



Stream Reasoning in DatalogMTL via Finite Materialisation

Przemysław (Przemek) Wałęga, Michał Zawidzki, and Bernardo Cuenca Grau

Stream Reasoning Workshop, 2022

Motivations

- ▶ DatalogMTL is a powerful extension of Datalog with metric temporal operators.
- ▶ DatalogMTL programs allow for **recursive propagation** of information (towards future and past) which is **problematic for stream reasoning**.

Question:

- ▶ How to *guarantee* that a DatalogMTL program has *no infinite recursion via time?*

Contributions

- ▶ We *define data dependent and data independent notions* of ‘no infinite recursion via time’.
- ▶ We show *algorithms* for checking if a DatalogMTL program satisfies above properties.
- ▶ We show *tight complexity bounds* for these problems.
- ▶ We show *sufficient conditions* which are easy to check.
- ▶ We show *tight complexity bounds* for reasoning in programs without recursion via time.

DatalogMTL

Dataset

A *dataset* consists of facts over rational intervals, e.g.:



Fraud detection example:

RedList(adam)@[0, 20]

HighRisk(david)@[0, 500]

HighRisk(ernesto)@[0, 500]

Transaction(adam, betty)@2.87

Transaction(betty, charlie)@12.15

Transaction(charlie, david)@17.5

Transaction(charlie, ernesto)@129.43

Program

Programs use MTL operators, e.g.,

$$\Diamond_{(0,100]} A \text{ at } t \Leftrightarrow A \text{ is true at } \textit{some moment in } (t+0, t+100]$$

$$\Box_{(0,100]} A \text{ at } t \Leftrightarrow A \text{ is true at } \textit{every moment in } (t+0, t+100]$$

A *program* is a set of rules $B \leftarrow A_1 \wedge \dots \wedge A_n$ where:

$$A := \top \mid \perp \mid P(\mathbf{c}) \mid \Diamond_\varrho A \mid \Diamond^\omega_\varrho A \mid \Box_\varrho A \mid \Box^\omega_\varrho A \mid A\mathcal{S}_\varrho A \mid A\mathcal{U}_\varrho A$$

$$B := \top \mid \perp \mid P(\mathbf{c}) \mid \Box_\varrho B \mid \Box^\omega_\varrho B$$

Fraud detection example:

$$\textit{TransactionChain}(x, y) \leftarrow \textit{Transaction}(x, y) \wedge \textit{RedList}(x)$$

$$\textit{TransactionChain}(x, z) \leftarrow \Diamond_{[0,24]} \textit{TransactionChain}(x, y) \wedge \textit{Transaction}(y, z)$$

$$\Box_{[0,100]} \textit{Suspect}(y) \leftarrow \textit{TransactionChain}(x, y) \wedge \textit{HighRisk}(y)$$

Reasoning

Main reasoning tasks (we consider the rational timeline):

- ▶ *Fact entailment*: do a program Π and a dataset \mathcal{D} entail a fact, e.g., $Suspect(adam)@100$?
- ▶ *Consistency checking*: do Π and \mathcal{D} have a model?

Theorem. Reasoning in DatalogMTL is:

- ▶ *ExpSpace-complete* for combined complexity,
- ▶ *PSpace-complete* for data complexity.

RECURSION VIA TIME

Canonical Interpretation

Unlike in Datalog:

- ▶ materialisation in DatalogMTL **can require infinitely many steps of rule application.**

Canonical Interpretation

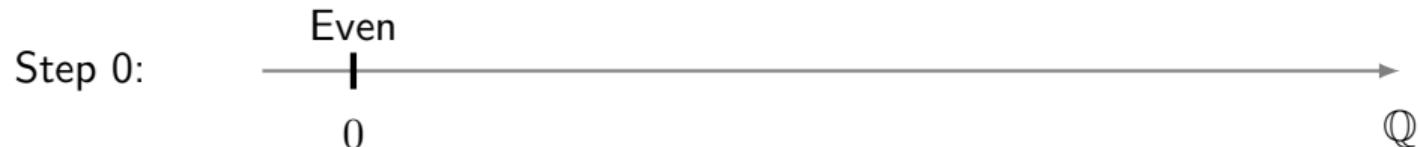
Unlike in Datalog:

