Proceedings of the ESSLLI 2022 Student Session, 121–127
doi.org/10.21942/uva.20368233

Completing the Picture: Complexity of the Ackermann Fragment

Daumantas Kojelis[0000−0002−1632−9498]

The University of Manchester
daumantas.kojelis@manchester.ac.uk

Abstract. We give a decision procedure for the satisfiability problem of the Ackermann fragment with equality, when the number of trailing existential quantifiers is bounded by some fixed integer m, and thus establish an ExpTime upper-bound. Taking the work of R. Jaakkola [2] into account, we conclude that any Ackermann (sub-)fragment must feature at least two leading as well as an unbounded number of trailing existential quantifiers to retain NExpTime-hardness.

Keywords: Computational Logic · Decidable Fragments of First-Order Logic · Complexity Theory · Ackermann Fragment with Equality · Quantifier Prefix Fragments

1 Introduction

The Ackermann fragment is one of the classical decidable quantifier prefix fragments of first-order logic. When referring to quantifier prefix fragments we take into account exactly those formulae of first-order logic that are of the form $O_1x_1 \dots O_nx_n\psi(x_1, \dots, x_n)$, where each O_i $(i \in [1, n])$ is either the universal or existential quantifier and ψ is quantifier-, constant- and function-free. For brevity, we will refer to classes of formulae featuring the prefix $O_1x_1 \dots O_nx_n$ as $[O_1 \dots O_n]$ (or $[O_1 \dots O_n]_=$ if the equality predicate is permitted). We take O_i^m to mean that the quantifier O_i is repeated (at most) m times, and O_i^* to mean that the number of repeated quantifiers O_i is unbounded.

An example of a specific quantifier prefix fragment is $[\forall\exists\forall\exists^5]$. This class, in fact, includes the sentence describing Turing's Halting Problem [4], and, therefore, has an undecidable satisfiability problem. Consider now the prefix class $[\exists^*\forall\exists^*]_=$ - commonly referred to as the Ackermann fragment (with equality). In [1, p. 228] the authors prove that the satisfiability problem for the fragment $[\exists^2\forall\exists^*]_=$ is NExpTime-hard and claim (without proof) that a reduction to $[\forall\exists^*]_=$ can be established. However, this is not possible as the fragment $[\exists\forall\exists^*]_=$ was recently shown to be ExpTime-complete [2].

In light of such discoveries, we are still left to wonder how many trailing existential quantifiers are needed to establish NExpTime-hardness? Whilst numerous restrictions (e.g. on the signature) are discussed by E. Börger et al.

[1, Chapter 6.3-6.4], conveniently, bounds on the number of trailing existential quantifiers are disregarded, and, to the best of the author's knowledge, an open problem. We will thus show that the decision problem for the fragment $[\exists^*\forall\exists^m]_=$ for all fixed m $(2 \le m < *)$ is in ExpTime (keeping in mind that $[\exists^*\forall\exists]_=$ is PSpace-complete [1, p. 283]).

2 Decision Procedure

In this section we will provide a deterministic decision procedure that runs in doubly exponential time relative to the size of the input. In Section 3 we will introduce a trick to lower the bound to singly exponential time and thus prove the following theorem:

Theorem 1. *The satisfiability problem of the fragment* $[\exists^*\forall\exists^m]_=$ *for all fixed* m $(2 \le m < *)$ *is in* ExpTime.

Let φ be a $[\exists^*\forall\exists^m]_=$ sentence of the form $\exists x_1 \dots \exists x_n \forall y \exists z_1 \dots \exists z_m \psi$ (n here is unbounded, but m is fixed), where ψ is quantifier-free and over a signature σ, which features no constants or function symbols. We then define Θ to be the set of all possible maps θ, s.t. $\theta : \{x_1, \dots, x_n, y, z_1, \dots, z_m\} \mapsto \{x_1, \dots, x_n, y, z_1, \dots, z_m\}$. Given an atom $\alpha(w_1, \dots, w_k)$, where each w_i is a variable in $\{x_1, \dots, x_n, y, z_1, \dots, z_m\}$, extend θ to be a substitution function in the obvious way: $\alpha(w_1, \dots, w_k)\theta \equiv \alpha(w_1\theta, \dots, w_k\theta)$.

