Rels}$ . Functions are used to encode relationships between tuples like for instance those implied by foreign-keys constraints. For instance, in the SmallBank example  $\mathbf{Funcs} = \{f_{A \rightarrow S}, f_{A \rightarrow C}\}$ ,  $\text{dom}(f_{A \rightarrow S}) = \text{dom}(f_{A \rightarrow C}) = A$ ,  $\text{range}(f_{A \rightarrow S}) = S$ , and  $\text{range}(f_{A \rightarrow C}) = C$ . A *database*  $\mathbf{D}$  over schema  $(\mathbf{Rels}, \mathbf{Funcs})$  assigns to every relation name  $R \in \mathbf{Rels}$  a finite set  $R^{\mathbf{D}} \subset \mathbf{Tuples}_R$  and to every function name  $f \in \mathbf{Funcs}$  a function  $f^{\mathbf{D}}$  from  $\text{dom}(f)^{\mathbf{D}}$  to  $\text{range}(f)^{\mathbf{D}}$ .

**3.2. Transactions and Schedules.** For a tuple  $\mathbf{t} \in \mathbf{Tuples}$ , we distinguish three operations  $R[\mathbf{t}]$ ,  $W[\mathbf{t}]$ , and  $U[\mathbf{t}]$  on  $\mathbf{t}$ , denoting that tuple  $\mathbf{t}$  is read, written, or updated, respectively. We say that the operation is on the tuple  $\mathbf{t}$ . The operation  $U[\mathbf{t}]$  is an atomic update and should be viewed as an atomic sequence of a read of  $\mathbf{t}$  followed by a write to  $\mathbf{t}$ . We will use the following terminology: a *read operation* is an  $R[\mathbf{t}]$  or a  $U[\mathbf{t}]$ , and a *write operation* is a  $W[\mathbf{t}]$  or a  $U[\mathbf{t}]$ . Furthermore, an R-operation is an  $R[\mathbf{t}]$ , a W-operation is a  $W[\mathbf{t}]$ , and a U-operation is a  $U[\mathbf{t}]$ . We also assume a special *commit* operation denoted  $C$ . To every operation  $o$  on a tuple of type  $R$ , we associate the set of attributes  $\text{ReadSet}(o) \subseteq \text{Attr}(R)$  and  $\text{WriteSet}(o) \subseteq \text{Attr}(R)$  containing, respectively, the set of attributes that  $o$  reads from and writes to. When  $o$  is a R-operation then  $\text{WriteSet}(o) = \emptyset$ . Similarly, when  $o$  is a W-operation then  $\text{ReadSet}(o) = \emptyset$ .

A *transaction*  $T$  is a sequence of read and write operations followed by a commit. We assume that a transactions starts when its first operation is executed, but no earlier. Formally, we model a transaction as a linear order  $(T, \leq_T)$ , where  $T$  is the set of (read, write and commit) operations occurring in the transaction and  $\leq_T$  encodes the ordering of the operations. As usual, we use  $<_T$  to denote the strict ordering.

When considering a set  $\mathcal{T}$  of transactions, we assume that every transaction in the set has a unique id  $i$  and write  $T_i$  to make this id explicit. Similarly, to distinguish the operations of different transactions, we add this id as a subscript to the operation. That is, we write  $W_i[\mathbf{t}]$ ,  $R_i[\mathbf{t}]$ , and  $U_i[\mathbf{t}]$  to denote a  $W[\mathbf{t}]$ ,  $R[\mathbf{t}]$ , and  $U[\mathbf{t}]$  occurring in transaction  $T_i$ ; similarly  $C_i$  denotes the commit operation in transaction  $T_i$ . This convention is consistent with the literature (see, *e.g.* [BBG<sup>+</sup>95, Fek05]). To avoid ambiguity of notation, we assume that a transaction performs at most one write, one read, and one update per tuple. The latter is a common assumption (see, *e.g.* [Fek05]). All our results carry over to the more general setting in which multiple writes and reads per tuple are allowed.

A (*multiversion*) *schedule*  $s$  over a set  $\mathcal{T}$  of transactions is a tuple  $(O_s, \leq_s, \ll_s, v_s)$  where  $O_s$  is the set containing all operations of transactions in  $\mathcal{T}$  as well as a special operation  $op_0$  conceptually writing the initial versions of all existing tuples,  $\leq_s$  encodes the ordering of these operations,  $\ll_s$  is a *version order* providing for each tuple  $\mathbf{t}$  a total order over all write operations on  $\mathbf{t}$  occurring in  $s$ , and  $v_s$  is a *version function* mapping each read operation  $a$  in  $s$  to either  $op_0$  or to a write<sup>1</sup> operation different from  $a$  in  $s$ . We require that  $op_0 \leq_s a$  for every operation  $a \in O_s$ ,  $op_0 \ll_s a$  for every write operation  $a \in O_s$ , and that  $a <_T b$  implies  $a <_s b$  for every  $T \in \mathcal{T}$  and every  $a, b \in T$ .<sup>2</sup> We furthermore require that for every read operation  $a$ ,  $v_s(a) <_s a$  and, if  $v_s(a) \neq op_0$ , then the operation  $v_s(a)$  is on the same

<sup>1</sup>Recall that a write operation is either a  $W[x]$  or a  $U[x]$ .

<sup>2</sup>Recall that  $<_T$  denotes the order of operations in transaction  $T$ .



tuple as  $a$ . Intuitively,  $op_0$  indicates the start of the schedule, the order of operations in  $s$  is consistent with the order of operations in every transaction  $T \in \mathcal{T}$ , and the version function maps each read operation  $a$  to the operation that wrote the version observed by  $a$ . If  $v_s(a)$  is  $op_0$ , then  $a$  observes the initial version of this tuple. The version order  $\ll_s$  represents the order in which different versions of a tuple are installed in the database. For a pair of write operations on the same tuple, this version order does not necessarily coincide with  $\leq_s$ . For example, under RC the version order is based on the commit order instead.

We say that a schedule  $s$  is a *single version schedule* if  $\ll_s$  coincides with  $\leq_s$  and every read operation always reads the last written version of the tuple. Formally, for each pair of write operations  $a$  and  $b$  on the same tuple,  $a \ll_s b$  iff  $a <_s b$ , and for every read operation  $a$  there is no write operation  $c$  on the same tuple as  $a$  with  $v_s(a) <_s c <_s a$ . A single version schedule over a set of transactions  $\mathcal{T}$  is *single version serial* if its transactions are not interleaved with operations from other transactions. That is, for every  $a, b, c \in O_s$  with  $a <_s b <_s c$  and  $a, c \in T$  implies  $b \in T$  for every  $T \in \mathcal{T}$ .

The absence of aborts in our definition of schedule is consistent with the common assumption [Fek05, BG16] that an underlying recovery mechanism will rollback aborted transactions. We only consider isolation levels that only read committed versions. Therefore there will never be cascading aborts.

**3.3. Conflict Serializability.** Let  $a_j$  and  $b_i$  be two operations on the same tuple from different transactions  $T_j$  and  $T_i$  in a set of transactions  $\mathcal{T}$ . We then say that  $a_j$  is *conflicting* with  $b_i$  if:

- (*ww-conflict*)  $\text{WriteSet}(a_j) \cap \text{WriteSet}(b_i) \neq \emptyset$ ; or,
- (*wr-conflict*)  $\text{WriteSet}(a_j) \cap \text{ReadSet}(b_i) \neq \emptyset$ ; or,
- (*rw-conflict*)  $\text{ReadSet}(a_j) \cap \text{WriteSet}(b_i) \neq \emptyset$ .

In this case, we also say that  $a_j$  and  $b_i$  are conflicting operations. Furthermore, commit operations and the special operation  $op_0$  never conflict with any other operation. When  $a_j$  and  $b_i$  are conflicting operations in  $\mathcal{T}$ , we say that  $a_j$  *depends on*  $b_i$  in a schedule  $s$  over  $\mathcal{T}$ , denoted  $b_i \rightarrow_s a_j$  if:<sup>3</sup>

- (*ww-dependency*)  $b_i$  is ww-conflicting with  $a_j$  and  $b_i \ll_s a_j$ ; or,
- (*wr-dependency*)  $b_i$  is wr-conflicting with  $a_j$  and  $b_i = v_s(a_j)$  or  $b_i \ll_s v_s(a_j)$ ; or,
- (*rw-antidependency*)  $b_i$  is rw-conflicting with  $a_j$  and  $v_s(b_i) \ll_s a_j$ .

Intuitively, a ww-dependency from  $b_i$  to  $a_j$  implies that  $a_j$  writes a version of a tuple that is installed after the version written by  $b_i$ . A wr-dependency from  $b_i$  to  $a_j$  implies that  $b_i$  either writes the version observed by  $a_j$ , or it writes a version that is installed before the version observed by  $a_j$ . A rw-antidependency from  $b_i$  to  $a_j$  implies that  $b_i$  observes a version installed before the version written by  $a_j$ .

Two schedules  $s$  and  $s'$  are *conflict equivalent* if they are over the same set  $\mathcal{T}$  of transactions and for every pair of conflicting operations  $a_j$  and  $b_i$ ,  $b_i \rightarrow_s a_j$  iff  $b_i \rightarrow_{s'} a_j$ .

**Definition 3.1.** A schedule  $s$  is *conflict serializable* if it is conflict equivalent to a single version serial schedule.

<sup>3</sup>Throughout the paper, we adopt the following convention: a  $b$  operation can be understood as a ‘before’ while an  $a$  can be interpreted as an ‘after’.

A *conflict graph*  $CG(s)$  for schedule  $s$  over a set of transactions  $\mathcal{T}$  is the graph whose nodes are the transactions in  $\mathcal{T}$  and where there is an edge from  $T_i$  to  $T_j$  if  $T_i$  has an operation  $b_i$  that conflicts with an operation  $a_j$  in  $T_j$  and  $b_i \rightarrow_s a_j$ .

**Theorem 3.2** [Pap86]. *A schedule  $s$  is conflict serializable iff the conflict graph for  $s$  is acyclic.*

**3.4. Multiversion Read Committed.** Let  $s$  be a schedule for a set  $\mathcal{T}$  of transactions. Then,  $s$  exhibits a *dirty write* iff there are two ww-conflicting operations  $a_j$  and  $b_i$  in  $s$  on the same tuple  $\mathbf{t}$  with  $a_j \in T_j$ ,  $b_i \in T_i$  and  $T_j \neq T_i$  such that  $b_i <_s a_j <_s C_i$ . That is, transaction  $T_j$  writes to an attribute of a tuple that has been modified earlier by  $T_i$ , but  $T_i$  has not yet issued a commit.

For a schedule  $s$ , the version order  $\ll_s$  corresponds to the commit order in  $s$  if for every pair of write operations  $a_j \in T_j$  and  $b_i \in T_i$ ,  $b_i \ll_s a_j$  iff  $C_i <_s a_j$ . We say that a schedule  $s$  is *read-last-committed (RLC)* if  $\ll_s$  corresponds to the commit order and for every read operation  $a_j$  in  $s$  on some tuple  $\mathbf{t}$  the following holds:

- $v_s(a_j) = op_0$  or  $C_i <_s a_j$  with  $v_s(a_j) \in T_i$ ; and
- there is no write<sup>4</sup> operation  $c_k \in T_k$  on  $\mathbf{t}$  with  $C_k <_s a_j$  and  $v_s(a_j) \ll_s c_k$ .

So,  $a_j$  observes the most recent version of  $\mathbf{t}$  (according to the order of commits) that is committed before  $a_j$ . Note in particular that a schedule cannot exhibit dirty reads, defined in the traditional way [BBG<sup>+</sup>95], if it is read-last-committed.

**Definition 3.3.** A schedule is *allowed under isolation level read committed (RC)* if it is read-last-committed and does not exhibit dirty writes.

Since a read operation in a schedule allowed under RC can access the most recently committed version immediately instead of waiting for an uncommitted version to be committed, our definition of read committed allows more schedules than the more restrictive lock-based implementation of read committed [BBG<sup>+</sup>95]. Furthermore, our definition of RC should be contrasted with more abstract specifications of Read Committed [ALO00] where read operations are only required to read a committed version, rather than the most recent one. We emphasize that our definition of RC is in line with practical implementations of read committed found in e.g. PostgreSQL.<sup>5</sup>

**3.5. Transaction Templates.** Transaction templates are transactions where operations are defined over typed variables together with functional constraints on these variables. Types of variables are relation names in **Rels** and indicate that variables can only be instantiated by tuples from the respective type. We fix an infinite set of variables **Var** that is disjoint from **Tuples**. Every variable  $X \in \mathbf{Var}$  has an associated relation name in **Rels** as type that we denote by  $\text{type}(X)$ . For an operation  $o_i$  in a template,  $\text{var}(o_i)$  denotes the variable in  $o_i$ . An *equality constraint* is an expression of the form  $X = f(Y)$  where  $X, Y \in \mathbf{Var}$ ,  $\text{dom}(f) = \text{type}(Y)$  and  $\text{range}(f) = \text{type}(X)$ . A *disequality constraint* is an expression of the form  $X \neq Y$  where  $\text{type}(X) = \text{type}(Y)$ .

<sup>4</sup>Recall that a write operation is either a W or a U-operation.

<sup>5</sup><https://www.postgresql.org/docs/15/transaction-iso.html>




FIGURE 3. Multiversion split schedule.

**Definition 3.4.** A *transaction template* is a transaction  $\tau$  over  $\mathbf{Var}$  together with a set  $\Gamma(\tau)$  of equality and disequality constraints. In addition, for every operation  $o$  in  $\tau$  over a variable  $\mathbf{X}$ ,  $\text{ReadSet}(o) \subseteq \text{Attr}(\text{type}(\mathbf{X}))$  and  $\text{WriteSet}(o) \subseteq \text{Attr}(\text{type}(\mathbf{X}))$ .

Recall that we denote variables by capital letters  $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$  and tuples by small letters  $\mathbf{t}, \mathbf{v}$ . The transaction templates derived from the SmallBank and TPC-C benchmarks are shown in Figure 1 and Figure 2, respectively. A variable assignment  $\mu$  is a mapping from  $\mathbf{Var}$  to  $\mathbf{Tuples}$  such that  $\mu(\mathbf{X}) \in \mathbf{Tuples}_{\text{type}(\mathbf{X})}$ . Furthermore,  $\mu$  *satisfies* a constraint  $\mathbf{X} = f(\mathbf{Y})$  (resp.,  $\mathbf{X} \neq \mathbf{Y}$ ) over a database  $\mathbf{D}$  when  $\mu(\mathbf{X}) = f^{\mathbf{D}}(\mu(\mathbf{Y}))$  (resp.,  $\mu(\mathbf{X}) \neq \mu(\mathbf{Y})$ ). A variable assignment  $\mu$  for a transaction template  $\tau$  is *admissible* for  $\mathbf{D}$  if it satisfies all constraints in  $\Gamma(\tau)$  over  $\mathbf{D}$ . By  $\mu(\tau)$ , we denote the transaction obtained by replacing each variable  $\mathbf{X}$  in  $\tau$  with  $\mu(\mathbf{X})$ .

A set of transactions  $\mathcal{T}$  is *consistent* with a set of transaction templates  $\mathcal{P}$  and database  $\mathbf{D}$ , if for every transaction  $T$  in  $\mathcal{T}$  there is a transaction template  $\tau \in \mathcal{P}$  and a variable mapping  $\mu_T$  that is admissible for  $\mathbf{D}$  such that  $\mu_T(\tau) = T$ . We refer to Section 2 for concrete examples based on transaction templates derived from the SmallBank and TPC-C benchmarks.

**3.6. Robustness.** We define the robustness property [BG16] (also called *acceptability* in [Fek05, FLO<sup>+</sup>05]), which guarantees serializability for all schedules of a given set of transactions for a given isolation level.

**Definition 3.5** (Transaction Robustness). A set  $\mathcal{T}$  of transactions is *robust* against RC if every schedule for  $\mathcal{T}$  that is allowed under RC is conflict serializable.

In the next definition, we represent conflicting operations from transactions in a set  $\mathcal{T}$  as quadruples  $(T_i, b_i, a_j, T_j)$  with  $b_i$  and  $a_j$  conflicting operations, and  $T_i$  and  $T_j$  their respective transactions in  $\mathcal{T}$ . We call these quadruples *conflicting quadruples* for  $\mathcal{T}$ . Further, for an operation  $b \in T$ , we denote by  $\text{prefix}_b(T)$  the restriction of  $T$  to all operations that are before or equal to  $b$  according to  $\leq_T$ . Similarly, we denote by  $\text{postfix}_b(T)$  the restriction of  $T$  to all operations that are strictly after  $b$  according to  $\leq_T$ . Throughout the paper, we interchangeably consider transactions both as linear orders as well as sequences. Therefore,  $T$  is then equal to the sequence  $\text{prefix}_b(T)$  followed by  $\text{postfix}_b(T)$  which we denote by  $\text{prefix}_b(T) \cdot \text{postfix}_b(T)$  for every  $b \in T$ .

**Definition 3.6** (Multiversion split schedule). Let  $\mathcal{T}$  be a set of transactions and  $C = (T_1, b_1, a_2, T_2), (T_2, b_2, a_3, T_3), \dots, (T_m, b_m, a_1, T_1)$  a sequence of conflicting quadruples for  $\mathcal{T}$  such that each transaction in  $\mathcal{T}$  occurs in at most two different quadruples. A *multiversion split schedule* for  $\mathcal{T}$  based on  $C$  is a multiversion schedule that has the following form:

$$\text{prefix}_{b_1}(T_1) \cdot T_2 \cdot \dots \cdot T_m \cdot \text{postfix}_{b_1}(T_1) \cdot T_{m+1} \cdot \dots \cdot T_n,$$

where

- (1) there is no write operation in  $\text{prefix}_{b_1}(T_1)$  ww-conflicting with a write operation in any of the transactions  $T_2, \dots, T_m$ ;
- (2)  $b_1 <_{T_1} a_1$  or  $b_m$  is rw-conflicting with  $a_1$ ; and,
- (3)  $b_1$  is rw-conflicting with  $a_2$ .

Furthermore,  $T_{m+1}, \dots, T_n$  are the remaining transactions in  $\mathcal{T}$  (those not mentioned in  $C$ ) in an arbitrary order.

Figure 3 depicts a schematic multiversion split schedule. The name stems from the fact that the schedule is obtained by splitting one transaction in two ( $T_1$  at operation  $b_1$  in Figure 3) and placing all other transactions in  $C$  in between. The figure does not display the trailing transactions  $T_{m+1}, T_{m+2}, \dots$  and assumes  $b_1 <_{T_1} a_1$ .

The following theorem characterizes non-robustness in terms of the existence of a multiversion split schedule.

**Theorem 3.7** [VKKN21]. *For a set of transactions  $\mathcal{T}$ , the following are equivalent:*

- (1)  $\mathcal{T}$  is not robust against RC;
- (2) there is a multiversion split schedule  $s$  for  $\mathcal{T}$  based on some  $C$ .

Let  $\mathcal{P}$  be a set of transaction templates and  $\mathbf{D}$  be a database. Then,  $\mathcal{P}$  is *robust against RC over  $\mathbf{D}$*  if for every set of transactions  $\mathcal{T}$  that is consistent with  $\mathcal{P}$  and  $\mathbf{D}$ , it holds that  $\mathcal{T}$  is robust against RC.

**Definition 3.8** (Template Robustness). A set of transaction templates  $\mathcal{P}$  is *robust against RC* if  $\mathcal{P}$  is robust against RC for every database  $\mathbf{D}$ .

We say that a transaction template  $(\tau, \Gamma)$  is a *variable transaction template* when  $\Gamma = \emptyset$  and an *equality transaction template* when all constraints in  $\Gamma$  are equalities. We denote these sets by **VarTemp** and **EqTemp**, respectively. For an isolation level  $\mathcal{I}$  and a class of transaction templates  $\mathcal{C}$ ,  $\text{T-ROBUSTNESS}(\mathcal{C}, \mathcal{I})$  is the problem to decide if a given set of transaction templates  $\mathcal{P} \in \mathcal{C}$  is robust against  $\mathcal{I}$ . When  $\mathcal{C}$  is the class of all transaction templates, we simply write  $\text{T-ROBUSTNESS}(\mathcal{I})$ .

**Theorem 3.9** [VKKN21].  $\text{T-ROBUSTNESS}(\mathbf{VarTemp}, \text{RC})$  is decidable in PTIME.

In Section 4 we start out with a negative result and argue that the addition of functional constraints in its most general form is undecidable by proving undecidability for  $\text{T-ROBUSTNESS}(\mathbf{EqTemp}, \text{RC})$ . Notice in particular that the undecidability result does not even require disequalities. To obtain decidable fragments, we introduce restrictions on the structure of functional constraints. The *schema graph*  $SG(\text{Rels}, \text{Funcs})$  of a schema  $(\text{Rels}, \text{Funcs})$  is a directed multigraph having the relations in  $\text{Rels}$  as nodes, and in which there are as many edges from a node  $R \in \text{Rels}$  to node  $S \in \text{Rels}$  as there are functions  $f \in \text{Funcs}$  with  $\text{dom}(f) = R$  and  $\text{range}(f) = S$ . We say that a schema  $(\text{Rels}, \text{Funcs})$  is *acyclic* if the multigraph  $SG(\text{Rels}, \text{Funcs})$  is acyclic and that it is a *multi-tree* if there is at most one directed path between any two nodes in  $SG(\text{Rels}, \text{Funcs})$ .

**Example 3.10.** Consider the schema  $(\{P, Q, R, S\}, \{f_{P,R}, f_{Q,R}, f_{R,S}\})$  with  $\text{dom}(f_{i,j}) = i$  and  $\text{range}(f_{i,j}) = j$  for each function  $f_{i,j}$ . The corresponding schema graph with solid lines is given in Figure 4. This schema is a multi-tree, as there is at most one path between any pair of nodes. Notice that the definition of a multi-tree is more general than a forest, as a node can still have multiple parents (e.g., node  $R$  in our example). Adding the function name  $f_{Q,S}$  with  $\text{dom}(f_{Q,S}) = Q$  and  $\text{range}(f_{Q,S}) = S$  results in the schema graph given in




FIGURE 4. Acyclic schema graph over the schema  $(\{P, Q, R, S\}, \{f_{P,R}, f_{Q,R}, f_{R,S}, f_{Q,S}\})$ . If we remove function name  $f_{Q,S}$  (dashed edge), the resulting schema graph is a multi-tree.

Figure 4 that is still acyclic, but no longer a multi-tree as there are now two paths from  $Q$  to  $S$ . □

The schema graph constructed in the proof of Theorem 4.1 contains several cycles (cf., Figure 6). We consider in Section 5 robustness for a fragment where a restricted form of cycles in the schema graph is allowed but where additional constraints on the templates are assumed. We consider robustness for acyclic schema graphs in Section 6.

#### 4. ROBUSTNESS FOR TEMPLATES

We start out with a negative result and show that the robustness problem in its most general form is undecidable (even when disequalities are not allowed). The proof is a reduction from *Post’s Correspondence Problem (PCP)* [Pos46] and relies on cyclic dependencies between functional constraints. The proof can be found in the remainder of this section and is quite elaborate but the basic intuition is simple: the counterexample split schedule will build up the two strings that need to be generated by the PCP instance by repeated application of functional constraints.

**Theorem 4.1.**  $T\text{-ROBUSTNESS}(\mathbf{EqTemp}, RC)$  is undecidable.

It might be tempting to relate the above result to the undecidability of the implication problem for functional and inclusion dependencies [CV85]. Functional constraints indeed allow to define inclusion dependencies (as in the SmallBank example) but they always relate complete tuples and are not suited to define functional dependencies. Furthermore, the proof of Theorem 4.1 makes use of only unary relations, for which the implication problem for functional dependencies and inclusion dependencies is known to be decidable.

The remainder of this section is devoted to proving the correctness of Theorem 4.1. We first present the reduction from the PCP problem in Section 4.1. Afterwards, we show that this reduction is indeed correct by proving both directions in respectively Section 4.2 and Section 4.4.

**4.1. Reduction.** The proof is based on a reduction from the *Post’s Correspondence Problem (PCP)*, which is known to be undecidable [Pos46]. A domino is a pair  $(\mathbf{a}, \mathbf{b})$  of two non-empty strings over  $\Sigma$ . Henceforth we call  $\mathbf{a}$  its *top value* and  $\mathbf{b}$  its *bottom value*. Given a set of dominoes  $\mathcal{D}$ , the PCP asks if a non-empty sequence  $d_1, d_2, \dots, d_r$  of dominoes in  $\mathcal{D}$  exists such that, with  $d_i = (\mathbf{a}_i, \mathbf{b}_i)$ , the strings  $\mathbf{a}_1\mathbf{a}_2 \dots \mathbf{a}_r$  and  $\mathbf{b}_1\mathbf{b}_2 \dots \mathbf{b}_r$  are identical.

Split(I):	First(S <sub>1</sub> ):	Last(B):
W[X <sub>1</sub> : Boolean]	W[X <sub>2</sub> : Boolean]	W[B : DominoSequence]
R[I : InitialConflict]	W[I : InitialConflict]	R[S <sub>t</sub> : String]
R[S <sub>1</sub> : String]	R[S <sub>0</sub> : String]	R[S <sub>b</sub> : String]
R[S <sub>e</sub> : String]	R[S <sub>e</sub> : String]	R[S <sub>1</sub> : String]
W[C : PCPSolution]	R[S <sub>1</sub> : String]	W[C : PCPSolution]
X <sub>1</sub> = f <sub>is-non-empty</sub> (S <sub>1</sub> )	W[B : DominoSequence]	S <sub>t</sub> = f <sub>top-string</sub> (B)
X <sub>1</sub> = f <sub>is-error</sub> (S <sub>e</sub> )	X <sub>2</sub> = f <sub>is-non-empty</sub> (S <sub>0</sub> )	S <sub>b</sub> = f <sub>bottom-string</sub> (B)
C = f <sub>final-domino-sequence</sub> (I)	X <sub>2</sub> = f <sub>is-error</sub> (S <sub>0</sub> )	S <sub>1</sub> = f <sub>future-solution-string</sub> (B)
S <sub>1</sub> = f <sub>final-dominoes-string</sub> (I)	S <sub>1</sub> = f <sub>final-dominoes-string</sub> (I)	C = f <sub>DS→PCP</sub> (B)
S <sub>e</sub> = f <sub>error-string</sub> (I)	S <sub>0</sub> = f <sub>top-string</sub> (B)	S <sub>t</sub> = f <sub>solution-string</sub> (C)
S <sub>1</sub> = f <sub>solution-string</sub> (C)	S <sub>0</sub> = f <sub>bottom-string</sub> (B)	S <sub>b</sub> = f <sub>solution-string</sub> (C)
I = f <sub>defines</sub> (X <sub>1</sub> )	S <sub>0</sub> = f <sub>empty-string</sub> (B)	S <sub>1</sub> = f <sub>solution-string</sub> (C)
	S <sub>1</sub> = f <sub>future-solution-string</sub> (B)	B = f <sub>PCP→DS</sub> (C)
	S <sub>e</sub> = f <sub>detach</sub> (S <sub>0</sub> )	
	S <sub>e</sub> = f <sub>detach</sub> (S <sub>e</sub> )	
	I = f <sub>defines</sub> (X <sub>2</sub> )	
	B = f <sub>empty-domino-sequence</sub> (S <sub>1</sub> )	

For every domino  $d_i = (a_1 a_2 \dots a_h, b_1 b_2 \dots b_k) \in \mathcal{D}$  a transaction template  $\text{Domino}_i(\mathbf{B})$ :

W[B : DominoSequence]	S <sub>t</sub> = f <sub>top-string</sub> (B)	S <sub>b</sub> = f <sub>bottom-string</sub> (B)
R[S <sub>0</sub> : String]	S <sub>ta<sub>1</sub></sub> = f <sub>append-a<sub>1</sub></sub> (S <sub>t</sub> )	S <sub>bb<sub>1</sub></sub> = f <sub>append-b<sub>1</sub></sub> (S <sub>b</sub> )
R[S <sub>1</sub> : String]	S <sub>ta<sub>1</sub>a<sub>2</sub></sub> = f <sub>append-a<sub>2</sub></sub> (S <sub>ta<sub>1</sub></sub> )	S <sub>bb<sub>1</sub>b<sub>2</sub></sub> = f <sub>append-b<sub>2</sub></sub> (S <sub>bb<sub>1</sub></sub> )
R[S <sub>t</sub> : String]	...	...
R[S <sub>ta<sub>1</sub></sub> : String]	S <sub>ta<sub>1</sub>a<sub>2</sub>...a<sub>h</sub></sub> =	S <sub>bb<sub>1</sub>b<sub>2</sub>...b<sub>k</sub></sub> =
R[S <sub>ta<sub>1</sub>a<sub>2</sub></sub> : String]	f <sub>append-a<sub>h</sub></sub> (S <sub>ta<sub>1</sub>...a<sub>h-1</sub></sub> )	f <sub>append-b<sub>k</sub></sub> (S <sub>bb<sub>1</sub>...b<sub>k-1</sub></sub> )
...	S <sub>t</sub> = f <sub>detach</sub> (S <sub>ta<sub>1</sub></sub> )	S <sub>b</sub> = f <sub>detach</sub> (S <sub>tb<sub>1</sub></sub> )
R[S <sub>ta<sub>1</sub>a<sub>2</sub>...a<sub>h</sub></sub> : String]	S <sub>ta<sub>1</sub></sub> = f <sub>detach</sub> (S <sub>ta<sub>1</sub>a<sub>2</sub></sub> )	S <sub>bb<sub>1</sub></sub> = f <sub>detach</sub> (S <sub>bb<sub>1</sub>b<sub>2</sub></sub> )
R[S <sub>b</sub> : String]	...	...
R[S <sub>bb<sub>1</sub></sub> : String]	S <sub>ta<sub>1</sub>a<sub>2</sub>...a<sub>h-1</sub></sub> =	S <sub>bb<sub>1</sub>b<sub>2</sub>...b<sub>k-1</sub></sub> =
R[S <sub>bb<sub>1</sub>b<sub>2</sub></sub> : String]	f <sub>detach</sub> (S <sub>ta<sub>1</sub>a<sub>2</sub>...a<sub>h</sub></sub> )	f <sub>detach</sub> (S <sub>bb<sub>1</sub>b<sub>2</sub>...b<sub>k</sub></sub> )
...	S <sub>ta<sub>1</sub>a<sub>2</sub>...a<sub>h</sub></sub> = f <sub>top-string</sub> (B <sub>next</sub> )	S <sub>bb<sub>1</sub>b<sub>2</sub>...b<sub>k</sub></sub> =
R[S <sub>bb<sub>1</sub>b<sub>2</sub>...b<sub>k</sub></sub> : String]	S <sub>a<sub>1</sub></sub> = f <sub>top</sub> (S <sub>ta<sub>1</sub></sub> )	f <sub>bottom-string</sub> (B <sub>next</sub> )
W[B <sub>next</sub> : DominoSequence]	S <sub>a<sub>2</sub></sub> = f <sub>top</sub> (S <sub>ta<sub>1</sub>a<sub>2</sub></sub> )	S <sub>b<sub>1</sub></sub> = f <sub>top</sub> (S <sub>bb<sub>1</sub></sub> )
	...	S <sub>b<sub>2</sub></sub> = f <sub>top</sub> (S <sub>bb<sub>1</sub>b<sub>2</sub></sub> )
	S <sub>a<sub>h</sub></sub> = f <sub>top</sub> (S <sub>ta<sub>1</sub>a<sub>2</sub>...a<sub>h</sub></sub> )	...
	S <sub>1</sub> = f <sub>future-solution-string</sub> (B)	S <sub>b<sub>k</sub></sub> = f <sub>top</sub> (S <sub>bb<sub>1</sub>b<sub>2</sub>...b<sub>k</sub></sub> )
	S <sub>0</sub> = f <sub>empty-string</sub> (B)	S <sub>1</sub> = f <sub>future-solution-string</sub> (B <sub>next</sub> )
		S <sub>0</sub> = f <sub>empty-string</sub> (B <sub>next</sub> )
		B <sub>next</sub> = f <sub>next-sequence</sub> (B)
		B = f <sub>previous-sequence</sub> (B <sub>next</sub> )

FIGURE 5. Transaction templates for the proof of Theorem 4.1.

