Weak Theories of Linear Algebra

Neil Thapen* and Michael Soltys†

July 27, 2005

Abstract

We investigate the theories $\mathrm{LA}, \mathrm{LAP}, \forall\mathrm{LAP}$ of linear algebra, which were originally defined to study the question of whether commutativity of matrix inverses has polysize Frege proofs. We give sentences separating quantified versions of these theories, and define a fragment $\exists\mathrm{LA}$ of $\forall\mathrm{LAP}$ in which we can interpret a weak theory V^1 of bounded arithmetic and carry out polynomial time reasoning about matrices - for example, we can formalize the Gaussian elimination algorithm. We show that, even if we restrict our language, $\exists\mathrm{LA}$ proves the commutativity of inverses.

1 Introduction

Linear algebra was first suggested by Bonet, Buss and Pitassi [2] as a source for hard tautologies separating the Frege and Extended Frege propositional proof systems. In this paper we are particularly interested in the implication $AB = I \supset BA = I$ of matrix algebra ("commutativity of inverses"), proposed by Cook. This is believed to be hard because all known proofs

*Department of Computer Science, University of Toronto, 10 King's College Road, Toronto, Ontario M5S 3G4, Canada; thapen@cs.toronto.edu

†Department of Computing and Software, McMaster University, 1280 Main Street West, Hamilton, Ontario L8S 4K1, Canada; soltys@mcmaster.ca

of it involve more than the simple manipulations of matrices in terms of their elements; they require, for example, going through an algorithm (such as Gaussian elimination or computing a determinant), or using a counting argument (about the size of a basis set).

In [7] Soltys studied this problem and gave a precise meaning to "simple" above. He defined a theory LA (for Linear Algebra) with three sorts - index elements, field elements and matrices. A matrix is treated as a table of field elements, coordinates in the table being given by index elements. It turns out that LA is strong enough to prove all the ring properties of the set of matrices (i.e., the associativity of matrix multiplication, commutativity of matrix addition, etc.). On the other hand, LA is weak enough that all the theorems of LA translate into families of boolean tautologies with polysize Frege proofs.

Soltys defined two extension of LA, to LAP which also has a symbol for matrix powering and to ∀LAP which additionally has induction for formulas with universal matrix quantifiers. LAP can formalize Berkowitz's algorithm [1] and prove the equivalence of many important universal principles of linear algebra. ∀LAP can prove the Cayley Hamilton theorem using Berkowitz's algorithm, and it seems that ∀LAP can prove all the universal principles of linear algebra (such as the commutativity of inverses and the multiplicativity of the determinant). See [4] for a survey of related results on the algorithmic versions of these problems.

In this paper we explore a little further the relative strengths of these theories and answer some of the questions posed in [7].

Section 2 gives the definitions of the theories mentioned above, together with new theories ∃LA and ∀LA with induction for formulas with existential or universal matrix quantifiers but without powering. If we have index subtraction in our language (see below) then these two theories are equivalent.

In section 3 we show that ∀LA proves the commutativity of inverses, and in section 4 we show that ∃LA is strong enough to interpret the theory V^1. We use the techniques of this section to show in section 5 that we can formalize in ∃LA complicated arguments about matrices, in particular the Gaussian elimination algorithm. In the last section we give sentences separating the theories LA, LAP and ∃LA.

We pay careful attention to the language in which we are working. In [7] LA and the theories extending it are defined with a very rich language for index elements, including addition, multiplication, quotient and remainder, and with a symbol for field inverse. However in [8] it was shown that ∀LAP proves the Cayley-Hamilton Theorem and the multiplicativity of determinant over commutative rings, without using inverse. Also, small adjustments to the arguments in [7] show that an index language with only successor, predecessor, ordering and a function $\max(i, j)$ is enough to develop the basic matrix properties in LA.

In this paper we try to work in this weaker version of LA, and in particular our matrix entries come in general from a ring with cancellation rather than a field, and we have changed the name of the sort accordingly. In section 2 we define our theories in the rich language, for completeness. The proofs in sections 3 and 6 work in the weak language. In section 4 we need the richer index language to prove that ∃LA interprets V^1, and in section 5 we need a term in the language for field inverse to do complicated reasoning about general matrices (rather than the binary matrices dealt with in section 4).

