Closures and Dichotomies for Quantified Constraints

Tomás Feder[1] and Phokion G. Kolaitis[2]

[1] 268 Waverley St.
Palo Alto, CA 94301, USA
tomas@theory.stanford.edu,
[2] University of California
Santa Cruz, CA 95064, USA
kolaitis@cse.ucsc.edu

Abstract. Quantified constraint satisfaction is the generalization of constraint satisfaction that allows for both universal and existential quantifiers over constrained variables, instead of just existential quantifiers. We study quantified constraint satisfaction problems $\mathrm{CSP}(Q, S)$, where Q denotes a pattern of quantifier alternation ending in exists or the set of all possible alternations of quantifiers, and S is a set of relations constraining the combinations of values that the variables may take. These problems belong to the corresponding level of the polynomial hierarchy or in PSPACE, depending on whether Q is a fixed pattern of quantifier alternation or the set of all possible alternations of quantifiers. We also introduce and study the quantified constraint satisfaction problems $\mathrm{CSP}'(Q, S)$ in which the universally quantified variables are restricted to range over given subsets of the domain. We first show that $\mathrm{CSP}(Q, S)$ and $\mathrm{CSP}'(Q, S)$ are polynomial-time equivalent to the problem of evaluating certain syntactically restricted monadic second-order formulas on finite structures. After this, we establish three broad sufficient conditions for polynomial-time solvability of $\mathrm{CSP}'(Q, S)$ that are based on closure functions; these results generalize and extend earlier results by other researchers about polynomial-time solvability of $\mathrm{CSP}(Q, S)$. Our study culminates with a dichotomy theorem for the complexity of list $\mathrm{CSP}'(Q, S)$, that is, $\mathrm{CSP}'(Q, S)$ where the relations of S include every subset of the domain of S. Specifically, list $\mathrm{CSP}'(Q, S)$ is either solvable in polynomial-time or complete for the corresponding level of the polynomial hierarchy, if Q is a fixed pattern of quantifier alternation (or PSPACE-complete if Q is the set of all possible alternations of quantifiers). The proofs are based on a more general unique sink property formulation.

1 Introduction

Constraint satisfaction problems are ubiquitous in several different areas of artificial intelligence and computer science, because constraints are widely used to specify design requirements (see [18]). Informally, an instance of a constraint-satisfaction problem consists of a set of variables, a set of possible values for the variables, and a set of constraints that restrict the combinations of values that certain tuples of variables may take; the question is to determine whether there is an assignment of values to the variables that satisfies the given constraints. As first articulated by Feder and Vardi [24], the Constraint Satisfaction Problem (CSP) can be formalized as the Homomorphism Problem: given two finite relational structures A and S, is there a homomorphism h from A to S? A different, but equivalent, way to formalize CSP is to identify it with the Conjunctive Query Evaluation Problem: given a conjunctive query q and a finite relational structure S (that is, a relational database), does S satisfy q? Here, a *conjunctive query* is a first-order formula of the form $\exists x_1 \ldots \exists x_m \psi$, where ψ is a conjunction of atomic formulas with predicates from the vocabulary B of the structure S. Thus, CSP is also a fundamental problems in database query processing, since conjunctive queries (also known as *select-project-join* queries), are the most frequently asked queries in databases.

In its full generality, CSP is an NP-complete problem; for this reason, there has been an extensive pursuit of polynomial-time cases of CSP that are often called "islands of tractability" of constraint satisfaction (see [31, 17, 32, 29]). An instance of CSP consists of two parts: in the formulation of CSP as the Homomorphism Problem, the two parts are a structure A and a structure S, while in the formulation of CSP as the Conjunctive Query Evaluation Problem the two parts are a conjunctive query q and a structure S. A *parametrized* family of constraint satisfaction problems $\mathrm{CSP}(S)$ can be obtained by fixing the structure S, where the input to each of these problems is a structure A in the first formulation, or a conjunctive query q in the second formulation. Much of the research on the complexity of $\mathrm{CSP}(S)$ is motivated from the *dichotomy conjecture* in [24], which asserts that for each structure S, the decision problem $\mathrm{CSP}(S)$ is either NP-complete or solvable in polynomial time. Schaefer's Dichotomy Theorem [33] for generalized satisfiability problems establishes this conjecture for Boolean structures, that is, when the universe of S is a two-element domain. More recently, Bulatov [4] established the dichotomy conjecture for structures S over a three element domain. This conjecture, however, remains open for structures with domains of cardinality bigger than three.

