

1 Robustness against Read Committed for 2 Transaction Templates with Functional 3 Constraints

4 **Brecht Vandervoort** ✉

5 Hasselt University and Transnational University of Limburg, Belgium

6 **Bas Ketsman** ✉

7 Vrije Universiteit Brussel, Belgium

8 **Christoph Koch** ✉

9 École Polytechnique Fédérale de Lausanne, Switzerland

10 **Frank Neven** ✉

11 Hasselt University and Transnational University of Limburg, Belgium

12 — Abstract —

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

26 **2012 ACM Subject Classification** Information systems → Database transaction processing

27 **Keywords and phrases** concurrency control, robustness, complexity

28 **Digital Object Identifier** 10.4230/LIPIcs.CVIT.2016.23

29 **1** Introduction

30 Many database systems implement several isolation levels, allowing users to trade isolation
31 guarantees for improved performance. The highest, serializability, projects the appearance
32 of a complete absence of concurrency, and thus perfect isolation. Executing transactions
33 concurrently under weaker isolation levels can introduce certain anomalies. Sometimes, a
34 transactional workload can be executed at an isolation level lower than serializability without
35 introducing any anomalies. This is a desirable scenario: a lower isolation level, usually
36 implementable with a cheaper concurrency control algorithm, yields the stronger isolation
37 guarantees of serializability for free. This formal property is called robustness [12, 7]: a set
38 of transactions \mathcal{T} is called *robust against a given isolation level* if every possible interleaving
39 of the transactions in \mathcal{T} that is allowed under the specified isolation level is serializable.

40 Robustness received quite a bit of attention in the literature. Most existing work focuses
41 on Snapshot Isolation (SI) [2, 4, 12, 13] or higher isolation levels [5, 7, 8, 10]. It is particularly
42 interesting to consider robustness against lower level isolation levels like multi-version Read
43 Committed (referred to as RC from now on). Indeed, RC is widely available, often the default



© Brecht Vandervoort, Bas Ketsman, Christoph Koch, and Frank Neven;
licensed under Creative Commons License CC-BY 4.0

42nd Conference on Very Important Topics (CVIT 2016).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:47

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

44 in database systems (see, e.g., [4]), and is generally expected to have better throughput than
45 stronger isolation levels.


46 In previous work [18], we provided a tractable decision procedure for robustness against
47 RC for workloads generated by transaction programs modeled as transaction templates. The
48 approach is centered on a novel characterization of robustness against RC in the spirit of
49 [12, 14] that improves over the sufficient condition presented in [3], and on a formalization
50 of transaction programs, called *transaction templates*, facilitating fine-grained reasoning
51 for robustness against RC. Conceptually, transaction templates as introduced in [18] are
52 functions with parameters, and can, for instance, be derived from stored procedures inside a
53 database system (c.f. Figure 1 for an example). The abstraction generalizes transactions as
54 usually studied in concurrency control research – sequences of read and write operations – by
55 making the objects worked on variable, determined by input parameters. Such parameters
56 are *typed* to add additional power to the analysis. They support *atomic updates* (that is,
57 a read followed by a write of the same database object, to make a relative change to its
58 value). Furthermore, database objects read and written are considered at the granularity of
59 fields, rather than just entire tuples, decoupling conflicts further and allowing to recognize
60 additional cases that would not be recognizable as robust on the tuple level.

61 An important insight obtained from [18] is that more accurate modeling of the workload
62 allows to recognize larger sets of transaction programs as robust. Processing workloads
63 under RC increases the throughput of the transactional database system compared to
64 when executing the workload under SI or serializable SI, so larger robust sets mean better
65 performance of the database system. In this work, we increase the modeling power of
66 transaction templates by extending them with *functional constraints*, which are useful for
67 capturing data dependencies like foreign keys (inclusion dependencies). This appears to be
68 a sweet spot for strengthening modelling power – as we show in this paper, it allows us to
69 remain with abstractions that have been well established within database theory, without
70 having to move to general program analysis, and it pushes the robustness frontier on popular
71 transaction processing benchmarks. Generally speaking, workloads can profit more from
72 richer modelling the larger and more complex they get, so the fact that adding functional
73 constraints yields larger robust sets already on these simple benchmarks suggests that these
74 techniques are practically useful. Our contributions can be summarized as follows:

- 75 ■ We argue in Section 2 through the SmallBank and TPC-C benchmarks that the incor-
76 poration of functional constraints can identify more workloads as robust that otherwise
77 would not be, and that they reduce the extent to which changes need to be made to
78 workloads to make them robust against RC.
- 79 ■ In Section 4, we establish that robustness in its most general form becomes undecidable.
80 The proof is a reduction from PCP and relies on cyclic dependencies between functions
81 allowing to connect data values through an unbounded application of functions.
- 82 ■ We consider a fragment in Section 5 that only allows a very limited form of cyclic
83 dependencies between functions and assumes additional constraints on templates that,
84 together, imply that functions behave as bijections. Robustness against RC can be decided
85 in NLOGSPACE and this fragment is general enough to model the SmallBank benchmark.
- 86 ■ In Section 6, we obtain an EXPSpace decision procedure when the schema graph is acyclic
87 (so, no cyclic dependencies between functions). Even for small input sizes, such a result
88 is not practical. We provide various restrictions that lower the complexity to PSPACE and
89 EXPTIME, and which allow to model the TPC-C benchmark as discussed. Notice that, for
90 robustness testing, an exponential time decision procedure is considered to be practical as
91 the size of the input is small and robustness is a static property that can be tested offline.

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 2** Multiversion split schedule.

■ **Figure 1** Transaction template.

92 Due to space constraints, proofs as well as a more complete description of the SmallBank
 93 and TPC-C benchmarks are moved to the appendix.

94 2 Application

95 We present a small extension of the SmallBank benchmark [2] to exemplify the modeling
 96 power of transaction templates and discuss how the addition of functional constraints can
 97 detect larger sets of transaction templates to be robust. Finally, we discuss in the context of
 98 the TPC-C benchmark how the incorporation of functional constraints requires less changes
 99 to templates in making them robust. A full description of these benchmarks is in Appendix D.

100 The SmallBank schema consists of three tables: `Account(`Name`,` `CustomerID,` `IsPremium)`,
 101 `Savings(`CustomerID`,` `Balance,` `InterestRate)`, and `Checking(`CustomerID`,` `Balance)`. Under-
 102 lined attributes are primary keys. The `Account` table associates customer names with IDs
 103 and keeps track of the premium status (Boolean); `CustomerID` is a `UNIQUE` attribute. The
 104 other tables contain the balance (numeric value) of the savings and checking accounts of
 105 customers identified by their ID. `Account (CustomerID)` is a foreign key referencing both the
 106 columns `Savings (CustomerID)` and `Checking (CustomerID)`. The interest rate on a savings
 107 account is based on a number of parameters, including the account status (premium or not).
 108 The application code can interact with the database through a fixed number of transaction
 109 programs presented in Appendix D.1: `Balance`, `TransactSavings`, `Amalgamate`, `WriteCheck`,
 110 `DepositChecking`, and `GoPremium`. We only discuss `GoPremium(N)`, given in Figure 1,
 111 which converts the account of the customer with name N to a premium account and updates
 112 the interest rate of the corresponding savings account.

113 In short, a transaction template is a sequence of read (R), write (W) and update (U)
 114 statements over typed variables (X, Y, \dots) with additional equality and disequality constraints.
 115 For instance, $R[Y : \text{Savings}\{C, I\}]$ indicates that a read operation is performed to a tuple in
 116 relation `Savings` on the attributes `CustomerID` and `InterestRate`. We abbreviate the names of
 117 attributes by their first letter to save space. The set $\{C, I\}$ is the read set. Write operations
 118 have an associated write set while update operations contain a read set followed by a write
 119 set: e.g., $U[Z : \text{Account}\{N, C\}\{I\}]$ first reads the `Name` and `CustomerID` of tuple X and then
 120 writes to the attribute `InterestRate`. To capture the dependencies between tuples induced
 121 by the foreign keys, we use two unary functions: $f_{A \rightarrow S}$ maps a tuple of type `Account` to a
 122 tuple of type `Savings`, while $f_{A \rightarrow C}$ maps a tuple of type `Account` to a tuple of type `Checking`.
 123 As `Account(CustomerID)` is `UNIQUE`, every savings and checking accounts is associated to
 124 a unique `Account` tuple. This is modelled through the functions $f_{C \rightarrow A}$ and $f_{S \rightarrow A}$ with an
 125 analogous interpretation. Notice that the equality constraints for `GoPremium` imply that
 126 these functions are bijections and each others inverses.

127 A transaction T over a database \mathbf{D} is an *instantiation* of a transaction template τ if
 128 there is a variable mapping μ from the variables in τ to tuples in \mathbf{D} that satisfies all the
 129 constraints in τ such that with $\mu(\tau) = T$. For instance, consider a database \mathbf{D} with tuples
 130 a_1, a_2, \dots of type `Account`, s_1, s_2, \dots of type `Savings`, and c_1, c_2, \dots of type `Checking` with

Delivery: $U[S : \text{Order}\{W, D, O\}\{\text{Sta}\}]$ $U[V_1 : \text{OrderLine}\{W, D, O, OL, Del\}\{\text{Del}\}]$ $U[V_2 : \text{OrderLine}\{W, D, O, OL, Del\}\{\text{Del}\}]$ $U[Z : \text{Customer}\{W, D, C, Bal\}\{\text{Bal}\}]$ $Z = f_{O \rightarrow C}(S), S = f_{L \rightarrow O}(V_1), S = f_{L \rightarrow O}(V_2)$	OrderStatus: $R[Z : \text{Customer}\{W, D, C, Inf, Bal\}]$ $R[S : \text{Order}\{W, D, O, C, Sta\}]$ $R[V_1 : \text{OrderLine}\{W, D, O, OL, I, Del, Qua\}]$ $R[V_2 : \text{OrderLine}\{W, D, O, OL, I, Del, Qua\}]$ $Z = f_{O \rightarrow C}(S), S = f_{L \rightarrow O}(V_1), S = f_{L \rightarrow O}(V_2)$
---	---

■ **Figure 3** Transaction templates Delivery and OrderStatus of the TPC-C benchmark.

131 $f_{A \rightarrow S}^D(a_i) = s_i, f_{A \rightarrow C}^D(a_i) = c_i, f_{S \rightarrow A}^D(s_i) = a_i, f_{C \rightarrow A}^D(c_i) = a_i$ for each i . Then, for
 132 $\mu_1 = \{X \rightarrow a_1, Y \rightarrow s_1\}, \mu_1(\text{GoPremium}) = U[a_1]R[s_1]U[c_1]$ is an instantiation of GoPremium
 133 whereas $\mu_2(\text{GoPremium})$ with $\mu_2 = \{X \rightarrow a_1, Y \rightarrow s_2\}$ is not as the functional constraint
 134 $Y = f_{A \rightarrow S}(X)$ is not satisfied. Indeed, $\mu_2(Y) = s_2 \neq s_1 = f_{A \rightarrow S}^D(a_1) = f_{A \rightarrow S}^D(\mu_2(X))$. We then
 135 say that a set of transactions is *consistent* with a set of templates if every transaction is an
 136 instantiation of a transaction template.

137 Our previous work [18], which did not consider functional constraints, has shown that
 138 $\{\text{Am,DC,TS}\}, \{\text{Bal,DC}\},$ and $\{\text{Bal,TS}\}$ are maximal robust sets of transaction templates.
 139 This means that for any database, for any set of transactions \mathcal{T} that is consistent with
 140 one of the three mentioned sets, any possible interleaving of the transactions in \mathcal{T} that is
 141 allowed under RC is *always* serializable! Using the results from Section 5, it follows that
 142 when functional constraints are taken into account GoPremium can be added to each of
 143 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 its 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 μ_1 and μ_2 as defined above for respectively T_1 and T_2 (we show the read and write sets to facilitate the discussion):

$$\begin{array}{ll}
 T_1 : U_1[a_1\{N,C\}\{I\}]R_1[s_1\{C,I\}] & U_1[s_1\{C\}\{I\}]C_1 \\
 T_2 : & U_2[a_2\{N,C\}\{I\}]R_2[s_1\{C,I\}]U_2[s_1\{C\}\{I\}]C_2
 \end{array}$$

144 The above schedule is allowed under RC as there is no dirty write, but it is not conflict
 145 serializable. Indeed, there is a rw-conflict between $R_1[s_1\{C,I\}]$ and $U_2[s_1\{C\}\{I\}]$ as the former
 146 reads the attribute I that is written to by the latter, which implies that T_1 should occur before
 147 T_2 in an equivalent serial schedule. But, there is a ww-conflict between $U_2[s_1\{C\}\{I\}]$ and
 148 $U_1[s_1\{C\}\{I\}]$ as both write to the common attribute I implying that T_2 should occur before
 149 T_1 in an equivalent serial schedule. Consequently, the schedule is not serializable. However,
 150 taking functional constraints into account, $\{T_1, T_2\}$ is not consistent with $\{\text{GoPremium}\}$
 151 as $\mu_2(Y) = s_1 \neq s_2 = f_{A \rightarrow S}(a_2) = f_{A \rightarrow S}(\mu_2(X))$ implying that the above schedule is *not* a
 152 counter example for robustness.

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 [18]. 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 $\mathcal{P} = \{\text{Delivery}, \text{OrderStatus}\}$ of the TPC-C benchmark, given in Figure 3, where functional constraints are added to express the fact that a tuple of type OrderLine implies the tuple of type Order, which in turn implies the tuple of type Customer. 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:

$$\begin{array}{l} T_1(\text{Orderstatus}) : \text{U}_1[\text{c}] \text{R}_1[\text{a}] \qquad \qquad \qquad \text{R}_1[\text{b}_1] \text{R}_1[\text{b}_2] \text{C}_1 \\ T_2(\text{Delivery}) : \qquad \qquad \qquad \text{U}_2[\text{a}] \text{U}_2[\text{b}_1] \text{U}_2[\text{b}_2] \text{U}_2[\text{c}'] \text{C}_2 \end{array}$$

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

156 3 Definitions

157 We recall the necessary definitions from [18] and extend them with functional constraints.

158 3.1 Databases

159 A *relational schema* is a pair $(\text{Rels}, \text{Funcs})$ where Rels is a set of relation names and Funcs is a
 160 set of function names. A finite set of attribute names $\text{Attr}(R)$ is associated to every relation
 161 $R \in \text{Rels}$. Relations will be instantiated by abstract objects that serve as an abstraction of
 162 relational tuples. To this end, for every relation $R \in \text{Rels}$, we fix an infinite set of tuples
 163 \mathbf{Tuples}_R . Furthermore, we assume that $\mathbf{Tuples}_R \cap \mathbf{Tuples}_S = \emptyset$ for all $R, S \in \text{Rels}$ with
 164 $R \neq S$. We then denote by \mathbf{Tuples} the set $\bigcup_{R \in \text{Rels}} \mathbf{Tuples}_R$ of all possible tuples. Notice
 165 that, by definition, for every $\mathbf{t} \in \mathbf{Tuples}$ there is a unique relation $R \in \text{Rels}$ such that
 166 $\mathbf{t} \in \mathbf{Tuples}_R$. In that case, we say that \mathbf{t} is of *type* R and denote the latter by $\text{type}(\mathbf{t}) = R$.
 167 Each function name $f \in \text{Funcs}$ has a domain $\text{dom}(f) \in \text{Rels}$ and a range $\text{range}(f) \in \text{Rels}$.
 168 Functions are used to encode relationships between tuples like for instance those implied by
 169 foreign-keys constraints. For instance, in the SmallBank example $\text{Funcs} = \{f_{A \rightarrow S}, f_{A \rightarrow C}\}$,
 170 $\text{dom}(f_{A \rightarrow S}) = \text{dom}(f_{A \rightarrow C}) = A$, $\text{range}(f_{A \rightarrow S}) = S$, and $\text{range}(f_{A \rightarrow C}) = C$. A *database* \mathbf{D}
 171 over schema $(\text{Rels}, \text{Funcs})$ assigns to every relation name $R \in \text{Rels}$ a finite set $R^{\mathbf{D}} \subset \mathbf{Tuples}_R$
 172 and to every function name $f \in \text{Funcs}$ a function $f^{\mathbf{D}}$ from $\text{dom}(f)^{\mathbf{D}}$ to $\text{range}(f)^{\mathbf{D}}$.

173 3.2 Transactions and Schedules

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

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

189 When considering a set \mathcal{T} of transactions, we assume that every transaction in the set has
 190 a unique id i and write T_i to make this id explicit. Similarly, to distinguish the operations
 191 of different transactions, we add this id as a subscript to the operation. That is, we write

192 $W_i[t]$, $R_i[t]$, and $U_i[t]$ to denote a $W[t]$, $R[t]$, and $U[t]$ occurring in transaction T_i ; similarly
 193 C_i denotes the commit operation in transaction T_i . This convention is consistent with the
 194 literature (see, e.g. [6, 12]). To avoid ambiguity of notation, we assume that a transaction
 195 performs at most one write, one read, and one update per tuple. The latter is a common
 196 assumption (see, e.g. [12]). All our results carry over to the more general setting in which
 197 multiple writes and reads per tuple are allowed.

198 A (*multiversion*) *schedule* s over a set \mathcal{T} of transactions is a tuple $(O_s, \leq_s, \ll_s, v_s)$ where
 199 O_s is the set containing all operations of transactions in \mathcal{T} as well as a special operation op_0
 200 conceptually writing the initial versions of all existing tuples, \leq_s encodes the ordering of
 201 these operations, \ll_s is a *version order* providing for each tuple \mathbf{t} a total order over all write
 202 operations on \mathbf{t} occurring in s , and v_s is a *version function* mapping each read operation a
 203 in s to either op_0 or to a write¹ operation different from a in s . We require that $op_0 \leq_s a$
 204 for every operation $a \in O_s$, $op_0 \ll_s a$ for every write operation $a \in O_s$, and that $a <_T b$
 205 implies $a <_s b$ for every $T \in \mathcal{T}$ and every $a, b \in T$. We furthermore require that for every
 206 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
 207 tuple as a . Intuitively, op_0 indicates the start of the schedule, the order of operations in s is
 208 consistent with the order of operations in every transaction $T \in \mathcal{T}$, and the version function
 209 maps each read operation a to the operation that wrote the version observed by a . If $v_s(a)$
 210 is op_0 , then a observes the initial version of this tuple. The version order \ll_s represents the
 211 order in which different versions of a tuple are installed in the database. For a pair of write
 212 operations on the same tuple, this version order does not necessarily coincide with \leq_s . For
 213 example, under RC the version order is based on the commit order instead.

214 We say that a schedule s is a *single version schedule* if \ll_s coincides with \leq_s and every
 215 read operation always reads the last written version of the tuple. Formally, for each pair of
 216 write operations a and b on the same tuple, $a \ll_s b$ iff $a <_s b$, and for every read operation
 217 a there is no write operation c on the same tuple as a with $v_s(a) <_s c <_s a$. A single
 218 version schedule over a set of transactions \mathcal{T} is *single version serial* if its transactions are
 219 not interleaved with operations from other transactions. That is, for every $a, b, c \in O_s$ with
 220 $a <_s b <_s c$ and $a, c \in T$ implies $b \in T$ for every $T \in \mathcal{T}$.

221 The absence of aborts in our definition of schedule is consistent with the common
 222 assumption [12, 7] that an underlying recovery mechanism will rollback aborted transactions.
 223 We only consider isolation levels that only read committed versions. Therefore there will
 224 never be cascading aborts.

225 3.3 Conflict Serializability

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

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

231 In this case, we also say that a_j and b_i are conflicting operations. Furthermore, commit
 232 operations and the special operation op_0 never conflict with any other operation. When a_j
 233 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} ,
 234 denoted $b_i \rightarrow_s a_j$ if:²

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

² Throughout the paper, we adopt the following convention: a b operation can be understood as a ‘before’

- 235 ■ (*ww-dependency*) b_i is ww-conflicting with a_j and $b_i \ll_s a_j$; or,
- 236 ■ (*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,
- 237 ■ (*rw-antidependency*) b_i is rw-conflicting with a_j and $v_s(b_i) \ll_s a_j$.

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

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

245 ► **Definition 1.** A schedule s is conflict serializable if it is conflict equivalent to a single
 246 version serial schedule.

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

250 ► **Theorem 2** ([15]). A schedule s is conflict serializable iff the conflict graph for s is acyclic.

251 3.4 Multiversion Read Committed

252 Let s be a schedule for a set \mathcal{T} of transactions. Then, s exhibits a *dirty write* iff there are
 253 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
 254 $T_j \neq T_i$ such that $b_i <_s a_j <_s C_i$. That is, transaction T_j writes to an attribute of a tuple
 255 that has been modified earlier by T_i , but T_i has not yet issued a commit.

256 For a schedule s , the version order \ll_s corresponds to the commit order in s if for every
 257 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
 258 s is *read-last-committed (RLC)* if \ll_s corresponds to the commit order and for every read
 259 operation a_j in s on some tuple \mathbf{t} the following holds:

- 260 ■ $v_s(a_j) = op_0$ or $C_i <_s a_j$ with $v_s(a_j) \in T_i$; and
- 261 ■ there is no write³ operation $c_k \in T_k$ on \mathbf{t} with $C_k <_s a_j$ and $v_s(a_j) \ll_s c_k$.

262 So, a_j observes the most recent version of \mathbf{t} (according to the order of commits) that is
 263 committed before a_j . Note in particular that a schedule cannot exhibit dirty reads, defined
 264 in the traditional way [6], if it is read-last-committed.

265 ► **Definition 3.** A schedule is allowed under isolation level *read committed (RC)* if it is
 266 read-last-committed and does not exhibit dirty writes.

267 3.5 Transaction Templates

268 Transaction templates are transactions where operations are defined over typed variables
 269 together with functional constraints on these variables. Types of variables are relation names
 270 in **Rels** and indicate that variables can only be instantiated by tuples from the respective
 271 type. We fix an infinite set of variables **Var** that is disjoint from **Tuples**. Every variable
 272 $X \in \mathbf{Var}$ has an associated relation name in **Rels** as type that we denote by $\text{type}(X)$. For an
 273 operation o_i in a template, $\text{var}(o_i)$ denotes the variable in o_i . An *equality constraint* is an

while an a can be interpreted as an ‘after’.

³ Recall that a write operation is either a **W** or a **U**-operation.

274 expression of the form $X = f(Y)$ where $X, Y \in \mathbf{Var}$, $\text{dom}(f) = \text{type}(Y)$ and $\text{range}(f) = \text{type}(X)$.
 275 A *disequality constraint* is an expression of the form $X \neq Y$ where $\text{type}(X) = \text{type}(Y)$.

276 ► **Definition 4.** A transaction template is a transaction τ over \mathbf{Var} together with a set $\Gamma(\tau)$
 277 of equality and disequality constraints. In addition, for every operation o in τ over a variable
 278 X , $\text{ReadSet}(o) \subseteq \text{Attr}(\text{type}(X))$ and $\text{WriteSet}(o) \subseteq \text{Attr}(\text{type}(X))$.

279 Recall that we denote variables by capital letters X, Y, Z and tuples by small letters \mathbf{t}, \mathbf{v} .
 280 A variable assignment μ is a mapping from \mathbf{Var} to \mathbf{Tuples} such that $\mu(X) \in \mathbf{Tuples}_{\text{type}(X)}$.
 281 Furthermore, μ satisfies a constraint $X = f(Y)$ (resp., $X \neq Y$) over a database \mathbf{D} when
 282 $\mu(X) = f^{\mathbf{D}}(\mu(Y))$ (resp., $\mu(X) \neq \mu(Y)$). A variable assignment μ for a transaction template
 283 τ is *admissible* for \mathbf{D} if it satisfies all constraints in $\Gamma(\tau)$ over \mathbf{D} . By $\mu(\tau)$, we denote the
 284 transaction obtained by replacing each variable X in τ with $\mu(X)$.

285 A set of transactions \mathcal{T} is *consistent* with a set of transaction templates \mathcal{P} and database
 286 \mathbf{D} , if for every transaction T in \mathcal{T} there is a transaction template $\tau \in \mathcal{P}$ and a variable
 287 mapping μ_T that is admissible for \mathbf{D} such that $\mu_T(\tau) = T$.

288 3.6 Robustness

289 We define the robustness property [7] (also called *acceptability* in [12, 13]), which guarantees
 290 serializability for all schedules of a given set of transactions for a given isolation level.

291 ► **Definition 5** (Transaction Robustness). A set \mathcal{T} of transactions is *robust against RC* if
 292 every schedule for \mathcal{T} that is allowed under RC is conflict serializable.

293 In the next definition, we represent conflicting operations from transactions in a set \mathcal{T} as
 294 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
 295 transactions in \mathcal{T} . We call these quadruples *conflicting quadruples* for \mathcal{T} . Further, for an
 296 operation $b \in T$, we denote by $\text{prefix}_b(T)$ the restriction of T to all operations that are before
 297 or equal to b according to \leq_T . Similarly, we denote by $\text{postfix}_b(T)$ the restriction of T to all
 298 operations that are strictly after b according to \leq_T . Throughout the paper, we interchangeably
 299 consider transactions both as linear orders as well as sequences. Therefore, T is then equal
 300 to the sequence $\text{prefix}_b(T)$ followed by $\text{postfix}_b(T)$ which we denote by $\text{prefix}_b(T) \cdot \text{postfix}_b(T)$
 301 for every $b \in T$.