For the reduction to non-robustness against RC, we construct a set  $\mathcal{P}$  of transaction templates consisting of the transaction templates in Figure 5 for  $\mathcal{D}$ . There are the transactions Split, First and Last (whose meaning will be explained next) and for every domino in  $\mathcal{D}$  there is a template in Figure 5 representing that domino and the action of appending that domino to a sequence of dominoes. The schema consists of the relations  $\{\text{Boolean}, \text{InitialConflict}, \text{String}, \text{PCPSolution}, \text{DominoSequence}\}$  whose meaning will be explained below together with a discussion of all the functions. The schema graph is presented in Figure 6 and contains various cycles.

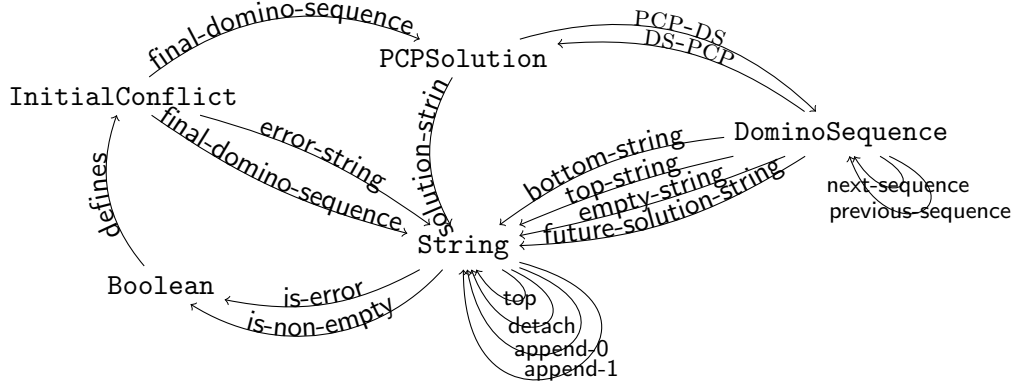


FIGURE 6. Schema graph for the transaction templates in Figure 5 (for any set of dominoes).

To prove Theorem 4.1, we will show that there is a solution for PCP if and only if  $\mathcal{P}$  is not robust against RC. For the only-if direction, we show that, if there is a solution  $\mathbf{d} = d_1, d_2, \dots, d_r$  for the PCP problem over  $\mathcal{D}$ , then there is a multiversion split schedule that encodes this solution in a particular way: in this schedule the split transaction is an instantiation of transaction template Split, the next transaction is an instantiation of First, then followed by instantiations of transaction templates  $\text{Domino}_{d_1}, \dots, \text{Domino}_{d_r}$ , representing the sequence of dominoes in solution  $\mathbf{d}$ , and finally an instantiation of transaction template Last. Henceforth, we call a schedule that encodes a sequence of dominoes  $\mathbf{d}$  in this way a *schedule-encoding of  $\mathbf{d}$* . For the if-direction, we first show that every multiversion split schedule consistent with the transaction templates in Figure 5 for some set  $\mathcal{D}$  of dominoes is a schedule-encoding for some sequence  $\mathbf{d}$  of dominoes from  $\mathcal{D}$ , and then that for every schedule-encoding of a sequence  $\mathbf{d}$  of dominoes,  $\mathbf{d}$  is always a solution for the PCP problem over a set of dominoes containing those in  $\mathbf{d}$ .

4.2. **Only-if direction.** We first prove the only-if direction of Theorem 4.1.

**Proposition 4.2** (Only-if part of Theorem 4.1). *Let  $\mathcal{D}$  be a set of dominoes with a solution  $\mathbf{d}$  for the PCP problem for  $\mathcal{D}$ . Then there exists a schedule-encoding of  $\mathbf{d}$  that is consistent with the transaction templates in Figure 5 and some database  $\mathbf{D}$ .*

*Proof.* Let  $\mathbf{d} = d_1, d_2, \dots, d_r$  be a solution to the PCP problem for  $\mathcal{D}$ . Let  $\mathbf{a}_1\mathbf{a}_2 \dots \mathbf{a}_r$  be the read of top values and  $\mathbf{b}_1\mathbf{b}_2 \dots \mathbf{b}_r$  be the read of bottom values, which thus represent an identical string  $\mathbf{c} = c_1 \dots c_n$ , with  $c_i \in \Sigma$ . We now construct a schedule  $s$  and database  $\mathbf{D}$  as in Definition 3.6 with transactions based on the transaction templates  $\mathcal{P}$  in Figure 5.

Relation `PCPSolution` contains a tuple that we interpret as the PCP solution  $\mathbf{d} = d_1, d_2, \dots, d_r$ . Relation `DominoSequence` contains  $r + 1$  tuples, one for every prefix of  $\mathbf{d}$ , including the empty sequence  $()$  and the PCP solution  $\mathbf{d}$  itself. For convenience of notation, we will henceforth often represent tuples by their interpretation, which is justified by the fact that every tuple in a particular relation will have a different interpretation, and the relation itself can always be derived from the context (e.g., the function signature).

Since the PCP solution has an interpretation in both the relations `PCPSolution` and `DominoSequence`, we assume two functions,  $f_{\text{PCP} \rightarrow \text{DS}} : \text{PCPSolution} \rightarrow \text{DominoSequence}$  and  $f_{\text{DS} \rightarrow \text{PCP}} : \text{DominoSequence} \rightarrow \text{PCPSolution}$  that relate these interpretations to each other. That is,  $f_{\text{PCP} \rightarrow \text{DS}}^{\mathbf{D}}(\mathbf{d}) = \mathbf{d}$  and  $f_{\text{typecase-to-C}}^{\mathbf{D}}(\mathbf{d}) = \mathbf{d}$ .

Further, we have functions  $f_{\text{next-sequence}} : \text{DominoSequence} \rightarrow \text{DominoSequence}$  and  $f_{\text{previous-sequence}} : \text{DominoSequence} \rightarrow \text{DominoSequence}$  with the following interpretation:

$$\begin{aligned} f_{\text{next-sequence}}^{\mathbf{D}}(\mathbf{d}') &= \mathbf{d}'d && \text{with } d' \text{ a strict prefix of } \mathbf{d} \text{ followed by domino } d \text{ in } \mathbf{d}, \\ f_{\text{next-sequence}}^{\mathbf{D}}(\mathbf{d}) &= \mathbf{d}, \\ f_{\text{previous-sequence}}^{\mathbf{D}}(\mathbf{d}'d) &= \mathbf{d}' && \text{with } d' \text{ a strict prefix of } \mathbf{d} \text{ followed by domino } d \text{ in } \mathbf{d}, \\ f_{\text{previous-sequence}}^{\mathbf{D}}(()) &= (). \end{aligned}$$

Intuitively, these functions relate each tuple in `DominoSequence` representing a prefix of  $\mathbf{d}$  to the prefixes obtained by adding or removing one domino in the sequence. That is, given a tuple in `DominoSequence` representing a strict prefix  $d_1, d_2, \dots, d_{i-1}, d_i$  of  $\mathbf{d}$ ,  $f_{\text{next-sequence}}$  returns the tuple representing  $d_1, d_2, \dots, d_{i-1}, d_i, d_{i+1}$  (i.e., the prefix of  $\mathbf{d}$  obtained by adding one domino), and  $f_{\text{previous-sequence}}$  returns the tuple representing  $d_1, d_2, \dots, d_{i-1}$  (i.e., the prefix of  $\mathbf{d}$  obtained by removing the last domino). We furthermore distinguish two special cases to guarantee that both functions are defined for all tuples in `DominoSequence`: if the tuple represents the solution  $\mathbf{d}$  itself, then  $f_{\text{next-sequence}}$  returns  $\mathbf{d}$ , and if the tuple represents the empty sequence  $()$ , then  $f_{\text{previous-sequence}}$  returns  $()$ .

Relation `StringD` contains a tuple representing the read  $\mathbf{c}$  of PCP-solution sequence  $\mathbf{d}$ , a tuple representing an error  $\langle \text{error} \rangle$ , and a tuple for every substring of  $\mathbf{c}$ , including the empty string  $\langle \rangle$ . We assume that all these tuples are different. We use notation  $\langle \rangle$  to denote the empty string to distinguish it from  $()$ , which denotes the empty sequence of dominoes.

Functions  $f_{\text{append-0}} : \text{String} \rightarrow \text{String}$ ,  $f_{\text{append-1}} : \text{String} \rightarrow \text{String}$ ,  $f_{\text{detach}} : \text{String} \rightarrow \text{String}$ , and  $f_{\text{top}} : \text{String} \rightarrow \text{String}$  simulate standard string operations for the interpretations



of tuples in relation **String**. Thus, tuples representing a (possibly empty) string  $\mathbf{e}$ :

$$f_{\text{append-}c}^{\mathbf{D}}(\langle \mathbf{e} \rangle) = \begin{cases} \langle \mathbf{e}c \rangle & \text{with } \mathbf{e} \text{ a (possibly empty) string over } \Sigma, c \in \Sigma, \text{ and} \\ & \langle \mathbf{e}c \rangle \text{ a substring of } \mathbf{c}, \\ \langle \text{error} \rangle & \text{otherwise,} \end{cases}$$

$$f_{\text{detach}}^{\mathbf{D}}(\langle \mathbf{e}c \rangle) = \langle \mathbf{e} \rangle \text{ with } \mathbf{e} \text{ a (possibly empty) string over } \Sigma, \text{ and } c \in \Sigma,$$

$$f_{\text{detach}}^{\mathbf{D}}(\langle \rangle) = f_{\text{detach}}^{\mathbf{D}}(\langle \text{error} \rangle) = \langle \text{error} \rangle,$$

$$f_{\text{top}}^{\mathbf{D}}(\langle \mathbf{e}c \rangle) = \langle c \rangle \text{ with } \mathbf{e} \text{ a (possibly empty) string over } \Sigma, \text{ and } c \in \Sigma,$$

$$f_{\text{top}}^{\mathbf{D}}(\langle \rangle) = f_{\text{top}}^{\mathbf{D}}(\langle \text{error} \rangle) = \langle \text{error} \rangle.$$

Notice that these function interpretations are closed under  $\mathbf{D}$ , that is, every tuple from relation  $\mathbf{String}^{\mathbf{D}}$  maps onto a tuple that is in relation  $\mathbf{String}^{\mathbf{D}}$ .

Every tuple in **DominoSequence** is associated with three tuples in **String** representing, respectively, the read of top values, the read of bottom values, and the empty string. The association is made via functions  $f_{\text{top-string}} : \mathbf{DominoSequence} \rightarrow \mathbf{String}$ ,  $f_{\text{bottom-string}} : \mathbf{DominoSequence} \rightarrow \mathbf{String}$ , and  $f_{\text{empty-string}} : \mathbf{DominoSequence} \rightarrow \mathbf{String}$  with following interpretations in  $\mathbf{D}$ :

$$f_{\text{top-string}}^{\mathbf{D}}(\mathbf{d}') = \mathbf{e}, \text{ with } \mathbf{e} \text{ the read of top values on dominoes in } \mathbf{d}',$$

$$f_{\text{bottom-string}}^{\mathbf{D}}(\mathbf{d}') = \mathbf{e}, \text{ with } \mathbf{e} \text{ the read of bottom values on dominoes in } \mathbf{d}', \text{ and}$$

$$f_{\text{empty-string}}^{\mathbf{D}}(\mathbf{d}') = \langle \rangle.$$

We emphasize that in the expressions above the read  $\mathbf{e}$  of top and bottom values on dominoes in  $\mathbf{d}'$  might be empty.

Finally, for function  $f_{\text{future-solution-string}} : \mathbf{DominoSequence} \rightarrow \mathbf{String}$  we consider the interpretation that associates every domino sequence  $\mathbf{d}'$  represented by a tuple in relation **DominoSequence** in  $\mathbf{D}$  to the final read  $f_{\text{future-solution-string}}^{\mathbf{D}}(\mathbf{d}') = \mathbf{c}$ . Function  $f_{\text{solution-string}} : \mathbf{PCPSolution} \rightarrow \mathbf{String}$  does the same for the single tuple representing  $\mathbf{d}$  in **PCPSolution**, thus with  $f_{\text{solution-string}}^{\mathbf{D}}(\mathbf{d}) = \mathbf{c}$ . Function  $f_{\text{empty-domino-sequence}} : \mathbf{String} \rightarrow \mathbf{DominoSequence}$  is interpreted to map every tuple in **String** onto the tuple from **DominoSequence** representing the empty sequence  $\langle \rangle$ .

All other relations and functions have as purpose to pass tuples from one transaction to another in a schedule and to enforce that certain tuples do not collide, which is useful for the (if)-part of the proof.

Relation **Boolean** <sup>$\mathbf{D}$</sup>  contains two tuples, which we interpret as Boolean values 0 and 1. Function  $f_{\text{is-non-empty}} : \mathbf{String} \rightarrow \mathbf{Boolean}$  and  $f_{\text{is-error}} : \mathbf{String} \rightarrow \mathbf{Boolean}$  are interpreted as follows:

$$f_{\text{is-non-empty}}^{\mathbf{D}}(s) = \begin{cases} 1 & \text{if } s \neq \langle \rangle, \\ 0 & \text{otherwise,} \end{cases} \text{ , and}$$

$$f_{\text{is-error}}^{\mathbf{D}}(s) = \begin{cases} 1 & \text{if } s = \langle \text{error} \rangle, \\ 0 & \text{otherwise,} \end{cases} \text{ , and.}$$

Finally, relation **InitialConflict** <sup>$\mathbf{D}$</sup>  contains a single tuple, which we refer to by  $\langle \text{init} \rangle$ . The interpretation of  $f_{\text{defines}} : \mathbf{Boolean} \rightarrow \mathbf{InitialConflict}$  maps 1 and 0 onto  $\langle \text{init} \rangle$ . Function  $f_{\text{error-string}} : \mathbf{InitialConflict} \rightarrow \mathbf{String}$  maps  $\langle \text{init} \rangle$  onto  $\langle \text{error} \rangle$ . Functions  $f_{\text{final-domino-string}} : \mathbf{InitialConflict} \rightarrow \mathbf{DominoSequence}$  and  $f_{\text{final-domino-sequence}} : \mathbf{InitialConflict} \rightarrow \mathbf{PCPSolution}$  map  $\langle \text{init} \rangle$  onto the solution domino sequence  $\mathbf{d}$ , respectively on the final read  $\mathbf{c}$  of  $\mathbf{d}$ .

Now the schedule  $\text{prefix}_{b_1}(T_1) \cdot T_2 \cdot \dots \cdot T_m \cdot \text{postfix}_{b_1}(T_1)$ , taking  $T_1 = \text{Split}(\langle \text{init} \rangle)$ ,  $T_2 = \text{First}(\langle \text{init} \rangle)$ , for  $i : 1 \leq i \leq r$ , transaction  $T_{i+2} = \text{Domino}_i((d_1, \dots, d_i))$ ,  $T_m = \text{Last}((d_1, \dots, d_r))$  and  $b_1 = \langle \text{init} \rangle$  has the conditions of Definition 3.6. Indeed, it is based on sequence of conflict quadruples  $(T_1, R_1[\langle \text{init} \rangle], W_2[\langle \text{init} \rangle], T_2)$ ,  $(T_2, W_2[()], W_3[()], T_3)$ ,  $(T_3, W_3[\mathbf{d}_1], W_4[\mathbf{d}_2], T_4), \dots, (T_{r+2}, W_{r+2}[\mathbf{d}_1, \dots, \mathbf{d}_r], W_{r+3}[\mathbf{d}_1, \dots, \mathbf{d}_r], T_{r+3})$ ,  $(T_{r+3}, W_{r+3}[\mathbf{d}], W_1[\mathbf{d}], T_1)$ .

Condition (1) is true because there is no ww-conflict between a write operation in  $\text{prefix}_{b_1}(T_1)$  and a write operation in any of the transactions  $T_2, \dots, T_m$ , since the first write operation, respectively second write operation, in  $\text{Split}(\langle \text{init} \rangle)$  has a type that only occurs before the conflict with  $\text{First}(\langle \text{init} \rangle)$ , and is the conflict with  $\text{Last}((d_1, \dots, d_r))$ , respectively. Furthermore (2) is true because  $b_1 <_{T_1} a_1$  and Condition (3) is true because  $b_1$  and  $a_2$  are rw-conflicting.  $\square$

**4.3. Helpful lemma.** Before proving the opposite direction of Theorem 4.1, we first establish the following Lemma.

**Lemma 4.3.** *If a set  $\mathcal{P}$  of transaction templates is not robust against RC then there is a multiversion split schedule  $\text{prefix}_{b_1}(T_1) \cdot T_2 \cdot \dots \cdot T_m \cdot \text{postfix}_{b_1}(T_1)$  for a set  $\mathcal{T} = \{T_1, \dots, T_m\}$  of transactions consistent with  $\mathcal{P}$  in which an operation from a transaction  $T_j$  depends on an operation from transaction  $T_i$  only if  $j = i + 1$  or  $i = m$  and  $j = 1$ .*

*Proof.* If  $\mathcal{P}$  is not robust against RC, then there is a database  $\mathbf{D}$  and a multiversion split schedule  $s = \text{prefix}_{b_1}(T_1) \cdot T_2 \cdot \dots \cdot T_m \cdot \text{postfix}_{b_1}(T_1) \cdot T_{m+1} \cdot \dots \cdot T_n$  based on a sequence of conflict quadruples  $C$  for a set of transactions  $\mathcal{T}$  that is consistent with  $\mathcal{P}$  and  $\mathbf{D}$  having the properties of Definition 3.6.

We can assume that  $n = m$ . Otherwise removing the transactions  $T_{m+1}, \dots, T_n$  from  $\mathcal{T}$ ,  $s$ , and  $C$ . We can also assume that  $s$  is read-last-committed. Otherwise, choosing an appropriate version order  $\ll_s$  and version function  $v_s$ .

Now suppose that there is a transaction  $T_j$  with an operation  $a'_j$  that depends on an operation  $b'_i$  from transaction  $T_i$  and with  $j \neq i + 1$  or  $i = m$  and  $j \neq 1$ . Clearly, by definition of dependency and the structure of a multiversion split schedule,  $i < j$  or  $j = 1$ .

We proceed the proof by a construction showing that under these assumptions there is an alternative schedule  $s'$  that is also a multiversion split schedule, but for a strict subset of transactions in  $\mathcal{T}$  (thus also still consistent with  $\mathcal{P}$  and  $\mathbf{D}$ ). The result of the lemma then follows from the observation that repeated application of this construction must lead to a schedule with the properties of the lemma, without existence of such a dependency.

For the construction, we proceed by case distinction.

IF  $i \neq 1$  AND  $j \neq 1$ , we construct a schedule  $s'$  from  $s$  by removing all operations from transactions  $T_h$  with  $i < h < j$ . Notice that we remove at least one transaction, since  $i < i + 1 < j$ . We can derive a sequence of conflict quadruples  $C'$  from  $C$  by removing all occurrences of these transactions  $T_h$  and adding the conflict quadruple  $(T_i, b'_i, a'_j, T_j)$  instead. By construction,  $s'$  is a multiversion split schedule based on  $C'$  over a set of transactions consistent with  $\mathcal{P}$  and  $\mathbf{D}$ . It remains to show that the newly constructed schedule  $s'$  has the properties of Definition 3.6. The latter is straightforward since  $C$  and  $C'$  agree on their first and last quadruple, due to assumption  $i \neq 1$  and  $j \neq 1$ .

IF  $i = 1$ , it follows that  $i < j$  and thus  $j \neq 1$ . Then, we construct a schedule  $s'$  from  $s$  by removing all operations from transactions  $T_h$  with  $i < h < j$  and updating the prefix and postfix of  $T_1$ , now based on  $b'_i$ . Notice that we again remove at least one transaction, since

$i < i + 1 < j$  and that we can derive a sequence of conflict quadruples  $C'$  from  $C$  in the same way as before, by removing all occurrences of these transactions  $T_h$  and adding the conflict quadruple  $(T_i, b'_i, a'_j, T_j)$  instead. By construction,  $s'$  is a multiversion split schedule based on  $C'$  over a set of transactions consistent with  $\mathcal{P}$  and  $\mathbf{D}$ . It remains to show that the newly constructed schedule  $s'$  has the properties of Definition 3.6.

First, we observe that  $b'_1$  and  $a'_j$  are rw-conflicting, which immediately implies that Condition (3) is true for  $s'$ . The argument is by exclusion. Indeed, if  $b'_1$  and  $a'_j$  would be ww-conflicting, then  $b'_1 \ll_s a'_j$  implying  $b'_1 <_s a'_j$  (due to the assumed read-last committed) and thus  $b'_1 \leq_s b_1$ , which is not allowed by condition (1) on  $s$ . It follows from a similar argument that  $b'_1$  and  $a'_j$  are not wr-conflicting: both  $b'_1 = v_s(a'_j)$  and  $b'_1 \ll_s v_s(a'_j)$  imply  $b'_1 <_s C_1 <_s a'_j$ , which contradicts with  $C_1$  being the last operation in  $s$ .

Since  $b'_1$  is rw-conflicting with  $a'_j$ , we have  $v_s(b'_1) \ll_s a'_j$ , implying  $b'_1 <_s a'_j$  (due to read-last-committed and the structure of a multiversion split schedule), thus  $b'_1 \leq_s b_1$ . Therefore, Condition (1) again transfers from  $s$  to  $s'$ . For similar reasons Condition (2) applies on  $s'$ : If  $b_1 <_{T_1} a_1$  then  $b'_1 \leq_{T_1} b_1 <_{T_1} a_1$ .

OTHERWISE, IF  $j = 1$ , it follows that  $1 < i$ . Then, we construct a schedule  $s'$  from  $s$  by removing all operations from transactions  $T_h$  with  $i < h$ . Notice that we remove at least one transaction, since  $i < m$ . We can derive a sequence of conflicting quadruples  $C'$  from  $C$  by removing all occurrences of these transactions  $T_h$  and adding the conflicting quadruple  $(T_i, b'_i, a'_j, T_j)$  instead.

In this schedule  $s'$ , Condition (1) and (3) transfer from  $s$  by its construction. To see that Condition (2) is true on  $s'$ , simply notice that if  $b'_i$  and  $a'_1$  are ww or wr-conflicting, then either  $b'_i \ll_s a'_j$  or  $b'_i = v_s(a'_j)$  or  $b'_i \ll_s v_s(a'_j)$ , which all imply  $b_i <_s C_i <_s a'_1$  and thus that  $b_1 <_s a'_1$ , implying  $b_1 <_{s'} a'_1$ .  $\square$

**4.4. If direction.** It remains to argue that the if direction of Theorem 4.1 is indeed correct.

Next, we show that, if there exists a multiversion split schedule for the set of transaction templates in Figure 5 for some set  $\mathcal{D}$  of dominoes, then this schedule is always a schedule-encoding of a sequence of dominoes in  $\mathcal{D}$ .

**Proposition 4.4.** *Let  $\mathcal{D}$  be a set of dominoes. If there is a multiversion split schedule  $s$  for a set of transactions consistent with the transaction template in Figure 5 for  $\mathcal{D}$  and some database  $\mathbf{D}$ , then this schedule  $s$  is a schedule-encoding of some sequence  $\mathbf{d}$  of dominoes in  $\mathcal{D}$ .*

For the proof, let  $\mathbf{D}$  be a database and  $s = \text{prefix}_{b_1}(T_1) \cdot T_2 \cdot \dots \cdot T_m \cdot \text{postfix}_{b_1}(T_1)$  a multiversion split schedule for a set of transactions  $\mathcal{T}$  consistent with  $\mathcal{P}$  and  $\mathbf{D}$ , with the conditions of Lemma 4.3 and based on some sequence of conflict quadruples  $C = (T_1, b_1, a_2, T_2), (T_2, b_2, a_3, T_3) \dots, (T_m, b_m, a_1, T_1)$ . We show through a sequence of properties (Lemmas 4.6, 4.7, 4.8, and 4.9), that  $s$  is a schedule-encoding of a sequence  $\mathbf{d}$  of dominoes in  $\mathcal{D}$ .

As a first property (Lemma 4.5), we observe that transaction templates in  $\mathcal{P}$  heavily constrain the possible variable instantiations. For transaction template Split, for example, a variable mapping depends entirely on the choice of the value for variable I. Since Lemma 4.3 forbids the presence of duplicate transactions in  $\mathcal{T}$ , two transactions  $T_i$  and  $T_j$  (with  $i \neq j$ ) based on transaction template Split cannot agree on their choice for variable I in  $s$ . By

applying this argument to other transaction templates, we obtain the following corollary of Lemma 4.3. Here, for each transaction  $T_i$  in  $s$ , we write  $\tau_i$  to denote the transaction template in  $\mathcal{P}$  that it is based on, and by  $\mu_i$  the associated variable mapping for  $\tau_i$ , with  $\mu_i(\tau_i) = T_i$ .

**Lemma 4.5.** *for two transactions  $T_i$  and  $T_j$  in  $s$ , with  $i \neq j$ :*

- *if  $T_i$  and  $T_j$  are based on Split, then  $\mu_i(\mathbf{I}) \neq \mu_j(\mathbf{I})$ ;*
- *if  $T_i$  and  $T_j$  are based on First then  $\mu_i(\mathbf{S}_1) \neq \mu_j(\mathbf{S}_1)$ ;*
- *if  $T_i$  and  $T_j$  are based on Last then  $\mu_i(\mathbf{B}) \neq \mu_j(\mathbf{B})$  and  $\mu_i(\mathbf{C}) \neq \mu_j(\mathbf{C})$ ;*
- *if  $T_i$  and  $T_j$  are based on domino transaction templates then  $\mu_i(\mathbf{B}) \neq \mu_j(\mathbf{B})$  and  $\mu_i(\mathbf{B}_{next}) \neq \mu_j(\mathbf{B}_{next})$ .*

We conclude the proof of Proposition 4.4 with the necessary arguments (Lemmas 4.6, 4.7, 4.8 and 4.9) that  $s$  is indeed a schedule-encoding for some sequence of dominoes.

**Lemma 4.6.** *Transaction  $T_1$  is based on Split,  $T_2$  is based on First, and  $\mu_1(\mathbf{I}) = \mu_2(\mathbf{I})$ ,  $\mu_1(\mathbf{X}_1) \neq \mu_2(\mathbf{X}_2)$ , and  $\mu_1(\mathbf{S}_1) = \mu_2(\mathbf{S}_1) \neq \mu_2(\mathbf{S}_0)$ .*

*Proof.* Since  $b_1$  and  $a_2$  are rw-conflicting (cf, Definition 3.6), and there are no updates in the considered transaction templates, operation  $b_1$  must be a read. Since InitialConflict is the only type allowing for conflicts involving a read, it is immediate that  $T_1$  must be based on Split and  $T_2$  based on First, with  $\mu_1(\mathbf{I}) = \mu_2(\mathbf{I})$ . From this equality and function  $f_{\text{final-dominos-string}}$  it follows that  $\mu_1(\mathbf{S}_1) = \mu_2(\mathbf{S}_1)$ . From Definition 3.6, particularly that there is no ww-conflict between a write operation in  $\text{prefix}_{b_1}(T_1)$  and a write operation in any of the transactions  $T_2, \dots, T_m$ , it follows that  $\mu_1(\mathbf{X}_1) \neq \mu_2(\mathbf{X}_2)$ . Finally, function  $f_{\text{is-non-empty}}$ , which maps  $\mathbf{S}_1$  onto  $\mathbf{X}_1$  in transaction template Split and  $\mathbf{S}_0$  onto  $\mathbf{X}_2$  in transaction template First, implies  $\mu_1(\mathbf{S}_1) \neq \mu_2(\mathbf{S}_0)$ .  $\square$

**Lemma 4.7.** *There is a transaction  $T_3$  in  $s$  and it is based on a domino transaction template, with  $\mu_3(\mathbf{S}_1) = \mu_2(\mathbf{S}_1)$ .*

*Proof.* First, suppose towards a contradiction that  $m = 2$ . We already know from Lemma 4.6 that  $\mu_1(\mathbf{I}) = \mu_2(\mathbf{I})$  and  $\mu_1(\mathbf{X}_1) \neq \mu_2(\mathbf{X}_2)$ , thus  $a_1 = b_1 = \mathbf{R}[\mu_1(\mathbf{I})]$  and  $b_2 = a_2 = \mathbf{W}[\mu_2(\mathbf{I})]$ , indicating  $b_1 \rightarrow_s a_2$ , particularly,  $v_s(b_1) \ll_s a_2$ , thus implying that  $a_1$  cannot depend on  $b_2$ , which is the desired contradiction.

