

Language-Integrated Query for Temporal Data

Simon Fowler

University of Glasgow
UK

simon.fowler@glasgow.ac.uk

Vashti Galpin

University of Edinburgh
UK

vashti.galpin@ed.ac.uk

James Cheney

University of Edinburgh
UK

jcheney@inf.ed.ac.uk

Abstract

Modern applications often manage time-varying data. Despite decades of research on temporal databases, which culminated in the addition of temporal data operations into the SQL:2011 standard, temporal data query and manipulation operations are unavailable in most mainstream database management systems, leaving developers with the unenviable task of implementing such functionality from scratch.

In this paper, we extend *language-integrated query* to support writing temporal queries and updates in a uniform host language, with the language performing the required underlying rewriting to emulate temporal capabilities automatically on any standard relational database. We introduce two core languages, λ_{TLINQ} and λ_{VLINQ} , for manipulating transaction time and valid time data respectively, and formalise existing implementation strategies by giving provably correct semantics-preserving translations into a non-temporal core language, λ_{LINQ} . We show how existing work on query normalisation supports a surprisingly simple implementation strategy for *sequenced joins*. We implement our approach in the Links programming language, and describe a non-trivial case study based on curating COVID-19 statistics.

CCS Concepts: • Software and its engineering → Domain specific languages; • Information systems → Structured Query Language; Temporal data.

Keywords: language-integrated query, temporal databases, domain-specific languages, multi-tier programming

ACM Reference Format:

Simon Fowler, Vashti Galpin, and James Cheney. 2022. Language-Integrated Query for Temporal Data. In *Proceedings of the 21st ACM SIGPLAN International Conference on Generative Programming: Concepts and Experiences (GPCE '22), December 06–07, 2022, Auckland, New Zealand*. ACM, New York, NY, USA, 31 pages. <https://doi.org/10.1145/3564719.3568690>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. GPCE '22, December 06–07, 2022, Auckland, New Zealand

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9920-3/22/12...\$15.00

<https://doi.org/10.1145/3564719.3568690>

1 Introduction

Most interesting programs access or query data stored persistently, often in a database. Relational database management systems (RDBMSs) are the most popular option and provide a standard domain-specific language, SQL, for querying and modifying the data. Ideally, programmers can express the desired queries or updates declaratively in SQL and leave the database to decide how to answer queries or perform updates efficiently and safely (e.g. in the presence of concurrent accesses), but there are many pitfalls arising from interfacing with SQL from a general-purpose language, leading to the well-known *impedance mismatch* problem [10]. These difficulties range from run-time failures due to the generation of queries as unchecked SQL strings at runtime, to serious security vulnerabilities like SQL injection attacks [30].

Among the most successful approaches to overcome the above challenges, and the approach we build upon in this paper, is *language-integrated query*, exemplified by Microsoft's popular LINQ for .NET [23, 34] and in a number of other language designs such as Links [9, 21] and libraries such as Quill [27]. Within this design space we focus on a family of techniques derived from foundational work by Buneman et al. [3] on the *nested relational calculus* (NRC), a core query language with monadic collection types; work by Wong [35] on rewriting NRC expressions to normal forms that can be translated to SQL, which forms the basis of the approach taken in Links [8, 21] and has been adapted to F# [5].

Many interesting database applications involve data that changes over time. Perhaps inevitably, *temporal data management* [17] has a long history. Temporal databases provide powerful features for querying and modifying data that changes over time, and are particularly suitable for modeling time-varying phenomena, such as enterprise data or evolving scientific knowledge, and supporting auditing and transparency about how the current state of the data was achieved, for example in financial or scientific settings.

To illustrate how temporal databases work and why they are useful, consider the following toy example: a temporal to-do list. A temporal database can be conceptualised abstractly as a *function* mapping each possible time instant (e.g. times of day) to a database state [18]. For efficiency in the common case where most of the data is unchanged most of the time, temporal databases are often represented by augmenting each row with an *interval timestamp* indicating the time

period when the row is present. For technical reasons, *closed-open* intervals $[start, end)$ representing the times $start \leq t < end$ are typically used [32].

In our temporal to-do list, the table at each time instant has fields “task”, a string field, and “done”, a Boolean field. Additional fields “start” and “end” record the endpoints of the time interval during which each row is to be considered part of the table. An end time of ∞ (“forever”) reflects that there is no (currently known) end time and in the absence of other changes, the row is considered present from the start time onwards. For example:

task	done	start	end
Go shopping	true	11:00	∞
Cook dinner	false	11:00	17:30
Walk the dog	false	11:00	∞
Cook dinner	true	17:30	∞
Watch TV	false	11:00	19:00

represents a temporal table where all four tasks were added at 11:00, with “Go shopping” being complete and the others incomplete; at 17:30 “Cook dinner” was marked “done” from then onwards, and at 19:00 “Watch TV” was removed from the table without being completed. Technically, note that this example interprets the time annotations as *transaction time*, that is, the times indicate when certain data was in the database; there is another dimension, *valid time*, and we will discuss both dimensions in greater detail later on.

The problems of querying and updating temporal databases have been well-studied, leading to the landmark language design TSQL2 [31] extending SQL. However, despite decades of effort, only limited elements of TSQL2 were eventually incorporated into the SQL:2011 standard [20] and these features have not yet been widely adopted. Directly implementing temporal queries in SQL is possible, but painful: a TSQL2-style query or update operation may grow by a factor of 10 or more when translated to plain SQL, which leaves plenty of scope for error, and thus these powerful capabilities remain outside the grasp of non-experts. In this paper we take first steps towards reconciling temporal data management with language-integrated query based on query normalisation. We propose supporting temporal capabilities by translation to ordinary language-integrated query and hypothesise that this approach can make temporal data management safer, easier to use and more accessible to non-experts than the current state of affairs. As an initial test of this hypothesis we present a high-level design, a working implementation, and detail our experience with a nontrivial case study.

Although both language-integrated query and temporal databases are now well-studied topics, we believe that their combination has never been considered before. Doing so has a number of potential benefits, including making the power of well-studied language designs such as TSQL2 more accessible to non-expert programmers, and providing a high-level abstraction that can be implemented efficiently in different ways. Our interest is particularly motivated by the needs

Types	A, B	$::= C \mid A \rightarrow^E B \mid \text{Bag}(A) \mid (\ell : A) \mid \text{Table}(A)$
Base types	C	$::= \text{String} \mid \text{Int} \mid \text{Bool} \mid \text{Time}$
Effects	e	$::= \text{read} \mid \text{write}$
Effect sets	E	
Terms	L, M, N	$::= x \mid c \mid t$ $\mid \lambda x.M \mid M N \mid \odot\{\vec{M}\}$ $\mid \text{if } L \text{ then } M \text{ else } N$ $\mid \int \mid \int M \mid M \uplus N \mid \text{for } (x \leftarrow M) N$ $\mid (\ell = M) \mid M.\ell \mid \text{now}$ $\mid \text{query } M \mid \text{get } M \mid \text{insert } M \text{ values } N$ $\mid \text{update } (x \leftarrow L) \text{ where } M \text{ set } (\ell = N)$ $\mid \text{delete } (x \leftarrow M) \text{ where } N$

Figure 1. Syntax of λ_{LINQ}

of scientific database development, where data versioning for accountability and research integrity are very important needs that are not well-supported by conventional database systems [4]. Temporal data management has the potential to become a “killer app” for language-integrated query, and this paper takes a first but significant step towards this goal.

The overarching contribution of this paper is the first extension of language-integrated query to support transaction time and valid time temporal data.

Concretely, we make three main contributions:

1. Based on λ_{LINQ} (§2), a formalism based on the Nested Relational Calculus (NRC) [29], we introduce typed λ -calculi to model queries and modifications on transaction time (§3) and valid time (§4) databases. We give semantics-preserving translations to λ_{LINQ} for both.
2. We show how existing work on query normalisation allows a surprisingly straightforward implementation strategy for *sequenced joins* (§5).
3. We implement our constructs in the Links functional web programming language, and describe a case study based on curating COVID-19 data (§6).

Although the concepts behind translating temporal queries and updates into non-temporal core languages are well known [32], our core calculi λ_{TLINQ} and λ_{VLINQ} are novel and aid us in showing (to the best of our knowledge) the first correctness results for the translations.

We relegate several details and all proofs to the appendices.

2 Background: Language-Integrated Query

We begin by introducing a basic λ -calculus, called λ_{LINQ} , to model language-integrated query in non-temporal databases. Our calculus is based heavily on the Nested Relational Calculus [29], with support for database modifications heavily inspired by the calculus of Fehrenbach and Cheney [13]. The calculus uses a type-and-effect system to ensure database accesses can occur in ‘safe’ places, i.e., that we do not attempt to perform a modification operation in the middle of a query. Effects include **read** (denoting a read from a database) and **write** (denoting an update to the database).

Types A, B include base types C , effect-annotated function types $A \rightarrow^E B$, unordered collection types $\text{Bag}(A)$, record

types $(\ell_i : A_i)_i$ denoting a record with labels ℓ_i and types A_i , and handles $\text{Table}(A)$ for tables containing records of type A . A record is a *base record* if it contains only fields of base type. We assume that the base types includes at least `Bool` and the `Time` type, which denotes (abstract) timestamps.

Basic terms include variables x , constants c , table handles t , functions $\lambda x.M$, application $M N$, n -ary operations $\odot \{\vec{M}\}$, and conditionals **if** L **then** M **else** N . We assume that the set of operations contains the usual unary and binary relation operators, as well as the n -ary operations **greatest** and **least** on timestamps which return their largest and smallest arguments respectively. We assume that the set of constants contains timestamps ι of type `Time`, and two distinguished timestamps $-\infty$ and ∞ of type `Time`, which denote the minimum and maximum timestamps respectively. The calculus also includes the empty multiset constructor $\{\}$; the singleton multiset constructor $\{M\}$; multiset union $M \uplus N$; and comprehensions **for** $(x \leftarrow M) N$. We write $\{M_1, \dots, M_n\}$ as sugar for $\{M_1\} \uplus \dots \uplus \{M_n\}$. We also have records $(\ell_i = M_i)_i$ and projection $M.\ell$. Term **now** retrieves the current time.

We write **let** $x = M$ **in** N as the usual syntactic sugar for $(\lambda x.N) M$, and $M; N$ as sugar for $(\lambda x.N)M$ for some fresh x . We also define **where** MN as sugar for **if** M **then** N **else** $\{\}$. We denote unordered collections with a tilde (e.g., \vec{M}), and ordered sequences with an arrow (e.g., \vec{M}).

The **get** M term retrieves the contents of a table into a bag; **insert** M **values** N inserts values N into table M ; **update** $(x \leftarrow L)$ **where** M **set** $(\ell_i = N_i)_i$ updates table L , updating the fields ℓ_i to N_i of each record x satisfying predicate M . The **delete** $(x \leftarrow M)$ **where** N term removes those records x in table M satisfying predicate N .

2.1 Typing rules

Figure 2 shows the typing rules for λ_{LINQ} . The rules are implicitly parameterised by a database schema Σ mapping table names to types of the form $\text{Bag}((\ell_i : C_i)_i)$. Many rules are similar to those of the simply-typed λ -calculus extended with monadic collection operations [3] and a set-based effect system [22], and such standard rules are relegated to Appendix A. The typing judgement has the shape $\Gamma \vdash M : A! E$, which can be read, “Under type environment Γ , term M has type A and produces effects E ”.

Rule T-QUERY states that a term **query** M is well-typed if M is of a *query type*: either a base type, a record type whose fields are query types, or a bag whose elements are query types. The term M must also only have **read** effects.

Rule T-GET states that **get** M has type $\text{Bag}(A)$ if M has type $\text{Table}(A)$, and produces the **read** effect. Rule T-TABLE states that a table reference follows the type of the table in the schema. T-INSERT types a database insert **insert** M **values** N , requiring M to be a table reference of type $\text{Table}(A)$, and the inserted values N to be a bag of type $\text{Bag}(A)$. T-UPDATE ensures the predicate and update terms

are typable under an environment extended with the row type, and ensures that all updated values match the type expected by the row. Rule T-DELETE is similar. All subterms used as predicates or used to calculate updated terms must be pure (that is, side-effect free), and all modifications have the **write** effect.

2.2 Semantics

Figure 3 shows the syntax and typing rules of values, and the big-step semantics of λ_{LINQ} . Most rules are standard, and presented in Appendix A. Values V, W include functions, constants, tables, fully-evaluated records, and fully-evaluated bags $\{\vec{V}\}$. Unlike the unary bag constructor $\{M\}$, fully-evaluated bags contain an unordered *collection* of values. All values are pure. We write \oplus for record extension, e.g., $(\ell_1 = M) \oplus (\ell_2 = N) = (\ell_1 = M, \ell_2 = N)$.

Since evaluation is effectful (as database operations can update the database), the evaluation judgement has the shape $M \Downarrow_{\Delta, \iota} (V, \Delta')$; this can be read “term M with current database Δ at time ι evaluates to value V with updated database Δ' ”. A database is a mapping from table names to bags of base records. To avoid additional complexity, we assume evaluation is atomic and does not update the time; one could straightforwardly update the semantics with a **tick** operation without affecting any results.

We use two further evaluation relations for terms which do not write to the database: $M \Downarrow_{\iota}^* V$ states that a pure term M (i.e., a term typable under an empty effect set) evaluates to V . Similarly, $M \Downarrow_{\Delta, \iota}^* V$ states that a term M which only performs the **read** effect evaluates to V . We omit the definitions, which are similar to the evaluation relation but do not propagate database changes (since no changes can occur).

Rule E-NOW returns the current timestamp. Rule E-QUERY evaluates the body of a query using the read-only evaluation relation. Rule E-GET evaluates its subject to a table reference, and then returns the contents of a table. Rule E-INSERT does similar, evaluating the values to insert, and then appending them to the contents of the table. Rule E-UPDATE iterates over a table, updating a record if the predicate matches, and leaving it unmodified if not. Finally, E-DELETE deletes those rows satisfying the deletion predicate.

λ_{LINQ} enjoys a standard type soundness property.

Proposition 2.1 (Type soundness). *If $\cdot \vdash M : A! E$ then there exists some V and Δ' such that $M \Downarrow_{\Delta, \iota} (V, \Delta')$ and $\cdot \vdash V : A! \emptyset$.*

More importantly, the type-and-effect system ensures that query and update expressions in λ_{LINQ} can be translated to SQL equivalent, even in the presence of higher-order functions and nested query results [5, 6, 8, 21]. This alternative implementation is equivalent to the semantics given here but usually much more efficient since the database query optimiser can take advantage of any available integrity constraints or statistics about the data.

Term typing

$$\begin{array}{c}
\text{T-QUERY} \\
\frac{\Gamma \vdash M:\text{Bag}(A)!E \quad A :: \text{QType} \quad E \subseteq \{\text{read}\}}{\Gamma \vdash \text{query } M:\text{Bag}(A)!E}
\end{array}
\quad
\begin{array}{c}
\text{T-NOW} \\
\frac{}{\Gamma \vdash \text{now}:\text{Time}!\emptyset}
\end{array}
\quad
\begin{array}{c}
\text{T-GET} \\
\frac{\Gamma \vdash M:\text{Table}(A)!E}{\Gamma \vdash \text{get } M:\text{Bag}(A)! \{\text{read}\} \cup E}
\end{array}
\quad
\begin{array}{c}
\text{T-INSERT} \\
\frac{\Gamma \vdash M:\text{Table}(A)!E \quad \Gamma \vdash N:\text{Bag}(A)! \emptyset}{\Gamma \vdash \text{insert } M \text{ values } N:()! \{\text{write}\} \cup E}
\end{array}
\quad
\boxed{\Gamma \vdash M:A!E}$$

$$\begin{array}{c}
\text{T-UPDATE} \\
\frac{\Gamma \vdash L:\text{Table}(A)!E \quad A = (\ell_i : B_i)_{i \in I} \quad \Gamma, x : A \vdash M:\text{Bool}!\emptyset \quad (j \in I \wedge \Gamma, x : A \vdash N_j:B_j!\emptyset)_{j \in J}}{\Gamma \vdash \text{update } (x \leftarrow L) \text{ where } M \text{ set } (\ell_j = N_j)_{j \in J}:()! \{\text{write}\} \cup E}
\end{array}
\quad
\begin{array}{c}
\text{T-DELETE} \\
\frac{\Gamma \vdash M:\text{Table}(A)!E \quad \Gamma, x : A \vdash N:\text{Bool}!\emptyset}{\Gamma \vdash \text{delete } (x \leftarrow M) \text{ where } N:()! \{\text{write}\} \cup E}
\end{array}$$

Query typing

$$\frac{}{C :: \text{QType}}
\quad
\frac{(A_i :: \text{QType})_i}{(\ell_i : A_i)_i :: \text{QType}}
\quad
\frac{A :: \text{QType}}{\text{Bag}(A) :: \text{QType}}
\quad
\boxed{A :: \text{QType}}$$

Figure 2. Typing rules for λ_{LINQ} (selected)

Syntax of values, operations on values, and value typing

$$\begin{array}{l}
\text{Values } V, W ::= \lambda x.M \mid c \mid t \mid (\ell_i = V_i)_i \mid \ulcorner \tilde{V} \urcorner \\
\ulcorner \tilde{V} \urcorner \hat{\cup} \ulcorner \tilde{W} \urcorner \triangleq \ulcorner \tilde{V} \cdot \tilde{W} \urcorner \\
((\ell_i = V_i)_{i \in I} \text{ with } (\ell_j = W_j)) \triangleq (\ell_i = V_i)_{i \in I \cup J} \oplus (\ell_j = W_j)_{j \in J}
\end{array}$$

$$\begin{array}{c}
\text{T-BAGV} \\
\frac{(\Gamma \vdash V_i:A!\emptyset)_i}{\Gamma \vdash \ulcorner \tilde{V} \urcorner:\text{Bag}(A)! \emptyset}
\end{array}$$

Big-step reduction rules

$$\begin{array}{c}
\text{E-NOW} \\
\frac{}{\text{now} \Downarrow_{\Delta, t} (t, \Delta)}
\end{array}
\quad
\begin{array}{c}
\text{E-QUERY} \\
\frac{M \Downarrow_{\Delta, t}^* V}{\text{query } M \Downarrow_{\Delta, t} (V, \Delta)}
\end{array}
\quad
\begin{array}{c}
\text{E-GET} \\
\frac{M \Downarrow_{\Delta, t} (t, \Delta')}{\text{get } M \Downarrow_{\Delta, t} (\Delta'(t), \Delta')}
\end{array}
\quad
\begin{array}{c}
\text{E-INSERT} \\
\frac{M \Downarrow_{\Delta, t} (t, \Delta') \quad N \Downarrow_t^* V}{\text{insert } M \text{ values } N \Downarrow_{\Delta, t} ((), \Delta'[t \mapsto \Delta'(t) \hat{\cup} V])}
\end{array}
\quad
\boxed{M \Downarrow_{\Delta, t} (V, \Delta')}$$

$$\begin{array}{c}
\text{E-UPDATE} \\
\frac{L \Downarrow_{\Delta, t} (t, \Delta_1) \quad \Delta_2 = \Delta_1 [t \mapsto \ulcorner \text{upd}(v) \urcorner \mid v \in \Delta_1(t)]}{\text{update } (x \leftarrow L) \text{ where } M \text{ set } (\ell = N) \Downarrow_{\Delta, t} ((), \Delta_2)}
\end{array}
\quad
\begin{array}{c}
\text{E-DELETE} \\
\frac{M \Downarrow_{\Delta, t} (t, \Delta_1) \quad \Delta_2 = \Delta_1 [t \mapsto \ulcorner v \in \Delta(t) \mid N\{v/x\} \Downarrow_t^* \text{false} \urcorner]}{\text{delete } (x \leftarrow M) \text{ where } N \Downarrow_{\Delta, t} ((), \Delta_2)}
\end{array}$$

Figure 3. Semantics of λ_{LINQ} (selected)

3 Transaction Time

The first dimension of time we investigate is *transaction time* [33], which records how the *state of the database* changes over time. The key idea behind transaction time databases is that update operations are *non-destructive*, so we can always view a database as it stood at a particular point in time.

