An Algebraic Approach to the Complexity of Generalized Conjunctive Queries*

Michael Bauland[1], Philippe Chapdelaine[2], Nadia Creignou[3], Miki Hermann[4], and Heribert Vollmer[1]

[1] Theoretische Informatik, Universität Hannover, Germany. bauland|vollmer@thi.uni-hannover.de
[2] GREYC (UMR 6072), Université de Caen, France. pchapdel@info.unicaen.fr
[3] LIF (UMR 6166), Univ. de la Méditerranée, France. creignou@lif.univ-mrs.fr
[4] LIX (FRE 2653), École Polytechnique, France. hermann@lix.polytechnique.fr

Abstract. Conjunctive-query containment is considered as a fundamental problem in database query evaluation and optimization. Kolaitis and Vardi pointed out that constraint satisfaction and conjunctive query containment are essentially the same problem. We study the Boolean conjunctive queries under a more detailed scope, where we investigate their counting problem by means of the algebraic approach through Galois theory, taking advantage of Post's lattice. We prove a trichotomy theorem for the generalized conjunctive query counting problem, showing this way that, contrary to the corresponding decision problems, constraint satisfaction and conjunctive-query containment differ for other computational goals. We also study the audit problem for conjunctive queries asking whether there exists a frozen variable in a given query. This problem is important in databases supporting statistical queries. We derive a dichotomy theorem for this audit problem that sheds more light on audit applicability within database systems.

1 Introduction

Constraint satisfaction is recognized as a fundamental problem in artificial intelligence, in automated deduction, in computer-aided verification, in operations research, etc. At the same time conjunctive-query containment is considered as a fundamental problem in database query evaluation and optimization [1]. Recent research points out that query containment is a central problem in several database and knowledge base applications, including data warehousing [26], data integration [15], query optimization, and (materialized) view maintenance [28]. Kolaitis and Vardi pointed out in [13] that constraint satisfaction and conjunctive-query containment are essentially the same problem. Constraints are usually specified by means of relations. The standard constraint satisfaction problem can therefore be parameterized by restricting the set of allowed relations. In particular, given a finite set S of Boolean relations, we consider conjunctive propositional formulas consisting of clauses built over relations

* Supported by ÉGIDE 05835SH, DAAD D/0205776 and DFG VO 630/5-1.

from S, also called S-formulas. Deciding the satisfiability of such an S-formula is known as the *generalized satisfiability problem*, denoted by $\text{SAT}(S)$, and was first investigated by Schaefer [20]. It turns out that the complexity of $\text{SAT}(S)$ can be characterized by closure properties of S. This correspondence is obtained through a generalization of Galois theory. In order to get complexity results via this algebraic approach, conjunctive queries $\text{COQ}(S)$ over a set of relations S turn out to be useful. Roughly speaking, a conjunctive query from $\text{COQ}(S)$ is an S-formula with distinguished variables, where all non-distinguished variables are existentially quantified. These queries play an important role in database theory, since they represent a broad class of queries and their expressive power is equivalent to select-join-project queries in relational algebra. Thus they are also of interest in their own right and we study the complexity of some related computational problems. The algebraic approach is particularly well adapted to this study, yielding short and elegant proofs.

We focus here on the counting and the audit problems for conjunctive queries. In the former the problem is to count the number of entries in the database that match the query, i.e., the number of satisfying assignments. In the latter the problem is to audit a database to ensure protection of sensitive data, where the goal is to decide whether the conjunctive query evaluates to false or whether there is some distinguished variable that is frozen, i.e., that takes the same value in all satisfying assignments. This frozen variable would then be considered as not protected. This is a generalization of the audit problem for Boolean attributes defined in [11] (see also [14]), which is particularly interesting in databases supporting statistical queries. For both considered problems we obtain a complete complexity classification that indicates a difference with respect to satisfiability problems of Boolean constraints. Peter Jonsson and Andrei Krokhin ([10] manuscript, submitted for publication) independently examined a variant of our audit problem. Our results can be shown to follow from theirs.

Measures such as conditional probability (confidence) and correlation have been used to infer rules of the form “buying diapers causes you to buy beer”. However, such rules indicate only a statistical relationship between items, but they do not specify the nature and causality of the relationship. In applications, knowing such causal relationship is extremely useful for enhancing understanding and effecting change. While distinguishing causality from correlation is a truly difficult problem, recent work in statistics and Bayesian learning provide some promising directions of attack. In this context, the ideas of Bayesian learning, where techniques are being developed to infer causal relationships from observational data, to mining large databases [21] trigger the necessity to study counting problems in connection with existing database applications. Yet another recent application of Bayesian learning based on counting is the task of spam elimination. Therefore we think that our results will have an impact on concrete database implementations and applications, since the considered formulas in our computational problems correspond better to the model of queries formulated within existing database systems than the so far mainly studied S-formulas.

2 Preliminaries

Throughout the paper we use the standard correspondence between predicates and relations. We use the same symbol for a predicate and its corresponding relation, since the meaning will always be clear from the context, and we say that the predicate *represents* the relation.

An n-ary *logical relation* R is a Boolean relation of arity n. Each element of a logical relation R is an n-ary Boolean vector $m = (m_1, \ldots, m_n) \in \{0,1\}^n$. Let V be a set of variables. A *constraint* is an application of R to an n-tuple of variables from V, i.e., $R(x_1, \ldots, x_n)$. An assignment $I \colon V \to \{0,1\}$ satisfies the constraint $R(x_1, \ldots, x_n)$ if $(I(x_1), \ldots, I(x_n)) \in R$ holds.

Example 1. Equivalence is the binary relation defined by $\mathrm{Eq} = \{(0,0),(1,1)\}$. Given the ternary relations

$$\begin{array}{rcl} R_{\mathrm{nae}} & = & \{0,1\}^3 \smallsetminus \{(0,0,0),(1,1,1)\} \qquad \text{and} \\ R_{1/3} & = & \{(1,0,0),(0,1,0),(0,0,1)\}, \end{array}$$

the constraint $R_{\mathrm{nae}}(x,y,z)$ is satisfied if not all variables are assigned the same value and the constraint $R_{1/3}(x,y,z)$ is satisfied if exactly one of the variables x, y, and z is assigned to 1.

Throughout the text we refer to different types of Boolean constraint relations following Schaefer's terminology [20]. We say that a Boolean relation R is

- *1-valid* if $(1,\ldots,1) \in R$ and it is *0-valid* if $(0,\ldots,0) \in R$,
- *Horn* (*dual Horn*) if R can be represented by a conjunctive normal form (CNF) formula having at most one unnegated (negated) variable in each clause,
- *bijunctive* if it can be represented by a CNF formula having at most two variables in each clause,
- *affine* if it can be represented by a conjunction of linear functions, i.e., a CNF formula with $\oplus$-clauses (XOR-CNF),
- *complementive* if for each $(\alpha_1, \ldots, \alpha_n) \in R$, also $(\neg\alpha_1, \ldots, \neg\alpha_n) \in R$.

A set S of Boolean relations is called 0-valid (1-valid, Horn, dual Horn, affine, bijunctive, complementive) if every relation in S is 0-valid (1-valid, Horn, dual Horn, affine, bijunctive, complementive).