The remainder of the proof is by exclusion. Transaction  $T_3$  is not based on transaction template First, because all possible conflicts between  $T_2$  and an instantiation of transaction template First (implying either  $\mu_2(\mathbf{X}_2) = \mu_3(\mathbf{X}_2)$ ,  $\mu_2(\mathbf{B}) = \mu_3(\mathbf{B})$ , or  $\mu_2(\mathbf{I}) = \mu_3(\mathbf{I})$ ) would imply the equality  $\mu_2(\mathbf{S}_1) = \mu_3(\mathbf{S}_1)$  (through functional constraints  $\mathbf{I} = f_{\text{defines}}(\mathbf{X}_2)$ ,  $\mathbf{S}_1 = f_{\text{future-solution-string}}(\mathbf{B})$ , and  $\mathbf{S}_1 = f_{\text{final-dominos-string}}(\mathbf{I})$ ), which is forbidden by Lemma 4.5. The argument that transaction  $T_3$  cannot be based on transaction template Split is similar: every possible conflict between  $T_2$  and an instantiation of Split implies  $\mu_1(\mathbf{I}) = \mu_2(\mathbf{I}) = \mu_3(\mathbf{I})$  either directly (taking  $\mu_2(\mathbf{I}_2) = \mu_3(\mathbf{I}_1)$  as conflict) or, when taking  $\mu_2(\mathbf{X}_2) = \mu_3(\mathbf{X}_1)$  as conflict, through constraints  $\mathbf{I} = f_{\text{defines}}(\mathbf{X}_1)$  and  $\mathbf{I} = f_{\text{defines}}(\mathbf{X}_2)$  in  $T_2$  and  $T_3$ , respectively. Either way,  $\mu_1(\mathbf{I}) = \mu_3(\mathbf{I})$  is forbidden by Lemma 4.5. Finally, to see that  $T_3$  is not based on transaction template Last, we observe that a conflict between  $T_2$  and an instantiation of transaction template Last must be ww-conflicting involving variables  $\mathbf{B}$ , thus with  $\mu_2(\mathbf{B}) = \mu_3(\mathbf{B})$ . Then,  $\mu_2(\mathbf{S}_0) = \mu_3(\mathbf{S}_t)$ , due to functional constraint  $\mathbf{S}_0 = f_{\text{top-string}}(\mathbf{B})$  in  $T_2$  and  $\mathbf{S}_t = f_{\text{top-string}}(\mathbf{B})$  in  $T_3$ , and  $\mu_2(\mathbf{S}_1) = \mu_3(\mathbf{S}_1)$ , due to functional constraint  $\mathbf{S}_1 = f_{\text{future-solution-string}}(\mathbf{B})$  in  $T_2$  and  $T_3$ . However, we also have  $\mu_3(\mathbf{S}_1) = \mu_3(\mathbf{S}_t)$ , due to constraints  $\mathbf{S}_1 = f_{\text{solution-string}}(\mathbf{B})$

and  $\mathbf{S}_t = f_{\text{solution-string}}(\mathbf{B})$ , thus implying  $\mu_2(\mathbf{S}_0) = \mu_3(\mathbf{S}_t) = \mu_3(\mathbf{S}_1) = \mu_2(\mathbf{S}_1)$ , which contradicts with earlier proven Lemma 4.6. We conclude that  $T_3$  is indeed based on a domino transaction template. Therefore, the conflict quadruple  $(T_2, b_2, a_3, T_3)$  must admit ww-conflicting operations over variable  $\mathbf{B}$  in  $T_2$  and either variable  $\mathbf{B}$  or  $\mathbf{B}_{\text{next}}$  in  $T_3$ . We notice that  $\mu_2(\mathbf{S}_1) = f_{\text{future-solution-string}}(\mu_2(\mathbf{B}))$ ,  $\mu_3(\mathbf{S}_1) = f_{\text{future-solution-string}}(\mu_3(\mathbf{B}))$ , and  $\mu_3(\mathbf{S}_1) = f_{\text{future-solution-string}}(\mu_3(\mathbf{B}_{\text{next}}))$ , thus independent of the variable  $\mathbf{B}_{\text{next}}$  or  $\mathbf{B}$  in  $T_3$ , we have  $\mu_2(\mathbf{S}_1) = \mu_3(\mathbf{S}_1)$ .  $\square$

**Lemma 4.8.** *For a transaction  $T_i$ , with  $i \geq 4$ , for which all  $T_j$ 's, with  $j \in \{3, \dots, i-1\}$ , are based on domino transaction templates, transaction  $T_i$  is based on a domino transaction template or on transaction template Last. Furthermore  $\mu_2(\mathbf{S}_1) = \mu_i(\mathbf{S}_1)$ .*

*Proof.* Since domino transaction templates do not mention variables of type InitialConflict and write only to variables of type DominoSequence, it remains to show that  $T_{i+1}$  is not based on transaction template First.

For this, observe that  $\mu_2(\mathbf{S}_1) = \mu_{i-1}(\mathbf{S}_1)$ . Indeed, every conflict quadruple  $(T_i, b_i, a_{i+1}, T_{i+1})$ , with  $i \in \{3, \dots, i-1\}$ , admits ww-conflicting operations with variables of type DominoSequence. No matter if the conflict is via a variable  $\mathbf{B}$  or  $\mathbf{B}_{\text{next}}$ , the constraints  $\mathbf{S}_1 = f_{\text{future-solution-string}}(\mathbf{B})$  and  $\mathbf{S}_1 = f_{\text{future-solution-string}}(\mathbf{B}_{\text{next}})$  ensure  $\mu_2(\mathbf{S}_1) = \mu_{i-1}(\mathbf{S}_1)$ .

Now, assume towards a contradiction that  $T_{i+1}$  is based on First, thus admitting a conflict quadruple  $(T_i, b_i, a_{i+1}, T_{i+1})$  in  $C$ . Then either  $b_i = \mu_i(\mathbf{B})$  and  $a_{i+1} = \mu_{i+1}(\mathbf{B})$  or  $b_i = \mu_i(\mathbf{B}_{\text{next}})$  and  $a_{i+1} = \mu_{i+1}(\mathbf{B})$ . Both of these equalities imply  $\mu_i(\mathbf{S}_1) = \mu_{i+1}(\mathbf{S}_1)$  due to constraints  $\mathbf{S}_1 = f_{\text{future-solution-string}}(\mathbf{B})$  and  $\mathbf{S}_1 = f_{\text{future-solution-string}}(\mathbf{B}_{\text{next}})$ , thus implying  $\mu_i(\mathbf{S}_1) = \mu_2(\mathbf{S}_1)$ , this contradict with Lemma 4.5. We conclude that  $T_i$  is indeed based on a domino transaction template or on transaction template Last. That  $\mu_{i-1}(\mathbf{S}_1) = \mu_i(\mathbf{S}_1)$  follows again from the constraints using function  $f_{\text{future-solution-string}}$ .  $\square$

**Lemma 4.9.** *If  $T_i$  is based on transaction template Last, then  $i = m$ .*

*Proof.* Let  $T_j$  be the transaction following  $T_i$ . We already know about  $T_j$  that either  $j = 1$  or must be a transaction that is different to all foregoing transactions  $T_1, \dots, T_i$  (due to Lemma 4.3).

We first show, by exclusion, that transaction  $T_j$  is based on Split: Transaction  $T_j$  cannot be based on Last, as then either  $\mu_i(\mathbf{B}) = \mu_j(\mathbf{B})$  or  $\mu_i(\mathbf{C}) = \mu_j(\mathbf{C})$ , which directly contradicts Lemma 4.5. Similarly, transaction  $T_j$  cannot be based on First, as then  $\mu_i(\mathbf{B}) = \mu_j(\mathbf{B})$  implying  $\mu_2(\mathbf{S}_1) = \mu_i(\mathbf{S}_1) = \mu_j(\mathbf{S}_1)$ , due to the constraints involving function  $f_{\text{future-solution-string}}$ . Finally, transaction  $T_j$  cannot be based on a domino transaction template, because then  $\mu_j(\mathbf{B}_{\text{next}}) = \mu_i(\mathbf{B}) = \mu_j(\mathbf{B})$  or  $\mu_{i-1}(\mathbf{B}_{\text{next}}) = \mu_i(\mathbf{B}) = \mu_j(\mathbf{B}_{\text{next}})$ , thus with  $T_i$  and  $T_j$  contradicting Lemma 4.5. We can thus indeed conclude that transaction  $T_j$  is based on Split.

To see that  $j = 1$ , recall that  $\mu_1(\mathbf{S}_1) = \mu_{i-1}(\mathbf{S}_1)$  and the only possible conflict between  $T_i$  and  $T_j$  implies  $\mu_i(\mathbf{C}) = \mu_j(\mathbf{C})$ . From the latter we obtain  $\mu_{i-1}(\mathbf{S}_1) = \mu_i(\mathbf{S}_1)$ , due to  $\mu_{i-1}(\mathbf{B}_{\text{next}}) = \mu_i(\mathbf{B})$  and function  $f_{\text{future-solution-string}}$ . From this it follows that  $\mu_1(\mathbf{I}) = \mu_j(\mathbf{I})$  through  $\mathbf{I} = (f_{\text{defines}} \circ f_{\text{is-non-empty}})(\mathbf{S}_1)$  in transaction template Split. That  $j = 1$  then follows from Lemma 4.5.  $\square$

Finally, we show that if there is a multiversion split schedule with the properties of Lemma 4.3 that is a schedule-encoding for a sequence of dominoes  $\mathbf{d}$ , then this sequence  $\mathbf{d}$  is also a solution to the respective PCP problem. The next Proposition thus finalizes the proof for the if-direction of Theorem 4.1.

**Proposition 4.10.** *Let  $\mathcal{D}$  be a set of dominoes. Let  $s$  be a multiversion split schedule with the properties of Lemma 4.3 that is consistent with the transaction templates in Figure 5 for  $\mathcal{D}$  and with some database  $\mathbf{D}$ . If  $s$  is a schedule-encoding of a sequence  $\mathbf{d}$  of dominoes in  $\mathcal{D}$ , then  $\mathbf{d}$  is a solution for the PCP problem on input  $\mathcal{D}$ .*

*Proof.* Let  $a_1a_2 \dots a_h$  and  $b_1b_2 \dots b_k$  be the two strings (with  $a_i, b_i \in \Sigma$ ) obtained by reading from left to right, symbol by symbol, the values on the top, respectively, the bottom of dominoes  $d_1, \dots, d_r$ . Let us say that  $a_1a_2 \dots a_h = \mathbf{a}_1\mathbf{a}_2 \dots \mathbf{a}_r$  and  $b_1b_2 \dots b_k = \mathbf{b}_1\mathbf{b}_2 \dots \mathbf{b}_r$ . Notice that  $h$  and  $k$  are not necessarily equal to  $r$  as the top and bottom value of an individual domino can be of different length.

For convenience of notation, we introduce for every  $i \in \{0, \dots, h\}$  and  $j \in \{1, \dots, k\}$  the following notation:

$$\begin{aligned}\alpha_i &:= (f_{\text{append-}a_i} \circ f_{\text{append-}a_{i-1}} \circ \dots \circ f_{\text{append-}a_1})(\mu_2(\mathbf{S}_0)). \\ \beta_j &:= (f_{\text{append-}b_j} \circ f_{\text{append-}b_{j-1}} \circ \dots \circ f_{\text{append-}b_1})(\mu_2(\mathbf{S}_0)).\end{aligned}$$

First, we show that

$$\alpha_h = \mu_2(\mathbf{S}_1) = \beta_k. \quad (4.1)$$

This result follows from the assumed structure of schedule  $s$ . More precisely, since an instantiation of First with an instantiation of  $\text{Domino}_{\mathbf{d}_1}$  can only have conflicts on instantiations of  $\mathbf{W}[\mathbf{B} : \text{DominoSequence}]$ , we have  $\mu_2(\mathbf{B}) = \mu_3(\mathbf{B})$ , from which it follows that  $\mu_2(\mathbf{S}_1) = \mu_3(\mathbf{S}_1)$ .

For every individual instantiation of  $\text{Domino}_{\mathbf{d}_i}$  in  $s$ , we have that  $f_{\text{append-}a_i^{\ell_a}} \circ \dots \circ f_{\text{append-}a_i^1}(\mu_i(\mathbf{S}_t)) = \mu_i(\mathbf{S}_{t\mathbf{a}_i})$  and  $f_{\text{append-}b_i^{\ell_b}} \circ \dots \circ f_{\text{append-}b_i^1}(\mu_i(\mathbf{S}_b)) = \mu_i(\mathbf{S}_{b\mathbf{b}_i})$ , with  $\mathbf{a}_i = a_1a_2 \dots a_{\ell_a}$  and  $\mathbf{b}_i = b_1b_2 \dots b_{\ell_b}$ .

For transactions  $T_i$ , with  $i \in \{3, \dots, m+1\}$ , (thus representing an instantiation of  $\text{Domino}_{\mathbf{d}_{i-2}}$  which is followed in  $s$  by an instantiation of  $\text{Domino}_{\mathbf{d}_{i-1}}$ ), the only possible conflict is between the instantiation of  $\mathbf{W}[\mathbf{B}_{\text{next}} : \text{DominoSequence}](T_i)$  and the instantiation of  $\mathbf{W}[\mathbf{B} : \text{DominoSequence}]$  in  $(T_{i+1})$  – notice that this is indeed the only option due to Lemma 4.5 – thus with  $\mu_i(\mathbf{B}_{\text{next}}) = \mu_{i+1}(\mathbf{B})$ , implying  $\mu_i(\mathbf{S}_{t\mathbf{a}_i}) = \mu_{i+1}(\mathbf{S}_t)$ ,  $\mu_i(\mathbf{S}_{b\mathbf{b}_i}) = \mu_{i+1}(\mathbf{S}_b)$ , and  $\mu_i(\mathbf{S}_1) = \mu_{i+1}(\mathbf{S}_1)$ .

Finally, transaction  $T_{m-1}$  (an instantiation of  $\text{Domino}_{\mathbf{d}_m}$ ) can only conflict with transaction  $T_m$  (an instantiation of Last) on instantiations of  $\mathbf{B}_{\text{next}}$  (in  $\text{Domino}_{\mathbf{d}_m}$ , and  $\mathbf{B}$  (in Last), thus with  $\mu_{m-1}(\mathbf{B}_{\text{next}}) = \mu_m(\mathbf{B})$ , implying  $\mu_{m-1}(\mathbf{S}_{t\mathbf{a}_r}) = \mu_m(\mathbf{S}_t) = \mu_m(\mathbf{S}_b) = \mu_{m-1}(\mathbf{S}_{b\mathbf{b}_r})$ .

Combining the above equalities indeed proves Condition (4.1).

From Condition (4.1) we can now derive that,

$$\alpha_i = \mu_2(\mathbf{S}_1) = \beta_i, \text{ for every } i \in \{1, \dots, \min\{h, k\}\}, \quad (4.2)$$

by following an analogous approach. Indeed, in every instantiation of  $\text{Domino}_i$ , there is a functional constraint for every application of the append function that requires its input to be the result of the detach function applied over its output, which indeed implies Condition (4.2).

To see that  $k = h$ , we observe that  $k \neq h$  implies an application of the detach function over the instantiation of  $\mathbf{S}_1$  (for which we already argued it has the same tuple assigned for every domino instantiation) for the shortest string, which contradicts with Condition (4.1) because such an application results in the same instantiation as  $\mathbf{S}_e$ , which can never equal the instantiation for  $\mathbf{S}_1$ .

The desired result that the individual symbols in the top and bottom reads of dominoes in sequence  $\mathbf{d}$  are the same now follows from the functional constraint that every interpretation of




FIGURE 7. Schema graph for the SmallBank benchmark. The dashed edges correspond to the multi-tree schema graph for the schema restricted to  $f_{A \rightarrow S}$  and  $f_{A \rightarrow C}$ .

a string sequence mapped via function  $f_{\text{top}}$  onto either the interpretation for  $\mathbf{S}_1$  (representing symbol  $1 \in \Sigma$ ) or  $\mathbf{S}_0$  (representing symbol  $0 \in \Sigma$ ).  $\square$

## 5. ROBUSTNESS FOR TEMPLATES ADMITTING MULTI-TREE BIJECTIVITY

We say that a set of transaction templates  $\mathcal{P}$  over a schema  $(\text{Rels}, \text{Funcs})$  admits *multi-tree bijectivity* if a disjoint partitioning of  $\text{Funcs}$  in pairs  $(f_1, g_1), (f_2, g_2), \dots, (f_n, g_n)$  exists such that  $\text{dom}(f_i) = \text{range}(g_i)$  and  $\text{dom}(g_i) = \text{range}(f_i)$  for every pair of function names  $(f_i, g_i)$ ; every schema graph  $SG(\text{Rels}, \{h_1, h_2, \dots, h_n\})$  over the schema restricted to function names  $\{h_1, h_2, \dots, h_n\}$  (with  $h_i = f_i$  or  $h_i = g_i$ ) is a multi-tree; and, for every pair of function names  $(f_i, g_i)$  and for every pair of variables  $\mathbf{X}, \mathbf{Y}$  occurring in a template  $\tau_j \in \mathcal{P}$ , we have  $f_i(\mathbf{X}) = \mathbf{Y} \in \Gamma_j$  iff  $g_i(\mathbf{Y}) = \mathbf{X} \in \Gamma_j$ . Intuitively, we can think of  $f_i$  as a bijective function, with  $g_i$  its inverse. We denote the class of all sets of templates admitting multi-tree bijectivity by **MTBTemp**. The SmallBank benchmark given in Figure 1 is in **MTBTemp**, witnessed by the partitioning  $\{(f_{A \rightarrow C}, f_{C \rightarrow A}), (f_{A \rightarrow S}, f_{S \rightarrow A})\}$ . For example, the schema graph restricted to  $f_{A \rightarrow C}$  and  $f_{A \rightarrow S}$  is a tree and therefore also a multi-tree, as illustrated in Figure 7.

The next theorem allows disequalities whereas Theorem 4.1 does not require them.

**Theorem 5.1.** T-ROBUSTNESS(**MTBTemp**, RC) is decidable in NLOGSPACE.

The approach followed in the proof of Theorem 5.1 is to repeatedly pick a transaction template while maintaining an overall consistent variable mapping in search for a counterexample multiversion split schedule that by Theorem 3.7 suffices to show that robustness does not hold. The main challenge is to show that a variable mapping consistent with all functional constraints can be maintained in logarithmic space and that all requirements for a multiversion split schedule can be verified in NLOGSPACE.

Central to our approach is a generalization of conflicting operations. Let  $\mathcal{P}$  be a set of transaction templates. For  $\tau_i$  and  $\tau_j$  in  $\mathcal{P}$ , we say that an operation  $o_i \in \tau_i$  is *potentially conflicting* with an operation  $o_j \in \tau_j$  if  $o_i$  and  $o_j$  are operations over a variable of the same type, and at least one of the following holds:

- $\text{WriteSet}(o_i) \cap \text{WriteSet}(o_j) \neq \emptyset$  (potentially ww-conflicting);
- $\text{WriteSet}(o_i) \cap \text{ReadSet}(o_j) \neq \emptyset$  (potentially wr-conflicting); or
- $\text{ReadSet}(o_i) \cap \text{WriteSet}(o_j) \neq \emptyset$  (potentially rw-conflicting).

Intuitively, potentially conflicting operations lead to conflicting operations when the variables of these operations are mapped to the same tuple by a variable assignment. In analogy to conflicting quadruples over a set of transactions as in Definition 3.6, we consider *potentially conflicting quadruples*  $(\tau_i, o_i, p_j, \tau_j)$  over  $\mathcal{P}$  with  $\tau_i, \tau_j \in \mathcal{P}$ , and  $o_i \in \tau_i$  an operation that is potentially conflicting with an operation  $p_j \in \tau_j$ . For a sequence of potentially conflicting

quadruples  $D = (\tau_1, o_1, p_2, \tau_2), \dots, (\tau_m, o_m, p_1, \tau_1)$  over  $\mathcal{P}$ , we write  $\text{Trans}(D)$  to denote the set  $\{\tau_1, \dots, \tau_m\}$  of transaction templates mentioned in  $D$ . For ease of exposition, we assume a variable renaming such that any pair of templates in  $\text{Trans}(D)$  uses a disjoint set of variables.<sup>6</sup> The sequence  $D$  induces a sequence of conflicting quadruples  $C = (T_1, b_1, a_2, T_2), \dots, (T_m, b_m, a_1, T_1)$  by applying a variable assignment  $\mu_i$  to each  $\tau_i$  in  $\text{Trans}(D)$ . We call such a set of variable assignments simply a *variable mapping* for  $D$ , denoted  $\bar{\mu}$ , and write  $\bar{\mu}(D) = C$ . For a variable  $X$  occurring in a template  $\tau_i$ , we write  $\bar{\mu}(X)$  as a shorthand notation for  $\mu_i(X)$ , with  $\mu_i$  the variable assignment over  $\tau_i$  in  $\bar{\mu}$ . This is well-defined as all templates in  $\text{Trans}(D)$  are variable-disjoint. Furthermore,  $\bar{\mu}(\text{var}(o_i)) = \bar{\mu}(\text{var}(p_j))$  for each potentially conflicting quadruple  $(\tau_i, o_i, p_j, \tau_j)$  in  $D$  as otherwise the induced quadruple  $(T_i, b_i, a_j, T_j)$  is not a valid conflicting quadruple in  $C$ . We say that a variable mapping  $\bar{\mu}$  is admissible for a database  $\mathbf{D}$  if every variable assignment  $\mu_i$  in  $\bar{\mu}$  is admissible for  $\mathbf{D}$ .

A basic insight is that if there is a multiversion split schedule  $s$  for some  $C$  over a set of transactions  $\mathcal{T}$  consistent with  $\mathcal{P}$  and a database  $\mathbf{D}$ , then there is a sequence of potentially conflicting quadruples  $D$  such that  $\bar{\mu}(D) = C$  for some  $\bar{\mu}$ . We will verify the existence of such a  $C$ , satisfying the properties of Definition 3.6, by nondeterministically constructing  $D$  on-the-fly together with a mapping  $\bar{\mu}$ . We show in Lemma 5.5 that when  $\mathcal{P} \in \mathbf{MTBTemp}$ ,  $\bar{\mu}$  is a collection of disjoint type mappings (that map variables of the same type to the same tuple) such that variables that are “connected” in  $D$  (in a way that we will make precise next) are mapped using the same type mapping. Lemma 5.6 then shows that already a constant number of those type mappings suffice.

We introduce the necessary notions to capture when two variables are connected in  $D$ . We can think of equality constraints  $Y = f(X)$  in a template  $\tau$  as constraints on the possible variable assignments  $\mu$  for  $\tau$  when a database  $\mathbf{D}$  is given. Indeed, if we fix  $\mu(X)$  to a tuple in  $\mathbf{D}$ , then  $\mu(Y) = f^{\mathbf{D}}(\mu(X))$  is immediately implied. These constraints can cause a chain reaction of implications. If for example  $Z = g(Y)$  is a constraint in  $\tau$  as well, then  $\mu(X)$  immediately implies  $\mu(Z) = g^{\mathbf{D}}(f^{\mathbf{D}}(\mu(X)))$ . We formalize this notion of implication next. We use sequences of function names  $F = f_1 \cdots f_n$ , denoting the empty sequence as  $\varepsilon$  and the concatenation of two sequences  $F$  and  $G$  by  $F \cdot G$ . For two variables  $X, Y$  occurring in a template  $\tau$  and a (possibly empty) sequence of function names  $F$ , we say that  $X$  *implies*  $Y$  *by*  $F$  *in*  $\tau$ , denoted  $X \overset{F}{\rightsquigarrow}_{\tau} Y$ , if  $X = Y$  and  $F = \varepsilon$  or if there is a variable  $Z$  such that  $Y = f(Z)$  is a constraint in  $\tau$ ,  $X \overset{F'}{\rightsquigarrow}_{\tau} Z$  and  $F = F' \cdot f$ . We next extend the notions of implication to sequences of potentially conflicting quadruples. Let  $D = (\tau_1, o_1, p_2, \tau_2), \dots, (\tau_m, o_m, p_1, \tau_1)$  be a sequence of potentially conflicting quadruples, and let  $X$  and  $Y$  be two variables occurring in templates  $\tau_i$  and  $\tau_j$  in  $\text{Trans}(D)$ , respectively. Then  $X$  *implies*  $Y$  *by a sequence of function names*  $F$  *in*  $D$ , denoted  $X \overset{F}{\rightsquigarrow}_D Y$  if

- $i = j$  and  $X \overset{F}{\rightsquigarrow}_{\tau_i} Y$  (implication within the same template);
- $F = \varepsilon$  and  $(\tau_i, o_i, p_j, \tau_j)$  or  $(\tau_j, o_j, p_i, \tau_i)$  is a potentially conflicting quadruple in  $D$  with  $o_i$  (respectively  $p_i$ ) an operation over  $X$  and  $p_j$  (respectively  $o_j$ ) an operation over  $Y$  (implication between templates, notice that  $X \overset{F}{\rightsquigarrow}_D Y$  iff  $Y \overset{F}{\rightsquigarrow}_D X$ ); or
- there exists a variable  $Z$  such that  $X \overset{F_1}{\rightsquigarrow}_D Z$  and  $Z \overset{F_2}{\rightsquigarrow}_D Y$  with  $F = F_1 \cdot F_2$ .

Two variables  $X$  and  $Y$  occurring in  $\text{Trans}(D)$  are *connected in*  $D$ , denoted  $X \approx_D Y$ , if  $X \overset{F}{\rightsquigarrow}_D Y$  or  $Y \overset{F}{\rightsquigarrow}_D X$ , or if there is a variable  $Z$  with  $X \approx_D Z$  and either  $Z \overset{F}{\rightsquigarrow}_D Y$  or  $Y \overset{F}{\rightsquigarrow}_D Z$  for some

<sup>6</sup>To be formally correct, the latter would require to add every such variable-renamed template to  $\mathcal{P}$  creating a larger set  $\mathcal{P}'$ . This does not influence the complexity of Theorem 5.1 as  $\text{Trans}(D)$  nor  $\mathcal{P}'$  are used in the algorithm. Their only purpose is to reason about properties of  $\bar{\mu}$ .



sequence  $F$ . Furthermore, two variables  $X$  and  $Y$  occurring in a template  $\tau$  are *connected in*  $\tau$ , denoted  $X \approx_{\tau} Y$ , if  $X \xrightarrow{F} Y$  or  $Y \xrightarrow{F} X$ , or if there is a variable  $Z$  with  $X \approx_{\tau} Z$  and either  $Z \xrightarrow{F} Y$  or  $Y \xrightarrow{F} Z$  for some sequence  $F$ . These definitions of connectedness can be trivially extended to operations over variables: two operations in  $D$  (respectively  $\tau$ ) are connected in  $D$  (respectively  $\tau$ ) if they are over variables that are connected in  $D$  (respectively  $\tau$ ). When  $F$  is not important we drop it from the notation. For instance, we denote by  $X \approx_D Y$  that there is an  $F$  with  $X \xrightarrow{F} Y$ .

**Lemma 5.2.** *Let  $D$  be a sequence of potentially conflicting quadruples over  $\mathcal{P} \in \mathbf{MTBTemp}$ . Then  $X \approx_D Y$  implies  $X \rightsquigarrow_D Y$  and  $Y \rightsquigarrow_D X$ . Furthermore, if  $\text{type}(X) = \text{type}(Y)$  then  $\bar{\mu}(X) = \bar{\mu}(Y)$  for every variable mapping  $\bar{\mu}$  for  $D$  that is admissible for some database  $D$ .*

Before proving the correctness of Lemma 5.2, we first present two additional lemmas that will be used in the correctness proof.

**Lemma 5.3.** *Let  $(\text{Rels}, \text{Funcs})$  be a schema for which a disjoint partitioning of  $\text{Funcs}$  in pairs  $P = (f_1, g_1), (f_2, g_2), \dots, (f_n, g_n)$  exists such that  $\text{dom}(f_i) = \text{range}(g_i)$  and  $\text{dom}(g_i) = \text{range}(f_i)$  for every  $(f_i, g_i) \in P$  and every schema graph  $SG(\text{Rels}, \{h_1, h_2, \dots, h_n\})$  over the schema restricted to function names  $\{h_1, h_2, \dots, h_n\}$  with  $h_i \in (f_i, g_i)$  is a multi-tree. Then:*

- (1) *there is no function name  $f \in \text{Funcs}$  with  $\text{dom}(f) = \text{range}(f)$ ; and*
- (2) *for every path in  $SG(\text{Rels}, \text{Funcs})$  visiting a sequence of nodes  $R_1, R_2, \dots, R_{m-1}, R_m$ , if  $R_1 = R_m$  and  $R_1 \neq R_i$  for every  $i \in [2, m-1]$ , then  $R_2 = R_{m-1}$  and  $(f, g)$  is a pair in  $P$  with  $f$  the edge from  $R_1$  to  $R_2$  and  $g$  the edge from  $R_{m-1}$  to  $R_m$ .*

*Proof.* Towards a contradiction, assume (1) does not hold. That is, there is a function name  $f_i$  with  $\text{dom}(f_i) = \text{range}(f_i) = R$  for some type  $R$ . Let  $g_i$  be the function name such that  $(f_i, g_i)$  is a pair in  $P$ . By definition,  $\text{dom}(g_i) = \text{range}(g_i) = R$ . But then we cannot pick a  $h_i \in (f_i, g_i)$  such that the resulting schema graph is a multi-tree. Indeed, in both cases, there is a self-loop on  $R$ , leading to the desired contradiction.

For (2), assume towards a contradiction that  $R_2 \neq R_{m-1}$ . Without loss of generality, we can assume that each node is visited only once in  $R_2, \dots, R_{m-1}$ . Otherwise,  $R_2, \dots, R_{m-1}$  contains a loop that can be removed from this sequence without altering  $R_2$  and  $R_{m-1}$ . Since  $R_1, R_2, \dots, R_{m-1}, R_m$  is a path in  $SG(\text{Rels}, \text{Funcs})$ , there is a sequence of function names  $F = e_1 \cdots e_{m-1}$  such that each  $e_i$  is an edge from  $R_i$  to  $R_{i+1}$  in  $SG(\text{Rels}, \text{Funcs})$ , implying  $\text{dom}(e_i) = R_i$  and  $\text{range}(e_i) = R_{i+1}$ . By assumption that each type  $R_i$  occurs only once in  $R_2, \dots, R_{m-1}$  (notice that, for  $i = 1$ , this follows from Condition (2) of the lemma) and type  $R_1 = R_m$  does not appear in  $R_2, \dots, R_{m-1}$ , there is no pair of function names  $e_i$  and  $e_j$  in  $F$  with  $i \neq j$ ,  $\text{dom}(e_i) = \text{range}(e_j)$  and  $\text{range}(e_i) = \text{dom}(e_j)$ . Therefore, at most one function name of each pair in  $P$  appears in  $F$ . But then we can choose  $h_i = e_i$  for each such pair in  $P$ , with  $e_i$  the function name appearing in  $F$ . Since  $F$  describes a cycle in  $SG(\text{Rels}, \text{Funcs})$ , the resulting schema graph restricted to these  $h_i$  cannot be a multi-tree, as it contains a cycle.