Let us illustrate with the to-do list example from the introduction. The original table is on the left. The table after making some changes is shown on the right.

task	done	task	done
Go shopping	true	Go shopping	true
Cook dinner	false	Cook dinner	true
Walk the dog	false	Walk the dog	false
Watch TV	false		

However, since updates and deletions in λ_{LINQ} are destructive, we have lost the original data. Instead, let us see how this could be handled by a transaction-time database:

task	done	start	end
Go shopping	true	11:00	∞
Cook dinner	false	11:00	∞
Walk the dog	false	11:00	∞
Watch TV	false	11:00	∞

There are several methods by which we can maintain the temporal information in the database: for example we could maintain a tracking log which records each entry, or we could use various *temporal partitioning* strategies [32]. In this paper, we use a *period-stamped* representation, where each record in the database is augmented with fields delimiting the interval when the record was present in the database.

The time period is a closed-open representation, meaning that each row is in the database from (and including) the time stated in the *start* column, and is in the database up to (but not including) the time stated in the *end* column. We also assume that $start < end$ always holds.

Here, our database states that all four tasks were entered into the database at 11:00. However, if we then decide to

check off “Cook dinner” at 17:30 and delete “Watch TV” at 19:00, we obtain the table shown in the introduction.

Since timestamps are either ∞ or only refer to the past; users do not modify period stamps directly; and the information in the database grows monotonically, we can reconstruct the state of the database at any given time.

3.1 Calculus

λ_{TLINQ} extends λ_{LINQ} with native support for transaction time operations; instead of performing destructive updates, we adjust the end timestamp of affected rows and, if necessary, insert updated rows. λ_{TLINQ} database entries are therefore of the form $V_1^{[V_2, V_3]}$, where V_1 is the record data and V_2 and V_3 are the start and end timestamps.

Figure 4 shows the syntax, typing rules, and semantics of λ_{TLINQ} ; for brevity, we show the main differences to λ_{LINQ} . Period-stamped transaction-time database rows are represented as triples $data^{[start, end]}$ with type $\text{TransactionTime}(A)$, where $data$ has type A (the type of each record), and both $start$ and end have type Time . A row is currently present in the database if its end value is ∞ . We introduce three accessors: **data** V extracts the data record from a transaction-time row; **start** V extracts the start time; and **end** V extracts the end time. The **get** construct has an updated type to show that it returns a bag of $\text{TransactionTime}(A)$ values, rather than the records directly. The typing rules for the other constructs remain the same as in λ_{LINQ} .

The accessor rules ET-DATA, ET-START, and ET-END project the expected component of the transaction-time row. Rule ET-INSERT period-stamps each record to begin at the current time, and sets the end time to be ∞ . Rule ET-DELETE records deletions for current rows satisfying the deletion predicate. Instead of being removed from the database, the end times of the affected rows are set to the current timestamp. Finally, rule ET-UPDATE performs updates for current rows satisfying the update predicate. Instead of changing a record directly, the upd definition generates *two* records: the previous record, closed off at the current timestamp, and the new record with updated values, starting from the current timestamp and with end field ∞ . Returning to our running example, define $at(tbl, time)$ to return all records in tbl starting before $time$ and ending after $time$. We can then query the database as it stood at 18:00:

```

at( $t, time$ )  $\triangleq$ 
  for ( $x \leftarrow \text{get } t$ )
    where ( $\text{start } x \leq time \wedge time < \text{end } x$ )
       $\lfloor \text{data } x \rfloor$ 
at( $tbl, 18:00$ )

```

task	done
Go shopping	true
Cook dinner	true
Walk the dog	false
Watch TV	false

Let $\text{current}(t) = at(t, \infty)$ return the current snapshot of t . We can then query the current snapshot of the database:

```

current( $tbl$ ) =


| task         | done         |
|--------------|--------------|
| Go shopping  | <b>true</b>  |
| Cook dinner  | <b>true</b>  |
| Walk the dog | <b>false</b> |


```

3.2 Translation

We can implement the native transaction-time semantics for λ_{TLINQ} database operations by translation to λ_{LINQ} . Our translation adapts the SQL implementations of temporal operations by Snodgrass [32] to a language-integrated query setting. We prove correctness relative to the semantics.

λ_{TLINQ} has a native representation of period-stamped data, whereas λ_{LINQ} requires table types to be flat. Consequently, the translations require knowledge of the types of each record. We therefore annotate each λ_{TLINQ} database term with the type of table on which it operates (this can be achieved through a standard type-directed translation pass).

The (omitted) translation of λ_{TLINQ} types into λ_{LINQ} types is straightforward, save for $\text{TransactionTime}(A)$ which is translated into a record ($data: \llbracket A \rrbracket, start: \text{Time}, end: \text{Time}$). The same is true for the basic λ -calculus terms. Timestamped rows $V_{data}^{[V_{start}, V_{end}]}$ are translated to fit the above record type; specifically, ($data = \llbracket V_{data} \rrbracket, start = \llbracket V_{start} \rrbracket, end = \llbracket V_{end} \rrbracket$).

We define the *flattening* of a λ_{TLINQ} row and database as:

$$\begin{aligned} \downarrow((\ell_i = V_i)_{i=1}^{[W_1, W_2]}) &\triangleq (\ell_i = V_i)_i \oplus (start = W_1, end = W_2) \\ \downarrow \Delta &\triangleq [t \mapsto \lfloor \downarrow D \rfloor \mid t \mapsto \lfloor \widetilde{D} \rfloor \in \Delta] \end{aligned}$$

Remark 3.1. We have chosen to represent a λ_{TLINQ} period-stamped record as a nested record in λ_{LINQ} , but we could equally adopt a flat representation. Since we build on the Nested Relational Calculus, we take advantage of the ability to return nested results; previous work on query shredding [6] allows us to flatten nested results in a later translation pass. A nested representation is more convenient for our implementation and makes the translation simpler and more compositional, so we mirror this choice in the formalism.

Figure 5 shows the translation of λ_{TLINQ} terms into λ_{LINQ} . The translation makes use of three auxiliary definitions. Eta-expansion $\eta(x, \vec{\ell})$ eta-expands a variable of record type with respect to a sequence of labels, and $\text{restrict}(x, \vec{\ell}, M)$ applies an eta-expanded record to a M under a λ -binder for x (required since the λ_{TLINQ} predicates expect just the data from the record, and not the period-stamping fields). Finally, $\text{isCurrent}(M)$ tests whether the *end* field of M is ∞ .

Since timestamped rows are translated to records, the temporal accessor functions are translated to record projection. The **get** function is translated to retrieve the flattened λ_{LINQ} representation of the table and pack it via η -expansion into a bag of nested records. The translation of **insert** extends the provided values with a *start* field referring to the current timestamp and an *end* field set to ∞ , before inserting them into the database.

A **delete** is translated as a λ_{LINQ} **update** operation, which sets the *end* record of each affected row to the current timestamp. An **update** is translated in three steps: querying the database to obtain the affected current records, with updated values and timestamps; updating the database to close off the existing affected rows; and materialising the insertion.

Additional Syntax for λ_{TLINQ}

Types $A, B ::= \dots \mid \text{TransactionTime}(A)$ Timestamped rows $D ::= V_1^{[V_2, V_3]}$
 Terms $L, M, N ::= \dots \mid \mathbf{data} M \mid \mathbf{start} M \mid \mathbf{end} M$ Values $V, W ::= \dots \mid D$

Modified Typing Rules for λ_{TLINQ}

T-ROW

$$\frac{\Gamma \vdash V_1:A! \emptyset \quad \Gamma \vdash V_2:\text{Time}! \emptyset \quad \Gamma \vdash V_3:\text{Time}! \emptyset}{\Gamma \vdash V_1^{[V_2, V_3]}:\text{TransactionTime}(A)! \emptyset}$$

$$\frac{\text{T-DATA} \quad \Gamma \vdash M:\text{TransactionTime}(A)! E}{\Gamma \vdash \mathbf{data} M:A! E}$$

$$\frac{\text{T-START} \quad \Gamma \vdash M:\text{TransactionTime}(A)! E}{\Gamma \vdash \mathbf{start} M:\text{Time}! E}$$

$$\frac{\text{T-END} \quad \Gamma \vdash M:\text{TransactionTime}(A)! E}{\Gamma \vdash \mathbf{end} M:\text{Time}! E}$$

T-GET

$$\frac{\Gamma \vdash M:\text{Table}(A)! E}{\Gamma \vdash \mathbf{get} M:\text{Bag}(\text{TransactionTime}(A))! \{\mathbf{read}\} \cup E}$$

Semantics for λ_{TLINQ} database operations

$$\frac{\text{ET-DATA} \quad M \Downarrow_{\Delta, t}^T (V_1^{[V_2, V_3]}, \Delta')}{\mathbf{data} M \Downarrow_{\Delta, t}^T (V_1, \Delta')}$$

$$\frac{\text{ET-START} \quad M \Downarrow_{\Delta, t}^T (V_1^{[V_2, V_3]}, \Delta')}{\mathbf{start} M \Downarrow_{\Delta, t}^T (V_2, \Delta')}$$

$$\frac{\text{ET-END} \quad M \Downarrow_{\Delta, t}^T (V_1^{[V_2, V_3]}, \Delta')}{\mathbf{end} M \Downarrow_{\Delta, t}^T (V_3, \Delta')}$$

ET-INSERT

$$\frac{M \Downarrow_{\Delta, t}^T (t, \Delta_1) \quad N \Downarrow_t^* \{\tilde{V}\} \quad vs = \{\nu^{[t, \infty)} \mid \nu \in \tilde{V}\} \quad \Delta_2 = \Delta' [t \mapsto \Delta_1(t) \hat{\cup} vs]}{\mathbf{insert} M \mathbf{values} N \Downarrow_{\Delta, t}^T ((), \Delta_2)}$$

ET-UPDATE

$$\frac{\text{ET-DELETE} \quad M \Downarrow_{\Delta, t}^T (t, \Delta_1) \quad \Delta_2 = \Delta_1 [t \mapsto \{\mathbf{del}(v) \mid v \in \Delta_1(t)\}]}{\mathbf{delete} (x \leftarrow M) \mathbf{where} N \Downarrow_{\Delta, t}^T ((), \Delta_2)}$$

$$\text{del}(data^{[start, end)}) = \begin{cases} data^{[start, t)} & \text{if } end = \infty \text{ and } N\{data/x\} \Downarrow_t^* \mathbf{true} \\ data^{[start, end)} & \text{otherwise} \end{cases}$$

$$\frac{L \Downarrow_{\Delta, t}^T (t, \Delta_1) \quad \Delta_2 = \Delta_1 [t \mapsto \{\hat{\oplus}\} \{\text{upd}(v) \mid v \in \Delta_1(t)\}]}{\mathbf{update} (x \leftarrow L) \mathbf{where} M \mathbf{set} (\ell = N) \Downarrow_{\Delta, t}^T ((), \Delta_2)}$$

$$\text{upd}(data^{[start, end)}) = \begin{cases} \{\text{data}^{[start, t)}, (data \mathbf{with} \ell = V)^{[t, \infty)}\} & \\ \{\text{if } M\{data/x\} \Downarrow_t^* \mathbf{true}, (N_i\{data/x\} \Downarrow_t^* V_i)_i, \text{ and } end = \infty \\ \text{data}^{[start, end)}\} & \text{otherwise} \end{cases}$$

Figure 4. Syntax, typing rules, and semantics of λ_{TLINQ} **3.3 Metatheory**

We restrict our attention to *well formed* rows and databases, where the *start* timestamp is less than the *end* timestamp.

Definition 3.1 (Well formed rows and databases). A database Δ is well formed, written $\text{wf}(\Delta)$, if for each timestamped row $V_{data}^{[V_{start}, V_{end})}$ in Δ we have that $V_{start} < V_{end}$.

Definition 3.2 (Maximum timestamp). The *maximum timestamp* of a collection of records $\tilde{D} = \{V_{data_i}^{[V_{start_i}, V_{end_i})}\}_i$, written $\max(\tilde{D})$, is either $-\infty$ if $\tilde{D} = \{\}$, or the greatest V_{end_i} such that $V_{end_i} \neq \infty$ otherwise. The maximum timestamp of a database Δ , written $\max(\Delta)$, is the maximum timestamp of all its constituent tables.

Again, λ_{TLINQ} enjoys type soundness.

Proposition 3.2 (Type soundness (λ_{TLINQ})). *If $\vdash M:A!E$, then given a $\text{wf}(\Delta)$ and t such that $\max(\Delta) \leq t$, then there exists some V and well formed Δ' such that $M \Downarrow_{\Delta, t}^T (V, \Delta')$.*

We can now show that the translation is correct:

Theorem 3.1. *If $\vdash M:A!E$ and $M \Downarrow_{\Delta, t}^T (V, \Delta')$ where $\text{wf}(\Delta)$ and $\max(\Delta) \leq t$, then $\llbracket M \rrbracket \Downarrow_{\Delta, t} \llbracket V \rrbracket, \llbracket \Delta' \rrbracket$*

4 Valid Time

The other dimension of time we will look at is *valid time*, which tracks when something is true in the *domain being modelled*. Each timestamp therefore defines the *period of validity* (PV) of each record.

Unlike in a transaction time database, the database does not necessarily grow monotonically since we can apply destructive updates and deletions. Furthermore, whereas in a transaction time database timestamps can only refer to the past (or ∞), in a valid time database we may state that a row is valid until some specific point in the future (for example, the end of a fixed-term employment contract). A further difference from transaction time databases is that users *can* modify timestamps directly, and can also apply updates and deletions over a time period. Let us illustrate with the ‘employees’ table of an HR database:

name	position	salary	start	end
Alice	Lecturer	40000	2010	2018
Alice	Senior Lecturer	50000	2018	∞
Bob	PhD Student	15000	2019	2023
Charles	PhD Student	15000	2018	2022

The first modification is to hire Dolores as a professor, on an open-ended contract. As this is an insertion operation on the database at the current moment in time, it is known as a *current insertion*. We can write the following query:

Auxiliary Definitions

$$\begin{aligned} \eta(x, \tilde{\ell}) &\triangleq (\ell_i = x.\ell_i)_{i \in I} \\ \text{restrict}(x, \tilde{\ell}, M) &\triangleq (\lambda x.M) \eta(\tilde{\ell}, x) \\ \text{isCurrent}(M) &\triangleq M.\text{end} = \infty \end{aligned}$$

Translation on database terms

```

[[data M]] = [[M]].data
[[start M]] = [[M]].start
[[end M]] = [[M]].end
[[get(ℓi:Ai)i M]] =
  query
  for (x ← get [[M]])
    ⌊(data = η(x, ℓ̃), start = x.start, end = x.end)⌋
[[insert(ℓi:Ai)i M values N]] =
  let rows =
    for (x ← [[N]])
      ⌊η(x, ℓ̃) ⊕ (start = now, end = ∞)⌋
  in
  insert [[M]] values rows
[[delete(ℓi:Ai)i (x ← M) where N]] =
  update (x ← [[M]])
  where (restrict(x, ℓ̃, [[N]]) ∧ isCurrent(x))
  set (end = now)
[[update(ℓi:Ai)i (x ← L) where M set (ℓ = Nj)j ∈ J]] =
  let tbl = [[L]] in
  let affected =
    query
    for (x ← get tbl)
      where ((restrict(x, {ℓi}i ∈ I, [[M]]) ∧ isCurrent(x))
        ⌊
          (ℓi = x.ℓi)i ∈ I ⊕
          (ℓj = restrict(x, {ℓi}i ∈ I, [[Nj]])j ∈ J ⊕
          (start = now, end = ∞)
        ⌋
    in
  update (x ← tbl)
  where (restrict(x, ℓ̃, [[M]]) ∧ isCurrent(x))
  set (end = now);
  insert tbl values affected

```

Figure 5. Translation from λ_{TLINQ} into λ_{LINQ} **insert employees values**

(name = "Dolores", position = "Professor", salary = 70000)

Next, we want to record that Alice has resigned. We can write the following *current deletion* query:

```
delete (x ← employees) where x.name = "Alice"
```

The resulting table state shows that Dolores is a Professor from the current time onwards, and that the ‘end’ field of Alice’s current row is updated to the current year:

name	position	salary	start	end
Alice	Lecturer	40000	2010	2018
Alice	Senior Lecturer	50000	2018	2022
Dolores	Professor	70000	2022	∞
...

A powerful feature of valid-time databases is the ability to perform *sequenced modifications*, which apply an update or deletion over a particular *period of applicability* (PA). In

fact, current modifications are a special case of sequenced modifications applied from **now** until ∞ . Suppose that Dolores has agreed to act as Head of School between 2023 and 2028. We can record this using a *sequenced update* query:

```

update sequenced (x ← employees)
  between 2023 and 2025 where (x.name = "Dolores")
  set (position = "Head of School")

```

with the resulting table being:

name	position	salary	start	end
Dolores	Professor	70000	2022	2023
Dolores	Head of School	70000	2023	2028
Dolores	Professor	70000	2028	∞
...

Since the period of applicability of the sequenced update is entirely contained within the period of validity of Dolores’s row, we end up with three rows: the unchanged record before and after the PA, and the updated record during the PA. We also allow a *sequenced deletion*, and a *sequenced insertion*, where each record’s period of validity is given explicitly.

Additionally, suppose that all PhD students are to be given a 1-year extension due to the disruption caused by the pandemic; in this case we want to change the period of validity directly. This is known as a *nonsequenced update*. We cannot express this modification using either current or sequenced modifications since we must calculate the each row’s new end date from its previous end date. We can write the modification as follows, noting that we can both read from, and write to, the period of validity directly:

```

update nonsequenced (x ← employees)
  where ((data x).position = "PhD student")
  set () valid from (start x) to (end x + 1)

```

The resulting table shows that the ‘end’ field of Bob’s and Charles’ records are updated to 2024 and 2023 respectively:

name	position	salary	start	end
Bob	PhD Student	15000	2019	2024
Charles	PhD Student	15000	2018	2023
...

