Theoretical Computer Science
13, (1981), 271–293

Generalized Galois Theory and its Application to Complexity

ABSTRACT We generalize those aspects of classical Galois theory that have to do with the discussion of solvability of problems (namely polynomial equations) relative to auxiliary procedures (e.g. radicals). The underlying structures need no longer be fields, and the problems and procedures more typically arise as algorithmic (e.g. combinatorial) problems. Some of the classical notions and results, e.g. resolvents and discriminants have their natural counterparts. We extend the classical theory mainly in the direction of relations between the group of a problem and the structure and complexity of its solution algorithm. The present paper gives a connected and detailed exposition of this theory, improving and considerably expanding our earlier reports [3, 4]. It now represents a tool for the systematic discussion of the solvability of algorithmic problems, their dependence on structural settings, and the relative merits of solution strategies.

Classical Galois theory is now primarily a theory of field extensions. Its main theorem connects the lattice of subgroups of the automorphism group of an extension field over the ground field (inversely) to the lattice of intermediate fields. Most of the present generalizations of Galois theory address themselves to a program of obtaining such a Galois correspondence in some, perhaps weakened, form; e.g. Jónsson [8]. The present chapter tries to generalize another aspect, one that lies at the historical roots of Galois theory, which is the theory of equations, mainly the problem of solvability of polynomial equations by repeatedly solving some special type of equations, e.g. solving for radicals.

Our goal is to investigate how far the paradigm of classical Galois theory can be applied to problems which are not necessarily about fields, but, for example, about structures suggested by algorithmic problems in finite combinatorics, graph theory, etc.

The way of describing such questions, which is to this author the most natural, uses the terminology and methods of first-order model theory, of which we use a very limited amount, sketched in the first section.

In the second section we make precise our notion of an algorithmic problem and the concept of solvability, resp. relative solvability, of a problem by means of an algorithm. This aspect has usually been left rather vague by traditional accounts of solvability, e.g. constructibility in geometry.

Next, we list the basic assumptions for the setting-up of a Galois theory: The notion of a well-posed problem, abstracting from some aspects of polynomial equations, and

the basic assumption of the amalgamation property.

The group of a problem is defined in Section 5, and in there and in Section 6 we see that some of the notions and results for the classical theory have their natural counterparts here; some results, e.g. about our notion of discriminants extend classical theory.

The connections between the structure and size of solution programs and the structure and order of the group of the problem are dealt with in the remaining sections and illustrated by diverse examples. In Section 7 we show that the group of a problem can be obtained as the group of any of its solution programs; in Section 8 we probe the basic relation between the factors in a subnormal series of the group of the problem to groups of the auxiliary problems that are solved in the relative solution algorithm. This yields tools for investigating the complexity of relative solution algorithms, expressed in terms of the structure of the group of the problem.

These results are detailed in Section 9. They allow the systematic discussion of the relative merits of solution strategies for algorithmic problems: A solution strategy is obtained by varying the language-setting of the problem; we get thereby varying possibilities for expressing auxiliary problems. The corresponding relative complexity of the problem is then compared to the complexity of the auxiliary problem. We obtain in this fashion lower bounds for the various solution strategies, both for getting all solutions (Theorem 24) and the fastest solution (Theorem 27).

In Section 10 we generalize the known sufficient condition for the computability of the Galois group of a polynomial, namely the existence of a reducibility algorithm.

The contents of this chapter were to a large extent presented in a course given at the Banach Center, Warsaw, during a Semester of Algebra and Application, 1978, and at the SIAM Spring Meeting, Madison, Wisconsin, 1978.

1 Some notions and tools from logic

Let $\mathcal{A} = \langle A, R, f, c\rangle$ be some *relational structure*, to be concrete, let R be a binary relation $R \subseteq A \times A$, f a binary operations $f : A \times A \to A$ and c a fixed special element $c \in A$. The reader may think of $\mathcal{A}$ as a field, a model of geometry, or any such structure he pleases. To adequately treat the descriptive and the algorithmic aspects of such a structure we shall now specify formal frameworks for mathematically expressing these aspects. Any such formalization brings with it some choices and restrictions in expressive power. So does ours, and the alert reader should become aware of the fact that not all our choices were necessary for the theory to be established. For example, already above we could have chosen $\mathcal{A}$ a many-sorted structure with any countable number of relations, functions and constants. The choice here was in favor of conciseness of notation.

The descriptive aspect of a relational structure is handled by the language of *first-order predicate logic*. In the present case, this language would provide variables $x_1, x_2, \ldots$ to run over the set A, and a constant symbol c to denote the distinguished element c in A. Corresponding to the binary operation f on A there is a function symbol f. With these language elements we can form *terms*; the class of terms is the smallest class of formal expressions (i.e. strings of symbols $(,), f, c, x_1, x_2, \ldots$) which contains the variables and the constant and is closed under the formation of the formal expression $f(\tau_1, \tau_2)$ from given terms τ_1 and τ_2. Corresponding to the relation R on A we introduce a binary predicate symbol R with which we are allowed to form the simplest type of *formulas*, namely atomic formulas $R(\tau_1, \tau_2)$ for any terms τ_1 and τ_2. Quantifier-free formulas are formed by taking Boolean combinations of atomic formulas, the corresponding connectives are $\wedge$(and), $\wedge$(or), $\neg$(not), $\supset$(implies), $\equiv$(if and only if, short "iff"). For most of the present chapter we shall be concerned with this language only, we shall call it L_0. Predicate logic, of course, extends L_0 to L by the addition of quantifiers $\exists x_i$ (there is an x_i such that $\cdots$) and $\forall x_i$ (for all $x_i \cdots$). We shall have occasion to extend L further by the addition of countable disjunctions: if δ_i are formulas of L for $i = 1, 2, \ldots$, then $\bigvee_{i=0}^{\infty} \delta_i$ is a formula of L_∞. Of course, one could go further in adding language constructs, finite or infinitary, but we shall have no need for this here.

Predicate logic is equipped with the basic notions of *satisfaction*, semantical *entailment* and formal *provability* which are in a particularly simple relationship. Let Γ be a set of closed formulas of L, the $\Gamma \vdash \phi$ for a closed formula ϕ means that ϕ is formally provable from $\Gamma, \Gamma \models \phi$ for a closed formula ϕ, means that ϕ is true in every model of Γ. If $\Gamma \vdash \phi$ is not the case for some logically false formula, then Γ is said to be consistent; in this case Γ has a model, i.e. a structure $\mathcal{A}$ which every formula ϕ in Γ is satisfied; for this we write $\mathcal{A} \models \phi$. We have the following basic fact:

Theorem 1 - Completeness. *For $\Gamma \subseteq L$ and $\phi \in L$ we have $\Gamma \models \phi$ iff $\Gamma \vdash \phi$.*

The following theorem is our main non-constructive tool.

Theorem 2 - Compactness. *If $\Gamma \subseteq L$ and $\phi \in L$ and $\Gamma \models \phi$, then $\Gamma_0 \models \phi$ for some finite $\Gamma_0 \subseteq \Gamma$.*

Corollary 3 *If $\Gamma \subseteq L, \delta_i(x) \in L$ for $i = 1, 2, 3, \ldots$ and $\Gamma \models \forall x \bigvee_{i=1}^{\infty} \delta_i(x)$ then there exists a finite subset $\{i_1, \ldots, i_n\} \subseteq \mathbb{N}$ such that*

$$\Gamma \models \forall x(\delta_{i_1}(x) \vee \cdots \vee \delta_{i_n}(x))$$

We shall repeatedly have occasion to introduce additional individual constants to L_0 (or L or L_∞); we call the corresponding extended language $L_0(a_1, \ldots, a_n)$, etc. as the case may be. If τ is a term, ϕ a formula of the extended language we shall indicate

these constants by $\tau(a_1, \ldots, a_n)$ resp. $\phi(a_1, \ldots, a_n)$. Let s be any *permutation* of $\{1, 2, \ldots, n\}$, then τ^s, ϕ^s denote the term, resp. the formula, obtained by replacing a_i with $a_{s(i)}$ at all occurrences. This notation is extended to sets of terms, resp. formulas: $\Delta^s = \{\delta^s : \delta \in \Delta\}$.

By a *diagram* Δ we understand a consistent set of closed (or "constant") negated and unnegated atomic formulas such that for all such α either $\alpha \in \Delta$ or $\neg\alpha \in \Delta$. Each structure $\mathcal{A}$ gives rise to such a diagram $\Delta(\mathcal{A})$, it consists of those α resp. $\neg\alpha$ which hold in $\mathcal{A}$. By the completeness theorem each consistent diagram has a model.

A *theory* is a consistent set Γ of closed formulas of L. A theory is called *universal theory* if Γ consists of the universal closures of formulas of L_0, i.e. of formulas of the form $\forall x_1 \forall x_2 \cdots \forall x_n \phi(x_1, \ldots x_n)$, where $\phi(x_1, \ldots x_n) \in L_0$ and $x_1, \ldots x_n$ are all the variables that occur in ϕ.

Lemma 4 *If Γ is a universal theory and $\mathcal{B}$ is a substructure of a model $\mathcal{A}$ of Γ, then $\mathcal{B}$ is also a model of Γ; in particular Γ has a minimal model.*

Proof. These facts are quite obvious from the definition; note that by a substructure of $\mathcal{A} = \langle A, R, f, c\rangle$ we mean a structure $\mathcal{B} = \langle B, S, g, d\rangle$, where $B \subseteq A, d = c \in B, S$ is the restriction $R \cap (B \times B)$ of R to B and B is closed under the operation f whose restriction to B is just g.

Let Γ be any theory, the universal theory of Γ is the set of universal closures of formulas of L_0 which are derivable from Γ; this theory is called *decidable* if there is a decision method (in the sense of recursion theory) which allows to decide, for every such ϕ, whether $\Gamma \vdash \phi$ or $\Gamma \nvdash \phi$. We shall use the following fact:

Lemma 5 *If the universal theory of Γ is decidable and if $\phi(a_1, \ldots, a_n) \in L_0(a_1, \ldots, a_n)$ is a closed formula, then it is decidable whether $\phi(a_1, \ldots, a_n)$ can be consistently added to Γ.*

Proof. By the completeness theorem we have $\Gamma \cup \{\phi(a_1, \ldots, a_n)\}$ is consistent iff $\Gamma \nvdash \forall x_1 \cdots \forall x_n \neg\phi(x_1, \ldots, x_n)$, which is decidable by assumption.