It remains to argue that if  $f$  is the edge from  $R_1$  to  $R_2$  and  $g$  is the edge from  $R_{m-1}$  to  $R_m$  on this path, then  $(f, g)$  is a pair in  $P$ . To this end, note that  $\text{dom}(f) = \text{range}(g)$  and  $\text{range}(f) = \text{dom}(g)$ , as  $R_1 = R_m$  and  $R_2 = R_{m-1}$ . If  $(f, g)$  is not a pair in  $P$ , then there are two pairs  $(f, f')$  and  $(g, g')$  in  $P$  with  $\text{dom}(f) = \text{dom}(g') = \text{range}(f') = \text{range}(g) = R_1$  and  $\text{range}(f) = \text{range}(g') = \text{dom}(f') = \text{dom}(g) = R_2$ . Then we can choose  $f$  in  $(f, f')$  and  $g$  in  $(g, g')$ . Since the resulting schema graph cannot be a multi-tree, as there is a cycle between  $R_1$  and  $R_2$ , this choice leads to a contradiction.  $\square$

**Lemma 5.4.** *Let  $D$  be a sequence of potentially conflicting quadruples over  $\mathcal{P} \in \mathbf{MTBTemp}$ . Then*

- (1)  $X \rightsquigarrow_{\tau} Y$  iff  $Y \rightsquigarrow_{\tau} X$  for every pair of variables  $X$  and  $Y$  occurring in a template  $\tau$ ; and
- (2)  $X \rightsquigarrow_D Y$  iff  $Y \rightsquigarrow_D X$  for every pair of variables  $X$  and  $Y$  occurring in  $D$ .

*Proof.* (1) We argue by induction on the definition of  $X \rightsquigarrow_{\tau} Y$  that  $X \rightsquigarrow_{\tau} Y$  implies  $Y \rightsquigarrow_{\tau} X$ . The other direction is analogous. The base case is immediate, as  $X = Y$  implies  $Y \rightsquigarrow_{\tau} X$  by definition. For the inductive case, assume a variable  $Z$  such that  $Y = f(Z)$  is a constraint in  $\tau$  and  $X \rightsquigarrow_{\tau} Z$ . By the induction hypothesis,  $Z \xrightarrow{F} X$  for some sequence of function names  $F$ . Since  $\mathcal{P} \in \mathbf{MTBTemp}$ , there is a constraint  $Z = f'(Y)$  in  $\tau$  as well. It follows that  $Y \xrightarrow{F'} X$  with  $F' = f' \cdot F$ .

(2) We argue by induction on the definition of  $X \rightsquigarrow_D Y$  that  $X \rightsquigarrow_D Y$  implies  $Y \rightsquigarrow_D X$ . The other direction is again analogous. The first base case is now immediate, as we already argued that  $X \rightsquigarrow_{\tau} Y$  implies  $Y \rightsquigarrow_{\tau} X$ . For the second base case, assume  $X \xrightarrow{\tau_i, o_i, p_j, \tau_j} Y$  and  $(\tau_i, o_i, p_j, \tau_j)$  is a potentially conflicting quadruple in  $D$  with  $\text{var}(o_i) = X$  and  $\text{var}(p_j) = Y$  (the case for  $(\tau_j, o_j, p_i, \tau_i)$  is analogous).  $Y \xrightarrow{\tau_j, o_j, p_i, \tau_i} X$  then follows by definition. For the inductive case, let  $Z$  be a variable such that  $X \xrightarrow{F_1} Z$  and  $Z \xrightarrow{F_2} Y$ . Then by induction hypothesis  $Z \xrightarrow{F_1} X$  and  $Y \xrightarrow{F_2} Z$  for some sequence of function names  $F_1'$  and  $F_2'$ . By definition,  $Y \xrightarrow{F'} X$  with  $F' = F_2' \cdot F_1'$ .  $\square$

We are now ready to prove the correctness of Lemma 5.2.

*Proof of Lemma 5.2.* Assuming  $X \approx_D Y$ , we first show by induction on the definition of connectedness that  $X \rightsquigarrow_D Y$ . By Lemma 5.4,  $Y \rightsquigarrow_D X$  then follows. For the base case, both  $X \rightsquigarrow_D Y$  and  $Y \rightsquigarrow_D X$  imply  $X \rightsquigarrow_D Y$ , where the former is immediate and the latter is by Lemma 5.4. For the inductive case, let  $Z$  be a variable with  $X \approx_D Z$  and either  $Z \rightsquigarrow_D Y$  or  $Y \rightsquigarrow_D Z$ . Again,  $Z \xrightarrow{F_2} Y$  for some sequence of function names  $F_2$  is implied in both cases. By induction hypothesis,  $X \xrightarrow{F_1} Z$  for some sequence of function names  $F_1$ . As a result,  $X \xrightarrow{F} Y$  with  $F = F_1 \cdot F_2$ .

Next, let  $X$  and  $Y$  be two variables occurring in  $\text{Trans}(D)$  with  $X \approx_D Y$  and  $\text{type}(X) = \text{type}(Y)$  and let  $\bar{\mu}$  be a variable mapping for  $D$  that is admissible for a database  $\mathbf{D}$ . We prove that  $\bar{\mu}(X) = \bar{\mu}(Y)$ .

We already argued that  $X \approx_D Y$  implies  $X \rightsquigarrow_D Y$ . By definition of  $X \rightsquigarrow_D Y$ , there is a sequence of variables  $X_1, X_2 \dots, X_n$  with  $X_1 = X$  and  $X_n = Y$  such that for each pair of adjacent variables  $X_i$  and  $X_{i+1}$ :

- (†)  $X_i$  and  $X_{i+1}$  both occur in the same template  $\tau \in \text{Trans}(D)$  and  $X_{i+1} = f(X_i) \in \Gamma(\tau)$  for some function name  $f$ ; or
- (‡)  $\text{type}(X_i) = \text{type}(X_{i+1})$  and there is a potentially conflicting quadruple  $(\tau_j, o_j, p_k, \tau_k)$  in  $D$  with either  $\text{var}(o_j) = X_i$  and  $\text{var}(p_k) = X_{i+1}$  or  $\text{var}(p_k) = X_i$  and  $\text{var}(o_j) = X_{i+1}$ .

In the remainder of this proof, we show that for each pair of variables  $X_i$  and  $X_j$  in this sequence with  $\text{type}(X_i) = \text{type}(X_j)$  that  $\bar{\mu}(X_i) = \bar{\mu}(X_j)$ . The desired  $\bar{\mu}(X) = \bar{\mu}(Y)$  then follows immediately as  $X = X_1$  and  $Y = X_n$ . Note that it suffices to show this property only for pairs of variables  $X_i$  and  $X_j$  for which no variable  $X_k$  exists with  $i < k < j$  and  $\text{type}(X_i) = \text{type}(X_j) = \text{type}(X_k)$ . Indeed, if such an  $X_k$  exists, we can recursively argue that  $\bar{\mu}(X_i) = \bar{\mu}(X_k)$  and  $\bar{\mu}(X_k) = \bar{\mu}(X_j)$ . The argument is by induction on the number of variables between  $X_i$  and  $X_j$ .

If  $j = i + 1$  (*base case*), then (†) applies to  $X_i$  and  $X_j$ . Indeed, if (†) would apply instead, then there would be a function name  $f$  with  $\text{dom}(f) = \text{type}(X_i) = \text{type}(X_j) = \text{range}(f)$ , contradicting Condition (1) of Lemma 5.3. By definition of  $\bar{\mu}$ , we have  $\bar{\mu}(X_i) = \bar{\mu}(X_j)$ .

Next, let  $i + 1 < j$  (*inductive case*), and assume that  $\bar{\mu}(\mathbf{X}_k) = \bar{\mu}(\mathbf{X}_\ell)$  for all  $\mathbf{X}_k$  and  $\mathbf{X}_\ell$  with  $i < k \leq \ell < j$  and  $\text{type}(\mathbf{X}_k) = \text{type}(\mathbf{X}_\ell)$  (*induction hypothesis*). From this sequence  $\mathbf{X}_i, \dots, \mathbf{X}_j$ , we derive a sequence of function names  $F = f_1 \cdots f_{m-1}$ , where each function name  $f_i$  is based on an application of  $(\dagger)$  on adjacent variables (notice that applications of  $(\ddagger)$  do not result in a function name being added to  $F$ ). By assumption on the types of variables  $\mathbf{X}_k$  with  $i < k < j$ , we have in particular  $\text{type}(\mathbf{X}_{i+1}) \neq \text{type}(\mathbf{X}_i)$  and  $\text{type}(\mathbf{X}_{j-1}) \neq \text{type}(\mathbf{X}_j)$ . This implies that  $(\dagger)$  is applicable for  $\mathbf{X}_i$  and  $\mathbf{X}_{i+1}$  (respectively  $\mathbf{X}_{j-1}$  and  $\mathbf{X}_j$ ). Furthermore,  $\mathbf{X}_i$  and  $\mathbf{X}_{i+1}$  appear in the same template, say  $\tau_i$  (respectively  $\tau_j$  for  $\mathbf{X}_{j-1}$  and  $\mathbf{X}_j$ ), and  $\mathbf{X}_{i+1} = f_1(\mathbf{X}_i) \in \Gamma(\tau_i)$  (respectively  $\mathbf{X}_j = f_{m-1}(\mathbf{X}_{j-1}) \in \Gamma(\tau_j)$ ). By construction,  $F$  then describes a path in  $SG(\text{Rels}, \text{Funcs})$  visiting the nodes  $R_1, R_2, \dots, R_{m-1}, R_m$  with  $\text{type}(\mathbf{X}_i) = R_1$ ,  $\text{type}(\mathbf{X}_{i+1}) = R_2$ ,  $\text{type}(\mathbf{X}_{j-1}) = R_{m-1}$  and  $\text{type}(\mathbf{X}_j) = R_m$ . Since this path satisfies Condition 2 in Lemma 5.3, it follows that  $\text{type}(\mathbf{X}_{i+1}) = \text{type}(\mathbf{X}_{j-1})$  and  $(f_1, f_{m-1})$  is a pair in the pairwise partitioning of  $\text{Funcs}$  witnessing  $\mathcal{P} \in \mathbf{MTBTemp}$ . By definition of  $\mathbf{MTBTemp}$ ,  $\mathbf{X}_{i+1} = f_1(\mathbf{X}_i) \in \Gamma(\tau_i)$  then implies  $\mathbf{X}_i = f_{m-1}(\mathbf{X}_{i+1}) \in \Gamma(\tau_i)$ . According to the induction hypothesis,  $\bar{\mu}(\mathbf{X}_{i+1}) = \bar{\mu}(\mathbf{X}_{j-1})$ . Since  $\bar{\mu}$  is admissible for  $\mathbf{D}$ , we conclude that  $\bar{\mu}(\mathbf{X}_i) = f_{m-1}^{\mathbf{D}}(\bar{\mu}(\mathbf{X}_{i+1})) = f_{m-1}^{\mathbf{D}}(\bar{\mu}(\mathbf{X}_{j-1})) = \bar{\mu}(\mathbf{X}_j)$ .  $\square$

It follows from Lemma 5.2 that, if we group connected variables, then the same tuple is assigned to all variables of the same type in this group. We encode this choice of tuples for variables through (total) functions  $c : \mathbf{Rels} \rightarrow \mathbf{Tuples}$  that we call *type mappings* and which map a relation onto a particular tuple of that relation's type. For instance, in *SmallBank*, a type mapping  $c$  is determined by an *Account* tuple  $\mathbf{a}$ , a *Savings* tuple  $\mathbf{s}$ , and a *Checking* tuple  $\mathbf{c}$ . The following Lemma makes explicit how  $\bar{\mu}$  can be decomposed into type mappings such that connected variables use the same type mapping and disequalities enforce the use of different type mappings.

**Lemma 5.5.** *For a multiversion split schedule  $s$  based on a sequence of conflicting quadruples  $C$  over a set of transactions  $\mathcal{T}$  consistent with a  $\mathcal{P} \in \mathbf{MTBTemp}$  and a database  $\mathbf{D}$ , let  $\bar{\mu}$  be the variable mapping for a sequence of potentially conflicting quadruples  $D$  over  $\mathcal{P}$  with  $\bar{\mu}(D) = C$ . Then, a set  $\mathcal{S}$  of type mappings over disjoint ranges and a function  $\varphi_{\mathcal{S}} : \mathbf{Var} \rightarrow \mathcal{S}$  exist with:*

- $\bar{\mu}(\mathbf{X}) = c(\text{type}(\mathbf{X}))$  for every variable  $\mathbf{X}$ , with  $c = \varphi_{\mathcal{S}}(\mathbf{X})$ ;
- $\varphi_{\mathcal{S}}(\mathbf{X}) = \varphi_{\mathcal{S}}(\mathbf{Y})$  whenever  $\mathbf{X} \approx_D \mathbf{Y}$ ; and,
- $\varphi_{\mathcal{S}}(\mathbf{X}) \neq \varphi_{\mathcal{S}}(\mathbf{Y})$  for every constraint  $\mathbf{X} \neq \mathbf{Y}$  occurring in a template  $\tau \in \text{Trans}(D)$ .

*Proof.* To aid the construction of  $\mathcal{S}$  and  $\varphi_{\mathcal{S}}$ , we first define a coloring function  $\lambda$  that assigns a color to each tuple occurring in the schedule  $s$  such that the following holds: for every pair of tuples  $\mathbf{t}$  and  $\mathbf{v}$  occurring in  $s$ :

- connected tuples are mapped to the same color: if  $\bar{\mu}(\mathbf{X}) = \mathbf{t}$ ,  $\bar{\mu}(\mathbf{Y}) = \mathbf{v}$  and  $\mathbf{X} \approx_D \mathbf{Y}$  for some variables  $\mathbf{X}, \mathbf{Y}$  occurring in  $\text{Trans}(D)$ , then  $\lambda(\mathbf{t}) = \lambda(\mathbf{v})$ ; and
- different tuples of the same type are mapped to different colors: if  $\text{type}(\mathbf{t}) = \text{type}(\mathbf{v})$  and  $\mathbf{t} \neq \mathbf{v}$ , then  $\lambda(\mathbf{t}) \neq \lambda(\mathbf{v})$ .

Note that we can always construct such a function  $\lambda$  as by Lemma 5.2, it cannot be the case that  $\text{type}(\mathbf{t}) = \text{type}(\mathbf{v})$ ,  $\mathbf{t} \neq \mathbf{v}$  and there is a pair of variables  $\mathbf{X}, \mathbf{Y}$  with  $\bar{\mu}(\mathbf{X}) = \mathbf{t}$ ,  $\bar{\mu}(\mathbf{Y}) = \mathbf{v}$ , and  $\mathbf{X} \approx_D \mathbf{Y}$ .

For  $\alpha \in \text{range}(\lambda)$ , define the type mapping  $c_\alpha$  as follows: for every type  $R \in \text{Rels}$ :

$$c_\alpha(R) = \begin{cases} \mathbf{t} & \text{if } \lambda(\mathbf{t}) = \alpha \text{ and } \text{type}(\mathbf{t}) = R, \\ \mathbf{v}_{c,R} & \text{otherwise,} \end{cases}$$

where  $\mathbf{v}_{c,R}$  is an arbitrary tuple of type  $R$  not occurring in  $s$  or any other type mapping  $c_\beta$  for  $\beta \in \text{range}(\lambda)$ . Define  $\mathcal{S} = \{c_\alpha \mid \alpha \in \text{range}(\lambda)\}$ . By construction, every type mapping in  $\mathcal{S}$  is well defined and all type mappings are over disjoint ranges. Furthermore,  $c_\alpha \neq c_\beta$  whenever  $\alpha \neq \beta$ .

We now construct  $\varphi_{\mathcal{S}}$  as follows:  $\varphi_{\mathcal{S}}(\mathbf{X}) = c_\alpha$  with  $\alpha = \lambda(\bar{\mu}(\mathbf{X}))$  for every variable  $\mathbf{X}$  occurring in  $\text{Trans}(D)$ . It remains to argue that  $\varphi_{\mathcal{S}}$  indeed satisfies all properties stated in Lemma 5.5. By construction of  $\mathcal{S}$  and  $\varphi_{\mathcal{S}}$ , we have  $\bar{\mu}(\mathbf{X}) = c(\text{type}(\mathbf{X}))$  for every variable  $\mathbf{X}$ , with  $c = \varphi_{\mathcal{S}}(\mathbf{X})$ . Towards the second property, notice that  $\mathbf{X} \approx_D \mathbf{Y}$  implies  $\varphi_{\mathcal{S}}(\mathbf{X}) = c_{\lambda(\bar{\mu}(\mathbf{X}))} = c_{\lambda(\bar{\mu}(\mathbf{Y}))} = \varphi_{\mathcal{S}}(\mathbf{Y})$  by definition of  $\lambda$  and  $\varphi_{\mathcal{S}}$ . For the last property, assume  $\mathbf{X} \neq \mathbf{Y}$  occurs in a template  $\tau \in \text{Trans}(D)$  and  $\text{type}(\mathbf{X}) = \text{type}(\mathbf{Y})$ . Since  $\bar{\mu}$  is admissible for database  $\mathbf{D}$ ,  $\bar{\mu}(\mathbf{X}) \neq \bar{\mu}(\mathbf{Y})$ . Then, by definition of  $\lambda$  and  $\varphi_{\mathcal{S}}$ , we have  $\varphi_{\mathcal{S}}(\mathbf{X}) = c_{\lambda(\bar{\mu}(\mathbf{X}))} \neq c_{\lambda(\bar{\mu}(\mathbf{Y}))} = \varphi_{\mathcal{S}}(\mathbf{Y})$ .  $\square$

From  $D = (\tau_1, o_1, p_2, \tau_2), \dots, (\tau_m, o_m, p_1, \tau_1)$  and  $\varphi_{\mathcal{S}}$  as in Lemma 5.5 we can derive a sequence of quintuples  $E = (\tau_1, o_1, c_{o_1}, p_1, c_{p_1}), \dots, (\tau_m, o_m, c_{o_m}, p_m, c_{p_m})$  such that  $c_{o_i} = \varphi_{\mathcal{S}}(\text{var}(o_i))$  and  $c_{p_i} = \varphi_{\mathcal{S}}(\text{var}(p_i))$  for  $i \in [1, m]$ . Intuitively, this sequence of quintuples can be used to reconstruct the original multiversion split schedule  $s$ . The next Lemma shows that we can decide robustness against RC over a set of transaction templates admitting multi-tree bijectivity by searching for a specific sequence of quintuples over at most four type mappings.

**Lemma 5.6.** *Let  $\mathcal{P} \in \text{MTBTemp}$  and let  $\mathcal{S} = \{c_1, c_2, c_3, c_4\}$  be a set consisting of four type mappings with disjoint ranges. Then,  $\mathcal{P}$  is not robust against RC iff there is a sequence of quintuples  $E = (\tau_1, o_1, c_{o_1}, p_1, c_{p_1}), \dots, (\tau_m, o_m, c_{o_m}, p_m, c_{p_m})$  with  $m \geq 2$  such that for each quintuple  $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$  in  $E$ :*

- (1)  $o_i$  and  $p_i$  are operations in  $\tau_i$ , and  $c_{o_i}, c_{p_i} \in \mathcal{S}$ ;
- (2)  $\mathbf{X}_i \not\approx_{\tau_i} \mathbf{Y}_i$  for each constraint  $\mathbf{X}_i \neq \mathbf{Y}_i$  in  $\tau_i$ ;
- (3)  $c_{o_i} = c_{p_i}$  if  $o_i \approx_{\tau_i} p_i$ ;
- (4)  $c_{o_i} \neq c_{p_i}$  if there is a constraint  $\mathbf{X}_i \neq \mathbf{Y}_i$  in  $\tau_i$  with  $\mathbf{X}_i \approx_{\tau_i} \text{var}(o_i)$  and  $\mathbf{Y}_i \approx_{\tau_i} \text{var}(p_i)$ ;
- (5) if  $i \neq 1$  and  $c_{q_i} = c_{q_1}$  for some  $q_i \in \{o_i, p_i\}$  and  $q_1 \in \{o_1, p_1\}$ , then there is no operation  $o'_i$  in  $\tau_i$  potentially ww-conflicting with an operation  $o'_1$  in  $\text{prefix}_{o_1}(\tau_1)$  with  $\text{var}(o'_i) \approx_{\tau_i} \text{var}(q_i)$  and  $\text{var}(o'_1) \approx_{\tau_1} \text{var}(q_1)$ .

Furthermore, for each pair of adjacent quintuples  $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$  and  $(\tau_j, o_j, c_{o_j}, p_j, c_{p_j})$  in  $E$  with  $j = i + 1$ , or  $i = m$  and  $j = 1$ :

- (6)  $o_i$  is potentially conflicting with  $p_j$  and  $c_{o_i} = c_{p_j}$ ;
- (7) if  $i = 1$  and  $j = 2$ , then  $o_1$  is potentially rw-conflicting with  $p_2$ ; and
- (8) if  $i = m$  and  $j = 1$ , then  $o_1 <_{\tau_1} p_1$  or  $o_m$  is potentially rw-conflicting with  $p_1$ .

The items have the following meaning: (2)  $\tau_i$  is satisfiable; (3) connected operations are assigned the same type mapping; (4) variables connected through an inequality are assigned a different type mapping; (5)  $\varphi_{\mathcal{S}}$  only assigns the same type mapping to  $o_1$  or  $p_1$  in  $\tau_1$  and  $o_i$  or  $p_i$  in  $\tau_i$  if it does not introduce a dirty write in the resulting multiversion split schedule (cf. Condition (1) in Definition 3.6); (6) each pair of variables in operations used for conflicts are assigned the same type mapping; (7, 8) the operations used for conflicts between  $\tau_1, \tau_2$

and  $\tau_m$  are restricted to satisfy respectively Condition (3) and (2) in Definition 3.6 in the resulting multiversion split schedule.

We first present an additional Lemma derived from Lemma 5.2 that will be used in the proof of Lemma 5.6.

**Lemma 5.7.** *Let  $D$  be a sequence of potentially conflicting quadruples. If  $\mathbf{X} \approx_D \mathbf{Y}$  and  $\mathbf{Y} \approx_D \mathbf{Z}$  then  $\mathbf{X} \approx_D \mathbf{Z}$  for every triple of variables  $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$  occurring in  $\text{Trans}(D)$ .*

*Proof.* According to Lemma 5.2,  $\mathbf{X} \approx_D \mathbf{Y}$  and  $\mathbf{Y} \approx_D \mathbf{Z}$  imply respectively  $\mathbf{X} \rightsquigarrow_D \mathbf{Y}$  and  $\mathbf{Y} \rightsquigarrow_D \mathbf{Z}$ . By definition,  $\mathbf{X} \rightsquigarrow_D \mathbf{Z}$  and hence  $\mathbf{X} \approx_D \mathbf{Z}$ .  $\square$

*Proof of Lemma 5.6. (if)* Let  $D = (\tau_1, o_1, p_1, \tau_1), \dots, (\tau_m, o_m, p_m, \tau_m)$  be the sequence of potentially conflicting quadruples derived from  $E$ . Notice in particular that  $D$  is indeed a sequence of potentially conflicting quadruples by (1) and (6). We construct a variable mapping  $\bar{\mu}$  for  $D$  admissible for a database  $\mathbf{D}$  such that the sequence of conflicting quadruples  $C = \bar{\mu}(D)$  satisfies the conditions in Definition 3.6, thereby proving that  $\mathcal{P}$  is not robust against RC.

Let  $\varphi_{\mathcal{S}} : \mathbf{Var} \rightarrow \mathcal{S}$  be the (partial) function assigning a type mapping in  $\mathcal{S}$  to each variable occurring in an operation in  $E$ :

$$\varphi_{\mathcal{S}}(\mathbf{X}) = \begin{cases} c_{o_i} & \text{if } \text{var}(o_i) = \mathbf{X} \text{ for some } (\tau_i, o_i, c_{o_i}, p_i, c_{p_i}) \in E, \\ c_{p_i} & \text{if } \text{var}(p_i) = \mathbf{X} \text{ for some } (\tau_i, o_i, c_{o_i}, p_i, c_{p_i}) \in E. \end{cases}$$

This function  $\varphi_{\mathcal{S}}$  is well defined: if there is a  $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i}) \in E$  with  $\text{var}(o_i) = \text{var}(p_i) = \mathbf{X}$ , then  $o_i \approx_{\tau_i} p_i$  and hence  $c_{o_i} = c_{p_i}$  by (3). Recall that we assume that templates in  $E$  are variable-disjoint. We argue that  $\varphi_{\mathcal{S}}(\mathbf{X}) = \varphi_{\mathcal{S}}(\mathbf{Y})$  if  $\mathbf{X} \approx_D \mathbf{Y}$  for each pair of variables  $\mathbf{X}$  and  $\mathbf{Y}$  for which  $\varphi_{\mathcal{S}}$  is defined. From Lemma 5.2, it follows that  $\mathbf{X} \rightsquigarrow_D \mathbf{Y}$  whenever  $\mathbf{X} \approx_D \mathbf{Y}$ . Let  $\tau_i$  and  $\tau_j$  be the template in which respectively  $\mathbf{X}$  and  $\mathbf{Y}$  occur. The argument is now by induction on the definition of  $\mathbf{X} \rightsquigarrow_D \mathbf{Y}$ :

- If  $i = j$  and  $\mathbf{X} \rightsquigarrow_{\tau_i} \mathbf{Y}$ , then  $\varphi_{\mathcal{S}}(\mathbf{X}) = \varphi_{\mathcal{S}}(\mathbf{Y})$  is immediate by (3);
- If  $(\tau_i, o_i, p_j, \tau_j) \in D$  with  $\text{var}(o_i) = \mathbf{X}$  and  $\text{var}(p_j) = \mathbf{Y}$  (respectively  $(\tau_j, o_j, p_i, \tau_i) \in D$  with  $\text{var}(o_j) = \mathbf{Y}$  and  $\text{var}(p_i) = \mathbf{X}$ ), then  $\varphi_{\mathcal{S}}(\mathbf{X}) = \varphi_{\mathcal{S}}(\mathbf{Y})$  is immediate by (6);
- Otherwise, if  $\mathbf{X} \rightsquigarrow_D \mathbf{Z}$  and  $\mathbf{Z} \rightsquigarrow_D \mathbf{Y}$  for some variable  $\mathbf{Z}$ , then by induction  $\varphi_{\mathcal{S}}(\mathbf{X}) = \varphi_{\mathcal{S}}(\mathbf{Z}) = \varphi_{\mathcal{S}}(\mathbf{Y})$ .

By Lemma 5.7,  $\approx_D$  is an equivalence relation. For  $\mathbf{X}$  occurring in  $\text{Trans}(D)$ , denote by  $[\mathbf{X}]$  the equivalence class of  $\mathbf{X}$ . Let  $\mathcal{S}'$  be obtained by extending  $\mathcal{S}$  with a type mapping  $c_{[\mathbf{X}]}$  for each equivalence class where no variable  $\mathbf{Y} \in [\mathbf{X}]$  is defined in  $\varphi_{\mathcal{S}}$ . Furthermore, each of the  $c_{[\mathbf{X}]}$  are picked such that all type mappings in  $\mathcal{S}'$  have disjoint ranges.

Next, we extend  $\varphi_{\mathcal{S}}$  to a function  $\varphi_{\mathcal{S}'} : \mathbf{Var} \rightarrow \mathcal{S}'$  assigning a type mapping to each variable  $\mathbf{X}$  occurring in  $\text{Trans}(D)$  as follows:

$$\varphi_{\mathcal{S}'}(\mathbf{X}) = \begin{cases} \varphi_{\mathcal{S}}(\mathbf{X}) & \text{if } \varphi_{\mathcal{S}} \text{ is defined for } \mathbf{X}, \\ \varphi_{\mathcal{S}}(\mathbf{Y}) & \text{if } \varphi_{\mathcal{S}} \text{ is defined for } \mathbf{Y} \text{ but not for } \mathbf{X} \text{ and } \mathbf{X} \approx_D \mathbf{Y}, \\ c_{[\mathbf{X}]} & \text{otherwise.} \end{cases}$$

Notice, furthermore, that in the second case  $\mathbf{X}$  might be connected in  $D$  to multiple variables for which  $\varphi_{\mathcal{S}}$  is defined, say  $\mathbf{Y}_1$  and  $\mathbf{Y}_2$ . Then, by Lemma 5.7,  $\mathbf{Y}_1 \approx_D \mathbf{Y}_2$  and hence  $\varphi_{\mathcal{S}}(\mathbf{Y}_1) = \varphi_{\mathcal{S}}(\mathbf{Y}_2)$ . We therefore conclude that  $\varphi_{\mathcal{S}'}(\mathbf{X})$  is well defined. We argue that  $\varphi_{\mathcal{S}'}(\mathbf{X}) = \varphi_{\mathcal{S}'}(\mathbf{Y})$  if  $\mathbf{X} \approx_D \mathbf{Y}$  for each pair of variables  $\mathbf{X}$  and  $\mathbf{Y}$ . If  $\varphi_{\mathcal{S}}$  is defined for both  $\mathbf{X}$  and  $\mathbf{Y}$ , then the result is immediate by  $\varphi_{\mathcal{S}}(\mathbf{X}) = \varphi_{\mathcal{S}}(\mathbf{Y})$ . If  $\varphi_{\mathcal{S}}$  is defined for one of these two variables, say

$\mathbf{X}$ , then  $\varphi_{S'}(\mathbf{X}) = \varphi_S(\mathbf{X}) = \varphi_{S'}(\mathbf{Y})$  by construction of  $\varphi_{S'}$ . If  $\varphi_S$  is not defined for both  $\mathbf{X}$  and  $\mathbf{Y}$ , then either there exists a variable  $\mathbf{Z}$  for which  $\varphi_S$  is defined and  $\mathbf{Z}$  is connected in  $D$  to one of these two variables, say  $\mathbf{X}$ , or no such variable  $\mathbf{Z}$  exists. In the former case,  $\mathbf{Z} \approx_D \mathbf{Y}$  follows from Lemma 5.7, implying  $\varphi_{S'}(\mathbf{X}) = \varphi_S(\mathbf{Z}) = \varphi_{S'}(\mathbf{Y})$ . In the latter case,  $\varphi_{S'}(\mathbf{X}) = c_{[\mathbf{X}]} = c_{[\mathbf{Y}]} = \varphi_{S'}(\mathbf{Y})$  by construction of  $\varphi_{S'}$ .