When formulated as the Conjunctive Query Evaluation Problem, constraint satisfaction can be viewed as an important special case of the Model Checking Problem for first-order logic: given a first-order formula φ and a structure S, does S satisfy φ? In turn, this suggests generalizing constraint satisfaction to *quantified constraint satisfaction* QCSP, as follows: given a first-order sentence of the form $Q_1 x_1 \ldots Q x_m \psi$ and a structure S, does S satisfy ψ? Here, each Q_i is an existential or a universal quantifier, and ψ is a conjunction of atomic formulas with predicates from S. In other words, QCSP is the evaluation problem for positive disjunction-free first-order queries. Quantified constraint satisfaction is a PSPACE-complete problem and contains QSAT (also known as QBF) as a special case. As with CSP, one can obtain a parametrized family $\mathrm{QCSP}(S)$ of quantified constraint satisfaction problems by fixing the structure S. For Boolean structures S, this family contains Q3SAT as a member, one of the prototypical PSPACE-complete problems. The following dichotomy theorem for the complexity of Boolean $\mathrm{QCSP}(S)$ was established in [14, 13]: for every Boolean structure S, either $\mathrm{QCSP}(S)$ is PSPACE-complete or $\mathrm{QCSP}(S)$ is solvable in polynomial-time. More recently, researchers have embarked on a systematic investigation of the computational complexity of $\mathrm{QCSP}(S)$, where S is a structure over a finite domain [1, 2, 8]. Specifically, Börner et al. [1, 2], using techniques from universal algebra, found sufficient conditions for tractability of QCSP(S) that are based on closure functions. Furthermore, they obtained a trichotomy theorem for the complexity of QCSP(S), when the relations in S include all graphs of permutations: each such QCSP(S) problem is polynomial-time solvable, NP-complete or PSPACE-complete. Additional broad sufficient conditions for tractability of QCSP(S) based on closure functions were obtained by Chen and others [8–11].

In this paper, we extend the study of quantified constraint satisfaction in two different, but related, directions. First, we consider the class of problems CSP(Q, S), where Q is a *pattern of quantifier alternation*. Thus, standard constraint satisfaction CSP(S) coincides with CSP($\exists, S$), while QCSP(S) coincides with CSP(Q, S) where Q ranges over all possible patterns of quantifier alternation. Note that if the pattern Q has k alternation of quantifiers for some fixed integer k, then CSP(Q, S) belongs to the k-th level of the polynomial hierarchy PH. We also introduce the class of problems CSP$'$(Q, S) where universally quantified variables are required to vary over given subsets of the domain of S. One could also consider the extension CSP$''$(Q, S) of quantified constraint satisfaction where both universally quantified variables and existentially quantified variables are required to vary over given subsets of the domain of S. It turns out, however, that this extension is subsumed by the previous one, since it is easy to see that CSP$''$(Q, S) coincides with CSP$'$($Q, S \cup R$) where R is the set of all subsets of the domain of S (that is, all possible unary constraints on the domain are present). Clearly, CSP$'$(Q, S) contains CSP(Q, S) as a subproblem. Consequently, tractability of CSP$'$(Q, S) implies tractability of CSP(Q, S), while C-hardness of CSP(Q, S) for a complexity class C implies C-hardness of CSP$'$(Q, S).

Our investigation of CSP(Q, S) and CSP$'$(Q, S) begins by showing that these problems are polynomial-time equivalent to the problem of evaluating certain syntactically restricted monadic second-order formulas on finite structures. After this, we establish three broad sufficient conditions for tractability of CSP$'$(Q, S) that are based on closure functions. Specifically, we show that CSP$'$(Q, S) (and, a fortiori, CSP(Q, S)) is solvable in polynomial time if one of the following conditions holds: (1) all relations of S are closed under a *commutative-conservative* operation; (2) all relations of S are closed under a *near-unanimity* operation; (3) all relations of S are closed under a *Maltsev* operation. We also point out that CSP(Q, S) may be tractable, while at the same time CSP$'$(Q, S) may be intractable. Our study of the complexity of CSP$'$(Q, S) culminates with a dichotomy theorem for *list* CSP$'$(Q, S), that is, CSP$'$(Q, S) where the relations of S include every subset of the domain of S. Specifically, in this case CSP$'$(Q, S) is solvable in polynomial time or is complete for the k-th level of the polynomial hierarchy PH, if Q has k alternations of quantifiers (or is PSPACE-complete, if Q ranges over all patterns of quantifier alternation). The proofs are based on a more general unique sink property formulation. Finally, we obtain a dichotomy theorem for CSP$'$(Q, S) where the relations in S include all graphs of permutations.

2 CSP(Q, S), CSP$'$(Q, S), and MMSNP(Q, F)

Let Q be an alternation of $\forall$ and $\exists$ quantifiers of length k ending in $\exists$, or let Q be the set of all possible alternations of quantifiers. We shall allow both of these cases for Q unless otherwise indicated. We let C_Q be the complexity class Σ_k^P if k is odd, the complexity class Π_k^P if k is even (that is, C_Q is one of the levels of the polynomial hierarchy PH), and the complexity class PSPACE if Q consists of all possible alternations of quantifiers. In particular $C_\exists = \Sigma_1^P$ =NP and $C_{\forall\exists} = \Pi_2^P$.

Let S be a fixed finite relational structure, that is, $S = (V, R_1, \ldots, R_m)$ where V is a finite set and each R_i is a relation on V. An instance of CSP(Q, S) is a first-order sentence $Q_1x_1 \ldots Q_lx_l(p_1 \wedge \ldots \wedge p_q)$, where each p_i is an atomic formula involving one of the relations R_j from S, $x_1, \ldots, x_l$ are the variables appearing in the p_i, and $Q_1, \ldots, Q_l$ is a sequence of quantifiers with alternation pattern Q, that is, after identifying consecutive identical Q_i and possibly adding some additional quantifiers to the sequence of Q_i, we can obtain Q. The question is whether the structure S satisfies the sentence $Q_1x_1 \ldots Q_lx_l(p_1 \wedge \ldots \wedge p_q)$. We note that every problem CSP(Q, S) is in the complexity class C_Q.

