Controlled Query Evaluation in OWL 2 QL: A "Longest Honeymoon" Approach

Piero Bonatti[1][0000−0003−1436−5660], Gianluca Cima[2][0000−0003−1783−5605], Domenico Lembo[3][0000−0002−0628−242X], Lorenzo Marconi[3][0000−0001−9633−8476], Riccardo Rosati[3][0000−0002−7697−4958], Luigi Sauro[1][0000−0001−6056−0868], and Domenico Fabio Savo[4][0000−0002−8391−8049]

[1] Università di Napoli Federico II
{pab,luigi.sauro}@unina.it
[2] University of Bordeaux, CNRS, Bordeaux INP, LaBRI
gianluca.cima@u-bordeaux.fr
[3] Sapienza Università di Roma
{lembo,marconi,rosati}@diag.uniroma1.it
[4] Università degli Studi di Bergamo
domenicofabio.savo@unibg.it

Abstract. Controlled Query Evaluation (CQE) has been recently studied in the context of Semantic Web ontologies. The goal of CQE is concealing some query answers so as to prevent external users from inferring confidential information. In general, there exist multiple, mutually incomparable ways of concealing answers, and previous CQE approaches choose in advance which answers are visible and which are not. In this paper, instead, we study a *dynamic* CQE method, namely, we propose to alter the answer to the current query based on the evaluation of previous ones. We aim at a system that, besides being able to protect confidential data, is maximally cooperative, which intuitively means that it answers affirmatively to as many queries as possible; it achieves this goal by delaying answer modifications as much as possible. We also show that the behavior we get cannot be intensionally simulated through a static approach, independent of query history. Interestingly, for OWL 2 QL ontologies and policy expressed through denials, query evaluation under our semantics is first-order rewritable, and thus in AC^0 in data complexity. This paves the way for the development of practical algorithms, which we also preliminarily discuss in the paper.

Keywords: Ontologies · Data Protection · Description Logics · First-order rewritability

1 Introduction

Semantic Web technologies are increasingly used to represent and link together different sources of information coming from public organizations as well as private citizens. This information may include sensitive knowledge, e.g. medical records or social network activities, whose disclosure may affect the privacy of

individuals if not adequately protected [8,16]. Furthermore, OWL 2 ontologies allow one to infer implicit information from explicit data, which amplifies the risk of information leakage.

One goal of confidentiality-preserving data publishing is to prevent the disclosure of sensitive information to unauthorized users while being as cooperative as possible, that is, answering queries honestly whenever this does not harm confidentiality. Specifically, in controlled query evaluation (CQE) [3,4] the data protection policy is declaratively specified through logical formulas and is enforced by altering query answers through so-called censors, which either refuse to answer some queries or lie when this is needed to protect some secrets. In general, there exist multiple, mutually incomparable ways of concealing answers, i.e., mutually incomparable censors. Different works have proposed static CQE methods, where a censor is constructed (or approximated) beforehand, establishing once and for all which queries should be answered truthfully [8,15,17,13,11]. In several cases, such approaches are not fully cooperative, because the secure view of the data is chosen without taking the users' interests into account.

Conversely, following the work of Biskup and Bonatti [5], in this paper we introduce a dynamic CQE (dynCQE) method that progressively decides whether being truthful or lying, based on the specific stream of queries. Roughly speaking, the dynamic CQE approach selects, at each step, as many censors as possible, coherently with the previous answers. By doing so, it maximizes the possibility of answering the next query honestly by choosing from the current pool of censors those that allow to answer the query truthfully (if any).

We will prove that this method satisfies the so-called "longest honeymoon" property, which means that, given a sequence of queries, dynCQE returns the longest possible sequence of honest answers before lying. This property can be supported with several arguments. First, without any specific model of the users' intentions, the order in which queries are posed allegedly reflects their importance. Secondly, since we cannot foresee which nor how many queries are coming in the future, answering honestly the current query (if possible) is the most cooperative possible strategy. We will prove also that dynCQE is optimal in a more classical sense: the set of queries honestly answered by dynCQE is always maximal under set containment.

After introducing the dynCQE framework and formally investigating its general properties (Section 3), the paper focuses on ontologies in OWL 2 QL [20], a tractable profile of OWL 2 designed for data-intensive applications. For this setting, in Section 4, we first show that the behavior of dynCQE cannot be simulated by static CQE through data-independent modifications of the intensional components of the framework, i.e., the ontology (TBox) and the formulas representing the data protection policy. It is thus necessary to devise specific techniques to implement the dynamic approach. To this aim, we provide a tailored query rewriting algorithm through which we show that dynCQE query processing in OWL 2 QL is *first-order rewritable*, which implies that its data complexity is in AC^0 (like the evaluation of first-order sentences, i.e., SQL, queries). Towards practical implementations, in Section 5, we present a first optimization

of the query reformulation technique used to prove the first-order rewritability result, based on the information acquired by the system during the interaction with users; we also present a possible approximation of the approach, should the sequence of queries become too long for our rewriting technique. A section on related work and one on final remarks conclude the paper.

2 Preliminaries

For the technical treatment we resort to Description Logics (DLs), which are decidable fragments of First-Order (FO) logic underpinning the OWL 2 standard. We introduce here the basic notions needed in this work and refer the reader to [1] for further details. The languages of our interest are built from an alphabet Γ that consists of unary predicates (a.k.a. *atomic concepts*), binary predicates (a.k.a. *atomic roles*), constants (a.k.a. *individual names*), and a countably infinite supply of variables. An atom is a formula of the form $A(t)$ or $P(t_1, t_2)$, where A is an atomic concept, P is an atomic role, and the terms t, t_1, t_2 are either variables or constants. An atom is *ground* if all its terms are constants.

A DL ontology $\mathcal{O} = \mathcal{T} \cup \mathcal{A}$ is constituted by a TBox $\mathcal{T}$ and an ABox $\mathcal{A}$, specifying intensional and extensional knowledge, respectively. In particular, in this paper we assume that the ABox is a set of ground atoms. A *model* of an ontology $\mathcal{O} = \mathcal{T} \cup \mathcal{A}$ is a FO interpretation that satisfies all axioms in $\mathcal{T}$ and $\mathcal{A}$. $\mathcal{O}$ is *consistent* if it has at least one model, *inconsistent* otherwise, and *entails* an FO sentence ϕ, denoted $\mathcal{O} \models \phi$, if ϕ is true in every model of $\mathcal{O}$. Given an ABox $\mathcal{A}$ and a FO sentence ϕ, we say that ϕ *evaluates to true in* $\mathcal{A}$ if the evaluation of ϕ in the Herbrand model of $\mathcal{A}$ is true [18], otherwise we say that ϕ *evaluates to false in* $\mathcal{A}$. In the paper, we often refer to the set of ground atoms entailed by $\mathcal{T} \cup \mathcal{A}$, which we denote with $\mathsf{cl}_{\mathcal{T}}(\mathcal{A})$.

In this work, we focus on ontologies expressed in *DL-Lite*$_R$ [9], which is the logical counterpart of OWL 2 QL [19]. In this DL, a role R is an atomic role P or its inverse P^-, whereas a concept B takes the form A, $\exists P$, or $\exists P^-$. The concepts $\exists P$ and $\exists P^-$ denote the domain and the range of a role P, respectively. A *DL-Lite*$_R$ TBox $\mathcal{T}$ is a set of *positive inclusions* of the form $B_1 \sqsubseteq B_2$ or $R_1 \sqsubseteq R_2$, and *negative inclusions* of the form $B_1 \sqsubseteq \neg B_2$ or $R_1 \sqsubseteq \neg R_2$.

By $conj(\vec{x})$ we mean a conjunction $\alpha_1 \wedge \ldots \wedge \alpha_n$ of atoms where $\vec{x}$ indicates all the variables occurring in it. Then, a Boolean Conjunctive Query (BCQ) is an existentially quantified conjunction of atoms $\exists \vec{x}(conj(\vec{x}))$ and a Boolean Union of Conjunctive Queries (BUCQ) is a disjunction $q_1 \vee \ldots \vee q_n$ of BCQs. Sometimes we write $q \in q'$ to indicate that the BCQ q is one of the BCQs of the BUCQ q'. Note that a ground atom can be seen as a BCQ with no variables, and that a BCQ is a BUCQ with only one disjunct.

Given a BCQ q, $Atoms(q)$ is the set of atoms occurring in q. Given two BUCQs $q_1 = q_1^1 \vee \ldots \vee q_1^n$ and $q_2 = q_2^1 \vee \ldots \vee q_2^m$, we denote by $q_1 \wedge q_2$ the BUCQ

$$\begin{array}{c} (q_1^1 \wedge q_2^1) \vee \ldots \vee (q_1^1 \wedge q_2^m) \vee \\ \vdots \\ (q_1^n \wedge q_2^1) \vee \ldots \vee (q_1^n \wedge q_2^m)\,. \end{array}$$

We recall that entailment of BUCQs in *DL-Lite*$_R$ is FO rewritable, i.e., for every *DL-Lite*$_R$ TBox $\mathcal{T}$ and BUCQ q, it is possible to compute an FO query q_r, called the *perfect reformulation of q with respect to $\mathcal{T}$*, such that, for each ABox $\mathcal{A}$, $\mathcal{T} \cup \mathcal{A} \models q$ iff q_r evaluates to true in $\mathcal{A}$. We will use the algorithm *PerfectRef* presented in [9], which uses only positive inclusions in $\mathcal{T}$ as rewriting rules to compute perfect reformulations. We point out that the reformulation returned by *PerfectRef* is a BUCQ. The following proposition is from [9].