We now define the variable mapping  $\bar{\mu}$  from  $\varphi_{S'}$  as  $\bar{\mu}(\mathbf{X}) = c(\text{type}(\mathbf{X}))$  for each variable  $\mathbf{X}$ , where  $c = \varphi_S(\mathbf{X})$ . Next, we construct the database  $\mathbf{D}$ . For each template  $\tau_i$  and corresponding variable mapping  $\mu_i$  in  $\bar{\mu}$ , we add all tuples in  $\mu_i(\tau_i)$  to the database  $\mathbf{D}$ . Furthermore, for each constraint  $\mathbf{X} = f(\mathbf{Y})$  in  $\Gamma(\tau_i)$ , we have  $f^{\mathbf{D}}(\mu_i(\mathbf{X})) = \mu_i(\mathbf{Y})$  in  $\mathbf{D}$ . This is well defined for each function  $f^{\mathbf{D}}$ . Towards a contradiction, assume we have  $f^{\mathbf{D}}(\bar{\mu}(\mathbf{X}_i)) = \bar{\mu}(\mathbf{Y}_i)$  witnessed by a template  $\tau_i$  and  $f^{\mathbf{D}}(\bar{\mu}(\mathbf{X}_j)) = \bar{\mu}(\mathbf{Y}_j)$  witnessed by a template  $\tau_j$ , where  $\bar{\mu}(\mathbf{X}_i) = \bar{\mu}(\mathbf{X}_j)$ , but  $\bar{\mu}(\mathbf{Y}_i) \neq \bar{\mu}(\mathbf{Y}_j)$ . Since  $\mathbf{X}_i \approx_D \mathbf{Y}_i$ ,  $\mathbf{X}_j \approx_D \mathbf{Y}_j$  and  $\bar{\mu}(\mathbf{X}_i) = \bar{\mu}(\mathbf{X}_j)$ , we have  $\varphi_{S'}(\mathbf{Y}_i) = \varphi_{S'}(\mathbf{X}_i) = \varphi_{S'}(\mathbf{X}_j) = \varphi_{S'}(\mathbf{Y}_j)$ . Then,  $\bar{\mu}(\mathbf{Y}_i) = \bar{\mu}(\mathbf{Y}_j)$ , leading to a contradiction. In order to argue that  $\bar{\mu}$  is indeed admissible for  $\mathbf{D}$ , it remains to show that for each constraint  $\mathbf{X} \neq \mathbf{Y}$  in a template  $\tau_i$ , we have  $\bar{\mu}(\mathbf{X}) \neq \bar{\mu}(\mathbf{Y})$ . Again towards a contradiction, assume  $\bar{\mu}(\mathbf{X}) = \bar{\mu}(\mathbf{Y})$ , and let  $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$  be the corresponding quintuple in  $E$ . By definition of  $\bar{\mu}$ , we have  $\varphi_{S'}(\mathbf{X}) = \varphi_{S'}(\mathbf{Y})$ . It follows from  $\varphi_{S'}$  that either  $\mathbf{X} \approx_{\tau_i} \mathbf{Y}$ ; or  $\mathbf{X} \approx_{\tau_i} \text{var}(o_i)$  (respectively  $\mathbf{Y} \approx_{\tau_i} \text{var}(o_i)$ ),  $\mathbf{Y} \approx_{\tau_i} \text{var}(p_i)$  (respectively  $\mathbf{X} \approx_{\tau_i} \text{var}(p_i)$ ) and  $c_{o_i} = c_{p_i}$ . However, the former is contradicted by (2) and the latter by (4). We therefore conclude that  $\bar{\mu}$  is admissible for  $\mathbf{D}$ , as it satisfies all constraints.

It remains to argue that the sequence of conflicting quadruples  $C = \bar{\mu}(D)$  satisfies all conditions stated in Definition 3.6. The second and third condition are immediate by respectively (8) and (7). Towards a contradiction, assume the first condition holds. Then, there is an operation  $b'_1$  in  $\text{prefix}_{\bar{\mu}o_1}(\bar{\mu}(\tau_1))$  ww-conflicting with an operation  $b'_i$  in a transaction  $\bar{\mu}\tau_i$ . Let  $b'_1 = \bar{\mu}(o'_1)$  with  $\text{var}(o'_1) = \mathbf{X}_1$  and  $b'_i = \bar{\mu}(o'_i)$  with  $\text{var}(o'_i) = \mathbf{X}$ , and let  $(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$  and  $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$  be the corresponding quintuples in  $E$ . Note that  $o'_1$  is potentially ww-conflicting with  $o'_i$  and  $\bar{\mu}(\mathbf{X}_1) = \bar{\mu}(\mathbf{X}_i)$ . Then,  $\varphi_{A'}(\mathbf{X}_1) = \varphi_{A'}(\mathbf{X}_i)$ . By construction of  $\varphi_{A'}$ , this can only hold if  $\mathbf{X}_1 \approx_{\tau_1} \text{var}(q_1)$ ,  $\mathbf{X}_i \approx_{\tau_i} \text{var}(q_i)$  and  $c_{q_1} = c_{q_i}$  for some  $q_1 \in o_1, p_1$  and  $q_i \in o_i, p_i$ , thereby contradicting (5).

(only if) If  $\mathcal{P}$  is not robust against RC, then there exists a multiversion split schedule  $s$  based on a sequence of conflicting quadruples  $C$  over a set of transactions  $\mathcal{T}$  consistent with  $\mathcal{P}$  and a database  $\mathbf{D}$ . Let  $\bar{\mu}$  be the variable mapping for a sequence of potentially conflicting quadruples  $D = (\tau_1, o_1, p_1, \tau_2), \dots, (\tau_m, o_m, p_1, \tau_1)$  over  $\mathcal{P}$  with  $\bar{\mu}(C) = D$ , and let  $\mathcal{S}$  and  $\varphi_S$  be as in Lemma 5.5.

From this sequence  $D$  and function  $\varphi_S$ , we derive the sequence of quintuples  $E = (\tau_1, o_1, \varphi_S(\text{var}(o_1)), p_1, \varphi_S(\text{var}(p_1))), \dots, (\tau_m, o_m, \varphi_S(\text{var}(o_m)), p_m, \varphi_S(\text{var}(p_m)))$ . Let  $\varphi'_S = \{c_1, c_2, c_3, c_4\}$  be a set consisting of four type mappings with disjoint ranges. We adapt each quintuple in  $E$  in order, thereby creating a sequence  $E'$  satisfying the properties stated in Lemma 5.6. First, we add  $(\tau_1, o_1, c_1, p_1, c_k)$  to  $E'$ , where  $c_k = c_1$  if  $\varphi_S(\text{var}(o_1)) = \varphi_S(\text{var}(p_1))$ , and  $c_k = c_2$  otherwise. For each of the remaining quintuples in  $E$ , let  $(\tau_{i-1}, o_{i-1}, c_{o_{i-1}}, p_{i-1}, c_{p_{i-1}})$  be the quintuple previously added to  $E'$ . We then add  $(\tau_i, o_i, c_{o_i},$

$p_i, c_{p_i}$ ) to  $E'$  where  $c_{o_i} = c_{p_{i-1}}$  and

$$c_{p_i} = \begin{cases} c_{o_i} & \text{if } \varphi_{\mathcal{S}}(\text{var}(o_i)) = \varphi_{\mathcal{S}}(\text{var}(p_i)), \\ c_1 & \text{if } \varphi_{\mathcal{S}}(\text{var}(o_i)) = \varphi_{\mathcal{S}}(\text{var}(o_1)), \\ c_2 & \text{if } \varphi_{\mathcal{S}}(\text{var}(o_i)) = \varphi_{\mathcal{S}}(\text{var}(p_1)) \text{ and } \varphi_{\mathcal{S}}(\text{var}(o_1)) \neq \varphi_{\mathcal{S}}(\text{var}(p_1)), \\ c_3 & \text{if } \varphi_{\mathcal{S}}(\text{var}(o_i)) \neq \varphi_{\mathcal{S}}(\text{var}(p_i)) \text{ and } c_{p_i} \neq c_3, \\ c_4 & \text{otherwise.} \end{cases}$$

By construction, for every quintuple  $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$  in  $E'$  we now have

- $c_{o_i} = c_{p_i}$  iff  $\varphi_{\mathcal{S}}(\text{var}(o_i)) = \varphi_{\mathcal{S}}(\text{var}(p_i))$ ; and
- $c_{q_i} = c_{q_1}$  iff  $\varphi_{\mathcal{S}}(\text{var}(q_i)) = \varphi_{\mathcal{S}}(\text{var}(q_1))$  for every  $q_i \in \{o_i, p_i\}$  and  $q_1 \in \{o_1, p_1\}$ .

It remains to argue that  $E'$  indeed satisfies all required properties. (1) is trivial by construction. (2) If  $X_i \neq Y_i$  is a constraint in  $\tau_i$ , then  $\varphi_{\mathcal{S}}(X) \neq \varphi_{\mathcal{S}}(Y)$  and  $X \not\approx_D Y$  according to Lemma 5.5. (3) If  $o_i \approx_{\tau_i} p_i$ , then  $\varphi_{\mathcal{S}}(\text{var}(o_i)) = \varphi_{\mathcal{S}}(\text{var}(p_i))$  by Lemma 5.5, and hence  $c_{o_i} = c_{p_i}$ . (4) Assume there is a constraint  $X_i \neq Y_i$  in a template  $\tau_i$  with  $X_i \approx_{\tau_i} \text{var}(o_i)$  and  $Y_i \approx_{\tau_i} \text{var}(p_i)$ . By Lemma 5.5,  $\varphi_{\mathcal{S}}(X_i) = \varphi_{\mathcal{S}}(\text{var}(o_i)) \neq \varphi_{\mathcal{S}}(\text{var}(p_i)) = \varphi_{\mathcal{S}}(Y_i)$ , and therefore  $c_{o_i} \neq c_{p_i}$ . (5) Let  $c_{q_i} = c_{q_1}$  for some  $q_i \in \{o_i, p_i\}$  and  $q_1 \in \{o_1, p_1\}$ , with  $i \neq 1$ . Assume towards a contradiction that there is an operation  $o'_i$  in  $\tau_i$  potentially ww-conflicting with an operation  $o'_1$  in  $\text{prefix}_{o_1}(\tau_1)$  with  $\text{var}(o'_i) \approx_{\tau_i} \text{var}(q_i)$  and  $\text{var}(o'_1) \approx_{\tau_1} \text{var}(q_1)$ . But then  $\varphi_{\mathcal{S}}(\text{var}(o'_i)) = \varphi_{\mathcal{S}}(\text{var}(q_i)) = \varphi_{\mathcal{S}}(\text{var}(q_1)) = \varphi_{\mathcal{S}}(\text{var}(o'_1))$ , implying that  $\bar{\mu}(o'_i)$  is ww-conflicting with  $\bar{\mu}(o'_1)$ , contradicting the properties of  $C$  stated in Definition 3.6. (6) is again trivial by construction. (7) By Definition 3.6,  $\bar{\mu}(o_1)$  is rw-conflicting with  $\bar{\mu}(p_2)$  in  $C$ . Therefore,  $o_1$  is potentially rw-conflicting with  $p_2$ . (8) By Definition 3.6,  $\bar{\mu}(o_1) <_{\bar{\mu}(\tau_1)} \bar{\mu}(p_1)$  or  $\bar{\mu}(o_m)$  is rw-conflicting with  $\bar{\mu}(p_1)$  in  $C$ . As a result,  $o_1 <_{\tau_1} p_1$  or  $o_m$  is rw-conflicting with  $p_1$ .  $\square$

The characterization for T-ROBUSTNESS(MTBTemp,RC) in Lemma 5.6 implies an NLOGSPACE algorithm guessing the counterexample sequence  $E$ , thereby proving Theorem 5.1. Indeed, the algorithm guesses the sequence of quintuples  $E$ , verifying all conditions for each newly guessed quintuple while only requiring logarithmic space. Notice in particular that we only need to keep track of two other quintuples when verifying all conditions for the newly guessed quintuple, namely the first quintuple over  $\tau_1$  and the quintuple immediately preceding the newly guessed one. As usual, we can think of the encoding of templates and operations mentioned in each quintuple as pointers referring to the corresponding templates and operations on the input tape. Furthermore, we do not encode the four type mappings explicitly as such a representation of a mapping might require polynomial space. Since we are only interested in (dis)equality between type mappings, an encoding where these four type mappings are represented by four arbitrary strings of constant size suffices.

*Proof of Theorem 5.1.* The NLOGSPACE algorithm goes as follows: We start by guessing three initial quintuples, representing respectively the first, second and last quintuple of the possible sequence of quintuples as in Lemma 5.6. Consistent with previously used notation, we refer to these quintuples by  $E_1$ ,  $E_2$ , and  $E_m$ , with  $E_i = (\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$ . Note that the indices we use here are not part of the algorithm. They are only used to distinguish between the different considered quintuples in the proof argument.

We store all three quintuples using a logarithmic amount of space, by storing pointers to the respective transaction templates in  $\mathcal{P}$ , the positions of operations in the respective transaction templates, and the number 1, 2, 3 or 4 for the type mappings.




FIGURE 8. Acyclic schema graph for the TPC-C benchmark.

At this point, we verify that Condition (7) and (8) are true, that Conditions (1-5) are true for all chosen transaction templates and operations, and that Condition (6) is true for  $\tau_1$  and  $\tau_2$ , and  $\tau_2$  and  $\tau_m$ . We reject the guessed quintuples if any of the conditions is false.

If all previous checks are true, we proceed by inserting another step. Let  $i = 2$ . We guess a new quintuple  $E_{i+1}$  and verify that Condition (5) is true for  $\tau_i$  and  $\tau_{i+1}$  and that Conditions (1-6) are true for  $\tau_{i+1}$  and reject the entire construction if one of these conditions failed. Notice that all Conditions, including Condition (5) can be checked easily, particularly because quintuple  $E_1$  is stored. To proceed, we discard quintuple  $E_i$  and store  $E_{i+1}$  instead, thus without increasing the amount of space we use.

If  $\tau_{i+1}$  and  $\tau_m$  (from quintuple  $E_m$ ) have Condition (6), the algorithm emits an accept. Indeed, then the sequence  $E_1, \dots, E_i, E_{i+1}, E_m$  of guessed quintuples has all the properties of Lemma 5.6. Otherwise, the algorithm proceeds with another insertion step, for  $i = i + 1$ .  $\square$

## 6. ROBUSTNESS FOR TEMPLATES OVER ACYCLIC SCHEMAS

We denote by **AcycTemp** the class of all sets of transaction templates over acyclic schemas. As a concrete example, the schema graph for the TPC-C benchmark is given in Figure 8. Since this schema graph does not contain any cycles, the TPC-C benchmark is situated within **AcycTemp**. Notice in particular how this acyclic schema graph corresponds to the hierarchical structure of many-to-one relationships inherent to the schema for this benchmark. For example, every orderline belongs to exactly one order, and every order is related to exactly one customer, but the opposite is never true (i.e., a customer can be related to multiple orders, each of which can be related to multiple orderlines). In general, the results presented in this section can be applied to all workloads over schemas with such a hierarchical structure.

**Theorem 6.1.**  $\text{T-ROBUSTNESS}(\mathbf{AcycTemp}, RC)$  is decidable in EXPSpace.

We first provide some intuition for the proof. For a given acyclic schema graph  $SG$ ,  $R \xrightarrow{F}_{SG} S$  denotes the directed path from node  $R$  to node  $S$  in  $SG$  with  $F$  the sequence of edge labels on the path. The next lemma relates implication between variables to paths in  $SG$ .

**Lemma 6.2.** Let  $D$  be a sequence of potentially conflicting quadruples over a set of transaction templates  $\mathcal{P} \in \mathbf{AcycTemp}$ . For every pair of variables  $X, Y$  occurring in  $\text{Trans}(D)$ , if  $X \xrightarrow{F}_D Y$ , then  $\text{type}(X) \xrightarrow{F}_{SG} \text{type}(Y)$ , with  $SG$  the corresponding schema graph.

*Proof.* Let  $\tau_i$  and  $\tau_j$  be the templates in  $D$  in which  $X$  and  $Y$  occur, respectively. The proof is by induction on the definition of  $X \xrightarrow{F}_D Y$ .

(*Implication within the same template*) If  $i = j$  and  $X \xrightarrow{F}_{\tau_i} Y$ , then either  $F = \varepsilon$  and  $X = Y$ , or there is a variable  $Z$  such that  $Y = f(Z)$  is a constraint in  $\Gamma(\tau_i)$ ,  $X \xrightarrow{F'}_{\tau_i} Z$  and  $F = F' \cdot f$ . In the former case,  $\text{type}(X) = \text{type}(Y)$ , so  $\text{type}(X) \xrightarrow{F}_{SG} \text{type}(Y)$  is immediate. In



the latter case, it follows by induction that  $\text{type}(\mathbf{X}) \xrightarrow{F}_{SG} \text{type}(\mathbf{Z})$ . Since  $\text{dom}(f) = \text{type}(\mathbf{Z})$  and  $\text{range}(f) = \text{type}(\mathbf{Y})$ , it follows by definition that  $\text{type}(\mathbf{Z}) \xrightarrow{f}_{SG} \text{type}(\mathbf{Y})$  and furthermore  $\text{type}(\mathbf{X}) \xrightarrow{f}_{SG} \text{type}(\mathbf{Z})$  holds.

(*Implication between templates*) If  $F = \varepsilon$  and  $(\tau_i, o_i, p_j, \tau_j)$  (respectively  $(\tau_j, o_j, p_i, \tau_i)$ ) is a potentially conflicting quadruple in  $D$ , with  $\text{var}(o_i) = \mathbf{X}$  and  $\text{var}(p_j) = \mathbf{Y}$  (respectively  $\text{var}(p_i) = \mathbf{X}$  and  $\text{var}(o_j) = \mathbf{Y}$ ), then  $\text{type}(\mathbf{X}) = \text{type}(\mathbf{Y})$  by definition of potentially conflicting operations. So,  $\text{type}(\mathbf{X}) \xrightarrow{F}_{SG} \text{type}(\mathbf{Y})$  is again immediate.

(*Inductive case*) If  $\mathbf{X} \xrightarrow{F_1}_D \mathbf{Z}$  and  $\mathbf{Z} \xrightarrow{F_2}_D \mathbf{Y}$  for some variable  $\mathbf{Z}$  with  $F = F_1 \cdot F_2$ , then  $\text{type}(\mathbf{X}) \xrightarrow{F_1}_{SG} \text{type}(\mathbf{Z})$  and  $\text{type}(\mathbf{Z}) \xrightarrow{F_2}_{SG} \text{type}(\mathbf{Y})$  follow by induction. We conclude that  $\text{type}(\mathbf{X}) \xrightarrow{F}_{SG} \text{type}(\mathbf{Y})$ .  $\square$

Notice that an assignment of a tuple to a variable  $\mathbf{X}$  determines the tuples assigned to all variables  $\mathbf{Y}$  with  $\mathbf{X} \xrightarrow{F}_D \mathbf{Y}$  for some sequence of function names  $F$ . From Lemma 6.2 it follows that each such implied tuple is witnessed by a path in the corresponding schema graph  $SG$ . Therefore, the maximal number of different tuples implied by  $\mathbf{X}$  corresponds to the number of paths in  $SG$  starting in  $\text{type}(\mathbf{X})$ , which is finite when  $SG$  is acyclic. Because there can be multiple paths between nodes in the schema graph, it is no longer the case as in the previous section that variables of the same type connected in  $D$  must be assigned the same value. So, instead of using type mappings, we introduce *tuple-contexts* to represent the sets of all tuples implied by the assignment of a given variable. Formally, a *tuple-context for a type*  $R \in \text{Rels}$  is a function from paths with source  $R$  in  $SG(\text{Rels}, \text{Funcs})$  to tuples in **Tuples** of the appropriate type. That is, for each tuple-context  $c$  for type  $R$  and for each path  $R \xrightarrow{F}_{SG} S$  in  $SG$ ,  $\text{type}(c(R \xrightarrow{F}_{SG} S)) = S$ .

Similar to Lemma 5.5, we show that we can represent a counterexample schedule based on  $D$  by assigning a tuple-context to each variable in  $\text{Trans}(D)$ , taking special care when assigning contexts to variables connected in  $D$  to make sure that they are properly related to each other. For this, we introduce a (partial) function  $\varphi_{\mathcal{A}} : \mathbf{Var} \rightarrow \mathcal{A}$  mapping (a subset of) variables in  $\text{Trans}(D)$  to tuple-contexts in  $\mathcal{A}$  (for  $\mathcal{A}$  a set of tuple-contexts) and refer to it as a (*partial*) *context assignment for  $D$  over  $\mathcal{A}$* . In a sequence of lemmas, we show that  $\varphi_{\mathcal{A}}$  can always be expanded into a total function and an approach based on enumeration of quintuples analogous to Lemma 5.6 suffices to decide robustness. A major difference with the previous section is that there is no longer a constant bound on the number of tuple-contexts that are needed and consistency between tuple-contexts in connected variables needs to be maintained.

We call node  $S$  a *descendant* of node  $R$  and  $R$  an *ancestor* of  $S$  in an acyclic schema graph  $SG$ . We write  $R \xrightarrow{F}_{SG} S$ , with  $\varepsilon$  denoting the empty labeling, for the case  $R = S$ . This means that a node is a descendant and ancestor of itself. When  $F$  is not relevant, we simply write  $R \rightsquigarrow_{SG} S$ .

Let  $c_R$  and  $c_S$  be two tuple-contexts for types  $R$  and  $S$ , respectively, such that  $S$  is a descendant of  $R$  in  $SG$ , witnessed by the path  $R \xrightarrow{F}_{SG} S$  in  $SG$ . We then say that  $c_S$  is a *tuple-subcontext of  $c_R$  witnessed by  $F$*  if  $c_S(S \xrightarrow{F'}_{SG} S') = c_R(R \xrightarrow{F \cdot F'}_{SG} S')$  for every path  $S \xrightarrow{F'}_{SG} S'$  in  $SG$ . It should be noted that  $R \xrightarrow{F \cdot F'}_{SG} S'$  is indeed a valid path in  $SG$ , as it concatenates the paths  $R \xrightarrow{F}_{SG} S$  and  $S \xrightarrow{F'}_{SG} S'$ . For a given tuple-context  $c$  for a type  $R$  in the schema graph  $SG$ , we will often write  $c(F)$  as a shorthand notation for  $c(R \xrightarrow{F}_{SG} S)$ .

Similar to Lemma 5.5 for sets of transaction templates admitting multi-tree bijectivity, Lemma 6.6 shows that we can represent a counterexample schedule based on a sequence of potentially conflicting quadruples  $D$  over an acyclic schema by assigning a tuple-context to

each variable in  $\text{Trans}(D)$ , taking special care when assigning contexts to variables connected in  $D$  to make sure that they are properly related to each other. For a set of tuple-contexts  $\mathcal{A}$ , we refer to a (partial) function  $\varphi_{\mathcal{A}} : \mathbf{Var} \rightarrow \mathcal{A}$  mapping (a subset of) variables in  $\text{Trans}(D)$  to tuple-contexts in  $\mathcal{A}$  as a (partial) *context assignment for  $D$  over  $\mathcal{A}$* . We furthermore say that  $\varphi_{\mathcal{A}}$  is a *total context assignment for  $D$  over  $\mathcal{A}$*  if  $\varphi_{\mathcal{A}}$  is defined for every variable in  $\text{Trans}(D)$ .

Two variables  $X$  and  $Y$  occurring in  $\text{Trans}(D)$  are *equivalent in  $D$* , denoted  $X \equiv_D Y$  if

- $X = Y$ ;
- there exists a pair of variables  $Z$  and  $W$  and a sequence of function names  $F$  with  $Z \equiv_D W$ ,  $Z \xrightarrow{F}_D X$  and  $W \xrightarrow{F}_D Y$ ; or
- there exists a variable  $Z$  with  $X \equiv_D Z$  and  $Z \equiv_D Y$ .

Similarly, two variables  $X$  and  $Y$  occurring in a transaction template  $\tau$  are *equivalent in  $\tau$* , denoted  $X \equiv_{\tau} Y$  if

- $X = Y$ ;
- there exists a pair of variables  $Z$  and  $W$  in  $\tau$  and a sequence of function names  $F$  with  $Z \equiv_{\tau} W$ ,  $Z \xrightarrow{F}_{\tau} X$  and  $W \xrightarrow{F}_{\tau} Y$ ; or
- there exists a variable  $Z$  with  $X \equiv_{\tau} Z$  and  $Y \equiv_{\tau} Z$ .

Intuitively, every variable mapping admissible for a given database will assign the same tuple to equivalent variables (see Lemma 6.4). Due to these equivalent variables, the assignment of a tuple to a variable  $X$  for a given database might imply the tuple assigned to a variable  $Y$ , even if  $X \rightsquigarrow_D Y$  does not hold. We capture this observation by introducing *variable determination*, which is stronger than the previously defined variable implication. Formally, a variable  $X$  *determines* a variable  $Y$  in  $D$  witnessed by a sequence of function names  $F$ , denoted  $X \xrightarrow{F}_D Y$  if:

- $X \xrightarrow{F}_D Y$ ;
- $F = \varepsilon$  and  $X \equiv_D Y$ ; or
- there exists a variable  $Z$  with  $X \xrightarrow{F}_D Z$ ,  $Z \xrightarrow{F}_D Y$  and  $F = F_1 \cdot F_2$ .

For two variables  $X$  and  $Y$  in a template  $\tau \in \text{Trans}(D)$  we furthermore say that  $X$  *determines*  $Y$  in  $\tau$  witnessed by a sequence of function names  $F$ , denoted  $X \xrightarrow{F}_{\tau} Y$  if:

- $X \xrightarrow{F}_{\tau} Y$ ;
- $F = \varepsilon$  and  $X \equiv_{\tau} Y$ ; or
- there exists a variable  $Z$  with  $X \xrightarrow{F}_{\tau} Z$ ,  $Z \xrightarrow{F}_{\tau} Y$  and  $F = F_1 \cdot F_2$ .

**Lemma 6.3.** *For a multiversion split schedule  $s$  based on a sequence of conflicting quadruples  $C$  over a set of transactions  $\mathcal{T}$  consistent with a set of transaction templates  $\mathcal{P}$  and a database  $D$ , let  $\bar{\mu}$  be the variable mapping for a sequence of potentially conflicting quadruples  $D$  over  $\mathcal{P}$  with  $\bar{\mu}(D) = C$ . Then, for every combination of variables  $W, X, Y, Z$  occurring in  $\text{Trans}(D)$ , if  $Z \xrightarrow{F}_D X$ ,  $W \xrightarrow{F}_D Y$  and  $\bar{\mu}(Z) = \bar{\mu}(W)$ , then  $\bar{\mu}(X) = \bar{\mu}(Y)$ .*

*Proof.* By definition of  $Z \rightsquigarrow_D X$ , there is a sequence of variables  $X_1, X_2, \dots, X_n$  with  $X_1 = Z$  and  $X_n = X$  such that for each pair of adjacent variables  $X_i$  and  $X_{i+1}$ :

- (†)  $X_i$  and  $X_{i+1}$  both occur in the same template  $\tau \in \text{Trans}(D)$  and  $X_{i+1} = f(X_i) \in \Gamma(\tau)$  for some function name  $f$ ; or
- (‡)  $\text{type}(X_i) = \text{type}(X_{i+1})$  and there is a potentially conflicting quadruple  $(\tau_j, o_j, p_k, \tau_k)$  in  $D$  with either  $\text{var}(o_j) = X_i$  and  $\text{var}(p_k) = X_{i+1}$  or  $\text{var}(p_k) = X_i$  and  $\text{var}(o_j) = X_{i+1}$ .

Furthermore, the sequence  $F$  corresponds to the function names used in applications of (†). Analogously,  $W \rightsquigarrow_D Y$ , implies a sequence of variables  $Y_1, Y_2, \dots, Y_n$  with  $Y_1 = W$  and  $Y_m = Y$

with the same properties. Notice that the lengths of these two sequences of variables, namely  $n$  and  $m$ , are not necessarily equal to each other and to the length of  $F$  due to possible applications of  $(\ddagger)$ . For a variable  $X_i$  in the sequence  $X_1, X_2, \dots, X_n$ , we denote the sequence of function names derived from applications of  $(\dagger)$  in the subsequence  $X_i, \dots, X_n$  by  $\text{suffix}_F(X_i)$ . Notice that  $\text{suffix}_F(X_i)$  is indeed always a suffix of  $F$ , and that  $\text{suffix}_F(X_1) = \text{suffix}_F(Y_1) = F$  and  $\text{suffix}_F(X_n) = \text{suffix}_F(Y_n) = \varepsilon$ .