Let S be a non-empty finite set of Boolean relations. An *S-formula* is a finite conjunction of *S-clauses*, $\varphi = c_1 \wedge \cdots \wedge c_k$, where each S-clause c_i is a constraint application of some logical relation $R \in S$. An assignment I satisfies the formula φ if it satisfies all clauses c_i. We denote by $\mathrm{sol}(\varphi)$ the set of satisfying assignments of a formula φ.

Schaefer in his seminal paper [20] developed a complexity classification of the satisfiability problem of S-formulas, denoted by $\textsc{sat}(S)$. *Conjunctive queries* turn out to be useful in order to obtain this result. Given a set S of Boolean relations, we denote by $\textsc{coq}(S)$ the set of all formulas of the form

$$F(x_1, \ldots, x_k) \quad = \quad \exists y_1 \exists y_2 \cdots \exists y_l \; \varphi(x_1, \ldots, x_k, y_1, \ldots, y_l),$$

where φ is an S-formula. These existentially quantified formulas are called *conjunctive queries over* S [13], with $\boldsymbol{x} = \{x_1, \ldots, x_k\}$ being the *distinguished variables*. We denote by SAT-COQ(S) the satisfiability problem of conjunctive queries over S.

3 Closure Properties of Constraints

There exist easy criteria to determine if a given relation is Horn, dual Horn, bijunctive, or affine. We recall these properties here briefly for completeness. An interested reader can find a more detailed description with proofs in the paper [20] or in the monograph [6]. The operations of conjunction, disjunction, majority, and addition applied coordinate-wise on n-ary Boolean vectors $m, m', m'' \in \{0,1\}^n$ are defined as follows:

$$\begin{aligned}
m \wedge m' &= (m[1] \wedge m'[1], \ldots, m[n] \wedge m'[n]) \\
m \vee m' &= (m[1] \vee m'[1], \ldots, m[n] \vee m'[n]) \\
\mathrm{maj}(m, m', m'') &= (m \vee m') \wedge (m' \vee m'') \wedge (m'' \vee m) \\
m \oplus m' &= (m[1] \oplus m'[1], \ldots, m[n] \oplus m'[n])
\end{aligned}$$

where $m[i]$ is the i-th coordinate of the vector m and $\oplus$ is the exclusive-or operator. Given a logical relation R, the following *closure properties* fully determine the structure of R.

- R is Horn if and only if $m, m' \in R$ implies $m \wedge m' \in R$.
- R is dual Horn if and only if $m, m' \in R$ implies $m \vee m' \in R$.
- R is bijunctive if and only if $m, m', m'' \in R$ implies $\mathrm{maj}(m, m', m'') \in R$.
- R is affine if and only if $m, m', m'' \in R$ implies $m \oplus m' \oplus m'' \in R$.

The notion of closure property of a Boolean relation has been defined more generally, see for instance [9, 16]. Let $f\colon \{0,1\}^k \to \{0,1\}$ be a Boolean function of arity k. We say that R is *closed under* f, or that f is a *polymorphism* of R, if for any choice of k vectors $m_1, \ldots, m_k \in R$, not necessarily distinct, we have that

$$\Big(f\big(m_1[1], \ldots, m_k[1]\big),\ f\big(m_1[2], \ldots, m_k[2]\big),\ \ldots,\ f\big(m_1[n], \ldots, m_k[n]\big)\Big) \in R,$$

i.e., that the new vector constructed coordinate-wise from $m_1, \ldots, m_k$ by means of f belongs to R.

We denote by $\mathrm{Pol}(R)$ the set of all polymorphisms of R and by $\mathrm{Pol}(S)$ the set of Boolean functions that are polymorphisms of every relation in S. It turns out that $\mathrm{Pol}(S)$ is a *closed set of Boolean functions* for every set of relations S. All closed classes of Boolean functions were identified by Post [19]. Post also detected the inclusion structure of these classes, which is now referred to as *Post's lattice*, presented in Fig. 2 with the notation from [2]. We did not use the previously accepted notation for the clones, as in [16, 18], since we think that the

$\mathrm{Pol}(R) \supseteq \mathrm{E}_2$	$\Leftrightarrow$	R is Horn	$\mathrm{Pol}(R) \supseteq \mathrm{V}_2$	$\Leftrightarrow$	R is dual Horn
$\mathrm{Pol}(R) \supseteq \mathrm{D}_2$	$\Leftrightarrow$	R is bijunctive	$\mathrm{Pol}(R) \supseteq \mathrm{L}_2$	$\Leftrightarrow$	R is affine
$\mathrm{Pol}(R) \supseteq \mathrm{N}_2$	$\Leftrightarrow$	R is complementive	$\mathrm{Pol}(R) \supseteq \mathrm{N}$	$\Leftrightarrow$	R is compl., 0- and 1-valid
$\mathrm{Pol}(R) \supseteq \mathrm{I}_0$	$\Leftrightarrow$	R is 0-valid	$\mathrm{Pol}(R) \supseteq \mathrm{I}_1$	$\Leftrightarrow$	R is 1-valid
$\mathrm{Pol}(R) \supseteq \mathrm{I}$	$\Leftrightarrow$	R is 0- and 1-valid	$\mathrm{Pol}(R) \supseteq \mathrm{I}_2$	$\Leftrightarrow$	R is Boolean

Fig. 1. Polymorphism correspondences

new one used in [2, 3] is better suited mnemotechnically and also scientifically than the old one. The correspondence of the most studied classes with respect to the polymorphisms of a relation R is presented in Fig. 1. The class I_2 is the closed class of Boolean functions generated by the identity function, thus for every Boolean relation R we have $\mathrm{Pol}(R) \supseteq \mathrm{I}_2$.

An interesting Galois correspondence has been exhibited between the sets of Boolean functions $\mathrm{Pol}(S)$ and the sets of Boolean relations S. A basic introduction to this correspondence can be found in [16, 17] and a comprehensive study in [18]. This theory helps us to get elegant and short proofs for results concerning the complexity of conjunctive queries. Indeed, it shows that the smaller the set of polymorphisms is, the more expressive the corresponding conjunctive queries are, which is the cornerstone for applying the algebraic method to complexity (see [3] for a survey). The following proposition can be found, e.g., in [16, 18].