The restricted index language seems to be sufficient for most arguments in linear algebra, and weakening the definitions of our theories in this way should make the main conjecture in this area, that commutativity of inverses is not provable in LA, easier to approach. However the restriction takes us away from one of the motivations for the definition of LA, that it should tell us something about provability in propositional calculus. Since in propositional proofs we can group together arbitrary collections of variables, it is natural to allow LA to define complicated sets of index elements; the simpler we force our definable sets to be, the less close is the match with propositional proof systems and the less a proof that LA does not prove commutativity of inverses would tell us about a propositional lower bound.

2 Definitions

We begin with the description of the logical theory LA. See [7] and [8] for a complete description. The three sorts of LA are *indices*, *ring elements* (we weaken the assumption in [7] that these must be from a field) and *matrices*. Where it is important to distinguish between the symbols in the language associated with different sorts, we will use subscripts ${}_{i},{}_{r},{}_{m}$. We will typically denote index elements by i, j, k, ring elements by a, b, c, and matrix elements by capital letters.

We have three equality symbols, one for each sort, and a relation $\leq_i$ for the ordering on the indices.

We define the function symbols in (the rich version of) our language: $(m +_i n), (m *_i n), (m -_i n), \mathrm{div}(m, n), \mathrm{rem}(m, n)$ are function symbols of type index that take inputs m, n of the index sort. $(t +_r u), (t *_r u), (-_r t), (t^{-1})$ are function symbols of type ring and take inputs t, u of the ring sort. Note that ${}^{-1}$ is defined everywhere, even at 0.

If T is a term of type matrix, then $r(T), c(T)$ are terms of type index which respectively denote the number of rows and columns of T, and $\Sigma(T)$ is a term of type ring that denotes the sum of all the entries of T, and if m, n are terms of type index, then $e(T, m, n)$ is a term of type ring which denotes the (m, n) entry of the matrix T. We also have a way of constructing matrices using some rudimentary λ-calculus: if m, n are terms of type index, and t is a term of type ring, then $\lambda ij\langle m, n, t\rangle$ is a *constructed term of type matrix* (note that the index variables i, j cannot occur free in m, n) with the properties $r(\lambda ij\langle m, n, t\rangle) = m$, $c(\lambda ij\langle m, n, t\rangle) = n$, and $e(\lambda ij\langle m, n, t\rangle, i, j) = t$.

Example 2.1 *We can define the sum of two $n \times n$ matrices A, B as follows: $\lambda ij\langle n, n, e(A, i, j) + e(B, i, j)\rangle$, where $e(A, i, j) + e(B, i, j)$ is a ring term in the language of LA, expressing the (i, j)-th entry of the matrix $A + B$. Thus, when we write $A+B$ we really mean the matrix $\lambda ij\langle n, n, e(A, i, j)+e(B, i, j)\rangle$.*

Finally, if α is a formula where all the atomic subformulas are of type index (that is, are of the form $m =_i n$ or $m \leq_i n$), then $\mathrm{cond}_i(\alpha, m, n)$ and

$\mathrm{cond}_f(\alpha, t, u)$ are terms of type index and ring respectively, and the idea is that $\mathrm{cond}_i(\alpha, m, n)$ has the value m if α is true, and n otherwise, and similarly for cond_f. The restriction that all the atomic subformulas of α are of type index is there because in the translations into propositional formulas all the free index variables get values, and therefore α will becomes true or false.

We now specify the axioms of the theory LA. When working in a restricted language, we use the suitable fragment of our axioms.

Firstly, all the usual axioms for equality are in LA, for each of our three sorts.

For elements of the index sort we have the usual axioms of Robinson's arithmetic together with axioms defining div, rem, and an axiom scheme for cond_i (with an axiom for each α):

$$j \neq 0 \supset \mathrm{rem}(i, j) < j$$
$$j \neq 0 \supset i = j * \mathrm{div}(i, j) + \mathrm{rem}(i, j)$$
$$(\alpha \supset \mathrm{cond}(\alpha, i, j) = i) \wedge (\neg\alpha \supset \mathrm{cond}(\alpha, i, j) = j)$$

The axioms for ring elements are the usual axioms for an integral domain (and, if we are including it in the language, an axiom for field inverse) together with axioms for cond_r, similar to those for cond_i.