► **Definition 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:

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

302 where

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

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

309 Figure 2 depicts a schematic multiversion split schedule. The name stems from the fact
 310 that the schedule is obtained by splitting one transaction in two (T_1 at operation b_1 in

311 Figure 2) and placing all other transactions in C in between. The figure does not display the
 312 trailing transactions T_{m+1}, T_{m+2}, \dots and assumes $b_1 <_{T_1} a_1$.

313 The following theorem characterizes non-robustness in terms of the existence of a mul-
 314 tiversion split schedule.

315 ▶ **Theorem 7** ([18]). *For a set of transactions \mathcal{T} , the following are equivalent:*

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

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

321 ▶ **Definition 8** (Template Robustness). *A set of transaction templates \mathcal{P} is robust against*
 322 *RC if \mathcal{P} is robust against RC for every database \mathbf{D} .*

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

329 ▶ **Theorem 9** ([18]). T-ROBUSTNESS(**VarTemp**, RC) is decidable in PTIME.

330 4 Robustness for Templates


331 We start out with a negative result and show that the robustness problem in its most general
 332 form is undecidable (even when disequalities are not allowed). The proof is a reduction
 333 from *Post's Correspondence Problem (PCP)* [16] and relies on cyclic dependencies between
 334 functional constraints. The proof can be found in Appendix A and is quite elaborate but the
 335 basic intuition is simple: the counterexample split schedule will build up the two strings that
 336 need to be generated by the PCP instance by repeated application of functional constraints.

337 ▶ **Theorem 10.** T-ROBUSTNESS(**EqTemp**, RC) is undecidable.


338 It might be tempting to relate the above result to the undecidability of the implication
 339 problem for functional and inclusion dependencies [11]. Functional constraints indeed allow
 340 to define inclusion dependencies (as in the SmallBank example) but they always relate
 341 complete tuples and are not suited to define functional dependencies. Furthermore, the proof
 342 of Theorem 10 makes use of only unary relations, for which the implication problem for
 343 functional dependencies and inclusion dependencies is known to be decidable.

344 To obtain decidable fragments, we introduce restrictions on the structure of functional
 345 constraints. The *schema graph* $SG(\text{Rels}, \text{Funcs})$ of a schema $(\text{Rels}, \text{Funcs})$ is a directed
 346 multigraph having the relations in Rels as nodes, and in which there are as many edges from
 347 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
 348 $\text{range}(f) = S$. We say that a schema $(\text{Rels}, \text{Funcs})$ is *acyclic* if the multigraph $SG(\text{Rels}, \text{Funcs})$
 349 is acyclic and that it is a *multi-tree* if there is at most one directed path between any two
 350 nodes in $SG(\text{Rels}, \text{Funcs})$.

351 ▶ **Example 11.** Consider the schema $(\{P, Q, R, S\}, \{f_{P,R}, f_{Q,R}, f_{R,S}\})$ with $\text{dom}(f_{i,j}) = i$
 352 and $\text{range}(f_{i,j}) = j$ for each function $f_{i,j}$. The corresponding schema graph with solid lines is



■ **Figure 4** Acyclic schema graph for 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 5** 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 $f_{A \rightarrow C}$.

353 given in Figure 4. This schema is a multi-tree, as there is at most one path between any pair
 354 of nodes. Notice that the definition of a multi-tree is more general than a forest, as a node
 355 can still have multiple parents (e.g., node R in our example). Adding the function name $f_{Q,S}$
 356 with $\text{dom}(f_{Q,S}) = Q$ and $\text{range}(f_{Q,S}) = S$ results in the schema graph given in Figure 4 that
 357 is still acyclic, but no longer a multi-tree as there are now two paths from Q to S . \square

358 The schema graph constructed in the proof of Theorem 10 contains several cycles (cf.,
 359 Figure 6 in Appendix A). We consider in Section 5 robustness for a fragment where a
 360 restricted form of cycles in the schema graph is allowed but where additional constraints on
 361 the templates are assumed. We consider robustness for acyclic schema graphs in Section 6.

5 Robustness for Templates admitting Multi-Tree Bijectivity

362 We say that a set of transaction templates \mathcal{P} over a schema $(\text{Rels}, \text{Funcs})$ admits *multi-tree*
 363 *bijectivity* if a disjoint partitioning of Funcs in pairs $(f_1, g_1), (f_2, g_2), \dots, (f_n, g_n)$ exists such
 364 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) ;
 365 every schema graph $SG(\text{Rels}, \{h_1, h_2, \dots, h_n\})$ over the schema restricted to function names
 366 $\{h_1, h_2, \dots, h_n\}$ (with $h_i = f_i$ or $h_i = g_i$) is a multi-tree; and, for every pair of function
 367 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
 368 $f_i(\mathbf{X}) = \mathbf{Y} \in \Gamma_j$ iff $g_i(\mathbf{Y}) = \mathbf{X} \in \Gamma_j$. Intuitively, we can think of f_i as a bijective function, with
 369 g_i its inverse. We denote the class of all sets of templates admitting multi-tree bijectivity by
 370 **MTBTemp**. The SmallBank benchmark given in Figure 8 is in **MTBTemp**, witnessed by
 371 the partitioning $\{(f_{A \rightarrow C}, f_{C \rightarrow A}), (f_{A \rightarrow S}, f_{S \rightarrow A})\}$. For example, the schema graph restricted
 372 to $f_{A \rightarrow C}$ and $f_{A \rightarrow S}$ is a tree and therefore also a multi-tree, as illustrated in Figure 5.

373 The next theorem allows disequalities whereas Theorem 10 does not require them.

374 ► **Theorem 12.** $\text{T-ROBUSTNESS}(\text{MTBTemp}, \text{RC})$ is decidable in NLOGSPACE .

375 The approach followed in the proof of Theorem 12 is to repeatedly pick a transaction
 376 template while maintaining an overall consistent variable mapping in search for a counter-
 377 example multiversion split schedule that by Theorem 7 suffices to show that robustness
 378 does not hold. The main challenge is to show that a variable mapping consistent with all
 379 functional constraints can be maintained in logarithmic space and that all requirements for a
 380 multiversion split schedule can be verified in NLOGSPACE .

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

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

389 Intuitively, potentially conflicting operations lead to conflicting operations when the variables
 390 of these operations are mapped to the same tuple by a variable assignment. In analogy to
 391 conflicting quadruples over a set of transactions as in Definition 6, we consider *potentially*
 392 *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
 393 potentially conflicting with an operation $p_j \in \tau_j$. For a sequence of potentially conflicting
 394 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
 395 set $\{\tau_1, \dots, \tau_m\}$ of transaction templates mentioned in D . For ease of exposition, we assume
 396 a variable renaming such that any pair of templates in $\text{Trans}(D)$ uses a disjoint set of
 397 variables.⁴ The sequence D induces a sequence of conflicting quadruples $C = (T_1, b_1, a_2, T_2),$
 398 $\dots, (T_m, b_m, a_1, T_1)$ by applying a variable assignment μ_i to each τ_i in $\text{Trans}(D)$. We call
 399 such a set of variable assignments simply a *variable mapping* for D , denoted $\bar{\mu}$, and write
 400 $\bar{\mu}(D) = C$. For a variable X occurring in a template τ_i , we write $\bar{\mu}(X)$ as a shorthand notation
 401 for $\mu_i(X)$, with μ_i the variable assignment over τ_i in $\bar{\mu}$. This is well-defined as all templates
 402 in $\text{Trans}(D)$ are variable-disjoint. Furthermore, $\bar{\mu}(\text{var}(o_i)) = \bar{\mu}(\text{var}(p_j))$ for each potentially
 403 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
 404 not a valid conflicting quadruple in C . We say that a variable mapping $\bar{\mu}$ is admissible for a
 405 database \mathbf{D} if every variable assignment μ_i in $\bar{\mu}$ is admissible for \mathbf{D} .

406 A basic insight is that if there is a multiversion split schedule s for some C over a set of
 407 transactions \mathcal{T} consistent with \mathcal{P} and a database \mathbf{D} , then there is a sequence of potentially
 408 conflicting quadruples D such that $\bar{\mu}(D) = C$ for some $\bar{\mu}$. We will verify the existence of
 409 such a C , satisfying the properties of Definition 6, by nondeterministically constructing D
 410 on-the-fly together with a mapping $\bar{\mu}$. We show in Lemma 14 that when $\mathcal{P} \in \mathbf{MTBTemp}$,
 411 $\bar{\mu}$ is a collection of disjoint type mappings (that map variables of the same type to the same
 412 tuple) such that variables that are “connected” in D (in a way that we will make precise next)
 413 are mapped using the same type mapping. Lemma 15 then shows that already a constant
 414 number of those type mappings suffice.

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

⁴ 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 12 as $\text{Trans}(D)$ nor \mathcal{P}' are used in the algorithm. Their only purpose is to reason about properties of $\bar{\mu}$.

430 ■ $i = j$ and $X \xrightarrow{F} \tau_i Y$ (implication within the same template);
 431 ■ $F = \varepsilon$ 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
 432 o_i (respectively p_i) an operation over X and p_j (respectively o_j) an operation over Y
 433 (implication between templates, notice that $X \xrightarrow{F} D Y$ iff $Y \xrightarrow{F} D X$); or
 434 ■ there exists a variable Z such that $X \xrightarrow{F} D Z$ and $Z \xrightarrow{F} D Y$ with $F = F_1 \cdot F_2$.
 435 Two variables X and Y occurring in $\text{Trans}(D)$ are *connected in D* , denoted $X \approx_D Y$, if
 436 $X \xrightarrow{F} D Y$ or $Y \xrightarrow{F} D X$, or if there is a variable Z with $X \approx_D Z$ and either $Z \xrightarrow{F} D Y$ or $Y \xrightarrow{F} D Z$ for
 437 some sequence F . Furthermore, two variables X and Y occurring in a template τ are *connected*
 438 *in τ* , denoted $X \approx_\tau Y$, if $X \xrightarrow{F} \tau Y$ or $Y \xrightarrow{F} \tau X$, or if there is a variable Z with $X \approx_\tau Z$ and either
 439 $Z \xrightarrow{F} \tau Y$ or $Y \xrightarrow{F} \tau Z$ for some sequence F . These definitions of connectedness can be trivially
 440 extended to operations over variables: two operations in D (respectively τ) are connected in
 441 D (respectively τ) if they are over variables that are connected in D (respectively τ). When
 442 F is not important we drop it from the notation. For instance, we denote by $X \rightsquigarrow_D Y$ that
 443 there is an F with $X \xrightarrow{F} D Y$.

444 ► **Lemma 13.** *Let D be a sequence of potentially conflicting quadruples over $\mathcal{P} \in \mathbf{MTBTemp}$.
 445 Then $X \approx_D Y$ implies $X \rightsquigarrow_D Y$ and $Y \rightsquigarrow_D X$. Furthermore, if $\text{type}(X) = \text{type}(Y)$ then $\bar{\mu}(X) = \bar{\mu}(Y)$
 446 for every variable mapping $\bar{\mu}$ for D that is admissible for some database \mathbf{D} .*

447 It follows from Lemma 13 that, if we group connected variables, then the same tuple is
 448 assigned to all variables of the same type in this group. We encode this choice of tuples for
 449 variables through (total) functions $c : \mathbf{Rels} \rightarrow \mathbf{Tuples}$ that we call *type mappings* and which
 450 map a relation onto a particular tuple of that relation's type. For instance, in `SmallBank`,
 451 a type mapping c is determined by an `Account` tuple \mathbf{a} , a `Savings` tuple \mathbf{s} , and a `Checking`
 452 tuple \mathbf{c} . The following lemma makes explicit how $\bar{\mu}$ can be decomposed into type mappings
 453 such that connected variables use the same type mapping and disequalities enforce the use of
 454 different type mappings.

455 ► **Lemma 14.** *For a multiversion split schedule s based on a sequence of conflicting quadruples
 456 C over a set of transactions \mathcal{T} consistent with a $\mathcal{P} \in \mathbf{MTBTemp}$ and a database \mathbf{D} , let $\bar{\mu}$
 457 be the variable mapping for a sequence of potentially conflicting quadruples D over \mathcal{P} with
 458 $\bar{\mu}(D) = C$. Then, a set \mathcal{S} of type mappings over disjoint ranges and a function $\varphi_{\mathcal{S}} : \mathbf{Var} \rightarrow \mathcal{S}$
 459 exist with:*

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

463 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 14 we can derive a
 464 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 $c_{o_i} =$
 465 $\varphi_{\mathcal{S}}(\text{var}(o_i))$ and $c_{p_i} = \varphi_{\mathcal{S}}(\text{var}(p_i))$ for $i \in [1, m]$. Intuitively, this sequence of quintuples can
 466 be used to reconstruct the original multiversion split schedule s . The next lemma shows that
 467 we can decide robustness against RC over a set of transaction templates admitting multi-tree
 468 bijectivity by searching for a specific sequence of quintuples over at most four type mappings.

469 ► **Lemma 15.** *Let $\mathcal{P} \in \mathbf{MTBTemp}$ and let $\mathcal{S} = \{c_1, c_2, c_3, c_4\}$ be a set consisting of four
 470 type mappings with disjoint ranges. Then, \mathcal{P} is not robust against RC iff there is a sequence
 471 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
 472 each quintuple $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$ in E :*

- 473 1. o_i and p_i are operations in τ_i , and $c_{o_i}, c_{p_i} \in \mathcal{S}$;
- 474 2. $X_i \not\approx_{\tau_i} Y_i$ for each constraint $X_i \neq Y_i$ in τ_i ;
- 475 3. $c_{o_i} = c_{p_i}$ iff $o_i \approx_{\tau_i} p_i$;

- 476 4. $c_{o_i} \neq c_{p_i}$ if there is a constraint $X_i \neq Y_i$ in τ_i with $X_i \approx_{\tau_i} \text{var}(o_i)$ and $Y_i \approx_{\tau_i} \text{var}(p_i)$;
 477 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
 478 in τ_i potentially ww-conflicting with an operation o'_1 in $\text{prefix}_{o_1}(\tau_1)$ with $\text{var}(o'_i) \approx_{\tau_i} \text{var}(q_i)$
 479 and $\text{var}(o'_1) \approx_{\tau_1} \text{var}(q_1)$.
 480 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
 481 E with $j = i + 1$, or $i = m$ and $j = 1$:
 482 6. o_i is potentially conflicting with p_j and $c_{o_i} = c_{p_j}$;
 483 7. if $i = 1$ and $j = 2$, then o_1 is potentially rw-conflicting with p_2 ; and
 484 8. if $i = m$ and $j = 1$, then $o_1 <_{\tau_1} p_1$ or o_m is potentially rw-conflicting with p_1 .

485 The items have the following meaning: (2) τ_i is satisfiable; (3) connected operations are
 486 assigned the same type mapping; (4) variables connected through an inequality are assigned
 487 a different type mapping; (5) φ_S only assigns the same type mapping to o_1 or p_1 in τ_1 and o_i
 488 or p_i in τ_i if it does not introduce a dirty write in the resulting multiversion split schedule (cf.
 489 Condition (1) in Definition 6); (6) each pair of variables in operations used for conflicts are
 490 assigned the same type mapping; (7, 8) the operations used for conflicts between τ_1 , τ_2 and
 491 τ_m are restricted to satisfy respectively Condition (3) and (2) in Definition 6 in the resulting
 492 multiversion split schedule.

493 The characterization for $\text{T-ROBUSTNESS}(\text{MTBTemp}, \text{RC})$ in Lemma 15 implies an NLOG-
 494 SPACE algorithm guessing the counterexample sequence E , thereby proving Theorem 12.
 495 Indeed, the algorithm guesses the sequence of quintuples E , verifying all conditions for each
 496 newly guessed quintuple while only requiring logarithmic space. Notice in particular that
 497 we only need to keep track of two other quintuples when verifying all conditions for the
 498 newly guessed quintuple, namely the first quintuple over τ_1 and the quintuple immediately
 499 preceding the newly guessed one. As usual, we can think of the encoding of templates and
 500 operations mentioned in each quintuple as pointers referring to the corresponding templates
 501 and operations on the input tape. Furthermore, we do not encode the four type mappings
 502 explicitly as such a representation of a mapping might require polynomial space. Since we
 503 are only interested in (dis)equality between type mappings, an encoding where these four
 504 type mappings are represented by four arbitrary strings of constant size suffices. More details
 505 can be found in Appendix B.4.

506 6 Robustness for Templates over Acyclic Schemas

507 We denote by **AcycTemp** the class of all sets of transaction templates over acyclic schemas.

508 ► **Theorem 16.** $\text{T-ROBUSTNESS}(\text{AcycTemp}, \text{RC})$ is decidable in EXPSpace.

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

512 ► **Lemma 17.** Let D be a sequence of potentially conflicting quadruples over a set of
 513 transaction templates $\mathcal{P} \in \text{AcycTemp}$. For every pair of variables X, Y occurring in $\text{Trans}(D)$,
 514 if $X \xrightarrow{F}_D Y$, then $\text{type}(X) \xrightarrow{F}_{SG} \text{type}(Y)$, with SG the corresponding schema graph.

515 Notice that an assignment of a tuple to a variable X determines the tuples assigned to
 516 all variables Y with $X \xrightarrow{F}_D Y$ for some sequence of function names F . From Lemma 17 it
 517 follows that each such implied tuple is witnessed by a path in the corresponding schema
 518 graph SG . Therefore, the maximal number of different tuples implied by X corresponds to
 519 the number of paths in SG starting in $\text{type}(X)$, which is finite when SG is acyclic. Because

520 there can be multiple paths between nodes in the schema graph, it is no longer the case as in
 521 the previous section that variables of the same type connected in D must be assigned the
 522 same value. So, instead of using type mappings, we introduce *tuple-contexts* to represent the
 523 sets of all tuples implied by the assignment of a given variable. Formally, a *tuple-context*
 524 *for a type* $R \in \text{Rels}$ is a function from paths with source R in $SG(\text{Rels}, \text{Funcs})$ to tuples in
 525 **Tuples** of the appropriate type. That is, for each tuple-context c for type R and for each
 526 path $R \xrightarrow{f}_{SG} S$ in SG , $\text{type}(c(R \xrightarrow{f}_{SG} S)) = S$.

527 Similar to Lemma 14, we show that we can represent a counterexample schedule based
 528 on D by assigning a tuple-context to each variable in $\text{Trans}(D)$, taking special care when
 529 assigning contexts to variables connected in D to make sure that they are properly related
 530 to each other. For this, we introduce a (partial) function $\varphi_{\mathcal{A}} : \mathbf{Var} \rightarrow \mathcal{A}$ mapping (a subset
 531 of) variables in $\text{Trans}(D)$ to tuple-contexts in \mathcal{A} (for \mathcal{A} a set of tuple-contexts) and refer to
 532 it as a *(partial) context assignment for D over \mathcal{A}* . In a sequence of lemma's, we show that
 533 $\varphi_{\mathcal{A}}$ can always be expanded into a total function and an approach based on enumeration of
 534 quintuples analogous to Lemma 15 suffices to decide robustness. A major difference with the
 535 previous section is that there is no longer a constant bound on the number of tuple-contexts
 536 that are needed and consistency between between tuple-contexts in connected variables needs
 537 to be maintained. A full proof can be found in Appendix C.

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

- 540 ■ $X = Y$;
- 541 ■ there exists a pair of variables Z and W in τ and a sequence of function names F with
 542 $Z \equiv_{\tau} W$, $Z \xrightarrow{f}_{\tau} X$ and $W \xrightarrow{f}_{\tau} Y$; or
- 543 ■ there exists a variable Z with $X \equiv_{\tau} Z$ and $Y \equiv_{\tau} Z$.

544 Then, a transaction template τ is *restricted* if for every combination of variables X, Y, W, Z in
 545 τ 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**
 546 the class of all sets of restricted transaction templates over acyclic schemas.

- 547 ► **Theorem 18. 1.** $\text{T-ROBUSTNESS}(\mathbf{AcycResTemp}, RC)$ is decidable in EXPTIME.
 548 **2.** $\text{T-ROBUSTNESS}(\mathbf{AcycTemp}, RC)$ is decidable in PSPACE when the number of paths between
 549 any two nodes in the schema graph is bounded by a constant k .

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

554 7 Related Work

555 **Transaction Programs.** Previous work on static robustness testing [13, 3] for transaction
 556 programs is based on the following key insight: when a *schedule* is not serializable, then the
 557 dependency graph constructed from that schedule contains a cycle satisfying a condition
 558 specific to the isolation level at hand (*dangerous structure* for SNAPSHOT ISOLATION and the
 559 presence of a *counterflow edge* for RC). That insight is extended to a workload of *transaction*
 560 *programs* through the construction of a so-called static dependency graph where each program
 561 is represented by a node, and there is a conflict edge from one program to another if there can
 562 be a schedule that gives rise to that conflict. The absence of a cycle satisfying the condition
 563 specific to that isolation level then guarantees robustness while the presence of a cycle does
 564 not necessarily imply non-robustness.

565 Other work studies robustness within a framework for uniformly specifying different
566 isolation levels in a declarative way [8, 7, 9]. A key assumption here is *atomic visibility*
567 requiring that either all or none of the updates of each transaction are visible to other
568 transactions. These approaches aim at higher isolation levels and cannot be used for RC, as
569 RC does not admit *atomic visibility*.

570 **Transaction Templates.** The static robustness approach based on transaction tem-
571 plates [18] differs in two ways. First, it makes more underlying assumptions explicit within
572 the formalism of transaction templates (whereas previous work departs from the static
573 dependency graph that should be constructed in some way by the dba). Second, it allows for
574 a decision procedure that is sound and complete for robustness testing against RC, allowing
575 to detect larger subsets of transactions to be robust [18].

576 The formalisation of transactions and conflict serializability in [18] and this paper is based
577 on [12], generalized to operations over attributes of tuples and extended with U-operations
578 that combine R- and W-operations into one atomic operation. These definitions are closely
579 related to the formalization presented by Adya et al. [1], but we assume a total rather than
580 a partial order over the operations in a schedule. There are also a few restrictions to the
581 model: there needs to be a fixed set of read-only attributes that cannot be updated and
582 which are used to select tuples for update. The most typical example of this are primary
583 key values passed to transaction templates as parameters. The inability to update primary
584 keys is not an important restriction in many workloads, where keys, once assigned, never get
585 changed, for regulatory or data integrity reasons.

586 In [18], a PTIME decision procedure is obtained for robustness against RC for templates
587 without functional constraints and the present paper improves that result to NLOGSPACE. In
588 addition, an experimental study was performed showing how an approach based on robustness
589 and making transactions robust through promotion can improve transaction throughput.

590 **Transactions.** Fekete [12] is the first work that provides a necessary and sufficient condition
591 for deciding robustness against SNAPSHOT ISOLATION for a workload of concrete transactions
592 (not transaction programs). That work provides a characterization for acceptable allocations
593 when every transaction runs under either SNAPSHOT ISOLATION or strict two-phase locking
594 (S2PL). The allocation then is acceptable when every possible execution respecting the alloc-
595 ated isolation levels is serializable. As a side result, this work indirectly provides a necessary
596 and sufficient condition for robustness against SNAPSHOT ISOLATION, since robustness against
597 SNAPSHOT ISOLATION holds iff the allocation where each transaction is allocated to SNAPSHOT
598 ISOLATION is acceptable. Ketsman et al. [14] provide full characterisations for robustness
599 against READ COMMITTED and READ UNCOMMITTED under lock-based semantics. In addition,
600 it is shown that the corresponding decision problems are complete for CONP and LOGSPACE,
601 respectively, which should be contrasted with the polynomial time characterization obtained
602 in [18] for robustness against *multiversion* read committed.

603 **8 Conclusion**

604 This paper falls within a more general research line investigating how transaction throughput
605 can be improved through an approach based on robustness testing that can be readily applied
606 without making any changes to the underlying database system. As argued in Section 2,
607 incorporating functional constraints can detect larger sets of templates to be robust and
608 requires less R-operations to be promoted to U-operations. In future work, we plan to look
609 at lower bounds, restrictions that lower complexity, and consider other referential integrity
610 constraints to further enlarge the modelling power of transaction templates.