2 Algorithmic problems

The algorithmic aspect of a relational structure $\mathcal{A}$ is handled by an appropriate programming language. The construction according to these programs take place within this structure; they are to be understood as contingent constructions, namely they are based on the hypothesis that the relation $R(a, b)$ can be decided for any given elements

$a, b \in A$ and that the basic operation $f(a, b)$ can be effectively performed whatever a and b. While this is unproblematic for finite structures or countable structures such as the field $\mathbb{Q}$ of rationals, the programs are truly contingent for some typical uncountable structures; in particular, the classical theory of geometrical constructions is contingent on the objects being algebraic over $\mathbb{Q}$. (Of course, we could take - and have taken - the platonic position of "non-constructive" basic constructions in their full generality. This gives rise to a nice theory, [3], see also [10] and [13].) Putting aside basic constructibility considerations, we now make precise a particular algorithmic language.

Let $x_1, x_2, x_3, \ldots$ be an infinite supply of variables, c a constant symbol. The atomic constituents of our *programming language* are assignment statements $x_i := x_j, x_i := f(x_j, x_k), x_i := c$. These commands signify the replacement of the value of x_i by the former value of x_j, the value of $f(x_j, x_k)$ for the former x_j and x_k, the value c respectively. The basic program connectives are $\pi_1; \pi_2$, **if** $R(x_i, x_j)$ **then** π **else** π_2 **fi** and **while** $R(x_i, x_j)$ **do** π **od**, where π, π_1, π_2 are any given programs, (e.g. assignment statements). The purpose of a program, as a formula in a particular programming language like the one just defined, is to provide a control structure for the sequencing of operations and decisions which transform stepwise a given valuation of the variables $x_1, x_2, \ldots$ by elements of A to some other valuation. There are many such types of control structures available in theory and practice, their basic aspect - and this is all that is actually needed in the sequel - is the following.

Lemma 6 *If $\pi(x_1, \ldots, x_n)$ is a program in which the variables $(x_1, \ldots, x_n)$ occur (and no others), then it is possible to effectively determine a countable sequence of n-tuples of terms $\langle \tau_{i1}(x_1, \ldots, x_n), \ldots, \tau_{in}(x_1, \ldots, x_n) \rangle$ and formulas $\phi_i(x_1, \ldots, x_n, \tau_{i1}(x_1, \ldots, x_n), \ldots, \tau_{in}(x_1, \ldots, x_n))$ such that all structures $\mathcal{A}$ and for all values $a_1, \ldots, a_n \in A$ the following is the case:*

> *the program π terminates on input $a_1, \ldots, a_n$ iff there exists $i \in \mathbb{N}$ such that $\phi_i(a_1, \ldots, a_n, \tau_{i1}(a_1, \ldots, a_n), \ldots, \tau_{in}(a_1, \ldots, a_n))$ holds in $\mathcal{A}$; the values of $x_1, \ldots, x_n$ at termination are those $\tau_{i1}(a_1, \ldots, a_n), \ldots, \tau_{in}(a_1, \ldots a_n)$ for which $\phi_i(a_1, \ldots, a_n, \tau_{i1}(a_1, \ldots, a_n), \ldots)$ holds.*

For what follows let Γ be a fixed universal theory.

Corollary 7 *If the program $\pi(x_1, \ldots, x_n)$ terminates in all models of Γ and for all inputs in that model, then there is a program $\pi'(x_1, \ldots, x_n)$ which is loop-free (i.e. composed only with ; and* **if** $\cdots$ **then** $\cdots$ **else** *from assignment statements) such that π and π' are equivalent (i.e. go through the same sequence of instructions on all inputs).*

Proof. The assumption reads $\Gamma \models \forall x_1 \cdots \forall x_n \bigvee_{i=1}^{\infty} \phi_i(x_1, \ldots, x_n), \tau_{i1}(x_1, \ldots x_n), \ldots, \tau_{in}$ $(x_1, \ldots, x_n))$ and the conclusion follows from Corollary 3, if we remark that any finite disjunction of the above form can obviously be programmed in the restricted programming language.

We state this result concisely by saying that universal algorithms have the unwind property. (In the case of real arithmetic it follows from [3] that therefore all programmable functions on the reals consist of finitely many pieces of rational functions.)

Typically, programs are written to solve problems. Our present position would be that a problem is presented by a problem statement, that is a formula $\phi(x_1, \ldots, x_n, y_1, \ldots, y_n)$ of L_0 (in which all variables that actually occur are among $x_1, \ldots, y_n$ but need not exhaust them). The problem is to find for any model $\mathcal{A}$ and values $a_1, \ldots, a_n$ such elements $b_1, \ldots, b_n$ for which $\phi(a_1, \ldots, a_n, b_1, \ldots, b_n)$ holds in $\mathcal{A}$. The problem is called an *algorithmic problem* if the question is to find a program $\pi(x_1, \ldots, x_n)$ which solves it. Formally, we can now express it as follows: a problem $\phi(x_1, \ldots, x_n, y_1, \ldots, y_n)$ is algorithmically solvable for a theory Γ iff there exists a program $\pi(x_1, \ldots, x_n)$ such that $\Gamma \models \forall x_1 \cdots \forall x_n(\phi_i(x_1, \ldots, x_n, \tau_{i1}(x_1, \ldots, x_n), \ldots) \supset \phi(x_1, \ldots, x_n, \tau_{i1}(x_1, \ldots, x_n), \ldots))$ for some $i \in I\!N$.

Our main concern in this chapter is with *relative algorithmic solvability*. In order to simplify notation, let $\phi(x)$ be a problem. (We have taken the input values as elements of $\mathcal{A}$ which can be denoted by a constant term of L_0 and are looking for a singleton as a solution x rather that the n-tuple $\langle y_1, \ldots, y_n \rangle$.) It may well be the case that x cannot be obtained by repeatedly applying the basic operations of $\mathcal{A}$ to the starting values, e.g. the solutions of $x^4 = 2$ cannot be obtained by the field operations starting with $0, 1$. However, (all) solutions can be obtained if we allow repeated extraction of square roots. In other words, if we have a solution procedure for all instances of the *auxiliary problem* $x^2 = a$, then these procedures can be composed to a solution program for $x^4 = 2$. More generally, an auxiliary problem is a formula $\psi(u, v)$ of L_0 and the working assumption is that we have a solution procedure for ψ, which in each model $\mathcal{A}$ of Γ produces all solutions of $\psi(a, v)$ for all $a \in A$; assume that there are only finitely many solutions possible for any a and $\mathcal{A}$. By another application of compactness, the number of possible solutions is then universally bounded (Corollary 8 below), say by m, and we shall abbreviate a procedure call for the solution of ψ by

$$v_1, \ldots, v_m := \psi(u, v)$$

where u stands for the variable that obtains the input value to the procedure. Note that the execution of this procedure call is not assumed to give us any specific order on the solutions, e.g. $v_1, v_2, v_3 := (x^3 = 1)$ will give us the third roots of unity in any order whatever.

The *relative complexity* of a problem $\phi(x)$ is the number of times the auxiliary procedure $v_1, \ldots, v_m := \psi(u, v)$ has to be called until *all* solutions of the problem have been obtained. This definition only makes sense if the solution is finite. To determine

lower bounds on the relative complexity will be the main goal of the theory. - Taking algebraic examples above is not a coincidence, our main paradigm for the theory is the solution of polynomial equations and the discussion of this problem by means of (classical) Galois theory.

3 Well-posed and irreducible problems

We have already observed that our main goal implies that we should look for problems (and auxiliary problems) which have always a finite number of solutions.

Corollary 8 *If $\psi(u,v)$ has in all models of Γ and for all values of u a finite number of solutions, then this number is universally bounded by a natural number m.*

Proof. The assumption reads

$$\Gamma \models \forall u \bigvee_{k=1}^{\infty} \exists v_1 \cdots \exists v_k (\bigwedge_{i=1}^{k} \psi(u,v_k) \wedge \bigwedge_{1 \leq i < j \leq k} v_i \neq v_j \wedge \forall v (\psi(u,v) \supset (\bigvee_{i=1}^{k} v = v_i)))$$

and the conclusion follows again from Corollary 3.

Let Δ be a diagram consistent with Γ and let $\phi(x)$ be a problem. A *splitting diagram* for $\phi(x)$ over Δ is a diagram $\Delta(a_1, \ldots, a_n) \supseteq \Delta$ such that Γ is consistent with $\Delta(a_1, \ldots, a_n), a_1, \ldots, a_n$ are different new individual constants and the following is the case:

$$\Gamma \cup \Delta(a_1, \ldots, a_n) \vdash \phi(a_i) \text{ for all } i = 1, \ldots, n;$$

$$\Gamma \cup \Delta(a_1, \ldots, a_n) \cup \{\phi(a)\} \vdash a = a_1 \vee \cdots \vee a = a_n$$

If $\mathcal{A}$ is a model of $\Gamma \cup \Delta$, and $\mathcal{B}$is a model of $\Gamma \cup \Delta(a_1, \ldots, a_n)$, then $\mathcal{B}$ is a splitting model of $\mathcal{A}$ for $\phi(x)$. If Γ is universal, then both $\Gamma \cup \Delta$ and $\Gamma \cup \Delta(a_1, \ldots, a_n)$ have minimal models; we shall later find conditions under which the minimal model of $\Gamma \cup \Delta(a_1, \ldots, a_n)$ is unique and can properly be called *the* splitting model (as in field theory).

A problem $\phi(x)$ is *well-posed* over Δ if it has a splitting diagram over Δ or equivalently if it has a splitting model for each model of $\Gamma \cup \Delta$. A problem $\phi(x)$ is *reducible*

over Δ if there exists a formula $\phi'(x) \in L_0$ such that $\Gamma \cup \Delta \cup \{\phi(a) \wedge \phi'(a)\}$ and $\Gamma \cup \Delta \cup \{\phi(a) \wedge \neg\phi'(a)\}$ are both consistent.

Example 1. The problem $x^2 = 2$ is well posed over $\Delta(\mathbb{Q})$, its splitting diagram is $\Delta(\mathbb{Q}(\sqrt{2}))$ and solutions are $\sqrt{2}, -\sqrt{2}$. Observe that $(x^2 = 2)$ is irreducible according to our definition as well as the classical one, but $(x^2 + 2x + 1 = 0)$ is classically reducible, but not reducible according to our definition. The reason is, strangely, that the classical definition is intensional: $(x^2+2x+1 = 0)$ is equivalent in $\mathbb{Q}$ to $(x+1 = 0)$. Our definition is extensional, which in our opinion is more in line with contemporary mathematics (and simplifies the theory considerably).

