

Data Provenance for Recursive SQL Queries

Benjamin Dietrich Tobias Müller Torsten Grust
University of Tübingen
Tübingen, Germany
[b.dietrich, to.mueller, torsten.grust]@uni-tuebingen.de

ABSTRACT

The adoption of recursion in SQL—framed either in terms of recursive common table expressions (CTEs) or recursive user-defined functions (UDFs)—marked a jump in the expressivity of the query language. The resulting queries can perform complex computation close to database-resident data but, at the same time, often prove challenging to understand and debug. We build on earlier work on the derivation of *where*- and *why*-provenance for complex (yet non-recursive) SQL queries to also embrace recursive SQL CTEs and UDFs. Fine-grained data provenance for recursive SQL is derived through language-level query rewriting and a two-phase evaluation strategy that does not invade the underlying RDBMS.

CCS CONCEPTS

• **Information systems** → **Structured Query Language**; • **Software and its engineering** → **Recursion**; • **Theory of computation** → **Data provenance**.

KEYWORDS

SQL, recursion, CTEs, UDFs, data provenance, query debugging

ACM Reference Format:

Benjamin Dietrich, Tobias Müller, and Torsten Grust. 2022. Data Provenance for Recursive SQL Queries. In *14th International Theory and Practice of Provenance (TaPP'22)*, June 17, 2022, Philadelphia, PA, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3530800.3534536>

1 INTRODUCTION

“*Rekursiv geht meistens schief*” (loosely translated: *recursion typically goes awry*) is a rhyming slogan widely known in German software development communities. While witty, this catchline expresses that it can indeed be challenging to fully wrap one’s head around the progress of a recursive computation. What are the intermediate states? Will this base case be reached eventually?

Since the advent of recursion in SQL—either in the form of SQL:1999’s *recursive common table expressions* (CTEs, WITH RECURSIVE) [7] or self-referential user-defined functions (recursive UDFs, as supported by PostgreSQL, for example [14])—these challenges also surface in the context of relational query languages:

1. The semantics of a recursive CTE (see Figure 1) is defined in terms of a fixpoint computation in which a query $q_{\cup}(T)$ is iteratively evaluated over its most recent result table T until no (previously unseen) rows are produced. The intermediate results of q_{\cup} are appended to form the overall result of the CTE. To fully understand query progress, developers would need to keep track of the states of the *working*, *intermediate*, and *union tables* that implement this iteration [8].
2. A recursive SQL UDF may invoke itself, possibly at multiple call sites (non-linear recursion) deeply embedded into a SQL query that post-processes the results once the recursive calls return (non-tail recursion). Here, keeping track requires an understanding of the UDF call stack and its frames that hold the UDF’s query context and current arguments, for example.

Should such queries indeed go awry, the developer’s debugging toolbox is rather empty. Recursive CTEs are not easily instrumented without destructive effects on their monotonicity or termination properties. Likewise, a recursive UDF would need to be rewritten to return information about intermediate states along with the actual function result. Such manual query instrumentation may affect or hide existing bugs and introduce entirely new ones.

Data provenance for recursive SQL queries. We propose to bank on *data provenance* as one tool that can provide insight into recursive query computation. Here, we focus on *where*- and *why*-provenance [3, 12] derived at the level of individual table cells. We aim to derive

- **where-provenance** which identifies those table cells that were *copied or transformed* to compute the next intermediate (or final) state of a recursive query, and
- **why-provenance** to locate those cells that were *inspected to decide* whether a particular value is part of the output at all.

In tandem, both provenance kinds paint a complete picture of which table cells guided the recursive computation or were sourced to construct the overall result. Below, we study recursive CTEs as well as UDFs to demonstrate how such provenance information may help to explain unexpected results or visualize relevant input cells.

The derivation of provenance for (very) complex queries has proven to be notoriously difficult [3]. Yet, to render provenance useful and practical for SQL query debugging, the derivation strategy is required to embrace constructs that constitute a potential hurdle for query authoring and understanding. This certainly includes constructs like subqueries (including correlation), grouping and aggregation, window functions, complex types (like row values or arrays), or scalar and table-valued built-in and user-defined functions.

```
WITH RECURSIVE T(...) AS (  
  q0           -- evaluated once  
  UNION ALL  
  q∪(T)       -- iterated until ∅  
)  
TABLE T;
```

Figure 1: Recursive CTE.

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.

TaPP'22, June 17, 2022, Philadelphia, PA, USA

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

ACM ISBN 978-1-4503-9349-2/22/06...\$15.00

<https://doi.org/10.1145/3530800.3534536>

The present work continues our earlier effort to derive *where*- and *why*-provenance for a SQL dialect that admits all language constructs listed above. To this we add the ability to process recursive CTEs and recursive SQL UDFs. To analyze SQL query q , we continue to pursue a two-phase evaluation process in which

- **Phase 1** evaluates a variant query q^1 that processes the original input tables while it also writes a protocol about the outcome of predicate evaluation, before
- **Phase 2** reads the interim protocol to evaluate variant q^2 which entirely operates over *sets of dependencies* between table cells. *Where*- and *why*-provenance for q may then be read off the resulting dependency sets.

The variant queries q^1 and q^2 are derived from q through systematic rewriting. Like q , both variants are regular SQL queries that can be evaluated on top of off-the-shelf RDBMSs. No kernel-level changes are called for.

We recapitulate the two phases when we turn to the sample recursive queries below (Sections 2 and 3). For a full review of the approach we refer to [12]. We have laid out the present paper as a “companion” to this earlier work—in particular, here we provide an addendum of inference rules for SQL query rewriting that form a coherent whole with the rules found in the companion paper [12]. Two-stage evaluation and the processing of—potentially large—dependency sets has an impact on query evaluation time. We quantify the slowdown in Section 3.2 but will also show how Phase 2 can peruse the mentioned interim protocol to actually perform significantly better than regular query evaluation. A review of related efforts is found in Section 4.

2 UNRAVELING RECURSIVE CTEs

When a SQL query yields unexpected results, provenance helps to zoom in on the relevant input data items and (potentially buggy) query portions. We have made this observation for non-recursive queries in [12]—here we extend it to recursive CTEs.

Let us focus on CTE `bom` of Figure 2 which recursively computes the bill of materials (or: parts explosion) of a humanoid robot. (We have adapted this query from the PostgreSQL manual [14]—the manual contains just the bug we discuss here.) Robot parts, along with their sub-parts and required quantities, are held in input table `parts` (see Figure 3(a)). The CTE yields the output table of Figure 3(a) which lists the quantities (column `qty`) of all parts (column `sub_part`) required to assemble the robot. The numbers of required fingers and feet (5 and 1, respectively) certainly look suspicious: we had expected 10 and 2. To understand how CTE `bom` arrived at these questionable results, we explore the provenance of output value 5 (5) in Figure 3(a)).