- ▶ materialisation in DatalogMTL **can require infinitely many steps of rule application.**

Example:

$$\mathcal{D} = \{\text{Even}@0\}$$

$$\Pi = \{\boxplus_{[1,1]} \text{Odd} \leftarrow \text{Even}, \quad \boxplus_{[1,1]} \text{Even} \leftarrow \text{Odd}\}$$



Canonical Interpretation

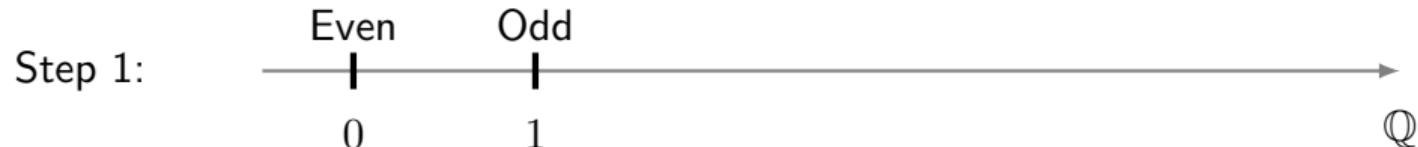
Unlike in Datalog:

- ▶ materialisation in DatalogMTL **can require infinitely many steps of rule application.**

Example:

$$\mathcal{D} = \{\text{Even}@0\}$$

$$\Pi = \{\boxplus_{[1,1]} \text{Odd} \leftarrow \text{Even}, \quad \boxplus_{[1,1]} \text{Even} \leftarrow \text{Odd}\}$$



Canonical Interpretation

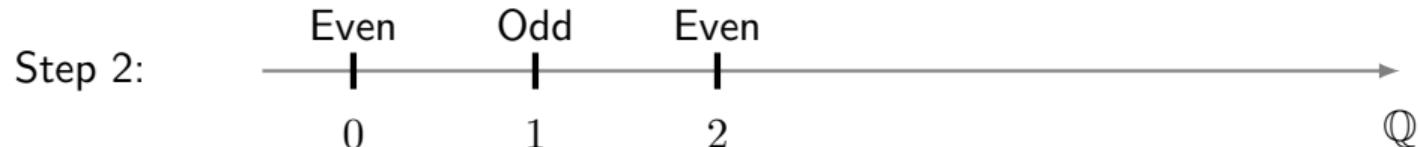
Unlike in Datalog:

- ▶ materialisation in DatalogMTL **can require infinitely many steps of rule application.**

Example:

$$\mathcal{D} = \{\text{Even}@0\}$$

$$\Pi = \{\boxplus_{[1,1]} \text{Odd} \leftarrow \text{Even}, \quad \boxplus_{[1,1]} \text{Even} \leftarrow \text{Odd}\}$$



Canonical Interpretation

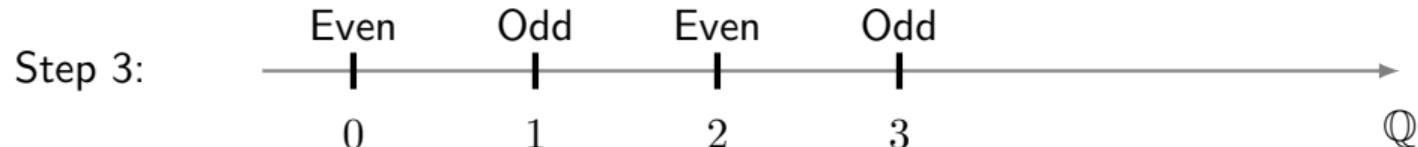
Unlike in Datalog:

- ▶ materialisation in DatalogMTL **can require infinitely many steps of rule application.**

Example:

$$\mathcal{D} = \{\text{Even}@0\}$$

$$\Pi = \{\boxplus_{[1,1]} \text{Odd} \leftarrow \text{Even}, \quad \boxplus_{[1,1]} \text{Even} \leftarrow \text{Odd}\}$$



Finitely Materialisable Programs

- ▶ How can we *guarantee that no infinite materialisation* occurs?