An instance of CSP$'(Q, S)$ is defined similarly, except that in the first-order sentence, if Q_i is a $\forall$ quantifier, we may use $Q_ix_i \in S_i$ instead of Q_ix_i, where S_i is an arbitrary subset of the elements of S. More formally, an instance of CSP$'(Q, S)$ has two parts: the first part is a first-order sentence $Q_1x_1 \ldots Q_lx_l(p_1 \wedge \ldots \wedge p_q)$, where each Q_i is the existential quantifier $\exists$ or a *bounded* universal quantifier $(\forall x_i \in S_i)$ for some unary relation symbol S_i; the second part is a collection of subsets $S'_{i_1}, \ldots, S'_{i_r}$ of the universe V of S that interpret the unary relation symbols $S_{i_1}, \ldots, S_{i_r}$ occurring in the bounded universal quantifiers of the sentence. The question is whether the expanded structure $S' = (V, R_1, \ldots, R_m, S'_{i_1}, \ldots, S'_{i_r})$ satisfies the sentence $Q_1x_1 \ldots Q_lx_l(p_1 \wedge \ldots \wedge p_q)$. Again, every problem CSP$'(Q, S)$ is in the complexity class C_Q.

Every instance of CSP(Q, S) is also an instance of CSP$'(Q, S)$, but not in the other direction. Indeed, an instance of CSP(Q, S) coincides with the instance of CSP$'(Q, S)$ in which every unary relation symbol S_i in the bounded quantification $(\forall x_i \in S_i)$ is interpreted by the universe V of S.

Feder and Vardi [24] studied the class of problems MMSNP, or monotone monadic strict NP without inequality. We study a more general class. Let Q be an alternation of quantifiers of length k ending in $\exists$. The problem MMSNP(Q, F) is defined by a second-order formula F with a second-order prefix that quantifies over monadic relations and has quantifier alternation Q as defined above, and a first-order universal part without the equality or inequality relations such that every input relation, not quantified in the second-order part, appears with negative polarity, that is, under the scope of an odd number of negations. An instance is a finite structure T having vocabulary given by the input relations in F, and the question is whether T satisfies F. Again, every problem MMSNP(Q, F) is in the complexity class C_Q. Feder and Vardi showed the following.

Theorem 1. *[24] Every* CSP$(\exists, S)$ *problem is an* MMSNP$(\exists, F)$ *problem for some* F *dependent on* S. *For every* MMSNP$(\exists, F)$ *problem, there exists a* CSP$(\exists, S)$ *problem such that the* MMSNP$(\exists, F)$ *problem reduces in polynomial time to the* CSP$(\exists, S)$ *problem, and the* CSP$(\exists, S)$ *problem reduces in randomized polynomial time to the* MMSNP$(\exists, F)$ *problem.*

We generalize the first statement of this theorem.

Theorem 2. *For every* $\mathrm{CSP}(Q,S)$ *or* $\mathrm{CSP}'(Q,S)$ *problem, where* Q *is an alternation of quantifiers of length* k *ending in* $\exists$*, there exists a polynomial time equivalent* $\mathrm{MMSNP}(Q,F)$ *problem.*

Proof. The input structure for the $\mathrm{MMSNP}(Q,F)$ problem is a given structure with relations of the same name and arity as S, plus additional monadic relations M_i to be added later.

For every quantifier type Q_i out of the k quantifier types in the alternation Q of length k and every element v_j in S, introduce a monadic relation R_{ij} which is quantified by Q_i. Require in the first-order part of F that for each existential Q_i and each variable x ranging over the input structure, exactly one $R_{ij}(x)$ holds, for exactly one choice of j, giving the corresponding choice v_j in S for x ranging over the input structure. Also condition the first-order part for each universal Q_i and each x to choices of R_{ij} such that exactly one $R_{ij}(x)$ holds, for exactly one choice of j, giving the corresponding choice of an element v_j in S for x ranging over the input structure.

For each quantifier type Q_i out of the k quantifier types in the alternation Q, add an input monadic relation M_i, and require that at most one $M_i(x)$ hold for each element x, indicating that quantification Q_i has been chosen for x ranging over the input structure. Finally, for each relation A in S of arity r, require that if $A(x_1,\ldots,x_r)$ holds and some corresponding $M_i(x_i)$ holds, then $A(v_1,\ldots,v_r)$ holds in S, where the v_j are given by the R_{ij} corresponding to the quantifier Q_i represented by M_i.

Thus a sentence $Q_1x_1\ldots Q_lx_l(p_1\wedge\ldots\wedge p_q)$ that gives an instance of $\mathrm{CSP}(Q,S)$ corresponds to an instance T of $\mathrm{MMSNP}(Q,F)$, where the structure T consists of the relations $p_1,\ldots,p_q$ on its elements, and the relations M_i on its elements to indicate the chosen quantification Q_i. Conversely, an instance T of $\mathrm{MMSNP}(Q,F)$ gives such a formula, since we may assume that exactly only one $M_i(x)$ holds by discarding elements that are not constrained by any M_i, since the formula is false when two $M_i(x)$ hold for the same x.