4.1 Calculus

The λ_{VLINQ} calculus gives a direct semantics to valid time operations. Like λ_{TLINQ} , λ_{VLINQ} has a native notion of a period-stamped database row, with accessors for the data and each timestamp; the typing rules, reduction rules, and translations are straightforward adaptations of those in λ_{TLINQ} .

Figure 6 shows how the syntax and typing rules for λ_{VLINQ} differ from those of λ_{LINQ} . Unlike in λ_{TLINQ} , we can use the term $M_1^{[M_2, M_3]}$ to construct a valid-time row. Ordinary insert, delete and update operations may be applied to valid time tables (the straightforward rules are omitted). Sequenced insertions are described by the term **insert sequenced M values N** where TV-SEQINSERT ensures that N is a bag of timestamped records. Sequenced updates are described by:

Syntax

Types	A, B	::=	$\text{ValidTime}(A)$
Terms	L, M, N	::=	$\dots \mid L^{M,N} \mid \mathbf{data} M \mid \mathbf{start} M \mid \mathbf{end} M \mid \mathbf{insert sequenced} M \mathbf{values} N$ $\mid \mathbf{update sequenced} (x \Leftarrow L) \mathbf{between} M_1 \mathbf{and} M_2 \mathbf{where} M_3 \mathbf{set} (\ell = \overline{N})$ $\mid \mathbf{update nonsequenced} (x \Leftarrow L) \mathbf{where} M \mathbf{set} (\ell = \overline{N}) \mathbf{valid from} N'_1 \mathbf{to} N'_2$ $\mid \mathbf{delete sequenced} (x \Leftarrow L) \mathbf{between} N_1 \mathbf{and} N_2 \mathbf{where} M$ $\mid \mathbf{delete nonsequenced} (x \Leftarrow M) \mathbf{where} N$
Values	V, W	::=	$\dots \mid V_1^{[V_2, V_3]}$

Typing rules

TV-GET	$\Gamma \vdash M:\text{Table}(A)!E$				$\boxed{\Gamma \vdash M : A!E}$	
	$\Gamma \vdash \mathbf{get} M:\text{Bag}(\text{ValidTime}(A))! \{\mathbf{read}\} \cup E$	TV-SEQINSERT	$\Gamma \vdash M:\text{Table}(A)!E$	$\Gamma \vdash N:\text{Bag}(\text{ValidTime}(A))! \emptyset$	$\Gamma \vdash \mathbf{insert sequenced} M \mathbf{values} N:(!)! \{\mathbf{write}\} \cup E$	
TV-SEQUPDATE	$\Gamma \vdash L:\text{Table}(A)!E$	$A = (\ell_i : B_i)_{i \in I}$	$\Gamma \vdash M_1:\text{Time}! \emptyset$	$\Gamma \vdash M_2:\text{Time}! \emptyset$	$\Gamma, x : A \vdash M_3:\text{Bool}! \emptyset$	$(j \in I \wedge \Gamma, x : A \vdash N_j:B_j! \emptyset)_{j \in J}$
	$\Gamma \vdash \mathbf{update sequenced} (x \Leftarrow L) \mathbf{between} M_1 \mathbf{and} M_2 \mathbf{where} M_3 \mathbf{set} (\ell_j = N_j)_{j \in J}(!) \{\mathbf{write}\} \cup E$					
TV-NONSEQUPDATE	$\Gamma \vdash L:\text{Table}(A)!E$	$A = (\ell_i : B_i)_{i \in I}$	$\Gamma, x : \text{ValidTime}(A) \vdash M:\text{Bool}! \emptyset$	$(j \in I \wedge \Gamma, x : \text{ValidTime}(A) \vdash N_j:B_j! \emptyset)_{j \in J}$	$\Gamma, x : \text{ValidTime}(A) \vdash N'_1:\text{Time}! \emptyset$	$\Gamma, x : \text{ValidTime}(A) \vdash N'_2:\text{Time}! \emptyset$
	$\Gamma \vdash \mathbf{update nonsequenced} (x \Leftarrow L) \mathbf{where} M \mathbf{set} (\ell_j = N_j)_{j \in J} \mathbf{valid from} N'_1 \mathbf{to} N'_2:(!)! \{\mathbf{write}\} \cup E$					
TV-SEQDELETE	$\Gamma \vdash L:\text{Table}(A)!E_1$	$\Gamma \vdash M_1:\text{Time}!E_2$	$\Gamma \vdash M_2:\text{Time}!E_3$	$\Gamma, x:A \vdash N:\text{Bool}! \emptyset$		
	$\Gamma \vdash \mathbf{delete sequenced} (x \Leftarrow L) \mathbf{between} M_1 \mathbf{and} M_2 \mathbf{where} N:(!)! \{\mathbf{write}\} \cup E_1 \cup E_2 \cup E_3$					
TV-NONSEQDELETE	$\Gamma \vdash M:\text{Table}(A)!E$	$\Gamma, x : \text{ValidTime}(A) \vdash N:\text{Bool}! \emptyset$				
	$\Gamma \vdash \mathbf{delete nonsequenced} (x \Leftarrow M) \mathbf{where} N:(!)! \{\mathbf{write}\} \cup E$					

Figure 6. Syntax and typing rules for λ_{VLINQ}

update sequenced $(x \Leftarrow L)$ **between** M_1 **and** M_2 **where** M_3 **set** $(\ell = \overline{N})$

Terms M_1 and M_2 must be of type Time , referring to the period of applicability of the sequenced update. Nonsequenced updates are described by the term:

update nonsequenced $(x \Leftarrow L)$
where M **set** $(\ell = \overline{N})$ **valid from** N'_1 **to** N'_2

with TV-NONSEQUPDATE stating that the database row (including period information) is bound as x in the predicate M , update terms N_j , and new time periods N'_1 and N'_2 .

Finally, the term:

delete sequenced $(x \Leftarrow L)$ **between** M_1 **and** M_2 **where** N

describes a sequenced deletion which removes the portion of each record satisfying N between times M_1 and M_2 .

Since current insertions, updates, and deletions are special cases of sequenced operations, we need not consider them explicitly; for completeness, direct semantics can be found in Appendix C. Instead, we show macro translations to the sequenced constructs. Current insertions can be implemented by desugaring to sequenced insertions, annotating each row with $[\mathbf{now}, \infty)$:

$\mathbf{insert} M \mathbf{values} N \rightsquigarrow \mathbf{let} \mathbf{rows} = \mathbf{for} (x \leftarrow N) \{x^{[\mathbf{now}, \infty)}\} \mathbf{in}$
 $\mathbf{insert sequenced} M \mathbf{values} \mathbf{rows}$

Current updates and deletions can be implemented as sequenced updates and deletions where the period of applicability spans from **now** until ∞ :

update $(x \Leftarrow L)$ **where** M **set** $(\ell_i = N_i)_i \rightsquigarrow$
 $\mathbf{update sequenced} (x \Leftarrow L)$
between now and ∞ **where** M **set** $(\ell_i = N_i)_i$
delete $(x \Leftarrow M)$ **where** $N \rightsquigarrow$
 $\mathbf{delete sequenced} (x \Leftarrow M)$ **between now and** ∞ **where** N

Fig. 7 shows selected reduction rules for sequenced operations. We show the rules for sequenced inserts and updates; the rules for other cases employ similar ideas and are included in Appendix A. Nonsequenced updates and deletes are similar to their analogues in λ_{LINQ} but allow access to, and modification of, row timestamps. For sequenced insertions, EV-SEQINSERT checks that the period of validity for each row is correct (i.e., that the *start* field is less than the *end* field) and appends the provided bag to the table. Sequenced updates and deletions must account for the various ways that the period of applicability can overlap the period of validity. There are five main cases, corresponding to the five ways two closed-open intervals can overlap (or fail to do so). Appendix B summarizes the five cases and describes how deletes and updates are handled in each case.

Reduction rules

$$M \Downarrow_{\Delta, t}^V (V, \Delta')$$

$$\begin{array}{c}
 \text{EV-ROW} \\
 \frac{M_1 \Downarrow_{\Delta, t}^V (V_1, \Delta_1) \quad M_2 \Downarrow_{\Delta, t}^V (V_2, \Delta_2) \quad M_3 \Downarrow_{\Delta, t}^V (V_3, \Delta_3)}{M_1^{[M_2, M_3]} \Downarrow_{\Delta, t}^V (V_1^{[V_2, V_3]}, \Delta_3)} \\
 \\
 \text{EV-SEQINSERT} \\
 \frac{M \Downarrow_{\Delta, t}^V (t, \Delta_1) \quad N \Downarrow_t^* \{\tilde{V}\} \quad \forall \text{data}^{[start, end]} \in \tilde{V}. \text{start} < \text{end} \quad \Delta_2 = \Delta_1 [t \mapsto \Delta_1(t) \hat{\cup} \{\tilde{V}\}]}{\text{insert sequenced } M \text{ values } N \Downarrow_{\Delta, t}^V ((), \Delta_2)} \\
 \\
 \text{EV-SEQUPDATE} \\
 \frac{L \Downarrow_{\Delta, t}^V (t, \Delta_1) \quad M_1 \Downarrow_t^* V_{start} \quad M_2 \Downarrow_t^* V_{end} \quad V_{start} < V_{end} \quad \Delta_2 = \Delta_1 [t \mapsto \hat{\cup} \{\uparrow \text{upd}(d) \mid d \in \Delta_1(t)\}]}{\text{update sequenced } (x \leftarrow L) \text{ between } M_1 \text{ and } M_2 \text{ where } M_3 \text{ set } (\ell_i = N_i)_i \Downarrow_{\Delta, t} ((), \Delta_2)} \\
 \text{upd}(v^{[start, end]}) = \begin{cases} \{W^{[start, end]}\} & \text{if } M_3\{v/x\} \Downarrow_t^* \text{true and } V_{start} \leq \text{start and } V_{end} \geq \text{end} \quad \text{(Case 1)} \\ \{W^{[start, V_{end}]}, v^{[V_{end}, end]}\} & \text{if } M_3\{v/x\} \Downarrow_t^* \text{true and } V_{start} \leq \text{start and } V_{end} < \text{end} \quad \text{(Case 2)} \\ \{v^{[start, V_{start}]}, W^{[V_{start}, V_{end}]}, v^{[V_{end}, end]}\} & \text{if } M_3\{v/x\} \Downarrow_t^* \text{true and } V_{start} > \text{start and } V_{end} < \text{end} \quad \text{(Case 3)} \\ \{v^{[start, V_{start}]}, W^{[V_{start}, end]}\} & \text{if } M_3\{v/x\} \Downarrow_t^* \text{true and } V_{start} > \text{start and } V_{end} \geq \text{end} \quad \text{(Case 4)} \\ \{v^{[start, end]}\} & \text{otherwise} \quad \text{(Case 5)} \end{cases} \\
 \text{where for all cases, } W = (v \text{ with } \ell = W') \text{ given } (N_i \Downarrow_t^* W'_i)_i \\
 \\
 \text{EV-NONSEQUPDATE} \\
 \frac{L \Downarrow_{\Delta, t}^V (t, \Delta_1) \quad \Delta_2 = \Delta_1 [t \mapsto \{\uparrow \text{upd}(d) \mid d \in \Delta_1(t)\}]}{\text{update nonsequenced } (x \leftarrow L) \text{ where } M \text{ set } (\ell_i = N_i)_i \text{ valid from } N'_1 \text{ to } N'_2 \Downarrow_{\Delta, t} ((), \Delta_2)} \\
 \text{upd}(D = v^{[start, end]}) = \begin{cases} (v \text{ with } \ell = W)^{[W_{start}, W_{end}]} & \text{if } M\{D/x\} \Downarrow_t^* \text{true and } (N_i\{D/x\} \Downarrow_t^* W_i)_i \text{ and} \\ & N'_1\{D/x\} \Downarrow_t^* W_{start} \text{ and } N'_2\{D/x\} \Downarrow_t^* W_{end} \text{ and } W_{start} < W_{end} \\ D & \text{if } M\{D/x\} \Downarrow_t^* \text{false} \end{cases}
 \end{array}$$

Figure 7. Reduction rules for λ_{VLINQ} (selected)

4.2 Translation

Figure 8 illustrates the translation from λ_{VLINQ} into λ_{LINQ} . We discuss the translations for sequenced inserts and both sequenced and nonsequenced updates; the other modifications are similar and included in Appendix A. As before, we require annotations on each of the database update terms.

Nonsequenced updates and deletions can be updated directly by their corresponding λ_{LINQ} operation; we use an auxiliary definition, lift, which lifts the flat representation into the nested representations expected by the predicate and update fields. Sequenced inserts flatten the contents of the provided bag and map directly to an **insert**.

The remaining sequenced operations are the most complex to translate. Since a sequenced modification may partition a row, the startRows and endRows functions calculate the records which must be inserted before and after the period of applicability. To translate a sequenced update, we calculate the rows to insert, perform an **update** to set the new values and set the new period of applicability to the overlap between the PA and PV using the **greatest** and **least** functions, and finally materialise the insertions. Sequenced deletions (shown in Appendix A) are similar but delete the rows that overlap the PA instead of updating them.

4.3 Metatheory

Evaluation preserves typing and well-formedness.

Proposition 4.1 (Preservation (λ_{VLINQ})). *If $\vdash M:A!E$ and $M \Downarrow_{\Delta, t}^V (V, \Delta')$ for some $\text{wf}(\Delta)$, then $\cdot \vdash V:A!\emptyset$ and $\text{wf}(\Delta')$.*

Unlike λ_{LINQ} and λ_{TLINQ} , evaluation in λ_{VLINQ} is *partial* in order to reflect the need for dynamic checks that start times precede end times. In practice, our implementation evaluates temporal updates as single transactions and raises an exception (aborting the transaction) when a well-formedness check fails, but our formalisation assumes updates preserve well-formedness in order to avoid clutter.

Our translation from λ_{VLINQ} into λ_{TLINQ} satisfies the following correctness property:

Theorem 4.1. *If $\cdot \vdash M:A!E$ and $M \Downarrow_{\Delta, t}^V (V, \Delta')$ for some $\text{wf}(\Delta)$, then $(M) \Downarrow_{\Delta, t} (V, \Delta')$*

5 Sequenced Joins

Queries that join multiple tables are straightforward to encode using language integrated query. Keeping with our employee database, say we wish to separate out the salary into a separate table. The non-temporal employee database might look as follows:

employees			salaries	
name	position	band	band	salary
Alice	Senior Lecturer	A08	A08	40000
Bob	PhD Student	B01	A09	50000
Charles	PhD Student	B01	A10	70000
Dolores	Professor	A10	B01	15000

We can get the salary for each employee as follows:

```

(insert( $\ell_i:A_i$ ) $i \in I$  sequenced  $M$  values  $N$ ) =
  let  $tbl = \langle M \rangle$  in
  let  $rows =$ 
    for ( $x \leftarrow \langle N \rangle$ )
       $\{ \eta(x.data, \ell) \oplus (start = x.start, end = x.end) \}$ 
  in
  insert  $tbl$  values  $rows$ 

(update( $\ell_i:A_i$ ) $i \in I$  sequenced ( $x \leftarrow L$ )
  between  $M_1$  and  $M_2$  where  $M_3$ 
  set ( $\ell_j = N_j$ ) $j \in J$ ) =
  let  $tbl = \langle L \rangle$  in
  let  $aStart = \langle M_1 \rangle$  in
  let  $aEnd = \langle M_2 \rangle$  in
  let  $lRows = startRows(tbl, pred, aStart)$  in
  let  $rRows = endRows(tbl, pred, aEnd)$  in
  update ( $x \leftarrow tbl$ )
    where ( $pred \wedge (x.start < aEnd) \wedge (x.end > aStart)$ )
    set  $\left( \begin{array}{l} (\ell_j = restrict(x, \{\ell_i\}_{i \in I}, \langle N_j \rangle))_{j \in J}, \\ start = \mathbf{greatest}(x.start, aStart), \\ end = \mathbf{least}(x.end, aEnd) \end{array} \right);$ 
  insert  $tbl$  values  $lRows$ ;
  insert  $tbl$  values  $rRows$ 

```

where

```

pred  $\triangleq$   $restrict(x, \{\ell_i\}_{i \in I}, \langle M_3 \rangle)$ 
startRows( $tbl, pred, aStart$ )  $\triangleq$  query
  for ( $x \leftarrow get\ tbl$ )
    where ( $pred \wedge (x.start < aStart) \wedge (x.end > aStart)$ )
     $\{ \eta(x, \{\ell_i\}_{i \in I}) \oplus (start = x.start, end = aStart) \}$ 
endRows( $tbl, pred, aEnd$ )  $\triangleq$  (symmetric)

```

```

(update( $\ell_i:A_i$ ) $i \in I$  nonsequenced ( $x \leftarrow L$ ) where  $M$ 
  set ( $\ell_j = N_j$ ) $j \in J$  valid from  $N'_1$  to  $N'_2$ ) =
  update( $\ell_i:A_i$ ) $i \in I$  ( $x \leftarrow \langle L \rangle$ )
  where ( $lift(x, \langle M \rangle)$ )
  set  $\left( \begin{array}{l} ((\ell_j = lift(x, \langle N_j \rangle))_{j \in J}), \\ start = lift(x, \langle N'_1 \rangle), \\ end = lift(x, \langle N'_2 \rangle) \end{array} \right)$ 

```

where $lift(x, f) \triangleq$

```

( $\lambda x.f$ ) ( $data = \eta(x, \{\ell_i\}_{i \in I}), start = x.start, end = x.end$ )

```

Figure 8. Translation from λ_{VLIQ} into λ_{LIQ} (selected cases)

```

query
for ( $e \leftarrow get\ employees$ )
for ( $s \leftarrow get\ salaries$ )
  where ( $e.band = s.band$ )
   $\{ (name = e.name, salary = s.salary) \}$ 

```

name	salary
Alice	40000
Bob	15000
Charles	15000
Dolores	70000

Joining a temporal table with a non-temporal table is also easily expressible. Consider a version of our previous temporal employees table from just after when Dolores joined:

name	position	band	start	end
Alice	Lecturer	A08	2010	2018
Alice	Senior Lecturer	A09	2018	∞
Bob	PhD Student	B01	2019	2023
Charles	PhD Student	B01	2018	2022
Dolores	Professor	A10	2022	∞

We can join this table with the non-temporal salaries table as follows; for clarity, we denote valid-time **get** as **get_v**:

```

query
for ( $e \leftarrow get_v\ employees$ )
for ( $s \leftarrow get_v\ salaries$ )
  where ( $(data\ e).band = s.band$ )
   $\{ (name = e.name, salary = s.salary) \}^{[start\ e.end\ e]}$ 

```

giving us the corresponding table in Section 4.

