

# Incremental Maintenance of DatalogMTL Materialisations

Kaiyue Zhao<sup>1\*</sup>, Dingqi Chen<sup>1\*</sup>, Shaoyu Wang<sup>1, 2\*</sup>, Pan Hu<sup>1†</sup>

<sup>1</sup>School of Computer Science, Shanghai Jiao Tong University, China

<sup>2</sup>Department of Computer Science, University of Oxford, UK

## Abstract

DatalogMTL extends the classical Datalog language with metric temporal logic (MTL), enabling expressive reasoning over temporal data. While existing reasoning approaches, such as materialisation-based and automata-based methods, offer soundness and completeness, they lack support for handling efficient dynamic updates—a crucial requirement for real-world applications that involve frequent data updates. In this work, we propose DRed<sub>MTL</sub>, an incremental reasoning algorithm for DatalogMTL with bounded intervals. Our algorithm builds upon the classical Delete/Rederive (DRed) algorithm, which incrementally updates the materialisation of a Datalog program. Unlike a Datalog materialisation which is in essence a finite set of facts, a DatalogMTL materialisation has to be represented as a finite set of facts plus periodic intervals indicating how the full materialisation can be constructed through unfolding. To cope with this, our algorithm is equipped with specifically designed operators to efficiently handle such periodic representations of DatalogMTL materialisations. We have implemented this approach and tested it on several publicly available datasets. Experimental results show that DRed<sub>MTL</sub> often significantly outperforms rematerialisation, sometimes by orders of magnitude.

## Code, datasets and instructions —

[github.com/Horizon12275/DREDmtl-for-DatalogMTL](https://github.com/Horizon12275/DREDmtl-for-DatalogMTL)

## Extended version with full proof —

[arxiv.org/abs/2511.12169](https://arxiv.org/abs/2511.12169)

## Introduction

DatalogMTL extends the well-known rule language Datalog (Ceri, Gottlob, and Tanca 1989) with metric temporal logic (MTL) (Koymans 1990). It has found applications in various domains, including ontology-based query answering (Brandt et al. 2018; Güzel Kalayci et al. 2018), stream reasoning (Wałęga et al. 2023; Wałęga, Kaminski, and Cuenca Grau 2019), and temporal reasoning in the financial sector (Colombo et al. 2023; Nissl and Sallinger 2022; Mori et al. 2022), among others.

To showcase the capabilities of DatalogMTL, consider an industrial application scenario in which DatalogMTL has

been applied by our research group to automatically detect anomaly of power transformers. More concretely, gas concentration data is first collected in real time using existing gas-in-oil sensors and then fed into a DatalogMTL rule engine, which fires alarms whenever appropriate. As an example, the following rule shows how oil thermal faults can be detected:

$$\begin{aligned} OTF(x) \leftarrow & HasEthylene(x, y) \\ & \wedge HasEthane(x, z) \\ & \wedge \diamond_{[0,10]} AboveThirty(y) \\ & \wedge \diamond_{[0,10]} AboveSeventy(z) \end{aligned} \quad (1)$$

This rule states that at any time point  $t$ , if ethylene and ethane have both been detected in the transformer oil, and their gas concentration values surpassed 30 ppm and 70 ppm at any time point in the past ten minutes, respectively, then an oil thermal fault (OTF) is detected.

Reasoning in DatalogMTL can be implemented using top-down or bottom-up approaches, or a combination of the two. One typical top-down approach is based on Büchi automata: it ensures correctness but incurs high reasoning costs. Therefore, efforts have been made to devise efficient bottom-up (or materialisation-based) approach for DatalogMTL reasoning. The vadalog system (Bellomarini, Nissl, and Sallinger 2022) implements such an approach, but it may not terminate due to recursion. The MeTeoR system (Wang et al. 2022) combines bottom-up and top-down approaches, but only resorts to the top-down approach when necessary. More recently, magic set rewriting, which simulates top-down evaluation via bottom-up reasoning, has been extended to support DatalogMTL reasoning (Wang et al. 2025). While materialisation-based approaches are popular for DatalogMTL reasoning, to the best of our knowledge, no incremental maintenance algorithm for DatalogMTL reasoning has been developed: when the set of explicit facts change, the above systems have no choice but to recompute the materialisation from scratch.

For plain Datalog, many incremental materialisation maintenance algorithms have been developed, including Delete/Rederive (DRed) and its variants (Gupta, Mumick, and Subrahmanian 1993; Staudt and Jarke 1996; Ren and Pan 2011; Urbani et al. 2013; Hu, Motik, and Horrocks 2018), Backward/Forward (B/F) (Motik et al. 2015),

\*These authors contributed equally.

†Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

FBF (Motik et al. 2019), among others. However, adapting an incremental materialisation maintenance algorithm for Datalog to support DatalogMTL reasoning is nontrivial. Unlike Datalog, DatalogMTL supports recursion over time, which easily leads to unbounded time intervals, making termination problematic. As a result, changes to the dataset may not only affect immediate derivations but also propagate across time through periodic patterns. Although it is known that under certain restrictions, materialisations in DatalogMTL exhibit repeating structures that can be finitely represented, how to correctly update such repeating structures is still highly challenging, especially due to the complex interplay between consequence propagation and termination condition checks.

In this paper, building upon the well-known DRed algorithm, we propose DRed<sub>MTL</sub>. It replaces standard Datalog materialisation maintenance with a novel set of operations over finite representations of DatalogMTL materialisations called *periodic materialisations*. Periodic materialisations compactly encode infinite sets of temporal facts by capturing their recurring periodic patterns, which allows us to reason over infinite time domains using finite representations.

To facilitate reasoning over periodic materialisations, we devised a novel seminaïve evaluation operator specifically designed to efficiently handle the propagation of facts in the new DatalogMTL setting. In addition, we developed a new period identification algorithm that effectively and correctly guarantees termination. Our experimental evaluation demonstrates that, compared to rematerialisation from scratch, our approach achieves significant performance improvements, especially for small updates.

## Preliminaries

In this section, we first briefly recapitulate the syntax and semantics of DatalogMTL. We focus on DatalogMTL with bounded intervals, as this restriction enables materialisation-based reasoning. Then, we briefly discuss how the materialisation of a DatalogMTL program can be finitely represented: our incremental maintenance algorithm will have to incrementally update such finite representations in response to changes in the explicitly given data.

**Syntax of DatalogMTL** Throughout the paper, we assume that the timeline consists of rational numbers. A time interval is a set of continuous time points  $\varrho$  of the form  $\langle t_1, t_2 \rangle$ , where  $t_1 \in \mathbb{Q} \cup \{-\infty\}$ ,  $t_2 \in \mathbb{Q} \cup \{\infty\}$ ,  $\langle$  is [ or ( and likewise  $\rangle$  is ] or ). An interval is bounded if both of its endpoints are rational numbers, i.e., neither  $\infty$  nor  $-\infty$ . When it is clear from the context, we may abuse the distinction between intervals (i.e. sets of time points) and their representation  $\langle t_1, t_2 \rangle$ . If  $\varrho = \langle t_1, t_2 \rangle$ , let  $\varrho^- = t_1$  and  $\varrho^+ = t_2$ .

A term is either a variable or a constant. A relational atom is an expression  $R(\mathbf{t})$ , where  $R$  is a predicate and  $\mathbf{t}$  is a tuple of terms of arity matching that of  $R$ . Metric atoms extend relational atoms by allowing operators from metric temporal logic (MTL), namely  $\boxplus_\varrho$ ,  $\boxminus_\varrho$ ,  $\boxtimes_\varrho$ ,  $\boxdiv_\varrho$ ,  $\mathcal{U}_\varrho$ , and  $\mathcal{S}_\varrho$ , where  $\varrho$  is an interval. Formally, metric atoms,  $M$ , are generated by