Finally, we have axioms for matrices: first of all, we want $e(A, i, j) = 0$ for all i, j which are out of bounds:

$$(i = 0 \vee r(A) < i \vee j = 0 \vee c(A) < j) \supset e(A, i, j) = 0.$$

Next, the definition for constructed matrices (with three axioms for each t):

$$r(\lambda ij\langle m, n, t\rangle) = m$$
$$c(\lambda ij\langle m, n, t\rangle) = n$$
$$1 \leq i \leq m \wedge 1 \leq j \leq n \supset e(i, j, \lambda ij\langle m, n, t\rangle) = t$$

Lastly we define how our term Σ behaves. There is nothing that prevents us from constructing a matrix with zero rows or zero columns, so we need the following axiom: $r(A) = 0 \vee c(A) = 0 \supset \Sigma A = 0$. There are four more axioms that give a recursive definition of Σ; we do not give them here, but the idea is simple:

$$\Sigma(A) = a_{11} + \Sigma(R) + \Sigma(S) + \Sigma(A[1|1])$$

where R, S are the first row and column of A (without the first entry), respectively, and $A[1|1]$ is the standard notation for the principal minor of A. The three submatrices $R, S, A[1|1]$ can be defined from A using the constructed matrices. See [7, Chapter 2.3] for all the details of this definition.

To make sure that our inductive definitions (in particular that of Σ) are well behaved, LA has the following induction axiom for open formulas ϕ possibly containing parameters (of any sort):

$$\phi(0) \wedge \forall i\, (\phi(i) \supset \phi(i+1)) \supset \forall i\, \phi(i).$$

We also have the obvious axiom for matrix equality:

$$c(A) = c(B) \wedge r(A) = r(B) \wedge \forall i, j\, e(A, i, j) = e(B, i, j) \supset A = B.$$

In [7] LA was introduced as a quantifier-free theory in the style of the sequent calculus, with special rules to deal with induction and matrix identity. Here we treat it as a first order theory with the normal logical rules and with induction and matrix identity axiom schemes. Standard soundness and completeness arguments show that the two are equivalent, in that they prove the same universal sentences.

We define the other theories we will use. LAP is LA together with a term $\mathrm{P}(i, X)$ of type matrix and axioms $\mathrm{P}(0, A) = I_{r(A)}$ and $\mathrm{P}(i+1, A) = \mathrm{P}(i, A)A$.

A bounded universal matrix quantifier has the form $\forall X \leq i \ldots$ meaning $\forall X\, (c(X) \leq i \wedge r(X) \leq i \supset \ldots)$. Bounded existential matrix quantifiers are defined dually. An $\forall$LA ($\exists$LA) formula consists of a block of bounded universal (existential) matrix quantifiers followed by a quantifier-free LA formula. The theory $\forall$LA is the extension of LA that also includes induction axioms for $\forall$LA formulas, and $\exists$LA is similar. $\forall$LAP is the analagous extension of LAP.

If we have index subtraction then $\forall$LA and $\exists$LA are equivalent by the standard argument, deriving induction on i in $\exists Y\, \phi(i, Y)$ from induction on j in $\forall Y\, \neg\phi(i-j, Y)$.

All the theorems of LA translate into families of boolean tautologies with polysize Frege proofs. Also, the theorems of LAP translate into quasi-poly-size Frege proofs (because the P function translates into NC^2-circuits, i.e., circuits of poly-size and $O(\log^2)$ depth), and the theorems of $\forall$LAP translate into poly-size Extended Frege proofs. See [7, Chapter 7] for the details on these translations (note that the underlying ring is one of the parameters of the translation).

For clarity of presentation we will write $X_{i,j}$ for $e(X, i, j)$ and W_j for $e(W, 1, j)$ if W is a row matrix. In the presence of the functions e and λ we can identify ring elements with 1×1 matrices, and for example quantify over ring elements wherever we are allowed to quantify over matrices.

3 Hard Matrix Identities

We show that $AB = I \supset BA = I$ can be proven in $\forall$LA, using an argument based on Gaussian elimination. We do not want to use $+, *, \mathrm{div}$ or rem in our index language, and without these we cannot talk about sequences of matrices, or even define an elementary matrix, so cannot carry out the text-book Gaussian elimination algorithm (although with them in the language, we can; see section 5). Instead we will do a simple induction, based on a computation step in Gaussian elimination.

Since all the hard matrix identities can be proven equivalent in LA (see [7] for details), it follows that all the hard matrix identities can be proven in $\forall$LA.