611 — **References** —

- 612 **1** Atul Adya, Barbara Liskov, and Patrick E. O’Neil. Generalized isolation level definitions. In
613 *ICDE*, pages 67–78, 2000.
- 614 **2** Mohammad Alomari, Michael Cahill, Alan Fekete, and Uwe Rohm. The cost of serializability
615 on platforms that use snapshot isolation. In *ICDE*, pages 576–585, 2008.
- 616 **3** Mohammad Alomari and Alan Fekete. Serializable use of read committed isolation level. In
617 *AICCSA*, pages 1–8, 2015.
- 618 **4** Sidi Mohamed Beillahi, Ahmed Bouajjani, and Constantin Enea. Checking robustness against
619 snapshot isolation. In *CAV*, pages 286–304, 2019.
- 620 **5** Sidi Mohamed Beillahi, Ahmed Bouajjani, and Constantin Enea. Robustness against transac-
621 tional causal consistency. In *CONCUR*, pages 1–18, 2019.
- 622 **6** Hal Berenson, Philip A. Bernstein, Jim Gray, Jim Melton, Elizabeth J. O’Neil, and Patrick E.
623 O’Neil. A critique of ANSI SQL isolation levels. In *SIGMOD*, pages 1–10, 1995.
- 624 **7** Giovanni Bernardi and Alexey Gotsman. Robustness against consistency models with atomic
625 visibility. In *CONCUR*, pages 7:1–7:15, 2016.
- 626 **8** Andrea Cerone, Giovanni Bernardi, and Alexey Gotsman. A framework for transactional
627 consistency models with atomic visibility. In *CONCUR*, pages 58–71, 2015.
- 628 **9** Andrea Cerone and Alexey Gotsman. Analysing snapshot isolation. *J.ACM*, 65(2):1–41, 2018.
- 629 **10** Andrea Cerone, Alexey Gotsman, and Hongseok Yang. Algebraic Laws for Weak Consistency.
630 In *CONCUR*, pages 26:1–26:18, 2017.
- 631 **11** Ashok K. Chandra and Moshe Y. Vardi. The implication problem for functional and inclusion
632 dependencies is undecidable. *SIAM J. Comput.*, 14(3):671–677, 1985. URL: <https://doi.org/10.1137/0214049>, doi:10.1137/0214049.
- 633 **12** Alan Fekete. Allocating isolation levels to transactions. In *PODS*, pages 206–215, 2005.
- 634 **13** Alan Fekete, Dimitrios Liarokapis, Elizabeth J. O’Neil, Patrick E. O’Neil, and Dennis E.
635 Shasha. Making snapshot isolation serializable. *ACM Trans. Database Syst.*, 30(2):492–528,
636 2005.
- 637 **14** Bas Ketsman, Christoph Koch, Frank Neven, and Brecht Vandevoort. Deciding robustness for
638 lower SQL isolation levels. In *PODS*, pages 315–330, 2020.
- 639 **15** Christos H. Papadimitriou. *The Theory of Database Concurrency Control*. Computer Science
640 Press, 1986.
- 641 **16** Emil L. Post. A variant of a recursively unsolvable problem. *Bull. Amer. Math. Soc.*, pages
642 264–268, 1946.
- 643 **17** TPC-C. On-line transaction processing benchmark. <http://www.tpc.org/tpcc/>.
- 644 **18** Brecht Vandevoort, Bas Ketsman, Christoph Koch, and Frank Neven. Robustness against
645 read committed for transaction templates (full version of submitted paper). <https://github.com/fneven-uh/paper>, 2021.
- 647

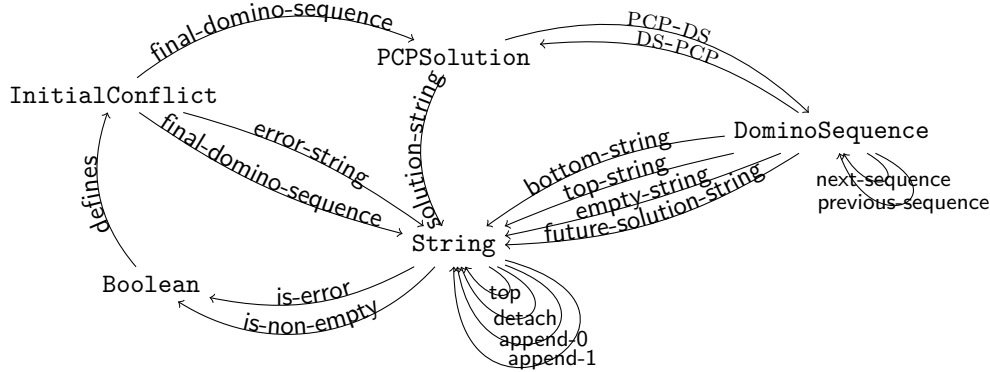
648 **A Proofs for Section 4**

649 Here, we present the proof of Theorem 10.

650 **► Theorem 10.** $T\text{-ROBUSTNESS}(EqTemp, RC)$ is undecidable.

651 A domino is a pair (\mathbf{a}, \mathbf{b}) of two non-empty strings over Σ . Henceforth we call \mathbf{a} its *top*
 652 *value* and \mathbf{b} its *bottom value*. Given a set of dominoes \mathcal{D} , the PCP asks if a non-empty
 653 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
 654 $\mathbf{a}_1\mathbf{a}_2 \dots \mathbf{a}_r$ and $\mathbf{b}_1\mathbf{b}_2 \dots \mathbf{b}_r$ are identical.

655 For the reduction to non-robustness against RC, we construct a set \mathcal{P} of transaction
 656 templates consisting of the transaction templates in Figure 7 for \mathcal{D} . There are the trans-
 657 actions Split, First and Last (whose meaning will be explained next) and for every domi-
 658 no in \mathcal{D} there is a template in Figure 7 representing that domino and the action of ap-
 659 pending that domino to a sequence of dominoes. The schema consists of the relations
 660 $\{\text{Boolean}, \text{InitialConflict}, \text{String}, \text{PCPSolution}, \text{DominoSequence}\}$ whose meaning will
 661 be explained below together with a discussion of all the functions. The schema graph is
 662 presented in Figure 6 and contains various cycles.



663 **■ Figure 6** Schema graph for the transaction templates in Figure 7 (for any set of dominoes).

664 To prove Theorem 10, we will show that there is a solution for PCP if and only if \mathcal{P}
 665 is not robust against RC. For the only-if direction, we show that, if there is a solution
 666 $\mathbf{d} = d_1, d_2, \dots, d_r$ for the PCP problem over \mathcal{D} , then there is a multiversion split schedule
 667 that encodes this solution in a particular way: in this schedule the split transaction is
 668 an instantiation of transaction template Split, the next transaction is an instantiation
 669 of First, then followed by instantiations of transaction templates $\text{Domino}_{\mathbf{a}_1}, \dots, \text{Domino}_{\mathbf{a}_r}$
 670 representing the sequence of dominoes in solution \mathbf{d} , and finally an instantiation of transaction
 671 template Last. Henceforth, we call a schedule that encodes a sequence of dominoes \mathbf{d} in this
 672 way a *schedule-encoding of \mathbf{d}* . For the if-direction, we first show that every multiversion split
 673 schedule consistent with the transaction templates in Figure 7 for some set \mathcal{D} of dominoes
 674 is a schedule-encoding for some sequence \mathbf{d} of dominoes from \mathcal{D} , and then that for every
 675 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} .

Split(I):	First(S ₁):	Last(B):
W[X ₁ : Boolean]	W[X ₂ : Boolean]	W[B : DominoSequence]
R[I : InitialConflict]	W[I : InitialConflict]	R[S _t : String]
R[S ₁ : String]	R[S ₀ : String]	R[S _b : String]
R[S _e : String]	R[S _e : String]	R[S ₁ : String]
W[C : PCPSolution]	R[S ₁ : String]	W[C : PCPSolution]
X ₁ = f _{is-non-empty} (S ₁)	W[B : DominoSequence]	S _t = f _{top-string} (B)
X ₁ = f _{is-error} (S _e)	X ₂ = f _{is-non-empty} (S ₀)	S _b = f _{bottom-string} (B)
C = f _{final-domino-sequence} (I)	X ₂ = f _{is-error} (S ₀)	S ₁ = f _{future-solution-string} (B)
S ₁ = f _{final-dominoes-string} (I)	S ₁ = f _{final-dominoes-string} (I)	C = f _{DS→PCP} (B)
S _e = f _{error-string} (I)	S ₀ = f _{top-string} (B)	S _t = f _{solution-string} (C)
S ₁ = f _{solution-string} (C)	S ₀ = f _{bottom-string} (B)	S _b = f _{solution-string} (C)
I = f _{defines} (X ₁)	S ₀ = f _{empty-string} (B)	S ₁ = f _{solution-string} (C)
	S ₁ = f _{future-solution-string} (B)	B = f _{PCP→DS} (C)
	S _e = f _{detach} (S ₀)	
	S _e = f _{detach} (S _e)	
	I = f _{defines} (X ₂)	
	B = f _{empty-domino-sequence} (S ₁)	

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 $\text{Domino}_i(\text{B})$:

W[B : DominoSequence]	S _t = f _{top-string} (B)	S _b = f _{bottom-string} (B)
R[S ₀ : String]	S _{ta₁} = f _{append-a₁} (S _t)	S _{bb₁} = f _{append-b₁} (S _b)
R[S ₁ : String]	S _{ta₁a₂} = f _{append-a₂} (S _{ta₁})	S _{bb₁b₂} = f _{append-b₂} (S _{bb₁})
R[S _t : String]
R[S _{ta₁} : String]	S _{ta₁a₂...a_h} = f _{append-a_h} (S _{ta₁...a_{h-1}})	S _{bb₁b₂...b_k} = f _{append-b_k} (S _{bb₁...b_{k-1}})
R[S _{ta₁a₂} : String]	S _t = f _{detach} (S _{ta₁})	S _b = f _{detach} (S _{bb₁})
...	S _{ta₁} = f _{detach} (S _{ta₁a₂})	S _{bb₁} = f _{detach} (S _{bb₁b₂})
R[S _{ta₁a₂...a_h} : String]
R[S _b : String]	S _{ta₁a₂...a_{h-1}} = f _{detach} (S _{ta₁a₂...a_h})	S _{bb₁b₂...b_{k-1}} = f _{detach} (S _{bb₁b₂...b_k})
R[S _{bb₁} : String]	S _{ta₁a₂...a_h} = f _{top-string} (B _{next})	S _{bb₁b₂...b_k} = f _{bottom-string} (B _{next})
R[S _{bb₁b₂} : String]	S _{a₁} = f _{top} (S _{ta₁})	S _{b₁} = f _{top} (S _{bb₁})
...	S _{a₂} = f _{top} (S _{ta₁a₂})	S _{b₂} = f _{top} (S _{bb₁b₂})
R[S _{bb₁b₂...b_k} : String]
W[B _{next} : DominoSequence]	S _{a_h} = f _{top} (S _{ta₁a₂...a_h})	S _{b_k} = f _{top} (S _{bb₁b₂...b_k})
	S ₁ = f _{future-solution-string} (B)	S ₁ = f _{future-solution-string} (B _{next})
	S ₀ = f _{empty-string} (B)	S ₀ = f _{empty-string} (B _{next})
		B _{next} = f _{next-sequence} (B)
		B = f _{previous-sequence} (B _{next})

■ **Figure 7** Transaction templates for the proof of Theorem 10.

A.1 Only-if direction for the proof of Theorem 10

► **Proposition 19** (Only-if part of Theorem 10). *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 7 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 6 with transactions based on the transaction templates \mathcal{P} in Figure 7.

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 functional constraints, $f_{\text{PCP} \rightarrow \text{DS}} : \text{PCPSolution} \rightarrow \text{DominoSequence}$ and $f_{\text{DS} \rightarrow \text{PCP}} : \text{DominoSequence} \rightarrow \text{PCPSolution}$ that map these interpretations on one another. That is, $f_{\text{PCP} \rightarrow \text{DS}}^{\mathbf{D}}(\mathbf{d}) = \mathbf{d}$ and $f_{\text{DS} \rightarrow \text{PCP}}^{\mathbf{D}}(\mathbf{d}) = \mathbf{d}$.

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

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

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} :

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

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

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 $f_{\text{top-string}} : \text{DominoSequence} \rightarrow \text{String}$, $f_{\text{bottom-string}} :$

719 DominoSequence \rightarrow String, and $f_{\text{empty-string}} : \text{DominoSequence} \rightarrow \text{String}$ with following
 720 interpretations in \mathbf{D} :

$$\begin{aligned}
 721 \quad & f_{\text{top-string}}^{\mathbf{D}}(\mathbf{d}') = \mathbf{e}, \text{ with } \mathbf{e} \text{ the (possibly empty) read of top values on dominoes in } \mathbf{d}', \\
 722 \quad & f_{\text{bottom-string}}^{\mathbf{D}}(\mathbf{d}') = \mathbf{e}, \text{ with } \mathbf{e} \text{ the (possibly empty) read of bottom values on dominoes in } \mathbf{d}', \text{ and} \\
 723 \quad & f_{\text{empty-string}}^{\mathbf{D}}(\mathbf{d}') = \langle \rangle.
 \end{aligned}$$

725 Finally, for function $f_{\text{future-solution-string}} : \text{DominoSequence} \rightarrow \text{String}$ we consider the
 726 interpretation that associates every domino sequence \mathbf{d}' represented by a tuple in relation
 727 DominoSequence in \mathbf{D} to the final read $f_{\text{future-solution-string}}^{\mathbf{D}}(\mathbf{d}') = \mathbf{c}$. Function $f_{\text{solution-string}} : \text{PCPSolution} \rightarrow \text{String}$
 728 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 $f_{\text{empty-domino-sequence}} : \text{String} \rightarrow \text{DominoSequence}$ is
 729 interpreted to map every tuple in String onto the tuple from DominoSequence representing
 730 the empty sequence $\langle \rangle$.
 731

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

735 Relation Boolean $^{\mathbf{D}}$ contains two tuples, which we interpret as Boolean values 0 and 1.
 736 Function $f_{\text{is-non-empty}} : \text{String} \rightarrow \text{Boolean}$ and $f_{\text{is-error}} : \text{String} \rightarrow \text{Boolean}$ are interpreted as
 737 follows:

$$\begin{aligned}
 738 \quad & f_{\text{is-non-empty}}^{\mathbf{D}}(s) = \begin{cases} 1 & \text{if } s \neq \langle \rangle, \\ 0 & \text{otherwise,} \end{cases} \text{ , and} \\
 739 \quad & f_{\text{is-error}}^{\mathbf{D}}(s) = \begin{cases} 1 & \text{if } s = \langle \text{error} \rangle, \\ 0 & \text{otherwise,} \end{cases} \text{ , and.} \\
 740
 \end{aligned}$$

741 Finally, relation InitialConflict $^{\mathbf{D}}$ contains a single tuple, which we refer to by $\langle \text{init} \rangle$. The
 742 interpretation of $f_{\text{defines}} : \text{Boolean} \rightarrow \text{InitialConflict}$ maps 1 and 0 onto $\langle \text{init} \rangle$. Function
 743 $f_{\text{error-string}} : \text{InitialConflict} \rightarrow \text{String}$ maps $\langle \text{init} \rangle$ onto $\langle \text{error} \rangle$. Functions $f_{\text{final-domino-string}} : \text{InitialConflict} \rightarrow \text{DominoSequence}$
 744 and $f_{\text{final-domino-sequence}} : \text{InitialConflict} \rightarrow \text{PCPSolution}$ map $\langle \text{init} \rangle$ onto the solution domino sequence \mathbf{d} , respectively on the final read \mathbf{c} of \mathbf{d} .
 745

746 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 $T_1 = \text{Split}(\langle \text{init} \rangle)$,
 747 $T_2 = \text{First}(\langle \text{init} \rangle)$, for $i : 1 \leq i \leq r$, transaction $T_{i+2} = \text{Domino}_i((d_1, \dots, d_i))$, $T_m =$
 748 $\text{Last}((d_1, \dots, d_r))$ and $b_1 = \langle \text{init} \rangle$ has the conditions of Definition 6. Indeed, it is based on se-
 749 quence of conflict quadruples $(T_1, \mathbf{R}_1[\langle \text{init} \rangle], \mathbf{W}_2[\langle \text{init} \rangle], T_2), (T_2, \mathbf{W}_2[\langle \rangle], \mathbf{W}_3[\langle \rangle], T_3), (T_3, \mathbf{W}_3[\langle \mathbf{d}_1 \rangle],$
 750 $\mathbf{W}_4[\langle \mathbf{d}_2 \rangle], T_4), \dots, (T_{r+2}, \mathbf{W}_{r+2}[\langle \mathbf{d}_1, \dots, \mathbf{d}_r \rangle], \mathbf{W}_{r+3}[\langle \mathbf{d}_1, \dots, \mathbf{d}_r \rangle], T_{r+3}), (T_{r+3}, \mathbf{W}_{r+3}[\langle \mathbf{d} \rangle], \mathbf{W}_1[\langle \mathbf{d} \rangle], T_1)$.
 751

752 Condition (1) is true because there is no ww-conflict between a write operation in
 753 $\text{prefix}_{b_1}(T_1)$ and a write operation in any of the transactions T_2, \dots, T_m , since the first write
 754 operation, respectively second write operation, in $\text{Split}(\langle \text{init} \rangle)$ has a type that only occurs
 755 before the conflict with $\text{First}(\langle \text{init} \rangle)$, and is the conflict with $\text{Last}((d_1, \dots, d_r))$, respectively.
 756 Furthermore (2) is true because $b_1 <_{T_1} a_1$ and Condition (3) is true because b_1 and a_2 are
 757 rw-conflicting. \blacktriangleleft

757 A.2 Helpful lemma

758 \blacktriangleright **Lemma 20.** *If a set \mathcal{P} of transaction templates is not robust against RC then there is a*
 759 *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 $\mathcal{T} = \{T_1, \dots, T_m\}$*
 760 *of transactions consistent with \mathcal{P} in which an operation from a transaction T_j depends on an*
 761 *operation from transaction T_i only if $j = i + 1$ or $i = m$ and $j = 1$.*

762 **Proof.** If \mathcal{P} is not robust against RC, then there is a database \mathbf{D} and a multiversion split
 763 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
 764 conflict quadruples C for a set of transactions \mathcal{T} that is consistent with \mathcal{P} and \mathbf{D} having the
 765 properties of Definition 6.

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

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

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

777 For the construction, we proceed by case distinction.

778 IF $i \neq 1$ AND $j \neq 1$, we construct a schedule s' from s by removing all operations from
 779 transactions T_h with $i < h < j$. Notice that we remove at least one transaction, since
 780 $i < i + 1 < j$. We can derive a sequence of conflict quadruples C' from C by removing all
 781 occurrences of these transactions T_h and adding the conflict quadruple (T_i, b'_i, a'_j, T_j) instead.
 782 By construction, s' is a multiversion split schedule based on C' over a set of transactions
 783 consistent with \mathcal{P} and \mathbf{D} . It remains to show that the newly constructed schedule s' has the
 784 properties of Definition 6. The latter is straightforward since C and C' agree on their first
 785 and last quadruple, due to assumption $i \neq 1$ and $j \neq 1$.

786 IF $i = 1$, it follows that $i < j$ and thus $j \neq 1$. Then, we construct a schedule s' from s by
 787 removing all operations from transactions T_h with $i < h < j$ and updating the prefix and
 788 postfix of T_1 , now based on b'_i . Notice that we again remove at least one transaction, since
 789 $i < i + 1 < j$ and that we can derive a sequence of conflict quadruples C' from C in the same
 790 way as before, by removing all occurrences of these transactions T_h and adding the conflict
 791 quadruple (T_i, b'_i, a'_j, T_j) instead. By construction, s' is a multiversion split schedule based
 792 on C' over a set of transactions consistent with \mathcal{P} and \mathbf{D} . It remains to show that the newly
 793 constructed schedule s' has the properties of Definition 6.

794 First, we observe that b'_i and a'_j are rw-conflicting, which immediately implies that
 795 Condition (3) is true for s' . The argument is by exclusion. Indeed, if b'_i and a'_j would be
 796 ww-conflicting, then $b'_i \ll_s a'_j$ implying $b'_i <_s a'_j$ (due to the assumed read-last committed)
 797 and thus $b'_i \leq_s b_1$, which is not allowed by condition (1) on s . It follows from a similar
 798 argument that b'_i and a'_j are not wr-conflicting: both $b'_i = v_s(a'_j)$ and $b'_i \ll_s v_s(a'_j)$ imply
 799 $b'_i <_s C_1 <_s a'_j$, which contradicts with C_1 being the last operation in s .

800 Since b'_i is rw-conflicting with a'_j , we have $v_s(b'_i) \ll_s a'_j$, implying $b'_i <_s a'_j$ (due to read-
 801 last-committed and the structure of a multiversion split schedule), thus $b'_i \leq_s b_1$. Therefore,
 802 condition (1) again transfers from s to s' . For similar reasons condition (2) applies on s' : If
 803 $b_1 <_{T_1} a_1$ then $b'_i \leq_{T_1} b_1 <_{T_1} a_1$,

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

809 In this schedule s' , condition (1) and (3) transfer from s by its construction. To see that

810 condition (2) is true on s' , simply notice that if b'_i and a'_1 are ww or wr-conflicting, then
 811 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
 812 $b_1 <_s a'_1$, implying $b_1 <_{s'} a'_1$. ◀

813 A.3 If direction for the proof of Theorem 10

814 A.3.1 First step

815 Next, we show that, if there exists a multiversion split schedule for the set of transaction
 816 templates in Figure 7 for some set \mathcal{D} of dominoes, then this schedule is always a schedule-
 817 encoding of a sequence of dominoes in \mathcal{D} .

818 ► **Proposition 21.** *Let \mathcal{D} be a set of dominoes. If there is a multiversion split schedule s for
 819 a set of transactions consistent with the transaction template in Figure 7 for \mathcal{D} and some
 820 database \mathbf{D} , then this schedule s is a schedule-encoding of some sequence \mathbf{d} of dominoes in \mathcal{D} .*

821 For the proof, let \mathbf{D} be a database and $s = \text{prefix}_{b_1}(T_1) \cdot T_2 \cdot \dots \cdot T_m \cdot \text{postfix}_{b_1}(T_1)$
 822 a multiversion split schedule for a set of transactions \mathcal{T} consistent with \mathcal{P} and \mathbf{D} , with
 823 the conditions of Lemma 20 and based on some sequence of conflict quadruples $C =$
 824 $(T_1, b_1, a_2, T_2), (T_2, b_2, a_3, T_3) \dots, (T_m, b_m, a_1, T_1)$. We show through a sequence of properties
 825 (Lemmas 23,24, 25, and 26), that s is a schedule-encoding of a sequence \mathbf{d} of dominoes in \mathcal{D} .

826 As a first property (22), we observe that transaction templates in \mathcal{P} heavily constrain
 827 the possible variable instantiations. For transaction template Split, for example, a variable
 828 mapping depends entirely on the choice of the value for variable I. Since Lemma 20 forbids
 829 the presence of duplicate transactions in \mathcal{T} , two transactions T_i and T_j (with $i \neq j$) based
 830 on transaction template Split cannot agree on their choice for variable I in s . By applying
 831 this argument to other transaction templates, we obtain the following corollary of Lemma 20.
 832 Here, for each transaction T_i in s , we write τ_i to denote the transaction template in \mathcal{P} that
 833 it is based on, and by μ_i the associated variable mapping for τ_i , with $\mu_i(\tau_i) = T_i$.

834 ► **Lemma 22.** *for two transactions T_i and T_j in s , with $i \neq j$:*

- 835 ■ *if T_i and T_j are based on Split, then $\mu_i(\mathbf{I}) \neq \mu_j(\mathbf{I})$;*
- 836 ■ *if T_i and T_j are based on First then $\mu_i(\mathbf{S}_1) \neq \mu_j(\mathbf{S}_1)$;*
- 837 ■ *if T_i and T_j are based on Last then $\mu_i(\mathbf{B}) \neq \mu_j(\mathbf{B})$ and $\mu_i(\mathbf{C}) \neq \mu_j(\mathbf{C})$;*
- 838 ■ *if T_i and T_j are based on domino transaction templates then $\mu_i(\mathbf{B}) \neq \mu_j(\mathbf{B})$ and $\mu_i(\mathbf{B}_{next}) \neq$
 839 $\mu_j(\mathbf{B}_{next})$.*

840 We conclude the proof of Proposition 21 with the necessary arguments that s is indeed a
 841 schedule-encoding for some sequence of dominoes.

842 ► **Lemma 23.** *Transaction T_1 is based on Split, T_2 is based on First, and $\mu_1(\mathbf{I}) = \mu_2(\mathbf{I})$,
 843 $\mu_1(\mathbf{X}_1) \neq \mu_2(\mathbf{X}_2)$, and $\mu_1(\mathbf{S}_1) = \mu_2(\mathbf{S}_1) \neq \mu_2(\mathbf{S}_0)$.*

844 **Proof.** Since b_1 and a_2 are rw-conflicting (cf, Definition 6), and there are no updates in the
 845 considered transaction templates, operation b_1 must be a read. Since InitialConflict is the only
 846 type allowing for conflicts involving a read, it is immediate that T_1 must be based on Split and
 847 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
 848 follows that $\mu_1(\mathbf{S}_1) = \mu_2(\mathbf{S}_1)$. From Definition 6, particularly that there is no ww-conflict
 849 between a write operation in $\text{prefix}_{b_1}(T_1)$ and a write operation in any of the transactions
 850 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
 851 onto \mathbf{X}_1 in transaction template Split and \mathbf{S}_0 onto \mathbf{X}_2 in transaction template First, implies
 852 $\mu_1(\mathbf{S}_1) \neq \mu_2(\mathbf{S}_0)$. ◀

853 ► **Lemma 24.** *There is a transaction T_3 in s and it is based on a domino transaction*
 854 *template, with $\mu_3(\mathbf{S}_1) = \mu_2(\mathbf{S}_1)$.*

855 **Proof.** First, suppose towards a contradiction that $m = 2$. We already know from Lemma 23
 856 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 $b_2 = a_2 = \mathbf{W}[\mu_2(\mathbf{I})]$,
 857 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 ,
 858 which is the desired contradiction.