We say that a program Π is *finitely materialisable* for a dataset \mathcal{D} if materialisation of Π and \mathcal{D} takes a finite number of steps.

We say that Π is *finitely materialisable* for all datasets if materialisation of Π and any dataset \mathcal{D} takes a finite number of steps.

ALGORITHMS AND COMPLEXITY

Checking Finite-materialisability for a Single Dataset

Theorem. If Π is finitely materialisable for \mathcal{D} , then all the facts they entail are inside $[t_{\mathcal{D}}^{\min} - \text{offset}(\Pi, \mathcal{D}), t_{\mathcal{D}}^{\max} + \text{offset}(\Pi, \mathcal{D})]$.

Algorithm 1: Checking finite materialisability for a single dataset

Input: A program Π and a dataset \mathcal{D}

Output: A Boolean value

- 1 $\mathcal{D}_{\text{new}} := \mathcal{D};$
 - 2 $\varrho := [t_{\mathcal{D}}^{\min} - \text{offset}(\Pi, \mathcal{D}), t_{\mathcal{D}}^{\max} + \text{offset}(\Pi, \mathcal{D})];$
 - 3 **repeat**
 - 4 $\mathcal{D}_{\text{old}} := \mathcal{D}_{\text{new}};$
 - 5 $\mathcal{D}_{\text{new}} := \text{ApplyRules}(\Pi, \mathcal{D}_{\text{old}});$
 - 6 **if** there is $M @ \varrho' \in \mathcal{D}_{\text{new}}$ with $\varrho' \not\subseteq \varrho$ **then** Return false;
 - 7 **until** $\mathfrak{I}_{\mathcal{D}_{\text{old}}} = \mathfrak{I}_{\mathcal{D}_{\text{new}}};$
 - 8 **return** true;
-

Checking Finite-materialisability for a Single Dataset

Theorem. Algorithm 1 returns ‘true’ iff Π is finitely materialisable for \mathcal{D} .

Theorem. Algorithm 1 works in doubly exponential time in the size of Π and \mathcal{D} .

Checking Finite-materialisability for a Single Dataset

Theorem. Algorithm 1 returns ‘true’ iff Π is finitely materialisable for \mathcal{D} .

Theorem. Algorithm 1 works in doubly exponential time in the size of Π and \mathcal{D} .

The algorithm is practically efficient, but **it is not worst-case optimal**:

Theorem. Checking finite materialisability for a given dataset is ExpSpace-complete for combined and PSpace-complete for data complexity.

Checking Finite-materialisability for All Datasets

- ▶ Checking data-independent finite-materialisability *reduces to checking the data-dependent variant* for a **critical dataset** defined as follows:

Definition. **Critical dataset** \mathcal{D}_Π for Π contains all facts $P(s)@[0, \text{depth}(\Pi)]$, where:

- ▶ P occurs in Π ,
- ▶ s mentions constants from Π and a single fresh constant.

Checking Finite-materialisability for All Datasets

- ▶ Checking data-independent finite-materialisability *reduces to checking the data-dependent variant* for a **critical dataset** defined as follows:

Definition. **Critical dataset** \mathcal{D}_Π for Π contains all facts $P(s)@[0, \text{depth}(\Pi)]$, where:

- ▶ P occurs in Π ,
- ▶ s mentions constants from Π and a single fresh constant.

Theorem. Π is finitely materialisable for all datasets iff it is for \mathcal{D}_Π .

- ▶ If materialisation of Π and some \mathcal{D} takes infinitely many steps, then the same holds for Π and \mathcal{D}_Π . This is guaranteed by the definition of \mathcal{D}_Π , in particular, by choosing the sufficiently long interval $[0, \text{depth}(\Pi)]$.