Example 2. The typical application of our theory cannot, as a rule, start with a ready-made and well-known structure and theory. These will have to be put together from case to case. Let us take the example of graph-embeddings.

Given are two finite directed graphs $\mathcal{E}_1$ and $\mathcal{E}_2, \mathcal{E}_1$ being the "smaller". We ask for all ways in which $\mathcal{E}_1$ can be embedded in $\mathcal{E}_2$. We must find a language and a formula $\phi(x)$ such that $\phi(f)$ expresses that the element f is an embedding. The obvious choice for a language is three-sorted: $x_1, x_2, x_3, \ldots$ for vertices $v \in V_1$ in $\mathcal{E}_1, y_1, y_2, \ldots$ for vertices $v \in V_2$ in $\mathcal{E}_2, f_1, f_2, \ldots$ for partial maps $f \in F$ from $\mathcal{E}_1$ into $\mathcal{E}_2$. For relations we take $x_i E_1 x_j$ to denote the edge-relationship in $\mathcal{E}_1, y_i E_2 y_j$ for the edge-relationship in $\mathcal{E}_2$ and $f_i M x_j y_k$ to denote the relation f_i maps x_j to y_k. In addition we introduce constants p_i for each vertex of $\mathcal{E}_1$ and q_i for each vertex in $\mathcal{E}_2$. The predicate $I(f_i)$ stands for f_i being a (partial) isomorphism, and $T(f_i)$ for f_i being total.

The corresponding universal theory axiomatizes all universally quantified formulas of this language which hold true in all intended models of Γ: these consist of V_1, V_2 and some arbitrary subset F' of F. Without having to write up the details, we see that for given $\mathcal{E}_1$ and $\mathcal{E}_2$ such Γ can easily be constructed. The diagram Δ describes the smallest model of Γ, namely the model with $F' = \emptyset$. Finally, the original embedding problem can now be formulated as $I(f) \wedge T(f)$; its splitting diagram is the diagram of the structure with $F' = F$. The problem is therefore well-posed over Δ. Note that we have not necessarily provided enough language facilities to describe all auxiliary problems that may become of interest in solving the embedding problem. (This was simpler in Example 1, where, by tradition, these auxiliary problems will always be chosen as polynomial equations.) The choice of appropriate notions corresponds closely to the choice of possible programming strategies; we will have more to say about this below.

Since there are only finitely many models of $\Gamma \cup \Delta$, it is obvious that we have a method for testing any problem $\phi(x)$ for reducibility and find the corresponding $\phi'(x)$; such is not the case for all instances of Example 1 (for $\mathbb{Q}$, however, a reducibility algorithm is well known). But even if we have such an algorithm "in principle", it could be extremely cumbersome.

4 The amalgamation property

Let Γ again be a universal theory. We say that Γ has the *amalgamation property*, if for all models $\mathcal{A}, \mathcal{B}_1, \mathcal{B}_2$ of Γ and all injections $f_1 : \mathcal{A} \to \mathcal{B}_1, f_2 : \mathcal{A} \to \mathcal{B}_2$, there is a model of $\mathcal{C}$ of Γ and injections $g_1 : \mathcal{B}_1 \to \mathcal{C}, g_2 : \mathcal{B}_2 \to \mathcal{C}$ such that $g_1 f_1 = g_2 f_2$. (The amalgamation property was first studied in full generality by Fraïssé [7] and A. Robinson [12]; use in connection with the notion of extension of relational systems and algebraic elements was extensively made in [8] - which was an early inspiration to some of the present work. There are some syntactic characterizations of amalgamation properties, in particular Bryars in [2].) We shall presently make the first use of it, namely to observe that the amalgamation property ensures the uniqueness of splitting diagrams for well-posed problems.

Lemma 9 *Let Γ have the amalgamation property, let Δ be a diagram consistent with Γ and let $\phi(x)$ be a well-posed problem. If $\Delta(a_1, \ldots, a_m)$ and $\Delta'(b_1, \ldots, b_n)$ are splitting diagrams of $\phi(x)$ with respect to Δ, then $m = n$ and there is a permutation s of $\{1, \ldots, n\}$ such that $\Delta'(a_1, \ldots, a_n) = \Delta^s(a_1, \ldots, a_n)$.*

Proof. Let $\mathcal{A}$ be the minimal model of $\Gamma \cup \Delta$ and let $\mathcal{B}_1, \mathcal{B}_2$ be the minimal models of $\Gamma \cup \Delta(a_1, \ldots, a_n), \Gamma \cup \Delta'(b_1, \ldots, b_m)$. Let $\mathcal{C}$ be the result of amalgamation. By the definition of splitting diagrams, each b_i is an a_j and vice versa. Namely $\mathcal{C}$ is a model of $\Gamma \cup \Delta(a_1, \ldots, a_n), \cup \Delta'(b_1, \ldots, b_n)$ and each a_i and b_j is solution of $\phi(x)$, of which there are exactly $n(= m)$. Thus, $(b_1, \ldots b_n)$ is simply a permutation of $(a_1, \ldots, a_n)$.

The *relative amalgamation property* concerns a sublanguage L^* of L and injections which need preserve only the operations and relations of L^*. Γ is said to have the L^*-amalgamation property if for all models $\mathcal{A}, \mathcal{B}_1, \mathcal{B}_2$ of Γ and all L^*-injections $f_1 : \mathcal{A} \to \mathcal{B}_1, f_2 : \mathcal{A} \to \mathcal{B}_2$ there exists a model $\mathcal{C}$ of Γ such that for some L^*-injections $g_1 : \mathcal{B}_1 \to \mathcal{C}$ and $g_2 : \mathcal{B}_2 \to \mathcal{C}$ we have $g_1 f_1 = g_2 f_2$. By the same proof as for Lemma 9 we get

Corollary 10 *If Γ has the L^*-amalgamation property and $\phi(x) \in L_0^*$ is a well-posed problem, then any two splitting diagrams $\Delta(a_1, \ldots, a_n)$ and $\Delta'(b_1, \ldots, b_m)$ over a diagram Δ have the property that $m = n$ and $(\Delta'(a_1, \ldots, a_n) \cap L_0^*) = (\Delta(a_1, \ldots, a_n) \cap L_0^*)^s$ for some $s \in S_n$.*

Since many of our intended applications of the general theory concern algorithmic problems of a finite combinatorial nature (such as Example 2 of Section 3), the following sufficient condition will often be the appropriate tool to establish the amalgamation property (respectively the relative amalgamation property).

Lemma 11 *If Γ has only finite models (their cardinality is then bounded), the L^*-amalgamation property follows from the following sufficient condition:*

(i) *all models of $\Gamma \cup \Delta$ are submodels of some maximal model $\mathcal{C}$;*

(ii) *in $\mathcal{C}$ each partial L^*-automorphism can be extended to an arbitrary additional element.*

Proof. Let $f_1 : \mathcal{A} \to \mathcal{B}_1, f_2 : \mathcal{A} \to \mathcal{B}_2$ be given injections. Since $\mathcal{B}_1$ and $\mathcal{B}_2$ are submodels of $\mathcal{C}$ by assumption, there exist injections $h_1 : \mathcal{B}_1 \to \mathcal{C}$ and $h_2 : \mathcal{B}_2 \to \mathcal{C}$. The images of $\mathcal{A}$ and $\mathcal{C}$ under $h_1 f_1$ and $h_2 f_2$ are isomorphic, by an isomorphism g, say. Use the assumption to extend g to all of $h_1(B_1)$, a partial automorphism. Take $g_1 = gh_1$ and $g_2 = h_2$. Then obviously $g_1 f_1 = g_2 f_2$.

The list of universal theories for which the amalgamation property is well known is rather long. It includes those theories for which a Galois theory is worked out: fields (see e.g. [1]), differential fields [9]; some where there are some rudiments established: Boolean algebra, various geometries; and finally many for which the amalgamation property is established in other contexts: partially ordered sets, cylindrical algebras, etc.).

5 The group of a problem

Before the advent of modern algebra, the group of a polynomial $p(x)$ over a field $\mathcal{F}$ was defined essentially as follows (e.g. [11]): Let $\mathcal{E} = \mathcal{F}(a_1, \ldots, a_n)$ be the splitting field of $p(x)$ over $\mathcal{F}$ and let $(a_1, \ldots, a_n)$ be the roots of $p(x) = 0$. Then the group of $p(x)$ consists of all permutations s of these roots which preserve all rational relations. That is, whenever $q(a_1, \ldots, a_n)$ is polynomial in $(a_1, \ldots, a_n)$ with coefficients in F then we ask that $q(a_1, \ldots, a_n) = 0$ iff $q(s(a_1), \ldots, s(a_n)) = 0$. - Our definition is a straightforward generalization of this definition.

Let Γ be a universal theory with the amalgamation property, let Γ be a diagram consistent with Γ, let $\phi(x)$ be well-posed with respect to Δ and let $\Delta(a_1, \ldots, a_n)$ be a splitting diagram. Let

$$\begin{aligned} G_\Delta(\phi) \;=\; & \{s \in S_n : \Gamma \cup \Delta(a_1, \ldots, a_n) \vdash \rho \equiv \rho^s \\ & \text{for all constant } \rho \in L_0(a_1, \ldots, a_n)\} \end{aligned}$$

Theorem 12 *$G_\Delta(\phi) = \{s \in S_n : \Delta(a_1, \ldots, a_n) = \Delta^s(a_1, \ldots, a_n)\}$ is a group and does not depend on the particular choice of a splitting diagram.*

Proof. The last two facts follow obviously from the characterization of $G_\Delta(\phi)$ as a group leaving the set $\Delta(a_1,\ldots,a_n)$ invariant, and from Lemma 9 according to which any other splitting diagram must be of the form $\Delta^t(a_1,\ldots,a_n)$ for some $t \in S_n$. This makes the corresponding group a conjugate (by t) subgroup of S_n to $G_\Delta(\phi)$.

To establish the second characterization we make the preliminary remark that by the definition of a diagram we have for all constant $\rho \in L_0(a_1,\ldots,a_n)$ either $\Gamma \cup \Delta(a_1,\ldots,a_n) \vdash \rho$ or $\Gamma \cup \Delta(a_1,\ldots,a_n) \vdash \neg\rho$.