Proposition 2. *Let S_1 and S_2 be two finite sets of Boolean relations. If the relation $\mathrm{Pol}(S_2) \subseteq \mathrm{Pol}(S_1)$ holds, then $\textsc{coq}(S_1 \cup \{\mathrm{Eq}\}) \subseteq \textsc{coq}(S_2 \cup \{\mathrm{Eq}\})$.*

4 Complexity Results

The only difference between conjunctive queries and S-formulas is that the former contain some existentially quantified variables, thus distinguishing the remaining ones. While this certainly does not lead to a different complexity of the satisfiability problem, this is not any more the case for other computational goals, such as counting the number of satisfying assignments. The algebraic correspondence described above is useful to determine the complexity of the satisfiability problem, since it proves that the complexity of $\textsc{sat-coq}(S)$ strongly depends on the set $\mathrm{Pol}(S)$, as shown in Proposition 2. It provides a polynomial-time reduction from the problem $\textsc{sat-coq}(S_1)$ to $\textsc{sat-coq}(S_2 \cup \{\mathrm{Eq}\})$ by locally replacing each S_1-clause by its equivalent constraint in $\textsc{coq}(S_2 \cup \{\mathrm{Eq}\})$. Moreover, the equivalence relation is actually superfluous. Indeed, from a set of equivalent variables we choose one variable, say z. Then we can delete the corresponding equivalence constraints and substitute the equivalent variables by z in the rest of the formula. Note that we must choose z to be a distinguished variable if an existentially quantified variable occurs in the equivalence set. This proves that $\textsc{sat-coq}(S_1)$ is polynomial-time reducible to $\textsc{sat-coq}(S_2)$. We will show in the sequel that the algebraic approach is helpful to study the complexity of the counting and the audit problems for conjunctive queries.

4.1 Introduction to Counting Problems and Their Reducibilities

A *counting problem* is typically presented using a suitable *witness* function which for every input x, returns a set of *witnesses* for x. Formally, a *witness* function is a function $w\colon \Sigma^* \longrightarrow \mathcal{P}^{<\omega}(\Gamma^*)$, where Σ and Γ are two alphabets, and $\mathcal{P}^{<\omega}(\Gamma^*)$ is the collection of all finite subsets of Γ^*. Every such witness function gives rise to the following *counting problem*: given a string $x \in \Sigma^*$, find the cardinality $|w(x)|$ of the *witness* set $w(x)$.

Let Σ, Γ be two alphabets and let $R \subseteq \Sigma^* \times \Gamma^*$ be a binary relation between strings such that, for each $x \in \Sigma^*$, the set $R(x) = \{y \in \Gamma^* \mid R(x, y)\}$ is finite. We write $\#R$ to denote the following counting problem: given a string $x \in \Sigma^*$, find the cardinality $|R(x)|$ of the witness set $R(x)$ associated with x. It is easy to see that every counting problem is of the form $\#R$ for some R.

Valiant [24, 25] was the first to investigate the computational complexity of counting problems. To this effect, he introduced the class #P of counting functions that count the number of accepting paths of nondeterministic polynomial-time Turing machines. The prototypical problem in #P is #SAT, which is the counting version of Boolean satisfiability. Valiant [24] showed that #SAT is #P-*complete* via *parsimonious* reductions, that is, every counting problem in #P can be reduced to #SAT via a polynomial-time reduction that preserves the cardinalities of the witness sets. Creignou and Hermann [5] proved that the complexity of the counting problem #SAT(S) of S-formulas is dichotomic: #SAT(S) is in FP if S is a set of affine relations, otherwise the problem is #P-complete under Turing reductions.

Hemaspaandra and Vollmer [8] have introduced higher complexity counting classes using a predicate-based framework that focuses on the complexity of membership in the witness sets. Specifically, if $\mathcal{C}$ is a complexity class of decision problems, then $\#\cdot\mathcal{C}$ is the class of all counting problems whose witness function w satisfies the following conditions:

1. There is a polynomial $p(n)$ such that for every x and every $y \in w(x)$, we have that $|y| \leq p(|x|)$, where $|x|$ is the length of x and $|y|$ is the length of y.
2. The witness recognition problem "given x and y, is $y \in w(x)$?" is in $\mathcal{C}$.

In particular, $\#\cdot\mathrm{NP}$ is the class of counting problems associated with decision problems, for which the witness size is polynomially bounded and the witness recognition problem is in NP. Following Toda [22], the inclusions $\#\cdot\Sigma_k\mathrm{P} \subseteq \#\cdot\Pi_k\mathrm{P}$ and $\#\cdot\Pi_k\mathrm{P} \subseteq \#\cdot\Sigma_{k+1}\mathrm{P}$ among counting classes hold for each k. In particular, we have the inclusion $\#\mathrm{P} \subseteq \#\cdot\mathrm{NP}$.

Following Valiant [24], we say that a reduction is *parsimonious* if it is a polynomial-time many-one reduction preserving the number of solutions. However, this reduction does not allow to prove completeness of many known #P-complete problems. Valiant [25] used counting reductions in his #P-completeness proofs, but the aforementioned counting classes are not closed under this reduction, following Toda and Watanabe [23]. Their result implies that every problem hard for #P under Turing reduction is also hard for $\#\cdot\mathrm{NP}$ under the same reduction. However, since the closure of #P under Turing reductions is the whole counting counterpart of the polynomial hierarchy, this does not say anything about the actual complexity of the problem in terms of counting classes. Therefore we have to aim at a result involving a reducibility that preserves (or almost preserves) the relevant classes. More useful for counting problems are *subtractive reductions* [7]. They allow us to obtain many completeness results and at the same time they leave the $\#\cdot\Pi_k\mathrm{P}$ classes closed. Nevertheless, these reductions do not seem to be well-suited for our purposes. Indeed, we need to express the operation of halving the witness set, which is quite delicate if we require the closure of the counting classes under these reductions. For this purpose, we define the *complementive reductions* which satisfy the aforementioned requirements, provided that every witness set of the target counting problem is complementive.

A finite alphabet Γ is called *even* if $|\Gamma| = 2k$ for some $k \in \mathbb{N}$. A permutation π on an even alphabet Γ is called *bipartite* if there exists a partition of Γ into two disjoint sets Γ_0 and Γ_1 such that the following conditions hold:

- $\Gamma = \Gamma_0 \cup \Gamma_1, \quad \Gamma_0 \cap \Gamma_1 = \emptyset, \quad$ and $\quad |\Gamma_0| = |\Gamma_1|$
- for all $x \in \Gamma_i$ we have $\pi(x) \in \Gamma_{1-i}$ for each $i = 0, 1$.

We homomorphically enlarge every permutation π on Γ to the strings in Γ^* by means of the identity $\pi(x_1 \cdots x_k) = \pi(x_1) \cdots \pi(x_k)$ for each string $x_1 \cdots x_k \in \Gamma^*$.

A set of strings $E \subseteq \Gamma^*$ over an even alphabet Γ is called *complementive* if there exists a bipartite permutation π_E on Γ such that $x \in E$ holds if and only if $\pi_E(x) \in E$. If we know that a set of strings E is complementive, we always assume that we are effectively given the permutation π_E. Given Σ, Γ two alphabets with Γ being even, a binary relation B between strings from Σ and Γ is said to be *complementive* if the sets $B(y)$ for each string $y \in \Sigma^*$ are complementive with respect to the same bipartite permutation π_B.

Definition 3. *Let Σ, Γ be two alphabets, Γ being even, and let $\#A$ and $\#B$ be two counting problems determined by the binary relations A and B between the strings from Σ and Γ, where B is complementive.*

- *We say that the counting problem $\#A$ reduces to the counting problem $\#B$ via a* ***strong complementive reduction****, if there exists two polynomial-time computable functions f and g such that for every string $x \in \Sigma^*$:*
  - $B(g(x)) \subseteq B(f(x))$
  - $2 \cdot |A(x)| = |B(f(x))| - |B(g(x))|$
- *A* ***complementive reduction*** *$\#A \leq_{cr} \#B$ from a counting problem $\#A$ to $\#B$ is a transitive closure of strong complementive and parsimonious reductions.*

It is clear that complementive reductions present a special case of counting reductions, the most frequently used reductions among counting problems.