Things get more interesting when *both* tables are temporal. Salaries are not static over time; bands go up with inflation, for example. Suppose we now have two temporal tables. Consider the above table along with a temporal salaries table showing a pay increase in 2015:

band	salary	start	end
A08	38000	2000	2015
A09	48000	2000	2015
A08	40000	2015	∞
A09	50000	2015	∞
...

What does it mean to join two *temporal* tables? In essence, we want to record *all* configurations of a particular joined record, creating new records with shorter periods of validity whenever data from either underlying table changes. Concretely, joining the above two temporal tables would give:

name	salary	start	end
Alice	38000	2010	2015
Alice	40000	2015	2018
Alice	50000	2018	∞
...

Now there are records for Alice for *three* different periods:

- The first when Alice was on salary band A08, conferring a salary of £38000.
- The second when band A08 increased to £40000.
- The third when Alice was promoted to band A09.

Such joins are called *sequenced* because they (conceptually) evaluate the join on the whole sequence of states encoded by each table. Manually writing the sequenced joins in SQL is error-prone. We instead introduce a construct, **join**, which allows us to write the following:

```

join
for ( $e \leftarrow get_v\ employees$ )
for ( $s \leftarrow get_v\ salaries$ )
  where ( $(data\ e).band = (data\ s).band$ )
   $\{ (name = (data\ e).name, salary = (data\ s).salary) \}$ 

```

Note that we *do not* need to calculate the period of validity for each resulting row; this is computed automatically.

Figure 9 shows how sequenced joins can be implemented; we show the constructs for valid time, but the same technique can be used for transaction time. The typing rule requires that the result of a **join** query is *flat* (nested sequenced queries are conceptually nontrivial).

As mentioned earlier, queries can be rewritten to normal forms for conversion to SQL, as shown in Figure 9. The structure of these normal forms allows sequenced joins to be implemented through a simple rewrite: the **greatest** and **least** functions are used to calculate the intersections of the periods of validity for each combination of records from

Typing rules

$$\begin{array}{c}
\boxed{A :: \text{FQType}} \quad \boxed{\Gamma \vdash M:A!E} \\
\hline
C :: \text{FQType} \qquad (\ell_i : C_i)_i :: \text{FQType} \\
\hline
\Gamma \vdash M:\text{Bag}(A)!E \quad A :: \text{FQType} \quad E \subseteq \{\text{read}\} \\
\hline
\Gamma \vdash \text{join } M:\text{Bag}(\text{ValidTime}(A))!E
\end{array}$$

Normal forms

Queries	$Q ::= K_1 \uplus \dots \uplus K_n$
Comprehensions	$K ::= \text{for } (\tilde{G}) \text{ where } P \wr S \wr$
Generators	$G ::= x \leftarrow \text{get } t \mid x \leftarrow \text{get}_v t$
Normalised terms	$S ::= P \mid R$
Base terms	$P ::= c \mid x.\ell \mid \odot\{\vec{P}\}$
Record terms	$R ::= (\ell_i = P_i)_i$

Translation on normal forms

$$\begin{array}{l}
\|\text{join } Q\| = \text{query } \|Q\| \qquad \|K_1 \uplus \dots \uplus K_n\| = \|K_1\| \uplus \dots \uplus \|K_n\| \\
\|P\| = P \qquad \|R\| = R
\end{array}$$

$$\left\| \text{for} \left(\begin{array}{l} x_1 \leftarrow \text{get}_v t_1, \dots, x_m \leftarrow \text{get}_v t_m, \\ y_1 \leftarrow \text{get } t'_1, \dots, y_n \leftarrow \text{get } t'_n \\ \text{where } (P) \end{array} \right) \wr R \wr \right\| = \left\| \text{for} \left(\begin{array}{l} x_1 \leftarrow \text{get}_v t_1, \dots, x_m \leftarrow \text{get}_v t_m, \\ y_1 \leftarrow \text{get } t'_1, \dots, y_n \leftarrow \text{get } t'_n \\ \text{where } (P \wedge \text{joinStart} < \text{joinEnd}) \end{array} \right) \wr R^{\text{joinStart, joinEnd}} \wr \right\|$$

where $\text{joinStart} \triangleq \text{greatest}(x_1.\text{start}, \dots, x_m.\text{start})$
 $\text{joinEnd} \triangleq \text{least}(x_1.\text{end}, \dots, x_m.\text{end})$

Figure 9. Sequenced joins

each generator, with the modified predicate ensuring that the periods of overlap make sense. The calculated overlapping periods of validity are then returned in the resulting row.

6 Implementation and Case Study

The Links programming language [9] is a statically-typed functional web programming language which allows client, server, and database code to be written in a uniform language. We have extended Links with support for the constructs described in Sections 3 and 4, as well as support for temporal joins as described in Section 5. In this section, we describe a case study based on curating COVID-19 data.

Our translations from λ_{VLINQ} into λ_{LINQ} are trivially realisable in SQL. Queries can be compiled using known techniques (e.g., [8]). The `startRows` and `endRows` functions can be compiled using an SQL `WITH` statement, and there is a direct correspondence between λ_{LINQ} modification operations and their SQL equivalents. Each translated temporal modification is executed as an SQL transaction, with primary key and referential integrity constraint checking deferred until the end of the transaction.

Case study. We have used the temporal features of Links in two prototypes based on curated scientific databases: curation of publicly available COVID-19 data, and storage and

curation of XML documents [16] using the Dynamic Dewey labelling algorithm [36].

We concentrate on the first application¹; a previous prototype which used a preliminary version of the language design has previously been presented as a short demo paper [15]. In 2020, the Scottish Government began releasing various data about the COVID-19 pandemic [24], which included weekly data of fatalities in Scotland in various categories (such as ‘Sex’). Each weekly release was a CSV file, with a row for each subcategory (e.g., ‘Sex’ has the subcategories ‘Male’ and ‘Female’) and a column for each week for which data was available. Each release included an additional week column with the latest data (see Figure 10). Importantly, each release could include revisions to data for previous weeks.

Information about the changes to the data over time is often desired to understand its provenance and assess its trustworthiness [4]. From a provenance point of view, this data is interesting because a column for an earlier week may contain updated data. We developed a web application for the querying of the data (“How do the Male and Female subcategories compare in terms of the change in fatalities from last week to this week?”) as well as querying the changes in the data (“How do the Male and Female subcategories compare in terms of number of updates to existing values?”).

Considering the non-temporal data, an entry in a database table would be a row consisting of the key fields `subcat` and `weekdate` and a value field giving the corresponding count. In the case of the temporal data, the key fields are insufficient to uniquely identify the value of the `count` because it may have different values over time. Thus the time validity fields are necessary to provide a key for the value.

The prototype uses a valid time table for fatality data to capture the notion that a count value, either brand new or an update, becomes valid as soon as the CSV is uploaded into the interface² (this can be a different time from when the new value is accepted and written to the database). In Links it is possible to specify the names of the period stamping fields, which have the built-in type `DateTime`. This table is defined using the following Links code; we have omitted some details in the code snippets for brevity.

```

var covid_data =
  table "covid_data"
  with (subcat: Int, weekdate: String, count: Int)
  using valid_time(valid_from, valid_to)
  from database "covid_curation";

```

The prototype’s upload workflow is as follows: the user uploads a new CSV file, and the count values for the new week are added to the database. For counts that pertain to earlier weeks and that now have different values, the user is shown these counts and can accept them, reject them or move them

¹Links code is available at <https://github.com/XAZKQV26IL>

²Other possibilities for the start time of validity are the date of the release of the CSV file or the start of the new week.

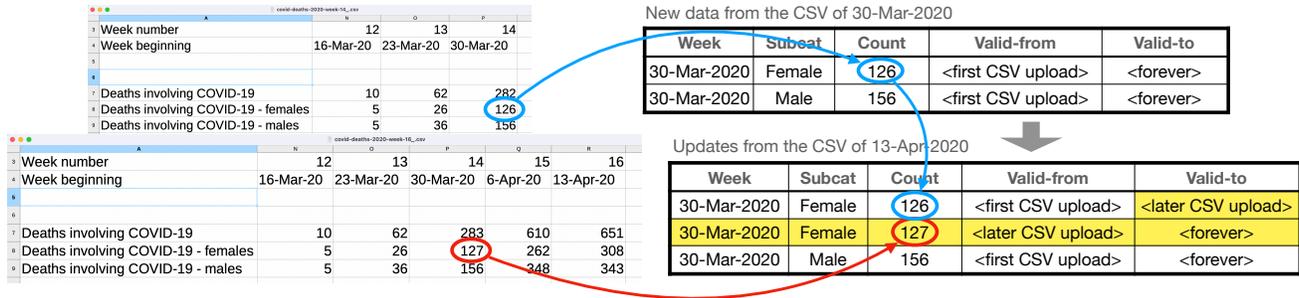


Figure 10. Example of data uploads, sequenced insertion and sequenced update

to a pending list for a later decision. In terms of implementation, brand new count values are added to the table with a *sequenced insert*, using the upload time as the start time. The Links code for this and other examples can be found in Appendix E. The process is more complex for updated count values, because the interface shows the user previous value, to support decision making. This requires a conditional join over the current state of the `covid_data` table and the count values from the CSV file. If a modification is accepted, it is added using a *sequenced update*. Figure 10 illustrates how the table changes as a result of a single update.

The prototype also provides functionality to query data, both as current data, and as data with information about changes. The current data is obtained using a *current query*. The result is a list of weeks and counts grouped by subcategory. This is repeated for each category.

```
fun getCurrentData (category) {
  query nested {
    for (x <- subcategory)
      where (x.cat == category)
      [(subcat_name = x.subcat_name, cat = x.cat,
        results =
          for (y <- vtCurrent(covid_data))
            for (z <- week)
              where (y.subcat == x.subcat &&
                y.weekdate == z.weekdate &&
                z.all_zero == false)
                [(count = y.count, weekdate = y.weekdate)]])]
  }
}
```

Instead of an explicit `get` construct, Links uses the ‘double arrow’ comprehension `<-` to represent a nontemporal database query, with `<-t-` and `<-v-` supporting transaction time and valid time queries respectively. The ‘single arrow’ comprehension `<-` denotes a list comprehension. Finally, `vtCurrent` is a standard library function which performs a valid time query to obtain the values valid at the current time.

For update provenance queries of individual counts, a self join is computed over the `subcategory` and `week` fields of the valid time table to provide a nested result table where each count is associated with a list of count values and their associated start and end time information. This is a nonsequenced

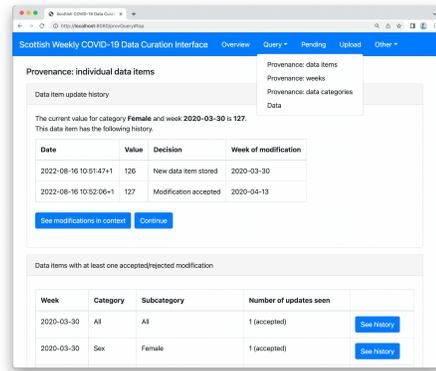


Figure 11. Interface screenshot: history of a count

query because the time period information is explicitly added to the result table. The user can specify the subcategory and week they are interested in, and obtain details of modifications. The interface also supports update provenance by week and by category. This is illustrated in Figure 11.

7 Related and Future Work

Most of the focus of effort on language-integrated query has been (perhaps unsurprisingly) on queries rather than updates, beginning with the foundational work on nested relational calculus by Buneman et al. [3] and on rewriting queries for translation to SQL by Wong [35]. Lindley and Cheney [21] presented a calculus including both query and update capabilities and our type and effect system for tracking database read and write access is loosely based on theirs. More recently a number of contributions extending the formal foundations of language-integrated query have appeared, including to handle higher-order functions [8], nested query results [6], sorting/ordering [19], grouping and aggregation [25, 26], and deduplication [28]. Our core calculus λ_{LINQ} only incorporates the first two of these, and developing a core calculus that handles more features, as well as translating temporal queries involving them, is an obvious future direction. To the best of our knowledge no previous work on language-integrated query has considered temporal data specifically.

The ubiquity and importance of time in applications of databases was appreciated from an early stage [7] and led to a significant community effort to standardize temporal extensions to SQL based on the TSQL2 language design in the 1990s and early 2000s [31]. This effort ultimately resulted in standardisation of a relatively limited subset of the original proposal in SQL:2011 [20]. Since then temporal database research has progressed steadily, including recent contributions showing how to implement temporal data management as a layer on top of a standard RDBMS [11], and establishing connections between temporal querying and data provenance and annotation models [12].

Snodgrass [32] describes how to implement TSQL2-style updates and queries by translation to SQL, but we are not aware of previous detailed formal proofs of correctness of translations for transaction time and valid time updates. Although timestamping rows with time intervals is among the most popular ways to represent temporal databases as flat relational tables, it is not the only possibility. Jensen et al. [18] proposed a *bitemporal conceptual data model* that captures the abstract meaning of a temporal table and used it to compare different representation strategies.

There can be multiple representations of the same abstract temporal data, leading to consideration of the problem of coalescing or normalizing the intervals to save space and avoid ambiguity. Nonsequenced updates can be used to perform modifications that have different effects on representations of the same conceptual table. We have not considered coalescing or other common issues such as how to handle operations such as deduplication, grouping and aggregation (including emptiness testing), or integrity constraints in a temporal setting. Some of these issues appear orthogonal to the high-level language design and could be incorporated “under the hood” into the implementation or even performed directly on the database.

One important future application is to retrofit temporal aspects to expert-curated databases, an example being the Guide to Pharmacology Database (GtoPdb) that summarises pharmacological targets and interactions [2]. Links has been used to implement a workalike version of GtoPdb [14] and we hope to build on this to provide a fully versioned implementation of GtoPdb. An important requirement here is to minimize changes to the existing system.

Finally we mention two immediate next steps. First, we plan to investigate *bitemporal* databases [33] allow transaction and valid time to be used together, allowing us to write queries such as “when was it recorded that Bob’s contract length was extended?”; *bitemporal* databases are considerably more difficult to formalise and reason about, so we aim to investigate how *bitemporal* support can be added in a compositional manner. Second, at present, the result of a sequenced join must be a flat record; further work is required to understand the semantics and implementation techniques for joins that produce nested results.

8 Discussion

Efficiency. Our main focus has been on *portability*: by distilling a language design and formalising and implementing a translation from temporal calculi to non-temporal calculi, we allow temporal functionality to be used on a mainstream DBMS. The cost of portability is that our translations will inevitably not perform as well as a native implementation. Although we do not make any specific claims about efficiency, we have no reason to believe that the performance is any different to hand-translated SQL.

In particular, as discussed in §6, all translated queries can be run directly on the database and do not require in-memory processing. Previous work on language-integrated query (e.g., [6, 9]) shows how the “nested loop” style of query is translated into efficient SQL, and our translation of queries happens prior to normalisation. Further optimisation is subsequently performed by the DBMS, and anecdotally we have not observed any performance issues in any applications we have written using the temporal extensions to Links.

Supporting existing native implementations. Some (mainly proprietary) DBMSs, for example Teradata [1], have native support for some temporal features inspired by TSQL2 or SQL:2011. Although we are yet to explore this, an advantage of the LINQ approach is that temporal SQL syntax could be generated for backends which support temporal operations directly, while maintaining functionality in mainstream backends without native temporal support.

The translation between our temporal modification constructs and SQL:2011 (and similarly, TSQL2) would be fairly direct. Consider a sequenced update in λ_{VLINQ} :

```
update sequenced (x ← employees)
  between 2022-01-1 and 2022-10-29 where (x.salary < 30000)
  set (salary = x.salary + 1000)
```

This could be implemented in SQL:2011 as follows:

```
UPDATE employees
FOR PORTION OF EPeriod
FROM DATE '2022-01-01' TO DATE '2022-10-29'
WHERE salary < 30000
SET salary = salary + 1000
```

Moving between the two DBMSs would not require any changes to application source code. However, not all operations can be as straightforwardly translated: in particular, SQL:2011 does not natively support sequenced joins.

9 Conclusions

In spite of decades of work on temporal databases and even an extension to the SQL standard, mainstream support for temporal data remains limited, requiring developers to implement temporal functionality from scratch. In this paper, we have shown how to extend language-integrated query to support transaction time and valid time data, making temporal data management accessible without explicit DBMS

support. We have formalised our constructs and translational implementation strategies based on those proposed by Snodgrass [32], and proved that the translations are semantics-preserving. We have implemented our approach in the Links programming language and assessed its value through a case study. Our work is a first but significant step towards fully supporting temporal data management at the language level.

Acknowledgements

We thank the anonymous reviewers for their helpful comments. This work is partially funded by EPSRC Grant EP/T014628/1 (STARDUST), ERC Consolidator Grant 682315 (Skye), and a UK Government ISCF Metrology Fellowship.