$\mathcal{I}, t \models \top$	for every $t \in \mathbb{Q}$
$\mathcal{I}, t \models \perp$	for no $t \in \mathbb{Q}$
$\mathcal{I}, t \models \boxplus_\varrho M$	iff $\mathcal{I}, t_1 \models M$ for all $t_1$ s.t. $t_1 - t \in \varrho$
$\mathcal{I}, t \models \boxminus_\varrho M$	iff $\mathcal{I}, t_1 \models M$ for all $t_1$ s.t. $t - t_1 \in \varrho$
$\mathcal{I}, t \models \boxtimes_\varrho M$	iff $\mathcal{I}, t_1 \models M$ for some $t_1$ s.t. $t_1 - t \in \varrho$
$\mathcal{I}, t \models \boxdiv_\varrho M$	iff $\mathcal{I}, t_1 \models M$ for some $t_1$ s.t. $t - t_1 \in \varrho$
$\mathcal{I}, t \models M_2 \mathcal{U}_\varrho M_1$	iff $\mathcal{I}, t_1 \models M_1$ for some $t_1$ s.t. $t_1 - t \in \varrho$ , and $\mathcal{I}, t_2 \models M_2$ for all $t_2 \in (t, t_1)$
$\mathcal{I}, t \models M_2 \mathcal{S}_\varrho M_1$	iff $\mathcal{I}, t_1 \models M_1$ for some $t_1$ s.t. $t - t_1 \in \varrho$ , and $\mathcal{I}, t_2 \models M_2$ for all $t_2 \in (t_1, t)$

Table 1: Semantics for ground metric atoms

the grammar

$$M ::= \perp \mid \top \mid R(\mathbf{t}) \mid \boxplus_\varrho M \mid \boxminus_\varrho M \mid \boxtimes_\varrho M \mid \boxdiv_\varrho M \mid MU_\varrho M \mid MS_\varrho M,$$

where  $\top$  and  $\perp$  are the constants representing truth and falsehood, respectively, and  $\varrho$  is any arbitrary interval containing only nonnegative rationals. A DatalogMTL rule,  $r$ , is of the form

$$M' \leftarrow M_1 \wedge M_2 \wedge \dots \wedge M_n, \quad \text{for } n \geq 1,$$

where each  $M_i$  is a metric atom and  $M'$  is a metric atom not mentioning  $\perp$ ,  $\boxtimes$ ,  $\boxdiv$ ,  $\mathcal{U}$ , and  $\mathcal{S}$ . Metric atom  $M'$  and the set  $\{M_i \mid i \in \{1, \dots, n\}\}$  are the head and body of  $r$ , denoted as  $head(r)$  and  $body(r)$ , respectively.

A rule  $r$  is safe if each variable in its head appears also in its body. A program  $\Pi$  is a finite set of safe rules. A substitution  $\sigma$  is a mapping of finitely many variables to constants. For  $\alpha$  an expression (e.g., an atom, a rule, or a program thereof),  $\alpha\sigma$  is the result of replacing each occurrence of a variable  $x$  in  $\alpha$  with  $\sigma(x)$ , if the latter is defined. An expression is ground if it mentions no variable. A fact is of the form  $R(\mathbf{t})@_\varrho$ , where  $R(\mathbf{t})$  is a ground relational atom,  $\varrho$  is an interval, and  $@$  indicates that the preceding atom holds over the time interval that follows. A dataset  $\mathcal{D}$  is a finite set of facts. If all the intervals a dataset (resp., a program) mentions are bounded, then the dataset (resp., the program) is bounded. Our work focuses on bounded datasets and programs.

**Semantics of DatalogMTL** A DatalogMTL interpretation  $\mathcal{I}$  is a function that maps each time point  $t \in \mathbb{Q}$  to a set of ground relational atoms (essentially to atoms that hold at  $t$ ). For a time point  $t \in \mathbb{Q}$ , if  $R(\mathbf{t})$  belongs to this set, we write  $\mathcal{I}, t \models R(\mathbf{t})$ . This extends to the ground metric atoms as presented in Table 1.

An interpretation  $\mathcal{I}$  satisfies a fact  $R(\mathbf{t})@_\varrho$ , denoted as  $\mathcal{I} \models R(\mathbf{t})@_\varrho$ , if  $\mathcal{I}, t \models R(\mathbf{t})$  for each  $t \in \varrho$ . An interpretation  $\mathcal{I}$  is a model of a dataset  $\mathcal{D}$ , written  $\mathcal{I} \models \mathcal{D}$ , if it satisfies every fact in  $\mathcal{D}$ . An interpretation  $\mathcal{I}$  satisfies a ground rule  $r_0$  if  $\mathcal{I}, t \models body(r_0)$  implies  $\mathcal{I}, t \models head(r_0)$  for each  $t \in \mathbb{Q}$ . A ground rule  $r_0$  is an instance of a rule  $r$  if there

is a substitution  $\sigma$  such that  $r_0 = r\sigma$ . An interpretation  $\mathcal{I}$  satisfies a rule  $r$ , if it satisfies all ground instances of  $r$ , and it is a model of a program  $\Pi$ , if it satisfies all rules in  $\Pi$ .  $\mathcal{I}$  is a model of a pair  $(\Pi, \mathcal{D})$  if  $\mathcal{I}$  is both a model of  $\mathcal{D}$  and a model of  $\Pi$ . A program-dataset pair  $(\Pi, \mathcal{D})$  entails a fact  $R(t)@_\varrho$ , written  $(\Pi, \mathcal{D}) \models R(t)@_\varrho$ , if all models of  $(\Pi, \mathcal{D})$  satisfy  $R(t)@_\varrho$ . An interpretation  $\mathcal{I}$  contains another interpretation  $\mathcal{I}'$ , written  $\mathcal{I}' \subseteq \mathcal{I}$ , if  $\mathcal{I}, t \models R(t)$  implies  $\mathcal{I}', t \models R(t)$  for each ground relational atom  $R(t)$  and each  $t \in \mathbb{Q}$ ; moreover,  $\mathcal{I} = \mathcal{I}'$  if they contain each other. An interpretation  $\mathcal{I}$  is the least in a set of interpretations  $\mathcal{I}$ , if  $\mathcal{I} \in \mathcal{I}$ , and  $\forall \mathcal{I}' \in \mathcal{I}, \mathcal{I} \subseteq \mathcal{I}'$ . Similarly,  $\mathcal{I}$  is the greatest in  $\mathcal{I}$ , if  $\mathcal{I} \in \mathcal{I}$ , and  $\forall \mathcal{I}' \in \mathcal{I}, \mathcal{I}' \subseteq \mathcal{I}$ . Each dataset  $\mathcal{D}$  has a unique least interpretation  $\mathcal{I}_{\mathcal{D}}$  such that  $\mathcal{I}_{\mathcal{D}}$  is the least in the set of all models of  $\mathcal{D}$ . For interpretations  $\mathcal{I}_1, \mathcal{I}_2$ , and  $\mathcal{I}_3$ , we write  $\mathcal{I}_3 = \mathcal{I}_1 \cup \mathcal{I}_2$  if  $\mathcal{I}_3$  is the least in  $\{\mathcal{I} \mid \mathcal{I}_1 \subseteq \mathcal{I} \text{ and } \mathcal{I}_2 \subseteq \mathcal{I}\}$ , and we write  $\mathcal{I}_3 = \mathcal{I}_1 \cap \mathcal{I}_2$ , if  $\mathcal{I}_3$  is the greatest in  $\{\mathcal{I} \mid \mathcal{I} \subseteq \mathcal{I}_1 \text{ and } \mathcal{I} \subseteq \mathcal{I}_2\}$ . The empty interpretation  $\mathcal{I}_{\emptyset}$  maps  $t$  to  $\emptyset$  for each  $t \in \mathbb{Q}$ , and it is contained by any interpretation. Finally, we write  $\mathcal{I}_3 = \mathcal{I}_1 - \mathcal{I}_2$ , if  $\mathcal{I}_3$  is the greatest in  $\{\mathcal{I} \mid \mathcal{I} \subseteq \mathcal{I}_1, \mathcal{I} \cap \mathcal{I}_2 = \mathcal{I}_{\emptyset}\}$ .

**Materialisation for DatalogMTL** We now briefly discuss how materialisation-based reasoning works for DatalogMTL with bounded intervals. To this end, we first introduce a few additional notations that will facilitate our discussion. For  $\mathcal{D}$  a dataset,  $t_{\mathcal{D}}^-$  and  $t_{\mathcal{D}}^+$  denote the minimal and maximal interval endpoints appearing in  $\mathcal{D}$ , respectively, and  $t_{\mathcal{D}}^- = t_{\mathcal{D}}^+ = 0$  if  $\mathcal{D}$  mentions no numbers. Moreover, the depth of a rule  $r$ , written  $\text{depth}(r)$ , is defined as the sum of right endpoints of all intervals appearing in the operators of  $r$ , and  $\text{depth}(r) = 0$  if  $r$  mentions no intervals; for  $\Pi$  a program,  $\text{depth}(\Pi)$  is the maximum depth of its rules.