The *why*-provenance of (5) reveals all input table cells that have been *inspected to decide* whether (5) occurs in the output table. We find the highlighted input cells (arm), (body), and (humanoid) which the CTE has used to recursively descend into the part hierarchy encoded by table `parts` (see predicate `p.part = b.sub_part` in Line 9 of Figure 2). The chain humanoid–body–arm describes the expected path from root part humanoid to sub-part finger: the recursive traversal expressed by CTE `bom` appears to be in order.

The *where*-provenance of output (5), however, is shown to only refer to the quantity (5) of fingers on an arm. This is unexpected:

```

1 WITH RECURSIVE bom(part, sub_part, qty) AS (
2   SELECT p.part, p.sub_part, p.qty
3   FROM parts AS p
4   WHERE p.part = 'humanoid'
5   UNION ALL
6   SELECT p.part, p.sub_part, p.qty -- BUG (should read p.qty * b.qty)
7   FROM bom AS b,
8       parts AS p
9   WHERE p.part = b.sub_part
10 )
11 SELECT b.sub_part, b.qty
12 FROM bom AS b
    
```

Figure 2: Recursive CTE to find the bill of materials for a humanoid robot. The computation of `qty` in Line 6 is buggy.

parts			output (buggy)	
part	sub_part	qty	sub_part	qty
humanoid	head	1	head	1
humanoid	body	1	body	1
body	arm	2	arm	2
body	leg	2	leg	2
arm	finger	5	finger	5
leg	foot	1	foot	1
chassis	wheel	4		

(a) Provenance derivation for the suspicious output value (5) yields *where*-provenance (●) and *why*-provenance (○).

parts			output	
part	sub_part	qty	sub_part	qty
humanoid	head	1	head	1
humanoid	body	1	body	1
body	arm	2	arm	2
body	leg	2	leg	2
arm	finger	5	finger	10
leg	foot	1	foot	2
chassis	wheel	4		

(b) Provenance after the bug in CTE `bom` has been fixed.

Figure 3: Input and output tables for CTE `bom` of Figure 2.

no other value in column `qty` has been accessed to compute the output (5). Indeed, the query disregards parent part quantities while it walks down the hierarchy (and thus misses the fact that a body has two arms, for example). To correctly compute the quantity of a part p , we need to factor in the quantity of its parent part b . Once we fix the quantity calculation to read `p.qty * b.qty` in Line 6, CTE `bom` yields the expected finger count (10), in output table of Figure 3(b). We also find the expected *where*-provenance now: during hierarchy traversal, we multiply quantities (1) * (2) * (5) to arrive at (10) fingers.

2.1 Values Here, Dependencies There

This work pursues an approach to provenance derivation [12] that strictly separates

1. the *realm of regular values* (e.g., the string or integer data found in table cells) from
2. the *realm of dependency sets* which describe the table cells that influenced the computation of those values.

This strict separation applies to both, tables and queries, used during provenance derivation. We turn to the tables first.

Tables and their mirror images. Given a table T , provenance derivation distinguishes between its

- variant T^1 which is an (almost exact) copy of T in which table cells hold regular values of the known SQL types, and
- its mirror image T^2 whose table cells hold dependencies of type \mathbb{P} (the type of sets of cell identifiers).

Throughout, we maintain a one-to-one correspondence between both: any value cell in T^1 is associated with its dedicated dependency set in T^2 . In consequence, T^2 has the same cardinality and columns as T^1 . Since T^1 , T^2 are relational tables (and thus are unordered), both carry a column ϱ of row identifiers which tie corresponding rows together.

Figure 4 shows the input and output tables of the bom query of Figure 2, both in their value and their dependency set variants (disregard the overlaid arrows for now). In the input tables, you will find that any cell value in parts^1 is associated with a *singleton* dependency set in parts^2 . To illustrate, value 5 of type `int` is associated with set $\{p_{15}\}$ of type \mathbb{P} , both in row ϱ_5 and column `qty` of their tables: cell identifier p_{15} represents the cell of value 5 and there is no dependency to any other cell.

In the (non-singleton, in general) dependency sets of output table output^2 , provenance derivation has accumulated all cell identifiers that influenced the computation of the associated value cell in output^1 . Row and cell identifiers suffice to trace the data provenance of any output table cell. For output value 10 in table output^1 , for example, we find the following (following the arrows overlaid on Figure 4):

1. The computation of 10 (row ϱ_{35} , column `qty`) depended on the 8 input cells with identifiers $\{p_6, p_9, \dots, p_{13}\}$, see table output^2 . (The grey p_{13} indicates *why*-provenance—we come back to the distinction between provenance kinds in Section 2.2.)
2. Among these, p_{13} designates row ϱ_5 , column `part` of input table parts^2 .
3. The value associated with cell p_{13} is string `arm` in row ϱ_5 , column `part` of table parts^1 .

(If we trace all 8 input cells influencing output 10, (if we obtain the 8 green highlights in Figure 3(b).)

We assume that input table T^1 is provided. Its column ϱ of row identifiers may either hold externalized RDBMS-internal row IDs or can be generated explicitly through row numbering. Table T^2 may be derived from T^1 in terms of a SQL view that invents singleton sets of arbitrary, yet unique cell identifiers (*i.e.*, the $\{p_i\}$ discussed above). Several options exist to represent these dependency sets in a SQL system: in [12], we discuss arrays as well as variable-length bit sets as two possible implementations of type \mathbb{P} .

Queries and their mirror images. The strict separation of the value and dependency set realms is reflected by queries as well. Given a SQL query q , we systematically rewrite it into a query pair $q \Rightarrow \langle q^1, q^2 \rangle$ that we evaluate in two phases (see Figure 5):

- In Phase 1, SQL query q^1 consumes tables T^1, \dots of cell values of the regular SQL types and emits table output^1 . Unlike q , q^1 additionally writes a protocol about value-based decisions performed during query evaluation.
- In Phase 2, SQL query q^2 entirely operates over tables of dependency sets T^2, \dots , *i.e.*, all table cells processed by q^2 are of type \mathbb{P} . While it executes, query q^2 consults the interim protocol and finally emits table output^2 .

While the paired queries operate in separate realms, both (1) read and emit tables of identical shape (cardinality and column width) and (2) adhere to a common syntactic structure: the rewrite \Rightarrow maps subexpressions e^1 of q^1 to their mirror image e^2 of q^2 in a compositional fashion. In a nutshell, both expressions relate as follows:

- Assume $e^1 \equiv x \otimes y$ (where \otimes denotes an arbitrary binary SQL operator). If e^1 yields value z , the dependencies of z comprise the dependencies of both argument values x and y .
- The mirror image of e^1 will be $e^2 \equiv x \cup y$ in which x and y are the dependency sets of the values x and y . e^2 thus reads and emits sets of type \mathbb{P} .