References

- [1] Mohammed Al-Kateb, Ahmad Ghazal, Alain Crolotte, Ramesh Bhashyam, Jaiprakash Chimanchode, and Sai Pavan Pakala. 2013. Temporal query processing in Teradata. In *EDBT*. ACM, 573–578.
- [2] J.F. Armstrong, E. Faccenda, S.D. Harding, A.J. Pawson, C. Southan, J.L. Sharman, B. Campo, D.R. Cavanagh, S.P.H. Alexander, A.P. Davenport, M. Spedding, and J.A. Davies. 2020. The IUPHAR/BPS Guide to PHARMACOLOGY in 2020: extending immunopharmacology content and introducing the IUPHAR/MMV Guide to MALARIA PHARMACOLOGY. *Nucleic Acids Research* 48 (2020), D1006–D1021. <https://doi.org/10.1093/nar/gkz951>
- [3] Peter Buneman, Shamim Naqvi, Val Tannen, and Limsoon Wong. 1995. Principles of programming with complex objects and collection types. *Theor. Comput. Sci.* 149, 1 (1995). [https://doi.org/10.1016/0304-3975\(95\)00024-Q](https://doi.org/10.1016/0304-3975(95)00024-Q)
- [4] Peter Buneman and Wang-Chiew Tan. 2018. Data Provenance: What next? *SIGMOD Rec.* 47, 3 (2018), 5–16. <https://doi.org/10.1145/3316416.3316418>
- [5] James Cheney, Sam Lindley, and Philip Wadler. 2013. A practical theory of language-integrated query. In *ICFP*. <https://doi.org/10.1145/2500365.2500586>
- [6] James Cheney, Sam Lindley, and Philip Wadler. 2014. Query shredding: efficient relational evaluation of queries over nested multisets. In *SIGMOD*. ACM, 1027–1038. <https://doi.org/10.1145/2588555.2612186>
- [7] James Clifford and David S. Warren. 1983. Formal Semantics for Time in Databases. *ACM Trans. Database Syst.* 8, 2 (1983), 214–254. <https://doi.org/10.1145/319983.319986>
- [8] Ezra Cooper. 2009. The Script-Writer’s Dream: How to Write Great SQL in Your Own Language, and Be Sure It Will Succeed. In *DBPL (Lecture Notes in Computer Science, Vol. 5708)*. Springer, 36–51. https://doi.org/10.1007/978-3-642-03793-1_3
- [9] Ezra Cooper, Sam Lindley, Philip Wadler, and Jeremy Yallop. 2006. Links: Web Programming Without Tiers. In *FMCO (Lecture Notes in Computer Science, Vol. 4709)*. Springer, 266–296. https://doi.org/10.1007/978-3-540-74792-5_12
- [10] George Copeland and David Maier. 1984. Making Smalltalk a database system. *SIGMOD Rec.* 14, 2 (1984). <https://doi.org/10.1145/602259.602300>
- [11] Anton Dignös, Michael H. Böhlen, Johann Gamper, and Christian S. Jensen. 2016. Extending the Kernel of a Relational DBMS with Comprehensive Support for Sequenced Temporal Queries. *ACM Trans. Database Syst.* 41, 4, Article 26 (2016), 46 pages. <https://doi.org/10.1145/2967608>
- [12] Anton Dignös, Boris Glavic, Xing Niu, Michael Böhlen, and Johann Gamper. 2019. Snapshot Semantics for Temporal Multiset Relations. *Proc. VLDB Endow.* 12, 6 (2019), 639–652. <https://doi.org/10.14778/3311880.3311882>
- [13] Stefan Fehrenbach and James Cheney. 2018. Language-integrated provenance. *Sci. Comput. Program.* 155 (2018), 103–145. <https://doi.org/10.1016/j.scico.2017.08.009>
- [14] Simon Fowler, Simon Harding, Joanna Sharman, and James Cheney. 2020. Cross-tier Web Programming for Curated Databases: A Case Study. *International Journal of Digital Curation* 16, 1 (2020). <https://doi.org/10.2218/ijdc.v15i1.717>
- [15] Vashti Galpin and James Cheney. 2021. Curating Covid-19 Data in Links. In *IPAW 2020 + IPAW 2021 (Lecture Notes in Computer Science, Vol. 12839)*. Springer, 237–243. https://doi.org/10.1007/978-3-030-80960-7_19
- [16] Vashti Galpin, Ian Smith, and Jean-Laurent Hippolyte. 2022. Supporting provenance of digital calibration certificates with temporal databases. In *IMEKO TC6 M4Dconf*. To appear.
- [17] Christian S. Jensen and Richard T. Snodgrass. 1999. Temporal Data Management. *IEEE Trans. Knowl. Data Eng.* 11, 1 (1999), 36–44. <https://doi.org/10.1109/69.755613>

- [18] Christian S. Jensen, Michael D. Soo, and Richard T. Snodgrass. 1994. Unifying Temporal Data Models via a Conceptual Model. *Inf. Syst.* 19, 7 (1994), 513–547. [https://doi.org/10.1016/0306-4379\(94\)90013-2](https://doi.org/10.1016/0306-4379(94)90013-2)
- [19] Oleg Kiselyov and Tatsuya Katsushima. 2017. Sound and Efficient Language-Integrated Query - Maintaining the ORDER. In *APLAS 2017*. 364–383. https://doi.org/10.1007/978-3-319-71237-6_18
- [20] Krishna Kulkarni and Jan-Eike Michels. 2012. Temporal Features in SQL:2011. *SIGMOD Rec.* 41, 3 (2012), 34–43. <https://doi.org/10.1145/2380776.2380786>
- [21] Sam Lindley and James Cheney. 2012. Row-based effect types for database integration. In *TLDI*. <https://doi.org/10.1145/2103786.2103798>
- [22] John M. Lucassen and David K. Gifford. 1988. Polymorphic Effect Systems. In *POPL*. ACM Press, 47–57. <https://doi.org/10.1145/73560.73564>
- [23] Erik Meijer, Brian Beckman, and Gavin M. Bierman. 2006. LINQ: reconciling object, relations and XML in the .NET framework. In *SIGMOD*. <https://doi.org/10.1145/1142473.1142552>
- [24] National Records of Scotland. 2021. Deaths involving coronavirus (COVID-19) in Scotland: archive. <https://www.nrscotland.gov.uk/statistics-and-data/statistics/statistics-by-theme/vital-events/general-publications/weekly-and-monthly-data-on-births-and-deaths/deaths-involving-coronavirus-covid-19-in-scotland/archive> Last accessed: 23 April 2021.
- [25] Rui Okura and Yuki Yoshi Kameyama. 2020. Language-Integrated Query with Nested Data Structures and Grouping. In *FLOPS*. 139–158. https://doi.org/10.1007/978-3-030-59025-3_9
- [26] Rui Okura and Yuki Yoshi Kameyama. 2020. Reorganizing Queries with Grouping. In *GPCE (Virtual, USA)*. ACM, 13 pages. <https://doi.org/10.1145/3425898.3426960>
- [27] Quill team. 2022. Quill: Compile-Time Language Integrated Queries for Scala. Open source project. <https://github.com/getquill/quill>.
- [28] Wilmer Ricciotti and James Cheney. 2021. Query Lifting: Language-integrated query for heterogeneous nested collections. In *ESOP*. 579–606. https://doi.org/10.1007/978-3-030-72019-3_21
- [29] Mark A. Roth, Henry F. Korth, and Abraham Silberschatz. 1988. Extended Algebra and Calculus for Nested Relational Databases. *ACM Trans. Database Syst.* 13, 4 (1988), 389–417. <https://doi.org/10.1145/49346.49347>
- [30] Lwin Khin Shar and Hee Beng Kuan Tan. 2013. Defeating SQL Injection. *Computer* 46, 3 (2013), 69–77. <https://doi.org/10.1109/MC.2012.283>
- [31] Richard T. Snodgrass (Ed.). 1995. *The TSQL2 Temporal Query Language*. Kluwer.
- [32] Richard T. Snodgrass. 1999. *Developing Time-Oriented Database Applications in SQL*. Morgan Kaufmann.
- [33] Richard T. Snodgrass and Ilsoo Ahn. 1985. A Taxonomy of Time in Databases. In *SIGMOD Conference*. ACM Press, 236–246.
- [34] Don Syme. 2006. Leveraging .NET meta-programming components from F#: integrated queries and interoperable heterogeneous execution. In *ML Workshop*.
- [35] Limsoon Wong. 1996. Normal Forms and Conservative Extension Properties for Query Languages over Collection Types. *J. Comput. Syst. Sci.* 52, 3 (1996). <https://doi.org/10.1006/jcss.1996.0037>
- [36] Liang Xu, Tok Wang Ling, Huayu Wu, and Zhifeng Bao. 2009. DDE: from dewey to a fully dynamic XML labeling scheme. In *SIGMOD*. ACM, 719–730. <https://doi.org/10.1145/1559845.1559921>

A Full definitions

In this appendix full definitions are provided for judgments and definitions where only selected rules could be included in the main body of the paper.

- Figure 12 presents the full definition of the typechecking judgment for λ_{LINQ} , complementing the partial definition in Figure 2.
- Figure 13 presents the full definition of the evaluation relation for λ_{LINQ} , including the rules omitted from Figure 3.
- Figure 14 shows the full typing rules for λ_{VLINQ} , complementing the selected rules in Figure 6.
- Figure 15 shows the full evaluation rules for λ_{VLINQ} , complementing the selected rules in Figure 7.
- Figure 16 presents the full translation rules for sequenced and nonsequenced operations in λ_{VLINQ} to λ_{LINQ} , completing the rules shown in Figure 8. The rules for translating current updates are special cases of those for sequenced updates and shown in Appendix C.

Term typing

				$\Gamma \vdash M : A ! E$
T-CONST	T-VAR	T-ABS	T-APP	T-IF
$\frac{c \text{ has base type } C}{\Gamma \vdash c : C ! \emptyset}$	$\frac{x : A \in \Gamma}{\Gamma \vdash x : A ! \emptyset}$	$\frac{\Gamma, x : A \vdash M : B ! E}{\Gamma \vdash \lambda x. M : A \rightarrow^E B ! \emptyset}$	$\frac{\Gamma \vdash M : A \rightarrow^{E_1} B ! E_2 \quad \Gamma \vdash N : A ! E_3}{\Gamma \vdash M N : B ! E_1 \cup E_2 \cup E_3}$	$\frac{\Gamma \vdash L : \text{Bool} ! E_1 \quad \Gamma \vdash M : A ! E_2 \quad \Gamma \vdash N : A ! E_3}{\Gamma \vdash \text{if } L \text{ then } M \text{ else } N : A ! E_1 \cup E_2 \cup E_3}$
T-OP	T-RECORD	T-PROJECT	T-EMPTYBAG	
$\frac{\odot : C_1 \times \dots \times C_n \rightarrow^E C \quad (\Gamma \vdash M_i : C_i ! E_i)_{i \in 1..n}}{\Gamma \vdash \odot \{\vec{M}\} : C ! E \cup \bigcup_{i \in 1..n} E_i}$	$\frac{(\Gamma \vdash M_i : A_i ! E_i)_i}{\Gamma \vdash (\ell_i = M_i)_i : (\ell_i : A_i)_i ! \bigcup_i E_i}$	$\frac{\Gamma \vdash M : (\ell_i : A_i)_{i \in I} ! E \quad j \in I}{\Gamma \vdash M. \ell_j : A_j ! E}$	$\frac{}{\Gamma \vdash \{\ \} : \text{Bag}(A) ! \emptyset}$	
T-BAG	T-BAGUNION	T-FOR	T-QUERY	
$\frac{\Gamma \vdash M : A ! E}{\Gamma \vdash \{M\} : \text{Bag}(A) ! E}$	$\frac{\Gamma \vdash M : \text{Bag}(A) ! E_1 \quad \Gamma \vdash N : \text{Bag}(A) ! E_2}{\Gamma \vdash M \uplus N : \text{Bag}(A) ! E_1 \cup E_2}$	$\frac{\Gamma \vdash M : \text{Bag}(A) ! E_1 \quad \Gamma, x : A \vdash N : \text{Bag}(B) ! E_2}{\Gamma \vdash \text{for } (x \leftarrow M) N : \text{Bag}(B) ! E_1 \cup E_2}$	$\frac{\Gamma \vdash M : \text{Bag}(A) ! E \quad A :: \text{QType} \quad E \subseteq \{\text{read}\}}{\Gamma \vdash \text{query } M : \text{Bag}(A) ! E}$	
T-TABLE	T-NOW	T-GET	T-INSERT	
$\frac{}{\Gamma \vdash t : \text{Table}(\Sigma(t)) ! \emptyset}$	$\frac{}{\Gamma \vdash \text{now} : \text{Time} ! \emptyset}$	$\frac{\Gamma \vdash M : \text{Table}(A) ! E}{\Gamma \vdash \text{get } M : \text{Bag}(A) ! \{\text{read}\} \cup E}$	$\frac{\Gamma \vdash M : \text{Table}(A) ! E \quad \Gamma \vdash N : \text{Bag}(A) ! \emptyset}{\Gamma \vdash \text{insert } M \text{ values } N : () ! \{\text{write}\} \cup E}$	
T-UPDATE			T-DELETE	
$\frac{A = (\ell_i : B_i)_{i \in I} \quad \Gamma \vdash L : \text{Table}(A) ! E \quad \Gamma, x : A \vdash M : \text{Bool} ! \emptyset \quad (j \in I \wedge \Gamma, x : A \vdash N_j : B_j ! \emptyset)_{j \in J}}{\Gamma \vdash \text{update } (x \leftarrow L) \text{ where } M \text{ set } (\ell_j = N_j)_{j \in J} : () ! \{\text{write}\} \cup E}$			$\frac{\Gamma \vdash M : \text{Table}(A) ! E \quad \Gamma, x : A \vdash N : \text{Bool} ! \emptyset}{\Gamma \vdash \text{delete } (x \leftarrow M) \text{ where } N : () ! \{\text{write}\} \cup E}$	

Figure 12. Typing rules for λ_{LINQ}

Big-step reduction rules

				$M \Downarrow_{\Delta, t} (V, \Delta')$
E-VAL	E-APP	E-OP	E-IFΓ	
$\frac{}{V \Downarrow_{\Delta, t} (V, \Delta)}$	$\frac{M \Downarrow_{\Delta, t} (\lambda x. L, \Delta_1) \quad N \Downarrow_{\Delta_1, t} (V, \Delta_2) \quad L\{V/x\} \Downarrow_{\Delta_2, t} (W, \Delta_3)}{MN \Downarrow_{\Delta, t} (W, \Delta_3)}$	$\frac{M_1 \Downarrow_{\Delta, t} (V_1, \Delta_1) \quad \dots \quad M_n \Downarrow_{\Delta_{n-1}, t} (V_n, \Delta_n)}{\odot \{\vec{M}\} \Downarrow_{\Delta, t} (\odot \{\vec{V}\}, \Delta_n)}$	$\frac{L \Downarrow_{\Delta, t} (\text{true}, \Delta_1) \quad M \Downarrow_{\Delta_1, t} (V, \Delta_2)}{\text{if } L \text{ then } M \text{ else } N \Downarrow_{\Delta, t} (V, \Delta_2)}$	
E-IFF	E-BAG	E-BAGUNION	E-RECORD	
$\frac{L \Downarrow_{\Delta, t} (\text{false}, \Delta_1) \quad N \Downarrow_{\Delta_1, t} (V, \Delta_2)}{\text{if } L \text{ then } M \text{ else } N \Downarrow_{\Delta, t} (V, \Delta_2)}$	$\frac{M \Downarrow_{\Delta, t} (V, \Delta')}{\{M\} \Downarrow_{\Delta, t} (\{V\}, \Delta')}$	$\frac{M \Downarrow_{\Delta, t} (V, \Delta_1) \quad N \Downarrow_{\Delta_1, t} (W, \Delta_2)}{M \uplus N \Downarrow_{\Delta, t} (V \uplus W, \Delta_2)}$	$\frac{M_1 \Downarrow_{\Delta, t} (V_1, \Delta_1) \quad \dots \quad M_n \Downarrow_{\Delta_{n-1}, t} (V_n, \Delta_n)}{(\ell = \vec{M}) \Downarrow_{\Delta, t} ((\ell = \vec{V}), \Delta_n)}$	
E-PROJECT	E-NOW	E-QUERY	E-FOREMPTY	
$\frac{M \Downarrow_{\Delta, t} ((\ell_i = V_i)_{i \in I}, \Delta') \quad j \in I}{M. \ell_j \Downarrow_{\Delta, t} (V_j, \Delta')}$	$\frac{}{\text{now} \Downarrow_{\Delta, t} (t, \Delta)}$	$\frac{M \Downarrow_{\Delta, t}^* V}{\text{query } M \Downarrow_{\Delta, t} (V, \Delta)}$	$\frac{M \Downarrow_{\Delta, t} (\{\ \}, \Delta')}{\text{for } (x \leftarrow M) N \Downarrow_{\Delta, t} (\{\ \}, \Delta')}$	
E-FOR	E-GET	E-INSERT		
$\frac{M \Downarrow_{\Delta, t} (\{\vec{V}_1 \cdot V \cdot \vec{V}_2\}, \Delta_1) \quad N\{V/x\} \Downarrow_{\Delta_1, t} (W_1, \Delta_2) \quad \text{for } (x \leftarrow \{\vec{V}_1 \cdot \vec{V}_2\}) M \Downarrow_{\Delta_2, t} (W_2, \Delta_3)}{\text{for } (x \leftarrow M) N \Downarrow_{\Delta, t} (W_1 \uplus W_2, \Delta_3)}$	$\frac{M \Downarrow_{\Delta, t} (t, \Delta')}{\text{get } M \Downarrow_{\Delta, t} (\Delta' (t), \Delta')}$	$\frac{M \Downarrow_{\Delta, t} (t, \Delta') \quad N \Downarrow_{\Delta, t}^* V}{\text{insert } M \text{ values } N \Downarrow_{\Delta, t} ((), \Delta' [t \mapsto \Delta' (t) \uplus V])}$		
E-UPDATE	E-DELETE			
$\frac{L \Downarrow_{\Delta, t} (t, \Delta_1) \quad \Delta_2 = \Delta_1 [t \mapsto \{\text{upd}(v) \mid v \in \Delta_1(t)\}] \quad \text{upd}(v) = \begin{cases} (v \text{ with } \ell = \vec{W}) & \text{if } M\{v/x\} \Downarrow_{\Delta_1, t}^* \text{true and } (N_i\{v/x\} \Downarrow_{\Delta_1, t}^* W_i)_i \\ v & \text{if } M\{v/x\} \Downarrow_{\Delta_1, t}^* \text{false} \end{cases}}{\text{update } (x \leftarrow L) \text{ where } M \text{ set } (\ell = \vec{N}) \Downarrow_{\Delta, t} ((), \Delta_2)}$	$\frac{M \Downarrow_{\Delta, t} (t, \Delta_1) \quad \Delta_2 = \Delta_1 [t \mapsto \{\ v \in \Delta(t) \mid N\{v/x\} \Downarrow_{\Delta_1, t}^* \text{false} \}]}{\text{delete } (x \leftarrow M) \text{ where } N \Downarrow_{\Delta, t} ((), \Delta_2)}$			

Figure 13. Semantics of λ_{LINQ}

Typing rules

$\frac{\text{TV-GET} \quad \Gamma \vdash M:\text{Table}(A)!E}{\Gamma \vdash \mathbf{get} \ M:\text{Bag}(\text{ValidTime}(A))! \{\mathbf{read}\} \cup E}$	$\frac{\text{TV-SEQINSERT} \quad \Gamma \vdash M:\text{Table}(A)!E \quad \Gamma \vdash M:\text{Bag}(\text{ValidTime}(A))! \emptyset}{\Gamma \vdash \mathbf{insert \ sequenced} \ M \ \mathbf{values} \ N:()! \{\mathbf{write}\} \cup E} \quad \boxed{\Gamma \vdash M : A! E}$
$\text{TV-NONSEQUPDATE} \quad \frac{\Gamma \vdash L:\text{Table}(A)!E \quad A = (\ell_i : B_i)_{i \in I} \quad \Gamma, x : \text{ValidTime}(A) \vdash M:\text{Bool}! \emptyset \quad (j \in I \wedge \Gamma, x : \text{ValidTime}(A) \vdash N_j:B_j! \emptyset)_{j \in J} \quad \Gamma, x : \text{ValidTime}(A) \vdash N'_1:\text{Time}! \emptyset \quad \Gamma, x : \text{ValidTime}(A) \vdash N'_2:\text{Time}! \emptyset}{\Gamma \vdash \mathbf{update \ nonsequenced} \ (x \leftarrow L) \ \mathbf{where} \ M \ \mathbf{set} \ (\ell_j=N_j)_{j \in J} \ \mathbf{valid \ from} \ N'_1 \ \mathbf{to} \ N'_2:()! \{\mathbf{write}\} \cup E}$	
$\text{TV-SEQUPDATE} \quad \frac{\Gamma \vdash L:\text{Table}(A)!E \quad A = (\ell_i : B_i)_{i \in I} \quad \Gamma \vdash M_1:\text{Time}! \emptyset \quad \Gamma \vdash M_2:\text{Time}! \emptyset \quad \Gamma, x : A \vdash M_3:\text{Bool}! \emptyset \quad (j \in I \wedge \Gamma, x : A \vdash N_j:B_j! \emptyset)_{j \in J}}{\Gamma \vdash \mathbf{update \ sequenced} \ (x \leftarrow L) \ \mathbf{between} \ M_1 \ \mathbf{and} \ M_2 \ \mathbf{where} \ M_3 \ \mathbf{set} \ (\ell_j = N_j)_{j \in J}:()! \{\mathbf{write}\} \cup E}$	
$\frac{\text{TV-NONSEQDELETE} \quad \Gamma \vdash M:\text{Table}(A)!E \quad \Gamma, x : \text{ValidTime}(A) \vdash N:\text{Bool}! \emptyset}{\Gamma \vdash \mathbf{delete \ nonsequenced} \ (x \leftarrow M) \ \mathbf{where} \ N:()! \{\mathbf{write}\} \cup E}$	$\frac{\text{TV-SEQDELETE} \quad \Gamma \vdash L:\text{Table}(A)!E \quad \Gamma \vdash M_1:\text{Time}! \emptyset \quad \Gamma \vdash M_2:\text{Time}! \emptyset \quad \Gamma, x:A \vdash N:\text{Bool}! \emptyset}{\Gamma \vdash \mathbf{delete \ sequenced} \ (x \leftarrow L) \ \mathbf{between} \ M_1 \ \mathbf{and} \ M_2 \ \mathbf{where} \ N:()! \{\mathbf{write}\} \cup E}$