We argue by induction that for every  $i \in [1, n]$  and  $j \in [1, m]$ , if  $\text{suffix}_F(X_i) = \text{suffix}_F(Y_j)$  then  $\bar{\mu}(X_i) = \bar{\mu}(Y_j)$ . This then implies  $\bar{\mu}(X) = \bar{\mu}(X_n) = \bar{\mu}(Y_m) = \bar{\mu}(Y)$ . (*base case*) Note that  $\bar{\mu}(X_1) = \bar{\mu}(Z) = \bar{\mu}(W) = \bar{\mu}(Y_1)$  and  $\text{suffix}_F(X_1) = F = \text{suffix}_F(Y_1)$ . (*inductive case*) Let  $\text{suffix}_F(X_{i+1}) = \text{suffix}_F(Y_{j+1})$ . Then, we distinguish the following cases:

- $\text{suffix}_F(X_i) = \text{suffix}_F(X_{i+1})$ : This means that  $(\ddagger)$  applies to  $X_i$  and  $X_{i+1}$ , and there is a potentially conflicting quadruple  $(\tau_k, o_k, p_\ell, \tau_\ell)$  in  $D$  with either  $\text{var}(o_k) = X_i$  and  $\text{var}(p_\ell) = X_{i+1}$  or  $\text{var}(p_\ell) = X_i$  and  $\text{var}(o_k) = X_{i+1}$ . By definition of  $\bar{\mu}$ , we have  $\bar{\mu}(X_{i+1}) = \bar{\mu}(X_i)$  and by induction that  $\bar{\mu}(X_{i+1}) = \bar{\mu}(Y_{j+1})$  implying that  $\bar{\mu}(X_{i+1}) = \bar{\mu}(X_{j+1})$ .
- $\text{suffix}_F(Y_j) = \text{suffix}_F(Y_{j+1})$ : similar as previous argument;
- $\text{suffix}_F(X_i) \neq \text{suffix}_F(X_{i+1})$  and  $\text{suffix}_F(Y_j) \neq \text{suffix}_F(Y_{j+1})$ : Then,  $(\dagger)$  applies to both  $X_i$  and  $X_{i+1}$ , and  $Y_j$  and  $Y_{j+1}$ . Furthermore,  $\text{suffix}_F(X_i) = \text{suffix}_F(Y_j) = f \cdot F'$  for some  $f$  and  $F'$ . By induction,  $\bar{\mu}(X_i) = \bar{\mu}(Y_j)$ . Then,  $X_{i+1} = f(X_i)$  is a constraint in some template  $\tau_k \in \text{Trans}(D)$  and  $Y_{j+1} = f(Y_j)$  is a constraint in some template  $\tau_\ell \in \text{Trans}(D)$ . Since  $\bar{\mu}$  is admissible for database  $\mathbf{D}$  and  $\bar{\mu}(X_i) = \bar{\mu}(Y_j)$ , it follows that  $\bar{\mu}(X_{i+1}) = f^{\mathbf{D}}(\bar{\mu}(X_i)) = f^{\mathbf{D}}(\bar{\mu}(Y_j)) = \bar{\mu}(Y_{j+1})$ .  $\square$

The next Lemma shows that every variable mapping admissible for a given database will assign the same tuple to equivalent variables.

**Lemma 6.4.** *For a multiversion split schedule  $s$  based on a sequence of conflicting quadruples  $C$  over a set of transactions  $\mathcal{T}$  consistent with a set of transaction templates  $\mathcal{P}$  and a database  $\mathbf{D}$ , let  $\bar{\mu}$  be the variable mapping for a sequence of potentially conflicting quadruples  $D$  over  $\mathcal{P}$  with  $\bar{\mu}(D) = C$ . Then, for every pair of variables  $X$  and  $Y$  occurring in templates  $\tau_i$  and  $\tau_j$  in  $\text{Trans}(D)$  respectively, if  $X \equiv_D Y$ , then  $\bar{\mu}(X) = \bar{\mu}(Y)$ .*

*Proof.* The proof is by induction on the definition of  $X \equiv_D Y$ . (*base case*) If  $X = Y$ , then the result is immediate. (*inductive cases*) If there are two variables  $Z$  and  $W$  and a sequence of function names  $F$  such that  $Z \equiv_D W$ ,  $Z \xrightarrow{F}_D X$  and  $W \xrightarrow{F}_D Y$ , then by induction we have  $\bar{\mu}(Z) = \bar{\mu}(W)$ . The proof that  $\bar{\mu}(X) = \bar{\mu}(Y)$  is now immediate by application of Lemma 6.3. If instead there is a variable  $Z$  with  $X \equiv_D Z$  and  $Y \equiv_D Z$ , then we can argue by induction that  $\bar{\mu}(X) = \bar{\mu}(Z)$  and  $\bar{\mu}(Y) = \bar{\mu}(Z)$ , and hence  $\bar{\mu}(X) = \bar{\mu}(Y)$ .  $\square$

**Definition 6.5.** Let  $D$  be a sequence of potentially conflicting quadruples,  $\mathcal{A}$  a set of tuple-contexts and  $\varphi_{\mathcal{A}}$  a partial context assignment for  $D$  over  $\mathcal{A}$ . We say that  $\varphi_{\mathcal{A}}$  respects the constraints of  $D$  if, for every two (not necessarily different) variables  $X$  and  $Y$  occurring in  $D$  that  $\varphi_{\mathcal{A}}$  is defined for, the following conditions are true, where  $c_X = \varphi_{\mathcal{A}}(X)$  and  $c_Y = \varphi_{\mathcal{A}}(Y)$ :

- (1)  $c_X$  is a tuple-context for  $\text{type}(X)$ ;
- (2) for every  $X \xrightarrow{F}_D Z$  and  $Y \xrightarrow{F}_D Z$ ,  $c_X(F_1) = c_Y(F_2)$ ;
- (3) for every  $X \xrightarrow{F}_D W$  and  $Y \xrightarrow{F}_D Z$  with  $W \neq Z$  a constraint in a template  $\tau$ ,  $c_X(F_1) \neq c_Y(F_2)$ ;
- (4) for every  $X \xrightarrow{F}_D W$  and  $Y \xrightarrow{F}_D Z$ , if  $c_X(F_1) = c_Y(F_2)$ , then there is no constraint  $W \neq Z$  in a template  $\tau \in \text{Trans}(D)$ ;
- (5) if  $X \xrightarrow{F}_D Y$ , then  $c_Y$  is a tuple-subcontext of  $c_X$  witnessed by  $F$ ; and

- (6) for every pair of tuple-subcontexts  $c'_X$  and  $c'_Y$  of  $c_X$  and  $c_Y$  witnessed by respectively  $F_X$  and  $F_Y$ , if  $c_X(F_X) = c_Y(F_Y)$ , then  $c'_X = c'_Y$ .

The next Lemma shows that we can represent a counterexample schedule based on a sequence of potentially conflicting quadruples  $D$  over an acyclic schema by assigning a tuple-context to each variable in  $\text{Trans}(D)$ .

**Lemma 6.6.** *For a multiversion split schedule  $s$  based on a sequence of conflicting quadruples  $C$  over a set of transactions  $\mathcal{T}$  consistent with a set of transaction templates  $\mathcal{P} \in \mathbf{AcycTemp}$  and a database  $\mathbf{D}$ , let  $\bar{\mu}$  be the variable mapping for a sequence of potentially conflicting quadruples  $D$  over  $\mathcal{P}$  with  $\bar{\mu}(D) = C$ . Then a set  $\mathcal{A}$  of tuple-contexts and a total context assignment  $\varphi_{\mathcal{A}}$  for  $D$  over  $\mathcal{A}$  exist with:*

- $\varphi_{\mathcal{A}}$  respects the constraints of  $D$ ; and
- $\bar{\mu}(X) = c_X(\varepsilon)$  for every variable  $X$ , with  $c_X = \varphi_{\mathcal{A}}(X)$ .

*Proof.* We first assign a tuple-context to each tuple in database  $\mathbf{D}$ , based on the functions in  $\mathbf{D}$ . Let  $(\text{Rels}, \text{Funcs})$  be the schema over which  $\mathcal{P}$  is defined. Since the schema graph  $SG(\text{Rels}, \text{Funcs})$  is acyclic, a total order  $<_{SG}$  over  $\text{Rels}$  exists such that there is no path from type  $R$  to type  $S$  in  $SG$  if  $R <_{SG} S$ . We now assign tuple-contexts to tuples based on the order implied by  $<_{SG}$ . That is, we first consider all tuples of the type that is ordered first by  $<_{SG}$ , then all tuples of the type that is ordered second, etc. If there are multiple tuples of the same type, the relative order in which we handle them is not important. For each tuple  $\mathbf{t}$ , we construct a tuple-context  $c_{\mathbf{t}}$  with  $c_{\mathbf{t}}(\varepsilon) = \mathbf{t}$ , and for each path  $F = f \cdot F'$  in  $SG$  starting in  $\text{type}(\mathbf{t})$ , set  $c_{\mathbf{t}}(F) = c_{\mathbf{v}}(F')$ , with  $\mathbf{v} = f^{\mathbf{D}}(\mathbf{t})$ . Notice that  $c_{\mathbf{v}}$  is already defined for  $\mathbf{v}$ , as there is a path from  $\text{type}(\mathbf{t})$  to  $\text{type}(\mathbf{v})$  in  $SG$  and, hence,  $\text{type}(\mathbf{v}) <_{SG} \text{type}(\mathbf{t})$ . By construction,  $c_{\mathbf{v}}$  is a tuple-subcontext of  $c_{\mathbf{t}}$  witnessed by  $f$ .

Next, we construct  $\varphi_{\mathcal{A}}$  as follows:  $\varphi_{\mathcal{A}}(X) = c_{\mathbf{t}}$  with  $\bar{\mu}(X) = \mathbf{t}$  for every variable  $X$  occurring in  $\text{Trans}(D)$ . We argue by induction on the definition of  $\xrightarrow{D}$  that

$$c_{\mathbf{t}}(F) = \bar{\mu}(Y) \text{ for every } X \xrightarrow{D} Y \text{ (with } c_{\mathbf{t}} = \varphi_{\mathcal{A}}(X)). \quad (\dagger)$$

If  $X \xrightarrow{D} Y$ , then by construction of  $\varphi_{\mathcal{A}}$  and since  $\bar{\mu}$  is admissible for  $\mathbf{D}$ , we have  $c_{\mathbf{t}}(F) = \bar{\mu}(Y)$ . If  $F = \varepsilon$  and  $X \equiv_D Y$ , then  $c_{\mathbf{t}}(\varepsilon) = \bar{\mu}(X) = \bar{\mu}(Y)$  by Lemma 6.4. Otherwise, if there exists a variable  $Z$  with  $X \xrightarrow{D} Z$ ,  $Z \xrightarrow{D} Y$  and  $F = F_1 \cdot F_2$ , then by induction  $c_{\mathbf{t}}(F_1) = \bar{\mu}(Z) = \mathbf{v}$  and  $c_{\mathbf{v}}(F_2) = \bar{\mu}(Y)$ , with  $\varphi_{\mathcal{A}}(Z) = c_{\mathbf{v}}$ . By construction of  $c_{\mathbf{t}}$  and  $c_{\mathbf{v}}$ , the desired  $c_{\mathbf{t}}(F) = \bar{\mu}(Y)$  now follows.

It remains to verify that  $\varphi_{\mathcal{A}}$  indeed satisfies all required properties. By construction,  $\bar{\mu}(X) = c_{\mathbf{t}}(\varepsilon)$  with  $c_{\mathbf{t}} = \varphi_{\mathcal{A}}(X)$ , so we only need to show that  $\varphi_{\mathcal{A}}$  respects the constraints of  $D$  by verifying all properties in Definition 6.5. To this end, let  $X$  and  $Y$  be two variables occurring in  $\text{Trans}(D)$ , and let  $c_{\mathbf{t}} = \varphi_{\mathcal{A}}(X)$  and  $c_{\mathbf{v}} = \varphi_{\mathcal{A}}(Y)$ . (1) By construction,  $c_{\mathbf{t}}$  is a tuple-context for  $\text{type}(X)$ . (2)  $c_{\mathbf{t}}(F_1) = \bar{\mu}(Z) = c_{\mathbf{v}}(F_2)$  by  $(\dagger)$ . (3) If  $W \neq Z$  is a constraint, then  $\bar{\mu}(W) \neq \bar{\mu}(Z)$  as  $\bar{\mu}$  is admissible for  $\mathbf{D}$ . By  $(\dagger)$ , it follows that  $c_{\mathbf{t}}(F_1) = \bar{\mu}(W) \neq \bar{\mu}(Z) = c_{\mathbf{v}}(F_2)$ . (4) If  $c_{\mathbf{t}}(F_1) = c_{\mathbf{v}}(F_2)$ , then  $\bar{\mu}(W) = c_{\mathbf{t}}(F_1) = c_{\mathbf{v}}(F_2) = \bar{\mu}(Z)$  by  $(\dagger)$ . Since  $\bar{\mu}$  is admissible for  $\mathbf{D}$ , there cannot be a constraint  $W \neq Z$ . (5) If  $X \xrightarrow{D} Y$ , then by construction of the tuple-contexts  $c_{\mathbf{t}}$  and  $c_{\mathbf{v}}$  it follows that  $c_{\mathbf{v}}$  is a tuple-subcontext of  $c_{\mathbf{t}}$  witnessed by  $F$ . (6) Let  $c_{\mathbf{t}'}$  and  $c_{\mathbf{v}'}$  be tuple-subcontexts of  $c_{\mathbf{t}}$  and  $c_{\mathbf{v}}$  witnessed by respectively  $F_X$  and  $F_Y$ . If  $c_{\mathbf{t}}(F_X) = c_{\mathbf{v}}(F_Y) = \mathbf{q}$  for some tuple  $\mathbf{q}$ , then by construction of  $c_{\mathbf{t}}$  and  $c_{\mathbf{v}}$  we have  $c_{\mathbf{t}'} = c_{\mathbf{v}'} = c_{\mathbf{q}}$ , with  $c_{\mathbf{q}}$  the tuple-context assigned to this tuple  $\mathbf{q}$ .  $\square$

From  $D = (\tau_1, o_1, p_2, \tau_2), \dots, (\tau_m, o_m, p_1, \tau_1)$  and  $\varphi_{\mathcal{A}}$  as in Lemma 6.6 we can derive a sequence of quintuples  $E = (\tau_1, o_1, c_{o_1}, p_1, c_{p_1}), \dots, (\tau_m, o_m, c_{o_m}, p_m, c_{p_m})$  such that  $c_{o_i} = \varphi_{\mathcal{A}}(\mathbf{X}_i)$  (respectively  $c_{p_i} = \varphi_{\mathcal{A}}(\mathbf{Y}_i)$ ) for  $i \in [1, m]$  with  $o_i$  (respectively  $p_i$ ) an operation over variable  $\mathbf{X}_i$  (respectively  $\mathbf{Y}_i$ ). Intuitively, this sequence of quintuples can be used to reconstruct the original multiversion split schedule  $s$ . To this end, notice that we can derive the original sequence of potentially conflicting quadruples  $D$  and a partial context assignment  $\varphi'_{\mathcal{A}}$  from  $E$  that is defined for each variable  $\mathbf{X}_i$  occurring in either an operation  $o_i$  or  $p_i$  in  $\tau_i$ . We first show that we can extend this partial context assignment  $\varphi'_{\mathcal{A}}$  to a total context assignment respecting the constraints in  $D$  (Lemma 6.7), and then prove that such a total context assignment respecting the constraints in  $D$  implies a variable assignment  $\bar{\mu}$  such that the  $C = \bar{\mu}(D)$  is a valid sequence of conflicting quadruples (Lemma 6.8).

**Lemma 6.7.** *Let  $D = (\tau_1, o_1, p_2, \tau_2), \dots, (\tau_m, o_m, p_1, \tau_1)$  be a sequence of potentially conflicting quadruples over a set of transaction templates  $\mathcal{P} \in \mathbf{AcycTemp}$  and  $\varphi_{\mathcal{A}}$  a partial context assignment defined for every variable  $\mathbf{X}_i$  of  $o_i$  and  $\mathbf{Y}_i$  of  $p_i$  in every  $\tau_i$ . If*

- $\varphi_{\mathcal{A}}$  respects the constraints of  $D$ ; and
- for every pair of variables  $\mathbf{X}$  and  $\mathbf{Y}$  in a template  $\tau_i$  with  $\mathbf{X} \equiv_{\tau_i} \mathbf{Y}$ , there is no constraint  $\mathbf{X} \neq \mathbf{Y}$  in  $\tau_i$ ;

then we can extend  $\varphi_{\mathcal{A}}$  to a total context assignment  $\varphi'_{\mathcal{A}}$  for  $D$  respecting the constraints of  $D$ .

*Proof.* By definition of equivalence,  $\equiv_D$  partitions all variables occurring in  $\text{Trans}(D)$  in equivalence classes. That is, two variables  $\mathbf{X}$  and  $\mathbf{Y}$  are in the same equivalence class iff  $\mathbf{X} \equiv_D \mathbf{Y}$ . For a given variable  $\mathbf{X}$ , we denote the equivalence class  $\mathbf{X}$  belongs to by  $[\mathbf{X}]$ . Note that for any pair of variables  $\mathbf{X}$  and  $\mathbf{Y}$  occurring in  $\text{Trans}(D)$ , if  $\mathbf{X} \xrightarrow{F}_D \mathbf{Y}$ , then  $\mathbf{X}' \xrightarrow{F}_D \mathbf{Y}'$  for any pair of variables  $\mathbf{X}' \in [\mathbf{X}]$  and  $\mathbf{Y}' \in [\mathbf{Y}]$ . By slight abuse of notation, we use  $\mathbf{X} \xrightarrow{F}_D [\mathbf{Y}]$  and  $[\mathbf{X}] \xrightarrow{F}_D \mathbf{Y}$  to denote that  $\mathbf{X} \xrightarrow{F}_D \mathbf{Y}'$  for every  $\mathbf{Y}' \in [\mathbf{Y}]$  and  $\mathbf{X}' \xrightarrow{F}_D \mathbf{Y}$  for every  $\mathbf{X}' \in [\mathbf{X}]$ , respectively.

Let  $(\text{Rels}, \text{Funcs})$  be the schema over which  $\mathcal{P}$  is defined. Since the schema graph  $SG(\text{Rels}, \text{Funcs})$  is acyclic, a total order  $<_{SG}$  over  $\text{Rels}$  exists such that there is no path from type  $R$  to type  $S$  in  $SG$  if  $R <_{SG} S$ . We now define  $\varphi'_{\mathcal{A}}$  for variables in  $\text{Trans}(D)$  according to the order implied by  $<_{SG}$ . If there are multiple variables of the same type, the relative order in which we handle them is not important.

The proof is as follows. Assume  $\varphi_{\mathcal{A}}$  respects the constraints of  $D$  and is at least defined for every variable  $\mathbf{X}_i$  of  $o_i$  and  $\mathbf{Y}_i$  of  $p_i$  in every  $\tau_i$ . We extend  $\varphi_{\mathcal{A}}$  towards  $\varphi'_{\mathcal{A}}$  by defining  $\varphi'_{\mathcal{A}}$  for the whole equivalence class  $[\mathbf{X}]$  of the first (according to  $<_{SG}$ ) variable  $\mathbf{X}$  for which  $\varphi_{\mathcal{A}}$  is not defined. The precise construction is by case. In the first case, the tuple-context that should be assigned to variables in  $[\mathbf{X}]$  is already implied, as it is the tuple-subcontext of an existing tuple-context. In the second case, we construct a fresh tuple-context, including existing tuple-contexts as tuple-subcontexts where we need to make sure that  $\varphi'_{\mathcal{A}}$  respects the constraints in  $D$ . In each case, we then argue that  $\varphi'_{\mathcal{A}}$  still respects the constraints in  $D$ . By repeating this argument, we can extend the context assignment to a total context assignment defined for all variables occurring in  $\text{Trans}(D)$ .

(Case 1) If a variable  $\mathbf{Y}$  exists with  $\varphi_{\mathcal{A}}$  defined for  $\mathbf{Y}$  and  $\mathbf{Y} \xrightarrow{F}_D [\mathbf{X}]$ , then  $\varphi'_{\mathcal{A}}(\mathbf{X}') = c_{\mathbf{X}'}$  for every variable  $\mathbf{X}' \in [\mathbf{X}]$ , with  $c_{\mathbf{X}'}$  the tuple-subcontext of  $c_{\mathbf{Y}} = \varphi_{\mathcal{A}}(\mathbf{Y})$  witnessed by  $F$ . Notice that this is well defined, even if there are multiple such  $\mathbf{Y}$ , as they all agree on  $c_{\mathbf{X}'}$  by Definition 6.5 (2, 6). Also note that the special case where  $\varphi_{\mathcal{A}}$  is already defined for at least one variable  $\mathbf{X}' \in [\mathbf{X}]$  is covered by this case as well, as  $\mathbf{X}' \xrightarrow{F}_D [\mathbf{X}]$  follows from  $\mathbf{X}' \in [\mathbf{X}]$ . In

this special case, the tuple-subcontext of  $\varphi_{\mathcal{A}}(\mathbf{X}')$  witnessed by  $\varepsilon$  (i.e.,  $\varphi_{\mathcal{A}}(\mathbf{X}')$  itself) will be assigned to each variable in  $[\mathbf{X}]$ .

We show that  $\varphi'_{\mathcal{A}}$  indeed respects the constraints in  $D$  according to the properties stated in Definition 6.5. To this end, let  $\mathbf{X}'$  and  $\mathbf{Y}'$  be two variables, with  $c_{\mathbf{X}'} = \varphi'_{\mathcal{A}}(\mathbf{X}')$  and  $c_{\mathbf{Y}'} = \varphi'_{\mathcal{A}}(\mathbf{Y}')$ . (1) By construction,  $c_{\mathbf{X}'}$  is a tuple-context for  $\text{type}(\mathbf{X}')$ . (2-4) Note that if  $\mathbf{X}'' \stackrel{\varepsilon}{\Rightarrow}_D \mathbf{Z}''$  with  $\mathbf{X}'' \in [\mathbf{X}]$ , then  $\mathbf{Y} \stackrel{F \cdot F'}{\Rightarrow}_D \mathbf{Z}''$  and  $c_{\mathbf{X}''} = \varphi'_{\mathcal{A}}(\mathbf{X}'')$  is the tuple-subcontext of  $c_{\mathbf{Y}} = \varphi_{\mathcal{A}}(\mathbf{Y})$  witnessed by  $F'$ , implying that  $c_{\mathbf{X}''}(F') = c_{\mathbf{Y}}(F \cdot F')$ . If  $\mathbf{X}'$  and/or  $\mathbf{Y}'$  are in  $[\mathbf{X}]$ , then we can apply this substitution and use the fact that  $\varphi_{\mathcal{A}}$  respects the constraints in  $\tau$  to conclude that the desired properties hold for  $\varphi'_{\mathcal{A}}$ . (5) Assume  $\mathbf{X}' \stackrel{\varepsilon}{\Rightarrow}_D \mathbf{Y}'$ . If both  $\mathbf{X}' \in [\mathbf{X}]$  and  $\mathbf{Y}' \in [\mathbf{X}]$ , then  $F' = \varepsilon$  as otherwise the schema graph is not acyclic. Since  $c_{\mathbf{Y}'} = c_{\mathbf{X}'}$ , it follows that  $c_{\mathbf{Y}'}$  is a tuple-subcontext of  $c_{\mathbf{X}'}$  witnessed by  $\varepsilon$ . If  $\mathbf{X}' \in [\mathbf{X}]$  and  $\mathbf{Y}' \notin [\mathbf{X}]$ , then  $\mathbf{Y} \stackrel{F \cdot F'}{\Rightarrow}_{\tau} \mathbf{Y}'$  and  $c_{\mathbf{Y}'}$  is a tuple-subcontext of  $c_{\mathbf{Y}}$  witnessed by  $F \cdot F'$  as  $\varphi_{\mathcal{A}}$  respects the constraints of  $\tau$ . Since  $c_{\mathbf{X}'}$  is the tuple-subcontext of  $c_{\mathbf{Y}}$  witnessed by  $F$ , it follows that  $c_{\mathbf{Y}'}$  is a tuple-subcontext of  $c_{\mathbf{X}'}$  witnessed by  $F'$ . If  $\mathbf{X}' \notin [\mathbf{X}]$  and  $\mathbf{Y}' \in [\mathbf{X}]$ , then  $\mathbf{Y} \stackrel{\varepsilon}{\Rightarrow}_{\tau} \mathbf{Y}'$ . Since  $\varphi_{\mathcal{A}}$  respects the constraints in  $D$ , we apply Definition 6.5 (2, 5, 6) to conclude that  $c_{\mathbf{Y}'}$  is a tuple-subcontext of  $c_{\mathbf{X}'}$  witnessed by  $F'$ . (6) Assume  $c_{\mathbf{X}'}(F_{\mathbf{X}'}) = c_{\mathbf{Y}'}(F_{\mathbf{Y}'})$ , and let  $c_{\mathbf{X}''}$  and  $c_{\mathbf{Y}''}$  be the tuple-subcontexts of respectively  $c_{\mathbf{X}'}$  witnessed by  $F_{\mathbf{X}'}$  and  $c_{\mathbf{Y}'}$  witnessed by  $F_{\mathbf{Y}'}$ . We argue that  $c_{\mathbf{X}''} = c_{\mathbf{Y}''}$ . Note that, if  $\mathbf{X}' \in [\mathbf{X}]$ , then  $c_{\mathbf{X}''}$  is the tuple-subcontext of  $c_{\mathbf{Y}}$  witnessed by  $F \cdot F_{\mathbf{X}'}$ . The reasoning for  $\mathbf{Y}' \in [\mathbf{X}]$  is analogous. Since  $\varphi_{\mathcal{A}}$  respects the constraints in  $D$ , it follows that  $c_{\mathbf{X}''} = c_{\mathbf{Y}''}$ .

(Case 2) Otherwise, we construct a fresh tuple-context  $c_{\mathbf{X}}$  and define  $\varphi'_{\mathcal{A}}(\mathbf{X}') = c_{\mathbf{X}}$  for every variable  $\mathbf{X}' \in [\mathbf{X}]$ . This tuple-context  $c_{\mathbf{X}}$  is constructed as follows:  $c_{\mathbf{X}}(\varepsilon) = \mathbf{t}_{\mathbf{X}}$ , with  $\mathbf{t}_{\mathbf{X}}$  a fresh tuple of the appropriate type. For every path  $F = f \cdot F'$  in  $SG$  starting in  $\text{type}(\mathbf{X})$ , if there is a variable  $\mathbf{Y}$  with  $[\mathbf{X}] \stackrel{f}{\Rightarrow}_D \mathbf{Y}$ , then  $c_{\mathbf{X}}(F) = c_{\mathbf{Y}}(F')$ , with  $c_{\mathbf{Y}} = \varphi_{\mathcal{A}}(\mathbf{Y})$ . In other words,  $c_{\mathbf{Y}}$  is the tuple-subcontext of  $c_{\mathbf{X}}$  witnessed by  $f$ . Note that due to the order  $<_{SG}$ ,  $\varphi_{\mathcal{A}}(\mathbf{Y})$  has to be defined already. Also note that this is well defined, even if multiple such  $\mathbf{Y}$  exist. In that case, all these  $\mathbf{Y}$  are equivalent to each other by definition of  $\equiv_D$ , and by construction of  $\varphi_{\mathcal{A}}$  they are assigned the same tuple-context. If instead no such variable  $\mathbf{Y}$  exists, we define  $c_{\mathbf{X}}(F) = \mathbf{t}_F$ , with  $\mathbf{t}_F$  a fresh tuple of the appropriate type.

We show that  $\varphi'_{\mathcal{A}}$  indeed respects the constraints in  $D$  according to the properties stated in Definition 6.5. To this end, let  $\mathbf{X}'$  and  $\mathbf{Y}'$  be two variables occurring in  $\text{Trans}(D)$ , with  $c_{\mathbf{X}'} = \varphi'_{\mathcal{A}}(\mathbf{X}')$  and  $c_{\mathbf{Y}'} = \varphi'_{\mathcal{A}}(\mathbf{Y}')$ . (1) By construction,  $c_{\mathbf{X}'}$  is a tuple-context for  $\text{type}(\mathbf{X}')$ . (2) Assume  $\mathbf{X}' \stackrel{\varepsilon}{\Rightarrow}_D \mathbf{Z}$  and  $\mathbf{Y}' \stackrel{\varepsilon}{\Rightarrow}_D \mathbf{Z}$  for some variable  $\mathbf{Z}$ . We argue that there exists a pair of variables  $\mathbf{X}''$  and  $\mathbf{Y}''$  and two sequences of function names  $F'_1$  and  $F'_2$  such that  $c_{\mathbf{X}'}(F_1) = c_{\mathbf{X}''}(F'_1) = c_{\mathbf{Y}''}(F'_2) = c_{\mathbf{Y}'}(F_2)$ , where  $c_{\mathbf{X}''} = \varphi'_{\mathcal{A}}(\mathbf{X}'')$  and  $c_{\mathbf{Y}''} = \varphi'_{\mathcal{A}}(\mathbf{Y}'')$ . If  $\mathbf{X}' \in [\mathbf{X}]$ , then either  $F_1 = f \cdot F'_1$  or  $F_1 = \varepsilon$ . In the former case there is a variable  $\mathbf{X}''$  with  $\mathbf{X}'' \stackrel{f}{\Rightarrow}_D \mathbf{Z}$  such that  $c_{\mathbf{X}'}(F_1) = c_{\mathbf{X}''}(F'_1)$ , where  $c_{\mathbf{X}''} = \varphi_{\mathcal{A}}(\mathbf{X}'')$ . In the latter case,  $\mathbf{Z} \in [\mathbf{X}]$ , and we simply take  $\mathbf{X}'' = \mathbf{X}'$  and  $F'_1 = F_1$ . If  $\mathbf{X}' \notin [\mathbf{X}]$ , we take  $\mathbf{X}'' = \mathbf{X}'$  and  $F'_1 = F_1$ . For  $\mathbf{Y}' \in [\mathbf{X}]$  and  $\mathbf{Y}' \notin [\mathbf{X}]$ , the reasoning is analogous. It remains to argue that  $c_{\mathbf{X}''}(F'_1) = c_{\mathbf{Y}''}(F'_2)$ . By choice of  $\mathbf{X}''$  and  $\mathbf{Y}''$ , either  $\mathbf{X}'' \notin [\mathbf{X}]$  and  $\mathbf{Y}'' \notin [\mathbf{X}]$ ; or  $\mathbf{Z} \in [\mathbf{X}]$ . In the former case,  $c_{\mathbf{X}''}(F'_1) = c_{\mathbf{Y}''}(F'_2)$  follows by the fact that  $\varphi_{\mathcal{A}}$  respects the constraints of  $D$ . In the latter case, both  $\mathbf{X}'' \in [\mathbf{X}]$  and  $\mathbf{Y}'' \in [\mathbf{X}]$ , as otherwise (Case 1) would apply to  $[\mathbf{X}]$  instead. Then,  $c_{\mathbf{X}''}(F'_1) = c_{\mathbf{Y}''}(F'_2) = c_{\mathbf{X}}(\varepsilon) = \mathbf{t}_{\mathbf{X}}$ . (3, 4) The reasoning is analogous to the previous property. Note in particular that by construction of the new  $c_{\mathbf{X}}$  we have  $c_{\mathbf{X}'}(F_1) = c_{\mathbf{Y}'}(F_2)$  if  $\mathbf{W} \equiv_D \mathbf{Z}$ . Since  $\mathbf{W} \equiv_D \mathbf{Z}$  implies that there is no constraint  $\mathbf{W} \neq \mathbf{Z}$  by the assumptions on  $\varphi_{\mathcal{A}}$  and on the disequality constraints in each template  $\tau \in \text{Trans}(D)$ , this does not lead to contradictions. (5) If  $\mathbf{X}' \stackrel{\varepsilon}{\Rightarrow}_{\mathbf{Y}'}$ , then  $\mathbf{Y}' \in [\mathbf{X}]$  only if  $\mathbf{X}' \in [\mathbf{X}]$ , as otherwise (Case 1) would apply to  $[\mathbf{X}]$  instead. We argue by case that  $c_{\mathbf{Y}'}$

is a tuple-subcontext of  $c_{X'}$  witnessed by  $F'$ . If  $X' \notin [X]$  and  $Y' \notin [X]$ , the result is immediate by the fact that  $\varphi_{\mathcal{A}}$  respects the constraints of  $D$ . If  $X' \in [X]$  and  $Y' \notin [X]$ , then  $c_{X'} = c_X$  and a variable  $Z$  exists such that  $F' = f \cdot F''$ ,  $X' \xrightarrow{f} Z$ ,  $Z \xrightarrow{F''} Y'$ , and, by construction of  $c_X$ ,  $c_{Y'}$  is a tuple-subcontext of  $\varphi_{\mathcal{A}}(Z)$  witnessed by  $F''$ . It now follows that  $c_{Y'}$  is a tuple-subcontext of  $c_X$  witnessed by  $F'$ . Lastly, If both  $X' \in [X]$  and  $Y' \in [X]$ , then  $F' = \varepsilon$ , as otherwise the schema graph is not acyclic. The result is immediate, as  $c_{Y'} = c_{X'} = c_X$  is by definition a tuple-subcontext of itself witnessed by  $\varepsilon$ . (6) Assume  $c_{X''}(F_1) = c_{Y''}(F_1)$  for some pair of tuple-contexts  $c_{X''}$  and  $c_{Y''}$  that are tuple-subcontexts of respectively  $c_{X'}$  witnessed by  $F_1$  and  $c_{Y'}$  witnessed by  $F_2$ . We argue that  $c_{X''} = c_{Y''}$ . If both  $c_{X'}$  and  $c_{Y'}$  are different from  $c_X$ , the result is immediate as  $\varphi_{\mathcal{A}}$  respects the constraints of  $D$ . Otherwise, since the construction of  $c_X$ , either copies existing tuple-contexts as tuple-subcontexts, or introduces fresh variables. the result holds if  $c_{X'}$  and/or  $c_{Y'}$  are equal to  $c_X$ .  $\square$

Given a total context assignment respecting the constraints in  $D$  as in Lemma 6.7, we show in the next Lemma that such a total context assignment implies a variable assignment  $\bar{\mu}$  such that the  $C = \bar{\mu}(D)$  is a valid sequence of conflicting quadruples.