Interim protocols. If q^1 contains a clause WHERE e^1 , the Boolean value of e^1 is used to perform value-based decisions during query evaluation in Phase 1. Note that q^2 will *not* be able to simply use WHERE e^2 to reenact these decisions in Phase 2, since e^2 is of type \mathbb{P} . Instead, we instrument query q^1 to invoke function $\text{write}_{\square}(\odot, v_1.\varrho, \dots, v_n.\varrho)$ to record the fact that e^1 evaluated to `true`. In this call,

- \odot identifies the WHERE clause's location in the SQL text (q^1 may contain multiple such clauses), and
- the v_1, \dots, v_n denote the SQL row variables that occur free in e^1 (the Boolean value of e^1 depends on the rows bound to these variables).

Function write_{\square} then (1) returns a unique row identifier ϱ representing the row that passed predicate e^1 under the current bindings of the v_i , and (2) saves $(\odot, v_1.\varrho, \dots, v_n.\varrho, \varrho)$ to a persistent protocol to record this value-based decision.

Phase 2 can reenact just this behavior through the invocation of $\text{read}_{\square}(\odot, v_1.\varrho, \dots, v_n.\varrho)$. read_{\square} returns ϱ if $(\odot, v_1.\varrho, \dots, v_n.\varrho, \varrho)$ is found in the protocol but yields \emptyset (represented as the empty table) otherwise. We have thus defined \Rightarrow to replace predicate evaluation in q^1 with corresponding read_{\square} invocations in q^2 (peek ahead at Lines 2 and 5 in Figures 7(a) and 7(b) to see how this plays out for the bom query). Analogously, the interim protocol can be used to communicate the formation of groups or the elimination of rows due to DISTINCT from Phase 1 to Phase 2.

Again, multiple implementation options exist to realize the protocol and the side effects performed by read_{\square} and write_{\square} . In [12], both functions were realized as SQL UDFs that write to and read from a common table. We will report on protocol sizes in Section 3.2.

The companion paper [12] defined syntax-directed rewrites $q \Rightarrow \langle q^1, q^2 \rangle$ for a broad range of SQL constructs. Below, we add rewriting rules for recursive CTEs (and recursive UDFs in Section 3).

2.2 Recursive Provenance Derivation

In tandem, the rewrite Rules WITH and UNIONALL of Figure 6 extend the definition of \Rightarrow to embrace the syntactic shape of recursive CTEs (see Figure 1). Both rules follow the principle of compositionality: the rewrite of a complex query construct is assembled from the rewrites of its (simpler) constituent queries—see the rewrites $q_i \Rightarrow \langle q_i^1, q_i^2 \rangle$ ($i = 0, \dots, n$) in Rule WITH as well as $q_0 \Rightarrow \langle q_0^1, q_0^2 \rangle$ and $q_{\cup} \Rightarrow \langle q_{\cup}^1, q_{\cup}^2 \rangle$ in Rule UNIONALL. Rule WITH extends the CTE's column list by column ϱ to preserve the row identifiers that tie values and their associated dependency sets together (recall Section 2.1). Other than that, both rules preserve the syntactic shape

(a) Input tables (original and dependency set mirror image).

parts ¹				parts ²			
part	sub_part	qty	ρ	part	sub_part	qty	ρ
humanoid	head	1	ρ ₁	{p ₁ }	{p ₂ }	{p ₃ }	ρ ₁
humanoid	body	1	ρ ₂	{p ₄ }	{p ₅ }	{p ₆ }	ρ ₂
body	arm	2	ρ ₃	{p ₇ }	{p ₈ }	{p ₉ }	ρ ₃
body	leg	2	ρ ₄	{p ₁₀ }	{p ₁₁ }	{p ₁₂ }	ρ ₄
arm	finger	5	ρ ₅	{p ₁₃ }	{p ₁₄ }	{p ₁₅ }	ρ ₅
leg	foot	1	ρ ₆	{p ₁₆ }	{p ₁₇ }	{p ₁₈ }	ρ ₆
chassis	wheel	4	ρ ₇	{p ₁₉ }	{p ₂₀ }	{p ₂₁ }	ρ ₇

(b) Output tables generated by provenance derivation.

output ¹			output ²		
sub_part	qty	ρ	sub_part	qty	ρ
head	1	ρ ₃₁	{p ₂ , p ₁ }	{p ₃ , p ₁ }	ρ ₃₁
body	1	ρ ₃₂	{p ₅ , p ₄ }	{p ₆ , p ₄ }	ρ ₃₂
arm	2	ρ ₃₃	{p ₈ , p ₄ , p ₅ , p ₇ }	{p ₆ , p ₉ , p ₄ , p ₅ , p ₇ }	ρ ₃₃
leg	2	ρ ₃₄	{p ₁₁ , p ₄ , p ₅ , p ₁₀ }	{p ₆ , p ₁₂ , p ₄ , p ₅ , p ₁₀ }	ρ ₃₄
finger	10	ρ ₃₅	{p ₁₄ , p ₄ , p ₅ , p ₇ , p ₈ , p ₁₃ }	{p ₆ , p ₉ , p ₁₅ , p ₄ , p ₅ , p ₇ , p ₈ , p ₁₃ }	ρ ₃₅
foot	2	ρ ₃₆	{p ₁₇ , p ₄ , p ₅ , p ₁₀ , p ₁₁ , p ₁₆ }	{p ₆ , p ₁₂ , p ₁₈ , p ₄ , p ₅ , p ₁₀ , p ₁₁ , p ₁₆ }	ρ ₃₆

Figure 4: Input and output tables of the bom query of Figure 2 with row identifiers ρ . Arrows trace output value 10 back to its dependency arm in input table parts¹. In output², p_i denotes *where*-provenance while p_j denotes *why*-provenance.




Figure 5: Two-phase provenance derivation.

of the original query, a salient feature of the two-phase approach that aids the efficient evaluation of the rewritten queries [12].

Given these extensions, \Rightarrow rewrites the bom CTE of Figure 2 into the CTE pair $(\text{bom}^1, \text{bom}^2)$ depicted in Figure 7. Where these generated CTEs exhibit relevant changes from the original bom, we have added blue highlights.

Phase 1, Figure 7(a). In CTE bom^1 , Line 2 calls $\text{write}_{\text{FILTER}}(\odot, p, \rho)$ to record the fact that row p has passed the predicate $p.\text{part} = \text{'humanoid'}$. Likewise, $\text{write}_{\text{JOIN}}(\odot, b, \rho, p, \rho)$ in Line 8 creates a protocol entry if the rows bound to row variables b and p joined under condition $p.\text{part} = b.\text{sub_part}$.