The immediate consequence operator  $T_{\Pi}$  for a program  $\Pi$  is the function that maps an interpretation  $\mathcal{I}$  to the least interpretation  $T_{\Pi}(\mathcal{I})$  containing  $\mathcal{I}$  and satisfying the following: for each  $r_0$  a ground rule instance of a rule in  $\Pi$  and each  $t$  a time point,  $\mathcal{I}, t \models \text{body}(r_0)$  implies  $\mathcal{I}, t \models \text{head}(r_0)$ . For a program  $\Pi$  and a dataset  $\mathcal{D}$ , a transfinite sequence of interpretations  $T_{\Pi}^{\alpha}(\mathcal{I}_{\mathcal{D}})$  can be defined for ordinals  $\alpha$  by successively applying  $T_{\Pi}$ , the immediate consequence operator for  $\Pi$ , to  $\mathcal{I}_{\mathcal{D}}$ , the unique least model for  $\mathcal{D}$ , as follows: (i)  $T_{\Pi}^0(\mathcal{I}_{\mathcal{D}}) = \mathcal{I}_{\mathcal{D}}$ , (ii)  $T_{\Pi}^{\alpha+1}(\mathcal{I}_{\mathcal{D}}) = T_{\Pi}(T_{\Pi}^{\alpha}(\mathcal{I}_{\mathcal{D}}))$  for  $\alpha$  an ordinal, and (iii)  $T_{\Pi}^{\alpha}(\mathcal{I}_{\mathcal{D}}) = \bigcup_{\beta < \alpha} T_{\Pi}^{\beta}(\mathcal{I}_{\mathcal{D}})$  for  $\alpha$  a limit ordinal; the canonical model  $\mathcal{C}_{\Pi, \mathcal{D}}$  of  $\Pi$  and  $\mathcal{D}$  is the interpretation  $T_{\Pi}^{\omega_1}(\mathcal{I}_{\mathcal{D}})$  with  $\omega_1$  the first uncountable ordinal. In fact,  $\mathcal{C}_{\Pi, \mathcal{D}}$  is the least model of  $\Pi$  and  $\mathcal{D}$  (Brandt et al. 2017).

For an interpretation  $\mathcal{I}$  and an interval  $\varrho$ , the projection  $\mathcal{I}|_{\varrho}$  of  $\mathcal{I}$  over  $\varrho$  is the interpretation that coincides with  $\mathcal{I}$  on  $\varrho$  and maps all relational atoms to false outside  $\varrho$ . Moreover, an interpretation  $\mathcal{I}$  is a shift of another interpretation  $\mathcal{I}'$ , if there is a rational number  $s$  such that for each  $R(t)$  a ground relational atom and each  $t$  a time point,  $\mathcal{I}', t \models R(t)$  implies  $\mathcal{I}, t + s \models R(t)$ , and vice versa. Finally, for a given program  $\Pi$  and a dataset  $\mathcal{D}$ , we are not interested in dealing with arbitrary rational number time points on the time line; instead, it is sufficient to handle time points that are somewhat related to those appearing in  $\Pi$  and  $\mathcal{D}$ . This is formalised by the notion of  $(\Pi, \mathcal{D})$ -ruler. Concretely,  $(\Pi, \mathcal{D})$ -ruler is the set of

time points of the value  $t + i \times \text{div}(\Pi)$ , where  $t$  is an endpoint mentioned in  $\mathcal{D}$  and  $i$  is an integer, and  $\text{div}(\Pi) = 1/k$ , with  $k$  being the product of all denominators in the rational endpoints mentioned in  $\Pi$ ; for generality,  $k = 1$  and  $\text{div}(\Pi) = 1$  if  $\Pi$  has no mention of rational endpoints.

We are now ready to define the notions of saturated interpretation and unfolding, which are key components for finitely representing the materialisation of a bounded program-dataset pair.

**Definition 1.** For a program  $\Pi$  and a dataset  $\mathcal{D}$ , interpretation  $T_{\Pi}^k(\mathcal{I}_{\mathcal{D}})$  is saturated if there exist closed intervals  $\varrho_1, \varrho_2, \varrho_3$  and  $\varrho_4$  of length  $2\text{depth}(\Pi)$ , whose endpoints are located on the  $(\Pi, \mathcal{D})$ -ruler and satisfy  $\varrho_1^+ < \varrho_2^+ < t_{\mathcal{D}}^-$  and  $t_{\mathcal{D}}^+ < \varrho_3^- < \varrho_4^-$ , and such that the following properties hold:

- $T_{\Pi}^k(\mathcal{I}_{\mathcal{D}})$  satisfies  $\Pi$  in  $[\varrho_1^-, \varrho_4^+]$ ;
- $T_{\Pi}^k(\mathcal{I}_{\mathcal{D}})|_{\varrho_1}$  and  $T_{\Pi}^k(\mathcal{I}_{\mathcal{D}})|_{\varrho_3}$  are shifts of  $T_{\Pi}^k(\mathcal{I}_{\mathcal{D}})|_{\varrho_2}$  and  $T_{\Pi}^k(\mathcal{I}_{\mathcal{D}})|_{\varrho_4}$ , respectively.

$[\varrho_1^-, \varrho_2^-)$  and  $(\varrho_3^+, \varrho_4^+]$  are often referred to as the left period, written  $\varrho_{\text{left}}$ , and the right period, written  $\varrho_{\text{right}}$ , of the interpretation  $T_{\Pi}^k(\mathcal{I}_{\mathcal{D}})$ , respectively.

**Definition 2.** The  $(\varrho_{\text{left}}, \varrho_{\text{right}})$ -unfolding of a saturated interpretation  $T_{\Pi}^k(\mathcal{I}_{\mathcal{D}})$  with periods  $(\varrho_{\text{left}}, \varrho_{\text{right}})$  is the interpretation  $\mathcal{C}$  such that:

- $\mathcal{C}|_{[\varrho_{\text{left}}^-, \varrho_{\text{right}}^+]} = T_{\Pi}^k(\mathcal{I}_{\mathcal{D}})|_{[\varrho_{\text{left}}^-, \varrho_{\text{right}}^+]}$ ,
- $\mathcal{C}|_{\varrho_{\text{left}} - n \cdot |\varrho_{\text{left}}|}$  is a shift of  $T_{\Pi}^k(\mathcal{I}_{\mathcal{D}})|_{\varrho_{\text{left}}}$ , for any  $n \in \mathbb{N}$ ,
- $\mathcal{C}|_{\varrho_{\text{right}} + n \cdot |\varrho_{\text{right}}|}$  is a shift of  $T_{\Pi}^k(\mathcal{I}_{\mathcal{D}})|_{\varrho_{\text{right}}}$ , for any  $n \in \mathbb{N}$ .

It has been shown that for an arbitrary pair of bounded program and dataset  $(\Pi, \mathcal{D})$ , there exists  $k \in \mathbb{N}$  and intervals  $\varrho_{\text{left}}$  and  $\varrho_{\text{right}}$  such that  $T_{\Pi}^k(\mathcal{I}_{\mathcal{D}})$  is a saturated interpretation with periods  $(\varrho_{\text{left}}, \varrho_{\text{right}})$ , and the  $(\varrho_{\text{left}}, \varrho_{\text{right}})$ -unfolding of  $T_{\Pi}^k(\mathcal{I}_{\mathcal{D}})$  coincides with  $\mathcal{C}_{\Pi, \mathcal{D}}$ , the canonical model of  $\Pi$  and  $\mathcal{D}$  (Wałęga et al. 2023). Intuitively speaking, a saturated interpretation finitely represents the canonical model of a program and a dataset. The goal of this work is to incrementally update a saturated interpretation in response to changes in the explicitly given dataset  $\mathcal{D}$ .

**The Delete/Rederive Algorithm** To incrementally update saturated interpretations for DatalogMTL, we draw inspirations from the Delete/Rederive (DRed) algorithm (Gupta, Mumick, and Subrahmanian 1993; Staudt and Jarke 1996), a well-known technique for maintaining Datalog materialisations. Given a Datalog program, a set of explicitly given facts, the original materialisation, and sets of facts to remove from and to add to the given facts, the DRed algorithm updates the materialisation to reflect changes in the explicitly given facts, without recomputing it from scratch. More specifically, the algorithm operates in three stages: in the overdeletion stage, the algorithm eagerly identifies all facts that depend on the set of deleted facts; it then enters the rederivation stage, in which it recognises which of the overdeleted facts can be rederived in one step; finally, during insertion, the algorithm computes the consequences of both the rederived facts and the newly inserted ones.