The same holds for $\mathrm{CSP}'(Q,S)$ by adding for each subset S' of the elements of S a monadic relation $M_{S'}$ such that $\neg M_{S'}(x)$ indicates that x is restricted to S', and adding a disjunct to the first order part that holds when $\neg M_{S'}(x)$ and the element v_i chosen for x with universal quantification Q_i is not in S'. $\square$

Courcelle's Theorem states that any problem in monadic second-order logic has a linear-time algorithm on instances of bounded treewidth [12, 19]. Thus for a fixed alternation of quantifiers Q of length k, all problems $\mathrm{MMSNP}(Q,F)$, and in particular problems $\mathrm{CSP}(Q,S)$ or $\mathrm{CSP}'(Q,S)$, have a linear-time algorithm on instances of bounded treewidth.

We may consider extensions of the class of all problems $\mathrm{MMSNP}(Q,F)$. A problem is $\mathrm{MMSNP}'(Q,F)$ if we also include the inequality relation as an input relation to an $\mathrm{MMSNP}(Q,F)$ problem. A problem is $\mathrm{MMSNP}''(Q,F)$ if it satisfies the requirements for $\mathrm{MMSNP}(Q,F)$ except for the condition that the formula be monotone in the input relations, so that input relations may appear with both positive and negative polarity. A problem is $\mathrm{MMSNP}'''(Q,F)$ if it satisfies the requirements for $\mathrm{MMSNP}(Q,F)$ except for the last existential second-order quantifier in Q, which may quantify over binary relations as well. Feder and Vardi [24] conjectured that every problem $\mathrm{MMSNP}(\exists, F)$ is either polynomial time solvable or NP-complete, and showed that this conjecture fails for the three extensions of $\mathrm{MMSNP}(\exists, F)$.

Theorem 3. *[24] Every problem in NP has a polynomial time equivalent problem in each of the three classes of problems* $\mathrm{MMSNP}'(\exists, F)$*,* $\mathrm{MMSNP}''(\exists, F)$ *and* $\mathrm{MMSNP}'''(\exists, F)$*. In particular, if P≠NP, these three classes contain problems that are neither in P nor NP-complete [30], and whether problems in each of these classes are polynomial or NP-complete is undecidable, even for problems guaranteed to be polynomial or NP-complete ahead of time.*

We obtain a similar result and conjecture for a quantification Q.

Theorem 4. *For a fixed alternation of quantifiers Q of length k ending in $\exists$, every problem in C_Q has a polynomial time equivalent problem in each of the three classes of problems* $\mathrm{MMSNP}'(Q, F)$*,* $\mathrm{MMSNP}''(Q, F)$ *and* $\mathrm{MMSNP}'''(Q, F)$*.*

Proof. As in [24] for $\mathrm{MMSNP}'(\exists, F)$, we use the result of Hillebrand, Kanellakis, Mairson and Vardi [26] showing that monadic Datalog with inequality (but without negation) can verify a polynomial time encoding of a Turing machine computation. As in [24], the Turing machine may be assumed to be oblivious, and we may use existentially quantified monadic relations to guess nondeterministic inputs to the machine and the resulting states for the machine reaching an accepting state. With additional monadic relations quantified according to Q, we may allow for both nondeterministic existential and universal inputs to the machine, thus encoding a C_Q machine as an $\mathrm{MMSNP}'(Q, F)$ formula.

For $\mathrm{MMSNP}''(Q, F)$, as in [24], we may consider the $\mathrm{MMSNP}'(Q, F)$ encoding of a C_Q machine, and since the formula is now not required to be monotone, we may replace the occurrences of equality with a binary input relation *eq*. We require that *eq* be an equivalence relation (reflexive, symmetric, transitive). For all other input input relations and existential monadic relations are required to hold on some elements if and only if they hold on other elements related to these via the equivalence relation *eq* (for example, if $M(x)$ and $eq(x, y)$ then $M(y)$ as well). Thus equality and inequality are simulated via *eq* and its negation, and therefore every C_Q machine is encoded as an $\mathrm{MMSNP}''(Q, F)$ formula.

For $\mathrm{MMSNP}'''(Q, F)$, as in [24], we include additional input and exsistential monadic and binary relations that force certain elements in the input structure to be marked and also force a uniquely defined existential binary equivalence relation *eq* on marked elements, so that the encoding of a C_Q machine as an $\mathrm{MMSNP}'(Q, F)$ formula yields an encoding as an $\mathrm{MMSNP}'''(Q, F)$ formula by representing inequality using *eq*. □

Conjecture 1. For a fixed alternation of quantifiers Q, every problem $\mathrm{MMSNP}(Q, F)$ is either solvable in polynomial time or complete for some class Σ_k^P or Π_k^P.

3 The Graph of Closure Functions and the Unique Sink Property

Let f be a function of r arguments on the domain of S. We say that S is f-closed if for all relations R in S, of arity p, if $R(a_{1j}, \ldots, a_{pj})$ holds in S for all $1 \leq j \leq r$, and $f(a_{i1}, \ldots, a_{ir}) = b_i$ for all $1 \leq i \leq p$, then $R(b_1, \ldots, b_p)$ also holds in S.

We shall consider problems such that if S is f-closed then f is *idempotent*, that is $f(x, x, \ldots, x) = x$ for all x in S. We define the *graph of closure functions of* S to be a digraph G given as follows. The vertices of G are the elements of the domain of S. Given two distinct elements x, y in S, we include in G a directed edge (x, y) if S is f-closed for some f of some arity p such that for some choice of $a_{ij} \in \{x, y\}$ for $1 \leq i, j \leq p$, we have $a_{ii} = y$ for all $1 \leq i \leq p$ and $f(a_{i1}, \ldots, a_{ip}) = x$ for all $1 \leq i \leq p$.