When query bom^1 is executed on input table parts¹ of Figure 4(a), the recursive CTE performs the initial query q_0 in Lines 2–5 once before it iterates query q_{\cup} in Lines 8–11 twice (recall Figure 1). Figure 8 shows how the intermediate results of these queries are assembled to form output¹. Side effects on the protocol are shown under heading protocol writes (read Figure 8 top-down). The protocol reflects the characteristic iterative nature of a SQL CTE:

- In the first iteration of q_{\cup} , rows ρ_3 and ρ_4 are joined with row ρ_{32} which has been generated by the initial query q_0 (establishing that arm and leg are parts of the robot body),
- in the second iteration of q_{\cup} , rows ρ_5 and ρ_6 join with rows ρ_{33} and ρ_{34} , respectively, which have just been generated by the first iteration (finger is part of arm, foot is part of leg).

The protocol contents thus contain a complete history of the iterated joins performed by the CTE.

Phase 2, Figure 7(b). CTE bom^2 operates over dependency sets and thus trades value-based expressions like $p.\text{qty} * b.\text{qty}$ for

- $p.\text{qty} \cup b.\text{qty}$ (see Line 9 in bom^1 and bom^2) to compute the *where*-provenance of the arithmetic operation. A predicate like $p.\text{part} = b.\text{sub_part}$ decides whether the current bindings for row variables p and b may contribute to the query result. We thus
1. derive the predicate's dependency set by $p.\text{part} \cup b.\text{sub_part}$, and
 2. use function $\Upsilon(\cdot)$ to mark all cell identifiers in that set as *why*-provenance, see Lines 6 and 12 in bom^2 (in table output² of Figure 4(b) we have colored these cell identifiers in grey: p_j). Since *why*-provenance affects all columns of a row that passed a predicate, we bind the resulting dependency set to wh.y once and then refer to that alias as needed to avoid recomputation effort.

Since we exclusively operate over dependency sets, function call $\text{read}_{\text{FILTER}}(\odot, p, \rho)$ in Line 5 assumes the role of predicate $p.\text{part} = \text{'humanoid'}$. Likewise, $\text{read}_{\text{JOIN}}(\odot, b, \rho, p, \rho)$ reenacts the value-based predicate $p.\text{part} = b.\text{sub_part}$. Both functions return the empty table if the corresponding protocol entries are missing. In that case, the current bindings for row variables p and b will not contribute to the query result (just like in Phase 1). If a protocol entry is found, the functions return row identifier $\log.\rho$ of the row that passed the predicate (e.g., $\text{read}_{\text{JOIN}}(\odot, \rho_{33}, \rho_5)$ yields ρ_{35} , see the last but one protocol row in Figure 8). In effect, bom^2 will show the exact filtering/join behavior like its mirror CTE bom^1 . In particular, the queries will perform the exact same number of CTE iterations in both phases.

WITH RECURSIVE ... (... UNION DISTINCT ...). This approach to provenance derivation can be adapted to cover recursive CTEs with set semantics in which iteration ends when *no new* result rows are produced by q_{\cup} . Since the generation of unique row identifiers by the write_{\square} functions can affect the detection of row duplicates, this calls for an appropriate implementation of equality on row identifiers (column ρ). We do not elaborate on the details here, but examples of such CTEs are found in the GitHub repository accompanying this paper (see Section 5).

$$\begin{array}{c}
q_i \Rightarrow \langle q_i^1, q_i^2 \rangle \Big|_{i=0 \dots n} \\
\hline
i^1 = \text{WITH RECURSIVE } t_1^1(\varrho, c_{11}, \dots, c_{1k_1}) \text{ AS } (q_1^1), \dots, \\
\quad t_n^1(\varrho, c_{n1}, \dots, c_{nk_n}) \text{ AS } (q_n^1) \\
q_0^1 \\
\hline
i^2 = \text{WITH RECURSIVE } t_1^2(\varrho, c_{11}, \dots, c_{1k_1}) \text{ AS } (q_1^2), \dots, \\
\quad t_n^2(\varrho, c_{n1}, \dots, c_{nk_n}) \text{ AS } (q_n^2) \\
q_0^2 \\
\hline
\text{WITH RECURSIVE } t_1(c_{11}, \dots, c_{1k_1}) \text{ AS } (q_1), \dots, \\
\quad t_n(c_{n1}, \dots, c_{nk_n}) \text{ AS } (q_n) \Rightarrow \langle i^1, i^2 \rangle \\
q_0 \\
\hline
q_0 \Rightarrow \langle q_0^1, q_0^2 \rangle \quad q_\cup \Rightarrow \langle q_\cup^1, q_\cup^2 \rangle \\
\frac{q_0 \text{ UNION ALL } q_\cup \Rightarrow \langle q_0^1 \text{ UNION ALL } q_\cup^1, q_0^2 \text{ UNION ALL } q_\cup^2 \rangle}{\text{(UNIONALL)}}
\end{array}
\tag{WITH}$$

Figure 6: Rewrite rules $q \Rightarrow \langle q^1, q^2 \rangle$ for recursive SQL CTEs. Combine with the rules of [12].

```

1 WITH RECURSIVE bom1(ϱ, part, sub_part, qty) AS (
2   SELECT writeFILTER(Ⓛ, p.ϱ) AS ϱ,
3     p.part, p.sub_part, p.qty
4   FROM parts1 AS p
5   WHERE p.part = 'humanoid'
6
7   UNION ALL
8   SELECT writeJOIN(Ⓛ, b.ϱ, p.ϱ) AS ϱ,
9     p.part, p.sub_part, p.qty * b.qty
10  FROM bom1 AS b, parts1 AS p
11  WHERE p.part = b.sub_part
12
13 )
14 SELECT b.ϱ, b.sub_part, b.qty
15 FROM bom1 AS b

```

(a) Phase 1 (CTE bom¹).

```

1 WITH RECURSIVE bom2(ϱ, part, sub_part, qty) AS (
2   SELECT log.ϱ,
3     p.part ∪ wh.y, p.sub_part ∪ wh.y, p.qty ∪ wh.y
4   FROM parts2 AS p,
5     readFILTER(Ⓛ, p.ϱ) AS log(ϱ),
6     Y(p.part) AS wh(y)
7
8   UNION ALL
9   SELECT log.ϱ, p.part ∪ wh.y,
10    p.sub_part ∪ wh.y, p.qty ∪ wh.y
11  FROM bom2 AS b, parts2 AS p,
12    readJOIN(Ⓛ, b.ϱ, p.ϱ) log(ϱ),
13    Y(p.part ∪ b.sub_part) AS wh(y)
14
15 SELECT b.ϱ, b.sub_part, b.qty
16 FROM bom2 AS p

```

(b) Phase 2 (CTE bom²).