## Motivation

Before we present the technical details of our incremental update approach for DatalogMTL, in this section, we discuss at a high level challenges that need to be addressed in devising such an incremental update approach. We also provide key insights behind our incremental update algorithm. Finally, we provide a concrete example that highlights the benefits of applying our incremental maintenance algorithm for DatalogMTL reasoning compared with recomputing the materialisation from scratch.

In extending DRed for Datalog to support DatalogMTL reasoning, two major challenges arise. First, DRed relies on the seminaïve evaluation strategy applied in both the overdeletion and the insertion stages to be efficient in dealing with updates. More specifically, in each round of rule application, at least one body atom is required to be evaluated in the set of facts deleted/inserted in the previous round; this strategy ensures that only consequences dependent on the deleted/inserted facts are considered. Although seminaïve evaluation has been considered for DatalogMTL materialisation, its interplay with the construction of saturated interpretations in the context of DatalogMTL reasoning has not been studied, let alone the application of seminaïve evaluation strategy in DatalogMTL materialisation maintenance. Second, for the Datalog setting, the bulk of the work is devoted to iterative rule application, so it is sufficient to make rule application process ‘incremental’ in order to obtain an algorithm that is efficient for updates; however, in the DatalogMTL setting, a significant amount of time needs to be dedicated to periods identification (i.e., termination checks). Intuitively, after each round of rule application, the period identification procedure enumerates all pairs of intervals of a certain length, and then it compares pairwise the facts in these intervals to detect the presence of a repeating pattern. Once two intervals coincide in their content, a period is identified. As such, the cost of this procedure naturally depends on the number of facts inside these intervals. Therefore, it is essential that the periods identification process is also designed to be somewhat ‘incremental’, so that the overall processing time aligns with the size of the updates rather than the size of the entire materialisation; this would require thorough understanding and careful treatment of the periodic structure of DatalogMTL materialisations.

Our incremental maintenance technique addresses both of these two challenges. For the first challenge, we devise a novel seminaïve evaluation operator that is tailored for incremental updates in the DatalogMTL setting: compared with the existing seminaïve operator, our new operator identifies more accurately the consequences affected by the update and also facilitates period identification. To tackle the second challenge, our algorithm is designed to identify periods in the updated facts and their consequences, rather than in the entire materialisation. Typically, deletions and insertions are small compared with the entire materialisation. As such, computing periods over the updated facts tend to be less costly than doing so for the entire materialisation; thus, our approach has the potential to significantly reduce redundant computations, especially in scenarios where updates only affect a limited portion of the materialisation. We use

the following example to highlight the benefits of our approach; throughout the technical section, this example will be expanded to illustrate key operations in our algorithm.

**Example 1.** Let program  $\Pi$  consist of a single rule (2), and let  $E$ ,  $E^-$ , and  $E^+$  be the original dataset, the set of facts to remove from  $E$ , and set of facts to add to  $E$ , respectively. Now consider the materialisation of  $\Pi$  over  $E$ . According to the definition of  $\text{depth}(\Pi)$  and Definition 1, it is easy to verify that  $\text{depth}(\Pi) = 11$ , and that when  $k = 4$ , interpretation  $T_{\Pi}^k(\mathcal{J}_E)$  is saturated, with  $\varrho_{\text{left}} = [-24, -23]$  and  $\varrho_{\text{right}} = [24, 34]$  being possible left and right periods, respectively. To arrive at the conclusion that  $\varrho_{\text{right}} = [24, 34]$ , the materialisation procedure needs to verify that  $T_{\Pi}^k(\mathcal{J}_E) \upharpoonright_{[2,24]}$  is a shift of  $T_{\Pi}^k(\mathcal{J}_E) \upharpoonright_{[12,34]}$ , and this clearly requires comparing  $O(n)$  facts. To summarise, the materialisation procedure requires  $O(n)$  time.

$$\boxplus_{[0,1]} R(x) \leftarrow \boxminus_{[9,10]} R(x) \quad (2)$$

$$E = \{R(a_i)@[0, 1] \mid 0 < i < n\}$$

$$E^- = \{R(a_1)@[0, 1]\}$$

$$E^+ = \{R(a_n)@[0, 1]\}$$

Now consider the removal of  $E^-$  and the insertion of  $E^+$ . If we recompute the materialisation from scratch over the entire set  $(E \setminus E^-) \cup E^+$ , the procedure again requires  $O(n)$  time. In contrast, our incremental update approach restricts the attention to only the updated facts and their consequences. For deletion, our approach tries to identify periodic structure only among facts of the form  $R(a_1)@[\varrho]$ . In our example, after a constant number of rule application, our algorithm will be able to identify the relevant left and right periods by only comparing  $O(1)$  facts instead of  $O(n)$  facts. In other words, (over)deletion requires only  $O(1)$  time. Insertion follows the same principle, and so we omit the details.

## DRed for DatalogMTL

Before we present the details of our algorithm, we first introduce the notion of periodic materialisations, which are finite representations of possibly unbounded models of DatalogMTL programs. Our update algorithm will have to incrementally maintain periodic materialisations in response to changes in the explicitly given data.

**Definition 3.** For  $(\Pi, E)$  a bounded DatalogMTL program-dataset pair, a periodic materialisation of  $(\Pi, E)$  is a triple  $\mathbb{I}$  of the form  $\langle I, \varrho_L, \varrho_R \rangle$ , where  $I$  is a saturated interpretation of  $\Pi$  and  $E$ , and  $\varrho_L$  and  $\varrho_R$  are periods of this saturated interpretation; moreover, the unfolding of  $\mathbb{I}$ , written  $\text{unfold}(\mathbb{I})$ , is defined as the  $(\varrho_L, \varrho_R)$ -unfolding of  $I$ .

By Definition 3 and properties already discussed in the Preliminary Section, if  $\mathbb{I}$  is a periodic materialisation of a program-dataset pair  $(\Pi, E)$ , then  $\text{unfold}(\mathbb{I})$  coincides with  $\mathcal{C}_{\Pi, E}$ , the canonical model of  $\Pi$  and  $E$ . Note that  $\mathbb{I}$  can capture the interpretation of a bounded dataset  $I$ , in which case  $\mathbb{I} = \langle I, \varrho_L, \varrho_R \rangle$  such that  $I \upharpoonright_{\varrho_L} = I \upharpoonright_{\varrho_R} = \mathcal{J}_{\emptyset}$ , and  $\text{unfold}(\mathbb{I}) = I$ .

---

**Algorithm 1: DRED<sub>MTL</sub>( $\Pi, E, \mathbb{I}, E^-, E^+$ )**

---

**Input:** program  $\Pi$ , original dataset  $E$ , periodic materialisation  $\mathbb{I}$  of  $(\Pi, E)$ , dataset to remove  $E^-$ , dataset to insert  $E^+$

```
1  $D := R := A := \emptyset, E^- := (E^- \cap E) \setminus E^+, E^+ := E^+ \setminus E$ 
2 OVERDELETE
3 REDERIVE
4 INSERT
5 procedure OVERDELETE
6    $N_D := E^-$ 
7   loop
8      $\Delta_D := N_D \setminus D$ 
9      $(\varrho_L^D, \varrho_R^D) := \text{PDS}(\Pi, E, \mathbb{I}, \varrho_L, \mathbb{I}, \varrho_R, D, \Delta_D)$ 
10    if  $(\varrho_L^D, \varrho_R^D) \neq (\emptyset, \emptyset)$  then break
11     $N_D := \Pi(\text{unfold}(\mathbb{I}) \setminus D : \Delta_D)$ 
12     $D := D \cup \Delta_D$ 
13   $\mathbb{D} := \langle D, \varrho_L^D, \varrho_R^D \rangle, \mathbb{I} := \mathbb{I} \setminus \mathbb{D}$ 
14 procedure REDERIVE
15   $N_R := \emptyset, k = 1$ 
16  loop
17     $t_L = (\mathbb{I}, \varrho_L)^+ - k \times \max(2\text{depth}(\Pi), |\mathbb{I}, \varrho_L|)$ 
18     $t_R = (\mathbb{I}, \varrho_R)^- + k \times \max(2\text{depth}(\Pi), |\mathbb{I}, \varrho_R|)$ 
19     $N_R := N_R \cup (\text{unfd}(\mathbb{D}))_{[t_L, t_R]} \cap T_\Pi(\text{unfd}(\mathbb{I}))$ 
20     $\Delta_R := N_R \setminus R$ 
21     $(\varrho_L^R, \varrho_R^R) := \text{PDS}(\Pi, E, \mathbb{I}, \varrho_L, \mathbb{I}, \varrho_R, R, \Delta_R)$ 
22    if  $(\varrho_L^R, \varrho_R^R) \neq (\emptyset, \emptyset)$  then break
23     $R := R \cup \Delta_R, k := k + 1$ 
24     $N_R := \Pi(\text{unfold}(\mathbb{I}) \cup R : \Delta_R)$ 
25   $\mathbb{R} := \langle R, \varrho_L^R, \varrho_R^R \rangle, \mathbb{I} := \mathbb{I} \cup \mathbb{R}$ 
26 procedure INSERT
27   $N_A := E^+, E = (E \setminus E^-) \cup E^+$ 
28  loop
29     $\Delta_A := N_A \setminus (\text{unfold}(\mathbb{I}) \cup A)$ 
30     $(\varrho_L^A, \varrho_R^A) := \text{PDS}(\Pi, E, \mathbb{I}, \varrho_L, \mathbb{I}, \varrho_R, A, \Delta_A)$ 
31    if  $(\varrho_L^A, \varrho_R^A) \neq (\emptyset, \emptyset)$  then break
32     $A := A \cup \Delta_A$ 
33     $N_A := \Pi(\text{unfold}(\mathbb{I}) \cup A : \Delta_A)$ 
34   $\mathbb{A} := \langle A, \varrho_L^A, \varrho_R^A \rangle, \mathbb{I} := \mathbb{I} \cup \mathbb{A}$ 
```