Now, for every atom α (possibly featuring the equality predicate) in ψ and substitution function $\theta \in \Theta$ we define $\alpha\theta$ to be a *ψ-atom*. Extending this notion, we define *ψ-literals* to be the set of *ψ-atoms* closed under negation.

In our procedure we will utilise model like structures called *weak sorts*. We define a *weak n-sort* to be a maximal consistent set of *ψ-literals* over the signature σ featuring (a subset of) variables in $\{x_1, \dots, x_n\}$. Similarly, define a *weak $(n+1)$-sort* and *weak $(n+m+1)$-sort* to be maximal consistent sets of *ψ-literals* featuring (a subset of) variables in $\{x_1, \dots, x_n, y\}$ and $\{x_1, \dots, x_n, y, z_1, \dots, z_m\}$ respectively. Notice that some *weak sorts* may feature the equality predicate. Thus, we would like to clarify that (in our definition) equality predicates are consulted when establishing consistency between literals over σ (e.g. the set of literals $\{P(x), x = y, \neg P(y)\}$ would be considered inconsistent).

Let $\|\psi\|$ be the number of (not necessarily distinct) symbols needed to write down ψ. The number of *ψ-atoms* is then at most $\|\psi\| \times |\Theta|$. Hence, the number of different *weak* n, $(n+1)$ and $(n+m+1)$-*sorts* is bounded by $2^{\|\psi\| \times |\Theta|}$. Our current naive approach considers all possible substitutions θ, thus setting $|\Theta| = (n+m+1)^{n+m+1}$. Since n is unbounded (and can grow to be as large as $\|\varphi\|$), we have that $|\Theta| \in O(2^{\|\varphi\|})$. In Section 3, however, we will show that fewer substitution functions are needed and provide a $O(\|\varphi\|^{m+1})$ bound on Θ, thus establishing a singly exponential bound on the number of *weak sorts*.

If ρ is an *weak sort* over σ, we take $\rho \restriction (w_1, \dots, w_k)$ to denote the unique *weak sort* obtained by deleting all literals featuring any of the variables in $vars(\varphi) \setminus$

$\{w_1, \ldots, w_k\}$, and $\rho[w/v]$ to be the same *weak sort* only with variables w renamed to v.

When provided with a model $\mathfrak{A}$ and a universe A, for each tuple $\bar{a} \in A^k$ we may find a unique *weak k-sort* ρ, s.t. $\mathfrak{A} \models \rho[\bar{a}]$. We call ρ *the weak sort realised by the tuple* $\bar{a}$.

We are now prepared to prove the following lemma:

Lemma 1. *Given a* $[\exists^*\forall\exists^m]_=$ *sentence* φ*, there exists a procedure that runs in deterministic time* $O(2^{\|\psi\| \times |\Theta|})$ *and accepts if and only if* φ *is satisfiable.*

Proof. We provide a procedure $\textsc{sat}[\exists^*\forall\exists^m]_=$ which, in short, computes all possible *weak* $(n+1)$*-sorts* (as well as *weak* n and $(n+m+1)$*-sorts*) and deletes the ones that cannot be extended to a *weak* $(n+m+1)$*-sort* that satisfies ψ.