Let $s \in S_n$ be such that $s \notin G_\Delta(\phi)$, hence $\Gamma \cup \Delta(a_1,\ldots,a_n) \nvdash \rho \equiv \rho^s$ for some $\rho \in L_0(a_1,\ldots,a_n)$. By our remark then either $\Gamma\cup\Delta(a_1,\ldots,a_n) \vdash \rho$ and $\Gamma\cup\Delta(a_1,\ldots,a_n) \vdash \neg\rho^s$ or $\Gamma \cup \Delta(a_1,\ldots,a_n) \vdash \neg\rho$ and $\Gamma \cup \Delta(a_1,\ldots,a_n) \vdash \rho^s$. In the first case, we would have $\Gamma \cup \Delta^s(a_1,\ldots,a_n) \vdash \rho^s$ by a permutation of the proof; if s were such that $\Delta(a_1,\ldots,a_n) = \Delta^s(a_1,\ldots,a_n)$, then we would get a contradiction. Similarly in the second case. Hence $\Delta(a_1,\ldots,a_n) = \Delta^s(a_1,\ldots,a_n)$ implies $s \in G_\Delta(\phi)$.

Conversely, let $s \in G_\Delta(\phi)$. Then, clearly, $\Gamma\cup\Delta(a_1,\ldots,a_n) \vdash \neg\rho$ and $\Gamma\cup\Delta(a_1,\ldots,a_n) \vdash \rho \supset \rho^s$ for all $\rho \in \Delta(a_1,\ldots,a_n)$ by logic and assumption respectively. Hence $\Gamma \cup \Delta(a_1,\ldots,a_n) \vdash \rho^s$ for all $\rho \in \Delta(a_1,\ldots,a_n)$. Since $\Delta(a_1,\ldots,a_n)$ is a diagram, we conclude $\Delta(a_1,\ldots,a_n) = \Delta^s(a_1,\ldots,a_n)$.

The second characterization of $G_\Delta(\phi)$ also allows to relate it to the automorphism group characterization of Galois groups. Let $\mathcal{A}$ be the minimal model (Lemma 4) of $\Gamma \cup \Delta$ and let $\mathcal{A}(a_1,\ldots,a_n)$ be the minimal model of $\Gamma \cup \Delta(a_1,\ldots,a_n)$.

Corollary 13 *Let $G(\mathcal{A}(a_1,\ldots,a_n)/\mathcal{A})$ be the group of all automorphisms of $\mathcal{A}(a_1,\ldots,a_n)$ leaving the submodel $\mathcal{A}$ pointwise fixed. Then its group is isomorphic to $G_\Delta(\phi)$.*

Let L^* be a sublanguage of L and let $\Gamma, \Delta, \phi(x), \Delta(a_1,\ldots,a_n)$ as before, but assume that $\phi(x) \in L_0^*$ and Γ has the L^*-amalgamation property. By the same reasoning as above we define the group

$$G^*_\Delta(\phi) = \{s \in S_n : \Gamma \cup \Delta(a_1,\ldots,a_n) \vdash \rho \equiv \rho^s \text{ for all constant } \rho \in L_0^*(a_1,\ldots,a_n)\}$$

and obtain the alternative characterizations:

Corollary 14 $G^*_\Delta(\phi) = \{s \in S_n : \Delta(a_1,\ldots,a_n) \cap L_0^* = \Delta^s(a_1,\ldots,a_n) \cap L_0^*\} \cong G^*(\mathcal{A}(a_1,\ldots,a_n)/\mathcal{A})$, *where in this latter group we take L^*-automorphisms.*

The main advantage of $G^*_\Delta(\phi)$ over $G_\Delta(\phi)$ will be that it is possibly larger. For example: if L is the language of ordered fields and L^* omits the ordering predicate, then $G_\Delta(p(x) = 0)$ is the trivial group, while $G^*_\Delta(p(x) = 0)$ is the group usually studied in Galois theory. Precisely:

Corollary 15 *$G_\Delta(\phi)$ is a subgroup of $G^*_\Delta(\phi)$.*

We shall be interested in the structure of the groups $G_\Delta(\phi)$ (resp. $G^*_\Delta(\phi)$). One simple result can be stated even here:

Lemma 16 *If $\phi(x)$ is reducible over Δ then $G_\Delta(\phi)$ is intransitive, and conversely; analogously for L^*-reducibility and $G^*_\Delta(\phi)$.*

Proof. Let $\Delta(a_1, \ldots, a_n, b_1, \ldots, b_m)$ be a splitting diagram of ϕ over Δ and assume that $\phi(a_i) \wedge \phi'(a_i)$ holds for all a_i, while $\phi(b_i) \wedge \neg\phi'(b_i)$ holds for all b_i. Let $t \in G_\Delta(\phi)$, then t cannot move an a_i to b_j by definition $G_\Delta(\phi)$.

Conversely, suppose $\phi(x)$ irreducible. Let a_i and a_j be any two solutions of ϕ. Since $\phi(x)$ is irreducible, we have $\Gamma \cup \Delta(a_1, \ldots, a_n) \vdash \rho(a_i) \equiv \rho(a_j)$ for all $\rho(x) \in L_0$. Hence $\Delta(a_i) = \Delta(a_j)$ and the two minimal models $\mathcal{A}(a_i)$ and $\mathcal{A}(a_j)$ are isomorphic by an isomorphism s which maps a_i to a_j and fixes $\mathcal{A}$ pointwise. Consider the diagram

$$\begin{array}{ccc} \mathcal{A}(a_i) & \longrightarrow & \mathcal{A}(a_1, \ldots, a_n) \\ \downarrow \mathrm{s} & & \vdots\ \mathrm{t} \\ \mathcal{A}(a_j) & \dashrightarrow & \mathcal{C} \end{array}$$

which exists by amalgamation and in which we may choose $\mathcal{C} = \mathcal{A}(a_1, \ldots, a_n)$. Thus, t extends s to an automorphism in $G(\mathcal{A}(a_1, \ldots, a_n)/\mathcal{A})$; t moves a_i to a_j and therefore, generally, $G_\Delta(\phi)$ is transitive.

6 Resolvents and discriminants

The group $G_\Delta(\phi)$ consists of all permutations of $s \in S_n$ which leave the validity of all formulas $\phi(a_1, \ldots, a_n) \in L_0$ unchanged. Since $G_\Delta(\phi)$ is finite, we expect that finitely many such ρ suffice, indeed one.

Theorem 17 *Let Γ be a universal theory with the amalgamation property, let Δ be a diagram consistent with Γ, let $\phi(x)$ be well-posed and $\Delta(a_1, \ldots, a_n)$ a splitting diagram for $\phi(x)$ over Δ. Then there exists a formula $\rho(a_1, \ldots, a_n) \in L_0(a_1, \ldots, a_n)$ such that $G_\Delta(\phi) = \{s \in S_n : \rho(a_1, \ldots, a_n) = \rho^s(a_1, \ldots, a_n)\}$.*

Proof. Consider $S - G_\Delta(\phi) = \{t_1, t_2, \ldots, t_k\}$. We build up formulas $\rho_i, i = 1, \ldots, k$, which falsify $t_j \in G_\Delta(\phi)$ for all $j < i$. Suppose $\rho_i \in L_0(a_1, \ldots, a_n)$ has been found and, without loss of generality $\Gamma \cup \Delta(a_1, \ldots, a_n) \vdash \rho_i$; we assume $\rho_i = \rho_i^s$ for all $s \in G_\Delta(\phi)$ but $\Gamma \cup \Delta(a_1, \ldots, a_n) \vdash \neg\rho_i^{t_j}$ for all $j < i$. Consider now t_i. If $\Gamma \cup$

$\Delta(a_1,\ldots,a_n) \vdash \neg\rho_i^{t_i}$ then let ρ_{i+1} be the old formula ρ_i, otherwise use the assumptions as follows: since $\Delta(a_1,\ldots,a_n) \neq \Delta^{t_i}(a_1,\ldots,a_n)$, there is $\alpha \in \Delta(a_1,\ldots,a_n)$ for which $\alpha^{t_i} \notin \Delta(a_1,\ldots,a_n)$. For such an α we let ρ_{i+1} be $\rho_i \wedge \bigwedge_{s\in G_\Delta(\phi)} \alpha^s$. Observe that $\Gamma \cup \Delta(a_1,\ldots,a_n) \vdash \rho_{i+1}$ since $\Gamma \cup \Delta(a_1,\ldots,a_n) \vdash \rho$ and $\Gamma \cup \Delta(a_1,\ldots,a_n) \vdash \alpha$ hence $\Gamma \cup \Delta^s(a_1,\ldots,a_n) \vdash \alpha^s$ for all $s \in G_\Delta(\phi)$. We have $\rho_{i+1}^s = \rho_{i+1}$ for all $s \in G_\Delta(\phi)$ by construction. Finally, we obviously have $\Gamma \cup \Delta(a_1,\ldots,a_n) \vdash \neg\rho_{i+1}^{t_i}$ because α^{t_i} is a conjunctive member of $\rho_{i+1}^{t_i}$, α being such in ρ_{i+1}. The formula whose existence we claimed can be taken as ρ_{k+1}.

We call a formula ρ for which $G_\Delta(\phi) = \{s \in S_n : \rho = \rho^s\}$ a *strong resolvent*; knowing one reduces the calculation of the group of ϕ to a finite combinatorial computation, very much like the computation of the automorphism group of a graph. In classical Galois theory, the resolvent $q(x_1,\ldots,x_n) = 0$ of a polynomial equation $p(x) = 0$ serves a very similar role: the group of $p(x)$ is the set of permutations $s \in S_n$ which leave the (numerical) value of q unchanged. In analogy, therefore, we call a formula $\rho \in L_0(a_1,\ldots,a_n)$ a *resolvent* of $\phi(x)$ over Δ if $G_\Delta(\phi) = \{s \in S_n : \Gamma \cup \Delta(a_1,\ldots,a_n) \vdash \rho \equiv \rho^s\}$.

In any case, we have so far established the existence of resolvents and strong resolvents only non-effectively, indeed under the assumption that we know the group $G_\Delta(\phi)$ already (which is somehow circular, since we would like to use resolvents to calculate the group!). Classically, resolvents can be computed for fields which have an effective reducibility algorithm, [11]. This generalizes (see Section 10).

Consider the equation $ax^2 + bx + c = 0, a, b, c \in \mathbb{Q}$. Its group is either the trivial group $\{e\}$ or S_2 according to the value of the discriminant $b^2 - 4ac$. Thus, we have for a family of parametrized problems a method to decide, for any given values of the parameters, which permutation group is the group of each problem, to "discriminate" between the various possible groups. Our next lemma establishes (non-effectively) the existence of a set of *discriminants* for parametrized problems.