Figure 14. Typing rules for λ_{VLINQ}

Reduction rules

$$M \Downarrow_{\Delta, t}^V (V, \Delta')$$

<p>EV-Row</p> $\frac{M_1 \Downarrow_{\Delta, t}^V (V_1, \Delta_1) \quad M_2 \Downarrow_{\Delta, t}^V (V_2, \Delta_2) \quad M_3 \Downarrow_{\Delta, t}^V (V_3, \Delta_3)}{M_1^{[M_2, M_3]} \Downarrow_{\Delta, t}^V (V_1^{[V_2, V_3]}, \Delta_3)}$	<p>EV-SEQINSERT</p> $\frac{M \Downarrow_{\Delta, t}^V (t, \Delta_1) \quad N \Downarrow_t^* \wr \tilde{V} \wr \quad \forall data^{[start, end]} \in \tilde{V}. start < end \quad \Delta_2 = \Delta_1 [t \mapsto \Delta_1(t) \hat{\cup} \wr \tilde{V} \wr]}{\text{insert sequenced } M \text{ values } N \Downarrow_{\Delta, t}^V ((), \Delta_2)}$	<p>EV-NONSEQDELETE</p> $\frac{M \Downarrow_{\Delta, t} (t, \Delta_1) \quad \Delta_2 = \Delta_1 [t \mapsto \wr d \in \Delta(t) \mid N \{d/x\} \Downarrow_t^* \text{false} \wr]}{\text{delete nonsequenced } (x \Leftarrow M) \text{ where } N \Downarrow_{\Delta, t} ((), \Delta_2)}$
<p>EV-SEQDELETE</p> $\frac{L \Downarrow_{\Delta, t} (t, \Delta_1) \quad M_1 \Downarrow_t^* V_{start} \quad M_2 \Downarrow_t^* V_{end} \quad V_{start} < V_{end} \quad \Delta_2 = \Delta_1 [t \mapsto \hat{\cup} \wr \wr \text{del}(d) \mid d \in \Delta(t) \wr]}{\text{delete sequenced } (x \Leftarrow L) \text{ between } M_1 \text{ and } M_2 \text{ where } N \Downarrow_{\Delta, t}^V ((), \Delta_2)}$ $\text{del}(v^{[start, end]}) \triangleq \begin{cases} \wr \wr & \text{if } N \{v/x\} \Downarrow_t^* \text{true and } V_{start} \leq start \text{ and } V_{end} \geq end \quad \text{(Case 1)} \\ \wr v^{[V_{end}, end]} \wr & \text{if } N \{v/x\} \Downarrow_t^* \text{true and } V_{start} \leq start \text{ and } V_{end} < end \quad \text{(Case 2)} \\ \wr v^{[start, V_{start}]} \wr, v^{[V_{end}, end]} \wr & \text{if } N \{v/x\} \Downarrow_t^* \text{true and } V_{start} > start \text{ and } V_{end} < end \quad \text{(Case 3)} \\ \wr v^{[start, V_{start}]} \wr & \text{if } N \{v/x\} \Downarrow_t^* \text{true and } V_{start} > start \text{ and } V_{end} \geq end \quad \text{(Case 4)} \\ \wr v^{[start, end]} \wr & \text{otherwise} \quad \text{(Case 5)} \end{cases}$		
<p>EV-NONSEQUPDATE</p> $\frac{L \Downarrow_{\Delta, t} (t, \Delta_1) \quad \Delta_2 = \Delta_1 [t \mapsto \wr \text{upd}(d) \mid d \in \Delta_1(t) \wr]}{\text{upd}(D = v^{[start, end]}) = \begin{cases} (v \text{ with } \ell = \widehat{W})^{[W_{start}, W_{end}]} & \text{if } M \{D/x\} \Downarrow_t^* \text{true and } (N_i \{D/x\} \Downarrow_t^* W_i)_i \text{ and} \\ & N'_1 \{D/x\} \Downarrow_t^* W_{start} \text{ and } N'_2 \{D/x\} \Downarrow_t^* W_{end} \text{ and } W_{start} < W_{end} \\ D & \text{if } M \{D/x\} \Downarrow_t^* \text{false} \end{cases}}$ <p style="text-align: center;">update nonsequenced $(x \Leftarrow L)$ where M set $(\ell_i = N_i)_i$ valid from N'_1 to $N'_2 \Downarrow_{\Delta, t} ((), \Delta_2)$</p>		
<p>EV-SEQUPDATE</p> $\frac{L \Downarrow_{\Delta, t}^V (t, \Delta_1) \quad M_1 \Downarrow_t^* V_{start} \quad M_2 \Downarrow_t^* V_{end} \quad V_{start} < V_{end} \quad \Delta_2 = \Delta_1 [t \mapsto \hat{\cup} \wr \wr \text{upd}(d) \mid d \in \Delta_1(t) \wr]}{\text{upd}(v^{[start, end]}) = \begin{cases} \wr W^{[start, end]} \wr & \text{if } M_3 \{v/x\} \Downarrow_t^* \text{true and } V_{start} \leq start \text{ and } V_{end} \geq end \quad \text{(Case 1)} \\ \wr W^{[start, V_{end}]} \wr, v^{[V_{end}, end]} \wr & \text{if } M_3 \{v/x\} \Downarrow_t^* \text{true and } V_{start} \leq start \text{ and } V_{end} < end \quad \text{(Case 2)} \\ \wr v^{[start, V_{start}]} \wr, W^{[V_{start}, V_{end}]} \wr, v^{[V_{end}, end]} \wr & \text{if } M_3 \{v/x\} \Downarrow_t^* \text{true and } V_{start} > start \text{ and } V_{end} < end \quad \text{(Case 3)} \\ \wr v^{[start, V_{start}]} \wr, W^{[V_{start}, end]} \wr & \text{if } M_3 \{v/x\} \Downarrow_t^* \text{true and } V_{start} > start \text{ and } V_{end} \geq end \quad \text{(Case 4)} \\ \wr v^{[start, end]} \wr & \text{otherwise} \quad \text{(Case 5)} \end{cases}$ <p style="text-align: center;">where for all cases, $W = (v \text{ with } \ell = \widehat{W})$ given $(N_i \Downarrow_t^* W'_i)_i$</p> <p style="text-align: center;">update sequenced $(x \Leftarrow L)$ between M_1 and M_2 where M_3 set $(\ell_i = N_i)_i \Downarrow_{\Delta, t} ((), \Delta_2)$</p>		

Figure 15. Reduction rules for λ_{VLIQ}

```

(delete( $\ell_i:A_i$ ) $i$  nonsequenced ( $x \leftarrow M$ ) where  $N$ ) =
  delete ( $x \leftarrow \langle M \rangle$ ) where  $\text{lift}(x, \langle N \rangle)$ 
where  $\text{lift}(x, f) \triangleq$ 
  ( $\lambda x.f$ ) ( $\text{data} = \eta(x, \{\ell_i\}_i)$ ,  $\text{start} = x.\text{start}$ ,  $\text{end} = x.\text{end}$ )

(update( $\ell_i:A_i$ ) $i \in I$  nonsequenced ( $x \leftarrow L$ ) where  $M$  set ( $\ell_j = N_j$ ) $_{j \in J}$  valid from  $N'_1$  to  $N'_2$ ) =
  update( $\ell_i:A_i$ ) $i \in I$  ( $x \leftarrow \langle L \rangle$ )
  where ( $\text{lift}(x, \langle M \rangle)$ )
  set ( $(\ell_j = \text{lift}(x, \langle N_j \rangle))_{j \in J}$ ,  $\text{start} = \text{lift}(x, \langle N'_1 \rangle)$ ,  $\text{end} = \text{lift}(x, \langle N'_2 \rangle)$ )
  where  $\text{lift}(x, f) \triangleq$ 
  ( $\lambda x.f$ ) ( $\text{data} = \eta(x, \{\ell_i\}_{i \in I})$ ,  $\text{start} = x.\text{start}$ ,  $\text{end} = x.\text{end}$ )

(insert( $\ell_i:A_i$ ) $i$  sequenced  $M$  values  $N$ ) =
  let  $\text{tbl} = \langle M \rangle$  in
  let  $\text{rows} =$ 
    for ( $x \leftarrow \langle N \rangle$ )
      [ $\eta(x, \tilde{\ell}) \oplus (\text{start} = x.\text{start}, \text{end} = x.\text{end})$ ]
  in
  insert  $\text{tbl}$  values  $\text{rows}$ 

(delete( $\ell_i:A_i$ ) $i$  sequenced ( $x \leftarrow L$ ) between  $M_1$  and  $M_2$  where  $N$ ) =
  let  $\text{tbl} = \langle L \rangle$  in
  let  $\text{aStart} = \langle M_1 \rangle$  in
  let  $\text{aEnd} = \langle M_2 \rangle$  in
  let  $\text{lRows} = \text{startRows}(\text{tbl}, \text{pred}, \text{aStart}, \tilde{\ell})$  in
  let  $\text{rRows} = \text{endRows}(\text{tbl}, \text{pred}, \text{aEnd}, \tilde{\ell})$  in
  delete ( $x \leftarrow \text{tbl}$ )
    where ( $\text{pred} \wedge (x.\text{start} < \text{aEnd}) \wedge (x.\text{end} > \text{aStart})$ )
  insert  $\text{tbl}$  values  $\text{lRows}$ ;
  insert  $\text{tbl}$  values  $\text{rRows}$ 
  where  $\text{pred} \triangleq \text{restrict}(x, \tilde{\ell}, \langle N \rangle)$ 

(update( $\ell_i:A_i$ ) $i \in I$  sequenced ( $x \leftarrow L$ ) between  $M_1$  and  $M_2$  where  $M_3$  set ( $\ell_j = N_j$ ) $_{j \in J}$ ) =
  let  $\text{tbl} = \langle L \rangle$  in
  let  $\text{aStart} = \langle M_1 \rangle$  in
  let  $\text{aEnd} = \langle M_2 \rangle$  in
  let  $\text{lRows} = \text{startRows}(\text{tbl}, \text{pred}, \text{aStart}, \{\ell_i\}_{i \in I})$  in
  let  $\text{rRows} = \text{endRows}(\text{tbl}, \text{pred}, \text{aEnd}, \{\ell_i\}_{i \in I})$  in
  update ( $x \leftarrow \text{tbl}$ )
    where ( $\text{pred} \wedge (x.\text{start} < \text{aEnd}) \wedge (x.\text{end} > \text{aStart})$ )
    set ( $(\ell_j = \text{restrict}(x, \{\ell_i\}_{i \in I}, \langle N_j \rangle))_{j \in J}$ ,  $\text{start} = \mathbf{greatest}(x.\text{start}, \text{aStart})$ ,  $\text{end} = \mathbf{least}(x.\text{end}, \text{aEnd})$ );
  insert  $\text{tbl}$  values  $\text{lRows}$ ;
  insert  $\text{tbl}$  values  $\text{rRows}$ 
  where  $\text{pred} \triangleq \text{restrict}(x, \{\ell_i\}_{i \in I}, \langle M_3 \rangle)$ 

 $\text{startRows}(\text{tbl}, \text{pred}, \text{aStart}, \tilde{\ell}) \triangleq$ 
  query
  for ( $x \leftarrow \mathbf{get} \text{tbl}$ )
    where ( $\text{pred} \wedge (x.\text{start} < \text{aStart}) \wedge (x.\text{end} > \text{aStart})$ )
    [ $\eta(x, \tilde{\ell}) \oplus (\text{start} = x.\text{start}, \text{end} = \text{aStart})$ ]

 $\text{endRows}(\text{tbl}, \text{pred}, \text{aEnd}, \tilde{\ell}) \triangleq$ 
  query
  for ( $x \leftarrow \mathbf{get} \text{tbl}$ )
    where ( $\text{pred} \wedge (x.\text{start} < \text{aEnd}) \wedge (x.\text{end} > \text{aEnd})$ )
    [ $\eta(x, \tilde{\ell}) \oplus (\text{start} = \text{aEnd}, \text{end} = x.\text{end})$ ]

```

Figure 16. Translation from λ_{VLINQ} into λ_{LINQ}

B Illustration of sequenced delete and update behavior

Figure 17 illustrates graphically the five different overlap relationships that can hold between the period of validity of an existing row to be updated or deleted (PV) and the period of applicability of a deletion or update. The five cases involve when one interval is totally contained in the other (1, 3), when there is overlap but neither containment relationship holds (2,4) and when the intervals are disjoint (5). In each case the effect of a deletion or update is described in the corresponding column of the table; generally the result is to replace the input row with zero, one, two or (exceptionally) three new rows. When the PV and PA intervals are disjoint no action needs to be taken.

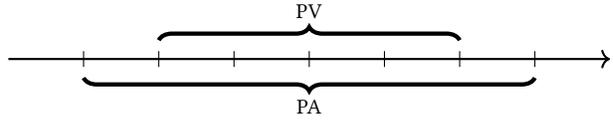
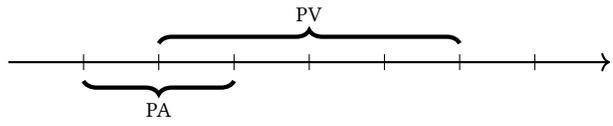
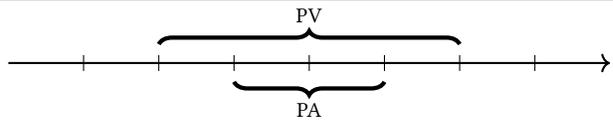
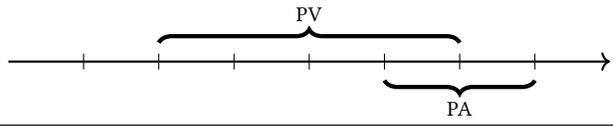
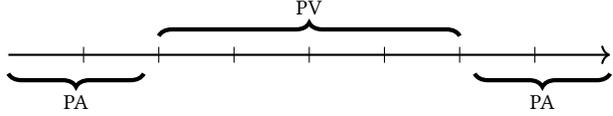
Case	Diagram	Delete behavior	Update behavior
1		the entire row will be deleted	the entire row will be updated and adjusted to cover PA
2		the overlapping portion will be deleted, leaving the rest of PV alone	insert new row with updated values covering PA and shorten existing row to only cover remainder of PV
3		split row into two, one covering the part of PV before PA and the other covering the part of PV after PA	split row into three, two as in the case of deletion and a third providing the updated values for PA
4		symmetric with case 2	symmetric with case 2
5		no effect	no effect

Figure 17. Timeline diagrams illustrating sequenced delete and update behavior

C Direct semantics and translations for λ_{VLINQ} current updates and deletions

In Section 4, we showed that current updates and deletions in λ_{VLINQ} are special cases of sequenced updates and deletions. For the sake of completeness, we include the direct semantics and translations here.

EV-UPDATE

$$\text{upd}(data^{[start, end]}) = \frac{\begin{array}{l} L \Downarrow_{\Delta, t}^V(t, \Delta') \quad \Delta'' = \Delta' [t \mapsto \hat{\oplus} \uparrow \uparrow \text{upd}(v) \mid v \in \Delta'(t) \downarrow] \\ \left\{ \begin{array}{l} \uparrow (data \text{ with } \widehat{\ell = V})^{[start, end]} \downarrow \quad \text{if } M\{data/x\} \Downarrow_t^* \text{ true, } t \leq start, \text{ and } (N_i\{data/x\} \Downarrow_t^* V_i)_i, \\ \uparrow data^{[start, t]} \downarrow, (data \text{ with } \widehat{\ell = V})^{[t, end]} \downarrow \quad \text{if } M\{data/x\} \Downarrow_t^* \text{ true, } start < t < end, \text{ and } (N_i\{data/x\} \Downarrow_t^* V_i)_i, \\ \uparrow data^{[start, end]} \downarrow \quad \text{if } M\{data/x\} \Downarrow_t^* \text{ false or } t \geq end \end{array} \right. \end{array}}{\text{update } (x \leftarrow L) \text{ where } M \text{ set } (\widehat{\ell = N}) \Downarrow_{\Delta, t}^V(\Delta'', ())}$$

EV-DELETE

$$\text{del}(data^{[start, end]}) = \frac{\begin{array}{l} M \Downarrow_{\Delta, t}^V(t, \Delta') \quad \Delta'' = \Delta' [t \mapsto \hat{\oplus} \uparrow \uparrow \text{del}(v) \mid v \in \Delta'(t) \downarrow] \\ \left\{ \begin{array}{l} \uparrow \downarrow \quad \text{if } N\{data/x\} \Downarrow_t^* \text{ true and } t \leq start \\ \uparrow data^{[start, t]} \downarrow \quad \text{if } N\{data/x\} \Downarrow_t^* \text{ true and } start < t < end \\ \uparrow data^{[start, end]} \downarrow \quad \text{if } N\{data/x\} \Downarrow_t^* \text{ false or } t \geq end \end{array} \right. \end{array}}{\text{delete } (x \leftarrow M) \text{ where } N \Downarrow_{\Delta, t}^V(\Delta'', ())}$$

Rules EV-DELETE and EV-UPDATE are different to their λ_{TLINQ} counterparts, taking into account that a row may be valid further into the future than the current time.