Proposition 1. *Let $\mathcal{T} \cup \mathcal{A}$ be a consistent DL-Lite$_R$ ontology and let q be a BUCQ. Then, $\mathcal{T} \cup \mathcal{A} \models q$ iff PerfectRef$(q, \mathcal{T})$ evaluates to true in $\mathcal{A}$.*

Furthermore, a *policy* $\mathcal{P}$ is a (finite) set of *denials*, i.e., sentences of the form $q \rightarrow \bot$, where q is a BCQ. An interpretation satisfies a denial $q \rightarrow \bot$ iff it does not satisfy the BCQ q. We denote by $q(\mathcal{P})$ the BUCQ $\bigvee_{q \rightarrow \bot \in \mathcal{P}} q$.

The following proposition follows from the definition of satisfaction of a denial and from Proposition 1.

Proposition 2. *Let $\mathcal{T} \cup \mathcal{A}$ be a consistent DL-Lite$_R$ ontology and let $\mathcal{P}$ be a policy. Then, $\mathcal{T} \cup \mathcal{P} \cup \mathcal{A}$ is a consistent FO theory iff PerfectRef$(q(\mathcal{P}), \mathcal{T})$ evaluates to false in $\mathcal{A}$.*

Our complexity results refer to data complexity, i.e., the complexity computed with respect to the size of the ABox only.

3 Framework

We now introduce our framework. All definitions and properties given in this section apply to any DL language.

A *CQE specification* is a pair $\langle \mathcal{T}, \mathcal{P} \rangle$, where $\mathcal{T}$ is a TBox and $\mathcal{P}$ is a policy, such that $\mathcal{T} \cup \mathcal{P}$ is consistent. A CQE instance is a triple $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$, where $\langle \mathcal{T}, \mathcal{P} \rangle$ is a CQE specification, and $\mathcal{A}$ is an ABox such that $\mathcal{T} \cup \mathcal{A}$ is consistent.

Censors specify which consequences of an ontology can be disclosed without violating the policy. The following definition is adapted from [11, Definition 1].[1]

Definition 1 (Censor). *Let $\mathcal{E} = \langle \mathcal{T}, \mathcal{A}, \mathcal{P} \rangle$ be a CQE instance. A* censor *for $\mathcal{E}$ is an ABox $\mathcal{C} \subseteq \mathsf{cl}_{\mathcal{T}}(\mathcal{A})$ such that $\mathcal{T} \cup \mathcal{P} \cup \mathcal{C}$ is consistent.*

Given a CQE instance $\mathcal{E}$ and a censor $\mathcal{C}$ for $\mathcal{E}$, we say that $\mathcal{C}$ is *optimal* if there exists no censor $\mathcal{C}'$ for $\mathcal{E}$ such that $\mathcal{C} \subset \mathcal{C}'$. We denote by *OptCens*$(\mathcal{E})$ the set of all the optimal censors for $\mathcal{E}$. We observe that a censor for a CQE instance $\mathcal{E}$ always exists,[2] and thus *OptCens*$(\mathcal{E}) \neq \emptyset$. Given a BUCQ q, we denote by *OptCens*$(\mathcal{E}, q)$ the set of optimal censors that, together with $\mathcal{T}$, entail q:

$$\mathit{OptCens}(\mathcal{E}, q) = \{\mathcal{C} \in \mathit{OptCens}(\mathcal{E}) \mid \mathcal{T} \cup \mathcal{C} \models q\}$$

[1] Other definitions of censors have been considered in the literature, for example in [15,17]. Definition 1 is chosen because it yields several important properties, such as *indistinguishability* (cf. Section 6), and it has been thoroughly investigated in various settings (e.g., in [10,11]).

[2] Trivially, the empty set is a censor for any CQE instance $\mathcal{E}$.

The following notion of *protection state* captures the history of queries submitted by the users to a CQE instance.

Definition 2 (State). *Let $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ be a CQE instance. A* protection state of $\mathcal{E}$ *(or simply state of $\mathcal{E}$) is a pair $\mathcal{S} = \langle \mathcal{E}, \mathcal{Q} \rangle$, where $\mathcal{Q} = \langle q_1, \ldots, q_n \rangle$ (with $n \geq 0$) is a sequence of BUCQs.*

Below we formalize our idea of dynamic CQE (dynCQE), i.e., a CQE that takes into account the sequence of queries that have been already processed. In what follows, given a CQE instance $\mathcal{E}$, a sequence $\mathcal{Q}_n = \langle q_1, \ldots, q_n \rangle$ of BUCQs, and any integer $i \in [0, n]$, we denote with $\mathcal{Q}_i$ the sequence $\langle q_1, \ldots, q_i \rangle$ and with $\mathcal{S}_i$ the state $\langle \mathcal{E}, \mathcal{Q}_i \rangle$ of $\mathcal{E}$, with the convention that $\mathcal{Q}_0$ is the empty sequence $\langle\rangle$.

Definition 3 (Dynamic CQE – dynCQE). *Let $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ be a CQE instance, and let $\mathcal{Q}_n = \langle q_1, \ldots, q_n \rangle$ (with $n \geq 0$) a sequence of BUCQs. The set $\mathit{StCens}(\mathcal{S}_n)$ of censors of $\mathcal{S}_n$ is inductively defined as follows:*

- $\mathit{StCens}(\mathcal{S}_0) = \mathit{OptCens}(\mathcal{E})$;
- $\mathit{StCens}(\mathcal{S}_{i+1}) = \begin{cases} \mathit{StCens}(\mathcal{S}_i) \text{ if } \mathit{StCens}(\mathcal{S}_i) \cap \mathit{OptCens}(\mathcal{E}, q_{i+1}) = \emptyset, \\ \mathit{StCens}(\mathcal{S}_i) \cap \mathit{OptCens}(\mathcal{E}, q_{i+1}) \text{ otherwise}, \end{cases}$
  for every $0 \leq i \leq n-1$.

For each BUCQ q_i occurring in $\mathcal{Q}_n$, we say that q_i is entailed by $\mathcal{S}_n$*, denoted by $\mathcal{S}_n \models q_i$, if $\mathcal{T} \cup \mathcal{C} \models q_i$ for every $\mathcal{C} \in \mathit{StCens}(\mathcal{S}_n)$. We denote by $\mathit{EntQ}(\mathcal{S}_n)$ the set of queries of $\mathcal{Q}_n$ entailed by $\mathcal{S}_n$, i.e., $\mathit{EntQ}(\mathcal{S}_n) = \{q \in \mathcal{Q}_n \mid \mathcal{S}_n \models q\}$.*

One can see that, for any $i = 1, \ldots, n$, the set of censors of a state $\mathcal{S}_i$ is always non-empty and consists of a subset of the set of censors in its predecessor state $\mathcal{S}_{i-1}$, i.e. $\mathit{StCens}(\mathcal{S}_{i-1}) \supseteq \mathit{StCens}(\mathcal{S}_i) \supset \emptyset$. This also means that $\mathit{EntQ}(\mathcal{S}_{i-1}) \subseteq \mathit{EntQ}(\mathcal{S}_i)$ holds for any $i = 1, \ldots, n$.

Informally speaking, each set $\mathit{StCens}(\mathcal{S}_i)$ (with $1 \leq i \leq n$) in the above definition progressively selects the optimal censors of $\mathcal{E}$ that agree with $\mathit{EntQ}(\mathcal{S}_i)$. If none of the surviving optimal censors in $\mathit{StCens}(\mathcal{S}_i)$ entails (together with $\mathcal{T}$) a query q_{i+1}, then $\mathcal{S}_{i+1} \not\models q_{i+1}$, so we have that $\mathit{StCens}(\mathcal{S}_{i+1}) = \mathit{StCens}(\mathcal{S}_i)$. Conversely, if at least one of the censors in $\mathit{StCens}(\mathcal{S}_i)$, together with the TBox, entails q_{i+1}, then, according to dynCQE, we have a positive answer, and $\mathit{StCens}(\mathcal{S}_{i+1})$ keeps only the censors in $\mathit{StCens}(\mathcal{S}_i)$ that agree with such answer.

As a result, the stream of queries is processed greedily, answering the truth as long as some of the censors in $\mathit{StCens}(\mathcal{S}_n)$ allows to do it (*longest honeymoon approach* [5]), as we will formally show below.

Note that, by Definition 3, given a state $\mathcal{S} = \langle \mathcal{E}, \mathcal{Q} \rangle$ and a query q occurring in $\mathcal{Q}$, we have that either $\mathcal{T} \cup \mathcal{C} \models q$ for every $\mathcal{C} \in \mathit{StCens}(\mathcal{S})$, or $\mathcal{T} \cup \mathcal{C} \not\models q$ for every $\mathcal{C} \in \mathit{StCens}(\mathcal{S})$. This means that $\mathcal{S} \models q$ if and only if there exists a censor $\mathcal{C} \in \mathit{StCens}(\mathcal{S})$ such that $\mathcal{T} \cup \mathcal{C} \models q$.