Theorem 4. $\#\mathrm{P}$ *and all higher complexity classes $\#\cdot\Pi_k\mathrm{P}$, $k \geq 1$, are closed under complementive reductions.*

Proof. Let k be a fixed nonnegative integer. We prove that the class $\#\cdot\Pi_k\mathrm{P}$ is closed under strong complementive reductions. The result will follow by induction on the number of strong complementive and parsimonious reductions used to compose the final complementive reduction. Recall that Toda [22] showed that $\#\cdot\Pi_k\mathrm{P} = \#\cdot\mathrm{P}^{\Sigma_k\mathrm{P}}$.

Let $\#A$ and $\#B$ be two counting problems such that $\#B \in \#\cdot\Pi_k\mathrm{P}$, B is complementive, and $\#A$ reduces to $\#B$ via a strong complementive reduction. We will show that $\#A$ belongs to $\#\cdot\Pi_k\mathrm{P}$ by constructing a predicate A' in $\mathrm{P}^{\Sigma_k\mathrm{P}}$ such that for each string x we have $2\cdot|A'(x)| = |B(f(x))| - |B(g(x))| = 2\cdot|A(x)|$, where f and g are the required polynomial-time computable functions.

Let $*$ be a delimiter symbol not in the alphabets Σ and Γ. Let Γ_0 and Γ_1 be the partition sets defined by the bipartite permutation π_B on Γ. The predicate A' consists of all pairs (x, y') of strings x and y', such that y' is of the form

$$f(x) * g(x) * y \text{ with } (f(x), y) \in B,\ (g(x), y) \notin B, \text{ and } \operatorname{last}(y) \in \Gamma_0,$$

where $\operatorname{last}(y)$ denotes the last symbol of the string y. Thus, a pair (x, y') belongs to A' if and only if (x, y') is accepted by the following algorithm:

1. extract $f(x)$, $g(x)$, and y from y';
2. check that $\operatorname{last}(y)$ belongs to Γ_0;
3. check that $(f(x), y)$ belongs to B;
4. check that $(g(x), y)$ does not belong to B.

Steps 1 and 2 take polynomial time. The test in Step 3 is in $\Pi_k\mathrm{P}$, therefore also in $\mathrm{P}^{\Sigma_k\mathrm{P}}$. The test in Step 4 is in $\Sigma_k\mathrm{P}$, hence it can be done in $\mathrm{P}^{\Sigma_k\mathrm{P}}$. Therefore the predicate A' is in $\mathrm{P}^{\Sigma_k\mathrm{P}}$. It is clear from the construction that the identity $2\cdot|A'(x)| = |B(f(x))| - |B(g(x))|$ holds, since B is complementive, and this implies $|A'(x)| = |A(x)|$. It follows that the counting problem $\#A$ is in the counting class $\#\cdot\mathrm{P}^{\Sigma_k\mathrm{P}} = \#\cdot\Pi_k\mathrm{P}$. If we take $k = 0$ in the proof, we get also the closure for the class $\#\mathrm{P}$. $\square$

In view of the preceding Theorem 4, it is quite natural to ask whether the classes $\#\cdot\Sigma_k\mathrm{P}$ are also closed under complementive reductions. The following proposition provides the evidence that *no* class $\#\cdot\Sigma_k\mathrm{P}$ is closed under complementive reductions.

Proposition 5. *For every $k \in \mathbb{N}$, the counting class $\#\cdot\Sigma_k\mathrm{P}$ is not closed under complementive reductions, unless $\#\cdot\Sigma_k\mathrm{P} = \#\cdot\Pi_k\mathrm{P}$.*

Proof. Following Wrathal [27], we must perform a case analysis, whether k is even or odd, To obtain completeness for levels of the polynomial hierarchy we have to use CNF or DNF, according to whether we are in an odd or even level. In the even case, take a $\Pi_{2i}\mathrm{P}$-formula $\varphi(x_1, \dots, x_n)$ and construct the formulas

$$\begin{aligned}
\tau(x_0, x_1, \dots, x_n) &= x_0 \lor x_1 \lor \cdots \lor x_n \lor \neg x_0 \lor \neg x_1 \lor \cdots \lor \neg x_n \\
\psi(x_0, x_1, \dots, x_n) &= \big(x_0 \land \neg\varphi(x_1, \dots, x_n)\big) \lor \big(\neg x_0 \land \neg\varphi(\neg x_1, \dots, \neg x_n)\big)
\end{aligned}$$

where $\neg\varphi$ is formed from φ by de Morgan's laws. For the odd case, take a Π_{2i+1}P-formula φ, maintain the same formula τ, and construct the formula

$$\psi(x_0, x_1, \ldots, x_n) \quad = \quad \big(x_0 \vee \neg\varphi(x_1, \ldots, x_n)\big) \wedge \big(\neg x_0 \vee \neg\varphi(\neg x_1, \ldots, \neg x_n)\big)$$

Both τ and ψ are complementive formulas, hence $\mathrm{sol}(\tau)$ and $\mathrm{sol}(\psi)$ are complementive sets of strings with $\mathit{\Gamma}_0 = \{0\}$ and $\mathit{\Gamma}_1 = \{1\}$.

The non-quantified part of the Π_{2i}P formula φ is in CNF, therefore the formulas $\neg\varphi(x_1, \ldots, x_n)$ and $\neg\varphi(\neg x_1, \ldots, \neg x_n)$ are in DNF. Using the distributive law, both formulas $x_0 \wedge \neg\varphi(x_1, \ldots, x_n)$ and $\neg x_0 \wedge \neg\varphi(\neg x_1, \ldots, \neg x_n)$ can be transformed into DNF in polynomial time and linear space. Hence, the formula ψ is equivalent to a DNF-formula, which can be obtained in polynomial time and linear space. Similarly for the odd case, the non-quantified part of the Π_{2i+1}P formula φ is in DNF, therefore the formulas $\neg\varphi(x_1, \ldots, x_n)$ and $\neg\varphi(\neg x_1, \ldots, \neg x_n)$ are in CNF. Using the distributive law, we can show that the final formula ψ can be transformed in polynomial time and linear space into an equivalent CNF formula.

In both cases, it is clear that $\mathrm{sol}(\psi) \subseteq \mathrm{sol}(\tau)$, $|\mathrm{sol}(\tau)| = 2 \cdot 2^n$, and $|\mathrm{sol}(\psi)| = 2 \cdot |\mathrm{sol}(\neg\varphi)| = 2 \cdot \big(2^n - |\mathrm{sol}(\varphi)|\big)$. Thus we conclude that $2 \cdot |\mathrm{sol}(\varphi)| = |\mathrm{sol}(\tau)| - |\mathrm{sol}(\psi)|$. Hence, we have a complementive reduction from a $\#\cdot\Pi_k$P-complete problem to a counting problem in $\#\cdot\Sigma_k$P. □

4.2 The Counting Problem of Conjunctive Queries

The counting problem associated with the satisfiability of generalized conjunctive queries is defined as follows.