Let $\psi(a,x)$ be a parametrized problem, with parameter a. Assume that $\psi(a,x)$ is well-posed in all models of $\Gamma \cup \Delta(a)$, where $\Delta(a) \supseteq \Delta$. Let $G_1, G_2, \ldots$ be the groups $G_{\Delta(a)}(\psi(a,x))$ for all $\Delta(a)$, (including $\Delta(a) = \Delta$, i.e. $a \in L_0$). Observe that there are only finitely many such groups: the assumption concerning well-posedness implies that

$$\Gamma \cup \Delta \vdash \forall x \bigwedge_k \exists x_1 \cdots \exists x_k \Big(\bigwedge_{i=1}^{k} \psi(x, x_1) \wedge \bigwedge_{1\leq i<j\leq k} x_i \neq x_j \Big)$$

which by Corollary 3 reduces to a finite disjunction. Hence all $G_{\Delta(a)}(\psi(a,x))$ are subgroups of some S_n with $n \leq m$ for some fixed finite m.

Let $\Delta(a) \supseteq \Delta$ be given. $G_{\Delta(a)}(\phi)$ is determined by its resolvent $\rho_{\Delta(a)}$, and we have

$$\Gamma \cup \Delta(a) \cup \{\psi(a,a_1),\ldots,\psi(a,a_n)\} \vdash \bigvee_G \Big(\bigwedge_{s\in G} \rho \equiv \rho^s \wedge \bigwedge_{s\notin G} \rho \not\equiv \rho^s \Big)$$

where G runs over all conjugates of $G_{\Delta(a)}(\psi(a,x))$ in S_n. Altogether there are only finitely many formulas from $\Delta(a)$ that are used in this proof, let us conjunctively collect them to a formula $\delta_{\Delta(a)}(a)$. Consider this to be done for each consistent $\Delta(a) \supseteq \Delta$. Since the language $L_0(a)$ is denumerable, there are at most countably many different such formulas, which fall into finitely many classes, according to the finitely many possible groups $G_{\Delta(a)}(\psi(a,x))$. In all models of $\Gamma \cup \Delta$, therefore, we have for each element b of that model at least one of the $\delta_{\Delta(a)}(b)$ holding, thus

$$\Gamma \cup \Delta \models \forall x \bigwedge_{\Delta(a) \supseteq \Delta} \delta_{\Delta(a)}(x)$$

and by Corollary 3 again, we get a finite set of formulas $\delta_{\Delta(a)}(x)$, classified according to the possible groups for which the disjunction is already true. For each group $G = G_{\Delta(a)}(\psi(a,x))$ let $\delta_G(x)$ disjunctively collect the corresponding (finitely many) $\delta_{\Delta(a)}(x)$. These are our discriminants.

Theorem 18 *Let Γ be a universal theory with the amalgamation property, let Δ be a diagram and let $\psi(a,x)$ be a parametrized problem which is well-posed in all consistent extensions $\Delta(a)$ of Δ. Then there exists a finite set of formulas $\delta_i(x) \in L_0$ and a corresponding set of groups G_i such that for each $\Delta(a)$ we have $G_{\Delta(a)}(\psi(a,x)) \cong G_i$ for exactly one i, namely the one for which $\Gamma \cup \Delta \vdash \delta_i(a)$. In short: discriminants exist.*

The fact seems to be new even to classical Galois theory, where discriminants (in our generalized sense) are known only for some special classes of problems and for polynomials of small degrees (for degrees two and three they are the classical discriminants).

There are two obvious problems left open from this section, namely to effectively determine resolvents and discriminants. These problems will be solved below under some hypotheses which are reasonable in our context (of algorithmic problems); as for general methods, much work remains to be done.

7 The group of a program

Let Γ as always be a universal theory and Δ a diagram consistent with Γ; suppose that Γ has the amalgamation property and that $\phi(x)$ is a well-posed problem with respect to Δ. We shall now describe an effective procedure, which allows us to compute the group $G_\Delta(\phi)$; this procedure is based upon two assumptions:

(i) Assume that the universal theory of $\Gamma \cup \Delta$ is decidable (hence, by Lemma 5 we can decide consistency of added constant formulas).

(ii) Assume that we have a solution program π for $\phi(x)$ relative to some auxiliary problem $\psi(u, v)$ which is assumed well-posed in all extensions Δ' of Δ and all values of the parameter u therein.

These assumptions may be modified for what follows by relativizing to a sublanguage L^* of L. In that case ψ and π are also to be restricted to the vocabulary of L^*; the computed group will then be $G^*_\Delta(\phi)$.

To determine the group of the program, we mentally follow a computation path according to π. The initial part of the path, up to the first time the auxiliary procedure for the solution of a problem ψ is called, is obviously determined by Δ alone. The first procedure call $v_1, \ldots, v_{m_1} := \psi(\tau_1, v)$ has as input a constant term $\tau_1 \in L_0$. We now determine the "degree" m_1 of $\psi(\tau_1, v)$; this is maximal m such that

$$\Gamma \cup \Delta \cup \bigcup_{j=1}^{m} \{\psi(\tau_1, a_{1j})\} \cup \bigcup_{j \neq k} \{a_{1j} \neq a_{1k}\}$$

is consistent. By assumption, m_1 can be effectively determined. Progressing further along the path (using the new constants a_{1j} symbolically to process the assignment statements, we may come to a point (in an **if** or **while**-statement), where we need to decide a relation $R(\mu_1, \mu_2)$ for some constant terms $\mu_1 \in L_0(a_{11}, \ldots, a_{1m_1})$. This is again done by adding a consistent formula (either $R(\mu_1, \mu_2)$ or $\neg R(\mu_1, \mu_2)$) to what we have collected so far, i.e. to

$$\Gamma \cup \Delta \cup \bigcup_{j=1}^{m_1} \{\psi(\tau_1, a_{1j})\} \cup \bigcup_{1 \leq j \leq k \leq m_1} \{a_{1j} \neq a_{1k}\}$$

After perhaps some more such decisions we get to the second procedure call, $v_1, \ldots, v_{m_2} := \psi(\tau_2, v)$, where now the input τ_2 is a constant term of $L_0(a_{11}, \ldots, a_{1m_1})$. Let, for each $t \in S_{m_1}$, the term τ_2^t be the result of the corresponding permutation of the a_{1j} in τ_2. We determine the degrees of $\psi(\tau_2^t, v)$ and add constants and consistent formulas as we did in the first procedure call. The language L_0 is at this stage extended to

$$L_0(a_{11}, \ldots, a_{1m_1}, a_{21}^{(t_1)}, \ldots, a_{2m_{21}}^{(t_1)}, a_{21}^{(t_2)}, \ldots, a_{2m_{22}}^{(t_2)}, \ldots, a_{21}^{(t_k)}, \ldots, a_{2m_{2k}}^{(t_k)})$$

where $t_1, \ldots, t_k$ enumerates S_{m_1}. We proceed as before, adding $R(\mu_1', \mu_2')$ or $\neg R(\mu_1', \mu_2')$ as the case may be till we get to the next procedure call $v_1, \ldots, v_m := \psi(\tau_3, v)$. Also this time we need to solve (symbolically) all problems $\psi(\tau_3^t, v)$, where now t is a permutation of $\{a_{11}, \ldots, a_{2m_{2k}}^{(t_k)}\}$. (We can actually restrict ourselves to such t as respects the levels, i.e. permute the a_{1j} among themselves and all the $a_{2j}^{(t)}$ among themselves.) In this fashion we proceed, adding constants and formulas until the symbolic execution of the program terminates (which it will after finitely many steps,

since π is by assumption an always terminating solution program). In the course of the execution we have also computed the solutions $a_1, \ldots, a_n$ of $\phi(x)$ symbolically, namely as constant terms

$$a_i = \sigma_i(a_{11}, \ldots, a_{1m_1}, a_{21}, \ldots, a_{2m_2}, \ldots, a_{k1}, \ldots, a_{km_k}), \qquad i = 1, \ldots, n,$$

where we have suppressed the superscripts (t_i) which in this case are actually always identity elements of the corresponding permutation group. Let T be the set of permutations of all the added constants, respecting levels. We now systematically try to add consistent formulas of the form $\sigma_i = \sigma_j^t$ or $\sigma_i \neq \sigma_j^t$ for all $t \in T$ and for all $i, j = 1, \ldots, n$. This can be done effectively by assumption and the process is finite. The *group of the program* is now

$$\begin{aligned} G_\Delta(\pi) \;\; = \;\; & \{s \in S_n : \exists t \in T, \forall i \text{ the formula } \sigma_{s(i)} = \sigma_i^t \text{ has been consistently} \\ & \text{added in the symbolic execution of the program } \pi\} \end{aligned}$$

It may not appear this way at the present point, but this set is actually a group and it is isomorphic to $G_\Delta(\phi)$, the group of the problem, and does not therefore depend on the many choices made during the symbolic execution of π. This will be proved below; but let us consider a simple example first. {We thank Prof. J. Flum, Freiburg for pointing out our omission of conjugate elements in an earlier version of this construction.}

Example 3. Let us solve the equation $x^4 = 2$ by radicals with groundfield $\mathbb{Q}$. Now, Γ is field theory, which is universal if we use $0, 1, +, -, \cdot, ^{-1}$, and has the amalgamation property. Let $\Delta = \Delta\,(\mathbb{Q})$ and consider the auxiliary problem $(u = v^2)$. The corresponding procedure, then, is $v_1, v_2 := (u = v^2)$. The following is a solution program, returning solutions at y_1, y_2, y_3, y_4:

$$x_1, x_2 := (2 = x^2); \qquad y_1, y_2 := (y = x_1^2); \qquad y_3, y_4 := (y = x_2^2)$$

The levels of constants are as follows:

$a_{11}, a_{12},$

$a_{21}, a_{22}, a_{21}^{(t)}, a_{22}^{(t)},$

$a_{31}, a_{32}, a_{31}^{(t)}, a_{32}^{(t)},$

where in this case t is the only non-trivial element of S_2 and does perform all the permutations necessary in the build-up. We have added to $\Gamma \cup \Delta(\mathbb{Q})$ the formulas

$$a_{11}^2 = 2, a_{12}^2 = 2, a_{21}^2 = a_{11}, a_{22}^2 = a_{11}, \ldots, (a_{32}^{(t)})^2 = a_{12}$$

The solutions are $b_1 = a_{21}, \ldots, b_4 = a_{32}$. In trying all permutations of $a_{11}, \ldots, a_{32}^{(t)}$ respecting levels, we find that we can consistently add $a_{21}^{(t)} = a_{31}, a_{22}^{(t)} = a_{32}, a_{31}^{(t)} = a_{21}, a_{32}^{(t)} = a_{22}$. Thus permutation can be reduced to $\{a_{11}, a_{12}, a_{21}, a_{31}, a_{32},\}$ and the group is then easily determined as a group of order 8.