859 The remainder of the proof is by exclusion. Transaction T_3 is not based on transaction
 860 template First, because all possible conflicts between T_2 and an instantiation of transaction
 861 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
 862 imply the equality $\mu_2(\mathbf{S}_1) = \mu_3(\mathbf{S}_1)$ (through functional constraints $\mathbf{I} = f_{\text{defines}}(\mathbf{X}_2)$, $\mathbf{S}_1 =$
 863 $f_{\text{future-solution-string}}(\mathbf{B})$, and $\mathbf{S}_1 = f_{\text{final-dominos-string}}(\mathbf{I})$), which is forbidden by Lemma 22.
 864 The argument that transaction T_3 cannot based on transaction template Split is similar: every
 865 possible conflict between T_2 and an instantiation of Split implies $\mu_1(\mathbf{I}) = \mu_2(\mathbf{I}) = \mu_3(\mathbf{I})$ either
 866 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,
 867 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,
 868 $\mu_1(\mathbf{I}) = \mu_3(\mathbf{I})$ is forbidden by Lemma 22. Finally, to see that T_3 is not based on transaction
 869 template Last, we observe that a conflict between T_2 and an instantiation of transaction
 870 template Last must be ww-conflicting involving variables \mathbf{B} , thus with $\mu_2(\mathbf{B}) = \mu_3(\mathbf{B})$. Then,
 871 $\mu_2(\mathbf{S}_0) = \mu_3(\mathbf{S}_t)$, due to functional constraint $\mathbf{S}_0 = f_{\text{top-string}}(\mathbf{B})$ in T_2 and $\mathbf{S}_t = f_{\text{top-string}}(\mathbf{B})$
 872 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
 873 and T_3 . However, we also have $\mu_3(\mathbf{S}_1) = \mu_3(\mathbf{S}_t)$, due to constraints $\mathbf{S}_1 = f_{\text{solution-string}}(\mathbf{B})$
 874 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
 875 contradicts with earlier proven Lemma 23. We conclude that T_3 is indeed based on a
 876 domino transaction template. Therefore, the conflict quadruple (T_2, b_2, a_3, T_3) must admit
 877 ww-conflicting operations over variable \mathbf{B} in T_2 and either variable \mathbf{B} or \mathbf{B}_{next} in T_3 . We
 878 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
 879 $\mu_3(\mathbf{S}_1) = f_{\text{future-solution-string}}(\mu_3(\mathbf{B}_{\text{next}}))$, thus independent of the variable \mathbf{B}_{next} or \mathbf{B} in T_3 ,
 880 we have $\mu_2(\mathbf{S}_1) = \mu_3(\mathbf{S}_1)$. ◀

881 ► **Lemma 25.** *For a transaction T_i , with $i \geq 4$, for which all T_j 's, with $j \in \{3, \dots, i-1\}$,*
 882 *are based on domino transaction templates, transaction T_i is based on a domino transaction*
 883 *template or on transaction template Last. Furthermore $\mu_2(\mathbf{S}_1) = \mu_i(\mathbf{S}_1)$.*

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

887 For this, observe that $\mu_2(\mathbf{S}_1) = \mu_{i-1}(\mathbf{S}_1)$. Indeed, every conflict quadruple $(T_i, b_i, a_{i+1}, T_{i+1})$,
 888 with $i \in \{3, \dots, i-1\}$, admits ww-conflicting operations with variables of type Domin-
 889 oSequence. No matter if the conflict is via a variable \mathbf{B} or \mathbf{B}_{next} , the constraints $\mathbf{S}_1 =$
 890 $f_{\text{future-solution-string}}(\mathbf{B})$ and $\mathbf{S}_1 = f_{\text{future-solution-string}}(\mathbf{B}_{\text{next}})$ ensure $\mu_2(\mathbf{S}_1) = \mu_{i-1}(\mathbf{S}_1)$.

891 Now, assume towards a contradiction that T_{i+1} is based on First, thus admitting a
 892 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
 893 $b_i = \mu_i(\mathbf{B}_{\text{next}})$ 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
 894 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
 895 $\mu_i(\mathbf{S}_1) = \mu_2(\mathbf{S}_1)$, this contradict with Lemma 22. We conclude that T_i is indeed based on
 896 a domino transaction template or on transaction template Last. That $\mu_{i-1}(\mathbf{S}_1) = \mu_i(\mathbf{S}_1)$
 897 follows again from the constraints using function $f_{\text{future-solution-string}}$. ◀

898 ► **Lemma 26.** *If T_i is based on transaction template Last, then $i = m$.*

899 **Proof.** Let T_j be the transaction following T_i . We already know about T_j that either $j = 1$
 900 or must be a transaction that is different to all foregoing transactions T_1, \dots, T_i (due to
 901 Lemma 20).

902 We first show, by exclusion, that transaction T_j is based on Split: Transaction T_j cannot
 903 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
 904 Lemma 22. Similarly, transaction T_j cannot be based on First, as then $\mu_i(\mathbf{B}) = \mu_j(\mathbf{B})$ imply-
 905 ing $\mu_2(\mathbf{S}_1) = \mu_i(\mathbf{S}_1) = \mu_j(\mathbf{S}_1)$, due to the constraints involving function $f_{\text{future-solution-string}}$.
 906 Finally, transaction T_j cannot be based on a domino transaction template, because then
 907 $\mu_j(\mathbf{B}_{\text{next}}) = \mu_i(\mathbf{B}) = \mu_j(\mathbf{B})$ 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 contra-
 908 dicting Lemma 22. We can thus indeed conclude that transaction T_j is based on Split.

909 To see that $j = 1$, recall that $\mu_1(\mathbf{S}_1) = \mu_{i-1}(\mathbf{S}_1)$ and the only possible conflict between
 910 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
 911 $\mu_{i-1}(\mathbf{B}_{\text{next}}) = \mu_i(\mathbf{B})$ and function $f_{\text{future-solution-string}}$. From this it follows that $\mu_1(\mathbf{I}) = \mu_j(\mathbf{I})$
 912 through $\mathbf{I} = (f_{\text{defines}} \circ f_{\text{is-non-empty}})(\mathbf{S}_1)$ in transaction template Split. That $j = 1$ then
 913 follows from Lemma 22. \blacktriangleleft

914 A.3.2 Final step

915 Finally, we show that if there is a multiversion split schedule with the properties of Lemma 20
 916 that is a schedule-encoding for a sequence of dominoes \mathbf{d} , then this sequence \mathbf{d} is also a
 917 solution to the respective PCP problem. The next Proposition thus finalized the proof for
 918 the if-direction of Theorem 10.

919 **► Proposition 27.** *Let \mathcal{D} be a set of dominoes. Let s be a multiversion split schedule with
 920 the properties of Lemma 20 that is consistent with the transaction templates in Figure 7 for
 921 \mathcal{D} and with some database \mathbf{D} . If s is a schedule-encoding of a sequence \mathbf{d} of dominoes in \mathcal{D} ,
 922 then \mathbf{d} is a solution for the PCP problem on input \mathcal{D} .*

923 **Proof.** Let $a_1 a_2 \dots a_h$ and $b_1 b_2 \dots b_k$ be the two strings (with $a_i, b_i \in \Sigma$) obtained by reading
 924 from left to right, symbol by symbol, the values on the top, respectively, the bottom of
 925 dominoes d_1, \dots, d_r . Let us say that $a_1 a_2 \dots a_h = \mathbf{a}_1 \mathbf{a}_2 \dots \mathbf{a}_r$ and $b_1 b_2 \dots b_k = \mathbf{b}_1 \mathbf{b}_2 \dots \mathbf{b}_r$.
 926 Notice that h and k are not necessarily equal to r as the top and bottom value of an individual
 927 domino can be of different length.

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

$$930 \quad \alpha_i := (f_{\text{append-}a_i} \circ f_{\text{append-}a_{i-1}} \circ \dots \circ f_{\text{append-}a_1})(\mu_2(\mathbf{S}_0)).$$

$$931 \quad \beta_j := (f_{\text{append-}b_j} \circ f_{\text{append-}b_{j-1}} \circ \dots \circ f_{\text{append-}b_1})(\mu_2(\mathbf{S}_0)).$$

933 First, we show that

$$934 \quad \alpha_h = \mu_2(\mathbf{S}_1) = \beta_k. \tag{1}$$

936 This result follows from the assumed structure of schedule s . More precisely, since an instan-
 937 tiation of First with an instantiation of Domino $_{\mathbf{d}_1}$ can only have conflicts on instantiations of
 938 $\mathbf{W}[\mathbf{B} : \text{DominoSequence}]$, we have $\mu_2(\mathbf{B}) = \mu_3(\mathbf{B})$, from which it follows that $\mu_2(\mathbf{S}_1) = \mu_3(\mathbf{S}_1)$.

939 For every individual instantiation of Domino $_{\mathbf{d}_i}$ in s , we have that $f_{\text{append-}a_i^{\ell_a}} \circ \dots \circ$
 940 $f_{\text{append-}a_i^1}(\mu_i(\mathbf{S}_t)) = \mu_i(\mathbf{S}_{t\mathbf{a}_i})$ and $f_{\text{append-}b_i^{\ell_b}} \circ \dots \circ f_{\text{append-}b_i^1}(\mu_i(\mathbf{S}_b)) = \mu_i(\mathbf{S}_{bb_i})$, with $\mathbf{a}_i =$
 941 $a_1 a_2 \dots a_{\ell_a}$ and $\mathbf{b}_i = b_1 b_2 \dots b_{\ell_b}$.

942 For transactions T_i , with $i \in \{3, \dots, m+1\}$, (thus representing an instantiation of
 943 $\text{Domino}_{\mathbf{d}_{i-2}}$ which is followed in s by an instantiation of $\text{Domino}_{\mathbf{d}_{i-1}}$), the only possible con-
 944 flict is between the instantiation of $\mathbb{W}[\mathbf{B}_{\text{next}} : \text{DominoSequence}](T_i)$ and of $\mathbb{W}[\mathbf{B} : \text{DominoSequence}]$
 945 in (T_{i+1}) – notice that this is indeed the only option due to Lemma 22 – thus with $\mu_i(\mathbf{B}_{\text{next}}) =$
 946 $\mu_{i+1}(\mathbf{B})$, implying $\mu_i(\mathbf{S}_{\mathbf{t}_{\mathbf{a}_i}}) = \mu_{i+1}(\mathbf{S}_t)$, $\mu_i(\mathbf{S}_{\mathbf{b}_{\mathbf{b}_i}}) = \mu_{i+1}(\mathbf{S}_b)$, and $\mu_i(\mathbf{S}_1) = \mu_{i+1}(\mathbf{S}_1)$.

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

950 Combining the above equalities indeed proves Condition (1).

951 From Condition (1) we can now derive that,

$$952 \quad \alpha_i = \mu_2(\mathbf{S}_1) = \beta_i, \text{ for every } i \in \{1, \dots, \min\{h, k\}\}, \quad (2)$$

954 by following an analogous approach. Indeed, in every instantiation of Domino_i , there is a
 955 functional constraint for every application of the append function that requires its input to be
 956 the result of the detach function applied over its output, which indeed implies Condition (2).

957 To see that $k = h$, we observe that $k \neq h$ implies an application of the detach function
 958 over the instantiation of \mathbf{S}_1 (for which we already argued it has the same tuple assigned for
 959 every domino instantiation) for the shortest string, which contradicts with Condition (1)
 960 because such an application results in the same instantiation as \mathbf{S}_e , which can never equal
 961 the instantiation for \mathbf{S}_1 .

962 The desired result that the individual symbols in the top and bottom reads of dominoes in
 963 sequence \mathbf{d} are the same now follows from the functional constraint that every interpretation of
 964 a string sequence mapped via function f_{top} onto either the interpretation for \mathbf{S}_1 (representing
 965 symbol $1 \in \Sigma$) or \mathbf{S}_0 (representing symbol $0 \in \Sigma$). ◀

966 **B** Proofs for Section 5

967 **B.1** Proof for Lemma 13

968 Before proving the correctness of Lemma 13, we first present two additional lemmas that
 969 will be used in the correctness proof.

970 ▶ **Lemma 28.** *Let $(\text{Rels}, \text{Funcs})$ be a schema for which a disjoint partitioning of Funcs in
 971 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 $\text{dom}(g_i) =$
 972 $\text{range}(f_i)$ for every $(f_i, g_i) \in P$ and every schema graph $\text{SG}(\text{Rels}, \{h_1, h_2, \dots, h_n\})$ over the
 973 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:*

- 974 1. *there is no function name $f \in \text{Funcs}$ with $\text{dom}(f) = \text{range}(f)$; and*
- 975 2. *for every path in $\text{SG}(\text{Rels}, \text{Funcs})$, say visiting the nodes $R_1, R_2, \dots, R_{m-1}, R_m$, if $R_1 =$
 976 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
 977 f the edge from R_1 to R_2 and g the edge from R_{m-1} to R_m .*

978 **Proof.** Towards a contradiction, assume (1) does not hold. That is, there is a function name
 979 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
 980 (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
 981 $h_i \in (f_i, g_i)$ such that the resulting schema graph is a multi-tree. Indeed, in both cases, there
 982 is a self-loop on R , leading to the desired contradiction.

983 For (2), assume towards a contradiction that $R_2 \neq R_{m-1}$. Without loss of generality, we
 984 can assume that each node is visited only once in R_2, \dots, R_{m-1} . Otherwise, R_2, \dots, R_{m-1}
 985 contains a loop that can be removed from this sequence without altering R_2 and R_{m-1} . Since

986 $R_1, R_2, \dots, R_{m-1}, R_m$ is a path in $SG(\text{Rels}, \text{Funcs})$, there is a sequence of function names
 987 $F = e_1 \cdots e_{m-1}$ such that each e_i is an edge from R_i to R_{i+1} in $SG(\text{Rels}, \text{Funcs})$, implying
 988 $\text{dom}(e_i) = R_i$ and $\text{range}(e_i) = R_{i+1}$. By assumption that each type R_i occurs only once in
 989 R_2, \dots, R_{m-1} (notice that, for $i = 1$, this follows from Condition (2) of the lemma) and type
 990 $R_1 = R_m$ does not appear in R_2, \dots, R_{m-1} , there is no pair of function names e_i and e_j in F
 991 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
 992 name of each pair in P appears in F . But then we can choose $h_i = e_i$ for each such pair in P ,
 993 with e_i the function name appearing in F . Since F describes a cycle in $SG(\text{Rels}, \text{Funcs})$, the
 994 resulting schema graph restricted to these h_i cannot be a multi-tree, as it contains a cycle.

995 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
 996 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
 997 $\text{range}(f) = \text{dom}(g)$, as $R_1 = R_m$ and $R_2 = R_{m-1}$. If (f, g) is not a pair in P , then there are
 998 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
 999 $\text{range}(f) = \text{range}(g') = \text{dom}(f') = \text{dom}(g) = R_2$. Then we can choose f in (f, f') and g in
 1000 (g, g') . Since the resulting schema graph cannot be a multi-tree, as there is a cycle between
 1001 R_1 and R_2 , this choice leads to a contradiction. \blacktriangleleft

1002 **► Lemma 29.** *Let D be a sequence of potentially conflicting quadruples over $\mathcal{P} \in \mathbf{MTBTemp}$.*
 1003 *Then*

- 1004 1. $X \rightsquigarrow_\tau Y$ iff $Y \rightsquigarrow_\tau X$ for every pair of variables X and Y occurring in a template τ ; and
- 1005 2. $X \rightsquigarrow_D Y$ iff $Y \rightsquigarrow_D X$ for every pair of variables X and Y occurring in D .

1006 **Proof.** (1) We argue by induction on the definition of $X \rightsquigarrow_\tau Y$ that $X \rightsquigarrow_\tau Y$ implies $Y \rightsquigarrow_\tau X$.
 1007 The other direction is analogous. The base case is immediate, as $X = Y$ implies $Y \rightsquigarrow_\tau X$ by
 1008 definition. For the inductive case, assume a variable Z such that $Y = f(Z)$ is a constraint in τ
 1009 and $X \rightsquigarrow_\tau Z$. By the induction hypothesis, $Z \overset{F}{\rightsquigarrow}_\tau X$ for some sequence of function names F .
 1010 Since $\mathcal{P} \in \mathbf{MTBTemp}$, there is a constraint $Z = f'(Y)$ in τ as well. It follows that $Y \overset{F'}{\rightsquigarrow}_\tau X$
 1011 with $F' = f' \cdot F$.

1012 (2) We argue by induction on the definition of $X \rightsquigarrow_D Y$ that $X \rightsquigarrow_D Y$ implies $Y \rightsquigarrow_D X$. The
 1013 other direction is again analogous. The first base case is now immediate, as we already argued
 1014 that $X \rightsquigarrow_\tau Y$ implies $Y \rightsquigarrow_\tau X$. For the second base case, assume $X \overset{F}{\rightsquigarrow}_D Y$ and $(\tau_i, o_i, p_j, \tau_j)$
 1015 is a potentially conflicting quadruple in D with $\text{var}(o_i) = X$ and $\text{var}(p_j) = Y$ (the case for
 1016 $(\tau_j, o_j, p_i, \tau_i)$ is analogous). $Y \overset{F}{\rightsquigarrow}_D X$ then follows by definition. For the inductive case, let Z be
 1017 a variable such that $X \overset{F_1}{\rightsquigarrow}_D Z$ and $Z \overset{F_2}{\rightsquigarrow}_D Y$. Then by induction hypothesis $Z \overset{F_1}{\rightsquigarrow}_D X$ and $Y \overset{F_2}{\rightsquigarrow}_D Z$
 1018 for some sequence of function names F_1' and F_2' . By definition, $Y \overset{F'}{\rightsquigarrow}_D X$ with $F' = F_2' \cdot F_1'$. \blacktriangleleft

1019 **► Lemma 13.** *Let D be a sequence of potentially conflicting quadruples over $\mathcal{P} \in \mathbf{MTBTemp}$.*
 1020 *Then $X \approx_D Y$ implies $X \rightsquigarrow_D Y$ and $Y \rightsquigarrow_D X$. Furthermore, if $\text{type}(X) = \text{type}(Y)$ then $\bar{\mu}(X) = \bar{\mu}(Y)$*
 1021 *for every variable mapping $\bar{\mu}$ for D that is admissible for some database \mathbf{D} .*

1022 **Proof.** (1) Assuming $X \approx_D Y$, we first show by induction on the definition of connectedness
 1023 that $X \rightsquigarrow_D Y$. By Lemma 29, $Y \rightsquigarrow_D X$ then follows. For the base case, both $X \rightsquigarrow_D Y$ and
 1024 $Y \rightsquigarrow_D X$ imply $X \rightsquigarrow_D Y$, where the former is immediate and the latter is by Lemma 29. For the
 1025 inductive case, let Z be a variable with $X \approx_D Z$ and either $Z \rightsquigarrow_D Y$ or $Y \rightsquigarrow_D Z$. Again, $Z \overset{F_2}{\rightsquigarrow}_D Y$
 1026 for some sequence of function names F_2 is implied in both cases. By induction hypothesis,
 1027 $X \overset{F_1}{\rightsquigarrow}_D Z$ for some sequence of function names F_1 . As a result, $X \overset{F}{\rightsquigarrow}_D Y$ with $F = F_1 \cdot F_2$.

1028 (2) Next, let X and Y be two variables occurring in $\text{Trans}(D)$ with $X \approx_D Y$ and $\text{type}(X) =$
 1029 $\text{type}(Y)$ and let $\bar{\mu}$ be a variable mapping for D that is admissible for a database \mathbf{D} . We prove
 1030 that $\bar{\mu}(X) = \bar{\mu}(Y)$.

1031 We already argued that $X \approx_D Y$ implies $X \rightsquigarrow_D Y$. By definition of $X \rightsquigarrow_D Y$, there is a
 1032 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
 1033 variables X_i and X_{i+1} :

- 1034 (†) 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
 1035 some function name f ; or
 1036 (‡) $\text{type}(X_i) = \text{type}(X_{i+1})$ and there is a potentially conflicting quadruple $(\tau_j, o_j, p_k, \tau_k)$ in D
 1037 with either $\text{var}(o_j) = X_i$ and $\text{var}(p_k) = X_{i+1}$ or $\text{var}(p_k) = X_i$ and $\text{var}(o_j) = X_{i+1}$.

1038 In the remainder of this proof, we show that for each pair of variables X_i and X_j in this
 1039 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
 1040 follows immediately as $X = X_1$ and $Y = X_n$. Note that it suffices to show this property
 1041 only for pairs of variables X_i and X_j for which no variable X_k exists with $i < k < j$ and
 1042 $\text{type}(X_i) = \text{type}(X_j) = \text{type}(X_k)$. Indeed, if such an X_k exists, we can recursively argue that
 1043 $\bar{\mu}(X_i) = \bar{\mu}(X_k)$ and $\bar{\mu}(X_k) = \bar{\mu}(X_j)$. The argument is by induction on the number of variables
 1044 between X_i and X_j .

1045 If $j = i + 1$ (*base case*), then (‡) applies to X_i and X_j . Indeed, if (†) would apply instead,
 1046 then there would be a function name f with $\text{dom}(f) = \text{type}(X_i) = \text{type}(X_j) = \text{range}(f)$,
 1047 contradicting Condition (1) of Lemma 28. By definition of $\bar{\mu}$, we have $\bar{\mu}(X_i) = \bar{\mu}(X_j)$.

1048 Next, let $i + 1 < j$ (*inductive case*), and assume that $\bar{\mu}(X_k) = \bar{\mu}(X_\ell)$ for all X_k and X_ℓ
 1049 with $i < k \leq \ell < j$ and $\text{type}(X_k) = \text{type}(X_\ell)$ (*induction hypothesis*). From this sequence
 1050 X_i, \dots, X_j , we derive a sequence of function names $F = f_1 \cdots f_{m-1}$, where each function
 1051 name f_i is based on an application of (†) on adjacent variables (notice that applications
 1052 of (‡) do not result in a function name being added to F). By assumption on the types of
 1053 variables X_k with $i < k < j$, we have in particular $\text{type}(X_{i+1}) \neq \text{type}(X_i)$ and $\text{type}(X_{j-1}) \neq$
 1054 $\text{type}(X_j)$. This implies that (†) is applicable for X_i and X_{i+1} (respectively X_{j-1} and X_j).
 1055 Furthermore, X_i and X_{i+1} appear in the same template, say τ_i (respectively τ_j for X_{j-1} and
 1056 X_j), and $X_{i+1} = f_1(X_i) \in \Gamma(\tau_i)$ (respectively $X_j = f_{m-1}(X_{j-1}) \in \Gamma(\tau_j)$). By construction,
 1057 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
 1058 $\text{type}(X_i) = R_1$, $\text{type}(X_{i+1}) = R_2$, $\text{type}(X_{j-1}) = R_{m-1}$ and $\text{type}(X_j) = R_m$. Since this path
 1059 satisfies Condition 2 in Lemma 28, it follows that $\text{type}(X_{i+1}) = \text{type}(X_{j-1})$ and (f_1, f_{m-1})
 1060 is a pair in the pairwise partitioning of Funcs witnessing $\mathcal{P} \in \text{MTBTemp}$. By definition
 1061 of MTBTemp , $X_{i+1} = f_1(X_i) \in \Gamma(\tau_i)$ then implies $X_i = f_{m-1}(X_{i+1}) \in \Gamma(\tau_i)$. According to
 1062 the induction hypothesis, $\bar{\mu}(X_{i+1}) = \bar{\mu}(X_{j-1})$. Since $\bar{\mu}$ is admissible for \mathbf{D} , we conclude that
 1063 $\bar{\mu}(X_i) = f_{m-1}^{\mathbf{D}}(\bar{\mu}(X_{i+1})) = f_{m-1}^{\mathbf{D}}(\bar{\mu}(X_{j-1})) = \bar{\mu}(X_j)$. ◀

1064 B.2 Proof for Lemma 14

1065 ▶ **Lemma 14.** *For a multiversion split schedule s based on a sequence of conflicting quadruples*
 1066 *C over a set of transactions \mathcal{T} consistent with a $\mathcal{P} \in \text{MTBTemp}$ and a database \mathbf{D} , let $\bar{\mu}$*
 1067 *be the variable mapping for a sequence of potentially conflicting quadruples D over \mathcal{P} with*
 1068 *$\bar{\mu}(D) = C$. Then, a set \mathcal{S} of type mappings over disjoint ranges and a function $\varphi_{\mathcal{S}} : \mathbf{Var} \rightarrow \mathcal{S}$*
 1069 *exist with:*

- 1070 ■ $\bar{\mu}(X) = c(\text{type}(X))$ for every variable X , with $c = \varphi_{\mathcal{S}}(X)$;
- 1071 ■ $\varphi_{\mathcal{S}}(X) = \varphi_{\mathcal{S}}(Y)$ whenever $X \approx_D Y$; and,
- 1072 ■ $\varphi_{\mathcal{S}}(X) \neq \varphi_{\mathcal{S}}(Y)$ for every constraint $X \neq Y$ occurring in a template $\tau \in \text{Trans}(D)$.