We generalize the problems $\mathrm{CSP}(Q, S)$ and $\mathrm{CSP}'(Q, S)$ by defining the problem $\mathrm{CSP}^P(Q, S)$ for a collection P of subsets P_i of the domain of S. In $\mathrm{CSP}^P(Q, S)$ the universally quantified variables are only allowed to range over chosen subsets P_i from P. Thus $\mathrm{CSP}(Q, S)$ has P consisting only of the domain P_0 of S, and $\mathrm{CSP}'(Q, S)$ has P consisting of all nonempty subsets P_i of P.

We say that a problem $\mathrm{CSP}^P(Q, S)$ has the *unique sink property* if S has G as the graph of its closure functions, and each subgraph G_i of G induced by a choice of P_i from P has a vertex v_i in P_i that is reachable from every vertex in G_i.

Theorem 5. *Let Q be either a fixed alternation or an arbitrary alternation of quantifiers, and suppose $\mathrm{CSP}^P(Q, S)$ has the unique sink property. then $\mathrm{CSP}^P(Q, S)$ is polynomially equivalent to $\mathrm{CSP}(\exists, S)$, that is, these two problems reduce in polynomial time to each other. Thus $\mathrm{CSP}^P(Q, S)$ is in NP (possibly polynomial or NP-complete).*

Proof. Clearly $\mathrm{CSP}(\exists, S)$ reduces to $\mathrm{CSP}^P(Q, S)$, so we reduce $\mathrm{CSP}^P(Q, S)$ to $\mathrm{CSP}(\exists, S)$. For each graph G_i corresponding to each P_i define an oriented tree T_i spanning G_i, rooted at v_i, and with the edges oriented towards the root v_i.

We would like to consider all possible values of universal quantifiers. However, we shall restrict attention to some specific choices of combinations of values for universal quantifiers. For the trees T_i, consider cosen subtrees U_i of T_i containing v_i. Then we consider assignments to universal variables such that if there are m universal variables ranging over P_i and assigned a particular value x of P_i not in U_i, and the edge (x, y) in T_i was obtained from a closure function f of arity p, then $m < p$.

If we let $U_i = T_i$, then we are considering all combinations of assignments to universally quantified values as needed. But if we only let $U_i = \{v_i\}$, then we are only considering a polynomial number of assignments, since if the maximum arity of functions f used is p, the maximum size of P_i is q, and the number of sets P_i in P is r, then only a constant number pqr of universally quantified variables could be assigned values different from the special value v_i, giving at most $n^{pqr} q^{pqr}$ possible assignments to universal variables. Thus this last case reduces to $\mathrm{CSP}(\exists, S)$ as a polynomial number of instances may be combined into a single instance. When combining instances for different assignments to universally quantified variables, we make sure that existentially quantified variables for different universal combinations are the same variable if the universally quantified variables preceeding them have the same assignments of values. This identification of existential variables in the multiple instances combined is needed as existential variables depend only on the values of preceeding universal variables.

We shall show that solving the problem with $U_i = \{v_i\}$ as we indicated also solves the problem of $U_i = T_i$. For this purpose, we start with $U_i = T_i$ and successively remove some leaf from some U_i, until we get all $U_i = \{v_i\}$. Consider say the step where we remove x from U_1 to obtain U_1'. We must then drop all assignments that assign value x to $p' \geq p$ universal variables ranging over P_1, where p is the arity of the closure function f that gave the edge (x, y) in T_1. We show how to take care of the removal of an assignment with $p' \geq p$ largest. Consider the first p variables $w_1, \ldots, w_p$ out of these p' variables assigned x, and consider the p substitutions that replace their values $x, \ldots, x$ with $a_{1j}, \ldots, a_{pj}$, for each $1 \leq j \leq p$. These assignments each have $a_{jj} = y$ and all $a_{ij} \in \{x, y\}$, thus reducing the number p' of occurrences of x to $p' - 1$ at least. From a solution for each of these p' derived assignments we may obtain a solution for $x, \ldots, x$ as in the original assignment by componentwise applying f, since S is f-closed and f is idempotent, by the definition of f.

Repeatedly taking care of the assignments with largest $p' \geq p$ will eventually guarantee $p' < p$ as required. We must also make sure when performing this substitution by p solutions with smaller p' that the values of existentially quantified variables only depend on preceeding existentially quantified variables, so a similar substitution is performed for each other assignment to universal variables that assign x to w_i for some $1 \leq i \leq p$. If after removing this assignment with large $p' \geq p$ the existential variables depend only on the universal variables preceeding them, then componentwise applying f to both universal and existential variables will preserve the property of having existential variables depending only on preceeding universal variables. Once all $p' < p$, we may substitute U_1 with U_1' consisting of removing x from U_1. Thus overall, we successfully solve the problem for all $U_i = T_i$ by solving the problem obtained when we reach all $U_i = \{v_i\}$, completing the proof. $\square$

The simplest problems having the unique sink property are the cases where the graph G of the closure functions contains a tournament, so that every pair of elements in the domain of S are joined by an edge in either or both directions. The following are examples of this.

A function f of two arguments is a *commutative conservative binary operation* on S if $f(a, b) = f(b, a) \in \{a, b\}$ for all $a, b \in S$. We note that the Horn and dual-Horn cases are f-closed for such a function f. Bulatov and Jeavons [7] showed that if S is f-closed for such a function f, then $\mathrm{CSP}(\exists, S)$ can be solved in polynomial time.