Problem: #SAT-COQ(S)
Input: A conjunctive query $F(\boldsymbol{x}) = \exists \boldsymbol{y}\, \varphi(\boldsymbol{x}, \boldsymbol{y})$ from COQ(S).
Output: Number of different satisfying assignments to the distinguished variables $\boldsymbol{x}$.

We used the notation #SAT-COQ to point out the importance of conjunctive queries, contrary to the cryptic notation $\#\Sigma_1$SAT used on a more theoretical level in [7]. Our ultimate goal is to determine the complexity of #SAT-COQ(S) for all possible sets S. Observe first that #SAT-COQ(S) is in $\#\cdot$NP for every set of Boolean relations S. A central result for our development is the following easy consequence of Proposition 2.

Proposition 6. *Let S_1 and S_2 be two finite sets of Boolean relations. If the inclusion* $\mathrm{Pol}(S_2) \subseteq \mathrm{Pol}(S_1)$ *holds, then there exists a parsimonious reduction from* #SAT-COQ(S_1) *to* #SAT-COQ(S_2).

This result, together with Post's lattice, allows us to prove the following trichotomy complexity classification. We need two propositions whose predecessors can already be found in a slightly different form in [4] and which provide two basic $\#\cdot$NP-complete problems.

Proposition 7. *$\#\text{SAT}(R_{1/3})$ is $\#$P-complete and $\#\text{SAT-COQ}(R_{1/3})$ is $\#\cdot$NP-complete, both via parsimonious reductions.*

Proof. From Valiant's original results [24] follows that $\#$SAT is the generic $\#$P-complete problem via parsimonious reductions. From the same reference and also from [7] it follows that $\#$SAT-COQ is the generic $\#\cdot$NP-complete counting problem under parsimonious reductions (see also [12]). It is clear that $\#\text{SAT}(R_{1/3})$ is in $\#$P and $\#\text{SAT-COQ}(R_{1/3})$ is in $\#\cdot$NP.

The standard reduction from SAT to 3SAT is also a parsimonious reduction from $\#$SAT to $\#$3SAT, and it gives rise to a parsimonious reduction from $\#$SAT-COQ to$\#$3SAT-COQ. Each clause $c = l_1 \vee l_2 \vee l_3$ of a 3SAT formula defines one of the following four relations.

$$\begin{aligned}
\mathrm{OR}_0(x_1,x_2,x_3) &= \mathrm{sol}(x_1 \vee x_2 \vee x_3) &&= \{0,1\}^3 \smallsetminus \{(0,0,0)\}\\
\mathrm{OR}_1(x_1,x_2,x_3) &= \mathrm{sol}(\neg x_1 \vee x_2 \vee x_3) &&= \{0,1\}^3 \smallsetminus \{(1,0,0)\}\\
\mathrm{OR}_2(x_1,x_2,x_3) &= \mathrm{sol}(\neg x_1 \vee \neg x_2 \vee x_3) &&= \{0,1\}^3 \smallsetminus \{(1,1,0)\}\\
\mathrm{OR}_3(x_1,x_2,x_3) &= \mathrm{sol}(\neg x_1 \vee \neg x_2 \vee \neg x_3) &&= \{0,1\}^3 \smallsetminus \{(1,1,1)\}
\end{aligned}$$

We will show that every relation OR_i can be represented as a conjunction of relations $R_{1/3}$. Note first that the relation $Z(v_1, v_2) = R_{1/3}(v_1, v_1, v_2)$ forces the variables v_1 to be assigned the value 0. Therefore the relation $N(x, y, v_1, v_2) = R_{1/3}(x, y, v_1) \wedge Z(v_1, v_2)$ forces y to be the negation of x. For each $c = \mathrm{OR}_i$ we construct now the corresponding formula $r(\mathrm{OR}_i)$ by means of $R_{1/3}$. We obtain the following constructions.

$$\begin{aligned}
r(\mathrm{OR}_0)(x_1,x_2,x_3) &= R_{1/3}(x_1,z_1,z_2) \wedge R_{1/3}(y_2,z_1,z_3) \wedge R_{1/3}(y_3,z_2,z_4) \wedge\\
&\qquad R_{1/3}(z_2,z_3,z_5) \wedge N(x_2,y_2,v_1,v_2) \wedge N(x_3,y_3,v_1,v_2)\\
r(\mathrm{OR}_1)(x_1,x_2,x_3) &= r(\mathrm{OR}_0)(u_1,x_2,x_3) \wedge N(x_1,u_1,v_1,v_2)\\
r(\mathrm{OR}_2)(x_1,x_2,x_3) &= r(\mathrm{OR}_1)(x_1,u_2,x_3) \wedge N(x_2,u_2,v_1,v_2)\\
r(\mathrm{OR}_3)(x_1,x_2,x_3) &= r(\mathrm{OR}_2)(x_1,x_2,u_3) \wedge N(x_3,u_3,v_1,v_2)
\end{aligned}$$

where $u_1, \ldots, u_3, v_1, v_2, y_2, y_3, z_1, \ldots, z_5$ are new variables. In the case of conjunctive queries, these new variables will be existentially quantified. The resulting formula is the conjunction of these partial formulas $r(c)$ for all clauses c. This proves the required parsimonious reductions from $\#$SAT to $\#\text{SAT}(R_{1/3})$ and from $\#$SAT-COQ to $\#\text{SAT-COQ}(R_{1/3})$ $\square$

Remark 8. There exists an alternative and shorter proof of Proposition 7 making use of algebraic arguments. We mention this proof here, since one of our goals is to promote the algebraic approach. The drawback of the proof is that it does not provide an explicit parsimonious reduction and that it is valid only for $\#$SAT-COQ.

Proof. Since $\mathrm{Pol}(R_{1/3}) = \mathrm{I}_2$ and $\mathrm{I}_2 \subseteq S$ for every clone S, we conclude by Proposition 6 that $\#\text{SAT-COQ}(S)$ reduces to $\#\text{SAT-COQ}(R_{1/3})$ via parsimonious reductions. $\square$

Proposition 9. *$\#\textsc{sat}(R_{\mathrm{nae}})$ is $\#\mathrm{P}$-complete and $\#\textsc{sat-coq}(R_{\mathrm{nae}})$ is $\#{\cdot}\mathrm{NP}$-complete, both via complementive reductions.*

Proof. It is clear that $\#\textsc{sat}(R_{\mathrm{nae}})$ is in $\#\mathrm{P}$ and $\#\textsc{sat-coq}(R_{\mathrm{nae}})$ is in $\#{\cdot}\mathrm{NP}$, respectively. To prove completeness, we will reduce $\#\textsc{sat}(R_{1/3})$ to $\#\textsc{sat}(R_{\mathrm{nae}})$. Observe that the algebraic approach is of no use here. Indeed, since R_{nae} is complementive, whereas $R_{1/3}$ is not, we have $\mathrm{Pol}(R_{1/3}) \subset \mathrm{Pol}(R_{\mathrm{nae}})$, which does not provide the desired reduction. Therefore we have to construct an explicit reduction. For each clause $c = R_{1/3}(x_1, x_2, x_3)$ of a $\{R_{1/3}\}$-formula φ, we construct the formula

$$q(c) \quad = \quad R_{\mathrm{nae}}(x_1,x_2,z) \wedge R_{\mathrm{nae}}(x_2,x_3,z) \wedge R_{\mathrm{nae}}(x_3,x_1,z) \wedge R_{\mathrm{nae}}(x_1,x_2,x_3)$$

where z is a new variable. The resulting formula $q(\varphi)$ is the conjunction of these partial formulas $q(c)$ for all clauses c. Observe that if an assignment I satisfies φ, then the dual assignment $\bar{I}$ does not. Observe also that the set of satisfying assignments for the formula $q(c)$ is complementive, therefore the resulting formula $q(\varphi)$ will have twice as many satisfying assignments as the original formula φ. This proves the required complementive reduction from $\#\textsc{sat}(R_{1/3})$ to $\#\textsc{sat}(R_{\mathrm{nae}})$.