Example 1. Some pharmaceutical products may reveal with high accuracy which kind of disease is affecting a person. For instance, drugs that contain phenytoin, or that are classified as anti-seizure medications, indicate some form of epilepsy.

Let $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ be a CQE instance, where:

$\mathcal{T} = \{\mathsf{Abc} \sqsubseteq \mathsf{Antiseizure}\}$;
$\mathcal{P} = \{\exists x, y(\mathsf{buy}(x, y) \wedge \mathsf{Antiseizure}(y)) \rightarrow \bot,$
$\quad \exists x, y(\mathsf{buy}(x, y) \wedge \mathsf{contain}(y, \mathsf{phenytoin})) \rightarrow \bot\}$;
$\mathcal{A} = \{\mathsf{buy}(\mathsf{john}, \mathsf{m}_a), \mathsf{Abc}(\mathsf{m}_a), \mathsf{buy}(\mathsf{alice}, \mathsf{m}_b), \mathsf{contain}(\mathsf{m}_b, \mathsf{phenytoin})\}$.

In words, the TBox states that Abc is an anti-seizure medication, while the policy conceals the presence of patients suffering from epilepsy.

Let us start by considering an empty sequence of BUCQs. By definition, we have that $\mathit{StCens}(\langle \mathcal{E}, \langle\rangle\rangle)$ coincides with the set of the optimal censors for $\mathcal{E}$:

- $\mathcal{C}_1 = \{\mathsf{buy}(\mathsf{john}, \mathsf{m}_a), \mathsf{buy}(\mathsf{alice}, \mathsf{m}_b)\}$;
- $\mathcal{C}_2 = \{\mathsf{buy}(\mathsf{john}, \mathsf{m}_a), \mathsf{contain}(\mathsf{m}_b, \mathsf{phenytoin})\}$;
- $\mathcal{C}_3 = \{\mathsf{Abc}(\mathsf{m}_a), \mathsf{Antiseizure}(\mathsf{m}_a), \mathsf{buy}(\mathsf{alice}, \mathsf{m}_b)\}$;
- $\mathcal{C}_4 = \{\mathsf{Abc}(\mathsf{m}_a), \mathsf{Antiseizure}(\mathsf{m}_a), \mathsf{contain}(\mathsf{m}_b, \mathsf{phenytoin})\}$.

Let $q_1 = \mathsf{buy}(\mathsf{john}, \mathsf{m}_a)$ be the first query. The censors $\mathcal{C}_1$ and $\mathcal{C}_2$ agree with answering *true* to this query. All the censors that disagree with such answer are then removed, obtaining $\mathit{StCens}(\langle \mathcal{E}, \langle q_1\rangle\rangle) = \mathit{StCens}(\langle \mathcal{E}, \langle\rangle\rangle) \cap \mathit{OptCens}(\mathcal{E}, q_1) = \{\mathcal{C}_1, \mathcal{C}_2\}$. Then, let $q_2 = \mathsf{Abc}(\mathsf{m}_a)$ be a new query in the sequence. Since neither $\mathcal{T} \cup \mathcal{C}_1$ nor $\mathcal{T} \cup \mathcal{C}_2$ entail q_2, then $\mathit{StCens}(\langle \mathcal{E}, \langle q_1, q_2\rangle\rangle) = \mathit{StCens}(\langle \mathcal{E}, \langle q_1\rangle\rangle)$. Now, consider to add $q_3 = \exists x \mathsf{buy}(x, \mathsf{m}_b)$ to the sequence. Since $\mathcal{T} \cup \mathcal{C}_1 \models q_3$ while $\mathcal{T} \cup \mathcal{C}_2 \not\models q_3$, we have $\mathit{StCens}(\mathcal{S}) = \{\mathcal{C}_1\}$, where $\mathcal{S} = \langle \mathcal{E}, \mathcal{Q}\rangle$ with $\mathcal{Q} = \langle q_1, q_2, q_3\rangle$. Clearly, $\mathcal{S} \models q_1$ and $\mathcal{S} \models q_3$, but $\mathcal{S} \not\models q_2$. □

Let $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A}\rangle$ be a CQE instance. For all states $\mathcal{S}$ of $\mathcal{E}$, our dynamic CQE method is *optimal with respect to* $\mathcal{S}$, in the sense that we have that $\mathit{EntQ}(\mathcal{S})$ is never strictly contained in the set of queries of $\mathcal{S}$ entailed by any censor $\mathcal{C}$ for $\mathcal{E}$. In order to formalize this property, for all states $\mathcal{S} = \langle \mathcal{E}, \mathcal{Q}\rangle$ and all censors $\mathcal{C}$ for $\mathcal{E}$, let $\mathit{EntQ}(\mathcal{Q}, \mathcal{C}, \mathcal{T})$ be the subset of queries of $\mathcal{Q}$ entailed by $\mathcal{C} \cup \mathcal{T}$, i.e. $\mathit{EntQ}(\mathcal{Q}, \mathcal{C}, \mathcal{T}) = \{q \in \mathcal{Q} \mid \mathcal{T} \cup \mathcal{C} \models q\}$.

Proposition 3. *Let $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A}\rangle$ be a CQE instance, $\mathcal{Q} = \langle q_1, \ldots, q_n\rangle$ (with $n \geq 0$) be a sequence of BUCQs, and $\mathcal{S} = \langle \mathcal{E}, \mathcal{Q}\rangle$. There exists no censor $\mathcal{C} \in \mathit{OptCens}(\mathcal{E})$ such that $\mathit{EntQ}(\mathcal{S}) \subset \mathit{EntQ}(\mathcal{Q}, \mathcal{C}, \mathcal{T})$.*

Proof. By contradiction, let such a censor $\mathcal{C}$ exist and let i be the least index such that $\mathcal{T} \cup \mathcal{C} \models q_i$ and $q_i \notin \mathit{EntQ}(\mathcal{S})$. By the minimality of i we have that, for all $j \in \{1, \ldots, i-1\}$, $\mathcal{T} \cup \mathcal{C} \models q_j$ iff $q_j \in \mathit{EntQ}(\langle \mathcal{E}, \langle q_1, \ldots, q_{i-1}\rangle\rangle)$. It follows that $\mathcal{C} \in \mathit{StCens}(\langle \mathcal{E}, \langle q_1, \ldots, q_{i-1}\rangle\rangle)$. But then, by definition, we should have that $\mathcal{C} \in \mathit{StCens}(\langle \mathcal{E}, \langle q_1, \ldots, q_i\rangle\rangle)$, and, consequently, that $q_i \in \mathit{EntQ}(\langle \mathcal{E}, \langle q_1, \ldots, q_i\rangle\rangle) \subseteq \mathit{EntQ}(\mathcal{S})$ (a contradiction). □

Moreover, dynCQE is the *only* way to guarantee that such optimality is preserved in the future. One might object that answering the current query q honestly may prevent the system from answering honestly another set of queries $\mathcal{Q}'$ in the future. However, the queries in $\mathcal{Q}'$ might never be submitted, so any censor that conceals the answer to q now might remain sub-optimal in the future. This may happen no matter how many additional queries are submitted by the users. Formally, we have:

Proposition 4. *Let $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ be a CQE instance, $\mathcal{Q} = \langle q_1, \ldots, q_n \rangle$ be a sequence of BUCQs, and $\mathcal{S} = \langle \mathcal{E}, \mathcal{Q} \rangle$. For all BUCQs q_{n+1}, and for all censors $\mathcal{C}$ in $StCens(\mathcal{S}) \setminus StCens(\langle \mathcal{E}, \mathcal{Q} \circ \langle q_{n+1} \rangle \rangle)$[3], there exist queries $q_{n+2}, q_{n+3}, \ldots, q_{n+k}, \ldots$ such that $EntQ(\langle q_1, \ldots, q_i \rangle, \mathcal{C}, \mathcal{T}) \subset EntQ(\langle \mathcal{E}, \langle q_1, \ldots, q_i \rangle \rangle)$ for all $i > n$.*

In the above proposition, the hypothesis $\mathcal{C} \in StCens(\mathcal{S}) \setminus StCens(\langle \mathcal{E}, \mathcal{Q} \circ \langle q_{n+1} \rangle \rangle)$ implies that q_{n+1} can be given a positive answer without disclosing any protected data, but $\mathcal{C}$ does not allow a positive answer to q_{n+1}.

Another property of dynamic CQE is that the first answer modification occurs as late as possible (*longest honeymoon* property). The following notion of maximal cooperativity implies and strengthens the longest honeymoon property.

Definition 4 (Cooperativity). *Let $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ be a CQE instance, $\mathcal{Q} = \langle q_1, \ldots, q_n \rangle$ (with $n \geq 0$) a sequence of BUCQs, and $\mathcal{C}$ and $\mathcal{C}'$ two censors for $\mathcal{E}$. We say that $\mathcal{C}$ is* more cooperative than *$\mathcal{C}'$ with respect to $\mathcal{Q}$ if there exists a non-negative natural number $m < n$ such that*

- *$\mathcal{T} \cup \mathcal{C} \models q_i \iff \mathcal{T} \cup \mathcal{C}' \models q_i$ for every $1 \leq i \leq m$, and*
- *$\mathcal{T} \cup \mathcal{C} \models q_{m+1}$ and $\mathcal{T} \cup \mathcal{C}' \not\models q_{m+1}$.*

We also say that $\mathcal{C}$ is maximally cooperative *with respect to $\mathcal{Q}$ if there does not exist any censor $\mathcal{C}''$ for $\mathcal{E}$ that is more cooperative than $\mathcal{C}$.*

The following intermediate result shows that a state of a CQE instance cannot discriminate between two optimal censors if they have answered all the queries posed so far in the same way.