1073 **Proof.** To aid the construction of \mathcal{S} and $\varphi_{\mathcal{S}}$, we first define a coloring function λ that assigns
 1074 a color to each tuple occurring in the schedule s such that the following holds: for every pair
 1075 of tuples \mathbf{t} and \mathbf{v} occurring in s :

- 1076 ■ connected tuples are mapped to the same color: if $\bar{\mu}(X) = \mathfrak{t}$, $\bar{\mu}(Y) = \mathfrak{v}$ and $X \approx_D Y$ for
 1077 some variables X, Y occurring in $\text{Trans}(D)$, then $\lambda(\mathfrak{t}) = \lambda(\mathfrak{v})$; and
 1078 ■ different tuples of the same type are mapped to different colors: if $\text{type}(\mathfrak{t}) = \text{type}(\mathfrak{v})$ and
 1079 $\mathfrak{t} \neq \mathfrak{v}$, then $\lambda(\mathfrak{t}) \neq \lambda(\mathfrak{v})$.

1080 Note that we can always construct such a function λ as by Lemma 13, it cannot be the case
 1081 that $\text{type}(\mathfrak{t}) = \text{type}(\mathfrak{v})$, $\mathfrak{t} \neq \mathfrak{v}$ and there is a pair of variables X, Y with $\bar{\mu}(X) = \mathfrak{t}$, $\bar{\mu}(Y) = \mathfrak{v}$,
 1082 and $X \approx_D Y$.

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

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

1083 where $\mathfrak{v}_{c,R}$ is an arbitrary tuple of type R not occurring in s or any other type mapping
 1084 c_β for $\beta \in \text{range}(\lambda)$. Define $\mathcal{S} = \{c_\alpha \mid \alpha \in \text{range}(\lambda)\}$. By construction, every type mapping
 1085 in \mathcal{S} is well defined and all type mappings are over disjoint ranges. Furthermore, $c_\alpha \neq c_\beta$
 1086 whenever $\alpha \neq \beta$.

1087 We now construct $\varphi_{\mathcal{S}}$ as follows: $\varphi_{\mathcal{S}}(X) = c_\alpha$ with $\alpha = \lambda(\bar{\mu}(X))$ for every variable X
 1088 occurring in $\text{Trans}(D)$. It remains to argue that $\varphi_{\mathcal{S}}$ indeed satisfies all properties stated in
 1089 Lemma 14. By construction of \mathcal{S} and $\varphi_{\mathcal{S}}$, we have $\bar{\mu}(X) = c(\text{type}(X))$ for every variable X ,
 1090 with $c = \varphi_{\mathcal{S}}(X)$. Towards the second property, notice that $X \approx_D Y$ implies $\varphi_{\mathcal{S}}(X) = c_{\lambda(\bar{\mu}(X))} =$
 1091 $c_{\lambda(\bar{\mu}(Y))} = \varphi_{\mathcal{S}}(Y)$ by definition of λ and $\varphi_{\mathcal{S}}$. For the last property, assume $X \neq Y$ occurs in a
 1092 template $\tau \in \text{Trans}(D)$ and $\text{type}(X) = \text{type}(Y)$. Since $\bar{\mu}$ is admissible for database \mathbf{D} , $\bar{\mu}(X) \neq$
 1093 $\bar{\mu}(Y)$. Then, by definition of λ and $\varphi_{\mathcal{S}}$, we have $\varphi_{\mathcal{S}}(X) = c_{\lambda(\bar{\mu}(X))} \neq c_{\lambda(\bar{\mu}(Y))} = \varphi_{\mathcal{S}}(Y)$. ◀

1094 B.3 Proof for Lemma 15

1095 ► **Lemma 30.** *Let D be a sequence of potentially conflicting quadruples. If $X \approx_D Y$ and*
 1096 *$Y \approx_D Z$ then $X \approx_D Z$ for every triple of variables X, Y, Z occurring in $\text{Trans}(D)$.*

1097 **Proof.** According to Lemma 13, $X \approx_D Y$ and $Y \approx_D Z$ imply respectively $X \rightsquigarrow_D Y$ and $Y \rightsquigarrow_D Z$.
 1098 By definition, $X \rightsquigarrow_D Z$ and hence $X \approx_D Z$. ◀

1099 ► **Lemma 15.** *Let $\mathcal{P} \in \text{MTBTemp}$ and let $\mathcal{S} = \{c_1, c_2, c_3, c_4\}$ be a set consisting of four*
 1100 *type mappings with disjoint ranges. Then, \mathcal{P} is not robust against RC iff there is a sequence*
 1101 *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*
 1102 *each quintuple $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$ in E :*

- 1103 1. o_i and p_i are operations in τ_i , and $c_{o_i}, c_{p_i} \in \mathcal{S}$;
- 1104 2. $X_i \not\approx_{\tau_i} Y_i$ for each constraint $X_i \neq Y_i$ in τ_i ;
- 1105 3. $c_{o_i} = c_{p_i}$ if $o_i \approx_{\tau_i} p_i$;
- 1106 4. $c_{o_i} \neq c_{p_i}$ if there is a constraint $X_i \neq Y_i$ in τ_i with $X_i \approx_{\tau_i} \text{var}(o_i)$ and $Y_i \approx_{\tau_i} \text{var}(p_i)$;
- 1107 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
 1108 in τ_i potentially ww-conflicting with an operation o'_1 in $\text{prefix}_{o_1}(\tau_1)$ with $\text{var}(o'_i) \approx_{\tau_i} \text{var}(q_i)$
 1109 and $\text{var}(o'_1) \approx_{\tau_1} \text{var}(q_1)$.

1110 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
 1111 E with $j = i + 1$, or $i = m$ and $j = 1$:

- 1112 6. o_i is potentially conflicting with p_j and $c_{o_i} = c_{p_j}$;
- 1113 7. if $i = 1$ and $j = 2$, then o_1 is potentially rw-conflicting with p_2 ; and
- 1114 8. if $i = m$ and $j = 1$, then $o_1 <_{\tau_1} p_1$ or o_m is potentially rw-conflicting with p_1 .

1115 **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 con-
 1116 flicting quadruples derived from E . Notice in particular that D is indeed a sequence of

1117 potentially conflicting quadruples by (1) and (6). We construct a variable mapping $\bar{\mu}$ for
 1118 D admissible for a database \mathbf{D} such that the sequence of conflicting quadruples $C = \bar{\mu}(D)$
 1119 satisfies the conditions in Definition 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 :

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

1120 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 $\text{var}(o_i) = \text{var}(p_i) = \mathbf{X}$,
 1121 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
 1122 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
 1123 \mathbf{Y} for which $\varphi_{\mathcal{S}}$ is defined. From Lemma 13, it follows that $\mathbf{X} \rightsquigarrow_D \mathbf{Y}$ whenever $\mathbf{X} \approx_D \mathbf{Y}$. Let
 1124 τ_i and τ_j be the template in which respectively \mathbf{X} and \mathbf{Y} occur. The argument is now by
 1125 induction on the definition of $\mathbf{X} \rightsquigarrow_D \mathbf{Y}$:

- 1126 ■ 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);
- 1127 ■ 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 $(\tau_j, o_j, p_i, \tau_i) \in D$
 1128 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);
- 1129 ■ Otherwise, if $\mathbf{X} \rightsquigarrow_D \mathbf{Z}$ and $\mathbf{Z} \rightsquigarrow_D \mathbf{Y}$ for some variable \mathbf{Z} , then by induction $\varphi_{\mathcal{S}}(\mathbf{X}) = \varphi_{\mathcal{S}}(\mathbf{Z}) =$
 1130 $\varphi_{\mathcal{S}}(\mathbf{Y})$.

1131 By Lemma 30, \approx_D is an equivalence relation. For \mathbf{X} occurring in $\text{Trans}(D)$, denote by
 1132 $[\mathbf{X}]$ the equivalence class of \mathbf{X} . Let \mathcal{S}' be obtained by extending \mathcal{S} with a type mapping $c_{[\mathbf{X}]}$
 1133 for each equivalence class where no variable $\mathbf{Y} \in [\mathbf{X}]$ is defined in $\varphi_{\mathcal{S}}$. Furthermore, each of
 1134 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:

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

1135 Notice, furthermore, that in the second case \mathbf{X} might be connected in D to multiple variables
 1136 for which $\varphi_{\mathcal{S}}$ is defined, say \mathbf{Y}_1 and \mathbf{Y}_2 . Then, by Lemma 30, $\mathbf{Y}_1 \approx_D \mathbf{Y}_2$ and hence $\varphi_{\mathcal{S}}(\mathbf{Y}_1) =$
 1137 $\varphi_{\mathcal{S}}(\mathbf{Y}_2)$. We therefore conclude that $\varphi_{\mathcal{S}'}(\mathbf{X})$ is well defined. We argue that $\varphi_{\mathcal{S}'}(\mathbf{X}) = \varphi_{\mathcal{S}'}(\mathbf{Y})$
 1138 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
 1139 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
 1140 \mathbf{X} , then $\varphi_{\mathcal{S}'}(\mathbf{X}) = \varphi_{\mathcal{S}}(\mathbf{X}) = \varphi_{\mathcal{S}'}(\mathbf{Y})$ by construction of $\varphi_{\mathcal{S}'}$. If $\varphi_{\mathcal{S}}$ is not defined for both \mathbf{X}
 1141 and \mathbf{Y} , then either there exists a variable \mathbf{Z} for which $\varphi_{\mathcal{S}}$ is defined and \mathbf{Z} is connected in
 1142 D to one of these two variables, say \mathbf{X} , or no such variable \mathbf{Z} exists. In the former case,
 1143 $\mathbf{Z} \approx_D \mathbf{Y}$ follows from Lemma 30, implying $\varphi_{\mathcal{S}'}(\mathbf{X}) = \varphi_{\mathcal{S}}(\mathbf{Z}) = \varphi_{\mathcal{S}'}(\mathbf{Y})$. In the latter case,
 1144 $\varphi_{\mathcal{S}'}(\mathbf{X}) = c_{[\mathbf{X}]} = c_{[\mathbf{Y}]} = \varphi_{\mathcal{S}'}(\mathbf{Y})$ by construction of $\varphi_{\mathcal{S}'}$.

1145 We now define the variable mapping $\bar{\mu}$ from $\varphi_{\mathcal{S}'}$ as $\bar{\mu}(\mathbf{X}) = c(\text{type}(\mathbf{X}))$ for each variable \mathbf{X} ,
 1146 where $c = \varphi_{\mathcal{S}}(\mathbf{X})$. Next, we construct the database \mathbf{D} . For each template τ_i and corresponding
 1147 variable mapping μ_i in $\bar{\mu}$, we add all tuples in $\mu_i(\tau_i)$ to the database \mathbf{D} . Furthermore,
 1148 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
 1149 defined for each function $f^{\mathbf{D}}$. Towards a contradiction, assume we have $f^{\mathbf{D}}(\bar{\mu}(\mathbf{X}_i)) = \bar{\mu}(\mathbf{Y}_i)$
 1150 witnessed by a template τ_i and $f^{\mathbf{D}}(\bar{\mu}(\mathbf{X}_j)) = \bar{\mu}(\mathbf{Y}_j)$ witnessed by a template τ_j , where
 1151 $\bar{\mu}(\mathbf{X}_i) = \bar{\mu}(\mathbf{X}_j)$, 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
 1152 $\varphi_{\mathcal{S}'}(\mathbf{Y}_i) = \varphi_{\mathcal{S}'}(\mathbf{X}_i) = \varphi_{\mathcal{S}'}(\mathbf{X}_j) = \varphi_{\mathcal{S}'}(\mathbf{Y}_j)$. Then, $\bar{\mu}(\mathbf{Y}_i) = \bar{\mu}(\mathbf{Y}_j)$, leading to a contradiction.

1153 In order to argue that $\bar{\mu}$ is indeed admissible for \mathbf{D} , it remains to show that for each
 1154 constraint $\mathbf{X} \neq \mathbf{Y}$ in a template τ_i , we have $\bar{\mu}(\mathbf{X}) \neq \bar{\mu}(\mathbf{Y})$. Again towards a contradiction,
 1155 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
 1156 definition of $\bar{\mu}$, we have $\varphi_{\mathcal{S}'}(\mathbf{X}) = \varphi_{\mathcal{S}'}(\mathbf{Y})$. It follows from $\varphi_{\mathcal{S}'}$ that either $\mathbf{X} \approx_{\tau_i} \mathbf{Y}$; or
 1157 $\mathbf{X} \approx_{\tau_i} \text{var}(o_i)$ (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
 1158 $c_{o_i} = c_{p_i}$. However, the former is contradicted by (2) and the latter by (4). We therefore
 1159 conclude that $\bar{\mu}$ is admissible for \mathbf{D} , as it satisfies all constraints.

1160 It remains to argue that the sequence of conflicting quadruples $C = \bar{\mu}(D)$ satisfies all
 1161 conditions stated in Definition 6. The second and third condition are immediate by respect-
 1162 ively (8) and (7). Towards a contradiction, assume the first condition holds. Then, there is
 1163 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 $\bar{\mu}\tau_i$.
 1164 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 $(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$
 1165 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
 1166 ww-conflicting with o'_i and $\bar{\mu}(\mathbf{X}_1) = \bar{\mu}(\mathbf{X}_i)$. Then, $\varphi_{\mathcal{A}'}(\mathbf{X}_1) = \varphi_{\mathcal{A}'}(\mathbf{X}_i)$. By construction of $\varphi_{\mathcal{A}'}$,
 1167 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
 1168 $q_i \in o_i, p_i$, thereby contradicting (5).

1169 (*only if*) If \mathcal{P} is not robust against RC, then there exists a multiversion split schedule s based
 1170 on a sequence of conflicting quadruples C over a set of transactions \mathcal{T} consistent with \mathcal{P}
 1171 and a database \mathbf{D} . Let $\bar{\mu}$ be the variable mapping for a sequence of potentially conflicting
 1172 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 $\varphi_{\mathcal{S}}$
 1173 be as in Lemma 14.

From this sequence D and function $\varphi_{\mathcal{S}}$, we derive the sequence of quintuples $E =$
 $(\tau_1, o_1, \varphi_{\mathcal{S}}(\text{var}(o_1)), p_1, \varphi_{\mathcal{S}}(\text{var}(p_1))), \dots, (\tau_m, o_m, \varphi_{\mathcal{S}}(\text{var}(o_m)), p_m, \varphi_{\mathcal{S}}(\text{var}(p_m)))$. Let $\varphi'_{\mathcal{S}} =$
 $\{c_1, c_2, c_3, c_4\}$ 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 prop-
 erties stated in Lemma 15. First, we add $(\tau_1, o_1, c_1, p_1, c_k)$ to E' , where $c_k = c_1$ if
 $\varphi_{\mathcal{S}}(\text{var}(o_1)) = \varphi_{\mathcal{S}}(\text{var}(p_1))$, 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

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

1174 By construction, for every quintuple $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$ in E' we now have

- 1175 ■ $c_{o_i} = c_{p_i}$ iff $\varphi_{\mathcal{S}}(\text{var}(o_i)) = \varphi_{\mathcal{S}}(\text{var}(p_i))$; and
- 1176 ■ $c_{q_i} = c_{q_1}$ 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 $q_1 \in \{o_1, p_1\}$.

1177 It remains to argue that E' indeed satisfies all required properties. (1) is trivial by
 1178 construction. (2) If $\mathbf{X}_i \neq \mathbf{Y}_i$ is a constraint in τ_i , then $\varphi_{\mathcal{S}}(\mathbf{X}) \neq \varphi_{\mathcal{S}}(\mathbf{Y})$ and $\mathbf{X} \approx_{\tau_i} \mathbf{Y}$ according
 1179 to Lemma 14. (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 14, and hence
 1180 $c_{o_i} = c_{p_i}$. (4) Assume there is a constraint $\mathbf{X}_i \neq \mathbf{Y}_i$ in a template τ_i with $\mathbf{X}_i \approx_{\tau_i} \text{var}(o_i)$
 1181 and $\mathbf{Y}_i \approx_{\tau_i} \text{var}(p_i)$. By Lemma 14, $\varphi_{\mathcal{S}}(\mathbf{X}_i) = \varphi_{\mathcal{S}}(\text{var}(o_i)) \neq \varphi_{\mathcal{S}}(\text{var}(p_i)) = \varphi_{\mathcal{S}}(\mathbf{Y}_i)$, and
 1182 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 $i \neq 1$.
 1183 Assume towards a contradiction that there is an operation o'_i in τ_i potentially ww-conflicting
 1184 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

1185 then $\varphi_S(\text{var}(o'_i)) = \varphi_S(\text{var}(q_i)) = \varphi_S(\text{var}(q_1)) = \varphi_S(\text{var}(o'_1))$, implying that $\bar{\mu}(o'_i)$ is ww-
 1186 conflicting with $\bar{\mu}(o'_1)$, contradicting the properties of C stated in Definition 6. (6) is again
 1187 trivial by construction. (7) By Definition 6, $\bar{\mu}(o_1)$ is rw-conflicting with $\bar{\mu}(p_2)$ in C . Therefore,
 1188 o_1 is potentially rw-conflicting with p_2 . (8) By Definition 6, $\bar{\mu}(o_1) <_{\bar{\mu}(\tau_1)} \bar{\mu}(p_1)$ or $\bar{\mu}(o_m)$ is
 1189 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 . ◀

1190 B.4 NLOGSPACE upper bound

1191 Theorem 12 relies on the characterization provided by Lemma 15 and the ability to guess a
 1192 sequence of quintuples and verify with logarithmic space if it has the properties of Lemma 15.

1193 The algorithm goes as follows: We start by guessing three initial quintuples, representing
 1194 respectively the first, second and last quintuple of the possible sequence of quintuples as in
 1195 Lemma 15. Consistent with previously used notation, we refer to these quintuples by E_1 ,
 1196 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
 1197 the algorithm. They are only used to distinguish between the different considered quintuples
 1198 in the proof argument.

1199 We store all three quintuples using a logarithmic amount of space, by storing pointers
 1200 to the respective transaction templates in \mathcal{P} , the positions of operations in the respective
 1201 transaction templates, and the number 1, 2, 3 or 4 for the type mappings.

1202 At this point, we verify that Condition (7) and (8) are true, that Conditions (1-5) are
 1203 true for all chosen transaction templates and operations, and that Condition (6) is true for
 1204 τ_1 and τ_2 , and τ_2 and τ_m . We reject the guessed quintuples if any of the conditions is false.

1205 If all previous checks are true, we proceed by inserting another step. Let $i = 2$. We
 1206 guess a new quintuple E_{i+1} and verify that Condition (5) is true for τ_i and τ_{i+1} and that
 1207 Conditions (1-6) are true for τ_{i+1} and reject the entire construction if one of these conditions
 1208 failed. Notice that all Conditions, including Condition (5) can be checked easily, particularly
 1209 because quintuple E_1 is stored. To proceed, we discard quintuple E_i and store E_{i+1} instead,
 1210 thus without increasing the amount of space we use.

1211 If τ_{i+1} and τ_m (from quintuple E_m) have Condition (6), the algorithm emits an accept.
 1212 Indeed, then the sequence $E_1, \dots, E_i, E_{i+1}, E_m$ of guessed quintuples has all the properties
 1213 of Lemma 15. Otherwise, the algorithm proceeds with another insertion step, for $i = i + 1$.

1214 C Proofs for Section 6

1215 C.1 Proof for Lemma 17

1216 ▶ **Lemma 17.** *Let D be a sequence of potentially conflicting quadruples over a set of*
 1217 *transaction templates $\mathcal{P} \in \mathbf{AcycTemp}$. For every pair of variables X, Y occurring in $\text{Trans}(D)$,*
 1218 *if $X \xrightarrow{F} D Y$, then $\text{type}(X) \xrightarrow{F}_{SG} \text{type}(Y)$, with SG the corresponding schema graph.*

1219 **Proof.** Let τ_i and τ_j be the templates in D in which X and Y occur, respectively. The proof
 1220 is by induction on the definition of $X \xrightarrow{F} D Y$.

1221 (*Implication within the same template*) If $i = j$ and $X \xrightarrow{F} \tau_i Y$, then either $F = \varepsilon$ and
 1222 $X = Y$, 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
 1223 $F = F' \cdot f$. In the former case, $\text{type}(X) = \text{type}(Y)$, so $\text{type}(X) \xrightarrow{F}_{SG} \text{type}(Y)$ is immediate. In
 1224 the latter case, it follows by induction that $\text{type}(X) \xrightarrow{F'}_{SG} \text{type}(Z)$. Since $\text{dom}(f) = \text{type}(Z)$
 1225 and $\text{range}(f) = \text{type}(Y)$, it follows by definition that $\text{type}(Z) \xrightarrow{F'}_{SG} \text{type}(Y)$ and furthermore
 1226 $\text{type}(X) \xrightarrow{F} \tau_i \text{type}(Z)$ holds.

1227 (*Implication between templates*) If $F = \varepsilon$ and $(\tau_i, o_i, p_j, \tau_j)$ (respectively $(\tau_j, o_j, p_i, \tau_i)$)
 1228 is a potentially conflicting quadruple in D , with $\text{var}(o_i) = X$ and $\text{var}(p_j) = Y$ (respectively

1229 $\text{var}(p_i) = X$ and $\text{var}(o_j) = Y$, then $\text{type}(X) = \text{type}(Y)$ by definition of potentially conflicting
 1230 operations. So, $\text{type}(X) \xrightarrow{\varepsilon}_{SG} \text{type}(Y)$ is again immediate.

1231 (*Inductive case*) If $X \xrightarrow{F_1}_D Z$ and $Z \xrightarrow{F_2}_D Y$ for some variable Z with $F = F_1 \cdot F_2$, then
 1232 $\text{type}(X) \xrightarrow{F}_{SG} \text{type}(Z)$ and $\text{type}(Z) \xrightarrow{F}_{SG} \text{type}(Y)$ follow by induction. We conclude that
 1233 $\text{type}(X) \xrightarrow{F}_{SG} \text{type}(Y)$. ◀

1234 C.2 Proof for Theorem 16

1235 We call node S a *descendant* of node R and R an *ancestor* of S in an acyclic schema graph
 1236 SG . We write $R \xrightarrow{\varepsilon}_{SG} S$, with ε denoting the empty labeling, for the case $R = S$. This means
 1237 that a node is a descendant and ancestor of itself. When F is not relevant, we simply write
 1238 $R \rightsquigarrow_{SG} S$.

1239 Let c_R and c_S be two tuple-contexts for types R and S , respectively, such that S is a
 1240 descendant of R in SG , witnessed by the path $R \xrightarrow{F}_{SG} S$ in SG . We then say that c_S is
 1241 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
 1242 $S \xrightarrow{F'}_{SG} S'$ in SG . It should be noted that $R \xrightarrow{F \cdot F'}_{SG} S'$ is indeed a valid path in SG , as it
 1243 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
 1244 the schema graph SG , we will often write $c(F)$ as a shorthand notation for $c(R \xrightarrow{F}_{SG} S)$.

1245 Similar to Lemma 14 for sets of transaction templates admitting multi-tree bijectivity,
 1246 Lemma 34 shows that we can represent a counterexample schedule based on a sequence of
 1247 potentially conflicting quadruples D over an acyclic schema by assigning a tuple-context to
 1248 each variable in $\text{Trans}(D)$, taking special care when assigning contexts to variables connected
 1249 in D to make sure that they are properly related to each other. For a set of tuple-contexts \mathcal{A} ,
 1250 we refer to a (partial) function $\varphi_{\mathcal{A}} : \mathbf{Var} \rightarrow \mathcal{A}$ mapping (a subset of) variables in $\text{Trans}(D)$
 1251 to tuple-contexts in \mathcal{A} as a (*partial*) *context assignment* for D over \mathcal{A} . We furthermore say
 1252 that $\varphi_{\mathcal{A}}$ is a *total context assignment* for D over \mathcal{A} if $\varphi_{\mathcal{A}}$ is defined for every variable in
 1253 $\text{Trans}(D)$.

1254 Two variables X and Y occurring in $\text{Trans}(D)$ are *equivalent in D* , denoted $X \equiv_D Y$ if
 1255 ■ $X = Y$;
 1256 ■ there exists a pair of variables Z and W and a sequence of function names F with $Z \equiv_D W$,
 1257 $Z \xrightarrow{F}_D X$ and $W \xrightarrow{F}_D Y$; or
 1258 ■ there exists a variable Z with $X \equiv_D Z$ and $Z \equiv_D Y$.