In case of conjunctive queries, z will be an existentially quantified variable. In order to be allowed to apply the same argument as above, we have to make sure that if an assignment I on the distinguished variables $\boldsymbol{x}$ satisfies the conjunctive query $F(\boldsymbol{x}) = \exists \boldsymbol{y} \varphi(\boldsymbol{x}, \boldsymbol{y})$, then the dual assignment $\bar{I}$ does not. Since it is not necessarily the case, we have to introduce two new variables u and v, and to consider first a new conjunctive query $F'(\boldsymbol{x}, u, v) = \exists \boldsymbol{y} \varphi(\boldsymbol{x}, \boldsymbol{y}) \wedge R_{1/3}(u, u, v)$. The number of satisfying assignments for F' is equal to the number of satisfying assignments for F. Moreover, F' has the desired property mentioned above. Therefore the previous construction, namely $q(F')$, provides a complementive reduction from $\#\textsc{sat-coq}(R_{1/3})$ to $\#\textsc{sat-coq}(R_{\mathrm{nae}})$. Using Proposition 7, this proves the result. $\square$

Theorem 10. *Let S be a non-empty finite set of Boolean relations.*

- *If S is affine, then $\#\textsc{sat-coq}(S)$ is in FP.*
- *Else if S is bijunctive, or Horn, or dual Horn, then $\#\textsc{sat-coq}(S)$ is $\#\mathrm{P}$-complete under counting reductions.*
- *Otherwise $\#\textsc{sat-coq}(S)$ is $\#{\cdot}\mathrm{NP}$-complete under complementive reductions.*

Proof. If S is affine, then the Gaussian elimination algorithm used in [5] for $\#\textsc{sat}(S)$ can also be used to construct a corresponding polynomial-time algorithm for $\#\textsc{sat-coq}(S)$.

If S is Horn, dual Horn, or bijunctive, then $\textsc{sat}(S)$ is in P following [20] and therefore $\#\textsc{sat-coq}(S)$ is in $\#\mathrm{P}$. Moreover, we know from [5] that in this case $\#\textsc{sat}(S)$ is $\#\mathrm{P}$-hard. Hence, the trivial (parsimonious) reduction from $\#\textsc{sat}(S)$ to $\#\textsc{sat-coq}(S)$ finally shows that $\#\textsc{sat-coq}(S)$ is $\#\mathrm{P}$-complete.

It remains to treat the case where $\mathrm{Pol}(S) = \mathrm{N}$. In fact, observe that all the other nonconsidered classes N_2, I, I_0, I_1 or I_2 are subsets of N. Therefore according to Proposition 6 and Post's lattice, it suffices to exhibit a set S of Boolean relations, such that $\mathrm{N} \subseteq \mathrm{Pol}(S)$ but $\#\textsc{sat-coq}(S)$ is $\#\cdot$NP-complete.

According to Proposition 9 we know that $\#\textsc{sat-coq}(R_{\mathrm{nae}})$ is $\#\cdot$NP-complete via complementive reductions. Construct now the relations

$$\begin{aligned} R''(u,v,x,y,z) &= (\neg u \wedge \neg v \wedge \neg x \wedge \neg y \wedge \neg z) \vee (u \wedge v \wedge x \wedge y \wedge z) \qquad \text{and} \\ R'(u,v,x,y,z) &= R''(u,v,x,y,z) \vee \\ &\quad (u \wedge \neg v \wedge R_{\mathrm{nae}}(x,y,z)) \vee (\neg u \wedge v \wedge R_{\mathrm{nae}}(x,y,z)). \end{aligned}$$

Consider now the formula $F(\boldsymbol{x}) = \exists \boldsymbol{y} \bigwedge_{i=1}^{m} R_{\mathrm{nae}}(x_1^i, x_2^i, x_3^i)$ being an instance of $\#\textsc{sat-coq}(R_{\mathrm{nae}})$, where x_1^i, x_2^i, x_3^i are variables from the vector $\boldsymbol{x}$. Build the formulas

$$\begin{aligned} F'(\boldsymbol{x},u,v) &= \exists \boldsymbol{y} \bigwedge_{i=1}^{m} R'(u,v,x_1^i,x_2^i,x_3^i) \qquad \text{and} \\ F''(\boldsymbol{x},u,v) &= \exists \boldsymbol{y} \bigwedge_{i=1}^{m} R''(u,v,x_1^i,x_2^i,x_3^i) \end{aligned}$$

from the relations R' and R''. The satisfying assignments of the query F' include those of F''. If q is the number of satisfying assignments of F then those of F' is $2q+2$ and those of F'' is 2. Hence, we have the equality $2\,|\mathrm{sol}(F)| = |\mathrm{sol}(F')| - |\mathrm{sol}(F'')|$, implying a complementive reduction from the counting problem $\#\textsc{sat-coq}(R_{\mathrm{nae}})$ to $\#\textsc{sat-coq}(\{R', R''\})$, proving that $\#\textsc{sat-coq}(\{R', R''\})$ is $\#\cdot$NP-complete. Moreover, both R' and R'' are 0-valid, 1-valid, and complementive, since R_{nae} is complementive. Hence $\mathrm{Pol}(\{R', R''\})$ contains N. $\square$

4.3 The Audit Problem

Another problem of interest, defined by Kleinberg *et al.* [11] and studied from a complexity standpoint by Jonsson and Krokhin [10, 14], is the *audit problem.* This problem is related to databases that support statistical queries. It can be generalized to conjunctive queries in the following way.

Problem: $\textsc{audit-coq}(S)$
Input: A conjunctive query $F(\boldsymbol{x}) = \exists \boldsymbol{y}\, \varphi(\boldsymbol{x}, \boldsymbol{y})$ from $\textsc{coq}(S)$.
Question: Is F unsatisfiable or is there some variable among $\boldsymbol{x}$ that is frozen, i.e., that takes the same value in all satisfying assignments?

Note that our $\textsc{audit-coq}(S)$ problem is different from the 1-$\textsc{audit}$ problem studied in [10], since we do not include the variable candidate to be frozen as part of the input. Nevertheless, our result can be shown to follow from those in [10]. We want to insist here on the clarity and simplicity of our proof.