Theorem 19 *If π is a solution program for ϕ, then $G_\Delta(\pi) \cong G_\Delta(\phi)$.*

Proof. The proof consists in following the above construction and (non-effectively) augmenting the consistent sets of formulas to splitting diagrams.

Consider a consistent path in π. Let $\tau_1 \in L_0, \tau_2 \in L_0(a_{11}, \ldots, a_{1m_1}), \ldots,$
$\tau_k \in L_0(a_{11}, \ldots, a_{k-1m_{k-1}})$ be the constant terms to which the auxiliary procedure $v_1, \ldots, v_{m_i} := \psi(\tau_i, v)$ is applied consecutively. Define $\Delta_0 = \Delta$ the given diagram. Then let $\Delta_1 = \Delta(a_{11}, \ldots, a_{1m_1})$ be a splitting diagram for $\psi(\tau_1, v)$ over Δ. Let $H_1 = \{h \in S_{m_1} : \Delta_1^h = \Delta_1\}$ be the group of $\psi(\tau_1, v)$ over Δ. The conjugates of τ_2 are $\tau_2^h, h \in H$. The groups of $\psi(\tau_2^h, v)$ are obviously isomorphic (being determined by a discriminant $\delta(\tau) \in L_0(a_{11}, \ldots, a_{1m_1})$ which is equivalent to $\delta^h(\tau)$, which is $\delta(\tau^h)$). Let $\Delta_2 = \Delta_1(a_{21}^{(h_1)}, \ldots, a_{2m_2}^{(h_1)}, \ldots, a_{21}^{(h_r)}, \ldots, a_{2m_2}^{(h_r)})$ be the joint splitting diagram for all $\psi(\tau_2^{(h_i)}, v), h_i \in H_1$. Let $H_2 = \{h \in S_{m_2 \cdot r} : \Delta_2^h = \Delta_2\}$ be the group of symmetries of this diagram. Let $\Delta_3 = \Delta_2(a_{31}^{(h_1)}, \ldots, a_{3m_3}^{(h_1)}, \ldots)$ be the joint splitting diagram of the $\psi(\tau_3^{(h_i)}, v), h_i \in H_2$, etc., until finally we obtain Δ_k and $H_k = \{h : h$ is a level-respecting permutation of the a_{ij}^h and $\Delta_k^h = \Delta_k\}$. Let $\mathcal{A}(b_1, \ldots, b_n)$ be the splitting model of $\phi(x)$ over the minimal model $\mathcal{A}$ of $\Gamma \cup \Delta$, and let $\mathcal{A}(\pi)$ be the minimal model of $\Gamma \cup \Delta_k$. Since, by assumption, π solves $\phi(x)$, the solutions $(b_1, \ldots, b_n)$ are terms $\sigma_1(a_{11}, \ldots, a_{km_k}), \ldots, \sigma_n(a_{11}, \ldots, a_{km_k})$. Hence $\mathcal{A}(b_1, \ldots, b_n)$ is a submodel of $\mathcal{A}(\pi)$. Let s be any automorphism of $\mathcal{A}(b_1, \ldots, b_n)$ leaving $\mathcal{A}$ pointwise fixed, then, by amalgamation, we can complete the following diagram

$$\begin{array}{ccc} \mathcal{A}(b_1, \ldots, b_n) & \longrightarrow & \mathcal{A}(\pi) \\ \downarrow s & & \vdots\ t \\ \mathcal{A}(b_1, \ldots, b_n) & \dashrightarrow & \mathcal{C} \end{array}$$

Now $\mathcal{C}$ can obviously be restricted to $\mathcal{A}(\pi)$ and hence t_s is an automorphism of $\mathcal{A}(\pi)$ leaving $\mathcal{A}$ pointwise fixed. Let $t \in H_k$ and $q \leq k$. Then t permutes $\{a_{q1}^{(h_1)}, \ldots, a_{qm_q}^{(h_1)},$ $a_{q1}^{(h_2)}, \ldots\}, h_i \in H_{q-1}$ since the $a_{qj}^{(h_i)}$ solve $\psi(\tau^{h_i}, v)$ and the τ^{h_i} are conjugates. Observe that $t(b_i) = t(\sigma_i(a_{11}, \ldots, a_{km_k})) = \sigma_i(t(a_{11}), \ldots, t(a_{km_k}))$ since t is an automorphism of $\mathcal{A}(\pi)$. Thus, the action of t on the b_i is completely determined by the action of t on the a_{ij}. Hence, the group $G_\Delta(\pi)$ is the group of H_k restricted to $\{b_1, \ldots, b_n\}$. Let now any $s \in G_\Delta(\phi)$ be given; it determines an automorphism s of $\mathcal{A}(b_1, \ldots, b_n)$ and also an automorphism t_s of $\mathcal{A}(\pi)$, indeed, obviously $t_s \in H_k$. It follows that the restriction of H_k to $\{b_1, \ldots, b_n\}$ is actually $G_\Delta(\phi)$ and we have proved the theorem.

The theorem shows that the method for computing the group $G_\Delta(\pi)$ described at the beginning of this section produces the right group, $G_\Delta(\phi)$. If we have this group explicitly, then the construction used for the existence proof for a resolvent (Theorem 17) can obviously be made effective, using assumption (i) above. This finishes the task set at the end of Section 6, and we have

Theorem 20 *Under assumption (i) and (ii) we can explicitly construct a strong resolvent for the problem $\phi(x)$ over Δ.*

8 Solvability and the structure of $G_\Delta(\phi)$

Let Γ be a universal theory with amalgamation, Δ a diagram, $\phi(x)$ a well-posed problem and π a solution algorithm for $\phi(x)$ relative to the well-posed auxiliary problem $\psi(u, v)$. We use this information to obtain results about the group $G_\Delta(\phi)$.

Consider a sequence of consistent extensions of Δ as it is produced along a consistent path through π. Along this path, we solve auxiliary problems $\psi(\tau_q, v)$ successively for terms τ_q as follows: $\tau_1 \in L_0, \tau_2 \in L_0(a_{11}, \ldots, a_{1m_1}), \tau_3 \in L_0(a_{11}, \ldots, a_{1m_1}, a_2, \ldots, a_{2m_2})$, $\ldots, \tau_k \in L_0(a_{11}, \ldots, a_{k-1m_{k-1}})$. The solutions of $\phi(x)$ are terms $b_i = \sigma_i(a_{11}, \ldots, a_{km_k})$ $\in L_0(a_{11}, \ldots, a_{km_k})$. Let the sequence of extensions be $\Delta \subseteq \Delta(a_{11}, \ldots, a_{1m_1}) \subseteq$ $\Delta(a_{11}, \ldots, a_{1m_1}, a_{21}, \ldots a_{2m_2}) \cdots \subseteq \Delta(a_{11}, \ldots, a_{km_k})$, where the extension each time is one to a splitting diagram: $\Delta(a_{11}, \ldots, a_{qm_q}, a_{q+11}, \ldots, a_{q+1,m_{q+1}})$ is a splitting diagram of $\psi(\tau_{q+1}, v)$ over $\Delta(a_{11}, \ldots, a_{qm_q})$. Let $\mathcal{A} \subseteq \mathcal{A}(a_{11}, \ldots, a_{1m_1}) \subseteq \cdots \subseteq \mathcal{A}(a_{11}, \ldots, a_{qm_q})$ $\subseteq \cdots \subseteq \mathcal{A}(a_{11}, \ldots, a_{km_k})$ be the corresponding sequence of minimal models.

For each $q = 0, \ldots, k$ we now define a group of automorphisms

$$G_q = G(\frac{\mathcal{A}(a_{11}, \ldots, a_{km_k})}{\mathcal{A}(a_{11}, \ldots, a_{qm_q})}), \quad G_0 = G(\mathcal{A}(a_{11}, \ldots, a_{km_k})/\mathcal{A}).$$

Thus $G_k = \{e\}$, the trivial group.

Lemma 21 *$G_{q+1} < G_q$ for $q = 0, \ldots, k-1$.*

Proof. Let $s \in G_q, t \in G_{q+1}$. We have to show $s^{-1}ts \in G_{q+1}$, i.e. $s^{-1}(ts(x) = x$ for all $x \in \mathcal{A}(a_{11}, \ldots, a_{q+1m_{q+1}})$.

By construction, $\mathcal{A}(a_{11}, \ldots, a_{q+1m_{q+1}})$ is the splitting model of $\mathcal{A}(a_{11}, \ldots, a_{qm_q})$ with respect to the problem $\psi(\tau_{q+1}, v), \tau_{q+1} \in L_0(a_{11}, \ldots, a_{qm_q})$. Therefore $s \in G_q$ leaves

τ_{q+1} fixed and hence must simply permute the solutions $a_{q+1}, \ldots, a_{q+1m_{q+1}}$. It follows that

$$s^{-1}ts(a_{q+1i}) = s^{-1}t(a_{q+1s(i)}) = s^{-1}(a_{q+1s(i)}) = a_{q+1i}$$

Therefore $s^{-1}ts(a_{q+1i}) = a_{q+1i}$ and $s^{-1}ts(\tau) = \tau$ for any term τ built up from these constants.

Lemma 22 $G_q/G_{q+1} \cong G_{\Delta(a_{11},\ldots,a_{qm_q})}(\psi(\tau_{q+1}, v))$.

Proof. We construct an isomorphism f of $G_{\Delta(a_{11},\ldots,a_{qm_q})}(\psi(\tau_{q+1}, v))$ onto G_q/G_{q+1} as follows: consider $s \in G_{\Delta(a_{11},\ldots,a_{qm_q})}(\psi(\tau_{q+1}, v))$ as an automorphism of $\mathcal{A}(a_{11}, \ldots, a_{q+1m_{q+1}})$ leaving $\mathcal{A}(a_{11}, \ldots, a_{qm_q})$ pointwise fixed; this is correct according to Corollary 5. Now use the amalgamation property as before to extend s to an automorphism t of $\mathcal{A}(a_{11}, \ldots, a_{km_k})$ leaving $\mathcal{A}(a_{11}, \ldots, a_{qm_q})$ still pointwise fixed. Then t is in G_q by definition, and we can define $f(s) = tG_{q+1} \in G_q/G_{q+1}$.