**Lemma 6.8.** *Let  $D = (\tau_1, o_1, p_2, \tau_2), \dots, (\tau_m, o_m, p_1, \tau_1)$  be a sequence of potentially conflicting quadruples over a set of transaction templates  $\mathcal{P}$  and  $\varphi_{\mathcal{A}}$  a total context assignment for  $D$  respecting the constraints of  $D$ . The variable mapping  $\bar{\mu}$  obtained by defining  $\bar{\mu}(X) = c_X(\varepsilon)$  for every variable  $X$  in  $\text{Trans}(D)$  with  $c_X = \varphi_{\mathcal{A}}(X)$  then is a valid variable mapping admissible for some database  $\mathbf{D}$ .*

*Proof.* We first argue that  $\bar{\mu}$  is valid by showing for each conflicting quadruple  $(\tau_i, o_i, p_j, \tau_j)$  in  $D$  that  $\bar{\mu}(X) = \bar{\mu}(Y)$  with  $X = \text{var}(o_i)$  and  $Y = \text{var}(p_j)$ . By definition,  $X \xrightarrow{\tau_i} Y$ , and hence  $X \xrightarrow{\varepsilon} Y$ . Let  $c_X = \varphi_{\mathcal{A}}(X)$  and  $c_Y = \varphi_{\mathcal{A}}(Y)$ . Since  $\varphi_{\mathcal{A}}$  respects the constraints of  $D$ ,  $c_Y$  is a tuple-subcontext of  $c_X$  witnessed by  $\varepsilon$ . By definition of tuple-subcontexts,  $c_X(\varepsilon) = c_Y(\varepsilon)$ , and, as a result,  $\bar{\mu}(X) = c_X(\varepsilon) = c_Y(\varepsilon) = \bar{\mu}(Y)$ .

Next, we construct a database  $\mathbf{D}$  and show that  $\bar{\mu}$  is admissible for  $\mathbf{D}$ . To this end, we add the tuple  $\bar{\mu}(X)$  to  $\mathbf{D}$  for each variable  $X$  occurring in  $\text{Trans}(D)$ . For each functional constraint  $Y = f(X)$  in a transaction template in  $\text{Trans}(D)$ , we define  $\bar{\mu}(Y) = f^{\mathbf{D}}(\bar{\mu}(X))$  for the corresponding function  $f^{\mathbf{D}}$  in  $\mathbf{D}$ . Note that this is well defined. Towards a contradiction, assume that we have both  $\bar{\mu}(Y) = f^{\mathbf{D}}(\bar{\mu}(X))$  and  $\bar{\mu}(W) = f^{\mathbf{D}}(\bar{\mu}(Z))$ , with  $\bar{\mu}(X) = \bar{\mu}(Z)$  but  $\bar{\mu}(Y) \neq \bar{\mu}(W)$ . Let  $c_X = \varphi_{\mathcal{A}}(X)$ ,  $c_Y = \varphi_{\mathcal{A}}(Y)$ ,  $c_Z = \varphi_{\mathcal{A}}(Z)$  and  $c_W = \varphi_{\mathcal{A}}(W)$ . By construction of  $\bar{\mu}$ , we have  $c_X(\varepsilon) = \bar{\mu}(X) = \bar{\mu}(Z) = c_Z(\varepsilon)$  and  $c_Y(\varepsilon) = \bar{\mu}(Y) \neq \bar{\mu}(W) = c_W(\varepsilon)$ . By Definition 6.5 (6), it now follows that  $c_X = c_Z$ , since  $c_X$  (respectively  $c_Z$ ) is a tuple-subcontext of itself witnessed by  $\varepsilon$  and  $c_X(\varepsilon) = c_Z(\varepsilon)$ . Since we defined  $\bar{\mu}(Y) = f^{\mathbf{D}}(\bar{\mu}(X))$ , there is a constraint  $Y = f(X)$  in some template in  $\text{Trans}(D)$ , and hence  $X \xrightarrow{f} Y$ . Analogously,  $Z \xrightarrow{f} W$ . By Definition 6.5 (5),  $c_Y$  and  $c_W$  are tuple-subcontexts of respectively  $c_X$  and  $c_Z$  witnessed by  $f$ . As  $c_X = c_Z$ , it immediately follows that  $c_Y = c_W$ , and in particular  $c_Y(\varepsilon) = c_W(\varepsilon)$ , leading to the desired contradiction.

To conclude the proof, we show that  $\bar{\mu}$  is indeed admissible for  $\mathbf{D}$ . By construction of  $\mathbf{D}$  based on  $\bar{\mu}$ ,  $\bar{\mu}(Y) = f^{\mathbf{D}}(\bar{\mu}(X))$  is immediate for each constraint  $Y = f(X)$  in a template  $\tau \in \text{Trans}(D)$ . We still need to argue that  $\bar{\mu}(X) \neq \bar{\mu}(Y)$  for each constraint  $X \neq Y$  in a template  $\tau \in \text{Trans}(D)$ . Let  $c_X = \varphi_{\mathcal{A}}(X)$  and  $c_Y = \varphi_{\mathcal{A}}(Y)$ . By construction of  $\bar{\mu}$  we have  $\bar{\mu}(X) = c_X(\varepsilon)$  and  $\bar{\mu}(Y) = c_Y(\varepsilon)$ . Note that  $X \xrightarrow{\tau} Y$  and  $Y \xrightarrow{\tau} X$ . Therefore, we can apply Definition 6.5 (3) to conclude that  $c_X(\varepsilon) \neq c_Y(\varepsilon)$ , and hence  $\bar{\mu}(X) \neq \bar{\mu}(Y)$ .  $\square$

In order to decide robustness against RC, one can now construct a sequence of quintuples  $E$  and derive the sequence of potentially conflicting quadruples  $D$  and partial context assignment  $\varphi_{\mathcal{A}}$  from it. If  $\varphi_{\mathcal{A}}$  respects the constraints in  $D$ , then it follows from Lemma 6.7 and Lemma 6.8 that we can construct a variable assignment  $\bar{\mu}$  such that  $C = \bar{\mu}(D)$  is a valid sequence of conflicting quadruples. However, in this construction of  $E$ , care should be taken to guarantee that  $\varphi_{\mathcal{A}}$  indeed respects the constraints in  $D$ , and that the resulting multiversion split schedule based on  $C$  indeed satisfies all properties in Definition 3.6.

In the algorithm that we are about to propose, we search for such a sequence of quintuples  $E$ , but without fixating all the tuples in each context. For this, we generalize our definition of tuple-contexts to allow variables: A *context for a type  $R$*  is a function from paths with source  $R$  in  $SG(\text{Rels}, \text{Funcs})$  to variables in **Var** and tuples in **Tuples** of the appropriate type. The purpose of variables is to encode equalities and disequalities within each context, without being explicit about the precise tuples. That is, if two paths ending in the same node in  $SG$  are mapped on the same variable, then they will represent the same tuple; if they are mapped on different variables, then they represent a different tuple. We remark that a same variable occurring in different contexts can still represent different tuples. Analogous to tuple-subcontexts, for two types  $R$  and  $S$  with  $R \xrightarrow{f} SG S$ , we say that a context  $c_S$  for type  $S$  is a *subcontext* of a context  $c_R$  for type  $R$  witnessed by  $F$  if:

- for every path  $S \xrightarrow{f} SG S'$  in  $SG$ , if  $c_R(F \cdot F')$  is a tuple, then  $c_S(F') = c_R(F \cdot F')$ ; otherwise,  $c_S(F')$  is a variable; and
- for every pair of paths  $S \xrightarrow{f_1} SG S_1$  and  $S \xrightarrow{f_2} SG S_2$  in  $SG$  with  $c_R(F \cdot F_1)$  and  $c_R(F \cdot F_2)$  variables,  $c_S(F_1) = c_S(F_2)$  iff  $c_R(F \cdot F_1) = c_R(F \cdot F_2)$ .

We call a context a *variable-context* if all paths are mapped on variables.

For a transaction template  $\tau$ , tuple-context  $c_p$  for  $p$  and  $c_o$  for  $o$  in  $\tau$ , we consider the set  $\text{Contexts}(SG, \tau, p, c_p, o, c_o)$  of all different (not-necessarily tuple-) contexts  $c$  (up to isomorphisms over the variables in  $c$ ) that can be obtained, starting from a variable-context  $c'$ , by performing substitutions of subcontexts of  $c'$  with subcontexts of  $c_p$  and/or  $c_o$ . More formally, these substitutions are of the form: For a path  $R_c \xrightarrow{f} SG S$  (here  $R_c$  is the type that  $c$  is for) and  $R_p \xrightarrow{f} SG S$  (with  $R_p$  the type that  $c_p$  is for) then  $c(F \cdot F'') = c_p(F' \cdot F'')$  for every path  $S \xrightarrow{f} SG S'$  in  $SG$  and with  $c(F \cdot F'') = c'(F \cdot F'')$  otherwise. (The substitution rule can be applied for  $c_p$  as well as for  $c_o$ .)

**Lemma 6.9.** *Let  $\mathcal{P}$  be a set of transaction templates over an acyclic schema. Then,  $\mathcal{P}$  is not robust against RC if, and only if, there is a sequence of quintuples  $E = (\tau_1, o_1, c_{o_1}, p_1, c_{p_1}), \dots, (\tau_m, o_m, c_{o_m}, p_m, c_{p_m})$  with  $m \geq 2$  such that for each quintuple  $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$  in  $E$ , with  $q_i$  and  $r_i$  two (not necessarily different) operations in  $\{o_i, p_i\}$ ,*

- (1) *if  $i = 1$ , then  $c_{o_1}$  and  $c_{p_1}$  are tuple-contexts for  $\text{type}(\text{var}(o_1))$  and  $\text{type}(\text{var}(p_1))$ . Furthermore, for every pair of tuple-subcontexts  $c'_{o_1}$  and  $c'_{p_1}$  of  $c_{o_1}$  and  $c_{p_1}$  witnessed by respectively  $F$  and  $F'$ , if  $c_{o_1}(F) = c_{p_1}(F')$ , then  $c'_{o_1} = c'_{p_1}$ ;*
- (2) *if  $i \neq 1$ , then  $c_{o_i}, c_{p_i} \in \text{Contexts}(SG, \tau_1, p_1, c_{p_1}, o_1, c_{o_1})$  are contexts for  $\text{type}(\text{var}(o_i))$  and  $\text{type}(\text{var}(p_i))$ ;*
- (3) *for every pair of variables  $W_i$  and  $Z_i$  in  $\tau_i$  with  $W_i \equiv_{\tau_i} Z_i$ , there is no constraint  $W_i \neq Z_i$  in  $\tau_i$ ;*
- (4) *for every  $\text{var}(q_i) \xrightarrow{f} \tau_i Z_i$  and  $\text{var}(r_i) \xrightarrow{f} \tau_i Z_i$ , the subcontext of  $c_{q_i}$  witnessed by  $F_1$  is equal (up to isomorphisms over variables) to the subcontext of  $c_{r_i}$  witnessed by  $F_2$ .*
- (5) *for every  $\text{var}(q_i) \xrightarrow{f} \tau_i W_i$  and  $\text{var}(r_i) \xrightarrow{f} \tau_i Z_i$  with  $W_i \neq Z_i$  a constraint in  $\tau_i$ ,  $c_{q_i}(F_1) \neq c_{r_i}(F_2)$  or  $c_{q_i}(F_1)$  and  $c_{r_i}(F_2)$  are both variables;*



- (6) for every  $\text{var}(q_i) \xrightarrow{\tau_i} W_i$  and  $\text{var}(r_i) \xrightarrow{\tau_i} Z_i$ , if  $c_{q_i}(F_1)$  and  $c_{r_i}(F_2)$  are the same tuple, then there is no constraint  $W_i \neq Z_i$  in  $\tau_i$ ;
- (7) if  $\text{var}(q_i) \xrightarrow{\tau_i} \text{var}(r_i)$ , then  $c_{r_i}$  is a subcontext of  $c_{q_i}$  witnessed by  $F$ ; and
- (8) If  $i \neq 1$  and  $c_{q_i}(F) = c_{q_1}(F')$  is a tuple for some  $q_1 \in \{o_1, p_1\}$  and some sequence of function names  $F$  and  $F'$ , then there is no operation  $o'_i \in \tau_i$  potentially ww-conflicting with an operation  $o'_1 \in \text{prefix}_{o_1}(\tau_1)$  with  $\text{var}(q_i) \xrightarrow{\tau_i} \text{var}(o'_i)$  and  $\text{var}(q_1) \xrightarrow{\tau_1} \text{var}(o'_1)$ .

Furthermore, for each pair of adjacent quintuples  $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$  and  $(\tau_j, o_j, c_{o_j}, p_j, c_{p_j})$  in  $E$  with  $j = i + 1$ , or  $i = m$  and  $j = 1$ :

- (9)  $o_i$  is potentially conflicting with  $p_j$  and  $c_{o_i} = c_{p_j}$ ;
- (10) if  $i = 1$  and  $j = 2$ , then  $o_1$  is potentially rw-conflicting with  $p_2$ ; and
- (11) if  $i = m$  and  $j = 1$ , then  $o_1 <_{\tau_1} p_1$  or  $o_m$  is potentially rw-conflicting with  $p_1$ .

*Proof.* (if) Let  $D = (\tau_1, o_1, p_2, \tau_2), \dots, (\tau_m, o_m, p_1, \tau_1)$  be the sequence of potentially conflicting quadruples derived from  $E$ . Note that each  $(\tau_i, o_i, p_j, \tau_j) \in D$  is indeed a valid sequence of potentially conflicting quadruples, as  $o_i$  is potentially conflicting with  $p_j$  by (9). We show in Lemma 6.10 that a partial context assignment  $\varphi_{\mathcal{A}}$  over a set of tuple-contexts  $\mathcal{A}$  exists such that

- for every pair of operations  $o_i$  and  $p_i$  occurring in  $D$ ,  $\varphi_{\mathcal{A}}$  is defined for  $\text{var}(o_i)$  and  $\text{var}(p_i)$ ;
- $\varphi_{\mathcal{A}}$  respects the constraints in  $D$ ; and
- for every template  $\tau_i$  in  $D$  with  $i \neq 1$  and for every  $q_i \in \{o_i, p_i\}$  and  $q_1 \in \{o_1, p_1\}$ , let  $c_{q_i} = \varphi_{\mathcal{A}}(\text{var}(q_i))$  and  $c_{q_1} = \varphi_{\mathcal{A}}(\text{var}(q_1))$ . If  $c_{q_i}(F) = c_{q_1}(F')$  for some sequence of function names  $F$  and  $F'$ , then there is no operation  $o'_i \in \tau_i$  potentially ww-conflicting with an operation  $o'_1 \in \text{prefix}_{o_1}(\tau_1)$  with  $\text{var}(q_i) \xrightarrow{\tau_i} \text{var}(o'_i)$  and  $\text{var}(q_1) \xrightarrow{\tau_1} \text{var}(o'_1)$ .

Because of (3), we can now apply Lemma 6.7 extending  $\varphi_{\mathcal{A}}$  to a total context assignment defined for all variables occurring in  $\text{Trans}(D)$ , without losing the property that  $\varphi_{\mathcal{A}}$  respects all constraints in  $D$ . Let  $\bar{\mu}$  be the variable mapping obtained by defining  $\bar{\mu}(X) = c_X(\varepsilon)$  for every variable  $X$  in  $\text{Trans}(D)$  with  $c_X = \varphi_{\mathcal{A}}(X)$ . By Lemma 6.8,  $\bar{\mu}$  is a valid variable mapping and a database  $\mathbf{D}$  exists such that  $\bar{\mu}$  is admissible for  $\mathbf{D}$ .

We now prove that the sequence of conflicting quadruples  $C = \bar{\mu}(D)$  satisfies the conditions stated in Definition 3.6 to show that  $\mathcal{P}$  is indeed not robust against RC. Condition (2) and Condition (3) are immediate by respectively (11) and (10). Towards a contradiction, assume Condition (1) does not hold. Then, there is an operation  $o'_i$  in a template  $\tau_i$  potentially ww-conflicting with an operation  $o'_1 \in \text{prefix}_{o_1}(\tau_1)$ , and  $\bar{\mu}(\text{var}(o'_i)) = \bar{\mu}(\text{var}(o'_1))$ . Let  $c_{o'_i} = \varphi_{\mathcal{A}}(\text{var}(o'_i))$  and  $c_{o'_1} = \varphi_{\mathcal{A}}(\text{var}(o'_1))$ . By construction of  $\bar{\mu}$ , we have  $c_{o'_i}(\varepsilon) = c_{o'_1}(\varepsilon)$ . By construction of the total context assignment in Lemma 6.7, this is only the case if for some  $q_i \in \{o_i, p_i\}$  and  $q_1 \in \{o_1, p_1\}$  it holds that  $\text{var}(q_i) \xrightarrow{\tau_i} \text{var}(o'_i)$  with  $c_{q_i}(F) = c_{o'_i}(\varepsilon)$  and  $\text{var}(q_1) \xrightarrow{\tau_1} \text{var}(o'_1)$  with  $c_{q_1}(F') = c_{o'_1}(\varepsilon)$ . Consequently,  $c_{q_i}(F) = c_{q_1}(F')$ , contradicting Lemma 6.10.

(only if) Since  $\mathcal{P}$  is not robust against RC, a multiversion split schedule  $s$  exists based on a sequence of conflicting quadruples  $C$  over a set of transactions  $\mathcal{T}$  consistent with a set of transaction templates  $\mathcal{P} \in \mathbf{AcycTemp}$  and a database  $\mathbf{D}$ . Let  $\bar{\mu}$  be the variable mapping for a sequence of potentially conflicting quadruples  $D = (\tau_1, o_1, p_2, \tau_2), \dots, (\tau_m, o_m, p_1, \tau_1)$  with  $\bar{\mu}(D) = C$ . By Lemma 6.6 a set  $\mathcal{A}$  of tuple-contexts and a total context assignment  $\varphi_{\mathcal{A}}$  for  $D$  over  $\mathcal{A}$  exist with:

- $\varphi_{\mathcal{A}}$  respects the constraints of  $D$ ; and
- $\bar{\mu}(X) = c_X(\varepsilon)$  for every variable  $X$ , with  $c_X = \varphi_{\mathcal{A}}(X)$ .

Let  $\lambda : \mathbf{Tuples} \rightarrow \mathbf{Tuples} \cup \mathbf{Var}$  with  $\lambda(\mathbf{t}) = \mathbf{t}$  if  $\mathbf{t}$  occurs in  $\varphi_{\mathcal{A}}(\text{var}(o_1))$  or  $\varphi_{\mathcal{A}}(\text{var}(p_1))$ ; and  $\lambda(\mathbf{t}) \in \mathbf{Var}$  otherwise, such that  $\lambda(\mathbf{t}) \neq \lambda(\mathbf{v})$  if  $\mathbf{t} \neq \mathbf{v}$ . From  $D$  and  $\varphi_{\mathcal{A}}$  we derive the sequence of quintuples  $E = (\tau_1, o_1, c_{o_1}, p_1, c_{p_1}), \dots, (\tau_m, o_m, c_{o_m}, p_m, c_{p_m})$  with  $c_{o_i} = \lambda \circ c'_{o_i}$  and  $c_{p_i} = \lambda \circ c'_{p_i}$  for each  $o_i$  and  $p_i$ , where  $c'_{o_i} = \varphi_{\mathcal{A}}(\text{var}(o_i))$  and  $c'_{p_i} = \varphi_{\mathcal{A}}(\text{var}(p_i))$ . Intuitively, we modify the tuple-contexts for each  $\text{var}(o_i)$  and  $\text{var}(p_i)$  as defined by  $\varphi_{\mathcal{A}}$  into contexts over tuples and variables by replacing all tuples that do not occur in  $\varphi_{\mathcal{A}}(\text{var}(o_1))$  and  $\varphi_{\mathcal{A}}(\text{var}(p_1))$  with unique variables. Note that by construction  $c_{o_1} = \varphi_{\mathcal{A}}(\text{var}(o_1))$  and  $c_{p_1} = \varphi_{\mathcal{A}}(\text{var}(p_1))$ .

Next, we show that  $E$  satisfies all properties. In the argumentation below, we denote  $\varphi_{\mathcal{A}}(\text{var}(o_i))$  by  $c'_{o_i}$  and  $\varphi_{\mathcal{A}}(\text{var}(p_i))$  by  $c'_{p_i}$  for each  $o_i$  and  $p_i$  in  $E$ . (1) By construction,  $c_{o_1} = c'_{o_1}$  is a tuple-context for  $\text{type}(\text{var}(o_1))$ , and  $c_{p_1} = c'_{p_1}$  is a tuple-context for  $\text{type}(\text{var}(p_1))$ . By Condition (6) of Definition 6.5, we have  $c'_{o_1} = c'_{p_1}$  if  $c_{o_1}(F) = c_{p_1}(F')$  for every pair of tuple-subcontexts  $c'_{o_1}$  and  $c'_{p_1}$  of  $c_{o_1}$  and  $c_{p_1}$  witnessed by respectively  $F$  and  $F'$ . (2) This property follows by construction of  $c_{o_i}$  and  $c_{p_i}$  based on  $\lambda$  and  $\varphi_{\mathcal{A}}$ . For completeness sake, one should note that for each group of contexts up to isomorphisms over variables,  $\text{Contexts}(SG, \tau_1, p_1, c_{p_1}, o_1, c_{o_1})$  contains only one context. W.l.o.g. we can implicitly assume that each  $c_{o_i}$  and  $c_{p_i}$  in  $E$  is replaced by the same context in  $\text{Contexts}(SG, \tau_1, p_1, c_{p_1}, o_1, c_{o_1})$  up to isomorphisms, as we will never directly test for equality or disequality between two variables of different contexts. (3) Assume  $\mathbb{W}_i \equiv_{\tau_i} \mathbb{Z}_i$ , then  $\mathbb{W}_i \xrightarrow{D} \mathbb{Z}_i$ . Since  $\varphi_{\mathcal{A}}$  respects the constraints in  $D$ ,  $c_{z_i}$  is a tuple-subcontext of  $c_{w_i}$  witnessed by  $\varepsilon$ , with  $c_{z_i} = \varphi_{\mathcal{A}}(\mathbb{Z}_i)$  and  $c_{w_i} = \varphi_{\mathcal{A}}(\mathbb{W}_i)$ . Then,  $c_{z_i} = c_{w_i}$ , and in particular  $c_{z_i}(\varepsilon) = c_{w_i}(\varepsilon)$ . By Lemma 6.6,  $\bar{\mu}(\mathbb{Z}_i) = c_{z_i}(\varepsilon) = c_{w_i}(\varepsilon) = \bar{\mu}(\mathbb{W}_i)$ . Since  $\bar{\mu}$  is admissible for  $\mathbf{D}$ , the constraint  $\mathbb{W}_i \neq \mathbb{Z}_i$  cannot exist. (4) Since  $\varphi_{\mathcal{A}}$  respects the constraints in  $D$ , it follows from Conditions (2) and (6) in Definition 6.5 that the tuple-subcontext of  $c'_{q_i}$  witnessed by  $F_1$  is equal to the tuple-subcontext of  $c'_{r_i}$  witnessed by  $F_2$ . As a result, the subcontext of  $c_{q_i} = \lambda \circ c'_{q_i}$  witnessed by  $F_1$  is equal (up to isomorphisms over variables) to the subcontext of  $c_{r_i} = \lambda \circ c'_{r_i}$  witnessed by  $F_2$ . (5) Analogous to the previous case, we can conclude that  $c_{q_i}(F_1) = \lambda \circ c'_{q_i}(F_1)$  and  $c_{r_i}(F_2) = \lambda \circ c'_{r_i}(F_2)$  are either two different tuples, or both a variable. (6) If  $c_{q_i}(F_1) = \lambda \circ c'_{q_i}(F_1) = \lambda \circ c'_{r_i}(F_2) = c_{r_i}(F_2)$  is a tuple, then  $c'_{q_i}(F_1) = c'_{r_i}(F_2)$ . Since  $\varphi_{\mathcal{A}}$  respects the constraints in  $D$ , it follows that there is no constraint  $\mathbb{W}_i \neq v z_i$ . (7) If  $\text{var}(q_i) \xrightarrow{\tau_i} \text{var}(r_i)$ , then  $c'_{r_i}$  is a tuple-subcontext of  $c'_{q_i}$ , as  $\varphi_{\mathcal{A}}$  respects the constraints of  $D$ . By construction of  $c_{q_i}$  and  $c_{r_i}$  based on respectively  $c'_{q_i}$  and  $c'_{r_i}$ , it immediately follows that  $c_{r_i}$  is a subcontext of  $c_{q_i}$ . The case for  $\mathbb{Y}_i \xrightarrow{\tau_i} \mathbb{X}_i$  is analogous. (8) Assume towards a contradiction that  $c_{q_i}(F) = c_{q_1}(F')$  is a tuple for some  $q_1 \in \{o_1, p_1\}$  and some sequence of function names  $F$  and  $F'$ , and there is an operation  $o'_i \in \tau_i$  potentially ww-conflicting with an operation  $o'_1 \in \text{prefix}_{o_1}(\tau_1)$  with  $\text{var}(q_i) \xrightarrow{\tau_i} \text{var}(o'_i)$  and  $\text{var}(q_1) \xrightarrow{\tau_1} \text{var}(o'_1)$ . Since  $c_{q_i}(F) = c_{q_1}(F')$  is a tuple,  $c'_{q_i}(F) = c_{q_i}(F) = c_{q_1}(F') = c'_{q_1}(F')$ . By definition of  $\bar{\mu}$  and since  $\varphi_{\mathcal{A}}$  respects the constraints in  $D$ , we conclude that  $\bar{\mu}(o'_i)$  in  $\bar{\mu}(\tau_i)$  is ww-conflicting with  $\bar{\mu}(o'_1)$  in  $\text{prefix}_{\bar{\mu}(o_1)}(\bar{\mu}(\tau_1))$ , thereby contradicting Condition (1) of Definition 3.6. (9) Since  $E$  is based on  $C$ , the operation  $o_i$  is potentially conflicting with  $p_j$ . Furthermore, since  $\text{var}(o_i) \equiv_D \text{var}(p_j)$  and since  $\varphi_{\mathcal{A}}$  respects the constraints of  $D$ ,  $c'_{o_i} = c'_{p_j}$ , and hence  $c_{o_i} = c_{p_j}$ . (10) Immediate by Condition (3) of Definition 3.6. (11) Immediate by Condition (2) of Definition 3.6.  $\square$

Central to the correctness of the if-direction of Lemma 6.9 is the observation that for any sequence  $E$  satisfying the stated conditions, we can assign tuples to variables in each context, thereby replacing contexts with tuple-contexts, in such a way that the resulting context

assignment over tuple-contexts respects the constraints of the corresponding sequence of potentially conflicting quadruples  $D$ . This observation is formalized in the following Lemma:

**Lemma 6.10.** *Let  $E = (\tau_1, o_1, c_{o_1}, p_1, c_{p_1}), \dots, (\tau_m, o_m, c_{o_m}, p_m, c_{p_m})$  be a sequence of quintuples satisfying the conditions stated in Lemma 6.9, and let  $D = (\tau_1, o_1, p_2, \tau_2), \dots, (\tau_m, o_m, p_1, \tau_1)$  be the sequence of potentially conflicting quadruples derived from  $E$ . Then a partial context assignment  $\varphi_{\mathcal{A}}$  over a set of tuple-contexts  $\mathcal{A}$  exists such that*

- for every pair of operations  $o_i$  and  $p_i$  occurring in  $D$ ,  $\varphi_{\mathcal{A}}$  is defined for  $\text{var}(o_i)$  and  $\text{var}(p_i)$ ;
- $\varphi_{\mathcal{A}}$  respects the constraints in  $D$ ; and
- for every template  $\tau_i$  in  $D$  with  $i \neq 1$  and for every  $q_i \in \{o_i, p_i\}$  and  $q_1 \in \{o_1, p_1\}$ , let  $c_{q_i} = \varphi_{\mathcal{A}}(\text{var}(q_i))$  and  $c_{q_1} = \varphi_{\mathcal{A}}(\text{var}(q_1))$ . If  $c_{q_i}(F) = c_{q_i}(F')$  for some sequence of function names  $F$  and  $F'$ , then there is no operation  $o'_i \in \tau_i$  potentially ww-conflicting with an operation  $o'_1 \in \text{prefix}_{o_1}(\tau_1)$  with  $\text{var}(q_i) \xrightarrow{F} \tau_i \text{var}(o'_i)$  and  $\text{var}(q_1) \xrightarrow{F'} \tau_1 \text{var}(o'_1)$ .