Figure 7: Rewriting CTE bom (Figure 2) for provenance derivation yields a pair of CTEs.

		output ¹		protocol writes		
		part	sub_part	qty	ϱ	
q_0	humanoid	head		1	ϱ ₃₁	writeFILTER(Ⓛ, ϱ ₁) = ϱ ₃₁
	humanoid	body		1	ϱ ₃₂	writeFILTER(Ⓛ, ϱ ₂) = ϱ ₃₂
q_\cup	body	arm		2	ϱ ₃₃	writeJOIN(Ⓛ, ϱ ₃₂ , ϱ ₃) = ϱ ₃₃
	body	leg		2	ϱ ₃₄	writeJOIN(Ⓛ, ϱ ₃₂ , ϱ ₄) = ϱ ₃₄
q_\cup	arm	finger		10	ϱ ₃₅	writeJOIN(Ⓛ, ϱ ₃₃ , ϱ ₅) = ϱ ₃₅
	leg	foot		2	ϱ ₃₆	writeJOIN(Ⓛ, ϱ ₃₄ , ϱ ₆) = ϱ ₃₆

Figure 8: Result rows and protocol entries generated during the evaluation of bom¹. Result and protocol grow top-down.

3 TRACING RECURSIVE SQL UDFS

CTEs are but one option to express recursive computation right inside the RDBMS (and thus close to the data). In what follows, we focus on *recursive user-defined functions* which provide an—often concise and readable, even elegant—alternative to CTEs. In a recursive UDF f , recursion is directly expressed through self-invocation of f , much like in a functional programming language. While recursive CTEs underly syntactic monotonicity and linearity restrictions [8], no such constraints apply to recursive UDFs. A number of RDBMSs, including PostgreSQL, Oracle, and Microsoft SQL Server, support recursive UDFs off-the-shelf. Provenance derivation helps to untangle, debug, or visualize the potentially complex recursive computation that such UDFs can perform.

Function dtw of Figure 9 follows the recursive UDF style. dtw uses three-fold recursion (see the recursive calls in Lines 8–10) to compute the *Dynamic Time Warping* (short: DTW) score of two time series $X = (x_i)$ and $Y = (y_j)$ [2]. DTW has the ability to locally stretch and compress the series to minimize the distance between both. The *warp* parameter w of dtw determines how far such stretching may go: x_i may be matched with any y_j such that $|i - j| \leq w$.

Applications of DTW abound and include speech recognition or feature comparison in machine-learning setups.

As input, dtw assumes a tabular encoding distances (i, j, δ) of the distance matrix of both time series: a row (i, j, δ) indicates that $\delta = |x_i - y_j|$. Figure 10 shows the distances tables and its matrix representation for two sample time series of length 6 depicted in Figure 10(c). A UDF call dtw(6,6,w=1) for these time series yields a distance score of 1.0.

If we derive *where*- and *why*-provenance for result value 1.0 to understand how UDF dtw arrived at this score, we find the dependencies marked ● and ○ in Figure 10. We find that DTW has inspected all entries in distances that are close to the matrix' main diagonal, *i.e.*, if their indices satisfy $|i - j| \leq w$ (*why-provenance* ○ along the diagonal). Inside this window around the diagonal, DTW has added the minimal distances to find the overall score: indeed, ● + ... + ● + ● + ● yields 1.0 (*where-provenance*). We also learn that a *warp* parameter of $w=2$ would have widened the

```

1 CREATE FUNCTION dtw(i int, j int, w int) RETURNS float AS $$
2 SELECT CASE
3   WHEN dtw.i=0 AND dtw.j=0 THEN 0.0
4   WHEN dtw.i=0 OR dtw.j=0 THEN ∞
5   WHEN abs(dtw.i-dtw.j)>dtw.w THEN ∞
6   ELSE (SELECT d.δ
7         + LEAST(
8           dtw(dtw.i-1, dtw.j-1, dtw.w),
9           dtw(dtw.i-1, dtw.j, dtw.w),
10          dtw(dtw.i, dtw.j-1, dtw.w))
11  FROM distances AS d
12  WHERE (d.i,d.j)=(dtw.i,dtw.j))
13 END
14 $$ LANGUAGE SQL STABLE;

```

Figure 9: A recursive SQL UDF that implements DTW.

distances		
i	j	δ
1	1	0
1	2	0
1	3	1
1	4	1
1	5	0
1	6	1
1	7	0
1	8	1

(a) Distance table.

6	0	0	1	0	0	0
5	1	1	0	1	1	1
4	1	1	0	1	1	1
3	1	1	0	1	1	1
2	0	0	1	0	0	0
1	0	0	1	0	0	0
	1	2	3	4	5	6


(b) Distance matrix $\delta(i, j)$.(c) Sample time series $X = (x_i)$ and $Y = (y_j)$.

Figure 10: Distance table, distance matrix, and visualization of the time series $X = (x_i)$, $Y = (y_j)$. Provenance derivation for the recursive UDF dtw of Figure 9 generates *where*-provenance \bullet and *why*-provenance \circ that explains the workings of DTW.