1: **procedure** $\textsc{sat}[\exists^*\forall\exists^m]_=(\exists x_1 \cdots \exists x_n \forall y \exists z_1 \cdots \exists z_m \psi)$
2: let Π be the set of all possible *weak* n*-sorts*
3: let T be the set of all possible *weak* $(n+1)$*-sorts*
4: let P be the set of all possible *weak* $(n+m+1)$*-sorts*
5:
6: **repeat**
7: **for all** $\tau \in T$ **do**
8: **find** $\rho \in P$ s.t $\left\{\begin{array}{l} 1)\ \tau \equiv \rho \restriction (x_1 \ldots x_n, y) \\ 2)\ \rho \restriction (x_1 \ldots x_n, z_i)[z_i/y] \in T, \text{ for all } i \in [1, m] \\ 3)\ \models \rho \to \psi \end{array}\right\}$
9:
10: **if** no such ρ found **then**
11: **remove** τ from T
12: **until** nothing was removed from T
13:
14: **for all** $\pi \in \Pi$ **do**
15: **if** for all $i \in [1, n]$ there exists $\tau \in T$ s.t. $\models \pi \cup \{y = x_i\} \to \tau$ **then**
16: **accept**
17:
18: **reject**.

Notice that Π, T and P in lines 2-4 have cardinality bounded by $O(2^{\|\psi\| \times |\Theta|})$ (and thus take no more time than that to generate). Keeping this in mind it is easy to verify that the code snippets in lines 6-12 and 14-16 also take $O(2^{\|\psi\| \times |\Theta|})$ number of steps to execute.

Now suppose that $\mathfrak{A} \models \varphi$. Let $\bar{a} \in A^n$ be a tuple of (not necessarily distinct) witnesses for the leading existential quantifiers. Now compute T' - a set of *weak* $(n+1)$*-sorts* over σ, s.t. for each $b_i \in A$ the *weak* $(n+1)$*-sort* τ_i realised by $(\bar{a}, b_i)$ is in T'. Notice that we can extend each $\tau_i \in T'$ to a *weak* $(n+m+1)$*-sort* ρ, s.t. $\models \rho \to \psi$ by simply picking the (trailing) existential witnesses $\bar{c_i} \in A^m$ for b_i. Furthermore, for each $c_{i,j} \in \bar{c_i}$ we have that the *weak* $(n+1)$*-sort* realised by $(\bar{a}, c_{i,j})$ is in T'. These properties comply with our criterion on line 8, thus ensuring that all $T \supseteq T'$ satisfy the acceptance condition of $\textsc{sat}[\exists^*\forall\exists^m]_=$ on line 15 (simply pick π to be the *weak n-sort* realised by $\bar{a}$).

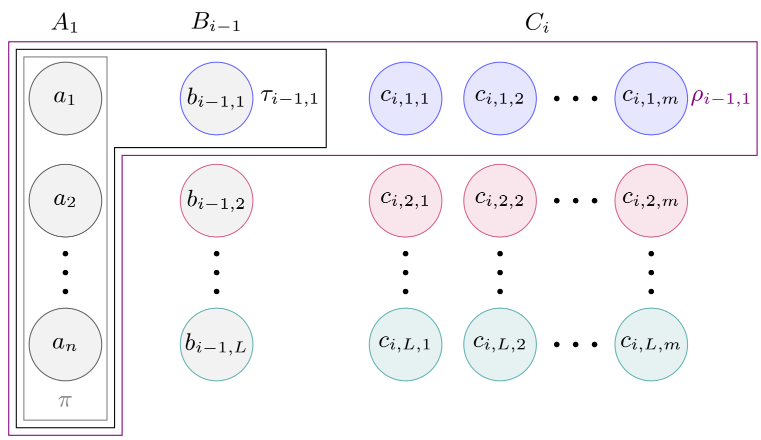

Fig. 1. Construction of the model $\mathfrak{A}_i$

Conversely, suppose $\textsc{sat}[\exists^*\forall\exists^m]_=$ accepts. Let us begin by defining a structure $\mathfrak{A}_0 = \emptyset$. We will now show by induction that we can build a further model $\mathfrak{A}_i$ that is an extension of $\mathfrak{A}_{i-1}$, and has the property we call *i-universality*: given that $B_i = A_i \setminus A_{i-1}$, each $b_{i,j} \in B_i$ has a *weak* $(n+1)$*-sort* $\tau_{i,j}$, realised by the witnesses of $x_1, \dots, x_n$ and $b_{i,j}$, s.t. $\tau_{i,j} \in T$, and *i-witnessing*: each $\tau_{i,j}$ can be extended to a *weak* $(n+m+1)$*-sort* $\rho_{i,j}$, s.t. $\models \rho_{i,j} \to \psi$.