Theorem 3.1 *$\forall\mathrm{LA} \vdash \forall A \forall B,\ BA = I \supset BA = I.$*

Proof The proof is by induction on the size of the matrices involved. Our induction hypothesis is:

$$\forall A \leqslant n\, \forall B \leqslant n\, \forall a,\ (a \neq 0 \wedge AB = aI_n) \supset BA = aI_n$$

where a is a ring element. We include the factor a because this way we do not need to use multiplicative inverses in the proof.

The base case $n = 1$ follows from the axiom for the commutativity of ring multiplication.

Now suppose that the hypothesis holds for matrices of size n and let A, B be matrices of size $n+1$ with $AB = I_{n+1}$. We will construct matrices E_1, E_2 (whose inverses are easy to compute) such that the first column of E_2E_1A consists only of 0s except possibly for the first entry. This will allow us to apply the inductive hypothesis.

If the first column of A is all 0s already, then we are done and we can set $E_1 = E_2 = I_{n+1}$. Otherwise let k be a row with a nonzero entry in the first column. We take E_1 to be the permutation matrix that moves the kth row to the top and moves rows $1, \ldots, k-1$ down one place. This matrix exists, since it is given by the term

$$\lambda ij\langle n{+}1, n{+}1, \mathrm{cond}((i = 1 \wedge j = k) \vee (1 < i \leqslant k \wedge j = i{-}1) \vee (i > k \wedge i = j), 0, 1)\rangle.$$

Let E_1^* be the inverse of E_1; it also exists, since it is just the transpose of E_1.

Now let $c_1, \ldots, c_n$ be the entries in the first column of E_1A (so $c_1 \neq 0$) and let E_2 and E_2^* be as follows:

$$E_2 = \begin{pmatrix} 1 & 0 & 0 & \ldots & 0 \\ -c_2 & c_1 & 0 & \ldots & 0 \\ -c_3 & 0 & c_1 & \ldots & 0 \\ \vdots & 0 & 0 & \ddots & 0 \\ -c_{n+1} & 0 & 0 & \ldots & c_1 \end{pmatrix} \qquad E_2^* = \begin{pmatrix} c_1 & 0 & 0 & \ldots & 0 \\ c_2 & 1 & 0 & \ldots & 0 \\ c_3 & 0 & 1 & \ldots & 0 \\ \vdots & 0 & 0 & \ddots & 0 \\ c_{n+1} & 0 & 0 & \ldots & 1 \end{pmatrix}$$

so that E_2E_1A has c_1 in the top-left position and has its first column otherwise zero and $E_2E_2^* = E_2^*E_2 = c_1I_{n+1}$. Clearly E_2 and E_2^* can be given by terms. Let $E = E_2E_1$ and $E^* = E_1^*E_2^*$. Now since $AB = aI_{n+1}$, it follows that $E(AB)E^* = c_1aI_{n+1}$ so by associativity $(EA)(BE^*) = c_1aI_{n+1}$.

Now let $b = c_1$, or if the first column of A was originally all 0s, let $b = 1$. In either case, we now have non-zero elements b, a and matrices E, E^* such that EA is in the desired form, $(EA)(BE^*) = baI_{n+1}$ and $E^*E = bI_{n+1}$.

Now we need to examine the entries of EA and BE^*, so let

$$EA = \begin{pmatrix} x_1 & R_1 \\ 0 & M_1 \end{pmatrix} \qquad BE^* = \begin{pmatrix} x_2 & R_2 \\ S_2 & M_2 \end{pmatrix}$$

where x_1, x_2 are ring elements, M_1, M_2 are $n \times n$ matrices and R_1, R_2, S_2 are row and column matrices. Using this notation, we have that

$$bBA = (BE^*)(EA) = \begin{pmatrix} x_1x_2 & x_2R_1 + R_2M_1 \\ x_1S_2 & S_2R_1 + M_2M_1 \end{pmatrix}$$

We show that this matrix is baI_{n+1}. we know that $(EA)(BE^*) = baI_{n+1}$ so we have

1. $M_1M_2 = baI_n$,
2. $x_1x_2 + R_1S_2 = ba$,
3. $M_1S_2 = 0$,
4. $x_1R_2 + R_1M_2 = 0$.