Checking Finite-materialisability for All Datasets

Theorem. If Π is finitely materialisable for all datasets, facts entailed by Π and any \mathcal{D} are in $[t_{\mathcal{D}}^{\min} - \text{offset}(\Pi), t_{\mathcal{D}}^{\max} + \text{offset}(\Pi)]$.

Algorithm 2: Checking finite materialisability for all datasets

Input: A program Π

Output: A Boolean value

- 1 $\mathcal{D}_{\text{new}} := \mathcal{D}_{\Pi};$
 - 2 $\varrho := [t_{\mathcal{D}_{\Pi}}^{\min} - \text{offset}(\Pi), t_{\mathcal{D}_{\Pi}}^{\max} + \text{offset}(\Pi)];$
 - 3 **repeat**
 - 4 $\mathcal{D}_{\text{old}} := \mathcal{D}_{\text{new}};$
 - 5 $\mathcal{D}_{\text{new}} := \text{ApplyRules}(\Pi, \mathcal{D}_{\text{old}});$
 - 6 **if** there is $M @ \varrho' \in \mathcal{D}_{\text{new}}$ with $\varrho' \not\subseteq \varrho$ **then** Return false;
 - 7 **until** $\mathfrak{I}_{\mathcal{D}_{\text{old}}} = \mathfrak{I}_{\mathcal{D}_{\text{new}}};$
 - 8 **return** true;
-

Checking Finite-materialisability for All Datasets

Theorem. Algorithm 2 returns ‘true’ iff Π is finitely materialisable for all datasets.

Theorem. Algorithm 2 runs in exponential time in the size of Π .

Checking Finite-materialisability for All Datasets

Theorem. Algorithm 2 returns ‘true’ iff Π is finitely materialisable for all datasets.

Theorem. Algorithm 2 runs in exponential time in the size of Π .

Moreover, our algorithm *is worst-case optimal*:

Theorem. Checking finite materialisability (for all datasets) is EXPTIME-complete.

- ▶ If materialisation of Π and \mathcal{D} takes finitely many steps, then the number of these steps is exponential in the size of Π . Thus, it suffices to check if the number of steps in materialisation of Π and \mathcal{D}_Π exceeds our exponential bound.

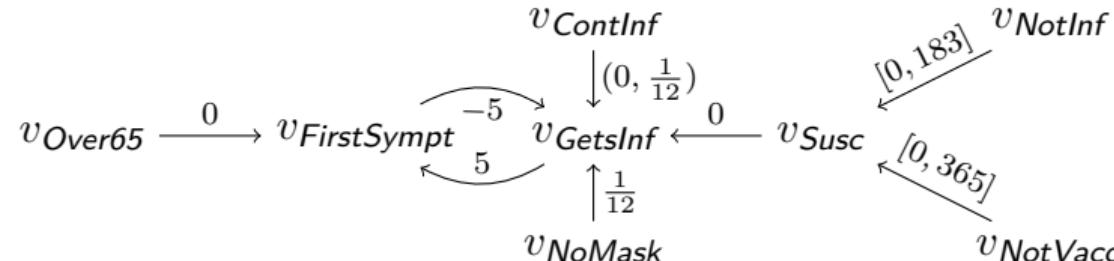
SUFFICIENT CONDITIONS

MTL-acyclicity

Consider the program:

$$\begin{aligned} \text{Susc}(x) &\leftarrow \Box_{[0,365]} \text{NotVacc}(x) \wedge \Box_{[0,183]} \text{NotInf}(x), \\ \text{GetsInf}(x) &\leftarrow \text{ContInf}(x, y) \mathcal{S}_{\frac{1}{12}} \text{NoMask}(x) \wedge \text{Susc}(x), \\ \text{FirstSympt}(x) &\leftarrow \Box_5 \text{GetsInf}(x) \wedge \text{Over65}(x), \\ \Box_5 \text{GetsInf}(x) &\leftarrow \text{FirstSympt}(x). \end{aligned}$$

It's *metric dependency graph* is:



There is no cycles with weight $\neq [0, 0]$, so the program is *MTL-acyclic*.