Base case: utilising line 15 we create a model $\mathfrak{A}_1$ which is exactly the *weak n-sort* π our procedure accepted on. We take $\bar{a} \in A_1^n$ to be a vector of exactly those (not necessarily distinct) elements realising the leading existentially quantified variables. Notice that line 15 establishes 1*-universality* as for all $j \in [1, n]$ the *weak* $(n+1)$*-sort* $\tau_{1,j}$ realised by the elements $(\bar{a}, a_j)$ is in T. 1*-witnessing* is then immediate by conditions 1 and 3 of line 8.

Inductive Step: Set B_{i-1} to be the elements added at step $i-1$ and L to be the cardinality of B_{i-1}, i.e. $B_{i-1} = A_{i-1} \setminus A_{i-2}$ and $L = |B_{i-1}|$. For each element $b_{i-1,j} \in B_{i-1}$ $(1 \le j \le L)$ we obtain a *weak* $(n+1)$*-sort* $\tau_{i-1,j}$ that $(\bar{a}, b_{i-1,j})$ realise. $(i-1)$*-witnessing* (*I. H.*) allows us to extend each $\tau_{i-1,j}$ to a legal *weak* $(n+m+1)$*-sort* $\rho_{i-1,j}$. Recall that $\rho_{i-1,j}$ features variables in $\{x_1, \dots, x_n, y, z_1, \dots, z_m\}$; the variables $x_1, \dots, x_n$ and y are associated with elements $\bar{a}$ and $b_{i-1,j}$ respectively, however, variables $z_1, \dots, z_m$ are not assigned any element. Thus, for each z_k $(1 \le k \le m)$ we allot an element $c_{i,j,k}$ as follows:

1. If z_k was featured in an equality predicate $z_k = x_l$ $(1 \leq l \leq n)$ or $z_k = y$, set $c_{i,j,k}$ to be the element $a_l \in \bar{a}$ or $b_{i-1,j}$ respectively.
2. If z_k was featured in an equality predicate $z_k = z_l$, where $l < k$, set $c_{i,j,k}$ to be the same element as $c_{i,j,l}$.
3. Else, create a new element $c_{i,j,k}$.

Let C_i be the set of these newly associated elements, i.e. $C_i = \bigcup_{j=1}^{L} \bigcup_{k=1}^{m} \{c_{i,j,k}\}$. We set the universe $A_i = A_{i-1} \cup C_i$ and define $\mathfrak{A}_i$ to be an extension of $\mathfrak{A}_{i-1}$ with *weak* $(n+m+1)$*-sorts* $\rho_{i-1,j}$, s.t. $\mathfrak{A}_i \models \rho_{i-1,j}[\bar{a}, b_{i-1,j}, c_{i,j,1}, \ldots c_{i,j,m}]$. It is worth noting that in case 1 (and subsequently case 2) no interpretation is ever redefined as the *weak sort* $\rho_{i-1,j}$ is, by definition, consistent and retains just enough relationship descriptions of $\bar{a}$ and $b_{i-1,j}$. Notice that (by condition 2 of line 8) the *weak* $(n+1)$*-sorts* $\tau_{i,(j,k)}$ realised by $(\bar{a}, c_{i,j,k})$ are in T. Taking note that $B_i = C_i = A_i \setminus A_{i-1}$, we established *i-universality*. Additionally, we obtain *i-witnessing* from the mere fact that $\tau_{i,(j,k)} \in T$: if there was no *weak* $(n+m+1)$*-sort* $\rho_{i,(j,k)}$ extending $\tau_{i,(j,k)}$, s.t. $\models \rho_{i,(j,k)} \to \psi$, then $\tau_{i,(j,k)}$ would not be in T (as condition 3 of line 8 would not be met), thus, contradicting *i-universality* just established.