This definition of f does not actually depend on the particular extension t chosen. Namely, let t_1, t_2 be any two such extensions. Consider $t^* = t_2^{-1}t_1$. It suffices to show $t^*(x) = x$ for all $x \in \mathcal{A}(a_{11}, \ldots, a_{q+1m_{q+1}})$, i.e. that $t^* \in G_{q+1}$. Both t_1 and t_2 leave $\mathcal{A}(a_{11}, \ldots, a_{qm_q})$ pointwise fixed, hence $t^*(\tau_{q+1}) = \tau_{q+1}$ as $\tau_{q+1} \in L_0(a_{11}, \ldots, a_{qm_q})$. It follows that t^* permutes the solutions $a_{q+11}, \ldots, a_{q+1m_{q+1}}$ by the law $t^*(a_{q+1i}) = t_2^{-1}t_1(a_{q+1i}) = t_2^{-1}s(a_{q+1i}) = s^{-1}s(a_{q+1i}) = a_{q+1i}$, that is, it leaves them pointwise fixed and therefore t^* fixes all of $\mathcal{A}(a_{11}, \ldots, a_{q+1m_{q+1}})$ pointwise.

We now prove this f to be the desired isomorphism between the two groups.

(a) f is one-to-one: suppose $f(s_1) = f(s_2)$, which means $t_1G_{q+1} = t_2G_{q+1}$. Then $t_1 \in t_2G_{q+1}$ and $t_1 = t_2t^*$ for some $t^* \in G_{q+1}$. Let $x \in \mathcal{A}(a_{11}, \ldots, a_{q+1m_{q+1}})$. Let us compute $s_1(x)$ as follows: $s_1(x) = t_1(x) = (t_2t^*)(x) = t_2(t^*(x)) = t_2(x)$ since t^* fixes $\mathcal{A}(a_{11}, \ldots, a_{q+1m_{q+1}})$ pointwise. But $t_2(x) = s_2(x)$ by construction of t_2 and hence $s_1(x) = s_2(x)$ as desired: $s_1 = s_2$.

(b) f is onto: let $tG_{q+1} \in G_q/G_{q+1}$ with $t \in G_q$ be given. Let t^q be the restriction of t to $\mathcal{A}(a_{11}, \ldots, a_{q+1_{q+1}})$. By construction t^q permutes the a_{q+1i} (since it fixes, as an element of G_q, the model $\mathcal{A}(a_{11}, \ldots, a_{qm_q})$ pointwise). Thus $t^q \in G(\mathcal{A}(a_{11}, \ldots, a_{q+1m_{q+1}})/\mathcal{A}(a_{11}, \ldots, a_{qm_q}))$ and we may take s as t^q for $f(s) = tG_{q+1}$.

(c) f is homomorphic: since $G_{q+1} \lhd G_q$ we may compute $f(s_1) \cdot f(s_2) = t_1G_{q+1}t_2G_{q+1}$ $= t_1(t_2G_{q+1}t_2^{-1})t_2G_{q+1}$. The latter simplifies to $t_1t_2G_{q+1}$. But $t_1t_2G_{q+1}$ can serve as $f(s_1s_2)$ because t_1t_2 agrees with s_1s_2 on $\mathcal{A}(a_{11}, \ldots, a_{q+1m_{q+1}})$ and leaves $\mathcal{A}(a_{11}, \ldots, a_{qm_q})$ pointwise fixed.

Lemma 23 $G_\Delta(\phi) \lhd G_0$.

Proof. We show $s^{-1}ts \in G_\Delta(\phi)$ for each $t \in G_\Delta(\phi)$ and $s \in G_0$. Note that t may be considered an automorphism in $G(\mathcal{A}(b_1, \ldots, b_n)/\mathcal{A})$ by Corollary 13 and extended to an automorphism in G_0 by amalgamation. Elements $s \in G_0$ respect solutions of $\phi(x)$, being automorphisms; hence $s^{-1}ts$ moves each solution b_i of $\phi(x)$ to some such solution and leaves $\mathcal{A}$ pointwise fixed; i.e. $s^{-1}ts \in G_\Delta(\phi)$.

We have already observed that the family of groups of the auxiliary problems $\psi(a, v)$ is finite (Theorem 18). In analogy to classical Galois theory we call a group G ψ-solvable, if G has a subnormal series

$$G = H_0 \rhd H_1 \rhd H_2 \rhd \cdots H_p \rhd = \{e\}$$

where each H_p/H_{p+1} is isomorphic to a group $G_{\Delta(a)}(\psi(a, v))$. However, the classical relation between the solvability of a problem and the solvability of its group does not hold in general: the following example shows a ψ-solvable group $G_\Delta(\phi)$, where ϕ is not solvable relative to ψ.

Example 4. Let W be a set of words over $\{a, b\}$ of length at most two. Consider the two-sorted theory describing such situations with variables $x, y, \ldots$ for letters, $u, v, \ldots$ for words, predicate $h(w, x)$ true if x is the first letter of w and w is not the empty word λ, predicate $t(w, x)$ correspondingly for the last letter. Let $s(w)$ be true of words of length 0 and 1 and of symmetric words of length two; let $l(w)$ be true only of words of length two. Let L^* consist of variables $u, v, \ldots$ and predicates s and l. Let Δ be the diagram for $w = \{\lambda\}$. The theory obviously has the L^*-amalgamation property. Consider $\phi(w)$ defined by $s(w) \wedge l(w)$ and $\psi(w)$ defined by $\neg s(w) \wedge l(w)$. Both are well-posed and have S_2 as their groups. But ϕ is obviously not solvable relative to ψ.

9 Lower bounds of complexity

Putting the lemmas of the previous sections together, we get the fundamental tool for our complexity results:

Theorem 24 $\mid G_\Delta(\phi) \mid \leq \Pi_{q=0}^{k-1} \mid G_{\Delta(a_{11},\ldots,a_{qm_q})}(\psi(\tau_{q+1}, v)) \mid$

Proof. By Lemma 22 we have

$$G_q/G_{q+1} \cong G_{\Delta(a_{11},\ldots,a_{qm_q})}(\psi(\tau_{q+1}, v))$$

hence by Lagrange's theorem

$$\mid G_q \mid = \mid G_{q+1} \mid \cdot \mid G_{\Delta(a_{11},\ldots,a_{qm_q})}(\psi(\tau_{q+1}, v)) \mid, \quad q = 0, 1, \ldots, k-1$$

Since $\mid G_k \mid = \mid \{e\} \mid = 1$, we get

$$\mid G_0 \mid = \Pi_{q=0}^{k-1} \mid G_{\Delta(a_{11},\ldots,a_{qm_q})}(\psi(\tau_{q+1}, v)) \mid$$

and the theorem follows from $\mid G_\Delta(\phi) \mid \leq \mid G_0 \mid$, which is a consequence of Lemma 23.

The number k in Theorem 24 is the relative complexity of $\phi(x)$ with respect to the auxiliary problem $\psi(u, v)$ as realized by the program π. The theorem gives us a tool to estimate a lower bound for k independent of the program π (whose only function could have been to obtain the group $G_\Delta(\phi)$, but for that any solution program would serve). The actual lower bounds will, of course, be computed for the individual problems. We mention here only two special cases, where a simple formula results.

Corollary 25 *If the groups $G_{\Delta(a)}(\psi(a, v))$ are always about the same size and if m is the maximum of these sizes, then $log_m \mid G_\Delta(\phi) \mid \leq k$.*

Proof. By Theorem 24 we have $\mid G_\Delta(\phi) \mid \leq \Pi_{q=0}^{k-1} m = m^k$.

Corollary 26 *If the groups $G_{\Delta(a)}(\psi(a, v))$ grow exponentially in size as the program progresses, say the qth group is of the size m^{q+1}, then $\sqrt{log_m \mid G_\Delta(\phi) \mid}^2 \leq k + 1$.*

Example 3 (continued). The group of $(x^4 - 2 = 0)$ over $\mathbb{Q}$ has order 8; the group of the auxiliary problem $(x^2 - a = 0)$ has at most order 2. Corollary 25 yields $k \geq \log_2 8 = 3$. This is actually the exact number of times our solution program does call the auxiliary procedure.

Observe that these lower bounds are bounds for the complexity of getting all solutions of the problem $\phi(x)$. Very often, actually, one is more interested in just getting one solution to a problem, and therefore in knowing bounds for the complexity of this restricted problem.

In a preliminary step we need to treat the case that the problem $\phi(x)$ is *reducible*. We would then, of course, only compute the solutions of an irreducible subproblem, preferably that in which the first solution is obtained with the least number of calls of the auxiliary procedure. By Lemma 16 a necessary and sufficient condition for the reducibility of $\phi(x)$ is the intransitivity of the group $G_\Delta(\phi)$. Thus, without even knowing the reduced problems we know their groups: these are the restrictions of $G_\Delta(\phi)$ to its orbits.

Example 5. The present example, besides giving an instance of an intransitive Galois group, is intended to demonstrate that a problem, which on the face of it seems to have but one solution (and therefore would not obviously lend itself to an application of Theorem 24 to get meaningful lower bounds), in fact very easily adapts to our scheme. The problem is: how many additions do we need at least in order to multiply two natural numbers? Let m and n be positive integers. The set of m-expressions is the smallest set of words containing $m, (m+m)$ and containing $(a+b)$ for any two m-expressions a and b. Let E_k be a unary predicate symbol and let $E_k(x)$ hold of an m-expression x iff x contains exactly k symbols m. For m-expressions a, b, c, d let $T(a, b, c, d)$ hold if either $d = ((a+b)+c)$ or $d = (a+(b+c))$. We shall use $T(u_1, u_2, u_3, v)$ as auxiliary problem and $E_n(x)$ as main problem. Observe that every m-expression can be obtained from m and $(m+m)$ by repeated application of the auxiliary problem. Namely: proceed by induction on the length of an m-expression d. It is of the form m or $(m+m)$ or $(e+f)$, where at least one of e and f is an m-expression of shorter length and not m; thus d is $(e+(f_1+f_2))$, say, and therefore a solution of $T(e, f_1, f_2, v)$. We now have to set up the appropriate universal theory Γ and diagram Δ, ensure the amalgamation property and discuss $G_\Delta(E_n(x))$. We make use of the additional unary operation symbols P_1, P_2 which each gives one of the two possible components a, b of the m-expression $(a+b)$. Let now L be the language with variables $x, y, z, u, v, \ldots$ for m-expressions, constants m and $(m+m)$, operations *add* (yielding $add(a, b) = (a+b)$ for m-expressions a and b), P_1, P_2, and predicates T and $E_1, \ldots, E_n$. Let Γ be the universal theory for all structures which contain at least the m-expressions m and $(m+m)$ and at most all m-expressions x for which $E_k(x)$ for some $k \leq n$. Let Δ be the diagram of the minimal such structure. Note that both the problem $E_n(x)$ and the auxiliary problems $T(u_1, u_2, u_3, v)$ are well-posed, and that $\Gamma \cup \Delta$ has the L^*-amalgamation property, where L^* is obtained from L by dropping the operation symbols *add*, P_1, P_2 (this is shown again by using Lemma 11). Observe that $T(a, b, c, v)$ has the group S_2 if it has a solution at all. The group $G^*_\Delta(E_n(x))$ is intransitive. We may for example take the orbit of the solution path $T(m, m, m+m, a_1)$; $T(m+m, m+m, a_1a_2)$; $T(a_1, a_1, a_2, a_3)$; $T(a_2, a_2, a_3, a_4)$; ... for n of the form $2^k, k \geq 1$. The group $G^*_\Delta(E_n(x))$ restricted to this orbit is obviously $S_2 * S_2 * \cdots * S_2 (k-1$ times the wreath-product). (For another example see [5], a further example is in [6].)