1259 Similarly, two variables X and Y occurring in a transaction template τ are *equivalent in τ* ,
 1260 denoted $X \equiv_{\tau} Y$ if

1261 ■ $X = Y$;
 1262 ■ there exists a pair of variables Z and W in τ and a sequence of function names F with
 1263 $Z \equiv_{\tau} W$, $Z \xrightarrow{F}_{\tau} X$ and $W \xrightarrow{F}_{\tau} Y$; or
 1264 ■ there exists a variable Z with $X \equiv_{\tau} Z$ and $Y \equiv_{\tau} Z$.

1265 Intuitively, every variable mapping admissible for a given database will assign the same
 1266 tuple to equivalent variables (see Lemma 32). Due to these equivalent variables, the
 1267 assignment of a tuple to a variable X for a given database might imply the tuple assigned
 1268 to a variable Y , even if $X \rightsquigarrow_D Y$ does not hold. We capture this observation by introducing
 1269 *variable determination*, which is stronger than the previously defined variable implication.
 1270 Formally, a variable X *determines* a variable Y in D witnessed by a sequence of function
 1271 names F , denoted $X \xrightarrow{F}_D Y$ if:

1272 ■ $X \xrightarrow{F}_D Y$;
 1273 ■ $F = \varepsilon$ and $X \equiv_D Y$; or
 1274 ■ there exists a variable Z with $X \xrightarrow{F_1}_D Z$, $Z \xrightarrow{F_2}_D Y$ and $F = F_1 \cdot F_2$.

1275 For two variables X and Y in a template $\tau \in \text{Trans}(D)$ we furthermore say that X *determines*
 1276 Y in τ witnessed by a sequence of function names F , denoted $X \stackrel{F}{\rightsquigarrow}_{\tau} Y$ if:

- 1277 ■ $X \stackrel{F}{\rightsquigarrow}_{\tau} Y$;
- 1278 ■ $F = \varepsilon$ and $X \equiv_{\tau} Y$; or
- 1279 ■ there exists a variable Z with $X \stackrel{F_1}{\rightsquigarrow}_{\tau} Z$, $Z \stackrel{F_2}{\rightsquigarrow}_{\tau} Y$ and $F = F_1 \cdot F_2$.

1280 ► **Lemma 31.** *For a multiversion split schedule s based on a sequence of conflicting quadruples*
 1281 *C over a set of transactions \mathcal{T} consistent with a set of transaction templates \mathcal{P} and a database*
 1282 *D , let $\bar{\mu}$ be the variable mapping for a sequence of potentially conflicting quadruples D over*
 1283 *\mathcal{P} with $\bar{\mu}(D) = C$. Then, for every combination of variables W, X, Y, Z occurring in $\text{Trans}(D)$,*
 1284 *if $Z \stackrel{F}{\rightsquigarrow}_D X$, $W \stackrel{F}{\rightsquigarrow}_D Y$ and $\bar{\mu}(Z) = \bar{\mu}(W)$, then $\bar{\mu}(X) = \bar{\mu}(Y)$.*

1285 **Proof.** By definition of $Z \rightsquigarrow_D X$, there is a sequence of variables X_1, X_2, \dots, X_n with $X_1 = Z$
 1286 and $X_n = X$ such that for each pair of adjacent variables X_i and X_{i+1} :

1287 (†) 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
 1288 some function name f ; or

1289 (‡) $\text{type}(X_i) = \text{type}(X_{i+1})$ and there is a potentially conflicting quadruple $(\tau_j, o_j, p_k, \tau_k)$ in D
 1290 with either $\text{var}(o_j) = X_i$ and $\text{var}(p_k) = X_{i+1}$ or $\text{var}(p_k) = X_i$ and $\text{var}(o_j) = X_{i+1}$.

1291 Furthermore, the sequence F corresponds to the function names used in applications of (†).
 1292 Analogously, $W \rightsquigarrow_D Y$, implies a sequence of variables Y_1, Y_2, \dots, Y_m with $Y_1 = W$ and $Y_m = Y$

1293 with the same properties. Notice that the lengths of these two sequences of variables, namely
 1294 n and m , are not necessarily equal to each other and to the length of F due to possible

1295 applications of (‡). For a variable X_i in the sequence X_1, X_2, \dots, X_n , we denote the sequence of
 1296 function names derived from applications of (†) in the subsequence X_i, \dots, X_n by $\text{suffix}_F(X_i)$.

1297 Notice that $\text{suffix}_F(X_i)$ is indeed always a suffix of F , and that $\text{suffix}_F(X_1) = \text{suffix}_F(Y_1) = F$
 1298 and $\text{suffix}_F(X_n) = \text{suffix}_F(Y_m) = \varepsilon$.

1299 We argue by induction that for every $i \in [1, n]$ and $j \in [1, m]$, if $\text{suffix}_F(X_i) = \text{suffix}_F(Y_j)$
 1300 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
 1301 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
 1302 $\text{suffix}_F(X_{i+1}) = \text{suffix}_F(Y_{j+1})$. Then, we distinguish the following cases:

1303 ■ $\text{suffix}_F(X_i) = \text{suffix}_F(X_{i+1})$: This means that (†) applies to X_i and X_{i+1} , and there
 1304 is a potentially conflicting quadruple $(\tau_k, o_k, p_\ell, \tau_\ell)$ in D with either $\text{var}(o_k) = X_i$ and
 1305 $\text{var}(p_\ell) = X_{i+1}$ or $\text{var}(p_\ell) = X_i$ and $\text{var}(o_k) = X_{i+1}$. By definition of $\bar{\mu}$, we have $\bar{\mu}(X_{i+1}) =$
 1306 $\bar{\mu}(X_i)$ and by induction that $\bar{\mu}(X_{i+1}) = \bar{\mu}(Y_{j+1})$ implying that $\bar{\mu}(X_{i+1}) = \bar{\mu}(X_{j+1})$.

1307 ■ $\text{suffix}_F(Y_j) = \text{suffix}_F(Y_{j+1})$: similar as previous argument;

1308 ■ $\text{suffix}_F(X_i) \neq \text{suffix}_F(X_{i+1})$ and $\text{suffix}_F(Y_j) \neq \text{suffix}_F(Y_{j+1})$: Then, (†) applies to both
 1309 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
 1310 some f and F' . By induction, $\bar{\mu}(X_i) = \bar{\mu}(Y_j)$. Then, $X_{i+1} = f(X_i)$ is a constraint
 1311 in some template $\tau_k \in \text{Trans}(D)$ and $Y_{j+1} = f(Y_j)$ is a constraint in some template
 1312 $\tau_\ell \in \text{Trans}(D)$. Since $\bar{\mu}$ is admissible for database D and $\bar{\mu}(X_i) = \bar{\mu}(Y_j)$, it follows that
 1313 $\bar{\mu}(X_{i+1}) = f^D(\bar{\mu}(X_i)) = f^D(\bar{\mu}(Y_j)) = \bar{\mu}(Y_{j+1})$.

1314 ◀

1315 ► **Lemma 32.** *For a multiversion split schedule s based on a sequence of conflicting quadruples*
 1316 *C over a set of transactions \mathcal{T} consistent with a set of transaction templates \mathcal{P} and a database*
 1317 *D , let $\bar{\mu}$ be the variable mapping for a sequence of potentially conflicting quadruples D over*
 1318 *\mathcal{P} with $\bar{\mu}(D) = C$. Then, for every pair of variables X and Y occurring in templates τ_i and τ_j*
 1319 *in $\text{Trans}(D)$ respectively, if $X \equiv_D Y$, then $\bar{\mu}(X) = \bar{\mu}(Y)$.*

1320 **Proof.** The proof is by induction on the definition of $X \equiv_D Y$. (*base case*) If $X = Y$, then
 1321 the result is immediate. (*inductive cases*) If there are two variables Z and W and a sequence

1322 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
 1323 $\bar{\mu}(Z) = \bar{\mu}(W)$. The proof that $\bar{\mu}(X) = \bar{\mu}(Y)$ is now immediate by application of Lemma 31. If
 1324 instead there is a variable Z with $X \equiv_D Z$ and $Y \equiv_D Z$, then we can argue by induction that
 1325 $\bar{\mu}(X) = \bar{\mu}(Z)$ and $\bar{\mu}(Y) = \bar{\mu}(Z)$, and hence $\bar{\mu}(X) = \bar{\mu}(Y)$. \blacktriangleleft

1326 **► Definition 33.** Let D be a sequence of potentially conflicting quadruples, \mathcal{A} a set of tuple-
 1327 contexts and $\varphi_{\mathcal{A}}$ a partial context assignment for D over \mathcal{A} . We say that $\varphi_{\mathcal{A}}$ respects the
 1328 constraints of D if, for every two (not necessarily different) variables X and Y occurring in D
 1329 that $\varphi_{\mathcal{A}}$ is defined for, the following conditions are true, where $c_X = \varphi_{\mathcal{A}}(X)$ and $c_Y = \varphi_{\mathcal{A}}(Y)$:

- 1330 1. c_X is a tuple-context for $\text{type}(X)$;
- 1331 2. for every $X \xrightarrow{F}_D Z$ and $Y \xrightarrow{F}_D Z$, $c_X(F_1) = c_Y(F_2)$;
- 1332 3. for every $X \xrightarrow{F}_D W$ and $Y \xrightarrow{F}_D Z$ with $W \neq Z$ a constraint in a template τ , $c_X(F_1) \neq c_Y(F_2)$;
- 1333 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
 1334 template $\tau \in \text{Trans}(D)$;
- 1335 5. if $X \xrightarrow{F}_D Y$, then c_Y is a tuple-subcontext of c_X witnessed by F ; and
- 1336 6. for every pair of tuple-subcontexts c'_X and c'_Y of c_X and c_Y witnessed by respectively F_X and
 1337 F_Y , if $c_X(F_X) = c_Y(F_Y)$, then $c'_X = c'_Y$.

1338 **► Lemma 34.** For a multiversion split schedule s based on a sequence of conflicting quadruples
 1339 C over a set of transactions \mathcal{T} consistent with a set of transaction templates $\mathcal{P} \in \mathbf{AcycTemp}$
 1340 and a database \mathbf{D} , let $\bar{\mu}$ be the variable mapping for a sequence of potentially conflicting
 1341 quadruples D over \mathcal{P} with $\bar{\mu}(D) = C$. Then a set \mathcal{A} of tuple-contexts and a total context
 1342 assignment $\varphi_{\mathcal{A}}$ for D over \mathcal{A} exist with:

- 1343 ■ $\varphi_{\mathcal{A}}$ respects the constraints of D ; and
- 1344 ■ $\bar{\mu}(X) = c_X(\varepsilon)$ for every variable X , with $c_X = \varphi_{\mathcal{A}}(X)$.

1345 **Proof.** We first assign a tuple-context to each tuple in database \mathbf{D} , based on the functions
 1346 in \mathbf{D} . Let $(\text{Rels}, \text{Funcs})$ be the schema over which \mathcal{P} is defined. Since the schema graph
 1347 $SG(\text{Rels}, \text{Funcs})$ is acyclic, a total order $<_{SG}$ over Rels exists such that there is no path from
 1348 type R to type S in SG if $R <_{SG} S$. We now assign tuple-contexts to tuples based on the
 1349 order implied by $<_{SG}$. That is, we first consider all tuples of the type that is ordered first
 1350 by $<_{SG}$, then all tuples of the type that is ordered second, etc. If there are multiple tuples
 1351 of the same type, the relative order in which we handle them is not important. For each
 1352 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
 1353 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
 1354 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
 1355 construction, $c_{\mathbf{v}}$ is a tuple-subcontext of $c_{\mathbf{t}}$ witnessed by f .

1356 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
 1357 in $\text{Trans}(D)$. We argue by induction on the definition of \xrightarrow{F}_D that

$$1358 \quad c_{\mathbf{t}}(F) = \bar{\mu}(Y) \text{ for every } X \xrightarrow{F}_D Y \text{ (with } c_{\mathbf{t}} = \varphi_{\mathcal{A}}(X)). \quad (\dagger)$$

1359 If $X \xrightarrow{F}_D Y$, then by construction of $\varphi_{\mathcal{A}}$ and since $\bar{\mu}$ is admissible for \mathbf{D} , we have $c_{\mathbf{t}}(F) = \bar{\mu}(Y)$.
 1360 If $F = \varepsilon$ and $X \equiv_D Y$, then $c_{\mathbf{t}}(\varepsilon) = \bar{\mu}(X) = \bar{\mu}(Y)$ by Lemma 32. Otherwise, if there exists a
 1361 variable Z with $X \xrightarrow{F_1}_D Z$, $Z \xrightarrow{F_2}_D Y$ and $F = F_1 \cdot F_2$, then by induction $c_{\mathbf{t}}(F_1) = \bar{\mu}(Z) = \mathbf{v}$ and
 1362 $c_{\mathbf{v}}(F_2) = \bar{\mu}(Y)$, 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
 1363 follows.

1364 It remains to verify that $\varphi_{\mathcal{A}}$ indeed satisfies all required properties. By construction,
 1365 $\bar{\mu}(X) = c_{\mathbf{t}}(\varepsilon)$ with $c_{\mathbf{t}} = \varphi_{\mathcal{A}}(X)$, so we only need to show that $\varphi_{\mathcal{A}}$ respects the constraints of D
 1366 by verifying all properties in Definition 33. To this end, let X and Y be two variables occurring
 1367 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

1368 for $\text{type}(X)$. (2) $c_t(F_1) = \bar{\mu}(Z) = c_v(F_2)$ by (\dagger) . (3) If $W \neq Z$ is a constraint, then $\bar{\mu}(W) \neq \bar{\mu}(Z)$
 1369 as $\bar{\mu}$ is admissible for \mathbf{D} . By (\dagger) , it follows that $c_t(F_1) = \bar{\mu}(W) \neq \bar{\mu}(Z) = c_v(F_2)$. (4) If
 1370 $c_t(F_1) = c_v(F_2)$, then $\bar{\mu}(W) = c_t(F_1) = c_v(F_2) = \bar{\mu}(Z)$ by (\dagger) . Since $\bar{\mu}$ is admissible for \mathbf{D} ,
 1371 there cannot be a constraint $W \neq Z$. (5) If $X \xrightarrow{F} Y$, then by construction of the tuple-contexts
 1372 c_t and c_v it follows that c_v is a tuple-subcontext of c_t witnessed by F . (6) Let $c_{t'}$ and $c_{v'}$ be
 1373 tuple-subcontexts of c_t and c_v witnessed by respectively F_X and F_Y . If $c_t(F_X) = c_v(F_Y) = \mathbf{q}$
 1374 for some tuple \mathbf{q} , then by construction of c_t and c_v we have $c_{t'} = c_{v'} = c_{\mathbf{q}}$, with $c_{\mathbf{q}}$ the
 1375 tuple-context assigned to this tuple \mathbf{q} . ◀

1376 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 34 we can derive a
 1377 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 $c_{o_i} =$
 1378 $\varphi_{\mathcal{A}}(X_i)$ (respectively $c_{p_i} = \varphi_{\mathcal{A}}(Y_i)$) for $i \in [1, m]$ with o_i (respectively p_i) an operation
 1379 over variable X_i (respectively Y_i). Intuitively, this sequence of quintuples can be used to
 1380 reconstruct the original multiversion split schedule s . To this end, notice that we can derive
 1381 the original sequence of potentially conflicting quadruples D and a partial context assignment
 1382 $\varphi'_{\mathcal{A}}$ from E that is defined for each variable X_i occurring in either an operation o_i or p_i in
 1383 τ_i . We first show that we can extend this partial context assignment $\varphi'_{\mathcal{A}}$ to a total context
 1384 assignment respecting the constraints in D (Lemma 35), and then prove that such a total
 1385 context assignment respecting the constraints in D implies a variable assignment $\bar{\mu}$ such that
 1386 the $C = \bar{\mu}(D)$ is a valid sequence of conflicting quadruples (Lemma 36).

1387 ▶ **Lemma 35.** *Let $D = (\tau_1, o_1, p_2, \tau_2), \dots, (\tau_m, o_m, p_1, \tau_1)$ be a sequence of potentially con-*
 1388 *flicting quadruples over a set of transaction templates $\mathcal{P} \in \mathbf{AcycTemp}$ and $\varphi_{\mathcal{A}}$ a partial*
 1389 *context assignment defined for every variable X_i of o_i and Y_i of p_i in every τ_i . If*

- 1390 ■ $\varphi_{\mathcal{A}}$ respects the constraints of D ; and
- 1391 ■ for every pair of variables X and Y in a template τ_i with $X \equiv_{\tau_i} Y$, there is no constraint
 1392 $X \neq Y$ in τ_i ;

1393 *then we can extend $\varphi_{\mathcal{A}}$ to a total context assignment $\varphi'_{\mathcal{A}}$ for D respecting the constraints of*
 1394 *D .*

1395 **Proof.** By definition of equivalence, \equiv_D partitions all variables occurring in $\text{Trans}(D)$ in
 1396 *equivalence classes*. That is, two variables X and Y are in the same equivalence class iff $X \equiv_D Y$.
 1397 For a given variable X , we denote the equivalence class X belongs to by $[X]$. Note that for
 1398 any pair of variables X and Y occurring in $\text{Trans}(D)$, if $X \xrightarrow{F} Y$, then $X' \xrightarrow{F} Y'$ for any pair of
 1399 variables $X' \in [X]$ and $Y' \in [Y]$. By slight abuse of notation, we use $X \xrightarrow{F} [Y]$ and $[X] \xrightarrow{F} Y$ to
 1400 denote that $X \xrightarrow{F} Y'$ for every $Y' \in [Y]$ and $X' \xrightarrow{F} Y$ for every $X' \in [X]$, respectively.

1401 Let $(\text{Rels}, \text{Funcs})$ be the schema over which \mathcal{P} is defined. Since the schema graph
 1402 $SG(\text{Rels}, \text{Funcs})$ is acyclic, a total order $<_{SG}$ over Rels exists such that there is no path
 1403 from type R to type S in SG if $R <_{SG} S$. We now define $\varphi'_{\mathcal{A}}$ for variables in $\text{Trans}(D)$
 1404 according to the order implied by $<_{SG}$. If there are multiple variables of the same type, the
 1405 relative order in which we handle them is not important.

1406 The proof is as follows. Assume $\varphi_{\mathcal{A}}$ respects the constraints of D and is at least defined
 1407 for every variable X_i of o_i and Y_i of p_i in every τ_i . We extend $\varphi_{\mathcal{A}}$ towards $\varphi'_{\mathcal{A}}$ by defining
 1408 $\varphi'_{\mathcal{A}}$ for the whole equivalence class $[X]$ of the first (according to $<_{SG}$) variable X for which
 1409 $\varphi_{\mathcal{A}}$ is not defined. The precise construction is by case. In the first case, the tuple-context
 1410 that should be assigned to variables in $[X]$ is already implied, as it is the tuple-subcontext of
 1411 an existing tuple-context. In the second case, we construct a fresh tuple-context, including
 1412 existing tuple-contexts as tuple-subcontexts where we need to make sure that $\varphi'_{\mathcal{A}}$ respects
 1413 the constraints in D . In each case, we then argue that $\varphi'_{\mathcal{A}}$ still respects the constraints in

1414 *D*. By repeating this argument, we can extend the context assignment to a total context
 1415 assignment defined for all variables occurring in $\text{Trans}(D)$.

1416 **(Case 1)** If a variable Y exists with $\varphi_{\mathcal{A}}$ defined for Y and $Y \xrightarrow{f}_D [X]$, then $\varphi'_{\mathcal{A}}(X') = c_{X'}$
 1417 for every variable $X' \in [X]$, with $c_{X'}$ the tuple-subcontext of $c_Y = \varphi_{\mathcal{A}}(Y)$ witnessed by F .
 1418 Notice that this is well defined, even if there are multiple such Y , as they all agree on $c_{X'}$
 1419 by Definition 33 (2, 6). Also note that the special case where $\varphi_{\mathcal{A}}$ is already defined for at
 1420 least one variable $X' \in [X]$ is covered by this case as well, as $X' \xrightarrow{\varepsilon}_D [X]$ follows from $X' \in [X]$.
 1421 In this special case, the tuple-subcontext of $\varphi_{\mathcal{A}}(X')$ witnessed by ε (i.e., $\varphi_{\mathcal{A}}(X')$ itself) will be
 1422 assigned to each variable in $[X]$.

1423 We show that $\varphi'_{\mathcal{A}}$ indeed respects the constraints in D according to the properties stated in
 1424 Definition 33. To this end, let X' and Y' be two variables, with $c_{X'} = \varphi'_{\mathcal{A}}(X')$ and $c_{Y'} = \varphi'_{\mathcal{A}}(Y')$.
 1425 (1) By construction, $c_{X'}$ is a tuple-context for $\text{type}(X')$. (2-4) Note that if $X'' \xrightarrow{f'}_D Z''$ with
 1426 $X'' \in [X]$, then $Y \xrightarrow{f'}_D Z''$ and $c_{X''} = \varphi'_{\mathcal{A}}(X'')$ is the tuple-subcontext of $c_Y = \varphi_{\mathcal{A}}(Y)$ witnessed
 1427 by F' , implying that $c_{X''}(F') = c_Y(F \cdot F')$. If X' and/or Y' are in $[X]$, then we can apply
 1428 this substitution and use the fact that $\varphi_{\mathcal{A}}$ respects the constraints in τ to conclude that the
 1429 desired properties hold for $\varphi'_{\mathcal{A}}$. (5) Assume $X' \xrightarrow{f'}_D Y'$. If both $X' \in [X]$ and $Y' \in [X]$, then
 1430 $F' = \varepsilon$ as otherwise the schema graph is not acyclic. Since $c_{Y'} = c_{X'}$, it follows that $c_{Y'}$ is
 1431 a tuple-subcontext of $c_{X'}$ witnessed by ε . If $X' \in [X]$ and $Y' \notin [X]$, then $Y \xrightarrow{f'}_{\tau} Y'$ and $c_{Y'}$ is a
 1432 tuple-subcontext of c_Y witnessed by $F \cdot F'$ as $\varphi_{\mathcal{A}}$ respects the constraints of τ . Since $c_{X'}$ is
 1433 the tuple-subcontext of c_Y witnessed by F , it follows that $c_{Y'}$ is a tuple-subcontext of $c_{X'}$
 1434 witnessed by F' . If $X' \notin [X]$ and $Y' \in [X]$, then $Y \xrightarrow{f'}_{\tau} Y'$. Since $\varphi_{\mathcal{A}}$ respects the constraints in
 1435 D , we apply Definition 33 (2, 5, 6) to conclude that $c_{Y'}$ is a tuple-subcontext of $c_{X'}$ witnessed
 1436 by F' . (6) Assume $c_{X'}(F_{X'}) = c_{Y'}(F_{Y'})$, and let $c_{X''}$ and $c_{Y''}$ be the tuple-subcontexts of
 1437 respectively $c_{X'}$ witnessed by $F_{X'}$ and $c_{Y'}$ witnessed by $F_{Y'}$. We argue that $c_{X''} = c_{Y''}$. Note
 1438 that, if $X' \in [X]$, then $c_{X''}$ is the tuple-subcontext of c_Y witnessed by $F \cdot F_{X'}$. The reasoning
 1439 for $Y' \in [X]$ is analogous. Since $\varphi_{\mathcal{A}}$ respects the constraints in D , it follows that $c_{X''} = c_{Y''}$.

1440 **(Case 2)** Otherwise, we construct a fresh tuple-context c_X and define $\varphi'_{\mathcal{A}}(X') = c_X$ for every
 1441 variable $X' \in [X]$. This tuple-context c_X is constructed as follows: $c_X(\varepsilon) = \mathbf{t}_X$, with \mathbf{t}_X a fresh
 1442 tuple of the appropriate type. For every path $F = f \cdot F'$ in SG starting in $\text{type}(X)$, if there
 1443 is a variable Y with $[X] \xrightarrow{f}_D Y$, then $c_X(F) = c_Y(F')$, with $c_Y = \varphi_{\mathcal{A}}(Y)$. In other words, c_Y is
 1444 the tuple-subcontext of c_X witnessed by f . Note that due to the order \prec_{SG} , $\varphi_{\mathcal{A}}(Y)$ has to be
 1445 defined already. Also note that this is well defined, even if multiple such Y exist. In that
 1446 case, all these Y are equivalent to each other by definition of \equiv_D , and by construction of
 1447 $\varphi_{\mathcal{A}}$ they are assigned the same tuple-context. If instead no such variable Y exists, we define
 1448 $c_X(F) = \mathbf{t}_F$, with \mathbf{t}_F a fresh tuple of the appropriate type.