An illustration of the construction process is depicted in Figure 1. Following the previously described rules, we are left with an infinite model $\mathfrak{B} = \mathfrak{A}_\omega$ or a finite model $\mathfrak{B} = \mathfrak{A}_k$ if at some step k no new elements were created (i.e. $B_k = A_k \setminus A_{k-1} = \emptyset$). Either way, utilising *i-witnessing* it is then trivial to verify that $\mathfrak{B} \models \varphi$.

3 Small Number of Substitution Functions

Recall that we defined Θ to be the set of all possible maps θ, s.t. $\theta : \{x_1, \ldots, x_n, y, z_1, \ldots, z_m\} \mapsto \{x_1, \ldots, x_n, y, z_1, \ldots, z_m\}$. In this section we will show that it suffices to consider the set Θ^*, which is comprised of substitution functions of the form $\theta^* : \{y, z_1, \ldots, z_m\} \mapsto \{x_1, \ldots, x_n, y, z_1, \ldots, z_m\}$.

To motivate this transition, consider a sentence of the form $\varphi \equiv \exists x_1 \cdots \exists x_n \forall y \exists z_1 \cdots \exists z_m \psi$ and fix a model $\mathfrak{A}$. Notice that $\mathfrak{A} \models \varphi$ iff (i) for each $1 \leq i \leq n$, there exists a mapping $c_i : \emptyset \mapsto A$; (ii) for each $1 \leq j \leq m$, there exists mappings f_j, s.t. $f_j : A \mapsto A$; (iii) and $\mathfrak{A} \models \forall y\, \psi_{sk}$, where $\psi_{sk} \equiv \psi[x_1/c_1, \ldots, x_n/c_n, y/y, z_1/f_1(y), \ldots, z_m/f_m(y)]$. Notice that whilst y and $f_j(y)$ range over elements of A (depending on the mapping of y to A), each c_i is constantly mapped to the same exact element.

Now suppose that ψ_{sk} features an atom $\alpha(w_1, \ldots, w_l, \ldots, w_k)$, where each $w_j \in \{c_1, \ldots, c_n, y, f_1(y), \ldots, f_m(y)\}$, but w_l is some c_i. We then consider the following cases:

Case 1: $\alpha(w_1, \ldots, c_i, \ldots, w_k)$ is the only occurrence of α in ψ_{sk}. Let ρ_{sk} be a *weak sort* (with variables replaced with functions in the same manner), s.t. $\models \rho_{sk} \to \psi_{sk}$. Notice that by our original construction, ρ_{sk} features all of the atoms $\alpha(w_1, \ldots, w_l, \ldots, w_k)$, where $w_l \in \{c_1, \ldots, c_n, y, f_1(y), \ldots, f_m(y)\}$, whilst ψ_{sk} only contains $\alpha(w_1, \ldots, c_i, \ldots, w_k)$. Since c_i is constantly mapped to the same single element (regardless of y), we may safely remove the atoms $\alpha(w_1, \ldots, w_l, \ldots, w_k)$ from ψ_{sk}, where w_l is not c_i and retain $\models \rho_{sk} \to \psi_{sk}$.

Some $f_j(a)$ ($a \in A$ and $1 \leq j \leq m$) may still map to the same element as c_i. However, we then simulate $\alpha(w_1, \ldots, f_j(y), \ldots, w_k)$ in ρ_{sk} by having the atom $c_i = f_j(y)$ in ρ_{sk}. Alternatively, if $c_i = f_j(y)$ is not among ψ_{sk}-*atoms*, by creating duplicate elements of c_i so that $c_i \neq f_j(a)$. The same argument applies to a and y in place of $f_j(a)$ and $f_j(y)$ respectively.