It is easy to see that this problem belongs to the class coNP. We prove that the algebraic approach applies to study the complexity of this problem. The following result follows again immediately from Proposition 2 (see also Proposition 6).

Proposition 11. *Let S_1 and S_2 be two finite sets of Boolean relations. If the inclusion* $\mathrm{Pol}(S_2) \subseteq \mathrm{Pol}(S_1)$ *holds, then* AUDIT-COQ(S_1) *is polynomial-time many-one reducible to* AUDIT-COQ(S_2).

Once more, this result together with Post's lattice allows us to get a complete complexity classification.

Theorem 12. *Let S be a non-empty finite set of Boolean relations.*

- *If S is both* 0- *and* 1-*valid, or affine, or Horn, or dual Horn or bijunctive, then* AUDIT-COQ(S) *is in* P.
- *Otherwise* AUDIT-COQ(S) *is* coNP-*complete.*

Proof. If S is both 0- and 1-valid, i.e., $\mathrm{I} \subseteq \mathrm{Pol}(S)$, then the problem is trivial.

If S is affine, Horn, dual Horn, or bijunctive, then observe that given an S-formula and a variable x, we can check in polynomial time whether both $F \wedge x$ and $F \wedge \neg x$ are satisfiable. Therefore, in this case AUDIT-COQ(S) is in P.

If S is complementive, but neither 0-valid, nor included in the four previous cases, i.e., $\mathrm{Pol}(S) = \mathrm{N}_2$, then no variable can be frozen. Therefore in this case the problem AUDIT-COQ(S) is equivalent to the coNP-complete problem UNSAT(S), asking whether an S-formula is unsatisfiable.

The remaining cases are those for which $\mathrm{Pol}(S) = \mathrm{I}_0$, I_1 or I_2. According to Proposition 6 and Post's lattice, in order to conclude the proof it suffices to exhibit a Boolean relation R_0 (resp. R_1) such that $\mathrm{I}_0 \subseteq \mathrm{Pol}(R_0)$ (resp. $\mathrm{I}_1 \subseteq \mathrm{Pol}(R_1)$) and AUDIT-COQ$(R_0)$ (resp. AUDIT-COQ(R_1)) is coNP-complete. Recall first that SAT$(R_{1/3})$ is NP-complete, so UNSAT$(R_{1/3})$ is coNP-complete. Consider an instance of UNSAT$(R_{1/3})$ defined by the formula $F(\boldsymbol{x}) = \bigwedge_{i=1}^{m} R_{1/3}(x_1^i, x_2^i, x_3^i)$. Construct the 0-valid relation

$$R_0(v, x, y, z) \quad = \quad (\neg v \wedge x \wedge y \wedge z) \vee (\neg v \wedge \neg x \wedge \neg y \wedge \neg z) \vee (v \wedge R_{1/3}(x, y, z))$$

and build the formula $F'(\boldsymbol{x}, v) = \bigwedge_{i=1}^{m} R_0(v, x_1^i, x_2^i, x_3^i)$. Clearly, the inclusion $\mathrm{I}_0 \subseteq \mathrm{Pol}(\{R_0\})$ holds since the relation R_0 is 0-valid.

Observe that F' is always satisfiable, that no variable among the $\boldsymbol{x}$ is frozen, and that F is unsatisfiable if and only if the variable v is frozen to 0 in F'. So, we have a reduction from UNSAT$(R_{1/3})$ to AUDIT-COQ(R_0), therefore the problem AUDIT-COQ(R_0) is coNP-complete. The proof is similar for $\mathrm{Pol}(S) = \mathrm{I}_1$, with a 1-valid relation R_1 similar to R_0, just flip the polarity of the variable v. ☐

5 Conclusion

While the complexity of conjunctive-query evaluation and constraint satisfaction is the same, we determined that this is not any more the case for other computational goals. We have shown that the counting problem for conjunctive queries has a different structure than that for conjunctive formulas. The latter displays a dichotomy behavior between the affine formulas in FP and the #P-complete other cases, as it was shown in [5], whereas the former presents a trichotomy structure between the affine cases in FP, the Horn, dual Horn, and bijunctive #P-complete cases, and finally the general #·NP-complete case. This shows that, under the more fine grained analysis presented by counting, the conjunctive queries present three different levels of (in)tractability. As a byproduct, we developed a new kind of reductions among counting problems, called the complementive reductions, that allow to use halving functions within the counting classes under certain circumstances , i.e., when every instance of the target set is complementive. Since there are many counting problems presenting this structure, we think that the complementive reductions will have a broader impact.

We have also shown that the corresponding audit problem for conjunctive queries displays a dichotomic behavior, where the cases of Horn, dual Horn, bijunctive, or both 0 and 1-valid constraints are in P, whereas the other cases are coNP-complete.

Acknowledgment. We thank Elmar Böhler, Matthias Galota, and Steffen Reith for helpful discussions.

References

1. S. Abiteboul, R. Hull, and V. Vianu. *Foundation of databases*. Addison-Wesley, 1995.
2. E. Böhler, N. Creignou, S. Reith, and H. Vollmer. Playing with Boolean blocks, part I: Post's lattice with applications to complexity theory. *SIGACT News, Complexity Theory Column 42*, 34(4):38–52, 2003.
3. E. Böhler, N. Creignou, S. Reith, and H. Vollmer. Playing with Boolean blocks, part II: Constraint satisfaction problems. *SIGACT News, Complexity Theory Column 43*, 35(1):22–35, 2004.
4. N. Creignou and M. Hermann. On #P-completeness of some counting problems. Research report 2144, Institut de Recherche en Informatique et en Automatique, December 1993. URL = `http://www.lix.polytechnique.fr/~hermann/publications/satcount.ps.gz`.
5. N. Creignou and M. Hermann. Complexity of generalized satisfiability counting problems. *Information and Computation*, 125(1):1–12, 1996.
6. N. Creignou, S. Khanna, and M. Sudan. *Complexity Classifications of Boolean Constraint Satisfaction Problems*, volume 7 of *SIAM Monographs on Discrete Mathematics and Applications*. SIAM, Philadelphia (PA), 2001.
7. A. Durand, M. Hermann, and P. G. Kolaitis. Subtractive reductions and complete problems for counting complexity classes. In M. Nielsen and B. Rovan, editors, *Proceedings 25th International Symposium on Mathematical Foundations of Computer Science (MFCS 2000), Bratislava (Slovakia)*, volume 1893 of *Lecture Notes in Computer Science*, pages 323–332. Springer-Verlag, August 2000. To appear in *Theoretical Computer Science*.
8. L. A. Hemaspaandra and H. Vollmer. The satanic notations: Counting classes beyond #P and other definitional adventures. *SIGACT News, Complexity Theory Column 8*, 26(1):2–13, March 1995.
9. P. Jeavons, D. Cohen, and M. Gyssens. Closure properties of constraints. *Journal of the Association for Computing Machinery*, 44(4):527–548, 1997.