A function g of $r \geq 3$ arguments is a *near-unanimity operation* on S if for a choice of elements a_i in S for $1 \leq i \leq r$ such that there exists an element a in S and some $1 \leq j \leq r$ such that $a_i = a$ for $i \neq j$, we have $f(a_1, \ldots, a_r) = a$. In the special case $r = 3$, the near-unanimity operation is called a *majority operation.* A special case of a majority operation is the *dual discriminator operation* defined by $d(x, y, z) = y$ if $y = z$ and $d(x, y, z) = x$ otherwise. We note that the bijunctive case is d-closed for the dual discriminator operation d. It is shown in [1, 2] that if S is d-closed for the dual discriminator operation d, then $\mathrm{CSP}(Q, S)$ is in nondeterministic logarithmic space, and thus in polynomial time. It is shown in [24] that if S is g-closed for a near-unanimity operation g, then $\mathrm{CSP}(\exists, S)$ can be solved in polynomial time.

A function h of three arguments is a *coset operation* on S if there exists a group operation on S such that $h(x, y, z) = xy^{-1}z$. We note that the affine case (the case of linear equations modulo p) is h-closed for a coset operation h. It is shown in [1, 2] that if S is h-closed for the coset operation of an Abelian group, then $\mathrm{CSP}(Q, S)$ can be solved in polynomial time. Without assuming the group to be abelian, it is shown in [24] that if S is h-closed for a coset operation h, then $\mathrm{CSP}(\exists, S)$ can be solved in polynomial time.

A function h of three arguments is a *Maltsev operation* on S if $h(x, x, y) = h(y, x, x) = y$ for all elements x, y in S. It is shown in [3] that if S is h-closed for a Maltsev operation h, then $\mathrm{CSP}(\exists, S)$ can be solved in polynomial time. An example of a Maltsev operation is the *switching operation* given by $s(x, x, y) = s(y, x, x) = y$ and $s(x, y, z) = y$ otherwise.

The following results were obtained by Bulatov, Chen [1, 8] and others for $\mathrm{CSP}(Q, S)$. Here we obtain immediate corollaries of Theorem refmain

Corollary 1. *If Q consists of all or some alternations of quantifiers, and S is f-closed for an operation that is commutative conservative, near unanimity, coset, or Maltsev, then $\mathrm{CSP}^P(Q, S)$ is polynomial time solvable for all choices of P.*

The preceeding corollary also applies to various forms of majority-minority operations, see e.g. Dalmau [16]. The most general result in this direction is a *k-edge operation*, which has recently been shown to be tractable by Idziak, Markovic, McKenzie, Valeriote, and Willard [27]. It is given as an operation g of arity $k+1$ that satisfies $g(x, x, y, \ldots, y) = g(x, y, x, y, \ldots, y) = g(y, y, y, x, y, \ldots, y) = g(y, y, y, y, x, y, \ldots, y) = \ldots = g(y, y, y, \ldots, y, x) = y$ for all x, y. Clearly this operation gives the unique sink property and is thus polynomial with arbitrary alternations of quantifiers as well as shown here.

Let R be the set of all subsets of the domain of S. Bulatov [5] showed that if the monadic relations constituting R are in S, that is, the fixed structure S contains all monadic relations, then $\mathrm{CSP}(\exists, S)$ is either polynomial time solvable or NP-complete.

Theorem 6. *Let R be the set of all subsets of S, and suppose $R \subseteq S$. Then $\mathrm{CSP}'(Q, S)$ is either polynomial time solvable or C_Q-complete; The criterion distinguishing polynomial and C_Q-complete cases is the same as the criterion distinguishing polynomial and NP-complete cases for $\mathrm{CSP}'(\exists, S)$ in [5].*

Proof. The NP-complete cases of $\mathrm{CSP}(\exists, S)$ in [5] simulate the one-in-three SAT relation $T = \{001, 010, 100\}$ on two particular elements 0 and 1 in S, so C_Q-completeness for $\mathrm{CSP}'(Q, S)$ follows the known classification of the Boolean domain case [25]. The polynomial cases follow from Theorem 5 since the graph G contains a tournament. $\square$

The result of [5] for the case where all subsets appear as monadic relations in S has been generalized to the case where subsets of size at most three appear as monadic relations in S. The proof involves a similar G containing a tournament derived from [5], so the preceeding theorem carries over to this case as well. A similar conjecture when only subsets of size at most two are required to appear as constraints remains open. We may also conjecture polynomiality more generally whenever G contains a tournament.

When G does not contain a tournament, one can construct examples of S for which $\mathrm{CSP}(\exists, S)$ is instead NP-complete. For this it suffices to include only relations involving pairs of values x, y, and if the edge (x, y) is directed only in one direction only use Horn clauses on this pair of values, and when no edge joins x, y include the NP-complete one-in-three SAT problem on these two values. However $\mathrm{CSP}(Q, S)$ remains NP-complete if the unique sink property does hold. However if the unique sink property does not hold then $\mathrm{CSP}(Q, S)$ can go up in complexity as high as C_Q-complete.