Case 2: there is an atom $\alpha(w_1, \ldots, w_l, \ldots, w_k)$ in ψ_{sk}, where $w_l \in \{y, f_1(y), \ldots, f_m(y)\}$. Let Γ be the set of atoms $\alpha(w_1, \ldots, w_l, \ldots, w_k)$, where $w_l \in \{c_1, \ldots, c_n, y, f_1(y), \ldots, f_m(y)\}$. Since w_l might range over multiple elements of A, the computed *weak sorts* must contain all the (possibly negated) atoms in Γ in order to avoid multiple (inconsistent) definitions. We may then disregard all other atoms $\alpha(w_1, \ldots, w_l, \ldots, w_k)$ in ψ_{sk} as they will appear in Γ.

Case 3: there are multiple atoms $\alpha(w_1, \ldots, w_l, \ldots, w_k)$ in ψ_{sk}, where $w_l \in \{c_1, \ldots, c_n\}$. We simply take all of these atoms when creating *weak sorts* and apply a similar argument as in Case 1.

For the sake of consistency, we again consider our original sentence $\varphi \equiv \exists x_1 \cdots \exists x_n \forall y \exists z_1 \cdots \exists z_m \psi$. We conclude from the cases above that the set Θ^* of all possible substitution functions $\theta^* : \{y, z_1, \ldots, z_m\} \mapsto \{x_1, \ldots, x_n, y, z_1, \ldots, z_m\}$ is sufficient to generate ψ-*atoms* and, in turn, *weak sorts*. Notice that $|\Theta^*| = (n + m + 1)^{m+1} \in O(\varphi^{m+1})$, where m is fixed integer.

Taking that the procedure in Lemma 1 runs in time $O(2^{\|\psi\| \times |\Theta|})$ and setting $\Theta = \Theta^*$, we have effectively proven Theorem 1.

Corollary 1. *The satisfiability problem of the fragment* $[\exists^*\forall\exists^m]_=$ *for all fixed* m $(2 \leq m < *)$ *is* ExpTime*-complete.*

Proof. ExpTime-hardness for the satisfiability problem of $[\forall\exists^2]$ was established by M. Fürer [3], whilst membership in ExpTime is given in Theorem 1.

4 Conclusions

We have shown that the satisfiability problem for the quantifier prefix fragment $[\exists^*\forall\exists^m]_=$ is ExpTime-complete for all fixed m $(2 \leq m < *)$. And, combined with the results of R. Jaakkola [2], filled long left complexity gaps in the Ackermann fragment (with equality), namely, we proved that the minimal NExpTime-complete fragment is $[\exists^2\forall\exists^*]_=$. We finish the article by providing an account of the complexity of the satisfiability problem for various Ackermann (sub-)fragments (with equality):

Fragment	Constraint	Complexity of Satisfiability	Reference
$[\exists^\alpha\forall\exists^*]_=$	$2 \leq \alpha \leq *$	NExpTime-complete	[1, p. 285]
$[\exists^n\forall\exists^*]_=$	$n < 2$	ExpTime-complete	[2]
$[\exists^*\forall\exists^m]_=$	$2 \leq m < *$	ExpTime-complete	Corollary 1
$[\exists^*\forall\exists]_=$		PSpace-complete	[1, p. 283]

This account is complete in a sense that the fragment $[\exists^2\forall\exists^*]_=$ is NExpTime-hard, whilst fragments with restrictions on quantifiers such as $[\exists\forall\exists^*]_=$, $[\exists^*\forall\exists^m]_=$ $(2 \leq m < *)$ are in ExpTime. These bounds are indeed optimal as $[\forall\exists^2]$ is ExpTime-hard. Likewise, the upper bound of $[\exists^*\forall\exists]_=$ is optimal as $[\forall\exists]_=$ is PSpace-hard.

References

1. Börger, E., Grädel E., Gurevich, Y.: The Classical Decision Problem. Springer Science & Business Media (1997).
2. Jaakkola, R.: Complexity of the Ackermann fragment with one leading existential quantifier (2022).
3. Fürer, M.: Alternation and the Ackermann case of the decision problem. L'Enseignement Math. 27, 137–162 (1981).
4. Turing, A.M.: On Computable Numbers, with an Application to the Entscheidungsproblem. Proceedings of the London Mathematical Society. 2, 42, 230–265 (1936).