Specifically, rule EV-DELETE has three cases: if the row being inspected is entirely in the future (i.e., its *start* field is greater than or equal to the current time), then it is deleted from the database. If the current time is greater than the start of the row, but less than the end of the row, then the row is closed off at the current time. The row is not modified if the current time is greater than the end time of the row, or the predicate evaluates to **false**. Current updates, described by EV-UPDATE, follow much the same pattern, but if the current time is between the start and end times of the row, then the row is split: the part up until the current time retains the previous values, and the part from the current time until the end of the row is set to the new values.

$$\text{currentAt}(x, time) \triangleq time \geq x.start \wedge time < x.end$$

$$\langle \text{delete}^{(\ell_i: A_i)_i} (x \leftarrow M) \text{ where } N \rangle =$$

```

let tbl = ⟨M⟩ in
let time = now in
update (x ← tbl)
  where (restrict(x, {ℓi }i, ⟨N⟩) ∧ x.start ≤ time ∧ x.end > time)
set (end = time);
delete (x ← tbl)
  where (restrict(x, {ℓi }i, ⟨N⟩) ∧ x.start ≥ time)

```

$$\langle \text{update}^{(\ell_i: A_i)_{i \in I}} (x \leftarrow L) \text{ where } M \text{ set } (\ell_j = N_j)_{j \in J} \rangle =$$

```

let tbl = ⟨L⟩ in
let time = now in
let affected =
  query
    for (x ← tbl)
      where (restrict(x, {ℓi }i \in I, ⟨M⟩) ∧ currentAt(x, time))
        (ℓi = x.ℓi)i \in I \setminus J ⊕
        [ (ℓj = restrict(x, {ℓi }i \in I, ⟨Nj⟩))j \in J ⊕ ]
        (start = time, end = x.end)
in
update (x ← tbl)
  where ( restrict(x, {ℓi }i \in I, ⟨M⟩) ∧ (x.start < time) ∧
        (x.end > time) )
set (end = time);
update (x ← tbl)
  where (restrict(x, {ℓi }i \in I, ⟨M⟩) ∧ x.start ≥ time)
set ((ℓj = restrict(x, {ℓi }i \in I, ⟨Nj⟩))j \in J);
insert tbl values affected

```

The meta-level *currentAt* function checks whether a row is valid at a particular timestamp.

Current deletions are implemented using two λ_{LINQ} operations: an **update** operation, which adjusts the *end* time of all matching, currently-valid rows to the current time; a **delete** operation, which entirely deletes all matching rows which begin after the current time.

To implement current updates, we firstly calculate the affected bag of rows (i.e., those rows which match the predicate and are valid at the current time), with their updated values. We then do three operations:

- An **update**, which adjusts the *end* time of all records which match the predicate, start before the current time, and end after the current time
- A second **update**, which updates all rows matching the predicate which both start and end after the current timestamp
- Finally, an **insert**, which inserts the contents of the *affected* bag.

D Proofs

Lemma D.1 (Pure comprehensions). If $V = \{V_1, \dots, V_n\}$ and we have a for comprehension **for** $(x \leftarrow V) M$ such that $\cdot \vdash \mathbf{for} (x \leftarrow V) M:A! \emptyset$, and $M\{W_i/x\} \Downarrow_i^* W'_i$ or $M\{W_i/x\} \Downarrow_{\Delta,i}^* W'_i$ for each $i \in 1..n$, then $\mathbf{for} (x \leftarrow V) M \Downarrow_i^* \{W'_1, \dots, W'_n\}$.

Proof. Follows from the comprehension evaluation rules. \square

Lemma D.2 (Lifting of pure reduction). If $\cdot \vdash M:A! \emptyset$ and $M \Downarrow_i^* V$ or $M \Downarrow_{\Delta,i}^* V$, then $M \Downarrow_{\Delta,i} (V, \Delta)$ for any Δ .

Proof. By induction on the derivation of $M \Downarrow_i^* V$. \square

Lemma D.3 (Correctness of λ_{TLINQ} translation (pure reduction)). If $M \Downarrow_i^* V$, then $\llbracket M \rrbracket \Downarrow_i^* \llbracket V \rrbracket$.

Proof. By induction on the derivation of $M \Downarrow_i^* V$. \square

Lemma D.4. If $\Gamma, x : (\ell_i : A_i)_{i \in I} \vdash M:B! \emptyset$ and $M \Downarrow_i^* V$, and $\llbracket \Gamma \rrbracket, x : ((\ell_i : \llbracket A_i \rrbracket)_{i \in I}) \oplus (\ell_j : A'_j)_{j \in J} \vdash \llbracket M \rrbracket : \llbracket A \rrbracket ! \emptyset$, then $\text{restrict}(x, (\ell_i)_{i \in I}, \llbracket M \rrbracket) \Downarrow_{\downarrow \Delta', i}^* \llbracket V \rrbracket$.

Proof. Follows from the definition of restrict and Lemma D.3. \square

Lemma D.5 (Correctness of λ_{VLINQ} translation (pure reduction)). If $M \Downarrow_i^* V$, then $\langle M \rangle \Downarrow_i^* \langle V \rangle$.

Proof. By induction on the derivation of $M \Downarrow_i^* V$. \square

Lemma D.6. If $\Gamma, x : (\ell_i : A_i)_{i \in I} \vdash M:B! \emptyset$ and $M \Downarrow_i^* V$, and $\langle \Gamma \rangle, x : ((\ell_i : \langle A_i \rangle)_{i \in I}) \oplus (\ell_j : A'_j)_{j \in J} \vdash \langle M \rangle : \langle A \rangle ! \emptyset$, then $\text{restrict}(x, (\ell_i)_{i \in I}, \langle M \rangle) \Downarrow_{\downarrow \Delta', i}^* \langle V \rangle$.

Proof. Follows from the definition of restrict and Lemma D.5. \square

Lemma D.7. If $\Gamma \vdash V:A! \emptyset$ where $A = C$ or $A = (\ell_i = C_i)_i$, then $\langle V \rangle = V$.

Proof. By induction on the derivation of $\Gamma \vdash V:A! \emptyset$. \square

Proposition D.1 (Type correctness of λ_{TLINQ} translation). If $\Gamma \vdash M:A! E$ in λ_{TLINQ} , then $\llbracket \Gamma \rrbracket \vdash \llbracket M \rrbracket : \llbracket A \rrbracket ! E$ in λ_{LINQ} .

Proof. By induction on the derivation of $\Gamma \vdash M:A! E$. \square

Proposition D.2 (Type correctness of λ_{VLINQ} translation). If $\Gamma \vdash M:A! E$ in λ_{VLINQ} , then $\langle \Gamma \rangle \vdash \langle M \rangle : \langle A \rangle ! E$ in λ_{LINQ} .

Proof. By induction on the derivation of $\Gamma \vdash M:A! E$. \square

Theorem 3.1. If $\cdot \vdash M:A! E$ and $M \Downarrow_{\Delta,i}^T (V, \Delta')$ where $\text{wf}(\Delta)$ and $\max(\Delta) \leq i$, then $\llbracket M \rrbracket \Downarrow_{\downarrow \Delta, i} (\llbracket V \rrbracket, \downarrow \Delta')$

Proof. By induction on the derivation of $M \Downarrow_{\Delta,i}^T V$. We show the cases for the database-relevant terms.

Case Accessor functions

Let us show the case for **data**; the others are similar.

Assumption:

$$\frac{M \Downarrow_{\Delta,i}^T (V_1^{[V_2, V_3]}, \Delta')}{\mathbf{data} M \Downarrow_{\Delta,i}^T (V_1, \Delta')}$$

Translation:

$$\llbracket M \rrbracket . \mathbf{data}$$

By the IH: $\llbracket M \rrbracket \Downarrow_{\downarrow \Delta, i} (\mathbf{data} = \llbracket V_1 \rrbracket, \mathbf{start} = \llbracket V_2 \rrbracket, \mathbf{end} = \llbracket V_3 \rrbracket)$

Evaluating in λ_{LINQ} :

$$(\mathbf{data} = \llbracket V_1 \rrbracket, \mathbf{start} = \llbracket V_2 \rrbracket, \mathbf{end} = \llbracket V_3 \rrbracket) . \mathbf{data} \Downarrow_{\downarrow \Delta, i} (\llbracket V_1 \rrbracket, \downarrow \Delta')$$

as required.

Case get

Assumption:

$$\frac{M \Downarrow_{\Delta, t}^{\top} (t, \Delta')}{\mathbf{get} M \Downarrow_{\Delta, t}^{\top} (\Delta'(t), \Delta')}$$

Translation:

for ($x \leftarrow \mathbf{get} \llbracket M \rrbracket$)
 $\lambda (data = (\ell_i = x.\ell_i)_i, start = x.start, end = x.end) \int$

Suppose $\Delta(t) = \lambda (\ell_i = V_{1_i})_{i \in 1..m}^{[W_{1_s}, W_{1_e}]}, \dots, (\ell_i = V_{n_i})_{i \in 1..m}^{[W_{n_s}, W_{n_e}]} \int$.

By the IH, $\llbracket M \rrbracket \Downarrow_{\Delta, t} (t, \downarrow \Delta')$

Then,

$$\begin{aligned} \downarrow \Delta'(t) = & \lambda (\ell_1 = \llbracket V_{1_1} \rrbracket, \dots, \ell_m = \llbracket V_{1_m} \rrbracket, start = \llbracket W_{1_s} \rrbracket, end = \llbracket W_{1_e} \rrbracket), \\ & \dots, \\ & (\ell_1 = \llbracket V_{n_1} \rrbracket, \dots, \ell_m = \llbracket V_{n_m} \rrbracket, start = \llbracket W_{n_s} \rrbracket, end = \llbracket W_{n_e} \rrbracket) \int \end{aligned}$$

In λ_{LINQ} , we have:

get $t \downarrow_{\Delta', t} \lambda (\ell_1 = \llbracket V_{1_1} \rrbracket, \dots, \ell_m = \llbracket V_{1_m} \rrbracket, start = \llbracket W_{1_s} \rrbracket, end = \llbracket W_{1_e} \rrbracket), \dots, (\ell_1 = \llbracket V_{n_1} \rrbracket, \dots, \ell_m = \llbracket V_{n_m} \rrbracket, start = \llbracket W_{n_s} \rrbracket, end = \llbracket W_{n_e} \rrbracket) \int$

By Lemmas D.1 and D.2

for ($x \leftarrow \mathbf{get} \llbracket t \rrbracket$)
 $\lambda (data = (\ell_i = x.\ell_i)_i, start = x.start, end = x.end) \int \Downarrow_{\Delta, t} (\lambda (data = (\ell_1 = \llbracket V_{1_1} \rrbracket, \dots, \ell_m = \llbracket V_{1_m} \rrbracket), start = \llbracket W_{1_s} \rrbracket, end = \llbracket W_{1_e} \rrbracket),$
 $\dots,$
 $(data = (\ell_1 = \llbracket V_{n_1} \rrbracket, \dots, \ell_m = \llbracket V_{n_m} \rrbracket), start = \llbracket W_{n_s} \rrbracket, end = \llbracket W_{n_e} \rrbracket) \int, \downarrow \Delta')$

as required.

Case insert

$$\frac{M \Downarrow_{\Delta, t}^{\top} (t, \Delta') \quad N \Downarrow_t^{\star} V \quad vs = \lambda (v^{[t, \infty)} \mid v \in V) \quad \Delta'' = \Delta' [t \mapsto \Delta'(t) \uplus vs]}{\mathbf{insert} M \mathbf{values} N \Downarrow_{\Delta, t}^{\top} ((, \Delta''))}$$

Translation:

$\llbracket \mathbf{insert}^{(\ell_i: A_i)_i} M \mathbf{values} N \rrbracket =$ **let** $rows =$
for ($x \leftarrow \llbracket N \rrbracket$)
 $\lambda (\eta(x, \bar{\ell}) \oplus (start = \mathbf{now}, end = \infty) \int)$
in
insert $\llbracket M \rrbracket$ **values** $rows$

By the typing rules and evaluation rule, V must be some bag $\lambda (\ell_{1_i} = W_{1_i})_i, \dots, (\ell_{n_i} = W_{n_i})_i \int$

By Lemma D.3 $\llbracket N \rrbracket \Downarrow_t^{\star} \llbracket V \rrbracket$ with $\llbracket V \rrbracket = \lambda (\ell_{1_i} = \llbracket W_{1_i} \rrbracket)_i, \dots, (\ell_{n_i} = \llbracket W_{n_i} \rrbracket)_i \int$.

By Lemmas D.1 and D.3, the **for** comprehension evaluates to $\llbracket V \rrbracket = \lambda (\ell_{1_i} = \llbracket W_{1_i} \rrbracket, start = \iota, end = \infty)_i, \dots, (\ell_{n_i} = \llbracket W_{n_i} \rrbracket, start = \iota, end = \infty)_i \int$. Let us call this value $rows$.

By the IH, $\llbracket M \rrbracket \Downarrow_{\Delta, t} (t, \downarrow \Delta')$.

Thus, we evaluate

insert $t \mathbf{values} rows \downarrow_{\Delta', t} ((, \downarrow \Delta' [t \mapsto \downarrow \Delta' \hat{\cup} rows])$

noting that $\llbracket \Delta'' \rrbracket = \downarrow \Delta' \hat{\cup} rows$, as required.

Case delete

Assumption:

$$\frac{M \Downarrow_{\Delta, t}^{\top} (t, \Delta') \quad \Delta'' = \Delta' [t \mapsto \lambda (del(d) \mid d \in \Delta'(t)) \int]}{del(data^{[start, end)}) = \begin{cases} data^{[start, t)} & \text{if } end = \infty \text{ and } N\{data/x\} \Downarrow_t^{\star} \mathbf{true} \\ data^{[start, end)} & \text{otherwise} \end{cases}}{\mathbf{delete} (x \leftarrow M) \mathbf{where} N \Downarrow_{\Delta, t}^{\top} ((, \Delta''))}$$

Translation:

$$\begin{aligned} \llbracket \text{delete } (x \leftarrow M) \text{ where } N \rrbracket = \\ \text{update } (x \leftarrow \llbracket M \rrbracket) \\ \text{where } (((\lambda x. \llbracket N \rrbracket) (\ell_i = x.\ell_i)_i) \wedge x.\text{end} = \infty) \\ \text{set } (\text{end} = \text{now}) \end{aligned}$$

By the IH, $\llbracket M \rrbracket \Downarrow_{\Delta, t} (t, \Downarrow \Delta')$

Suppose $\Delta'(t) = \wr (\ell_i = V_{1_i})_{i \in 1..m}^{[W_{1_s}, W_{1_e}]}, \dots, (\ell_i = V_{n_i})_{i \in 1..m}^{[W_{n_s}, W_{n_e}]} \wr$.

Then,

$$\begin{aligned} \Downarrow \Delta'(t) = \\ \wr (\ell_1 = \llbracket V_{1_1} \rrbracket, \dots, \ell_m = \llbracket V_{1_m} \rrbracket, \text{start} = \llbracket W_{1_s} \rrbracket, \text{end} = \llbracket W_{1_e} \rrbracket), \\ \dots, \\ (\ell_1 = \llbracket V_{n_1} \rrbracket, \dots, \ell_m = \llbracket V_{n_m} \rrbracket, \text{start} = \llbracket W_{n_s} \rrbracket, \text{end} = \llbracket W_{n_e} \rrbracket) \wr \end{aligned}$$

Recall the definition of **update** in λ_{LINQ} :

$$\begin{aligned} L \Downarrow_{\Delta, t} (t, \Delta_1) \quad \Delta_2 = \Delta_1 [t \mapsto \wr \text{upd}(v) \mid v \in \Delta_1(t) \wr] \\ \text{upd}(v) = \begin{cases} (v \text{ with } \ell = \widetilde{W}) & \text{if } M\{v/x\} \Downarrow_i^* \text{true and } (N_i\{v/x\} \Downarrow_i^* W_i)_i \\ v & \text{if } M\{v/x\} \Downarrow_i^* \text{false} \end{cases} \\ \hline \text{update } (x \leftarrow L) \text{ where } M \text{ set } (\ell = \widetilde{N}) \Downarrow_{\Delta, t} ((\cdot), \Delta_2) \end{aligned}$$

Let us reason on each row individually. Say we have a row $(\ell_i = V_i)^{[W_{start}, W_{end}]}$ in Δ' . Thus, in $\Downarrow \Delta'$ that row will be $(\ell_i = \llbracket V_i \rrbracket)_i \oplus (\text{start} = \llbracket W_{start} \rrbracket, \text{end} = \llbracket W_{end} \rrbracket)$; let us call this x .

Now note that $(\ell_i = x.\ell_i)_i \Downarrow_i^* (\ell_i = \llbracket V_i \rrbracket)_i$ which is equal to the translation of the data component of the record in the transaction time database.

Therefore by Lemma D.3, we have that $\llbracket N \rrbracket \{(\ell_i = \llbracket V_i \rrbracket)_i / x\} \Downarrow_i^* \text{true}$ whenever $N\{data/x\} \Downarrow_i^* \text{true}$. Since $\llbracket \infty \rrbracket = \infty$ we have that $\text{end} = \infty$ if $x.\text{end} = \infty$. Thus the **update** affects the same rows.

We have that **now** $\Downarrow_i^* \iota$ and thus $((\ell_i = \llbracket V_i \rrbracket)_i \oplus (\text{start} = \llbracket W_{start} \rrbracket, \text{end} = \llbracket W_{end} \rrbracket)) \text{ with } \text{end} = \iota = \downarrow (\ell_i = \llbracket V_i \rrbracket)_i^{[W_{start}, \iota]}$ for each affected record, as required.

Case update

Assumption:

$$\begin{aligned} L \Downarrow_{\Delta, t}^T (t, \Delta') \quad \Delta'' = \Delta' [t \mapsto \downarrow \wr u(v) \mid v \in \Delta'(t) \wr] \\ u(data^{[start, end]}) = \begin{cases} \wr data^{[start, \iota]}, (data \text{ with } \ell = \widetilde{V}')^{[t, \infty)} \wr & \text{if } M\{data/x\} \Downarrow_i^* \text{true,} \\ & \text{end} = \infty, \\ & (N_i\{data/x\} \Downarrow_i^* V'_i)_i \\ \wr data^{[start, end]} \wr & \text{otherwise} \end{cases} \\ \hline \text{update } (x \leftarrow L) \text{ where } M \text{ set } (\ell = \widetilde{N}) \Downarrow_{\Delta, t}^T ((\cdot), \Delta'') \end{aligned}$$

Translation:

$$\begin{aligned} \llbracket \text{update}^{(\ell_i: A_i)_{i \in I}} (x \leftarrow L) \text{ where } M \text{ set } (\ell = N_j)_{j \in J} \rrbracket = \\ \text{let } tbl = \llbracket L \rrbracket \text{ in} \\ \text{let } affected = \\ \text{for } (x \leftarrow \text{get } tbl) \\ \text{where } ((\text{restrict}(x, \widetilde{\ell}, \llbracket M \rrbracket) \wedge \text{isCurrent}(x))) \\ \wr \left(\begin{array}{l} data = (\ell_i = x.\ell_i)_{i \in I \setminus J} \oplus \\ (\ell_j = \text{restrict}(x, (\ell_i)_{i \in I}, \llbracket N_j \rrbracket)), \\ start = \text{now}, \text{end} = \infty \end{array} \right) \wr \\ \text{in} \\ \text{update } (x \leftarrow tbl) \\ \text{where } (\text{restrict}(x, \widetilde{\ell}, \llbracket M \rrbracket) \wedge \text{isCurrent}(x)) \\ \text{set } (\text{end} = \text{now}); \\ \text{insert } tbl \text{ values } affected \end{aligned}$$

By the IH, $\llbracket L \rrbracket \Downarrow_{\Delta, t} (t, \downarrow \Delta')$.