10. P. Jonsson and A. Krokhin. Computational complexity of auditing finite attributes in statistical databases. In *Proceedings Structural Theory of Automata, Semigroups and Universal Algebra, Montreal (Canada)*, July 2003.
11. J. Kleinberg, C. Papadimitriou, and P. Raghavan. Auditing Boolean attributes. *Journal of Computer and System Science*, 66(1):244–253, 2003.
12. J. Köbler, U. Schöning, and J. Torán. On counting and approximation. *Acta Informatica*, 26(4):363–379, 1989.
13. P. G. Kolaitis and M. Y. Vardi. Conjunctive-query containment and constraint satisfaction. *Journal of Computer and System Science*, 61(2):302–332, 2000.
14. A. Krokhin and P. Jonsson. Recognizing frozen variables in constraint satisfaction problems. Technical Report TR03-062, Electronic Colloquium on Computational Complexity, 2003.
15. M. Lenzerini. Data integration: a theoretical perspective. In *Proceeding 21st Symposium on Principles of Database Systems (PODS 2002), Madison (Wisconsin, USA)*, pages 233–246. SIGACT-SIGMOD-SIGART, ACM Press, June 2002.
16. N. Pippenger. *Theories of Computability*. Cambridge University Press, Cambridge, 1997.
17. R. Pöschel. Galois connection for operations and relations. Technical Report MATH-AL-8-2001, Technische Universität Dresden, 2001.
18. R. Pöschel and L. A. Kalužnin. *Funktionen- und Relationenalgebren*. Deutscher Verlag der Wissenschaften, Berlin, 1979.
19. E. L. Post. The two-valued iterative systems of mathematical logic. *Annals of Mathematical Studies*, 5:1–122, 1941.
20. T. J. Schaefer. The complexity of satisfiability problems. In *Proceedings 10th Symposium on Theory of Computing (STOC'78), San Diego (California, USA)*, pages 216–226, 1978.
21. C. Silberstein, S. Brin, R. Motwani, and J. D. Ullman. Scalable techniques for mining causal structures. *Data Mining and Knowledge Discovery*, 4(2-3):163–192, 2000.
22. S. Toda. *Computational complexity of counting complexity classes*. PhD thesis, Tokyo Institute of Technology, Department of Computer Science, Tokyo, Japan, 1991.
23. S. Toda and O. Watanabe. Polynomial-time 1-Turing reductions from #PH to #P. *Theoretical Computer Science*, 100(1):205–221, 1992.
24. L. G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8(2):189–201, 1979.
25. L. G. Valiant. The complexity of enumeration and reliability problems. *SIAM Journal on Computing*, 8(3):410–421, 1979.
26. J. Widom. Research problems in data warehousing. In *Proceedings 4th International Conference on Information and Knowledge Management (CIKM'95), Baltimore (Maryland, USA)*, pages 25–30. Association for Computing Machinery, 1995.
27. C. Wrathall. Complete sets and the polynomial-time hierarchy. *Theoretical Computer Science*, 3(1):23–33, 1976.
28. Y. Zhuge, H. Garcia-Molina, J. Hammer, and J. Widom. View maintenance in a warehousing environment. In M. J. Carey and D. A. Schneider, editors, *Proceedings SIGMOD International Conference on Management of Data, San Jose (California, USA)*, pages 316–327. ACM Press, May 1995.

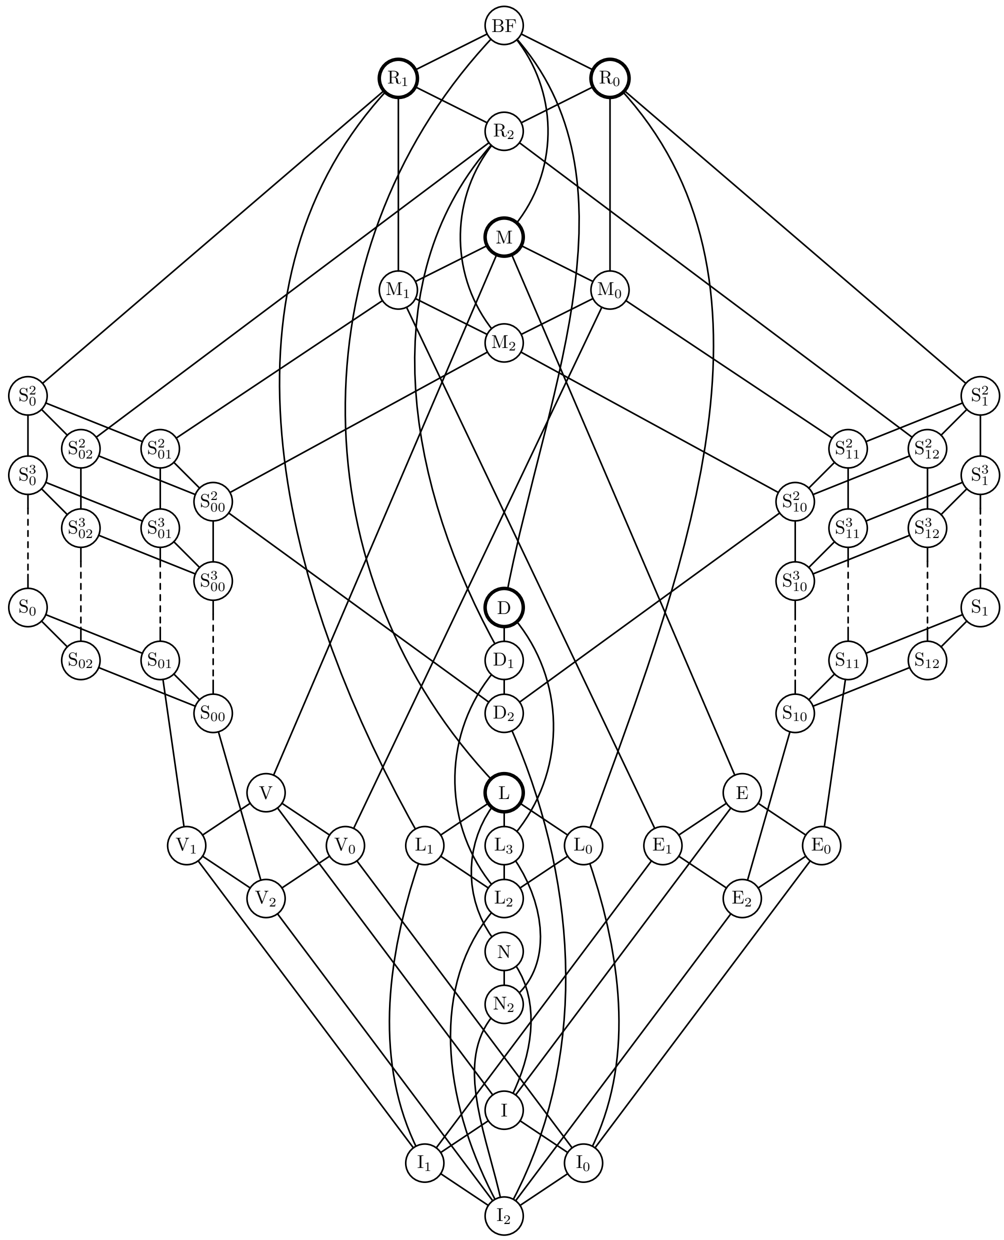

Fig. 2. Graph of all closed classes of Boolean functions