$e_i \Rightarrow \langle e_i^1, e_i^2 \rangle \Big _{i=1..n} \quad \textcircled{\text{C}} := \text{unique call site identifier}$ $\{v_1, \dots, v_m\} := \text{free row variables in } e_1, \dots, e_n$ $i^1 := f^1(\text{writeCALL}(\textcircled{\text{C}}, v_1.\rho, \dots, v_m.\rho), e_1^1, \dots, e_n^1)$ $i^2 := f^2(\text{readCALL}(\textcircled{\text{C}}, v_1.\rho, \dots, v_m.\rho), e_1^2, \dots, e_n^2) \quad \text{(UDF-CALL)}$ $f(e_1, \dots, e_n) \Rightarrow \langle i^1, i^2 \rangle$	$q \Rightarrow \langle q^1, q^2 \rangle$ <table border="0" style="width: 100%;"> <tr> <td style="width: 50%; border-right: 1px solid black; padding-right: 10px;"> CREATE FUNCTION udf¹(ρ rowID, $a_1 \tau_1, \dots, a_n \tau_n$) RETURNS τ_{res} AS \$\$ SELECT q^1 \$\$ LANGUAGE SQL VOLATILE </td> <td style="width: 50%; padding-left: 10px;"> CREATE FUNCTION udf²(ρ rowID, $a_1 \mathbb{P}, \dots, a_n \mathbb{P}$) RETURNS \mathbb{P} AS \$\$ SELECT q^2 \$\$ LANGUAGE SQL STABLE </td> </tr> <tr> <td colspan="2" style="text-align: center; border-top: 1px solid black; border-bottom: 1px solid black;"> CREATE FUNCTION udf($a_1 \tau_1, \dots, a_n \tau_n$) RETURNS τ_{res} AS \$\$ SELECT q \$\$ LANGUAGE SQL STABLE </td> </tr> <tr> <td colspan="2" style="text-align: right; padding-right: 10px;"> $\Rightarrow \langle i^1, i^2 \rangle$ </td> </tr> </table>	CREATE FUNCTION udf ¹ (ρ rowID, $a_1 \tau_1, \dots, a_n \tau_n$) RETURNS τ_{res} AS \$\$ SELECT q^1 \$\$ LANGUAGE SQL VOLATILE	CREATE FUNCTION udf ² (ρ rowID, $a_1 \mathbb{P}, \dots, a_n \mathbb{P}$) RETURNS \mathbb{P} AS \$\$ SELECT q^2 \$\$ LANGUAGE SQL STABLE	CREATE FUNCTION udf($a_1 \tau_1, \dots, a_n \tau_n$) RETURNS τ_{res} AS \$\$ SELECT q \$\$ LANGUAGE SQL STABLE		$\Rightarrow \langle i^1, i^2 \rangle$	
CREATE FUNCTION udf ¹ (ρ rowID, $a_1 \tau_1, \dots, a_n \tau_n$) RETURNS τ_{res} AS \$\$ SELECT q^1 \$\$ LANGUAGE SQL VOLATILE	CREATE FUNCTION udf ² (ρ rowID, $a_1 \mathbb{P}, \dots, a_n \mathbb{P}$) RETURNS \mathbb{P} AS \$\$ SELECT q^2 \$\$ LANGUAGE SQL STABLE						
CREATE FUNCTION udf($a_1 \tau_1, \dots, a_n \tau_n$) RETURNS τ_{res} AS \$\$ SELECT q \$\$ LANGUAGE SQL STABLE							
$\Rightarrow \langle i^1, i^2 \rangle$							

Figure 11: Rewrite rules $q \Rightarrow \langle q^1, q^2 \rangle$ for the definition and invocation of recursive UDFs. Combine with the rules of [12].

<pre> 1 CREATE FUNCTION dtw¹(ρ rowID, i int, j int, w int) RETURNS float AS \$\$ 2 SELECT 3 CASE writeCASE($\textcircled{\text{C}}$, dtw¹.ρ, CASE 4 WHEN dtw¹.i=0 AND dtw¹.j=0 THEN 1 5 WHEN dtw¹.i=0 OR dtw¹.j=0 THEN 2 6 WHEN abs(dtw¹.i-dtw¹.j)>dtw¹.w THEN 3 7 ELSE 0 END) 8 WHEN 1 THEN 0.0 9 WHEN 2 THEN '∞' 10 WHEN 3 THEN '∞' 11 WHEN 0 THEN 12 (SELECT t.score FROM (13 SELECT writeJOIN($\textcircled{\text{C}}$, dtw¹.ρ, d.ρ) AS ρ, 14 d.δ + LEAST¹(15 writeCALL($\textcircled{\text{C}}$, dtw¹.ρ, 16 dtw¹(writeCALL($\textcircled{\text{C}}$, dtw¹.ρ, dtw¹.i-1, dtw¹.j-1, dtw¹.w), 17 dtw¹(writeCALL($\textcircled{\text{C}}$, dtw¹.ρ, dtw¹.i-1, dtw¹.j, dtw¹.w), 18 dtw¹(writeCALL($\textcircled{\text{C}}$, dtw¹.ρ, dtw¹.i, dtw¹.j-1, dtw¹.w) 19) AS score 20 FROM distances¹ AS d 21 WHERE (d.i,d.j)=(dtw¹.i,dtw¹.j) 22) AS t(ρ, score)) 23 END 24 \$\$ LANGUAGE SQL VOLATILE; </pre>	<pre> 1 CREATE FUNCTION dtw²(ρ rowID, i \mathbb{P}, j \mathbb{P}, w \mathbb{P}) RETURNS \mathbb{P} AS \$\$ 2 SELECT 3 CASE readCASE($\textcircled{\text{C}}$, dtw².ρ) 4 WHEN 1 THEN $\emptyset \cup Y(\text{dtw}^2.i \cup \text{dtw}^2.j)$ 5 WHEN 2 THEN $\emptyset \cup Y(\text{dtw}^2.i \cup \text{dtw}^2.j)$ 6 WHEN 3 THEN $\emptyset \cup Y(\text{dtw}^2.i \cup \text{dtw}^2.j \cup \text{dtw}^2.w)$ 7 WHEN 0 THEN 8 (SELECT t.score FROM (9 SELECT log.ρ AS ρ, 10 d.δ \cup LEAST²(11 readCALL($\textcircled{\text{C}}$, dtw².ρ), 12 dtw²(readCALL($\textcircled{\text{C}}$, dtw².ρ), 13 dtw².i \cup \emptyset, dtw².j \cup \emptyset, dtw².w), 14 dtw²(readCALL($\textcircled{\text{C}}$, dtw².ρ), 15 dtw².i \cup \emptyset, dtw².j, dtw².w), 16 dtw²(readCALL($\textcircled{\text{C}}$, dtw².ρ), 17 dtw².i, dtw².j \cup \emptyset, dtw².w) 18) \cup Y(d.i \cup d.j \cup dtw².i \cup dtw².j) AS score 19 FROM distances² AS d, 20 readJOIN($\textcircled{\text{C}}$, dtw².ρ, d.ρ) AS log(ρ) 21) AS t(ρ, score)) \cup Y(dtw².i \cup dtw².j \cup dtw².w) 22 ELSE \emptyset 23 END 24 \$\$ LANGUAGE SQL STABLE; </pre>
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

(a) Phase 1 (UDF dtw¹).(b) Phase 2 (UDF dtw²).

Figure 12: Generated pair of recursive UDFs to derive *where*- and *why*-provenance for SQL UDF dtw (Figure 9).

window around the diagonal to include the distances at (i, j) coordinates (3, 5) and (4, 6) to obtain a perfect DTW score of 0.0. (Widening the window inspects more data as *why*-provenance has shown but may lead to better matches—a typical tradeoff in the use of DTW.)

3.1 Rewriting UDFs for Provenance Derivation

We stick to the principle of two-phase provenance derivation: the extended rules for \Rightarrow in Figure 11 rewrite a recursive UDF f into a pair of UDFs $\langle f^1, f^2 \rangle$. The application of these rules to UDF dtw of Figure 9 yields the UDF pair $\langle \text{dtw}^1, \text{dtw}^2 \rangle$ shown in Figure 12.