1449 We show that $\varphi'_{\mathcal{A}}$ indeed respects the constraints in D according to the properties stated
 1450 in Definition 33. To this end, let X' and Y' be two variables occurring in $\text{Trans}(D)$, with
 1451 $c_{X'} = \varphi'_{\mathcal{A}}(X')$ and $c_{Y'} = \varphi'_{\mathcal{A}}(Y')$. (1) By construction, $c_{X'}$ is a tuple-context for $\text{type}(X')$.
 1452 (2) Assume $X' \xrightarrow{f}_D Z$ and $Y' \xrightarrow{f}_D Z$ for some variable Z . We argue that there exists a
 1453 pair of variables X'' and Y'' and two sequences of function names F'_1 and F'_2 such that
 1454 $c_{X'}(F_1) = c_{X''}(F'_1) = c_{Y''}(F'_2) = c_{Y'}(F_2)$, where $c_{X''} = \varphi'_{\mathcal{A}}(X'')$ and $c_{Y''} = \varphi'_{\mathcal{A}}(Y'')$. If $X' \in [X]$,
 1455 then either $F_1 = f \cdot F'_1$ or $F_1 = \varepsilon$. In the former case there is a variable X'' with $X'' \xrightarrow{f}_D Z$ such
 1456 that $c_{X'}(F_1) = c_{X''}(F'_1)$, where $c_{X''} = \varphi_{\mathcal{A}}(X'')$. In the later case, $Z \in [X]$, and we simply take
 1457 $X'' = X'$ and $F'_1 = F_1$. If $X' \notin [X]$, we take $X'' = X'$ and $F'_1 = F_1$. For $Y' \in [X]$ and $Y' \notin [X]$, the
 1458 reasoning is analogous. It remains to argue that $c_{X''}(F'_1) = c_{Y''}(F'_2)$. By choice of X'' and Y'' ,
 1459 either $X'' \notin [X]$ and $Y'' \notin [X]$; or $Z \in [X]$. In the former case, $c_{X''}(F'_1) = c_{Y''}(F'_2)$ follows by the
 1460 fact that $\varphi_{\mathcal{A}}$ respects the constraints of D . In the latter case, both $X'' \in [X]$ and $Y'' \in [X]$, as
 1461 otherwise (Case 1) would apply to $[X]$ instead. Then, $c_{X''}(F'_1) = c_{Y''}(F'_2) = c_X(\varepsilon) = \mathbf{t}_X$. (3, 4)

1462 The reasoning is analogous to the previous property. Note in particular that by construction
 1463 of the new c_X we have $c_{X'}(F_1) = c_{Y'}(F_2)$ if $W \equiv_D Z$. Since $W \equiv_D Z$ implies that there is
 1464 no constraint $W \neq Z$ by the assumptions on $\varphi_{\mathcal{A}}$ and on the disequality constraints in each
 1465 template $\tau \in \text{Trans}(D)$, this does not lead to contradictions. (5) If $X' \xrightarrow{f} Y'$, then $Y' \in [X]$ only
 1466 if $X' \in [X]$, as otherwise (Case 1) would apply to $[X]$ instead. We argue by case that $c_{Y'}$ is a
 1467 tuple-subcontext of $c_{X'}$ witnessed by F' . If $X' \notin [X]$ and $Y' \notin [X]$, the result is immediate by
 1468 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
 1469 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
 1470 tuple-subcontext of $\varphi_{\mathcal{A}}(Z)$ witnessed by F'' . It now follows that $c_{Y'}$ is a tuple-subcontext
 1471 of c_X witnessed by F' . Lastly, If both $X' \in [X]$ and $Y' \in [X]$, then $F' = \varepsilon$, as otherwise the
 1472 schema graph is not acyclic. The result is immediate, as $c_{Y'} = c_{X'} = c_X$ is by definition a
 1473 tuple-subcontext of itself witnessed by ε . (6) Assume $c_{X''}(F_1) = c_{Y''}(F_1)$ for some pair of
 1474 tuple-contexts $c_{X''}$ and $c_{Y''}$ that are tuple-subcontexts of respectively $c_{X'}$ witnessed by F_1 and
 1475 $c_{Y'}$ witnessed by F_2 . We argue that $c_{X''} = c_{Y''}$. If both $c_{X'}$ and $c_{Y'}$ are different from c_X , the
 1476 result is immediate as $\varphi_{\mathcal{A}}$ respects the constraints of D . Otherwise, since the construction of
 1477 c_X , either copies existing tuple-contexts as tuple-subcontexts, or introduces fresh variables.
 1478 the result holds if $c_{X'}$ and/or $c_{Y'}$ are equal to c_X . \blacktriangleleft

1479 **► Lemma 36.** *Let $D = (\tau_1, o_1, p_2, \tau_2), \dots, (\tau_m, o_m, p_1, \tau_1)$ be a sequence of potentially con-*
 1480 *flicting quadruples over a set of transaction templates \mathcal{P} and $\varphi_{\mathcal{A}}$ a total context assignment for*
 1481 *D respecting the constraints of D . The variable mapping $\bar{\mu}$ obtained by defining $\bar{\mu}(X) = c_X(\varepsilon)$*
 1482 *for every variable X in $\text{Trans}(D)$ with $c_X = \varphi_{\mathcal{A}}(X)$ then is a valid variable mapping admissible*
 1483 *for some database D .*

1484 **Proof.** We first argue that $\bar{\mu}$ is valid by showing for each conflicting quadruple $(\tau_i, o_i, p_j, \tau_j)$
 1485 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
 1486 $X \xrightarrow{\tau_i} Y$. 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
 1487 tuple-subcontext of c_X witnessed by ε . By definition of tuple-subcontexts, $c_X(\varepsilon) = c_Y(\varepsilon)$, and,
 1488 as a result, $\bar{\mu}(X) = c_X(\varepsilon) = c_Y(\varepsilon) = \bar{\mu}(Y)$.

1489 Next, we construct a database \mathbf{D} and show that $\bar{\mu}$ is admissible for \mathbf{D} . To this end, we
 1490 add the tuple $\bar{\mu}(X)$ to \mathbf{D} for each variable X occurring in $\text{Trans}(D)$. For each functional
 1491 constraint $Y = f(X)$ in a transaction template in $\text{Trans}(D)$, we define $\bar{\mu}(Y) = f^{\mathbf{D}}(\bar{\mu}(X))$ for
 1492 the corresponding function $f^{\mathbf{D}}$ in \mathbf{D} . Note that this is well defined. Towards a contradiction,
 1493 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
 1494 $\bar{\mu}(Y) \neq \bar{\mu}(W)$. 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 $\bar{\mu}$,
 1495 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 33 (6),
 1496 it now follows that $c_X = c_Z$, since c_X (respectively c_Z) is a tuple-subcontext of itself witnessed
 1497 by ε 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
 1498 some template in $\text{Trans}(D)$, and hence $X \xrightarrow{f} Y$. Analogously, $Z \xrightarrow{f} W$. By Definition 33 (5),
 1499 c_Y and c_W are tuple-subcontexts of respectively c_X and c_Z witnessed by f . As $c_X = c_Z$, it
 1500 immediately follows that $c_Y = c_W$, and in particular $c_Y(\varepsilon) = c_W(\varepsilon)$, leading to the desired
 1501 contradiction.

1502 To conclude the proof, we show that $\bar{\mu}$ is indeed admissible for \mathbf{D} . By construction of
 1503 \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
 1504 $\tau \in \text{Trans}(D)$. We still need to argue that $\bar{\mu}(X) \neq \bar{\mu}(Y)$ for each constraint $X \neq Y$ in a
 1505 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
 1506 $\bar{\mu}(X) = c_X(\varepsilon)$ and $\bar{\mu}(Y) = c_Y(\varepsilon)$. Note that $X \xrightarrow{\tau} X$ and $Y \xrightarrow{\tau} Y$. Therefore, we can apply
 1507 Definition 33 (3) to conclude that $c_X(\varepsilon) \neq c_Y(\varepsilon)$, and hence $\bar{\mu}(X) \neq \bar{\mu}(Y)$. \blacktriangleleft

1508 In order to decide robustness against RC, one can now construct a sequence of quintuples

1509 E and derive the sequence of potentially conflicting quadruples D and partial context
 1510 assignment $\varphi_{\mathcal{A}}$ from it. If $\varphi_{\mathcal{A}}$ respects the constraints in D , then it follows from Lemma 35
 1511 and Lemma 36 that we can construct a variable assignment $\bar{\mu}$ such that $C = \bar{\mu}(D)$ is a
 1512 valid sequence of conflicting quadruples. However, in this construction of E , care should
 1513 be taken to guarantee that $\varphi_{\mathcal{A}}$ indeed respects the constraints in D , and that the resulting
 1514 multiversion split schedule based on C indeed satisfies all properties in Definition 6.

1515 In the algorithm that we are about to propose, we search for such a sequence of quintuples
 1516 E , but without fixating all the tuples in each context. For this, we generalize our definition
 1517 of tuple-contexts to allow variables: A *context for a type R* is a function from paths with
 1518 source R in $SG(\text{Rels}, \text{Funcs})$ to variables in **Var** and tuples in **Tuples** of the appropriate
 1519 type. The purpose of variables is to encode equalities and disequalities within each context,
 1520 without being explicit about the precise tuples. That is, if two paths ending in the same
 1521 node in SG are mapped on the same variable, then they will represent the same tuple; if they
 1522 are mapped on different variables, then they represent a different tuple. We remark that a
 1523 same variable occurring in different contexts can still represent different tuples. Analogous
 1524 to tuple-subcontexts, for two types R and S with $R \xrightarrow{\sim}_{SG} S$, we say that a context c_S for
 1525 type S is a *subcontext* of a context c_R for type R witnessed by F if:

- 1526 ■ for every path $S \xrightarrow{\sim}_{SG} S'$ in SG , if $c_R(F \cdot F')$ is a tuple, then $c_S(F') = c_R(F \cdot F')$; otherwise,
 1527 $c_S(F')$ is a variable; and
- 1528 ■ for every pair of paths $S \xrightarrow{\sim}_{SG} S_1$ and $S \xrightarrow{\sim}_{SG} S_2$ in SG with $c_R(F \cdot F_1)$ and $c_R(F \cdot F_2)$
 1529 variables, $c_S(F_1) = c_S(F_2)$ iff $c_R(F \cdot F_1) = c_R(F \cdot F_2)$.

1530 We call a context a *variable-context* if all paths are mapped on variables.

1531 For a transaction template τ , tuple-context c_p for p and c_o for o in τ , we consider
 1532 the set $\text{Contexts}(SG, \tau, p, c_p, o, c_o)$ of all different (not-necessarily tuple-) contexts c (up to
 1533 isomorphisms over the variables in c) that can be obtained, starting from a variable-context
 1534 c' , by performing substitutions of subcontexts of c' with subcontexts of c_p and/or c_o . More
 1535 formally, these substitutions are of the form: For a path $R_c \xrightarrow{\sim}_{SG} S$ (here R_c is the type that
 1536 c is for) and $R_p \xrightarrow{\sim}_{SG} S$ (with R_p the type that c_p is for) then $c(F \cdot F'') = c_p(F' \cdot F'')$ for
 1537 every path $S \xrightarrow{\sim}_{SG} S'$ in SG and with $c(F \cdot F'') = c'(F \cdot F'')$ otherwise. (The substitution
 1538 rule can be applied for c_p as well as for c_o .)

1539 ► **Lemma 37.** *Let \mathcal{P} be a set of transaction templates over an acyclic schema. Then, \mathcal{P} is
 1540 not robust against RC if, and only if, there is a sequence of quintuples $E = (\tau_1, o_1, c_{o_1}, p_1, c_{p_1}),$
 1541 $\dots, (\tau_m, o_m, c_{o_m}, p_m, c_{p_m})$ with $m \geq 2$ such that for each quintuple $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$ in E ,
 1542 with q_i and r_i two (not necessarily different) operations in $\{o_i, p_i\}$,*

- 1543 1. *if $i = 1$, then c_{o_1} and c_{p_1} are tuple-contexts for $\text{type}(\text{var}(o_1))$ and $\text{type}(\text{var}(p_1))$. Fur-*
 1544 *thermore, for every pair of tuple-subcontexts c'_{o_1} and c'_{p_1} of c_{o_1} and c_{p_1} witnessed by*
 1545 *respectively F and F' , if $c_{o_1}(F) = c_{p_1}(F')$, then $c'_{o_1} = c'_{p_1}$;*
- 1546 2. *if $i \neq 1$, then $c_{o_i}, c_{p_i} \in \text{Contexts}(SG, \tau_i, p_i, c_{p_i}, o_i, c_{o_i})$ are contexts for $\text{type}(\text{var}(o_i))$ and*
 1547 *$\text{type}(\text{var}(p_i))$;*
- 1548 3. *for every pair of variables W_i and Z_i in τ_i with $W_i \equiv_{\tau_i} Z_i$, there is no constraint $W_i \neq Z_i$ in*
 1549 *τ_i ;*
- 1550 4. *for every $\text{var}(q_i) \xrightarrow{\sim}_{\tau_i} Z_i$ and $\text{var}(r_i) \xrightarrow{\sim}_{\tau_i} Z_i$, the subcontext of c_{q_i} witnessed by F_1 is equal*
 1551 *(up to isomorphisms over variables) to the subcontext of c_{r_i} witnessed by F_2 .*
- 1552 5. *for every $\text{var}(q_i) \xrightarrow{\sim}_{\tau_i} W_i$ and $\text{var}(r_i) \xrightarrow{\sim}_{\tau_i} Z_i$ with $W_i \neq Z_i$ a constraint in τ_i , $c_{q_i}(F_1) \neq$
 1553 $c_{r_i}(F_2)$ or $c_{q_i}(F_1)$ and $c_{r_i}(F_2)$ are both variables;*
- 1554 6. *for every $\text{var}(q_i) \xrightarrow{\sim}_{\tau_i} W_i$ and $\text{var}(r_i) \xrightarrow{\sim}_{\tau_i} Z_i$, if $c_{q_i}(F_1)$ and $c_{r_i}(F_2)$ are the same tuple,*
 1555 *then there is no constraint $W_i \neq Z_i$ in τ_i ;*
- 1556 7. *if $\text{var}(q_i) \xrightarrow{\sim}_{\tau_i} \text{var}(r_i)$, then c_{r_i} is a subcontext of c_{q_i} witnessed by F ; and*

- 1557 **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
 1558 function names F and F' , then there is no operation $o'_i \in \tau_i$ potentially ww-conflicting
 1559 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 $\text{var}(q_1) \xrightarrow{F'} \tau_1 \text{var}(o'_1)$.
 1560 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
 1561 E with $j = i + 1$, or $i = m$ and $j = 1$:
 1562 **9.** o_i is potentially conflicting with p_j and $c_{o_i} = c_{p_j}$;
 1563 **10.** if $i = 1$ and $j = 2$, then o_1 is potentially rw-conflicting with p_2 ; and
 1564 **11.** if $i = m$ and $j = 1$, then $o_1 <_{\tau_1} p_1$ or o_m is potentially rw-conflicting with p_1 .

1565 **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 conflict-
 1566 ing quadruples derived from E . Note that each $(\tau_i, o_i, p_j, \tau_j) \in D$ is indeed a valid sequence
 1567 of potentially conflicting quadruples, as o_i is potentially conflicting with p_j by (9). We show
 1568 in Lemma 38 that a partial context assignment $\varphi_{\mathcal{A}}$ over a set of tuple-contexts \mathcal{A} exists such
 1569 that

- 1570 ■ for every pair of operations o_i and p_i occurring in D , $\varphi_{\mathcal{A}}$ is defined for $\text{var}(o_i)$ and
 1571 $\text{var}(p_i)$;
- 1572 ■ $\varphi_{\mathcal{A}}$ respects the constraints in D ; and
- 1573 ■ for every template τ_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
 1574 $c_{q_i} = \varphi_{\mathcal{A}}(\text{var}(q_i))$ and $c_{q_1} = \varphi_{\mathcal{A}}(\text{var}(q_1))$. If $c_{q_i}(F) = c_{q_1}(F')$ for some sequence of
 1575 function names F and F' , then there is no operation $o'_i \in \tau_i$ potentially ww-conflicting
 1576 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 $\text{var}(q_1) \xrightarrow{F'} \tau_1 \text{var}(o'_1)$.

1577 Because of (3), we can now apply Lemma 35 extending $\varphi_{\mathcal{A}}$ to a total context assignment
 1578 defined for all variables occurring in $\text{Trans}(D)$, without losing the property that $\varphi_{\mathcal{A}}$ respects
 1579 all constraints in D . Let $\bar{\mu}$ be the variable mapping obtained by defining $\bar{\mu}(X) = c_X(\varepsilon)$ for
 1580 every variable X in $\text{Trans}(D)$ with $c_X = \varphi_{\mathcal{A}}(X)$. By Lemma 36, $\bar{\mu}$ is a valid variable mapping
 1581 and a database \mathbf{D} exists such that $\bar{\mu}$ is admissible for \mathbf{D} .

1582 We now prove that the sequence of conflicting quadruples $C = \bar{\mu}(D)$ satisfies the conditions
 1583 stated in Definition 6 to show that \mathcal{P} is indeed not robust against RC. Condition (2) and
 1584 Condition (3) are immediate by respectively (11) and (10). Towards a contradiction, assume
 1585 Condition (1) does not hold. Then, there is an operation o'_i in a template τ_i potentially
 1586 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
 1587 $c_{o'_i} = \varphi_{\mathcal{A}}(\text{var}(o'_i))$ 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
 1588 construction of the total context assignment in Lemma 35, this is only the case if for some
 1589 $q_i \in \{o_i, p_i\}$ and $q_1 \in \{o_1, p_1\}$ it holds that $\text{var}(q_i) \xrightarrow{F} \tau_i \text{var}(o'_i)$ with $c_{q_i}(F) = c_{o'_i}(\varepsilon)$ and
 1590 $\text{var}(q_1) \xrightarrow{F'} \tau_1 \text{var}(o'_1)$ with $c_{q_1}(F') = c_{o'_1}(\varepsilon)$. Consequently, $c_{q_i}(F) = c_{q_1}(F')$, contradicting
 1591 Lemma 38.

1592 (only if) Since \mathcal{P} is not robust against RC, a multiversion split schedule s exists based on
 1593 a sequence of conflicting quadruples C over a set of transactions \mathcal{T} consistent with a set of
 1594 transaction templates $\mathcal{P} \in \mathbf{AcycTemp}$ and a database \mathbf{D} . Let $\bar{\mu}$ be the variable mapping
 1595 for a sequence of potentially conflicting quadruples $D = (\tau_1, o_1, p_2, \tau_2), \dots, (\tau_m, o_m, p_1, \tau_1)$
 1596 with $\bar{\mu}(D) = C$. By Lemma 34 a set \mathcal{A} of tuple-contexts and a total context assignment $\varphi_{\mathcal{A}}$
 1597 for D over \mathcal{A} exist with:

- 1598 ■ $\varphi_{\mathcal{A}}$ respects the constraints of D ; and
- 1599 ■ $\bar{\mu}(X) = c_X(\varepsilon)$ for every variable X , with $c_X = \varphi_{\mathcal{A}}(X)$.

1600 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 $\varphi_{\mathcal{A}}(\text{var}(p_1))$;
 1601 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
 1602 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 $c_{o_i} = \lambda \circ c'_{o_i}$
 1603 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 $c'_{p_i} = \varphi_{\mathcal{A}}(\text{var}(p_i))$.
 1604 Intuitively, we modify the tuple-contexts for each $\text{var}(o_i)$ and $\text{var}(p_i)$ as defined by $\varphi_{\mathcal{A}}$ into

1605 contexts over tuples and variables by replacing all tuples that do not occur in $\varphi_{\mathcal{A}}(\text{var}(o_1))$
 1606 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
 1607 $c_{p_1} = \varphi_{\mathcal{A}}(\text{var}(p_1))$.

1608 Next, we show that E satisfies all properties. In the argumentation below, we denote
 1609 $\varphi_{\mathcal{A}}(\text{var}(o_i))$ 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, $c_{o_1} =$
 1610 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 $\text{type}(\text{var}(p_1))$.
 1611 By Condition (6) of Definition 33, we have $c'_{o_1} = c'_{p_1}$ if $c_{o_1}(F) = c_{p_1}(F')$ for every pair of
 1612 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
 1613 property follows by construction of c_{o_i} and c_{p_i} based on λ and $\varphi_{\mathcal{A}}$. For completeness
 1614 sake, one should note that for each group of contexts up to isomorphisms over variables,
 1615 $\text{Contexts}(SG, \tau_1, p_1, c_{p_1}, o_1, c_{o_1})$ contains only one context. W.l.o.g. we can implicitly assume
 1616 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
 1617 to isomorphisms, as we will never directly test for equality or disequality between two variables
 1618 of different contexts. (3) Assume $\mathbb{W}_i \equiv_{\tau_i} \mathbb{Z}_i$, then $\mathbb{W}_i \xrightarrow{\varepsilon} \mathbb{Z}_i$. Since $\varphi_{\mathcal{A}}$ respects the constraints
 1619 in D , $c_{\mathbb{Z}_i}$ is a tuple-subcontext of $c_{\mathbb{W}_i}$ witnessed by ε , with $c_{\mathbb{Z}_i} = \varphi_{\mathcal{A}}(\mathbb{Z}_i)$ and $c_{\mathbb{W}_i} = \varphi_{\mathcal{A}}(\mathbb{W}_i)$.
 1620 Then, $c_{\mathbb{Z}_i} = c_{\mathbb{W}_i}$, and in particular $c_{\mathbb{Z}_i}(\varepsilon) = c_{\mathbb{W}_i}(\varepsilon)$. By Lemma 34, $\bar{\mu}(\mathbb{Z}_i) = c_{\mathbb{Z}_i}(\varepsilon) = c_{\mathbb{W}_i}(\varepsilon) =$
 1621 $\bar{\mu}(\mathbb{W}_i)$. Since $\bar{\mu}$ is admissible for \mathbf{D} , the constraint $\mathbb{W}_i \neq \mathbb{Z}_i$ cannot exist. (4) Since $\varphi_{\mathcal{A}}$
 1622 respects the constraints in D , it follows from Conditions (2) and (6) in Definition 33 that the
 1623 tuple-subcontext of c'_{q_i} witnessed by F_1 is equal to the tuple-subcontext of c'_{r_i} witnessed by
 1624 F_2 . As a result, the subcontext of $c_{q_i} = \lambda \circ c'_{r_i}$ witnessed by F_1 is equal (up to isomorphisms
 1625 over variables) to the subcontext of $c_{r_i} = \lambda \circ c'_{r_i}$ witnessed by F_2 . (5) Analogous to the
 1626 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
 1627 two different tuples, or both a variable. (6) If $c_{q_i}(F_1) = \lambda \circ c'_{q_i}(F_1) = \lambda \circ c'_{r_i}(F_2) = c_{r_i}(F_2)$
 1628 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
 1629 there is no constraint $\mathbb{W}_i \neq v_{z_i}$. (7) If $\text{var}(q_i) \xrightarrow{\varepsilon} \tau_i \text{var}(r_i)$, then c'_{r_i} is a tuple-subcontext of
 1630 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
 1631 c'_{q_i} and c'_{r_i} , it immediately follows that c_{r_i} is a subcontext of c_{q_i} . The case for $\mathbb{Y}_i \xrightarrow{\varepsilon} \tau_i \mathbb{X}_i$ is
 1632 analogous. (8) Assume towards a contradiction that $c_{q_i}(F) = c_{q_1}(F')$ is a tuple for some
 1633 $q_1 \in \{o_1, p_1\}$ and some sequence of function names F and F' , and there is an operation $o'_i \in \tau_i$
 1634 potentially ww-conflicting with an operation $o'_1 \in \text{prefix}_{o_1}(\tau_1)$ with $\text{var}(q_i) \xrightarrow{\varepsilon} \tau_i \text{var}(o'_i)$ and
 1635 $\text{var}(q_1) \xrightarrow{\varepsilon} \tau_1 \text{var}(o'_1)$. Since $c_{q_i}(F) = c_{q_1}(F')$ is a tuple, $c'_{q_i}(F) = c_{q_i}(F) = c_{q_1}(F') = c'_{q_1}(F')$.
 1636 By definition of $\bar{\mu}$ and since $\varphi_{\mathcal{A}}$ respects the constraints in D , we conclude that $\bar{\mu}(o'_i)$ in
 1637 $\bar{\mu}(\tau_i)$ is ww-conflicting with $\bar{\mu}(o'_1)$ in $\text{prefix}_{\bar{\mu}(o_1)}(\bar{\mu}(\tau_1))$, thereby contradicting Condition (1)
 1638 of Definition 6. (9) Since E is based on C , the operation o_i is potentially conflicting with p_j .
 1639 Furthermore, since $\text{var}(o_i) \equiv_D \text{var}(p_j)$ and since $\varphi_{\mathcal{A}}$ respects the constraints of D , $c'_{o_i} = c'_{p_j}$,
 1640 and hence $c_{o_i} = c_{p_j}$. (10) Immediate by Condition (3) of Definition 6. (11) Immediate by
 1641 Condition (2) of Definition 6. \blacktriangleleft

1642 Central to the correctness of the if-direction of Lemma 37 is the observation that for any
 1643 sequence E satisfying the stated conditions, we can assign tuples to variables in each context,
 1644 thereby replacing contexts with tuple-contexts, in such a way that the resulting context
 1645 assignment over tuple-contexts respects the constraints of the corresponding sequence of
 1646 potentially conflicting quadruples D . This observation is formalized in the following Lemma:

1647 **► Lemma 38.** *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 quin-*
 1648 *tuples satisfying the conditions stated in Lemma 37, and let $D = (\tau_1, o_1, p_2, \tau_2), \dots, (\tau_m, o_m, p_1, \tau_1)$*
 1649 *be the sequence of potentially conflicting quadruples derived from E . Then a partial context*
 1650 *assignment $\varphi_{\mathcal{A}}$ over a set of tuple-contexts \mathcal{A} exists such that*

1651 **■** *for every pair of operations o_i and p_i occurring in D , $\varphi_{\mathcal{A}}$ is defined for $\text{var}(o_i)$ and*

1652 $var(p_i)$;
 1653 ■ $\varphi_{\mathcal{A}}$ respects the constraints in D ; and
 1654 ■ for every template τ_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
 1655 $c_{q_i} = \varphi_{\mathcal{A}}(var(q_i))$ and $c_{q_1} = \varphi_{\mathcal{A}}(var(q_1))$. If $c_{q_i}(F) = c_{q_1}(F')$ for some sequence of
 1656 function names F and F' , then there is no operation $o'_i \in \tau_i$ potentially ww-conflicting
 1657 with an operation $o'_1 \in \text{prefix}_{o_1}(\tau_1)$ with $var(q_i) \xrightarrow{F} \tau_i var(o'_i)$ and $var(q_1) \xrightarrow{F'} \tau_1 var(o'_1)$.