MTL-acyclicity

Theorem. MTL-acyclic programs are finitely materialisable (for all datasets).

Theorem. Checking if a program is MTL-acyclic is NL-complete.

- ▶ Hence, MTL-acyclicity is an *easy to check sufficient condition* for finite materialisability.

COMPLEXITY OF REASONING

Reasoning in Finitely Materialisable Programs

Theorem. Fact entailment in finitely materialisable programs is EXPTIME-complete for combined and PSpace-complete for data complexity.

Observe that:

- ▶ DatalogMTL is ExpSpace-complete for combined and PSpace-complete for data complexity.
- ▶ Datalog is ExpTime-complete for combined and P-complete for data complexity.

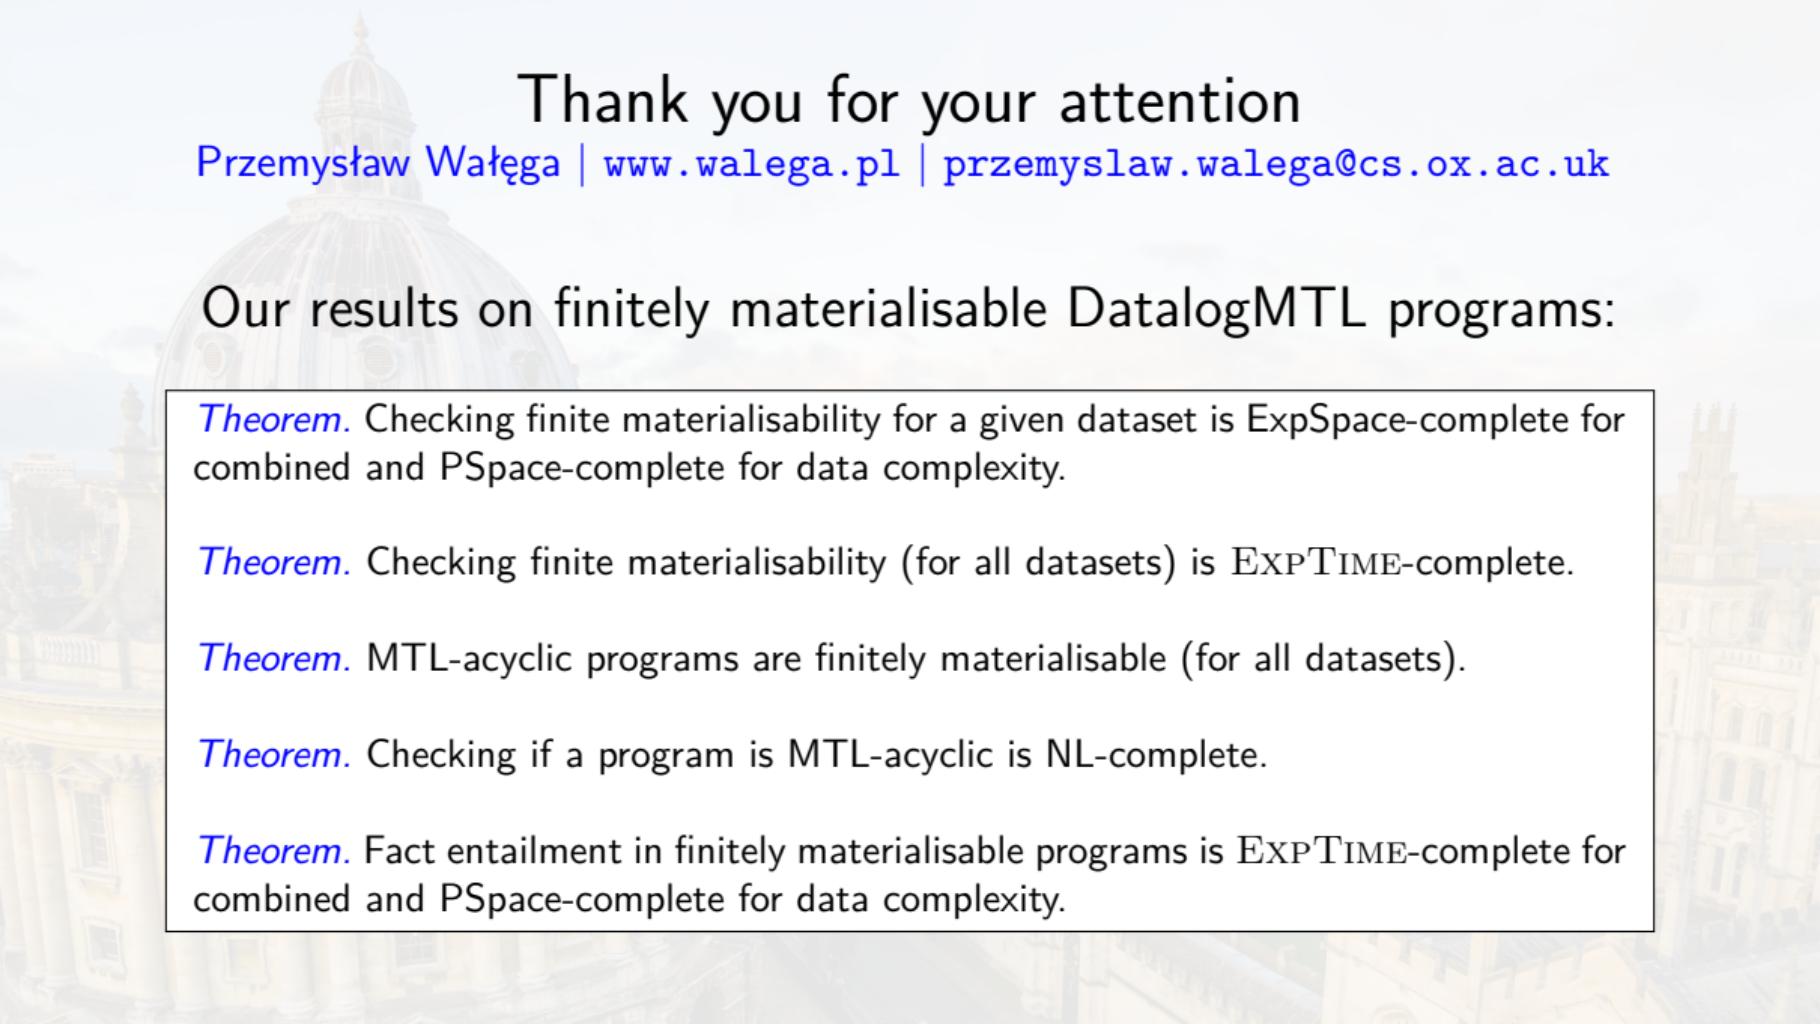
Moreover:

- ▶ Finitely materialisable DatalogMTL programs *strictly contain Datalog* programs.
- ▶ However, reasoning in finitely materialisable DatalogMTL programs has the *same combined complexity as Datalog*.

Conclusions

We introduced a class of *finitely materialisable DatalogMTL programs*:

- ▶ that are naturally amenable to materialisation-based reasoning via scalable forward chaining techniques,
- ▶ for which we provided a membership check and an easy to verify sufficient condition,
- ▶ in which reasoning is no harder (for combined complexity) than in pure Datalog.



Thank you for your attention

Przemysław Wałęga | www.walega.pl | przemyslaw.walega@cs.ox.ac.uk

Our results on finitely materialisable DatalogMTL programs:

Theorem. Checking finite materialisability for a given dataset is ExpSpace-complete for combined and PSpace-complete for data complexity.

Theorem. Checking finite materialisability (for all datasets) is EXPTIME-complete.

Theorem. MTL-acyclic programs are finitely materialisable (for all datasets).

Theorem. Checking if a program is MTL-acyclic is NL-complete.

Theorem. Fact entailment in finitely materialisable programs is EXPTIME-complete for combined and PSpace-complete for data complexity.