Lemma 1. *Let $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ be a CQE instance, $\mathcal{Q} = \langle q_1, \ldots, q_n \rangle$ (with $n \geq 0$) be a sequence of BUCQs, and $\mathcal{C}$ and $\mathcal{C}'$ be two optimal censors for $\mathcal{E}$ such that $\mathcal{T} \cup \mathcal{C} \models q_i \iff \mathcal{T} \cup \mathcal{C}' \models q_i$, for all $i \in \{1, \ldots, n\}$. Then, $\mathcal{C} \in StCens(\langle \mathcal{E}, \mathcal{Q} \rangle)$ iff $\mathcal{C}' \in StCens(\langle \mathcal{E}, \mathcal{Q} \rangle)$.*

Proof. The proof is by induction on the length of $\mathcal{Q}$.

Case $n = 0$. Since $\mathcal{Q}$ is empty, both $\mathcal{C}$ and $\mathcal{C}'$ are in $StCens(\langle \mathcal{E}, \mathcal{Q} \rangle)$.

Case $n \geq 1$. In this case $\mathcal{Q} = \mathcal{Q}' \circ \langle q_n \rangle$, where $\mathcal{Q}' = \langle q_1, \ldots, q_{n-1} \rangle$. From the assumption $\mathcal{T} \cup \mathcal{C} \models q_i$ iff $\mathcal{T} \cup \mathcal{C}' \models q_i$, for all $i \in \{1, \ldots, n\}$, the following two facts hold: (*i*) $\mathcal{C} \in OptCens(\mathcal{E}, q_n)$ iff $\mathcal{C}' \in OptCens(\mathcal{E}, q_n)$; (*ii*) by IH, $\mathcal{C} \in StCens(\langle \mathcal{E}, \mathcal{Q}' \rangle)$ iff $\mathcal{C}' \in StCens(\langle \mathcal{E}, \mathcal{Q}' \rangle)$. Then, since $StCens(\langle \mathcal{E}, \mathcal{Q} \rangle)$ is by Definition 3 equal either to $StCens(\langle \mathcal{E}, \mathcal{Q}' \rangle)$ or to $StCens(\langle \mathcal{E}, \mathcal{Q}' \rangle) \cap OptCens(\mathcal{E}, q_n)$, we have the thesis. □

Then, we prove that for all states $\mathcal{S} = \langle \mathcal{E}, \mathcal{Q} \rangle$ of a CQE instance, the set $StCens(\mathcal{S})$ coincides with the set of all censors that are maximally cooperative with respect to $\mathcal{Q}$.

Theorem 1. *Let $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ be a CQE instance, and $\mathcal{Q} = \langle q_1, \ldots, q_n \rangle$ (with $n \geq 0$) be a sequence of BUCQs. A censor $\mathcal{C}$ for $\mathcal{E}$ is maximally cooperative with respect to $\mathcal{Q}$ iff $\mathcal{C} \in StCens(\langle \mathcal{E}, \mathcal{Q} \rangle)$.*

[3] With $\mathcal{Q} \circ \langle q_{n+1} \rangle$ we denote the sequence $\langle q_1, \ldots, q_n, q_{n+1} \rangle$.

Proof. We start by showing that every $\mathcal{C} \in \mathit{StCens}(\langle \mathcal{E}, \mathcal{Q} \rangle)$ is maximally cooperative with respect to $\mathcal{Q}$. Let $\mathcal{S}_h = \langle \mathcal{E}, \langle q_1, \ldots, q_h \rangle \rangle$, with $h \leq n$, and assume by contradiction that, for some $\mathcal{C} \in \mathit{StCens}(\mathcal{S}_n)$, there exists an optimal censor $\mathcal{C}'$ and a number $m < n$ such that (i) $\mathcal{T} \cup \mathcal{C} \models q_i \iff \mathcal{T} \cup \mathcal{C}' \models q_i$, for each $i \leq m$, and (ii) $\mathcal{T} \cup \mathcal{C} \not\models q_{m+1}$ and $\mathcal{T} \cup \mathcal{C}' \models q_{m+1}$.

Note that the sets $\mathit{StCens}(\mathcal{S}_h)$ form by construction a descending $\subseteq$-chain, hence $\mathcal{C}$ is in $\mathit{StCens}(\mathcal{S}_m)$. Then, from (i) and Lemma 1, $\mathcal{C}' \in \mathit{StCens}(\mathcal{S}_m)$ too.

From (ii) we have that $\mathcal{C}'$ occurs in $\mathit{OptCens}(\mathcal{E}, q_{m+1})$ whereas $\mathcal{C}$ does not. Then, on the one hand, since $\mathcal{C}' \in \mathit{StCens}(\mathcal{S}_m) \cap \mathit{OptCens}(\mathcal{E}, q_{m+1})$, $\mathit{StCens}(\mathcal{S}_{m+1})$ is equal by definition to $\mathit{StCens}(\mathcal{S}_m) \cap \mathit{OptCens}(\mathcal{E}, q_{m+1})$. On the other hand, $\mathit{StCens}(\mathcal{S}_{m+1})$ does not contain $\mathcal{C}$, as $\mathcal{C}$ is not in $\mathit{OptCens}(\mathcal{E}, q_{m+1})$. But this means that also $\mathit{StCens}(\mathcal{S}_n)$ does not contain $\mathcal{C}$, a contradiction.

Now, we show that if a censor $\mathcal{C}$ for $\mathcal{E}$ is maximally cooperative w.r.t. $\mathcal{Q}$, then $\mathcal{C} \in \mathit{StCens}(\langle \mathcal{E}, \mathcal{Q} \rangle)$. By contradiction, assume that $\mathcal{C} \notin \mathit{StCens}(\langle \mathcal{E}, \mathcal{Q} \rangle)$. So, there exists in $\mathcal{Q} = \langle q_1, \ldots, q_n \rangle$ a query q_i such that $\mathcal{C} \in \mathit{StCens}(\langle \mathcal{E}, \langle q_1, ..., q_{i-1} \rangle \rangle) \setminus \mathit{StCens}(\langle \mathcal{E}, \langle q_1, ..., q_i \rangle \rangle)$. Hence, there exists a censor $\mathcal{C}' \in \mathit{StCens}(\langle \mathcal{E}, \langle q_1, ..., q_i \rangle \rangle)$ such that $\mathcal{T} \cup \mathcal{C}' \models q_i$, while $\mathcal{T} \cup \mathcal{C} \not\models q_i$ and such that $\mathcal{T} \cup \mathcal{C}' \models q_j \iff \mathcal{T} \cup \mathcal{C} \models q_j$ for every $1 \leq j \leq i-1$. So, by Definition 4, $\mathcal{C}'$ is more cooperative than $\mathcal{C}$, which contradicts the fact that $\mathcal{C}$ is maximally cooperative. □

We conclude this section by comparing our new semantics of entailment with some other semantics from the literature. A first proposed strategy is arbitrarily choosing an optimal censor [6,13,14]. In this case, it might happen, as also stated by Proposition 4, that one looses optimality with respect to the state $\mathcal{S}$. For instance, if one arbitrarily picks censor $\mathcal{C}_2$ in Example 1, then $\mathit{EntQ}(\mathcal{Q}, \mathcal{C}_2, \mathcal{T}) \subset \mathit{EntQ}(\mathcal{S})$. On the other hand, when the chosen censor $\mathcal{C}$ turns out to be optimal with respect to a state $\mathcal{S}$, then, due to Theorem 1, either $\mathcal{C} \in \mathit{StCens}(\mathcal{S})$ or $\mathcal{C}$ is not maximally cooperative with respect to $\mathcal{Q}$.

Other two CQE semantics proposed in literature are: (i) *skeptical reasoning* [13,17], where a query q is entailed by a CQE instance $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$, denoted by $\mathcal{E} \models q$, if it is entailed by all the optimal censors for $\mathcal{E}$ together with the TBox, i.e., $\mathcal{T} \cup \mathcal{C} \models q$ for each $\mathcal{C} \in \mathit{OptCens}(\mathcal{E})$, and (ii) its approximation, called IGA semantics [10], under which q is entailed – in symbols, $\mathcal{E} \models_{\mathsf{IGA}} q$ – if it is entailed by $\mathcal{T} \cup \mathcal{C}_{\mathsf{IGA}}$, where $\mathcal{C}_{\mathsf{IGA}}$ is the intersection of all the optimal censors for $\mathcal{E}$, i.e, $\mathcal{C}_{\mathsf{IGA}} = \bigcap_{\mathcal{C} \in \mathit{OptCens}(\mathcal{E})} \mathcal{C}$. The following proposition shows that skeptically reasoning over all optimal censors is always a sound approximation of dynCQE.