---

**Algorithm Overview** Our incremental materialisation maintenance algorithm for DatalogMTL is formalised in Algorithm 1. The algorithm takes as input a program  $\Pi$ , a dataset  $E$ , a periodic materialisation  $\mathbb{I}$  of  $(\Pi, E)$ , a dataset  $E^-$  to be removed from  $E$ , and a dataset  $E^+$  to be added to  $E$ , and it updates  $\mathbb{I}$  such that after the execution of the algorithm,  $\mathbb{I}$  becomes a periodic materialisation of  $(\Pi, (E \setminus E^-) \cup E^+)$ , and this is achieved without recomputing the periodic materialisation from scratch. Similar to the DRed algorithm for plain Datalog, Algorithm 1 consists of three stages, which we outline below.

In the overdeletion stage, the dataset  $D$  is extended with all facts that depend on the deleted facts. The algorithm then

applies the rules of  $\Pi$  iteratively (line 11) until a pair of periods are found (line 9–10). Line 11 makes use of our newly devised seminaive evaluation operator  $\Pi(I : \Delta)$ , which computes the union of all facts derived by  $(r\sigma, t)(I : \Delta)$ , where  $r\sigma$  is a ground rule instance with  $r$  a rule in  $\Pi$  and  $\sigma$  a substitution mapping each variable appearing in  $r$  to a constant appearing in  $I$ , and  $t$  is a rational time point; operator  $(r', t)(I : \Delta)$  where  $r'$  is a ground rule instance and  $t$  is a rational time point is in turn defined as the minimal set of punctual facts  $N$  such that  $\mathfrak{J}_{I, t} \models \text{body}(r')$ ,  $\mathfrak{J}_{I \setminus \Delta, t} \not\models \text{body}(r')$ , and  $\mathfrak{J}_N, t \models \text{head}(r')$ . When  $I$  and  $\Delta$  are finite,  $\Pi(I : \Delta)$  can be efficiently evaluated by instantiating the query from facts inside  $\Delta$ , which is typically much smaller than  $I$ . However, in line 21, by slight abuse of notation, we use  $\text{unfold}(\mathbb{I}) \setminus D$  as an operand of the operator: this does not mean that we have to compute the entire unfolding of  $\mathbb{I}$  prior to the execution of the operator. In contrast, the evaluation is still instantiated from facts in  $\Delta$ , while interpretation  $\mathbb{I}$  can be unfolded lazily as required.

In the rederivation stage, the algorithm recovers the facts that were overdeleted but should actually still hold after the update. It should be noted that overdeleted facts may span the entire timeline, and so it may not be sufficient to recover facts within a fixed interval. Lines 17–19 address this issue, in each round of the loop of lines 16–24, we extend the interval of interest  $[t_L, t_R]$  so that overdeleted facts inside this interval could be successfully recovered (line 19); this is done in parallel to the propagation of already recovered facts (line 24). The loop terminates when a pair of periods are successfully identified (lines 21–22), and the periodic materialisation  $\mathbb{I}$  is updated in line 25.

Insertion is analogous to deletion: the dataset  $A$  is populated with new facts that are derivable from  $E^+$ . In parallel to consequence propagation, the algorithm tries to detect the periodic structure within the dataset  $A$ . Ultimately, the algorithm updates the periodic materialisation in line 34 to incorporate the inserted facts and their consequences.

**Period Identification** The PDS procedure formalised in Algorithm 2 identifies the repeating patterns within a dataset being populated during iterative application of rules. More specifically, it takes as input a program  $\Pi$ , a dataset  $E$ , intervals  $\varrho_L$  and  $\varrho_R$ , the dataset being populated  $D$ , and a set of new facts  $\Delta_D$  that are to be integrated into  $D$  after period detection. The procedure first examines facts in  $\Delta_D$ : if there is a fact in  $\Delta_D$  that overlaps with the interval  $[t_E^-, t_E^+]$ , it means facts inside this interval have not stabilised, and so the procedure terminates. Otherwise,  $u$  and  $v$  will be computed separately (lines 5–6) such that  $[u, v]$  is largest interval containing  $[t_E^-, t_E^+]$  and does not contain a fact satisfied by  $\mathfrak{J}_{\Delta_D}$ . Now if  $[u, v]$  is nonempty, we search for the left and right periods with endpoints on the  $(\Pi, E)$ -ruler in  $\gamma_L = (u, \varrho_L^+ + 2\text{depth}(\Pi))$  and  $\gamma_R = [\varrho_{\text{right}}^- - 2\text{depth}(\Pi), v)$ , respectively (lines 9–16). The procedure returns the pair of periods  $(\varrho_L', \varrho_R')$ , or  $(\emptyset, \emptyset)$  if no valid period is detected at either end (lines 17–18). Intuitively, the algorithm requires  $\varrho_1, \varrho_2$  (resp.,  $\varrho_3, \varrho_4$ ) to be a multiple of  $|\varrho_L|$  (resp.,  $|\varrho_R|$ ) apart so that it would be convenient align the new periods with the given ones.

---

**Algorithm 2:** PDS( $\Pi, E, \varrho_L, \varrho_R, D, \Delta_D$ )

---

**Input:** program  $\Pi$ , dataset  $E$ , intervals  $\varrho_L$  and  $\varrho_R$ , dataset  $D$  of facts derived so far, and dataset  $\Delta_D$  of facts derived in the last round

```
1  $\varrho'_L := \varrho'_R := \emptyset, \quad P := (\Pi, E)$ -ruler
2 foreach  $M\sigma @ \varrho \in \Delta_D$  do
3   if  $\varrho \cap [t_E^-, t_E^+] \neq \emptyset$  then
4     return  $(\emptyset, \emptyset)$ 
5  $u = \max(\{x < t_E^- \mid \exists y, M @ \langle y, x \rangle \in \Delta_D\} \cup \{-\infty\})$ 
6  $v = \min(\{x > t_E^+ \mid \exists y, M @ \langle x, y \rangle \in \Delta_D\} \cup \{\infty\})$ 
7  $\gamma_L := (u, \varrho_L^+ + 2\text{depth}(\Pi))$ 
8  $\gamma_R := [\varrho_R^- - 2\text{depth}(\Pi), v)$ 
9 foreach  $\varrho_1, \varrho_2 \subseteq \gamma_L$  with  $|\varrho_1| = |\varrho_2| = 2\text{depth}(\Pi)$  do
10  if  $\varrho_1^- < \varrho_2^-, \varrho_1^- \in P$  and  $\varrho_2^- \equiv \varrho_1^- \pmod{|\varrho_L|}$  then
11    if  $\mathcal{J}_D|_{\varrho_1}$  is a shift of  $\mathcal{J}_D|_{\varrho_2}$  then
12       $\varrho'_L := [\varrho_1^-, \varrho_2^-]$ 
13 foreach  $\varrho_3, \varrho_4 \subseteq \gamma_R$  with  $|\varrho_3| = |\varrho_4| = 2\text{depth}(\Pi)$  do
14  if  $\varrho_3^+ < \varrho_4^+, \varrho_3^+ \in P$  and  $\varrho_3^+ \equiv \varrho_4^+ \pmod{|\varrho_R|}$  then
15    if  $\mathcal{J}_D|_{\varrho_3}$  is a shift of  $\mathcal{J}_D|_{\varrho_4}$  then
16       $\varrho'_R := (\varrho_3^+, \varrho_4^+]$ 
17 if  $\varrho'_L = \emptyset$  or  $\varrho'_R = \emptyset$  then return  $(\emptyset, \emptyset)$ 
18 else return  $(\varrho'_L, \varrho'_R)$ 
```