Let us now return to the problem of the complexity of obtaining a single solution in terms of the group of the original problem. By what was said above, we may, without loss of generality, make the simplifying assumption that $\phi(x)$ is irreducible. Let π be a solution program of $\phi(x)$ relative to $\psi(u, v)$, and let a path through π be given along which the first time solution of $\phi(x)$ appears is after p procedure calls for $\psi(u, v)$, i.e. in the model $\mathcal{A}(a_{11,\ldots,}a_{km_k})$, say $b_1 = \sigma_1(a_{11,\ldots,}a_{km_k}), \ldots, b_n = \sigma_n(a_{11,\ldots,}a_{km_k}), \sigma_i \in L_0(a_{11,\ldots,}a_{km_k})$. To shorten notation, we again write $\mathcal{A}_q$ for $\mathcal{A}(a_{11,\ldots,}a_{qm_q})$ for $q = 1, \ldots, k$, and recall $G_q = G(\mathcal{A}_k/\mathcal{A}_q), G_0 = G(\mathcal{A}_k/\mathcal{A}) \rhd G_\Delta(\phi)$.

Let h be the homomorphism $h : G_0 \to G_\Delta(\phi)$ defined by: $h(t)(i) = j$, where j is the index of that σ_j for which $\sigma_i^t = \sigma_j$ holds in $\mathcal{A}_k$. This homomorphism maps the subgroup $G(\mathcal{A}_k/\mathcal{A}_p) = G_p$ of G_0 onto a subgroup G^* of the stabilizer subgroup

$$G_\Delta(\phi)_B = \{s \in G_\Delta(\phi) : s(i) = i \text{ for all } i \in B\}$$

where B is the set of (indices of) solutions b_i of $\phi(x)$ which belong to $\mathcal{A}_p$. Namely: if $t \in G(\mathcal{A}_k/\mathcal{A}_p)$, then $h(t)$ leaves these b_i fixed, thus $h(t) \in G_\Delta(\phi)_B$. The homomorphism h transforms the co-set representation

$$G(\mathcal{A}_k/\mathcal{A}) = t_1 G(\mathcal{A}_k/\mathcal{A}_p) + \cdots + t_r G(\mathcal{A}_k/\mathcal{A}_p)$$

with index $r = \mid G_0 \mid / \mid G_p \mid$, into $G_\Delta(\phi) = h(t_1)G^* + \cdots + h(t_r)G^*$, where $G^* \subseteq G_\Delta(\phi)_B$. The index of G^* in $G_\Delta(\phi)$ is therefore at most r and at least as large as the index of $G_\Delta(\phi)_B$ in $G_\Delta(\phi)$. Thus $G_0 \mid / \mid G_p \mid \geq \mid G_\Delta(\phi) \mid / \mid G_\Delta(\phi)_B$, and therefore as before:

Theorem 27 $\Pi_{q=0}^{p} \mid G_{\Delta(\mathcal{A}_p)}(\psi(\tau_{q+1}, v)) \mid \geq \mid G_\Delta(\phi) \mid / \mid G_\Delta(\phi)_B \mid$

From this theorem we obtain a lower bound for the first solution in the same way as in Corollaries 25 and 26; for example:

Corollary 28 *If the groups $G_{\Delta(a)}(\psi(a, v))$ are all about the same small size, and if m is the maximum of these sizes, then*

$$p \geq log_m \mid G_\Delta(\phi) \mid - log_m deg(G_\Delta(\phi))$$

Example 3 (continued). For the problem $x^4 = 2$ over $\mathbb{Q}$ the only nontrivial groups $G_\Delta(\phi)_B$ are of order 2; hence $p \geq 3 - 1 = 2$, which is the actual minimal number of square roots for the first solution of the given equation.

10 Reducibility and resolvents

So far, we have been able to determine the group of a problem under the assumption that the universal theory of Γ is decidable and we have some solution algorithm for the problem. In classical Galois theory, there is a well-known direct way to determine the group, namely via obtaining a resolvent. The method hinges on the availability of an algorithm which produces an irreducible factor for every polynomial. This method generalizes.

Let $\mu(a_1, \ldots, a_n) \in L_0(a_1, \ldots, a_n)$ be called *irreducible* over $\Gamma \cup \Delta$ if there is no $\rho(a_1, \ldots, a_n) \in L_0(a_1, \ldots, a_n)$ such that both $\rho \wedge \mu$ and $\neg\rho \wedge \mu$ are consistent with $\Gamma \cup \Delta$. A *reducibility algorithm* for $\Gamma \cup \Delta$ produces for every consistent $\eta(a_1, \ldots, a_k) \in L_0(a_1, \ldots, a_k)$ an irreducible formula $\mu(a_1, \ldots, a_k) \in L_0(a_1, \ldots, a_k)$ such that $\Gamma \cup \Delta \cup \mu$ is consistent and $\Gamma \cup \Delta \neg\mu \supset \eta$.

Theorem 29 *If Γ has the amalgamation property and $\Gamma \cup \Delta$ admits a reducibility algorithm, then there is an algorithm which determines a resolvent for each well-posed problem.*

Proof. Let $\phi(x)$ be well-posed over $\Gamma \cup \Delta$, suppose that Γ has the amalgamation property, and that $\phi(x)$ has degree n. Let $\eta(a_1, \ldots, a_n)$ be $\phi(a_1) \wedge \cdots \wedge \phi(a_n) \wedge \bigwedge_{i<j} a_i \neq a_j$. Apply the reduction algorithm to η, producing $\mu(a_1, \ldots, a_k) \in L_0(a_1, \ldots, a_n)$. We claim that μ is resolvent for ϕ. Namely: by construction, $\Gamma \cup \Delta \cup \mu \cup \eta$ is a consistent subset of $L_0(a_1, \ldots, a_n)$, which we then may extend to a splitting diagram $\Delta(a_1, \ldots, a_n)$; thus $\mu(a_1, \ldots, a_n) \in \Delta(a_1, \ldots, a_n)$. We show that for all $s \in S_n$ we have $\Gamma \cup \Delta(a_1, \ldots, a_n) \vdash \mu \equiv \mu^s$ iff for all $\rho \in L_0(a_1, \ldots, az_n)$ we have $\Gamma \cup \Delta(a_1, \ldots, a_n) \vdash \rho \equiv \rho^s$. For the nontrivial direction, let $\rho \in L_0(a_1, \ldots, a_n)$. Observe that either $\Gamma \cup \Delta \vdash \mu \supset \rho$, or $\Gamma \cup \Delta \vdash \mu \supset \neg\rho$. For if not, then both $\mu \wedge \neg\rho$ and $\mu \wedge \rho$ would be consistent with $\Gamma \cup \Delta$, contrary to the irreducibility of μ. Let us assume then $\Gamma \cup \Delta \vdash \mu \supset \rho$, (the other case is handled the same way). If $\Gamma \cup \Delta(a_1, \ldots, a_n) \vdash \mu \equiv \mu^s$, then $\mu^s \in \Delta(a_1, \ldots, a_n)$ (since $\mu \in \Delta(a_1, \ldots, a_n)$). Because $\Gamma \cup \Delta \vdash \mu \supset \rho$, we have $\Gamma \cup \Delta^s \vdash \mu^s \supset \rho^s$, and, since $\Delta^s = \Delta$, also $\Gamma \cup \Delta \vdash \mu^s \supset \rho^s$. Thus, both ρ and ρ^s are in $\Delta(a_1, \ldots, a_n)$ and a fortiori $\Gamma \cup \Delta(a_1, \ldots, a_n) \vdash \rho \equiv \rho^s$.

Corollary 30 *If Γ has the amalgamation property, $\Gamma \cup \Delta$ admits a reducibility algorithm and the universal theory of $\Gamma \cup \Delta$ is decidable, then $G_\Delta(\phi)$ can be effectively computed for every well-posed problem ϕ.*

Bibliography

[1] E. Artin. *Galois Theory* (University of Notre Dame, 1942).

[2] P. D. Bacsich and D. R. Hughes. Syntactic characterizations of amalgamation, convexity and related properties, *J. Symbolic Logic* 39 (1974), 433–451.

[3*] E. Engeler. On the solvability of algorithmic problems. In: *Logic Colloquium 73* (eds J. C. Sheperdson and H. E. Rose), North-Holland, Amsterdam (1975), 231–251.

[4*] E. Engeler. Structural relations between programs and problems. *Logic, Foundations of Mathematics and Computability Theory* (eds Butts and Hintikka), Reidel, Dordrecht (1977), 267–280.

[5*] E, Engeler. Lower bounds of Galois theory, *Asterisque*, 38/39 (1976), 45–52.

[6] R. Fraïssé. Sur l'extension aux relations de quelque proprietes des ordres, *Ann. Sci. Ecole Norm. Sup.* 71(1954), 363–388.

[7] G. Gati. Some elements of a Galois theory of the structure and complexity of the Tree Automorphism Problem, *Theoret. Computer Science* 14 (1981), 1–17.

[8] B. Jónsson. Algebraic extensions of relational systems, *Math. Scand.* 11 (1962), 179–205.

[9] I. Kaplansky. An introduction to differential algebra, *Actualites Sci. Indust.* 1251 (1957).

[10] A. Kreczmar. Programmability in fields, *Fund. Informaticae* 1 (1977), 195–230.

[11] O. Perron. *Algebra, Vol. 2*, Göschen-de Gruyter, Berlin (1933).

[12] A. Robinson. *On the Metamathematics of Algebra*, North-Holland, Amsterdam (1951).

[13] H. Seeland. *Algorithmische Theorien und konstruktive Geometrie*, Stuttgart (1978).