Proposition 5. *Let $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ be a CQE instance, $\mathcal{Q} = \langle q_1, \ldots, q_n \rangle$ (with $n \geq 0$) be a sequence of BUCQs, and q be a BCQ in $\mathcal{Q}$. We have that $\mathcal{E} \models_{\mathsf{IGA}} q \implies \mathcal{E} \models q \implies \langle \mathcal{E}, \mathcal{Q} \rangle \models q$. The converse does not necessarily hold.*

Proof. Suppose that $\mathcal{E} \models_{\mathsf{IGA}} q$. By [10, Proposition 1], we already know that $\mathcal{E} \models q$. Now, since $\mathcal{E} \models q$ by definition means that $\mathcal{T} \cup \mathcal{C} \models q$ holds for each $\mathcal{C} \in \mathit{StCens}(\mathcal{S}) \subseteq \mathit{OptCens}(\mathcal{E})$, we trivially have that $\langle \mathcal{E}, \mathcal{Q} \rangle \models q$.

As for the converse, consider Example 1. We have that $\langle \mathcal{E}, \mathcal{Q} \rangle \models q_1$ but $\mathcal{E} \not\models q_1$ (and thus, also $\mathcal{E} \not\models_{\mathsf{IGA}} q_1$) because $\mathcal{T} \cup \mathcal{C}_3 \not\models q_1$. □

4 First-order Rewritability of Query Entailment

We now move to the study of computational complexity of query entailment. In this investigation, we focus on *DL-Lite*$_R$ CQE specifications, i.e., whose TBox and ABox are expressed in *DL-Lite*$_R$.

A first way to solve query entailment in a state might consist in finding a reduction to the stateless CQE approach, for which algorithms are already known. It turns out, however, that the behavior of dynCQE cannot be intensionally simulated by a stateless CQE instance, independent of query history.

Theorem 2. *There exist a DL-Lite*$_R$ *CQE specification* $\langle \mathcal{T}, \mathcal{P} \rangle$ *and a BUCQ* q *such that there exist no DL-Lite*$_R$ *CQE specification* $\langle \mathcal{T}', \mathcal{P}' \rangle$ *such that, for every ABox* $\mathcal{A}$, $\mathit{OptCens}(\langle \mathcal{T}', \mathcal{P}', \mathcal{A} \rangle) = \mathit{StCens}(\mathcal{S})$, *where* $\mathcal{S} = \langle \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle, \langle q \rangle \rangle$.

Proof. Let $\mathcal{T} = \emptyset$, let $\mathcal{P} = \{C(x) \wedge D(x) \rightarrow \bot\}$, and let $q = \exists x C(x)$. By contradiction, suppose there exist a TBox $\mathcal{T}'$ and a policy $\mathcal{P}'$ such that, for every ABox $\mathcal{A}$, $\mathit{OptCens}(\langle \mathcal{T}', \mathcal{P}', \mathcal{A} \rangle) = \mathit{StCens}(\mathcal{S})$.

Now consider the ABox $\mathcal{A} = \{C(a_1), C(a_2), D(a_1), D(a_2)\}$, where a_1, a_2 are individual names that do not appear in $\mathcal{P}'$. The optimal censors for $\langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ are $\mathcal{C}_1 = \{C(a_1), C(a_2)\}$, $\mathcal{C}_2 = \{C(a_1), D(a_2)\}$, $\mathcal{C}_3 = \{D(a_1), C(a_2)\}$, $\mathcal{C}_4 = \{D(a_1), D(a_2)\}$. Among such optimal censors, only $\mathcal{C}_4$ does not satisfy q. Therefore, $\mathit{StCens}(\mathcal{S}) = \{\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3\}$. Since by hypothesis $\mathit{StCens}(\mathcal{S}) = \mathit{OptCens}(\langle \mathcal{T}', \mathcal{P}', \mathcal{A} \rangle)$, it follows that $\mathcal{T}' \cup \mathcal{P}' \cup \mathcal{C}_4$ is inconsistent and $\mathcal{T}' \cup \mathcal{P}' \cup \mathcal{C}_3$ is consistent. Consequently, by Proposition 2, $\mathit{PerfectRef}(q(\mathcal{P}'), \mathcal{T}')$ evaluates to true in $\mathcal{C}_4$ and evaluates to false in $\mathcal{C}_3$.

On the other hand, it is immediate to see that, for every BUCQ q that does not mention individual names in $\mathcal{A}$, q evaluates to true in $\mathcal{C}_4$ only if q evaluates to true in $\mathcal{C}_3$. Consequently, $\mathit{PerfectRef}(q(\mathcal{P}'), \mathcal{T}')$ evaluates to true in $\mathcal{C}_4$ only if $\mathit{PerfectRef}(q(\mathcal{P}'), \mathcal{T}')$ evaluates to true in $\mathcal{C}_3$. Thus we get a contradiction. □

We now study the data complexity of the query entailment problem in a state, i.e., given a state $\mathcal{S} = \langle \mathcal{E}, \mathcal{Q} \rangle$ of a CQE instance $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$, the problem of checking whether a BUCQ q in $\mathcal{Q}$ belongs to $\mathit{EntQ}(\mathcal{S})$. In particular, we prove that this problem is FO rewritable, and, so, that it is in AC^0 in data complexity.

We start by showing a fundamental property of query entailment in a state, which holds for all DLs.

Theorem 3. *Let* $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ *be a CQE instance,* $\mathcal{Q} = \langle q_1, \ldots, q_n \rangle$ *be a sequence of BUCQs, and let* $\mathcal{S} = \langle \mathcal{E}, \mathcal{Q} \rangle$. *For every* i *such that* $1 \leq i \leq n$, $q_i \in \mathit{EntQ}(\mathcal{S})$ *iff there exists a censor* $\mathcal{C}$ *for* $\mathcal{E}$ *such that*

$$\mathcal{T} \cup \mathcal{C} \models \Big(\bigwedge_{q \in \mathit{EntQ}(\mathcal{S}_{i-1})} q \Big) \wedge q_i$$

Proof. ($\Leftarrow$:) Suppose there exists a censor $\mathcal{C}$ for $\mathcal{E}$ such that $\mathcal{T} \cup \mathcal{C} \models (\bigwedge_{q \in \mathit{EntQ}(\mathcal{S}_{i-1})} q) \wedge q_i$. Then, it follows immediately that there exists an optimal censor $\mathcal{C}'$ for $\mathcal{E}$ such that $\mathcal{C}' \supset \mathcal{C}$, consequently $\mathcal{T} \cup \mathcal{C}' \models (\bigwedge_{q \in \mathit{EntQ}(\mathcal{S}_{i-1})} q) \wedge q_i$. Hence, by Definition 3, $\mathcal{C}' \in \mathit{StCens}(\langle \mathcal{E}, \langle q_1, \ldots, q_i \rangle \rangle)$. Therefore, $q_i \in \mathit{EntQ}(\mathcal{S})$.

($\Rightarrow$:) Suppose $q_i \in \mathit{EntQ}(\mathcal{S})$. Now, let $\mathcal{C}'$ be an optimal censor for $\mathcal{E}$ such that $\mathcal{C}' \in \mathit{StCens}(\mathcal{S})$. We have that $\mathcal{T} \cup \mathcal{C}' \models q$ for every $q \in \mathit{EntQ}(\mathcal{S})$, and since $q_i \in \mathit{EntQ}(\mathcal{S})$ and $\mathit{EntQ}(\mathcal{S}_{i-1}) \subseteq \mathit{EntQ}(\mathcal{S})$, it follows that $\mathcal{T} \cup \mathcal{C}' \models (\bigwedge_{q \in \mathit{EntQ}(\mathcal{S}_{i-1})} q) \wedge q_i$, thus proving the thesis. □

Given a BUCQ q and an ABox $\mathcal{A}$, we say that an *image of q in $\mathcal{A}$* is a minimal subset $\mathcal{A}'$ of $\mathcal{A}$ such that $\mathcal{A}' \models q$. Furthermore, given a BUCQ q, a TBox $\mathcal{T}$ and an ABox $\mathcal{A}$, we say that an *image of q in $\mathcal{A}$ with respect to $\mathcal{T}$* is a minimal subset $\mathcal{A}'$ of $\mathcal{A}$ such that $\mathcal{T} \cup \mathcal{A}' \models q$.

Theorem 4. *Let $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ be a DL-Lite$_R$ CQE instance and $\mathcal{Q} = \langle q_1, \ldots, q_n \rangle$ (with $n \geq 0$) be a sequence of BUCQs. For every i such that $1 \leq i \leq n$, $q_i \in \mathit{EntQ}(\mathcal{S})$ iff there exists an image IM of $\mathit{PerfectRef}((\bigwedge_{q \in \mathit{EntQ}(S_{i-1})} q) \wedge q_i, \mathcal{T})$ in $\mathsf{cl}_\mathcal{T}(\mathcal{A})$ such that $\mathit{PerfectRef}(q(\mathcal{P}), \mathcal{T})$ evaluates to false in IM.*

Now observe that: (*i*) $\mathsf{cl}_\mathcal{T}(\mathcal{A})$ can be computed in PTIME w.r.t. data complexity; (*ii*) every image of a BUCQ q has a size that is not larger than the length of the longest BCQ in q; (*iii*) such a maximum length is a constant w.r.t. data complexity; (*iv*) all the conditions in the theorem can be verified in PTIME with respect to data complexity [9]. This implies that the entailment problem in a state can be decided in PTIME w.r.t. data complexity.