---

---

**Algorithm 3:** Implementation of Periodic Minus

---

**Input:** Two periodic materialisations  $\mathbb{D}_1$  and  $\mathbb{D}_2$

```
1 if  $\mathbb{D}_2.\varrho_L$  then
2    $(\mathbb{D}_1, \mathbb{D}_2) := \text{Ext}(\mathbb{D}_1, \mathbb{D}_2, L)$ 
3    $(\mathbb{D}_1, \mathbb{D}_2, \varrho_L) := \text{Aln}(\mathbb{D}_1, \mathbb{D}_2, L)$ 
4 else  $\varrho_L := \mathbb{D}_1.\varrho_L$ 
5 if  $\mathbb{D}_2.\varrho_R$  then
6    $(\mathbb{D}_1, \mathbb{D}_2) := \text{Ext}(\mathbb{D}_1, \mathbb{D}_2, R)$ 
7    $(\mathbb{D}_1, \mathbb{D}_2, \varrho_R) := \text{Aln}(\mathbb{D}_1, \mathbb{D}_2, R)$ 
8 else  $\varrho_R := \mathbb{D}_1.\varrho_R$ 
9  $I := \mathbb{D}_1.I \setminus \mathbb{D}_2.I$ 
10 return  $\langle I, \varrho_L, \varrho_R \rangle$ 
```

---

**Implementation of Periodic Operators** Algorithm 1 has made frequent use of operators  $\bowtie$  and  $\cup$ , which are responsible for taking the difference and union of two periodic materialisations, respectively. In practice, these two operators can be implemented arbitrarily, so long as  $\text{unfold}(\mathbb{D}_1) - \text{unfold}(\mathbb{D}_2) = \text{unfold}(\mathbb{D}_1 \bowtie \mathbb{D}_2)$  and  $\text{unfold}(\mathbb{D}_1) \cup \text{unfold}(\mathbb{D}_2) = \text{unfold}(\mathbb{D}_1 \cup \mathbb{D}_2)$ . Next we describe our implementation of these operators, which utilises two auxiliary functions,  $\text{Ext}$  and  $\text{Aln}$ , responsible for extending and aligning periodic materialisations, respectively.

Consider two periodic materialisations  $\mathbb{D}_1$  and  $\mathbb{D}_2$ , and let  $\text{end} \in \{L, R\}$  denote which end of the periodic materialisation to operate on (either left or right). Given periodic materialisations  $\mathbb{D}_1$  and  $\mathbb{D}_2$  which both have valid (but potentially different) periods  $\varrho_{\text{end}}$ , function  $\text{Ext}$  computes a pair of periodic materialisations  $(\mathbb{D}'_1, \mathbb{D}'_2) = \text{Ext}(\mathbb{D}_1, \mathbb{D}_2, \text{end})$  such that the periodic regions are extended to have the same

---

**Algorithm 4:** Implementation of Periodic Union

---

**Input:** Two periodic materialisations  $\mathbb{D}_1$  and  $\mathbb{D}_2$

```
1  $\varrho_L := \varrho_R := \emptyset$ 
2 if  $\mathbb{D}_1.\varrho_L$  and  $\mathbb{D}_2.\varrho_L$  then
3    $(\mathbb{D}_1, \mathbb{D}_2) := \text{Ext}(\mathbb{D}_1, \mathbb{D}_2, L)$ 
4 if  $\mathbb{D}_1.\varrho_L$  or  $\mathbb{D}_2.\varrho_L$  then
5    $(\mathbb{D}_1, \mathbb{D}_2, \varrho_L) := \text{Aln}(\mathbb{D}_1, \mathbb{D}_2, L)$ 
6 if  $\mathbb{D}_1.\varrho_R$  and  $\mathbb{D}_2.\varrho_R$  then
7    $(\mathbb{D}_1, \mathbb{D}_2) := \text{Ext}(\mathbb{D}_1, \mathbb{D}_2, R)$ 
8 if  $\mathbb{D}_1.\varrho_R$  or  $\mathbb{D}_2.\varrho_R$  then
9    $(\mathbb{D}_1, \mathbb{D}_2, \varrho_R) := \text{Aln}(\mathbb{D}_1, \mathbb{D}_2, R)$ 
10  $I := \mathbb{D}_1.I \cup \mathbb{D}_2.I$ 
11 return  $\langle I, \varrho_L, \varrho_R \rangle$ 
```

---

length, which is the least common multiple (LCM) of the two; this can be easily achieved by copying the relevant periodic segments of the timeline. In contrast, function  $\text{Aln}$  tries to align the start and end points of periodic intervals of the two periodic materialisations at the target end, again through facts copying. If only one of the two periodic operators has a valid  $\varrho_{\text{end}}$ , an empty period is introduced for the other periodic materialisation so that alignment can still be performed. It should be noted that both  $\text{Ext}$  and  $\text{Aln}$  operations preserve the semantics of the interpretations: they do not change which facts hold at any time point but only alters the finite representations of the input periodic materialisations.

Theorem 1 states that Algorithms 3 and 4 correctly implement operators  $\bowtie$  and  $\cup$ , respectively; Theorem 2 states that Algorithm 1 is correct. The full proofs for these theorems are lengthy, so we only provide proof sketches; the full proofs are given in the online technical report.

**Theorem 1.** *If given input  $\mathbb{D}_1, \mathbb{D}_2$ , Algorithm 3 and 4 output  $\mathbb{D}_m$  and  $\mathbb{D}_u$ , respectively, then  $\text{unfold}(\mathbb{D}_1) - \text{unfold}(\mathbb{D}_2) = \text{unfold}(\mathbb{D}_m)$  and  $\text{unfold}(\mathbb{D}_1) \cup \text{unfold}(\mathbb{D}_2) = \text{unfold}(\mathbb{D}_u)$*

*Proof Sketch.* For minus, we show that the equation holds by leveraging that the periods of  $\text{unfold}(\mathbb{D}_1)$  and  $\text{unfold}(\mathbb{D}_2)$  can be extended and aligned to form a new pair of periods, and their difference also abides by the new periods. The case for union is similar.  $\square$

**Theorem 2.** *For a bounded DatalogMTL program  $\Pi$ , a bounded dataset  $E$ , a bounded deleting dataset  $E^-$ , a bounded inserting dataset  $E^+$ , a periodic materialisation  $\mathbb{I}$  such that  $\text{unfold}(\mathbb{I}) = \mathcal{C}_{\Pi, E}$ , let  $E' = E \setminus E^- \cup E^+$ , then after calling  $\text{DRED}_{\text{MTL}}(\Pi, E, E^-, E^+, \mathbb{I})$ , we have*

$$\text{unfold}(\mathbb{I}) = \mathcal{C}_{\Pi, E'}$$

*Proof Sketch.* We first show that  $\text{unfold}(\mathbb{I})$  computed by the  $\text{OVERDELETE}$  contains  $\mathcal{C}_{\Pi, E} - \mathcal{C}_{\Pi, E \setminus E^- \cup E^+}$ , which means every fact that no longer holds because of deleting  $E^-$  will be removed, and then we show that the union of  $\text{unfold}(\mathbb{R}), \text{unfold}(\mathbb{A})$  computed by the  $\text{REDERIVE, INSERT}$  and  $\text{unfold}(\mathbb{I})$  after deletion equals  $\mathcal{C}_{\Pi, E \setminus E^- \cup E^+}$ , which involves proving mistakenly deleted facts are rederived and every new fact that holds because of  $E^+$  are inserted.  $\square$

## Evaluation

We implemented the proposed  $\text{DRed}_{\text{MTL}}$  algorithm based on an efficient DatalogMTL reasoner MeTeoR (Wang et al. 2022) and empirically tested the performance of our implementation on three publicly available datasets. We chose MeTeoR as its latest version supports materialisation for DatalogMTL with bounded intervals (Wałęga et al. 2023), which allows us to directly compare the performance of  $\text{DRed}_{\text{MTL}}$  with rematerialisation from scratch. The source code of our implementation, the benchmarks we used, as well as an extended technical report containing all detailed proofs, are available online.