Suppose

$$\Delta(t') = \wr (\ell_i = V_{1_i})_{i \in 1..l}^{[W_{1_s}, W_{1_e})}, \dots, (\ell_i = V_{m_i})_{i \in 1..l}^{[W_{m_s}, W_{m_e})}, \\ (\ell_i = V_{m+1_i})_{i \in 1..l}^{[W_{m+1_s}, W_{m+1_e})}, \dots, (\ell_i = V_{n_i})_{i \in 1..l}^{[W_{n_s}, W_{n_e})} \wr$$

Then

$$\downarrow \Delta'(t) = \wr (\ell_i = \llbracket V_{1_i} \rrbracket)_{i \in 1..l} \oplus (start = \llbracket W_{1_s} \rrbracket, end = \llbracket W_{1_e} \rrbracket), \dots, (\ell_i = \llbracket V_{m_i} \rrbracket)_{i \in 1..l} \oplus (start = \llbracket W_{m_s} \rrbracket, end = \llbracket W_{m_e} \rrbracket), \\ (\ell_i = \llbracket V_{m+1_i} \rrbracket)_{i \in 1..l} \oplus (start = \llbracket W_{m+1_s} \rrbracket, end = \llbracket W_{m+1_e} \rrbracket), \dots, (\ell_i = \llbracket V_{n_i} \rrbracket)_{i \in 1..l} \oplus (start = \llbracket W_{n_s} \rrbracket, end = \llbracket W_{n_e} \rrbracket) \wr$$

Suppose that the guard returns **true** for records $1..m$, i.e., $M\{(\ell_i = V_{1_i})_{i \in 1..l}\} \Downarrow_i^* \mathbf{true}$, \dots , $M\{(\ell_i = V_{m_i})_{i \in 1..l}\} \Downarrow_i^* \mathbf{true}$ with each $W_{i_e} = \infty$.

Suppose that the guard returns false for the remainder.

By Proposition D.1 and Lemma D.4 we have that

$$\llbracket M \rrbracket \{(\ell_i = \llbracket V_{1_i} \rrbracket)_{i \in 1..l}\} \Downarrow_i^* \mathbf{true}, \dots, \llbracket M \rrbracket \{(\ell_i = \llbracket V_{m_i} \rrbracket)_{i \in 1..l}\} \Downarrow_i^* \mathbf{true}.$$

and $W_{1_e} = \infty, \dots, W_{m_e} = \infty$

Thus we have by Lemma D.1, Proposition D.1, and Lemma D.4 that

$$affected = \\ (\ell_i = \llbracket V'_{1_i} \rrbracket)_{i \in 1..l} \oplus (start = \iota, end = \infty), \dots, (\ell_i = \llbracket V'_{m_i} \rrbracket)_{i \in 1..l} \oplus (start = \iota, end = \infty)$$

Note that $affected = \downarrow \wr (\ell_i = V'_{1_i})_{i \in 1..l}^{[\iota, \infty)}, \dots, (\ell_i = V'_{m_i})_{i \in 1..l}^{[\iota, \infty)} \wr$.

By applying the **update** operation (with appropriate uses of Proposition D.1 and Lemma D.4), we have some Δ'' with table t containing:

$$\Delta'' = \wr (\ell_i = \llbracket V_{1_i} \rrbracket)_{i \in 1..l} \oplus (start = \llbracket W_{1_s} \rrbracket, end = \iota), \dots, (\ell_i = \llbracket V_{m_i} \rrbracket)_{i \in 1..l} \oplus (start = \llbracket W_{m_s} \rrbracket, end = \iota), \\ (\ell_i = \llbracket V_{m+1_i} \rrbracket)_{i \in 1..l} \oplus (start = \llbracket W_{m+1_s} \rrbracket, end = \llbracket W_{m+1_e} \rrbracket), \dots, (\ell_i = \llbracket V_{n_i} \rrbracket)_{i \in 1..l} \oplus (start = \llbracket W_{n_s} \rrbracket, end = \llbracket W_{n_e} \rrbracket) \wr$$

and after performing the **insert** of $affected$, the table stands

$$\Delta''(t) = \wr (\ell_i = \llbracket V_{1_i} \rrbracket)_{i \in 1..l} \oplus (start = \llbracket W_{1_s} \rrbracket, end = \iota), \dots, (\ell_i = \llbracket V_{m_i} \rrbracket)_{i \in 1..l} \oplus (start = \llbracket W_{m_s} \rrbracket, end = \iota), \\ (\ell_i = \llbracket V_{m+1_i} \rrbracket)_{i \in 1..l} \oplus (start = \llbracket W_{m+1_s} \rrbracket, end = \llbracket W_{m+1_e} \rrbracket), \dots, (\ell_i = \llbracket V_{n_i} \rrbracket)_{i \in 1..l} \oplus (start = \llbracket W_{n_s} \rrbracket, end = \llbracket W_{n_e} \rrbracket), \\ (\ell_i = \llbracket V'_{1_i} \rrbracket)_{i \in 1..l} \oplus (start = \iota, end = \infty), \dots, (\ell_i = \llbracket V'_{m_i} \rrbracket)_{i \in 1..l} \oplus (start = \iota, end = \infty) \wr$$

which is equal to the database obtained by performing the λ_{TLINQ} update, as required. \square

Theorem 4.1. If $\cdot \vdash M:A!E$ and $M \Downarrow_{\Delta, t}^V (V, \Delta')$ for some $\text{wf}(\Delta)$, then $\llbracket M \rrbracket \Downarrow_{\Delta, t} (\llbracket V \rrbracket, \downarrow \Delta')$

Proof. By induction on the derivation of $M \Downarrow_{\Delta, t}^V (V, \Delta_V)$.

Case EV-SeqInsert

Assumption:

$$M \Downarrow_{\Delta, t}^V (t, \Delta_1) \quad N \Downarrow_i^* \wr \tilde{V} \wr \\ \forall data^{[start, end)} \in \tilde{V}. start < end \\ \Delta_2 = \Delta_1 [t \mapsto \Delta_1(t) \hat{\cup} \wr \tilde{V} \wr]$$

$$\text{insert sequenced } M \text{ values } N \Downarrow_{\Delta, t}^V ((), \Delta_2)$$

Translation:

```
(insert sequenced M values N)
  let tbl =  $\llbracket M \rrbracket$  in
  let rows =
    for (x ←  $\llbracket N \rrbracket$ )
       $\wr \eta(x, \tilde{V}) \oplus (start = x.start, end = x.end) \wr$ 
  in
  insert tbl values rows
```

By the IH, $\llbracket M \rrbracket \Downarrow_{\Delta, t} (t, \downarrow \Delta')$.

By Lemma D.5, $\langle N \rangle \Downarrow_{\Delta, t}^* \{ \widetilde{V} \}$; since each V_i is a value of type $\text{ValidTime}(A)$, we have that each V_i must be of the form $W_{data}^{[W_{start}, W_{end}]}$.

Thus we have that by the definition of the translation on database rows and Lemma D.7, each $\langle V_i \rangle$ must be of the form: $(data = W_{data}, start = W_{start}, end = W_{end})$

By Lemma D.1 the **for** comprehension evaluates to a bag of the form $(\eta(x, \widetilde{\ell})) \oplus (start = W_{start}, end = W_{end})$.

Note that this matches the definition of $\Downarrow V_i$.

Therefore by E-INSERT:

$$\frac{\langle M \rangle \Downarrow_{\Delta, t} (t, \Delta_1) \quad \langle N \rangle \Downarrow_t^* \{ \widetilde{V} \}}{\text{insert } M \text{ values } N \Downarrow_{\Delta, t} ((, \Delta_1 [t \mapsto \Delta_1(t) \hat{\cup} \widetilde{V}])}$$

as required.

Case EV-NonseqDelete

Assumption:

$$\frac{M \Downarrow_{\Delta, t} (t, \Delta_1) \quad \Delta_2 = \Delta_1 [t \mapsto \{ d \in \Delta(t) \mid N\{d/x\} \Downarrow_t^* \text{false} \}]}{\text{delete nonsequenced } (x \leftarrow M) \text{ where } N \Downarrow_{\Delta, t} ((, \Delta_2)}$$

Translation:

$$\begin{aligned} & (\text{delete}^{(\ell_i: A_i)_i} \text{nonsequenced } (x \leftarrow M) \text{ where } N) \\ & \quad \text{delete } (x \leftarrow \langle M \rangle) \text{ where lift}(x, \langle N \rangle) \\ & \text{where lift}(x, f) \triangleq \\ & \quad (\lambda x.f) (data = \eta(x, \{\ell_i\}_i), start = x.start, end = x.end) \end{aligned}$$

We want to show:

$$\frac{M \Downarrow_{\Delta, t} (t, \Delta_1) \quad \Delta_2 = \Delta_1 [t \mapsto \{ d \in \Delta_1(t) \mid \text{lift}(x, \langle N \rangle)\{d/x\} \Downarrow_t^* \text{false} \}]}{\text{delete nonsequenced } (x \leftarrow \langle M \rangle) \text{ where lift}(x, \langle N \rangle) \Downarrow_{\Delta, t} ((, \Delta_2)}$$

By the IH we have that $M \Downarrow_{\Delta, t} (t, \Delta_1)$.

By TV-NONSEQDELETE, we have that $\Gamma, x : \text{ValidTime}(A) \vdash N : \text{Bool} ! \emptyset$.

Since $N\{d/x\} \Downarrow_t^* V_{res}$ where d is some row $(\ell_i = V_i)_i^{[W_{start}, W_{end}]}$, it follows by Lemma D.5 that $\langle N \rangle\{d/x\} \Downarrow_t^* \langle V_{res} \rangle$.

By Lemma D.7, $\langle V_{res} \rangle = V_{res}$.

Consider the definition of lift. We can show:

$$\frac{(data = \eta(x, \{\ell_i\}_i), start = x.start, end = x.end) \Downarrow_t^* (data = (\ell_i = V_i)_i, start = W_{start}, end = W_{end}) \quad \langle N \rangle\{d/x\} \Downarrow_t^* V_{res}}{(\lambda x. \langle N \rangle) (data = \eta(x, \{\ell_i\}_i), start = x.start, end = x.end) \Downarrow_t^* V_{res}}$$

since $(data = (\ell_i = V_i)_i, start = W_{start}, end = W_{end}) = \langle d \rangle$.

Therefore we know that the lifted translated predicate will evaluate to the same value as the source predicate in λ_{VLINQ} . It follows that the **delete** operation will affect the same rows, resulting in the same (flattened) database, as required.

Case EV-NonseqUpdate

Follows the same reasoning as EV-NONSEQDELETE.

Case EV-SeqUpdate

As before, it suffices to reason about each case of the update in turn.

By the IH, we have that $\langle L \rangle \Downarrow_{\Delta, t} (t, \Delta')$.

By Lemma D.5, we have that:

- $M_1 \Downarrow_t^* \langle V_{start} \rangle$
- $M_2 \Downarrow_t^* \langle V_{end} \rangle$

and by Lemma D.7, $(V_{start}) = V_{start}$ and $(V_{end}) = V_{end}$.

Therefore let $aStart = V_{start}$ and $aEnd = V_{end}$.

Since the two **insert** statements occur after the **update** statement, the records in *startRows* and *endRows* will not affect the **update**. It therefore suffices to prove the more general statement by considering a single row at a time.

Therefore, we now reason by cases on upd with a database of the form: $\Delta' = [t \mapsto \downarrow v^{[start, end]} \int]$.

Since we know databases are well formed, we know that for each row, $start < end$. By the premises, we also know that $V_{start} < V_{end}$.

Subcase Case 1: $M_3\{v/x\} \Downarrow_t^* \mathbf{true}$ and $V_{start} \leq start$ and $V_{end} \geq end$

Here we have that $V_{start} \leq start < end \leq V_{end}$.

Since $V_{start} \leq start$, we have that $x.start \not\prec V_{start}$. Therefore, $startRows = \int \int$.

Since $V_{end} \geq end$, we have that $x.end \not\prec V_{end}$. Therefore, $endRows = \int \int$.

The **update** predicate will match since by Proposition D.2 and Lemma D.6, $\text{restrict}(x, \{\ell_i\}_{i \in I}, (M_3)) \Downarrow_t^* \mathbf{true}$.

By Proposition D.2, and Lemmas D.6 and D.7, $\text{restrict}(x, \{\ell_i\}_{i \in I}, (N_j)) \Downarrow_t^* W_j$ for all $j \in J$.

In this case:

- $\mathbf{greatest}(x.start, V_{start}) = x.start$
- $\mathbf{least}(x.end, V_{end}) = x.end$

Thus after the update and two null inserts we have:

$$\Delta'' = [t \mapsto \int (\ell_i = V_i)_{i \in I} \int \oplus (\ell_j = W_j)_{j \in J} \oplus (start = start, end = end) \int].$$

which is equal to $[t \mapsto \downarrow (v \mathbf{with} \ell = W)^{[start, end]} \int]$ as required.

Subcase Case 2: $M_3\{v/x\} \Downarrow_t^* \mathbf{true}$ and $V_{start} \leq start$ and $V_{end} < end$

In this case we have that $V_{start} \leq start \leq V_{end} < end$.

Since $x.start < V_{start}$, the *startRows* predicate does not hold and so $startRows = \int \int$.

However, $x.start < V_{end}$ and $x.end > V_{end}$, so $\int \eta(x, \{\ell_j\}_{j \in J}) \oplus (start = V_{end}, end = x.end) \int \Downarrow_t \int v \oplus (start = V_{end}, end = end) \int = endRows$.

Again the **update** predicate matches; this time we have:

- $\mathbf{greatest}(start, V_{start}) = start$
- $\mathbf{least}(end, V_{end}) = V_{end}$

and after the **update** and **insert** our database is:

$$[t \mapsto \int (v \oplus (start = V_{end}, end = end)), ((\ell_i = V_i)_{i \in I} \int \oplus (\ell_j = W_j)_{j \in J} \oplus (start = start, end = V_{end})) \int]$$

which is equal to

$$[t \mapsto \downarrow \int (v \mathbf{with} (\ell_j = W_j)_{j \in J})^{[start, V_{end}]}, v^{[start=V_{end}, end=end]} \int]$$

as required.

Subcase Case 3: $M_3\{v/x\} \Downarrow_t^* \mathbf{true}$ and $V_{start} > start$ and $V_{end} < end$

Here we have that $start < V_{start} < V_{end} < end$.

By similar reasoning to the above case (we omit references to specific lemmas, which are as above), we have that:

- The where clause of *startRows* evaluates to **true** as both $x.start < V_{start}$ and $x.end > V_{start}$. Thus $startRows = \int v \oplus (start = start, end = V_{start}) \int$
- The where clause of *endRows* evaluates to **true** as both $x.start < V_{end}$ and $x.end > V_{end}$. Thus $endRows = \int v \oplus (start = V_{end}, end = end) \int$.

Similarly, the **update** will apply since $x.start < V_{end}$ and $x.end > V_{start}$.

We have in this case that $\mathbf{greatest}(start, V_{start}) = V_{start}$ and $\mathbf{least}(end) V_{end} = V_{end}$.

Thus after the update and two inserts, our final database is equal to:

$$[t \mapsto \downarrow \int v^{[start, V_{start}]}, W^{[V_{start}, V_{end}]}, v^{[V_{end}, end]} \int]$$

as required.

Subcase Case 4: $M_3\{v/x\} \Downarrow_t^* \mathbf{true}$ and $V_{start} > start$ and $V_{end} \geq end$

Similar to Case 2, except there is a record produced in *startRows* instead of *endRows*.

Subcase Case 5: other cases

This ‘catch all’ case boils down to three sub-subcases:

Subsubcase $M_3 \Downarrow_t^* \mathbf{false}$

In this case, the where clause of the two queries would evaluate to **false**, so we would have $startRows = endRows = \{\}$. Furthermore, for the same reason, the **update** would not apply. Therefore, the database would be unaltered.

Subsubcase $V_{end} < start$

Here we have that $V_{start} < V_{end} \leq start < end$.

In this case, the where clause for *startRows* would evaluate to **false** as $x.start \not\prec V_{start}$, and the where clause for *endRows* would evaluate to **false** since $x.start \not\prec V_{end}$. Thus, $startRows = endRows = \{\}$.

The **update** would not apply since $x.start \not\prec V_{end}$.

Therefore, the database will be unaltered.

Subsubcase $V_{start} > end$

Here we have that $start < end \leq V_{start} < V_{end}$.

Again, the where clause of *startRows* would evaluate to **false** since $x.end \not\prec V_{start}$, and the where clause of *endRows* would evaluate to **false** since $x.end \not\prec V_{end}$. Thus, $startRows = endRows = \{\}$.

The update would not apply since $x.end \not\prec V_{start}$.

Therefore, the database will be unaltered.

Case EV-SeqDelete

Follows the same reasoning as EV-SEQUPDATE. The main difference is that the final **update** is replaced with a **delete**, so where the translation results in an updated record in EV-SEQUPDATE, the translation results in the absence of a row here. \square

E Links examples

New count values are added to the table with a *sequenced insert*, using the upload time as the start time.

```
fun insertNewData (new) {
  vt_insert sequenced covid_data
  values (subcat, weekdate, count)
  [withValidity(
    (subcat = new.subcat, weekdate = new.weekdate,
     count = new.count),
    upload_time, forever]]
}
```

Accepted updates are added using sequenced updates. Here `accepted_updates` is a list of updated values that have been approved by the user. For each element of the list, a sequenced update is made.

```
fun updateData (accepted_updates) {
  for (x <- accepted_updates)
    [update sequenced (y <-v- covid_data)
     between (x.time_added, forever)
     where (x.subcat==y.subcat && x.weekdate==y.weekdate)
     set (count = x.new_value) ]
}
```

For update provenance queries of individual counts, a self join is computed over the subcategory and week fields of the valid time table to provide a nested result table where each count is associated with a list of count values and their associated start and end time information. This is a nonsequenced query because the time period information is explicitly added to the result table.

```
query nested {for (x <- vtCurrent(covid_data))
  [(subcat = x.subcat, weekdate = x.weekdate,
   count = x.count, mods =
   for (y <-v- covid_data)
     where (x.weekdate==vtData(y).weekdate &&
            x.subcat==vtData(y).subcat &&
            [(vtf=vtFrom(y),vtt=vtTo(y),
             count=vtData(y).count)]))]}
```

Received 2022-08-12; accepted 2022-10-10