In the following, we provide a tighter upper bound, showing that this entailment problem is in AC^0 in data complexity. We do so by proving that the problem is FO rewritable. That is, for every BUCQ q of the state, there exists an FO query q' that does not depend on the ABox and is such that q is entailed in the state iff q' evaluates to true in the ABox.

To this purpose, we will find an FO query that depends on the intensional part of the state, i.e., the TBox, the policy and the sequence of queries, and such that its evaluation on the ABox is true if and only if the condition expressed in Theorem 4 holds (Theorem 7). We will make two intermediate steps towards this result: first (Theorem 5), given a query q on a *DL-Lite$_R$* CQE specification $\langle \mathcal{T}, \mathcal{P} \rangle$, we will find a query denoted by $\mathit{BraveRef}(q, \mathcal{T}, \mathcal{P})$ whose evaluation on $\mathsf{cl}_\mathcal{T}(\mathcal{A})$ corresponds to checking the existence of an optimal censor $\mathcal{C}$ for the CQE instance $\langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ such that $\mathcal{T} \cup \mathcal{C} \models q$; then (Theorem 6), we will find an FO query such that its evaluation on $\mathsf{cl}_\mathcal{T}(\mathcal{A})$ is true if and only if the condition expressed in Theorem 4 holds.

Given two BCQs q and q', a *mapping of q' into q* is a function $h : \mathit{Atoms}(q') \to \mathit{Atoms}(q)$ such that there exists a most general unifier σ_h such that, for every atom $\alpha \in \mathit{Atoms}(q')$, $\sigma_h(\alpha) = \sigma_h(h(\alpha))$. Such a most general unifier (variable substitution) assigns variables occurring either in q' or in q to either variables of q or constants. We denote by $\mathit{Map}(q', q)$ the set of all the mappings of q' into q.

Furthermore, we denote by $\sigma_h[q]$ the variable substitutions of σ_h limited to variables occurring in q. For instance, if $q = \exists x, y, z R(x, y, z)$, $q' = \exists x' R(x', x', a)$ (where a is a constant and all other arguments are variables), then $\sigma_h = \{x' \leftarrow x, y \leftarrow x, z \leftarrow a\}$ and $\sigma_h[q] = \{y \leftarrow x, z \leftarrow a\}$.

Given two BCQs q and q', we denote by $\mathit{Unify}(q, q')$ the formula:

$$\bigvee_{h \in Map(q',q)} \Big(\bigwedge_{x \leftarrow t \, \in \sigma_h[q]} x = t \Big)$$

Definition 5. *Given a BUCQ q, a DL-Lite$_R$ TBox $\mathcal{T}$ and a policy $\mathcal{P}$, we define $\mathit{BraveRef}(q, \mathcal{T}, \mathcal{P})$ as the FO sentence:*

$$\bigvee_{q_r \in PerfectRef(q,\mathcal{T})} \exists \vec{x}_r \big(conj_r(\vec{x}_r) \wedge \neg \big(\bigwedge_{q_d \in PerfectRef(q(\mathcal{P}),\mathcal{T})} \mathit{Unify}(q_r, q_d) \big) \big)$$

(where we assume $q_r = \exists \vec{x}_r (conj_r(\vec{x}_r))$).

We now establish the fundamental property of the above query reformulation function *BraveRef*.

Theorem 5. *Let $\langle \mathcal{T}, \mathcal{P} \rangle$ be a DL-Lite$_R$ CQE specification. For every ABox $\mathcal{A}$, there exists an optimal censor $\mathcal{C}$ for $\langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ such that $\mathcal{T} \cup \mathcal{C} \models q$ iff $\mathit{BraveRef}(q, \mathcal{T}, \mathcal{P})$ evaluates to true in $\mathsf{cl}_{\mathcal{T}}(\mathcal{A})$.*

Then, we use *BraveRef* to define the new query reformulation function *StateRef* as follows.

Definition 6. *Let $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ be a DL-Lite$_R$ CQE instance, $\mathcal{Q} = \langle q_1, \ldots, q_n \rangle$ (with $n \geq 0$) be a sequence of BUCQs, let i be such that $1 \leq i \leq n$, and let $I \subseteq \{1, \ldots, i-1\}$: I represents the set of indexes of the queries that precede query q_i in $\mathcal{Q}$ and that are guessed to be true in the state $\mathcal{S} = \langle \mathcal{E}, \mathcal{Q} \rangle$. We define $\mathit{StateRef}(\mathcal{S}, i, I)$ as the FO sentence:*

$$\Big(\bigwedge_{\substack{1 \leq j \leq i-1 \\ \wedge \, j \notin I}} \neg \mathit{BraveRef}\big(\big(\bigwedge_{\ell \in I \, \wedge \, \ell < j} q_\ell \big) \wedge q_j, \mathcal{T}, \mathcal{P} \big) \Big) \wedge \mathit{BraveRef}\big(\big(\bigwedge_{\ell \in I} q_\ell \big) \wedge q_i, \mathcal{T}, \mathcal{P} \big)$$

As an example, consider the *DL-Lite$_R$* CQE instance $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ and the query sequence $\mathcal{Q} = \langle q_1, q_2, q_3 \rangle$ of Example 1, and let us set $i = 3$ and $I = \{1\}$. We have that $\mathit{StateRef}(\langle \mathcal{E}, \mathcal{Q} \rangle, i, I)$ is the FO sentence $\neg \mathit{BraveRef}(q_1 \wedge q_2, \mathcal{T}, \mathcal{P}) \wedge \mathit{BraveRef}(q_1 \wedge q_3, \mathcal{T}, \mathcal{P}) = \neg(\mathsf{buy}(\mathsf{john}, \mathsf{m}_a) \wedge \mathsf{Abc}(\mathsf{m}_a) \wedge \neg(\exists z, w(\mathsf{buy}(z, w) \wedge \mathsf{Abc}(w) \wedge z = \mathsf{john} \wedge w = \mathsf{m}_a))) \wedge \exists x(\mathsf{buy}(\mathsf{john}, \mathsf{m}_a) \wedge \mathsf{buy}(x, \mathsf{m}_b))$.

The query reformulation function *StateRef* allows for reducing query entailment in a state to evaluating an FO query, as stated by the following property.

Theorem 6. *Let $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ be a DL-Lite$_R$ CQE instance, $\mathcal{Q} = \langle q_1, \ldots, q_n \rangle$ (with $n \geq 0$) be a sequence of BUCQs. For every i such that $1 \leq i \leq n$, $q_i \in \mathit{EntQ}(\mathcal{S})$ iff the following FO sentence evaluates to true in $\mathsf{cl}_{\mathcal{T}}(\mathcal{A})$:*

$$\bigvee_{I \in \wp(\{1,\ldots,i-1\})} \mathit{StateRef}(\mathcal{S}, i, I),$$

where $\wp(\{1, \ldots, i-1\})$ denotes the powerset of $\{1, \ldots, i-1\}$.

The last two theorems show the FO rewritability of the problems studied on $\mathsf{cl}_{\mathcal{T}}(\mathcal{A})$. We now modify the respective reformulations to evaluate them directly on the ABox $\mathcal{A}$ and thus produce "genuine" FO rewritability results.

In what follows we will make use of the algorithm *AtomRewr* provided in [13], that we now briefly describe. Given an FO sentence ϕ and a *DL-Lite*$_R$ TBox $\mathcal{T}$, $AtomRewr(\phi, \mathcal{T})$ computes the FO sentence obtained from ϕ by replacing every atom $\alpha = p(\vec{x})$ (where $\vec{x}$ are all the variables occurring in α) with the disjunction of atoms corresponding to the perfect rewriting of the non-Boolean atomic query $q_\alpha = \{\vec{x} \mid p(\vec{x})\}$ with respect to $\mathcal{T}$.

For our purposes, we recall the key property of *AtomRewr* provided in [13].

Proposition 6. *For every FO sentence ϕ, DL-Lite$_R$ TBox $\mathcal{T}$, and ABox $\mathcal{A}$, ϕ evaluates to true in $\mathsf{cl}_{\mathcal{T}}(\mathcal{A})$ iff $AtomRewr(\phi, \mathcal{T})$ evaluates to true in $\mathcal{A}$.*

Now, Proposition 6 and Theorem 6 immediately imply the next property.

Theorem 7. *Let $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ be a DL-Lite$_R$ CQE instance, $\mathcal{Q} = \langle q_1, \ldots, q_n \rangle$ be a sequence of BUCQs. For every i such that $1 \leq i \leq n$, $q_i \in EntQ(\mathcal{S})$ iff the following FO sentence evaluates to true in $\mathcal{A}$:*

$$AtomRewr\Big(\bigvee_{I \in \wp(\{1,\ldots,i-1\})} StateRef(\mathcal{S}, i, I), \mathcal{T}\Big)$$

The previous theorem shows the FO rewritability of the problem of entailment of BUCQs in a state.