*Proof.* The general proof idea is as follows. We iteratively extend  $\varphi_{\mathcal{A}}$  by deriving  $\varphi_{\mathcal{A}}(\text{var}(o_i))$  and  $\varphi_{\mathcal{A}}(\text{var}(p_i))$  from contexts  $c_{o_i}$  and  $c_{p_i}$  for each quintuple  $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$  in the order that they appear in  $E$ . Afterwards, we argue that  $\varphi_{\mathcal{A}}$  respects the constraints in  $D$  and for every  $q_i \in \{o_i, p_i\}$  with  $i \neq 1$  and  $q_1 \in \{o_1, p_1\}$  with  $c_{q_i} = \varphi_{\mathcal{A}}(\text{var}(q_i))$  and  $c_{q_1} = \varphi_{\mathcal{A}}(\text{var}(q_1))$ , and for every pair of sequences of function names  $F$  and  $F'$  with  $c_{q_i}(F) = c_{q_i}(F')$ , there is no operation  $o'_i \in \tau_i$  potentially ww-conflicting with an operation  $o'_1 \in \text{prefix}_{o_1}(\tau_1)$  with  $\text{var}(q_i) \xrightarrow{F} \tau_i \text{var}(o'_i)$  and  $\text{var}(q_1) \xrightarrow{F'} \tau_1 \text{var}(o'_1)$ .

Let  $(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$  be the first quintuple in the sequence  $E$ . We initiate  $\varphi_{\mathcal{A}}$  by defining  $\varphi_{\mathcal{A}}(\text{var}(o_1)) = c_{o_1}$  and  $\varphi_{\mathcal{A}}(\text{var}(p_1)) = c_{p_1}$ . Next, we iteratively extend  $\varphi_{\mathcal{A}}$  by considering the remaining quintuples in  $E$  in order. To this end, let  $(\tau_{i-1}, o_{i-1}, c_{o_{i-1}}, p_{i-1}, c_{p_{i-1}})$  be the last considered quintuple and  $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$  the next quintuple in  $E$ . We define  $\varphi_{\mathcal{A}}(\text{var}(p_i)) = c'_{p_i} = \varphi_{\mathcal{A}}(\text{var}(o_{i-1}))$  and  $\varphi_{\mathcal{A}}(\text{var}(o_i)) = c'_{o_i} = \lambda_i \circ c_{o_i}$ , where  $\lambda_i : \mathbf{Tuples} \cup \mathbf{Var} \rightarrow \mathbf{Tuples}$  is a function mapping tuples and variables occurring in  $c_{o_i}$  to tuples such that<sup>7</sup>

- $\lambda_i(\mathbf{t}) = \mathbf{t}$  for each tuple  $\mathbf{t}$  occurring in  $c_{o_i}$ ;
- $\lambda_i(\mathbf{Q}) = \mathbf{v}$  for each variable  $\mathbf{Q}$  occurring in  $c_{o_i}$  for which there is a variable  $\mathbf{Z}$  in  $\tau_i$  and sequences of function names  $F, F'$  and  $F''$  with  $\text{var}(p_i) \xrightarrow{F} \tau_i \mathbf{Z}$ ,  $\text{var}(o_i) \xrightarrow{F'} \tau_i \mathbf{Z}$ ,  $c'_{p_i}(F \cdot F'') = \mathbf{v}$  and  $c_{o_i}(F' \cdot F'') = \mathbf{Q}$ ; and
- $\lambda_i(\mathbf{Q}) = \mathbf{t}_{i,\mathbf{Q}}$ , for the remaining variables  $\mathbf{Q}$  in  $c_{o_i}$  where  $\mathbf{t}_{i,\mathbf{Q}}$  is a fresh tuple.

Note that this  $\lambda_i$  is well defined. In particular, the second rule intuitively states that the tuple-subcontext of the resulting  $c'_{o_i}$  witnessed by  $F'$  is equal to the tuple-subcontext of  $c'_{p_i}$  witnessed by  $F$ , given that there is a variable  $\mathbf{Z}$  with  $\text{var}(o_i) \xrightarrow{F'} \tau_i \mathbf{Z}$  and  $\text{var}(p_i) \xrightarrow{F} \tau_i \mathbf{Z}$ . This substitution is well defined since in this case the subcontext of  $c_{o_i}$  witnessed by  $F'$  is equal (up to isomorphisms over variables) to the subcontext of  $c_{p_i}$  witnessed by  $F$  according to Condition 4 of Lemma 6.9.

It remains to show that  $\varphi_{\mathcal{A}}$  indeed satisfies the conditions stated in Lemma 6.10. To this end, note that by definition of variable determination, if  $\mathbf{X} \xrightarrow{F} D \mathbf{Y}$  with  $\mathbf{X}$  in a template  $\tau_i$  and  $\mathbf{Y}$  in a template  $\tau_j$  in  $\text{Trans}(D)$ , then a sequence of variables  $\mathbf{X}_{k_1}, \mathbf{Y}_{k_1}, \dots, \mathbf{X}_{k_m}, \mathbf{Y}_{k_m}$  exists such that ( $\dagger$ ):

- $\mathbf{X}_{k_1} = \mathbf{X}$  and  $\mathbf{Y}_{k_m} = \mathbf{Y}$ ;

<sup>7</sup>Note that  $\lambda_i$  is defined over variables in the context  $c_{o_i}$ . These variables are unrelated to the variables occurring in  $\text{Trans}(D)$ . We therefore denote these variables by  $\mathbf{Q}$  instead of the usual  $\mathbf{W}, \mathbf{X}, \mathbf{Y}, \mathbf{Z}$  to avoid confusion.

- each pair of (not necessarily different) variables  $X_{k_i}, Y_{k_i}$  occur in the same template  $\tau_{k_i}$  in  $\text{Trans}(D)$  and  $X_{k_i} \xrightarrow{E} \tau_{k_i} Y_{k_i}$  with  $F = F_1 \cdot \dots, F_m$ ;
- in the implied sequence of templates  $\tau_{k_1}, \dots, \tau_{k_m}$ , these  $\tau_{k_i}, \dots, \tau_{k_{i+1}}$  are neighbouring in  $E$  (where we assume that  $\tau_1$  is neighbouring to  $\tau_n$  in  $E$ ); and
- for each pair of variables  $Y_{k_i}, X_{k_{i+1}}$ , there is a sequence of function names  $F'$  such that  $\text{var}(o_{k_i}) \xrightarrow{F'} \tau_{k_i} Y_{k_i}$  and  $\text{var}(p_{k_{i+1}}) \xrightarrow{F'} \tau_{k_{i+1}} X_{k_{i+1}}$  (i.e., equivalence of  $Y_{k_i}$  and  $X_{k_{i+1}}$  in  $D$  is implied by equivalence of  $\text{var}(o_{k_i})$  and  $\text{var}(p_{k_{i+1}})$ ).

In other words,  $X \xrightarrow{E} D Y$  can be broken down into a sequence of  $X_{k_i} \xrightarrow{E} \tau_{k_i} Y_{k_i}$  through a sequence of neighbouring templates, where equivalence between each  $Y_{k_i}$  and  $X_{k_{i+1}}$  is implied by the variables in the potentially conflicting operations  $o_{k_i}$  and  $p_{k_{i+1}}$ . For ease of exposition, we implicitly assumed that  $\tau_{k_1}, \dots, \tau_{k_m}$  agrees with the order in  $E$ . If the order is opposite to the order in  $E$  instead, the above still holds, but the occurrences of  $o_{k_i}$  and  $p_{k_{i+1}}$  should be replaced with  $p_{k_i}$  and  $o_{k_{i+1}}$ .

We argue by construction of  $\varphi_{\mathcal{A}}$  that for every pair of variables  $X$  and  $Y$  for which  $\varphi_{\mathcal{A}}$  is defined with  $c_X = \varphi_{\mathcal{A}}(X)$  and  $c_Y = \varphi_{\mathcal{A}}(Y)$ , if  $c_X(F) = c_Y(F')$  for some sequence of function names  $F$  and  $F'$ , then  $c'_X = c'_Y$ , where  $c'_X$  and  $c'_Y$  are the tuple-subcontexts of  $c_X$  witnessed by  $F$  and  $c_Y$  witnessed by  $F'$ , respectively ( $\dagger$ ). If  $X = \text{var}(o_1)$  and  $Y = \text{var}(p_1)$  (or the other way around), the result is immediate by Lemma 6.9 (1). Otherwise, let  $X$  be the variable in a template  $\tau_i$  and  $Y$  the variable in a template  $\tau_j$  such that  $j \leq i$  (i.e.,  $\tau_i$  does not occur before  $\tau_j$  in  $E$ ). W.l.o.g., we assume that  $X = \text{var}(o_i)$  with  $i \neq 1$  (the case where  $X = \text{var}(p_i)$  is analogous, as  $\varphi_{\mathcal{A}}(\text{var}(p_i)) = \varphi_{\mathcal{A}}(\text{var}(o_{i-1}))$  by construction). By construction of each  $c'_{o_i}$  based on  $c'_{p_i}$  and  $\lambda_i$ , if  $c'_{o_i}(F_i) = c'_{p_i}(F'_i)$ , then the whole tuple-subcontext of  $c'_{o_i}$  witnessed by  $F_i$  is copied over from the tuple-subcontext of  $c'_{p_i}$  witnessed by  $F'_i$ . Indeed  $\lambda_i$  introduces fresh tuples whenever the tuple for  $c'_{o_i}(F_i)$  is not implied by  $c'_{p_i}$ .

The desired properties now follow from ( $\dagger$ ) and ( $\ddagger$ ) as well as the conditions in Lemma 6.9. In particular,  $\varphi_{\mathcal{A}}$  respecting the constraints of  $D$  can now be derived from Conditions (1, 2, 4-7) in Lemma 6.9, and the last condition of Lemma 6.10 follows from Condition (8) in Lemma 6.9.  $\square$

*Proof of Theorem 6.1.* A NEXSPACE algorithm proving the correctness of Theorem 6.1 is now immediate by Lemma 6.9, as we can iteratively guess and verify quintuples in  $E$  while only keeping track of the very first quintuple and the previous quintuple. Since in an acyclic schema graph the number of paths starting in a given type is at most exponential in the total number of types, each context is defined over at most an exponential number of paths. However, to formally argue that these contexts can be encoded in exponential space, we still need to show that each tuple or variable used in a context can be encoded in at most exponential space. Since the only tuples used are those mentioned in  $c_{o_1}$  and  $c_{p_1}$ , and since we can reuse the same variables over all contexts, both the maximal number of tuples and the maximal number of variables needed are exponential in the total number of types.  $\square$

**6.1. Lowering complexity.** Next, we consider restrictions that lower the complexity. To this end, we say that two variables  $X$  and  $Y$  occurring in a transaction template  $\tau$  are *equivalent in  $\tau$* , denoted  $X \equiv_{\tau} Y$  if

- $X = Y$ ;
- there exists a pair of variables  $Z$  and  $W$  in  $\tau$  and a sequence of function names  $F$  with  $Z \equiv_{\tau} W$ ,  $Z \xrightarrow{F} \tau X$  and  $W \xrightarrow{F} \tau Y$ ; or

- there exists a variable  $Z$  with  $X \equiv_{\tau} Z$  and  $Y \equiv_{\tau} Z$ .

Then, a transaction template  $\tau$  is *restricted* if for every combination of variables  $X, Y, W, Z$  in  $\tau$  with  $X \rightsquigarrow_{\tau} W$  and  $Y \rightsquigarrow_{\tau} Z$ , either  $W \equiv_{\tau} Z$ ,  $W \rightsquigarrow_{\tau} Z$  or  $Z \rightsquigarrow_{\tau} W$ . We denote by **AcycResTemp** the class of all sets of restricted transaction templates over acyclic schemas.

**Theorem 6.11.** (1)  $T\text{-ROBUSTNESS}(\mathbf{AcycResTemp}, RC)$  is decidable in EXPTIME.  
 (2)  $T\text{-ROBUSTNESS}(\mathbf{AcycTemp}, RC)$  is decidable in PSPACE when the number of paths between any two nodes in the schema graph is bounded by a constant  $k$ .

Regarding (1), all templates in TPC-C with the exception of NewOrder are restricted. Regarding (2), when the schema graph is a multi-tree then  $k = 1$  and for TPC-C  $k = 2$  (recall that in general there can be an exponential number of paths), leading to a more practical algorithm for robustness in those cases.

The PSPACE result in Theorem 6.11 for workloads over a schema (Rels, Funcs) where the number of paths between any two nodes in the schema graph is bounded by a constant  $k$  is immediate by the nondeterministic algorithm based on Lemma 6.9 presented for Theorem 6.1. Indeed, in this case, the total number of paths starting in a given type is at most  $k \cdot |\text{Rels}|$  and therefore each context is defined over at most a polynomial number of paths, instead of an exponential number of paths for the general case in Theorem 6.1.

The EXPTIME result in Theorem 6.11 for workloads in **AcycResTemp** follows from a deterministic algorithm based on Lemma 6.9. In the remainder of this section, we first present the algorithm, and then discuss its complexity.

**A deterministic algorithm.** Towards a deterministic algorithm, assume the first quintuple  $(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$  of  $E$  is fixed. We now translate the problem of deciding whether we can extend  $E$  such that it satisfies all properties to a graph problem over a graph  $G(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$ . This graph is constructed as follows:

- each quintuple  $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$  satisfying Conditions (2-8) of Lemma 6.9 is added as a node to  $G(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$ ; and
- there is an edge from a node  $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$  to a node  $(\tau_j, o_j, c_{o_j}, p_j, c_{p_j})$  if  $o_i$  is potentially conflicting with  $p_j$  and  $c_{o_i} = c_{p_j}$  (c.f. Condition (9) of Lemma 6.9).

By construction, it is now easy to see that there is a sequence  $E$  satisfying Lemma 6.9 if there is a path from a quintuple  $(\tau_2, o_2, c_{o_2}, p_2, c_{p_2})$  to a quintuple  $(\tau_m, o_m, c_{o_m}, p_m, c_{p_m})$  in  $G(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$  (where we allow a zero-length path with  $2 = m$ ), such that ( $\dagger$ )

- $c_{o_1} = c_{p_2}$  and  $c_{o_m} = c_{p_1}$  (c.f. Condition (9) of Lemma 6.9);
- $o_1$  is potentially rw-conflicting with  $p_2$  (c.f. Condition (10) of Lemma 6.9); and
- $o_1 <_{\tau_1} p_1$  or  $o_m$  is potentially rw-conflicting with  $p_1$  (c.f. Condition (11) of Lemma 6.9).

Given a set of transaction templates  $\mathcal{P}$  over a schema (Rels, Funcs), the algorithm iterates over all possible quintuples  $(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$  satisfying Condition (1, 3-7) of Lemma 6.9, where we consider all possible tuple-contexts  $c_{o_1}$  and  $c_{p_1}$  up to isomorphisms. For each such quintuple, the graph  $G(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$  is constructed. Let  $TC$  be the reflexive-transitive closure of  $G$ . If there is a pair of quintuples  $(\tau_2, o_2, c_{o_2}, p_2, c_{p_2})$  and  $(\tau_m, o_m, c_{o_m}, p_m, c_{p_m})$  in  $TC$  satisfying ( $\dagger$ ), the algorithm emits a reject, indicating that  $\mathcal{P}$  is not robust against RC. Otherwise, it proceeds with a new choice for  $(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$ . If, the algorithm didn't reject after considering all such quintuples, it accepts, indicating that  $\mathcal{P}$  is indeed robust against RC. The correctness of this algorithm is immediate by Lemma 6.9.

**Complexity analysis.** We show the complexity of the presented algorithm. For this, first, notice that we have defined contexts  $c$  based on a type  $S$  (with  $S$  not necessarily a root of  $SG$ ). For encoding purposes it makes sense to encode these as contexts for a root type  $R$  in combination with the intended type  $S$ . The context as defined in the previous section can then be derived by taking the left-most subtree with root  $S$ . Notice that this is purely an encoding choice that will simplify the analysis.

For a schema graph  $SG(\text{Rels}, \text{Funcs})$  the total number of non-isomorphic tuple-contexts can be expressed using Bell's number  $B(n)$ , denoting the number of partitions for a set of size  $n$ , and the set  $\text{Path}_{SG}(R, S) = \{(R, F, S) \mid R \xrightarrow{F} SG S\}$  expressing the different paths from one node  $R$  to another node  $S$  in  $SG$ . Concretely,

$$|\text{TupleContexts}(SG)| \leq \sum_{R \in \text{roots}(SG)} \prod_{S \in \text{Rels}} B(|\text{Path}_{SG}(R, S)|) = B^*$$

where  $\text{TupleContexts}(SG)$  denotes the set containing all different tuple-contexts (up to isomorphisms).

Now let  $c_1$  and  $c_2$  be two fixed contexts for types that are descendants of roots  $R_1$  and root  $R_2$ , respectively, in  $SG$ , and let  $c$  be a context for a type descending from root  $R$ . To express a bound on the number of substitutions in  $c$  from (parts of)  $c_1$  and  $c_2$ , we need some additional terminology: Let  $\text{Path}_{SG}(R, *) = \bigcup_{S \in \text{Rels}} \text{Path}_{SG}(R, S)$ . We say that a path  $R \xrightarrow{F} SG S$  is a *prefix* of a path  $R \xrightarrow{F'} SG S'$  in  $SG$  if there is a (possibly empty) sequence of function names  $F_2$  with  $F = F_1 \cdot F_2$ . The number of substitutions in  $c$  from (parts of)  $c_1$  and  $c_2$  is now bounded by

$$\begin{aligned} \sum_{\text{Part} \subseteq \text{Path}_{SG}(R, *)} \mathbf{1}_{PFP} \prod_{R \xrightarrow{F} SG S \in \text{Part}} (|\text{Path}_{SG}(R_1, S)| + |\text{Path}_{SG}(R_2, S)|) &\leq T^\ell \cdot (2P)^\ell \\ &\leq (2TP)^\ell \end{aligned}$$

In the above expression,  $\mathbf{1}_{PFP}$  is an indicator variable that equals 1 if no path in  $\text{Part}$  is a prefix of another path in  $\text{Part}$  and that equals 0 otherwise. Further:  $P$  denotes the maximum number of different paths between a particular root and a particular node in  $SG$ ,  $T$  denotes the maximum number of different paths from a particular root to nodes in  $SG$ , and  $\ell$  denotes the maximal size of a set in which no path is a prefix of another path in the set. The latter is trivially bounded by  $T$ .

A special cases exists if all templates  $\tau$  in  $\mathcal{P}$  are restricted. In that case, the size of sets  $\text{Part}$  is bounded by 2, hence  $\ell \leq 2$ .

With the above bounds, the complexity of the presented algorithm is rather straightforward. The iteration over all possible quintuples  $(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$  requires at most  $|\mathcal{P}| \cdot t^2 \cdot (B^*)^2$  iterations, with  $t$  denoting the maximal number of operations in a transaction template of  $\mathcal{P}$ . The remainder of the computation is dominated by the transitive-closure computation. Since the constructed graph  $G(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$  has at most

$$|\mathcal{P}| \cdot t^2 \cdot (B^* \cdot (2TP)^\ell)^2$$

nodes, the transitive closure computation requires

$$\left( |\mathcal{P}| \cdot t^2 \cdot (B^* \cdot (2TP)^\ell)^2 \right)^3$$

steps. Putting these numbers together, we obtain:

$$\mathcal{O}(|\mathcal{P}|^4 \cdot t^8 \cdot (B^*)^8 \cdot (2TP)^{6\ell}).$$

Since  $\ell$  is bounded by a constant if all template are restricted, and since  $B^*$ ,  $T$  and  $P$  can be exponential in the size of the input, the presented algorithm indeed decides  $\text{T-ROBUSTNESS}(\mathbf{AcycResTemp}, \text{RC})$  in EXPTIME.

## 7. RELATED WORK

**Transaction Programs.** Previous work on static robustness testing [FLO<sup>+</sup>05, AF15] for transaction programs is based on the following key insight: when a *schedule* is not serializable, then the dependency graph constructed from that schedule contains a cycle satisfying a condition specific to the isolation level at hand (*dangerous structure* for SNAPSHOT ISOLATION and the presence of a *counterflow edge* for RC). That insight is extended to a workload of *transaction programs* through the construction of a so-called static dependency graph where each program is represented by a node, and there is a conflict edge from one program to another if there can be a schedule that gives rise to that conflict. The absence of a cycle satisfying the condition specific to that isolation level then guarantees robustness while the presence of a cycle does not necessarily imply non-robustness.

Other work studies robustness within a framework for uniformly specifying different isolation levels in a declarative way [CBG15, BG16, CG18]. A key assumption here is *atomic visibility* requiring that either all or none of the updates of each transaction are visible to other transactions. These approaches aim at higher isolation levels and cannot be used for RC, as RC does not admit *atomic visibility*.

**Transaction Templates.** The static robustness approach based on transaction templates [VKKN21] differs in two ways. First, it makes more underlying assumptions explicit within the formalism of transaction templates (whereas previous work departs from the static dependency graph that should be constructed in some way by the dba). Second, it allows for a decision procedure that is sound and complete for robustness testing against RC, allowing to detect larger subsets of transactions to be robust [VKKN21].

The formalization of transactions and conflict serializability in [VKKN21] and this paper is based on [Fek05], generalized to operations over attributes of tuples and extended with U-operations that combine R- and W-operations into one atomic operation. These definitions are closely related to the formalization presented by Adya et al. [ALO00], but we assume a total rather than a partial order over the operations in a schedule. There are also a few restrictions to the model: there needs to be a fixed set of read-only attributes that cannot be updated and which are used to select tuples for update. The most typical example of this are primary key values passed to transaction templates as parameters. The inability to update primary keys is not an important restriction in many workloads, where keys, once assigned, never get changed, for regulatory or data integrity reasons.

In [VKKN21], a PTIME decision procedure is obtained for robustness against RC for templates without functional constraints and the present paper improves that result to NLOGSPACE. In addition, an experimental study was performed showing how an approach based on robustness and making transactions robust through promotion can improve transaction throughput.

**Transactions.** The work by Fekete [Fek05] is the first work that provides a necessary and sufficient condition for deciding robustness against SNAPSHOT ISOLATION for a workload of concrete transactions (not transaction programs). That work provides a characterization for acceptable allocations when every transaction runs under either SNAPSHOT ISOLATION or strict two-phase locking (S2PL). The allocation then is acceptable when every possible execution respecting the allocated isolation levels is serializable. As a side result, this work indirectly provides a necessary and sufficient condition for robustness against SNAPSHOT ISOLATION, since robustness against SNAPSHOT ISOLATION holds iff the allocation where each transaction is allocated to SNAPSHOT ISOLATION is acceptable. Ketsman et al. [KKNV20] provide full characterizations for robustness against READ COMMITTED and READ UNCOMMITTED under lock-based semantics. In addition, it is shown that the corresponding decision problems are complete for CONP and LOGSPACE, respectively, which should be contrasted with the polynomial time characterization obtained in [VKKN21] for robustness against *multiversion* read committed.

## 8. CONCLUSION

This paper falls within a more general research line investigating how transaction throughput can be improved through an approach based on robustness testing that can be readily applied without making any changes to the underlying database system. As argued in Section 2, incorporating functional constraints can detect larger sets of templates to be robust and requires less R-operations to be promoted to U-operations. In future work, we plan to look at lower bounds, restrictions that lower complexity, and consider other referential integrity constraints to further enlarge the modelling power of transaction templates. For example, the current formalism allows to express dependencies between tuples, but it is not suited to express the fact that a transaction accesses *all* tuples depending on a specific tuple, such as all *Order* tuples depending on a specific *Customer* tuple.

## ACKNOWLEDGMENT

This work is partly funded by FWO-grant G019921N.

## REFERENCES

- [ACFR08] Mohammad Alomari, Michael Cahill, Alan Fekete, and Uwe Rohm. The cost of serializability on platforms that use snapshot isolation. In *ICDE*, pages 576–585, 2008.
- [AF15] Mohammad Alomari and Alan Fekete. Serializable use of read committed isolation level. In *AICCSA*, pages 1–8, 2015.
- [ALO00] Atul Adya, Barbara Liskov, and Patrick E. O’Neil. Generalized isolation level definitions. In *ICDE*, pages 67–78, 2000.
- [BBE19a] Sidi Mohamed Beillahi, Ahmed Bouajjani, and Constantin Enea. Checking robustness against snapshot isolation. In *CAV*, pages 286–304, 2019.
- [BBE19b] Sidi Mohamed Beillahi, Ahmed Bouajjani, and Constantin Enea. Robustness against transactional causal consistency. In *CONCUR*, pages 1–18, 2019.
- [BBG<sup>+</sup>95] Hal Berenson, Philip A. Bernstein, Jim Gray, Jim Melton, Elizabeth J. O’Neil, and Patrick E. O’Neil. A critique of ANSI SQL isolation levels. In *SIGMOD*, pages 1–10, 1995.
- [BDF<sup>+</sup>13] Peter Bailis, Aaron Davidson, Alan Fekete, Ali Ghodsi, Joseph M. Hellerstein, and Ion Stoica. Highly available transactions: Virtues and limitations. *PVLDB*, 7(3):181–192, 2013.
- [BG16] Giovanni Bernardi and Alexey Gotsman. Robustness against consistency models with atomic visibility. In *CONCUR*, pages 7:1–7:15, 2016.



- [CBG15] Andrea Cerone, Giovanni Bernardi, and Alexey Gotsman. A framework for transactional consistency models with atomic visibility. In *CONCUR*, pages 58–71, 2015.
- [CG18] Andrea Cerone and Alexey Gotsman. Analysing snapshot isolation. *J.ACM*, 65(2):1–41, 2018.
- [CGY17] Andrea Cerone, Alexey Gotsman, and Hongseok Yang. Algebraic Laws for Weak Consistency. In *CONCUR*, pages 26:1–26:18, 2017.
- [CV85] Ashok K. Chandra and Moshe Y. Vardi. The implication problem for functional and inclusion dependencies is undecidable. *SIAM J. Comput.*, 14(3):671–677, 1985.
- [Fek05] Alan Fekete. Allocating isolation levels to transactions. In *PODS*, pages 206–215, 2005.
- [FLO<sup>+</sup>05] Alan Fekete, Dimitrios Liarokapis, Elizabeth J. O’Neil, Patrick E. O’Neil, and Dennis E. Shasha. Making snapshot isolation serializable. *ACM Trans. Database Syst.*, 30(2):492–528, 2005.
- [KKNV20] Bas Ketsman, Christoph Koch, Frank Neven, and Brecht Vandevoort. Deciding robustness for lower SQL isolation levels. In *PODS*, pages 315–330, 2020.
- [Pap86] Christos H. Papadimitriou. *The Theory of Database Concurrency Control*. Computer Science Press, 1986.
- [Pos46] Emil L. Post. A variant of a recursively unsolvable problem. *Bull. Amer. Math. Soc.*, pages 264–268, 1946.
- [TC] TPC-C. On-line transaction processing benchmark. <http://www.tpc.org/tpcc/>.
- [VKKN21] Brecht Vandevoort, Bas Ketsman, Christoph Koch, and Frank Neven. Robustness against read committed for transaction templates. *PVLDB*, 14(11):2141–2153, 2021.
- [VKKN22] Brecht Vandevoort, Bas Ketsman, Christoph Koch, and Frank Neven. Robustness against read committed for transaction templates with functional constraints. In *ICDT*, pages 16:1–16:17, 2022.

## APPENDIX A. SMALLBANK BENCHMARK SQL CODE

```

Balance(N):
    SELECT CustomerId INTO :x
      FROM Account
     WHERE Name=:N;

    SELECT Balance INTO :a
      FROM Savings
     WHERE CustomerId=:x;

    SELECT Balance + :a
      FROM Checking
     WHERE CustomerId=:x;
    COMMIT;

TransactSavings(N,V):
    SELECT CustomerId INTO :x
      FROM Account
     WHERE Name=:N;

    UPDATE Savings
      SET Balance = Balance + :V
     WHERE CustomerId=:x;
    COMMIT;

WriteCheck(N,V):
    SELECT CustomerId INTO :x
      FROM Account
     WHERE Name=:N;

    SELECT Balance INTO :a
      FROM Savings
     WHERE CustomerId=:x;

    SELECT Balance INTO :b
      FROM Checking
     WHERE CustomerId=:x;

    IF (:a + :b) < :V THEN
      UPDATE Checking
        SET Balance = Balance - (:V+1)
       WHERE CustomerId=:x;
    ELSE
      UPDATE Checking
        SET Balance = Balance - :V
       WHERE CustomerId=:x;
    END IF;
    COMMIT;

Amalgamate(N1,N2):
    SELECT CustomerId INTO :x1
      FROM Account
     WHERE Name=:N1;

    SELECT CustomerId INTO :x2
      FROM Account
     WHERE Name=:N2;

    UPDATE Savings AS new
      SET Balance = 0
      FROM Savings AS old
     WHERE new.CustomerId=:x1
           AND old.CustomerId
           = new.CustomerId
    RETURNING old.Balance INTO :a;

    UPDATE Checking AS new
      SET Balance = 0
      FROM Checking AS old
     WHERE new.CustomerId=:x1
           AND old.CustomerId
           = new.CustomerId
    RETURNING old.Balance INTO :b;

    UPDATE Checking
      SET Balance = Balance + :a + :b
     WHERE CustomerId=:x2;

GoPremium(N):
    UPDATE Account
      SET IsPremium = TRUE
     WHERE Name=:N
    RETURNING CustomerId INTO :x;

    SELECT InterestRate INTO :a
      FROM Savings
     WHERE CustomerId=:x;

    :rate = computePremiumRate(:x,:a);

    UPDATE Savings
      SET InterestRate = :rate
     WHERE CustomerId=:x;
    COMMIT;

DepositChecking(N,V):
    SELECT CustomerId INTO :x
      FROM Account
     WHERE Name=:N;

    UPDATE Checking
      SET Balance = Balance + :V
     WHERE CustomerId=:x;
    COMMIT;

```