1658 **Proof.** The general proof idea is as follows. We iteratively extend $\varphi_{\mathcal{A}}$ by deriving $\varphi_{\mathcal{A}}(var(o_i))$
 1659 and $\varphi_{\mathcal{A}}(var(p_i))$ 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
 1660 that they appear in E . Afterwards, we argue that $\varphi_{\mathcal{A}}$ respects the constraints in D and for
 1661 every $q_i \in \{o_i, p_i\}$ with $i \neq 1$ and $q_1 \in \{o_1, p_1\}$ with $c_{q_i} = \varphi_{\mathcal{A}}(var(q_i))$ and $c_{q_1} = \varphi_{\mathcal{A}}(var(q_1))$,
 1662 and for every pair of sequences of function names F and F' with $c_{q_i}(F) = c_{q_1}(F')$, there
 1663 is no operation $o'_i \in \tau_i$ potentially ww-conflicting with an operation $o'_1 \in \text{prefix}_{o_1}(\tau_1)$ with
 1664 $var(q_i) \xrightarrow{F} \tau_i var(o'_i)$ and $var(q_1) \xrightarrow{F'} \tau_1 var(o'_1)$.

1665 Let $(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$ be the first quintuple in E . We initiate $\varphi_{\mathcal{A}}$ by defining $\varphi_{\mathcal{A}}(var(o_1)) =$
 1666 c_{o_1} and $\varphi_{\mathcal{A}}(var(p_1)) = c_{p_1}$. Next, we iteratively extend $\varphi_{\mathcal{A}}$ by considering the remaining
 1667 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
 1668 quintuple and $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$ the next quintuple in E . We define $\varphi_{\mathcal{A}}(var(p_i)) = c'_{p_i} =$
 1669 $\varphi_{\mathcal{A}}(var(o_{i-1}))$ and $\varphi_{\mathcal{A}}(var(o_i)) = c'_{o_i} = \lambda_i \circ c_{o_i}$, where $\lambda_i : \mathbf{Tuples} \cup \mathbf{Var} \rightarrow \mathbf{Tuples}$ is a
 1670 function mapping tuples and variables occurring in c_{o_i} to tuples such that⁵

- 1671 ■ $\lambda_i(\mathbf{t}) = \mathbf{t}$ for each tuple \mathbf{t} occurring in c_{o_i} ;
- 1672 ■ $\lambda_i(\mathbf{Q}) = \mathbf{v}$ for each variable \mathbf{Q} occurring in c_{o_i} for which there is a variable \mathbf{Z} in τ_i
 1673 and sequences of function names F, F' and F'' with $var(p_i) \xrightarrow{F} \tau_i \mathbf{Z}$, $var(o_i) \xrightarrow{F'} \tau_i \mathbf{Z}$,
 1674 $c'_{p_i}(F \cdot F'') = \mathbf{v}$ and $c_{o_i}(F' \cdot F'') = \mathbf{Q}$; and
- 1675 ■ $\lambda_i(\mathbf{Q}) = \mathbf{t}_{i,\mathbf{Q}}$, for the remaining variables \mathbf{Q} in c_{o_i} where $\mathbf{t}_{i,\mathbf{Q}}$ is a fresh tuple.

1676 Note that this λ_i is well defined. In particular, the second rule intuitively states that the
 1677 tuple-subcontext of the resulting c'_{o_i} witnessed by F' is equal to the tuple-subcontext of c'_{p_i}
 1678 witnessed by F , given that there is a variable \mathbf{Z} with $var(o_i) \xrightarrow{F'} \tau_i \mathbf{Z}$ and $var(p_i) \xrightarrow{F} \tau_i \mathbf{Z}$. This
 1679 substitution is well defined since in this case the subcontext of c_{o_i} witnessed by F' is equal
 1680 (up to isomorphisms over variables) to the subcontext of c_{p_i} witnessed by F according to
 1681 Condition 4 of Lemma 37.

1682 It remains to show that $\varphi_{\mathcal{A}}$ indeed satisfies the conditions stated in Lemma 38. To this
 1683 end, note that by definition of variable determination, if $\mathbf{X} \xrightarrow{F} \tau_i \mathbf{Y}$ with \mathbf{X} in a template τ_i
 1684 and \mathbf{Y} in a template τ_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
 1685 such that (\dagger):

- 1686 ■ $\mathbf{X}_{k_1} = \mathbf{X}$ and $\mathbf{Y}_{k_m} = \mathbf{Y}$;
- 1687 ■ each pair of (not necessarily different) variables $\mathbf{X}_{k_i}, \mathbf{Y}_{k_i}$ occur in the same template τ_{k_i} in
 1688 $\text{Trans}(D)$ and $\mathbf{X}_{k_i} \xrightarrow{F} \tau_{k_i} \mathbf{Y}_{k_i}$ with $F = F_1 \cdot \dots \cdot F_m$;
- 1689 ■ 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
 1690 E (where we assume that τ_1 is neighbouring to τ_n in E); and
- 1691 ■ for each pair of variables $\mathbf{Y}_{k_i}, \mathbf{X}_{k_{i+1}}$, there is a sequence of function names F' such that
 1692 $var(o_{k_i}) \xrightarrow{F'} \tau_{k_i} \mathbf{Y}_{k_i}$ and $var(p_{k_{i+1}}) \xrightarrow{F'} \tau_{k_{i+1}} \mathbf{X}_{k_{i+1}}$ (i.e., equivalence of \mathbf{Y}_{k_i} and $\mathbf{X}_{k_{i+1}}$ in D is
 1693 implied by equivalence of $var(o_{k_i})$ and $var(p_{k_{i+1}})$).

1694 In other words, $\mathbf{X} \xrightarrow{F} \tau_i \mathbf{Y}$ can be broken down into a sequence of $\mathbf{X}_{k_i} \xrightarrow{F} \tau_{k_i} \mathbf{Y}_{k_i}$ through a
 1695 sequence of neighbouring templates, where equivalence between each \mathbf{Y}_{k_i} and $\mathbf{X}_{k_{i+1}}$ is implied

⁵ Note that λ_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.

1696 by the variables in the potentially conflicting operations o_{k_i} and $p_{k_{i+1}}$. For ease of exposition,
 1697 we implicitly assumed that $\tau_{k_1}, \dots, \tau_{k_m}$ agrees with the order in E . If the order is opposite
 1698 to the order in E instead, the above still holds, but the occurrences of o_{k_i} and $p_{k_{i+1}}$ should
 1699 be replaced with p_{k_i} and $o_{k_{i+1}}$.

1700 We argue by construction of $\varphi_{\mathcal{A}}$ that for every pair of variables \mathbf{X} and \mathbf{Y} for which $\varphi_{\mathcal{A}}$ is
 1701 defined with $c_{\mathbf{X}} = \varphi_{\mathcal{A}}(\mathbf{X})$ and $c_{\mathbf{Y}} = \varphi_{\mathcal{A}}(\mathbf{Y})$, if $c_{\mathbf{X}}(F) = c_{\mathbf{Y}}(F')$ for some sequence of function
 1702 names F and F' , then $c'_{\mathbf{X}} = c'_{\mathbf{Y}}$, where $c'_{\mathbf{X}}$ and $c'_{\mathbf{Y}}$ are the tuple-subcontexts of $c_{\mathbf{X}}$ witnessed
 1703 by F and $c_{\mathbf{Y}}$ witnessed by F' , respectively (\ddagger). If $\mathbf{X} = \text{var}(o_1)$ and $\mathbf{Y} = \text{var}(p_1)$ (or the other
 1704 way around), the result is immediate by Lemma 37 (1). Otherwise, let \mathbf{X} be the variable in a
 1705 template τ_i and \mathbf{Y} the variable in a template τ_j such that $j \leq i$ (i.e., τ_i does not occur before
 1706 τ_j in E). W.l.o.g., we assume that $\mathbf{X} = \text{var}(o_i)$ with $i \neq 1$ (the case where $\mathbf{X} = \text{var}(p_i)$ is
 1707 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}
 1708 based on c'_{p_i} and λ_i , if $c'_{o_i}(F_i) = c'_{p_i}(F'_i)$, then the whole tuple-subcontext of c'_{o_i} witnessed
 1709 by F_i is copied over from the tuple-subcontext of c'_{p_i} witnessed by F'_i . Indeed λ_i introduces
 1710 fresh tuples whenever the tuple for $c'_{o_i}(F_i)$ is not implied by c'_{p_i} .

1711 The desired properties now follow from (\dagger) and (\ddagger) as well as the conditions in Lemma 37.
 1712 In particular, $\varphi_{\mathcal{A}}$ respecting the constraints of D can now be derived from Conditions (1, 2, 4-7)
 1713 in Lemma 37, and the last condition of Lemma 38 follows from Condition (8) in Lemma 37. \blacktriangleleft

1714 A NEXPSPACE algorithm proving the correctness of Theorem 16 is now immediate by
 1715 Lemma 37, as we can iteratively guess and verify quintuples in E while only keeping track of
 1716 the very first quintuple and the previous quintuple. Since in an acyclic schema graph the
 1717 number of paths starting in a given type is at most exponential in the total number of types,
 1718 each context is defined over at most an exponential number of paths. However, to formally
 1719 argue that these contexts can be encoded in exponential space, we still need to show that
 1720 each tuple or variable used in a context can be encoded in at most exponential space. Since
 1721 the only tuples used are those mentioned in c_{o_1} and c_{p_1} , and since we can reuse the same
 1722 variables over all contexts, both the maximal number of tuples and the maximal number of
 1723 variables needed are exponential in the total number of types.

1724 C.3 Proof for Theorem 18

1725 The PSPACE result for workloads over a schema (**Rels**, **Funcs**) where the number of paths
 1726 between any two nodes in the schema graph is bounded by a constant k is immediate by the
 1727 nondeterministic algorithm based on Lemma 37 presented for Theorem 16. Indeed, in this
 1728 case, the total number of paths starting in a given type is at most $k \cdot |\text{Rels}|$ and therefore each
 1729 context is defined over at most a polynomial number of paths, instead of an exponential
 1730 number of paths for the general case in Theorem 16.

1731 The EXPTIME result for workloads in **AcycResTemp** follows from a deterministic al-
 1732 gorithm based on Lemma 37. In the remainder of this section, we first present the algorithm,
 1733 and then discuss its complexity.

1734 C.3.1 A deterministic algorithm

1735 Towards a deterministic algorithm, assume the first quintuple $(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$ of E is fixed.
 1736 We now translate the problem of deciding whether we can extend E such that it satisfies all
 1737 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
 1738 as follows:

- 1739 ■ each quintuple $(\tau_i, o_i, c_{o_i}, p_i, c_{p_i})$ satisfying Conditions (2-8) of Lemma 37 is added as a
- 1740 node to $G(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$; and

1741 ■ 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
 1742 potentially conflicting with p_j and $c_{o_i} = c_{p_j}$ (c.f. Condition (9) of Lemma 37).

1743 By construction, it is now easy to see that there is a sequence E satisfying Lemma 37 if
 1744 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
 1745 $G(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$ (where we allow a zero-length path with $2 = m$), such that (\dagger)

- 1746 ■ $c_{o_1} = c_{p_2}$ and $c_{o_m} = c_{p_1}$ (c.f. Condition (9) of Lemma 37);
- 1747 ■ o_1 is potentially rw-conflicting with p_2 (c.f. Condition (10) of Lemma 37); and
- 1748 ■ $o_1 <_{\tau_1} p_1$ or o_m is potentially rw-conflicting with p_1 (c.f. Condition (11) of Lemma 37).

1749 **(Algorithm)** Given a set of transaction templates \mathcal{P} over a schema $(\text{Rels}, \text{Funcs})$, the
 1750 algorithm iterates over all possible quintuples $(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$ satisfying Condition (1, 3-7)
 1751 of Lemma 37, where we consider all possible tuple-contexts c_{o_1} and c_{p_1} up to isomorphisms.
 1752 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
 1753 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
 1754 $(\tau_m, o_m, c_{o_m}, p_m, c_{p_m})$ in TC satisfying (\dagger) , the algorithm emits a reject, indicating that \mathcal{P} is
 1755 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,
 1756 the algorithm didn't reject after considering all such quintuples, it accepts, indicating that \mathcal{P}
 1757 is indeed robust against RC . The correctness of this algorithm is immediate by Lemma 37.

1758 C.3.2 Complexity analysis

1759 We show the complexity of the presented algorithm. For this, first, notice that we have
 1760 defined contexts c based on a type S (with S not necessarily a root of SG). For encoding
 1761 purposes it makes sense to encode these as contexts for a root type R in combination with
 1762 the intended type S . The context as defined in the previous section can then be derived by
 1763 taking the left-most subtree with root S . Notice that this is purely an encoding choice that
 1764 will simplify the analysis.

1765 For a schema graph $SG(\text{Rels}, \text{Funcs})$ the total number of non-isomorphic tuple-contexts
 1766 can be expressed using Bell's number $B(n)$, denoting the number of partitions for a set of
 1767 size n , and the set $\text{Paths}_{SG}(R, S) = \{(R, F, S) \mid R \xrightarrow{F}_{SG} S\}$ expressing the different paths
 1768 from one node R to another node S in SG . Concretely,

$$1769 \quad |\text{TupleContexts}(SG)| \leq \sum_{R \in \text{roots}(SG)} \prod_{S \in \text{Rels}} B(|\text{Paths}_{SG}(R, S)|) = B^*$$

1770 where $\text{TupleContexts}(SG)$ denotes the set containing all different tuple-contexts (up to
 1771 isomorphisms).

1772 Now let c_1 and c_2 be two fixed contexts for types that are descendants of roots R_1 and
 1773 root R_2 , respectively, in SG , and let c be a context for a type descending from root R . To
 1774 express a bound on the number of substitutions in c from (parts of) c_1 and c_2 , we need some
 1775 additional terminology: Let $\text{Paths}_{SG}(R, *) = \bigcup_{S \in \text{Rels}} \text{Paths}_{SG}(R, S)$. We say that a path
 1776 $R \xrightarrow{F_1}_{SG} S$ is a *prefix* of a path $R \xrightarrow{F_2}_{SG} S'$ in SG if there is a (possibly empty) sequence of
 1777 function names F_2 with $F = F_1 \cdot F_2$. The number of substitutions in c from (parts of) c_1
 1778 and c_2 is now bounded by

$$1779 \quad \sum_{\text{Part} \subseteq \text{Paths}_{SG}(R, *)} \mathbf{1}_{PFP} \prod_{R \xrightarrow{F}_{SG} S \in \text{Part}} (|\text{Paths}_{SG}(R_1, S)| + |\text{Paths}_{SG}(R_2, S)|) \leq T^\ell \cdot (2P)^\ell \leq (2TP)^\ell$$

1780 In the above expression, $\mathbf{1}_{PFP}$ is an indicator variable that equals 1 if no path in Part
 1781 is a prefix of another path in Part and that equals 0 otherwise. Further: P denotes the
 1782 maximum number of different paths between a particular root and a particular node in SG ,

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

■ **Figure 8** Transaction templates for SmallBank.

Account(Name, CustomerID, IsPremium)
 Savings(CustomerID, Balance, InterestRate)
 Checking(CustomerID, Balance)

■ **Figure 9** Tables of the SmallBank benchmark. Underlined attributes are primary keys.

1783 T denotes the maximum number of different paths from a particular root to nodes in SG,
 1784 and ℓ denotes the maximal size of a set in which no path is a prefix of another path in the
 1785 set. The latter is trivially bounded by T .

1786 A special cases exists if all templates τ in \mathcal{P} are restricted. In that case, the size of sets
 1787 Part is bounded by 2, hence $\ell \leq 2$.

1788 With the above bounds, the complexity of the presented algorithm is rather straight-
 1789 forward. The iteration over all possible quintuples $(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$ requires at most
 1790 $|\mathcal{P}| \cdot t^2 \cdot (B^*)^2$ iterations, with t denoting the maximal number of operations in a transaction
 1791 template of \mathcal{P} . The remainder of the computation is dominated by the transitive-closure com-
 1792 putation. Since the constructed graph $G(\tau_1, o_1, c_{o_1}, p_1, c_{p_1})$ has at most $|\mathcal{P}| \cdot t^2 \cdot (B^* \cdot (2TP)^\ell)^2$
 1793 nodes, the transitive closure computation requires $(|\mathcal{P}| \cdot t^2 \cdot (B^* \cdot (2TP)^\ell)^2)^3$ steps. Putting
 1794 these numbers together, we obtain:

$$1795 \quad \mathcal{O}(|\mathcal{P}|^4 \cdot t^8 \cdot (B^*)^8 \cdot (2TP)^{6\ell}).$$

1796 Since ℓ is bounded by a constant if all template are restricted, and since B^* , T and P
 1797 can be exponential in the size of the input, the presented algorithm indeed decides T-
 1798 ROBUSTNESS(**AcycResTemp**,RC) in EXPTIME.

1799 **D** SmallBank and TPC-C benchmarks

1800 **D.1** SmallBank Benchmark

1801 The SmallBank schema consists of three tables as given in Figure 9. The Account table
 1802 associates customer names with IDs and keeps track of the premium status (Boolean);

1803 CustomerID is a UNIQUE attribute. The other tables contain the balance (numeric value) of the
 1804 savings and checking accounts of customers identified by their ID. Account (CustomerID) is a
 1805 foreign key referencing both the columns Savings (CustomerID) and Checking (CustomerID).
 1806 The interest rate on a savings account is based on a number of parameters, including the
 1807 account status (premium or not). The application code can interact with the database only
 1808 through the following transaction programs:

- 1809 ■ Balance(N): returns the total balance (savings & checking) for a customer with name N .
- 1810 ■ DepositChecking(N, V): makes a deposit of amount V on the checking account of the
 1811 customer with name N .
- 1812 ■ TransactSavings(N, V): makes a deposit or withdrawal V on the savings account of the
 1813 customer with name N .
- 1814 ■ Amalgamate(N_1, N_2): transfers all the funds from N_1 to N_2 .
- 1815 ■ WriteCheck(N, V): writes a check V against the account of the customer with name N ,
 1816 penalizing if overdrawing.
- 1817 ■ GoPremium(N): converts the account of the customer with name N to a premium account
 1818 and updates the interest rate of the corresponding savings account. This transaction
 1819 program is an extension w.r.t. [2].

1820 Figure 10 contains the SQL code for the SmallBank transaction templates presented in
 1821 Figure 8.

1822 D.2 TPC-C Benchmark

1823 This benchmark is based on the TPC-C benchmark [17]. We modified the schema and
 1824 templates to turn all predicate reads into key-based accesses. The schema consists of six
 1825 relations:

- 1826 ■ Warehouse(WarehouseID, Info, YTD),
- 1827 ■ District(WarehouseID, DistrictID, Info, YTD, NextOrderID),
- 1828 ■ Customer(WarehouseID, DistrictID, CustID, Info, Balance),
- 1829 ■ Order(WarehouseID, DistrictID, OrderID, CustID, Status),
- 1830 ■ OrderLine(WarehouseID, DistrictID, OrderID, OrderLineID, ItemID, DeliveryInfo, Quant-
 1831 ity), and
- 1832 ■ Stock(WarehouseID, ItemID, Quantity).

1833 The function names belonging to this schema are given in Table 1.

1834 We focus on five different transaction templates:

- 1835 1. NewOrder($W, D, C, I_1, Q_1, I_2, Q_2, \dots$): creates a new order for the customer identified
 1836 by (W, D, C) . The id for this order is obtained by increasing the NextOrderID attribute
 1837 of the District tuple identified by (W, D) by one. Each order consists of a number of
 1838 items I_1, I_2, \dots with respectively quantities Q_1, Q_2, \dots . For each of these items, a new
 1839 OrderLine tuple is created and the related stock quantity is decreased.
- 1840 2. Payment(W, D, C, A): represents a customer identified by (W, D, C) paying an amount
 1841 A . This payment is reflected in the database by increasing the balance of this customer
 1842 by A . This amount is furthermore added to the YearToDate (YTD) income of both the
 1843 related warehouse and district.
- 1844 3. OrderStatus(W, D, C, O): requests information about the current status of the order
 1845 identified by (W, D, O) . This transaction template collects information of the customer
 1846 identified by (W, D, C) who created the order, the order itself, and the different OrderLine
 1847 tuples related to this order.
- 1848 4. Delivery(W, D, C, O): delivers the order represented by (W, D, O) . The status of the
 1849 order is updated, as well as the DeliveryInfo attribute of each OrderLine tuple related to

23:46 Robustness against Read Committed for Transaction Templates with Functional Constraints

```
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;

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;

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;

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

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

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;
```

■ **Figure 10** SmallBank SQL Transaction Templates.

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.

NewOrder:

$R[X : \text{Warehouse}\{W, \text{Inf}\}]$
 $U[Y : \text{District}\{W, D, \text{Inf}, N\}\{N\}]$
 $R[Z : \text{Customer}\{W, D, C, \text{Inf}\}]$
 $W[S : \text{Order}\{W, D, O, C, \text{Sta}\}]$
 $U[T_1 : \text{Stock}\{W, I, \text{Qua}\}\{\text{Qua}\}]$
 $W[V_1 : \text{OrderLine}\{W, D, O, \text{OL}, I, \text{Del}, \text{Qua}\}]$
 $U[T_2 : \text{Stock}\{W, I, \text{Qua}\}\{\text{Qua}\}]$
 $W[V_2 : \text{OrderLine}\{W, D, O, \text{OL}, I, \text{Del}, \text{Qua}\}]$
 $X = f_{D \rightarrow W}(Y), Y = f_{C \rightarrow D}(Z), Z = f_{O \rightarrow C}(S)$
 $S = f_{L \rightarrow O}(V_1), S = f_{L \rightarrow O}(V_2)$
 $T_1 = f_{L \rightarrow S}(V_1), T_2 = f_{L \rightarrow S}(V_2)$
 $X = f_{S \rightarrow W}(T_1), X = f_{S \rightarrow W}(T_2)$

Delivery:

$U[S : \text{Order}\{W, D, O\}\{\text{Sta}\}]$
 $U[V_1 : \text{OrderLine}\{W, D, O, \text{OL}, \text{Del}\}\{\text{Del}\}]$
 $U[V_2 : \text{OrderLine}\{W, D, O, \text{OL}, \text{Del}\}\{\text{Del}\}]$
 $U[Z : \text{Customer}\{W, D, C, \text{Bal}\}\{\text{Bal}\}]$
 $Z = f_{O \rightarrow C}(S), S = f_{L \rightarrow O}(V_1), S = f_{L \rightarrow O}(V_2)$

Payment:

$U[X : \text{Warehouse}\{W, \text{YTD}\}\{\text{YTD}\}]$
 $U[Y : \text{District}\{W, D, \text{YTD}\}\{\text{YTD}\}]$
 $U[Z : \text{Customer}\{W, D, C, \text{Bal}\}\{\text{Bal}\}]$
 $X = f_{D \rightarrow W}(Y), Y = f_{C \rightarrow D}(Z)$

OrderStatus:

$R[Z : \text{Customer}\{W, D, C, \text{Inf}, \text{Bal}\}]$
 $R[S : \text{Order}\{W, D, O, C, \text{Sta}\}]$
 $R[V_1 : \text{OrderLine}\{W, D, O, \text{OL}, I, \text{Del}, \text{Qua}\}]$
 $R[V_2 : \text{OrderLine}\{W, D, O, \text{OL}, I, \text{Del}, \text{Qua}\}]$
 $Z = f_{O \rightarrow C}(S), S = f_{L \rightarrow O}(V_1), S = f_{L \rightarrow O}(V_2)$

StockLevel:

$R[T : \text{Stock}\{W, I, \text{Qua}\}]$

■ **Figure 11** Abstraction for the TPC-C transaction templates. Attribute names are abbreviated.

1850 this order. The total price of the order is deduced from the balance of the customer who
 1851 made this order, identified by (W, D, C) .

1852 5. StockLevel(W, I): returns the current stock level of item I in warehouse W .

1853 A detailed abstraction of each transaction template is given in Figure 11. To shorten the
 1854 presentation, we only show two orderlines per order.