**Benchmarks**  $\text{LUBM}_t$  (Wang et al. 2022) is a temporal version of the well-known LUBM benchmark (Guo, Pan, and Heflin 2005). It has a recursive program consisting of 85 rules, of which 29 have temporal operators in them (denoted as  $|\Pi_{\text{mtl}}|$ ) and 56 do not (denoted as  $|\Pi_{\text{no\_mtl}}|$ ). The iTemporal dataset is generated from a temporal benchmark generator developed by Bellomarini, Nissl, and Sallinger (2022). Its program is highly recursive and consists of 12 rules. Finally, the Meteorological dataset (Maurer et al. 2002) contains long-term meteorological observations. It has a non-recursive program consisting of four rules, of which two contain temporal operators. These datasets were collected and made publicly available by Wang et al. (2025), and we used them without any modification. The statistics of these datasets is given in Table 2, where  $|E|$  and  $|I|$  are the number of explicitly given facts and the number of facts in the saturated interpretation, respectively. The ratio between  $|I|$  and  $|E|$  is usually a good indicator of the recursiveness of the corresponding rule set: the larger the ratio is, the more recursive and complex the rule set is and the more likely it is to generate a greater number of facts. Indeed, our choice of benchmarks is a nice mixture of highly recursive (iTemporal), mildly recursive ( $\text{LUBM}_t$ ) and nonrecursive datasets, which allows us to test the potential of our reasoning algorithm in different possible application scenarios.

**Test Settings** We conducted our experiment on a server with 256GB RAM and an Intel Xeon Platinum 8269CL 2.50GHz CPU, running Fedora Linux 40, kernel version Linux 6.10.10-200.fc40.x86\_64. Our evaluation primarily examines the capabilities of  $\text{DRed}_{\text{MTL}}$  in handling both deletions and insertions. Table 3 summarises the performance

	$\text{LUBM}_t$	iTemporal	Meteorological
$ \Pi_{\text{mtl}} $	29	3	2
$ \Pi_{\text{no\_mtl}} $	56	8	2
$ \Pi $	85	11	4
<i>Recursive</i>	Yes	Yes	No
$ E $	630.5k	46.41k	62.01M
$ I $	1.426M	30.62M	62.48M

Table 2: Dataset Statistics

comparison of our algorithm against rematerialisation. For our approach, we record the wall-clock time of running Algorithm 1 to handle the updates. For the baseline, we record the wall-clock time that MeTeoR spends on computing the canonical representation over the updated set of explicitly given facts from scratch. In addition to the time metrics, to better reflect the workload of  $\text{DRed}_{\text{MTL}}$ , we also record the number of facts derived in the three stages of our algorithm, i.e., overdeletion ( $|D|$ ), rederivation ( $|R|$ ) and insertion ( $|A|$ ). Note that for insertion, the sets of overdeleted and rederived facts are always empty, so we only record the number of inserted facts,  $|A|$ .

Our evaluation considers both small-scale and large-scale updates. For small-scale update tests, we randomly selected 100 facts and ran  $\text{DRed}_{\text{MTL}}$  to deal with the deletion; then we added these facts back and recorded the time  $\text{DRed}_{\text{MTL}}$  took to handle the insertion. Large-scale tests were performed in a similar fashion, except that 10% of the original dataset were removed and added back; the exact numbers of deleted facts are shown in Table 3. For each test case, three test runs with different (randomly selected) updates were performed; the results showed no significant variation in terms of running time. Therefore, for the ease of presentation we only reported the results of one run for each test case. Finally, for each test run, we made sure that the periodic materialisation produced by  $\text{DRed}_{\text{MTL}}$  ( $\mathbb{I}_1$ ) is equivalent to that produced by rematerialisation ( $\mathbb{I}_2$ ): this is achieved by verifying that both  $\mathbb{I}_1 \setminus \mathbb{I}_2$  and  $\mathbb{I}_2 \setminus \mathbb{I}_1$  are empty.

**Results** Our evaluation shows that  $\text{DRed}_{\text{MTL}}$  consistently outperforms rematerialisation for all small deletions and insertions. As shown in Table 3, on  $\text{LUBM}_t$ ,  $\text{DRed}_{\text{MTL}}$  is 69.4 times faster than the baseline for small deletion, and 121.3 times faster for small insertion. On the nonrecursive Meteorological dataset,  $\text{DRed}_{\text{MTL}}$  achieves similar performance improvement for small updates. On iTemporal, the performance improvement is more modest: the speedup is around six times. This is so since the program of iTemporal is highly recursive, making incremental maintenance especially challenging: as one can see, deleting only 100 facts leads to the overdeletion of over 244 thousand facts.

For large updates,  $\text{DRed}_{\text{MTL}}$  achieved a significant performance boost on the  $\text{LUBM}_t$  and iTemporal datasets. When deleting 10% of the data,  $\text{DRed}_{\text{MTL}}$  is 13.2 times faster on  $\text{LUBM}_t$  and 4.2 times faster on iTemporal. For large insertion, the improvements are 43.8 and 4.2 times, respectively. Interestingly, on the Meteorological dataset,  $\text{DRed}_{\text{MTL}}$  is slower than rematerialisation in dealing with large updates. Notice that for this dataset, the ratio between  $|I|$  and  $|E|$  is rather small, as depicted in Table 2, indicating that only a small fraction of the materialisation are inferred facts; moreover, the program is nonrecursive, so no effort is required for identifying the repeating pattern inside the materialisation. As such, for large updates on this dataset,  $\text{DRed}_{\text{MTL}}$ , which is based on efficient seminaive reasoning and period identification, loses its advantage.

To further analyse the performance of the proposed algorithm, we profiled our system on the deletion test cases and reported the runtime breakdown (by stage of reason-

‘Remat’ stands for Rematerialisation from the updated set of explicitly given facts.

Dataset	$ E^\pm $	Deletion				Insertion			
		DRED <sub>MTL</sub>				Remat	DRED <sub>MTL</sub>		Remat
		Time(s)	$ D $	$ R $	$ A $	Time(s)	Time(s)	$ A $	Time(s)
LUBM <sub>t</sub>	100	0.7k	11.5k	1.5k	11.4k	48.6k	0.4k	0.2k	48.5k
	63.1k	3.4k	372.1k	143.5k	251.9k	45.0k	1.1k	190.2k	48.2k
iTemporal	100	8.1k	244.5k	274.0k	274.4k	52.7k	8.7k	231.2k	52.8k
	4.6k	12.6k	8.963M	1.286M	1.288M	52.3k	12.4k	7.829M	52.6k
Meteorological	100	10	102	1	1	0.7k	11	101	0.7k
	6.201M	1.4k	6.271M	1105	1105	0.7k	1.2k	6.270M	0.7k

Table 3: Evaluation Results

Dataset	$ E^- $	Overdeletion	Rederivation	Insertion
LUBM <sub>t</sub>	100	59.0%	2.4%	38.6%
	63.1k	15.7%	67.9%	16.4%
iTemporal	100	70.5%	10.6%	18.9%
	4.6k	64.2%	29.9%	5.9%
Meteorological	100	62.8%	21.8%	15.4%
	6.201M	92.2%	7.7%	0.1%

Table 4: Deletion Test Runtime Breakdown

ing) in Table 4. It could be readily observed that in most cases, overdeletion is the most time consuming step of the algorithm. This is so since overdeletion produces the largest number of facts across the three stages. The only exception is the case of large deletion for LUBM<sub>t</sub>: rederivation consumes more time than overdeletion and insertion combined. Indeed, although rederivation produces less facts than overdeletion, it involves evaluating rules ‘backwards’ from head to body, which, as observed by Hu, Motik, and Horrocks (2018), could be more expensive than the seminaïve evaluation taking place in overdeletion and insertion. In the Datalog setting, enhancing rederivation with counting could help alleviate this issue, but extending the counting technique to DatalogMTL is beyond the scope of this paper.

Overall our test results suggest that the proposed algorithm improves substantially and consistently over rematerialisation, for both small and large updates. In some test cases, the running time decreased from several hours to several minutes, or from an hour to a few seconds, demonstrating the potential of deploying the proposed approach in industrial applications where short service response time is highly desirable.