A *semilattice operation* is a binary operation satisfying $f(a, a) = a$, $f(a, b) = f(b, a)$, $f(a, f(b, c)) = f(f(a, b), c)$. Bulatov and Chen [6] have shown that for the problem of f-closed structures for a semilattice operation f, then $\mathrm{CSP}(Q, S)$ and $\mathrm{CSP}'(Q, S)$ are polynomial if the unique sink property holds in either case, otherwise the problems are coNP-complete for any alternation of quantifiers involving at least one for-all (while still polynomial with just one exists), and for arbitrary alternations they become PSPACE-complete. It is easy to extend these results to $\mathrm{CSP}^P(Q, S)$ in the obvious manner.

4 Trichotomies for Graphs of Permutations

Assume S contains at least three elements. Let Δ be the set of all graphs of permutations on the elements of S, that is, the set of all binary relations $\{(x, y) : y = \pi(x)\}$ for permutations π. Börner, Krokhin, Bulatov, and Jeavons [1, 2] classified the complexity of $\mathrm{CSP}(Q, S)$ in the cases where Q consists of all alternations of quantifiers and $\Delta \subseteq S$, obtaining a trichotomy as polynomial, NP-complete, or PSPACE-complete. The polynomial cases remain polynomial if Q is a fixed alternation of quantifiers ending in $\exists$. The NP-complete cases are such that $\mathrm{CSP}(\exists, S)$ is also NP-complete, and thus remain NP-complete if Q is a fixed alternation of quantifiers ending in $\exists$. The PSPACE-complete cases are such that S is f-closed for a surjective function f if and only if f is a projection. Note if S is f-closed for a function f, then f must be surjective because S contains all graphs of permutations, so that if $f(a_1, \dots, a_k) = b$ then $f(\pi(a_1), \dots, \pi(a_k)) = \pi(b)$ for every permutation π. Thus the PSPACE-complete cases are such that S is f-closed for a function f (surjective or not) if and only if f is a projection. Thus by a result of Post (see also Jeavons [28]), $\mathrm{CSP}(\exists, S)$ simulates all relations, and in particular simulates the binary relation W of a complete irreflexive graph on the elements in S. The proof in [1] that $\mathrm{CSP}(Q, \{W\})$ is PSPACE-complete when Q consists of all alternations of quantifiers gives a reduction from $\mathrm{CSP}(Q, \{\mathrm{NAE}\})$ for the not-all-equal relation NAE, which is PSPACE-complete by Theorem 6, and adapts to a proof that $\mathrm{CSP}(Q, \{W\})$ is C_Q-complete for a fixed alternation of quantifiers Q ending in $\exists$ with a corresponding reduction from $\mathrm{CSP}(Q, \{\mathrm{NAE}\})$ and Theorem 7. Thus the PSPACE-complete cases become C_Q-complete in this case. A similar proof can be found in [2]

Theorem 7. *If W is an irreflexive clique on $r \geq 3$ elements, then CSP$(Q, \{W\})$ is C_Q-complete.*

Proof. The proof that $\mathrm{CSP}(Q, \{W\})$ is C_Q-complete for a clique W of size $r \geq 3$ encodes an instance of $\mathrm{CSP}(Q, \{\mathrm{NAE}\})$ by making $\binom{r}{2}$ copies of the variables v_i and constraints $\mathrm{NAE}(v_i, v_{i'}, v_{i''})$, then replaces each universally quantified v_i in the resulting instance with two variables x_i, y_i with an edge (x_i, y_i), each existentially quantified v_i with a variable y_i, and also replaces each constraint $\mathrm{NAE}(v_i, v_{i'}, v_{i''})$ with a triangle $(t_i, t_{i'}, t_{i''})$ and edges (y_i, t_i), $(y_{i'}, t_{i'})$, $(y_{i''}, t_{i''})$. Finally, a clique $(u_1, \ldots, u_{r-2})$ is added with edges joining the y_i to all u_j and edges joining the t_i to all u_j other than u_{r-2}. The quantification for the v_i is replicated as a universal quantification for the corresponding x_i or as an existential quantification for the corresponding y_i, followed by existential quantification on the y_i corresponding to universally quantified variables and on the t_i at the end of the formula, and with the first existential quantifiers applied to the u_j.

If we consider the $\binom{r}{2}$ choices of pairs of possible values a, b for the first universally quantified x_i and use them in the $\binom{r}{2}$ corresponding copies of the original instance, we may assume that variables $u_1, \ldots, u_{r-2}$ do not take the two values a, b in one of these copies, in which case the variables y_i take values a, b, corresponding to values $0, 1$ on the v_i that must satisfy the NAE constraints as required by the triangles of t_i. The cases where some x_i do not take values among a, b are more easily satisfied. □

As argued above, Theorem 17 implies the following.

Theorem 8. *Suppose S contains the set Δ of all graphs of permutations. Then* $\mathrm{CSP}(Q, S)$ *is either polynomial, NP-complete, or C_Q-complete.*

The polynomial cases in Theorem 18 are those that are closed under the dual discriminator d, or the switching operation s, or an affine operation, so these cases remain polynomial for $\mathrm{CSP}'(Q, S)$ as well by Theorems 9 and 11. The NP-complete cases in Theorem 18 are such that $\mathrm{CSP}(\exists, S)$ simulates the not-all-equal relation N on a 2-element subdomain [15], and so Theorems 6 and 7 imply that $\mathrm{CSP}'(Q, S)$ is C_Q-complete. The C_Q-complete cases in Theorem 18 remain C_Q-complete for $\mathrm{CSP}'(Q, S)$.