Example 2. Let $\mathcal{E}$ and $\mathcal{Q} = \langle q_1, q_2, q_3 \rangle$ be as in Example 1. According to Theorem 7, the query $q_3 = \exists x \mathsf{buy}(x, \mathsf{m}_b)$ belongs to $EntQ(\langle \mathcal{E}, \mathcal{Q} \rangle)$ if and only if the FO sentence below evaluates to true in $\mathcal{A}$ (f_I denotes the sub-formula considering the guess I of the indexes of the queries that precede the query q_3):

$$\begin{array}{l|l}
 & AtomRewr(\bigvee_{I \in \wp(\{1,2\})} StateRef(\langle \mathcal{E}, \mathcal{Q} \rangle, i, I), \mathcal{T}) = \\
f_{I=\emptyset} & \neg BraveRef(q_1, \mathcal{T}, \mathcal{P}) \wedge \neg BraveRef(q_2, \mathcal{T}, \mathcal{P}) \wedge BraveRef(q_3, \mathcal{T}, \mathcal{P}) \vee \\
f_{I=\{1\}} & \neg BraveRef(q_1 \wedge q_2, \mathcal{T}, \mathcal{P}) \wedge BraveRef(q_1 \wedge q_3, \mathcal{T}, \mathcal{P}) \vee \\
f_{I=\{2\}} & \neg BraveRef(q_1, \mathcal{T}, \mathcal{P}) \wedge BraveRef(q_2 \wedge q_3, \mathcal{T}, \mathcal{P}) \vee \\
f_{I=\{1,2\}} & BraveRef(q_1 \wedge q_2 \wedge q_3, \mathcal{T}, \mathcal{P}) = \\
f_{I=\emptyset} & \neg \mathsf{buy}(\mathsf{john}, \mathsf{m}_a) \wedge \neg \mathsf{Abc}(\mathsf{m}_a) \wedge \exists x \mathsf{buy}(x, \mathsf{m}_b) \vee \\
f_{I=\{1\}} & \neg(\mathsf{buy}(\mathsf{john}, \mathsf{m}_a) \wedge \mathsf{Abc}(\mathsf{m}_a) \wedge \neg(\exists z, w(\mathsf{buy}(z, w) \wedge \mathsf{Abc}(w) \wedge \\
 & \quad z = \mathsf{john} \wedge w = \mathsf{m}_a))) \wedge \exists x(\mathsf{buy}(\mathsf{john}, \mathsf{m}_a) \wedge \mathsf{buy}(x, \mathsf{m}_b)) \vee \\
f_{I=\{2\}} & \neg(\mathsf{buy}(\mathsf{john}, \mathsf{m}_a)) \wedge (\exists x(\mathsf{Abc}(\mathsf{m}_a) \wedge \mathsf{buy}(x, \mathsf{m}_b))) \vee \\
f_{I=\{1,2\}} & \exists x(\mathsf{buy}(\mathsf{john}, \mathsf{m}_a) \wedge \mathsf{Abc}(\mathsf{m}_a) \wedge \mathsf{buy}(x, \mathsf{m}_b)) \wedge \\
 & \quad \neg(\exists z, w(\mathsf{buy}(z, w) \wedge \mathsf{Abc}(w) \wedge z = \mathsf{john} \wedge w = \mathsf{m}_a))
\end{array}$$

which, indeed, evaluates to true in $\mathcal{A}$ thanks to $f_{I=\{1\}}$. □

5 Towards Practical Techniques and Approximations

We now provide a simplification of the query rewriting presented in Theorem 7. In particular, in a real maximally collaborative CQE system, the answers to

the queries already executed (i.e., the queries belonging to the state) can obviously be stored and re-used when the next query is submitted. This allows for greatly simplifying the structure of the FO reformulation of the query defined in Theorem 7, as shown in the following.

Theorem 8. *Let $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ be a DL-Lite$_R$ CQE instance, $\mathcal{Q} = \langle q_1, \ldots, q_n \rangle$ be a sequence of BUCQs, let $\mathcal{S} = \langle \mathcal{E}, \mathcal{Q} \rangle$, let q_{n+1} be a BUCQ, and let $\mathcal{S}' = \langle \mathcal{E}, \langle q_1, \ldots, q_n, q_{n+1} \rangle \rangle$. Then, q_{n+1} is entailed by $\mathcal{S}'$ iff the following FO sentence evaluates to true in $\mathcal{A}$:*

$$\mathit{AtomRewr}(\mathit{BraveRef}((\bigwedge_{q_i \in \mathit{EntQ}(\mathcal{S})} q_i) \wedge q_{n+1}, \mathcal{T}, \mathcal{P}), \mathcal{T})$$

Proof. Suppose $\mathcal{S}' \models q_{n+1}$, i.e. $\mathit{EntQ}(\mathcal{S}') = \mathit{EntQ}(\mathcal{S}) \cup \{q_{n+1}\}$. By Theorem 7, the sentence $\psi = \mathit{AtomRewr}(\mathit{StateRef}(\mathcal{S}', n+1, I), \mathcal{T})$ evaluates to true in $\mathcal{A}$, where $I = \{i \mid q_i \in \mathit{EntQ}(\mathcal{S}')\}$. Consequently, the sentence $\mathit{AtomRewr}(\mathit{BraveRef}((\bigwedge_{q_i \in \mathit{EntQ}(\mathcal{S})} q_i) \wedge q_{n+1}, \mathcal{T}, \mathcal{P}), \mathcal{T})$ is equal to the last conjunct of ψ, and therefore evaluates to true in $\mathcal{A}$ as well.

Suppose now $\mathcal{S}' \not\models q_{n+1}$. From Theorem 7, we have that the sentence $\mathit{AtomRewr}(\mathit{StateRef}(\mathcal{S}', n+1, I), \mathcal{T})$ evaluates to false in $\mathcal{A}$, where $I = \{i \mid q_i \in \mathit{EntQ}(\mathcal{S})\}$. Since $\mathit{EntQ}(\mathcal{S})$ is the set of BUCQ from $\langle q_1, \ldots, q_n \rangle$ entailed by $\mathcal{S}$, all the conjuncts of $\mathit{AtomRewr}(\mathit{StateRef}(\mathcal{S}', n+1, I), \mathcal{T})$ except the last one evaluate to true in $\mathcal{A}$. This means that its last conjunct evaluates to false in $\mathcal{A}$. Such a conjunct is equal to the sentence $\mathit{AtomRewr}(\mathit{BraveRef}((\bigwedge_{q_i \in \mathit{EntQ}(\mathcal{S})} q_i) \wedge q_{n+1}, \mathcal{T}, \mathcal{P}), \mathcal{T})$, which proves the thesis. □

Example 3. Let $\mathcal{E}$ and the queries q_1, q_2, and q_3 be as in Example 1. Consider the sequence of queries $\mathcal{Q} = \langle q_1, q_2 \rangle$. From Example 1, we know that only $q_1 = \mathsf{buy}(\mathsf{john}, \mathsf{m}_a)$ belongs to $\mathit{EntQ}(\langle \mathcal{E}, \mathcal{Q} \rangle)$. Hence, according to Theorem 8, the query $q_3 = \exists x \mathsf{buy}(x, \mathsf{m}_b)$ is entailed by the state $\langle \mathcal{E}, \mathcal{Q} \circ \{q_3\} \rangle$ if and only if the FO sentence $\exists x(\mathsf{buy}(\mathsf{john}, \mathsf{m}_a) \wedge \mathsf{buy}(x, \mathsf{m}_b))$ evaluates to true in $\mathcal{A}$. □

An issue that the query rewriting technique of Theorem 8 does not solve is the scalability w.r.t. the number of submitted queries, which might become too large to make the FO query produced by the rewriting executable in practice. On the other hand, Theorem 2 shows that it is not always possible to intensionally simulate dynCQE by using a stateless CQE specification, i.e., through an ABox-independent transformation of the intensional part of a CQE instance.

To overcome the above issue, a possible approach is to materialize a censor $\mathcal{C}$ of the current state $\mathcal{S}$ of the CQE instance, and then evaluate the next queries over the ontology $\mathcal{T} \cup \mathcal{C}$. If the current state $\mathcal{S}$ has multiple censors, evaluating a query over $\mathcal{T} \cup \mathcal{C}$ is only an approximation of the query entailment through dynCQE, i.e., in the corresponding state. More precisely: as long as the materialized system processes only queries entailed by $\mathcal{T} \cup \mathcal{C}$ (i.e., it always answers "yes"), it returns exactly the same answers provided by dynCQE. The first time it processes a query q non-entailed by $\mathcal{T} \cup \mathcal{C}$ (i.e., it answers "no"), its behaviour might differ from the dynamic approach, where q might be either entailed or not

entailed (depending on how the censors of the states evolve). After the first negative answer, the system using $\mathcal{C}$ might answer "yes" (resp. "no") to a subsequent query q even if the state does not entail (resp. entail) q. Obviously, if the state $\mathcal{S}$ has the only censor $\mathcal{C}$, then $\mathcal{T} \cup \mathcal{C}$ and the dynCQE system will have the same behaviour. Below we describe how to materialize a censor of a state.