As expected, dtw^1 computes over regular values (consuming int parameters, yielding the score of type float), while dtw^2 receives and returns dependency sets of type \mathbb{P} .

The body of a (recursive) UDF is a regular SQL query q . Rule UDF-DEF thus recursively applies rewrite \Rightarrow to q to obtain function bodies $\langle q^1, q^2 \rangle$ that perform provenance derivation. These rewrites handle SQL constructs like CASE...WHEN or the invocation of non-recursive functions like LEAST (n -way minimum). Note how literals like 0 or 1 are represented as empty dependency sets \emptyset in Phase 2: a literal does not depend on any input table cell [12]. Unlike the original UDF which we assume to be a true read-only (or: STABLE)




Figure 13: Quantifying the runtime and protocol size overhead of provenance derivation (recursive UDF dtw).

function, the body of dtw^1 will perform protocol writing (recall the discussion in Section 2.1 as well as Figure 8). Rule UDF-DEF changes the function volatility category of dtw^1 to VOLATILE to announce this to the RDBMS [14].

In the case of recursive CTEs, extra row identifiers ρ were introduced to tie values and their dependency sets together. Here, Rule UDF-DEF introduces an *extra UDF parameter* ρ of type rowID to serve the same purpose: if the invocation $f^1(\rho^*, \dots)$ yields a value z , the invocation of the mirror UDF $f^2(\rho^*, \dots)$ using the same call identifier ρ^* derives z 's dependency set. To rewrite a UDF call $f(e_1, \dots, e_n)$, Rule UDF-CALL thus collects the free row variables v_1, \dots, v_m among the function arguments e_i and proceeds as follows:

1. In Phase 1, instrument the UDF to invoke $write_{CALL}(\mathcal{C}, v_1 \cdot \rho, \dots, v_m \cdot \rho)$ to map the row identifiers of the v_j to a unique call identifier ρ^* . This mapping is also saved in the interim protocol as entry $(\mathcal{C}, v_1 \cdot \rho, \dots, v_m \cdot \rho, \rho^*)$. ρ^* is then passed as the first parameter to f^1 to identify this particular call.
2. In Phase 2, read the protocol via $read_{CALL}(\mathcal{C}, v_1 \cdot \rho, \dots, v_m \cdot \rho)$ to retrieve ρ^* and pass it to f^2 . Functions f^1 and f^2 will thus agree on the call identifier as required.

Note how the body queries of dtw^1 and dtw^2 pass the call identifier (referred to as $dtw^1 \cdot \rho$ and $dtw^2 \cdot \rho$ in Figures 12(a) and 12(b)) to all embedded $write_{CALL}/read_{CALL}$ invocations: at query runtime, this ensures that protocol entries can be unambiguously associated with the current UDF call in progress.

3.2 Provenance: Priceless, but Not for Free

It is expected that provenance derivation introduces tangible query runtime overhead:

- provenance derivation requires *two* query evaluation phases,
- the Phases 1 and 2 write and then read the interim protocol, and
- dependency sets may be of substantial cardinality and will certainly exceed the size of regular 1NF table cell values.

Below, we quantify this overhead for the recursive UDF dtw of Figure 9. We will also learn, however, that the availability of the protocol can be a true advantage.

These experiments were performed on PostgreSQL v14.1, hosted on a Linux-based Intel Xeon™ machine with 72 GB of RAM, 32 GB of which were used for PostgreSQL's buffer. We extended the stack size to 64 MB to support UDFs that recurse deeply. Dependency sets were represented using bit sets which efficiently support duplicate

elimination and the central \cup operation (the associated PostgreSQL extension is discussed in [12]). We report the average runtime of five query runs (worst and best runs discarded).

Recursive call counts. UDF dtw indeed constitutes a true stress test for recursive query evaluation in general and provenance derivation in particular. The UDF features three-fold recursion which leads to a number of recursive calls that is exponential in the length of the time series [6]: the naive, non-memoizing implementation of dtw in Figure 9 already performs about 800 000 such calls for time series of length 10. In all plots of Figure 13, the x-axis range over the recursive call count. The 10 data points of each curve report on a series of dtw runs over time series lengths of 1 ... 10. The largest input instance consists of $10 \times 10 = 100$ rows stored in the `distances` table, effectively the cross product of both time series. (Note that this instance already leads to millions of recursive UDF invocations.)

Runtime overhead of provenance derivation. In the proposed approach, provenance derivation for a SQL query q is complete only once both rewritten variants q^1 and q^2 have been evaluated. The interim protocol forces the Phases 1 and 2 to be performed sequentially. If $t(q)$ denotes the runtime required to evaluate query q , we thus observe a provenance derivation overhead of $(t(q^1) + t(q^2))/t(q)$ relative to the original subject query q .

The curves of Figure 13(a) report on this overhead for a SQL query that invokes UDF dtw (and performs no other computation). If we derive *where-* as well as *why*-provenance (see the \bullet curve), we observe slowdowns that range between $3.2\times$ and $13.4\times$. As we travel the x-axis left to right and thus perform more and more recursive calls while we process longer time series, the overhead increases: each call requires protocol access in both phases ($write_{CALL}/read_{CALL}$), protocol size grows, and more dependency set operations have to be performed. Beyond 10 000 UDF calls, the slowdown reaches a plateau indicating that the cost of UDF invocation itself becomes significant (q pays this price just like q^1 and q^2 do).

If we derive *where*-provenance only (curve Δ), the overhead never exceeds $4.0\times$. We certainly expect the construction of smaller dependency sets (all cell identifiers p_j are missing), but this alone cannot explain the significant overhead reduction. Since q and q^1 remain identical, savings must occur in Phase 2 and q^2 . That is exactly what we find below.

Wall clock times. The original query q runs fastest, q^1 requires additional time for protocol access, on top of that q^2 juggles potentially large dependency sets and will be slowest. This is just what the plot of query execution times in Figure 13(b) shows. Yet, the curves \bullet , \circ , and \times largely run in parallel which indicates that PostgreSQL is able to find query plans for the q^1 and q^2 that—despite the side-effecting operations of protocol writing and reading—exhibit runtime characteristics comparable to the plans for q (this also explains the overhead plateau in Figure 13(a)).