## Conclusion and Future Work

In this paper, we have introduced a new technique for incrementally updating DatalogMTL materialisations. Compared with recomputing the materialisation from scratch, our technique achieves significant performance improvements, especially when the updates are small.

We see many exciting future research directions. From a

practical perspective, it would be interesting to see how well the proposed approach works in industrial scenarios such as IoT anomaly detection (Zhang et al. 2024): in typical such applications, DatalogMTL can be used to model anomaly detection rules that are triggered by streams composed of timestamped sensor data; the ability to reason incrementally offered by our method is crucial for ensuring efficient and reliable fault alerts. Moreover, we shall consider comparing the performance of our system with well-established stream reasoning frameworks such as C-SPARQL (Barbieri et al. 2009) and CQELS (Phuoc et al. 2011). While the underlying languages are quite different, it should be possible to transform our reasoning workload (at least the nonrecursive cases) to a C-SPARQL or CQELS workload, and test it on the corresponding engine. Last but not least, DRed is not the only incremental update algorithm for Datalog materialisations; algorithms such as B/F, FBF, and DRed<sup>c</sup> are popular alternatives of DRed and sometimes offer superior performance. It would thus be useful to consider extending these incremental maintenance algorithms for Datalog to support DatalogMTL reasoning, and to compare the performance of the adapted algorithms with that of ours.

## Acknowledgements

This work was generously funded by National Science and Technology Major Project of China under grant number 2025ZD1600800 and National Natural Science Foundation of China under grant number 62206169. We thank the anonymous reviewers for their constructive comments that helped greatly in shaping the final version of this paper.

## References

- Barbieri, D. F.; Braga, D.; Ceri, S.; Della Valle, E.; and Grossniklaus, M. 2009. C-SPARQL: SPARQL for continuous querying. In *Proceedings of the 18th International Conference on World Wide Web, WWW '09*, 1061–1062. New York, NY, USA: Association for Computing Machinery. ISBN 9781605584874.
- Bellomarini, L.; Nissl, M.; and Sallinger, E. 2022. iTemporal: An Extensible Generator of Temporal Benchmarks. In *International Conference on Data Engineering, ICDE 2022*, 2021–2033.
- Brandt, S.; Kalayci, E. G.; Kontchakov, R.; Ryzhikov, V.; Xiao, G.; and Zakharyashev, M. 2017. Ontology-based data access with a horn fragment of metric temporal logic. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, AAAI'17*, 1070–1076. AAAI Press.
- Brandt, S.; Kalayci, E. G.; Ryzhikov, V.; Xiao, G.; and Zakharyashev, M. 2018. Querying Log Data with Metric Temporal Logic. *J. Artif. Intell. Res.*, 62: 829–877.
- Ceri, S.; Gottlob, G.; and Tanca, L. 1989. What you Always Wanted to Know About Datalog (And Never Dared to Ask). *IEEE Trans. Knowl. Data Eng.*, 1(1): 146–166.
- Colombo, A.; Bellomarini, L.; Ceri, S.; and Laurenza, E. 2023. Smart Derivative Contracts in DatalogMTL. In *International Conference on Extending Database Technology, EDBT, 773–781*.
- Guo, Y.; Pan, Z.; and Heflin, J. 2005. LUBM: A benchmark for OWL knowledge base systems. *J. Web Semant.*, 3(2-3): 158–182.
- Gupta, A.; Mumick, I. S.; and Subrahmanian, V. S. 1993. Maintaining views incrementally. In *Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data, SIGMOD '93*, 157–166. New York, NY, USA: Association for Computing Machinery. ISBN 0897915925.
- Güzel Kalayci, E.; Xiao, G.; Ryzhikov, V.; Kalayci, T. E.; and Calvanese, D. 2018. Ontop-temporal: a tool for ontology-based query answering over temporal data. In *ACM International Conference on Information and Knowledge Management, 1927–1930*.
- Hu, P.; Motik, B.; and Horrocks, I. 2018. Optimised Maintenance of Datalog Materialisations. In McIlraith, S. A.; and Weinberger, K. Q., eds., *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, New Orleans, Louisiana, USA, February 2-7, 2018*, 1871–1879. AAAI Press.
- Koymans, R. 1990. Specifying Real-Time Properties with Metric Temporal Logic. *Real Time Syst.*, 2(4): 255–299.
- Maurer, E. P.; Wood, A. W.; Adam, J. C.; Lettenmaier, D. P.; and Nijssen, B. 2002. A Long-Term Hydrologically Based Dataset of Land Surface Fluxes and States for the Conterminous United States. *Journal of Climate*, 15(22): 3237 – 3251.
- Mori, M.; Papotti, P.; Bellomarini, L.; and Giudice, O. 2022. Neural Machine Translation for Fact-checking Temporal Claims. In *Fact Extraction and VERification Workshop*, 78–82.
- Motik, B.; Nenov, Y.; Piro, R.; and Horrocks, I. 2015. Incremental update of datalog materialisation: the backward/forward algorithm. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI'15*, 1560–1568. AAAI Press. ISBN 0262511290.
- Motik, B.; Nenov, Y.; Piro, R.; and Horrocks, I. 2019. Maintenance of datalog materialisations revisited. *Artificial Intelligence*, 269: 76–136.
- Nissl, M.; and Sallinger, E. 2022. Modelling Smart Contracts with DatalogMTL. In *Workshops of the EDBT/ICDT*, volume 3135.
- Phuoc, D. L.; Dao-Tran, M.; Parreira, J. X.; and Hauswirth, M. 2011. A Native and Adaptive Approach for Unified Processing of Linked Streams and Linked Data. In Aroyo, L.; Welty, C.; Alani, H.; Taylor, J.; Bernstein, A.; Kagal, L.; Noy, N. F.; and Blomqvist, E., eds., *The Semantic Web - ISWC 2011 - 10th International Semantic Web Conference, Bonn, Germany, October 23-27, 2011, Proceedings, Part I*, volume 7031 of *Lecture Notes in Computer Science*, 370–388. Springer.
- Ren, Y.; and Pan, J. Z. 2011. Optimising ontology stream reasoning with truth maintenance system. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management, CIKM '11*, 831–836. New York, NY, USA: Association for Computing Machinery. ISBN 9781450307178.
- Staudt, M.; and Jarke, M. 1996. Incremental Maintenance of Externally Materialized Views. In *Proceedings of the 22th International Conference on Very Large Data Bases, VLDB '96*, 75–86. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc. ISBN 1558603824.
- Urbani, J.; Margara, A.; Jacobs, C.; van Harmelen, F.; and Bal, H. 2013. DynamiTE: Parallel Materialization of Dynamic RDF Data. In Alani, H.; Kagal, L.; Fokoue, A.; Groth, P.; Biemann, C.; Parreira, J. X.; Aroyo, L.; Noy, N.; Welty, C.; and Janowicz, K., eds., *The Semantic Web – ISWC 2013*, 657–672. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-642-41335-3.
- Wałęga, P.; Kaminski, M.; and Cuenca Grau, B. 2019. Reasoning over Streaming Data in Metric Temporal Datalog. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, 3092–3099. AAAI Press.
- Wałęga, P. A.; Zawidzki, M.; Wang, D.; and Cuenca Grau, B. 2023. Materialisation-Based Reasoning in DatalogMTL with Bounded Intervals. In *AAAI Conference on Artificial Intelligence*, 6566–6574.
- Wałęga, P. A.; Kaminski, M.; Wang, D.; and Grau, B. C. 2023. Stream reasoning with DatalogMTL. *Journal of Web Semantics*, 76: 100776.
- Wang, D.; Hu, P.; Wałęga, P. A.; and Cuenca Grau, B. 2022. MeTeoR: Practical Reasoning in Datalog with Metric Temporal Operators. In *AAAI Conference on Artificial Intelligence*, 5906–5913.

Wang, S.; Zhao, K.; Wei, D.; Walega, P. A.; Wang, D.; Cai, H.; and Hu, P. 2025. Goal-Driven Reasoning in Data-logMTL with Magic Sets. In Walsh, T.; Shah, J.; and Kolter, Z., eds., *The Thirty-Ninth AAAI Conference on Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA*, 15203–15211. AAAI Press.

Zhang, F.; Hu, P.; Cai, H.; and Jiang, L. 2024. Parallel Collaborative Reasoning Approaches Based on Data-logMTL in IoT Scenarios. In *2024 27th International Conference on Computer Supported Cooperative Work in Design (CSCWD)*, 1055–1060.