1. Split the FO query of Theorem 8, execute only one $q' \in \mathit{PerfectRef}(\mathit{EntQ}(\mathcal{S}) \cup \{q\}, \mathcal{T})$ at a time, and turn all the variables appearing in $\mathit{AtomRewr}(q', \mathcal{T})$ as free variables.
2. As soon as one of such queries is true in $\mathcal{A}$, we can construct (through the corresponding binding of the free variables of the query) an image of this query in $\mathcal{A}$. Let $\mathcal{A}'$ be such a subset of $\mathcal{A}$.
3. $\mathcal{P} \cup \mathcal{A}'$ is consistent, so there exists at least one censor $\mathcal{C}$ of $\mathcal{S}$ that contains $\mathcal{A}'$. One such censor can be computed by first setting $\mathcal{C} = \mathcal{A}'$, and then, as long as it possible, by iteratively adding to $\mathcal{C}$ ground atoms γ from $\mathsf{cl}_{\mathcal{T}}(\mathcal{A}) \setminus \mathcal{A}'$ such that $\mathcal{T} \cup \mathcal{P} \cup \mathcal{C} \cup \{\gamma\}$ is consistent.

6 Related Work

As shown in [11], the censors introduced in Definition 1 enjoy the *indistinguishability property*, that is, for all CQE instances $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A} \rangle$ and all censors $\mathcal{C}$ for $\mathcal{E}$, there exists an ABox $\mathcal{A}'$ that entails no secrets, such that $\mathcal{C}$ is also a censor for $\mathcal{E} = \langle \mathcal{T}, \mathcal{P}, \mathcal{A}' \rangle$. Such censors are called *indistinguishability-based* (IB) because the instances with $\mathcal{A}$ and $\mathcal{A}'$ cannot be distinguished based on the answers allowed by $\mathcal{C}$. IB censors are secure against attackers that know the censor's algorithm. In particular, even if the attackers could compute the ABoxes that yield $\mathcal{C}$, using their knowledge about the algorithm, the ABox $\mathcal{A}'$ would prevent them from inferring any secret.

Benedikt et al. [2] provide, for OBDA settings, a systematic complexity analysis of confidentiality preserving query answering based on indistinguishability. They do not address the issue of selecting a secure data disclosure among the available ones. IB censors in OBDA are also considered in [10], where a practical approach to skeptical reasoning in CQE is presented. Differently from our approach, in [10] censors do not take into account the history of the users' queries.

In [13], IB censors are compared with so-called *confidentiality preserving* (CP) censors, that in general do not enjoy the indistinguishability property. Moreover, [13] introduces algorithms and complexity results for skeptical reasoning in CQE, i.e., the problem of computing only the query answers that are returned by *all* IB censors. By definition, the skeptical CQE method is generally less cooperative than the dynamic method introduced and analyzed in this paper (Theorem 5). In [12], policies have been extended with numerical restrictions, and it is proved that this extension preserves FO rewritability.

The first IB CQE method for Description Logics was introduced in [8]. Its confidentiality model is more robust and general, as it takes into account both object-level and meta-level background knowledge of the attacker. However, CQ answering and FO rewritability are not addressed. Moreover, the *secure views* of

[8] are constructed from a sequence of queries that covers *all* possible relevant queries, while the properties we investigate here hold for arbitrary (possibly non-exhaustive) sequences of queries submitted by the users.

The issue of how to select an optimal censor has been tackled in [11]. The selection criterion is based on explicit preferences over predicates, that are specified together with the CQE instance. This approach, in general, is neither maximally cooperative nor optimal w.r.t. a given state, because the optimal censor is selected statically, in a stateless fashion. Moreover, the given preferences are not always able to select a single optimal censor.

Other CQE approaches based on censors, such as CP censors, in general do not enjoy the indistinguishability property [15,17], which makes them vulnerable to attacks based on knowledge of the CQE algorithm. Moreover, they do not address dynamic query-based censor selection. See [8] for a list of earlier approaches with similar features focused on publishing secure subsets of the ontology. Two nice abstract analyses of censors properties can be found in [22,21].

Finally, Cuenca Grau et al. [16] introduce and investigate an anonymization framework for knowledge graphs based on substituting nodes with blanks.

7 Conclusions

In this paper, we have presented a maximally cooperative approach to controlled query evaluation in OWL and Description Logic ontologies. We have shown that the approach is computationally not harder than the previous static and less cooperative approaches to CQE. Moreover, we have defined a new query rewriting algorithm to solve the query entailment problem in this framework.

The present work can be extended in several interesting directions. First, while the presented results indicate the possibility of a query rewriting approach to dynamic CQE, more work is still needed to define a practical query answering technique and to extend it to non-Boolean UCQs.

Then, the policy language adopted in this paper (set of denials) can be extended to encompass more expressive data protection policies. One step towards this direction, although in the context of static CQE, has been presented e.g. in [12]: it would be interesting to see whether dynCQE can also be extended in a similar way. Finally, it would be interesting to study the computational properties of dynamic CQE in ontology languages different from OWL 2 QL and *DL-Lite*$_R$, in particular in the other lightweight profiles of OWL 2.

Supplemental Material Statement: For complete proofs of our results we refer the reader to an extended version of the present paper [7].

Acknowledgements This work was partly supported by the EU within the Horizon Europe Programme under the Glaciation project (ref. no. 101070141) and within the H2020 Programme - ERA-NET Cofund ICT-AGRI-FOOD under the ADCATER Project (ref. no. 862665).

References

1. F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, 2003.
2. M. Benedikt, B. Cuenca Grau, and E. V. Kostylev. Logical foundations of information disclosure in ontology-based data integration. *Artif. Intell.*, 262:52–95, 2018.
3. J. Biskup. For unknown secrecies refusal is better than lying. *Data and Knowledge Engineering*, 33(1):1–23, 2000.
4. J. Biskup and P. A. Bonatti. Controlled query evaluation for enforcing confidentiality in complete information systems. *Int. J. of Information Security*, 3(1):14–27, 2004.
5. J. Biskup and P. A. Bonatti. Controlled query evaluation for known policies by combining lying and refusal. *Ann. Math. Artif. Intell.*, 40(1-2):37–62, 2004.
6. J. Biskup and P. A. Bonatti. Controlled query evaluation with open queries for a decidable relational submodel. *Ann. Math. Artif. Intell.*, 50(1–2):39–77, 2007.
7. P. Bonatti, G. Cima, D. Lembo, L. Marconi, R. Rosati, L. Sauro, and D. F. Savo. CQE in OWL 2 QL: A "longest honeymoon" approach (extended version). arXiv:2207.11155, 2022.
8. P. A. Bonatti and L. Sauro. A confidentiality model for ontologies. In *Proc. of ISWC 2013*, volume 8218 of *LNCS*, pages 17–32. Springer, 2013.
9. D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. Tractable reasoning and efficient query answering in description logics: The *DL-Lite* family. *J. of Automated Reasoning*, 39(3):385–429, 2007.
10. G. Cima, D. Lembo, L. Marconi, R. Rosati, and D. F. Savo. Controlled query evaluation in ontology-based data access. In *Proc. of ISWC 2020*, volume 12506 of *LNCS*, pages 128–146. Springer, 2020.
11. G. Cima, D. Lembo, L. Marconi, R. Rosati, and D. F. Savo. Controlled query evaluation over prioritized ontologies with expressive data protection policies. In *Proc. of ISWC 2021*, volume 12922 of *LNCS*, pages 374–391. Springer, 2021.
12. G. Cima, D. Lembo, L. Marconi, R. Rosati, D. F. Savo, and D. Sinibaldi. Controlled query evaluation over ontologies through policies with numerical restrictions. In *Proc. of AIKE 2021*, pages 33–36. IEEE, 2021.
13. G. Cima, D. Lembo, R. Rosati, and D. F. Savo. Controlled query evaluation in description logics through instance indistinguishability. In *Proc. of IJCAI 2020*, pages 1791–1797, 2020.
14. B. Cuenca Grau, E. Kharlamov, E. V. Kostylev, and D. Zheleznyakov. Controlled query evaluation over OWL 2 RL ontologies. In *Proc. of ISWC 2013*, volume 8218 of *LNCS*, pages 49–65. Springer, 2013.
15. B. Cuenca Grau, E. Kharlamov, E. V. Kostylev, and D. Zheleznyakov. Controlled query evaluation for datalog and OWL 2 profile ontologies. In *Proc. of IJCAI 2015*, pages 2883–2889, 2015.
16. B. Cuenca Grau and E. V. Kostylev. Logical foundations of linked data anonymisation. *J. Artif. Intell. Res.*, 64:253–314, 2019.
17. D. Lembo, R. Rosati, and D. F. Savo. Revisiting controlled query evaluation in description logics. In *Proc. of IJCAI 2019*, pages 1786–1792, 2019.
18. J. W. Lloyd. *Foundations of Logic Programming (Second, Extended Edition)*. Springer, Berlin, Heidelberg, 1987.

19. B. Motik, B. Cuenca Grau, I. Horrocks, Z. Wu, A. Fokoue, and C. Lutz. OWL 2 Web Ontology Language profiles (second edition). W3C Recommendation, W3C, Dec. 2012. Available at `http://www.w3.org/TR/owl2-profiles/`.
20. B. Motik, A. Fokoue, I. Horrocks, Z. Wu, C. Lutz, and B. Cuenca Grau. OWL Web Ontology Language profiles. W3C Recommendation, W3C, Oct. 2009. Available at `http://www.w3.org/TR/owl-profiles/`.
21. T. Studer. No-go theorems for data privacy. *CoRR*, abs/2005.13811, 2020.
22. T. Studer and J. Werner. Censors for boolean description logic. *Trans. Data Privacy*, 7(3):223–252, 2014.