Interestingly, the protocol can save Phase 2 and thus q^2 from the exponential complexity of three-fold recursion if we derive *where*-provenance only (see curve \triangle). In the original query q as well q^1 , the evaluation of the three-way minimum $\text{LEAST}(e_1, e_2, e_3)$ requires the evaluation of *all* arguments e_i —for dtw, this leads to three recursive calls (see Lines 16–18 in Figure 12(a)). When q^2 evaluates LEAST , however, the protocol already reveals which of the three branches (say $e_m, m \in \{1, 2, 3\}$) yielded the minimum. (Internally, LEAST uses $\text{write}_{\text{CASE}}/\text{read}_{\text{CASE}}$; refer to Lines 3–7 in Figures 12(a) and 12(b) but disregard the terms $Y(\cdot)$). Pursuing branch e_m only leads to a single recursive call, effectively turning dtw into a linear recursive function. Indeed, we find q^2 to use negligible time of no more than 26 ms across all time series lengths. Similar, yet less dramatic, Phase 2 savings have also been found in [12].

Protocol size. The interim protocol provides essential glue between Phases 1 and 2. All timings reported in this section are based on a tabular implementation of the protocol that $\text{write}_{\square}/\text{read}_{\square}$ write to and read from. Figure 13(c) reports that both side-effecting functions move between 120 kB and 160 MB of protocol data (between 16 and 3 400 000 protocol entries), depending on the number of recursive dtw calls.

4 MORE RELATED WORK

We view this work as a proof of the versatility of the two-phase approach to cell-level provenance derivation for SQL. Our focus in the companion paper [12] has been on the derivation of *where*- and *why*-provenance for a practically relevant (yet non-recursive) dialect of SQL, including correlation, grouping and aggregation, or window functions. The present findings blend with that work and also fit with the computation of *how*-provenance for SQL [13].

With *Perm* [9, 10], we share the goal to bring data provenance to rich SQL subsets. *Perm* derives provenance for aggregates, set operations, or correlated subqueries, but has not embraced recursion. Its implementation is based on a PostgreSQL kernel that has been modified to support data provenance. We have deliberately designed Phases 1 and 2 to build on non-invasive, language-level query transformations that apply to a variety of RDBMS backends. In this respect, we are closer to Glavic’ *GProM* [1].

Provenance support for *Datalog*—in a sense *the* prototypical relational query language to support recursion—has already been established by the early works of the field, notably in the pervasive semirings approach [11]. These foundations are turned into an efficient implementation by [15] which instruments the internal relational algebra machine program for a given Datalog query to derive its semiring-based provenance. Indeed, query instrumentation is a common theme that spans our efforts, the just mentioned work of Senellart and his colleagues, as well as the effort of Deutch *et*

al. [4]. The latter realizes a notion of *how-provenance* for recursive Datalog in which the relevant intensional facts (or: derivation trees) of a query are computed. The returned provenance can be queried—also to control its potentially huge size—an idea that we fully subscribe to: much like *Perm*, we return a relational representation of dependency sets (recall Figure 4) that is subject to exploration via SQL itself.

Recursive queries may yield infinite provenance if *all* possible derivations of an output row are considered [11]. Deutch *et al.* employ *Boolean circuits* [5] to mitigate this. In contrast, the present approach derives a *single* dependency set per output cell.

5 WRAP-UP

Recursion is key to bring complex computation close to database-resident data. We pursue this thread of work to bring the derivation of data provenance in line with the recursive SQL query constructs found “in the wild,” recursive common table expressions and user-defined functions, in particular. Provenance derivation is especially valuable when the runtime behavior of a query can be intricate (or even puzzling), *e.g.*, in the presence of fixpoint computation or *n*-fold, non-tail recursive functions.

While our discussion has revolved around CTE *bom* and UDF *dtw*, we have prepared a GitHub repository¹ that holds the original as well as instrumented SQL sources for a series of further recursive queries. Applications include parsing, string matching, 2D pixel processing (*Marching Squares*), and graph algorithms. All queries are ready to run, include sample data, and can be evaluated on any contemporary PostgreSQL instance.

REFERENCES

- [1] S.A. Arab, S. Feng, B. Glavic, S. Lee, X. Niu, and Q. Zeng. 2018. GProM - A Swiss Army Knife for Your Provenance Needs. *IEEE Data Eng. Bulletin* 41, 1 (2018).
- [2] D.J Berndt and J. Clifford. 1994. Using Dynamic Time Warping to Find Patterns in Time Series. In *AAAI Workshop*. Seattle, WA, USA.
- [3] J. Cheney, L. Chiticariu, and W.-C. Tan. 2007. Provenance in Databases: Why, How, and Where. *Foundations and Trends in Databases* 1, 4 (2007).
- [4] D. Deutch, A. Gilad, and Y. Moskovitch. 2018. Efficient Provenance Tracking for Datalog Using Top-k Queries. *VLDB J.* 27, 2 (2018).
- [5] D. Deutch, T. Milo, S. Roy, and V. Tannen. 2014. Circuits for Datalog Provenance. In *Proc. ICDT*. Athens, Greece.
- [6] C. Duta and T. Grust. 2020. Functional-Style SQL UDFs with a Capital ‘F’. In *Proc. SIGMOD*. Portland, OR, USA.
- [7] A. Eisenberg and J. Melton. 1999. SQL:1999, Formerly Known as SQL3. *ACM SIGMOD Record* 28, 1 (March 1999).
- [8] S.J. Finkelstein, N. Mattos, I. Mumick, and H. Pirahesh. 1996. Expressive Recursive Queries in SQL. Joint Technical Committee ISO/IEC JTC 1/SC 21 WG 3, Document X3H2-96-075r1.
- [9] B. Glavic and G. Alonso. 2009. PERM: Processing Provenance and Data on the Same Data Model through Query Rewriting. In *Proc. ICDE*. Shanghai, China.
- [10] B. Glavic and G. Alonso. 2009. Provenance for Nested Subqueries. In *Proc. EDBT*. Saint Petersburg, Russia.
- [11] T.J. Green, G. Karvounarakis, and V. Tannen. 2007. Provenance Semirings. In *Proc. PODS*. Beijing, China.
- [12] T. Müller, B. Dietrich, and T. Grust. 2018. You Say ‘What’, I Hear ‘Where’ and ‘Why’? (Mis-)Interpreting SQL to Derive Fine-Grained Provenance. *PVLDB* 11, 11 (2018).
- [13] T. Müller and P. Engel. 2022. How, Where, and Why Data Provenance Improves Query Debugging. In *Proc. ICDE*. Kuala Lumpur, Malaysia.
- [14] PostgreSQL [n.d.]. *PostgreSQL 14 Documentation*. <http://www.postgresql.org/docs/14/>.
- [15] Y. Ramusat, S. Maniu, and P. Senellart. 2021. A Practical Dynamic Programming Approach to Datalog Provenance Computation. *CoRR* (2021).

¹<https://github.com/DBatUTuebingen/provenance-for-recursive-queries>