Theorem 9. *Suppose S contains the set Δ of all graphs of permutations. Then* $\mathrm{CSP}'(Q, S)$ *is either polynomial or C_Q-complete.*

References

1. F. Börner, A. Krokhin, A. Bulatov, and P. Jeavons, Quantified constraints: algorithms and complexity, in Proceedings of 17th International Workshop Computer Science Logic, CSL '03, LNCS 2803, Vienna, Austria (2003) 58–70.
2. F. Börner, A. Krokhin, A. Bulatov, and P. Jeavons, Quantified constraints and surjective polymorphisms, Technical report PRG-RR-02-11, Oxford University, 2002, 25pp.

3. A. Bulatov, Malt'sev constraints are tractable, Technical report PRG-RR-02-05 Oxford University, 2002, 36pp., see also ECCC TR 03-034.
4. A. Bulatov, A dichotomy theorem for constraints on a three-element set, Proc. 43rd IEEE Symposium on Foundations of Computer Science, 649–658, 2002.
5. A. Bulatov, Tractable conservative constraint satisfaction problems, Proc. Eighteenth Annual IEEE Symposium on Logic in Computer Science, 321-330, 2003.
6. A. Bulatov and H. Chen, personal communication.
7. A. Bulatov and P. Jeavons, Tractable constraints closed under a binary operation, Technical Report PRG-TR-12-00, Computing Laboratroy, University of Oxford, Oxford, UK, 2000.
8. H. Chen, Quantified constraint satisfaction problems: closure properties, complexity, and algorithms, manuscript.
9. H. Chen, The computational complexity of quantified constraint satisfaction, Ph.D. thesis, Cornell University, August 2004.
10. H. Chen, Quantified constraint satisfaction and 2-semilattice polymorphisms, Tenth International Conference on Principles and practice of constraint programming, 2004, Toronto, Canada.
11. H. Chen and V. Dalmau, From pebble games to tractability: an ambidextrous consistency algorithm for quantified constraint satisfaction, Computer Science Logic (CSL), 2005, Oxford, England.
12. B. Courcelle, The monadic second-order logic of graphs I; Recognizable sets of finite graphs, Information and Computation 85 (1990) 12–75.
13. N. Creignou, S. Khanna, and M. Sudan, Complexity Classification of Boolean Constraint Satisfaction Problems, volume 7 of SIAM Monographs on Discrete Mathematics and Applications, 2001.
14. V. Dalmau, Some dichotomy theorems on constant-free quantified Boolean formulas, Technical Report TR LSI-97-43-R, Department LSI, Universitat Politecnica de Catalunya, 1997.
15. V. Dalmau, A new tractable class of constraint satisfaction problems, in Proceedings 6th International Symposium on Artificial Intelligence and Mathematics, 2000.
16. V. Dalmau, Generalized majority-minority operations are tractable, in Proceedings of Logic in Computer Science, LICS 2005.
17. R. Dechter, Constraint networks, Encyclopedia of Artificial Intelligence, Wiley, 1992, 276–185.
18. R. Dechter, Constraint Processing, Morgan Kauffman, 2003
19. R. G. Downey and M. R. Fellows, Parametrized Complexity, Springer-Verlag, 1999.
20. T. Feder, Constraint satisfaction on finite groups with near subgroups, ECCC TR05-005, 2005.
21. T. Feder, P. Hell, and J. Huang, Bi-arc graphs and the complexity of list homomorphisms, Journal of Graph Theory 42 (1999) 61–80.
22. T. Feder, P. Hell, and J. Huang, List homomorphisms to reflexive digraphs, manuscript.
23. T. Feder, F. Madelaine, and I. A. Stewart, Dichotomies for classes of homomorphism problems involving unary functions, Theoretical Computer Science, to appear.
24. T. Feder and M. Y. Vardi, The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory. SIAM Journal of Computing 28 (1998), 57–104.
25. E. Hemaspaandra, Dichotomy theorems for alternation-bounded quantified boolean formulas, CoRR, cs.CC/0406006, 2004.
26. G. G. Hillebrand, P. C. Kanellakis, H. G. Mairson, and M. Y. Vardi, Tools for Datalog boundedness, Proc. 10th ACM SIGACT-SIGMOD-SIGART Symp. on Principles of Database Systems (1991), 1–12.

27. P. Idziak, P. Markovic, R. McKenzie, M. Valeriote, and R. Willard, Tractability and learnability arising from algebras with few subpowers, manuscript
28. P. G. Jeavons, On the algebraic structure of combinatorial problems, Theoretical Computer Science 200 (1998), 185–204.
29. Ph.G. Kolaitis, Constraint satisfaction, databases, and logic, Proc. of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI 2003), 1587–1595, 2003.
30. R. E. Ladner, On the structure of polynomial time reducibility, J. Assoc. Comput. Mach. 22 (1975) 155–171.
31. A.K. Mackworth and E.C. Freuder, The complexity of constraint satisfaction revisited, Artificial Intelligence, 1993, 59, 57–62.
32. J. Pearson and P. Jeavons, A survey of tractable constraint satisfaction problems, Technical Report, Royal Holloway University of London, CSD-TR-97-15, 1997.
33. T. J. Schaefer, The complexity of satisfiability problems, Proc. 10th ACM Symp. on Theory of Computing (1978) 216–226.