From 1. and the inductive hypothesis $M_2M_1 = baI_n$. From 3. we know that $M_2M_1S_2 = 0$ hence $baS_2 = 0$ so S_2 is the zero vector by cancellation. Thus from 2. we know that $x_1x_2 = ba$. Multiplying 4. on the left by x_2 and on the right by M_1 we get $x_2x_1R_2M_1 + x_2R_1M_2M_1 = 0$ and so $baR_2M_1 + x_2R_1ba = 0$; hence by cancellation $x_2R_1 + R_2M_1 = 0$. Therefore $bBA = baI_{n+1}$ and with one more cancellation we are done. □

4 Recursion and bounded index quantifiers

We can use the axiom for matrix equality to write certain statements involving index quantifiers as quantifier free formulas in LA. For example, suppose that we have a formula $\phi(i)$ that can be represented by a term t, by which we mean that $t(i) = 0$ when $\phi(i)$ is false and $t(i) = 1$ when $\phi(i)$ is true. Then LA proves

$$(\forall 1 \leqslant i \leqslant n\, \phi(i)) \leftrightarrow \lambda ij\langle n, 1, t(i)\rangle = \lambda ij\langle n, 1, 1\rangle.$$

We can use this idea to express that a matrix satisfies a recursive property, then use $\exists$LA induction to show that such a matrix exists and thus to carry out the recursion. For example, suppose we have a term t taking ring elements to ring elements, and we want to iteratively apply this n times to a ring element a. Then any $1 \times n$ matrix X with the following property will encode the correct sequence of values:

$$X = \lambda ij\langle 1, n, \mathrm{cond}(j = 1, a, t(X_{j-1})\rangle.$$

In words the first entry in X is a and every other entry is t applied to the preceding entry. $\exists$LA-induction is sufficient to prove that for every a a unique such X exists.

The most useful example of this is the following: let $\mathrm{mult}(X, W)$ be the formula

$$W = \lambda ij\langle 1, c(X), \mathrm{cond}(j = 1, X_1, X_j W_{j-1})\rangle.$$

Then $\mathrm{mult}(X, W)$ holds if for each j, W_j is the product $X_1 \cdot \ldots \cdot X_j$. In particular the open induction available in LA is enough to prove that if X is a $0-1$ matrix then

$$\exists W \leqslant c(X)\, (\mathrm{mult}(X, W) \wedge W_{c(X)} = 1) \leftrightarrow \forall 1 \leqslant j \leqslant c(X)\, X_j = 1.$$

$\exists$LA-induction further proves that for every X some W exists with $\mathrm{mult}(X, W)$. Hence we can use an $\exists$LA formula to evaluate a single index quantifier.

We can extend this to deal with nested quantifiers. We will show this for one alternation, it naturally extends to any number of alternations.

Suppose $\phi(i,j,\bar{x})$ is represented by a term t_ϕ and we want to evaluate

$$\forall 1 \leqslant i \leqslant s(\bar{x})\, \exists 1 \leqslant j \leqslant s(\bar{x})\, \phi(i,j,\bar{x})$$

where s is an index term and $\bar{x}$ is a tuple of parameters of any type, which we will suppress in what follows.

Define the formula $\mathrm{Mult}(X,W,n)$ as:

$$W = \lambda ij\langle 1, c(Y), \mathrm{cond}(\mathrm{rem}(j,n)=1, X_j, X_j W_{j-1})\rangle.$$

This is similar to mult except that W, rather than multiplying together all the elements of X, multiplies each block of n elements of X separately. As with mult, $\exists$LA proves that if X is a $0-1$ matrix then for each n there exists W such that $\mathrm{Mult}(X,W,n)$ and for every i,

$$W_{n\cdot i} = 1 \leftrightarrow \forall (i-1)n < j \leqslant in\, X_j = 1.$$

So that we can deal with the existential quantifier as the dual of the universal quantifier, write $\bar{Y}$ for the matrix obtained by subtracting the elements of Y from 1 (ie. inverting Y). Now let $\Phi(X,W,X',W')$ be the conjunction of:

1. $X = \lambda kl\langle 1, s^2, t_\phi(\mathrm{div}(l-1,s)+1, \mathrm{rem}(l-1,s)+1)\rangle$;
2. $\mathrm{Mult}(\bar{X}, W, s)$;
3. $X' = \lambda kl\langle 1, s, 1 - W_{l\cdot s}\rangle$;
4. $\mathrm{mult}(X', W')$;
5. $W'_s = 1$.

Here 1. expresses that X is a $1 \times s^2$ matrix coding the truth values of $\phi(i,j)$ in the natural way, and 2. and 3. together mean that X' is a $1 \times s$ matrix coding the truth values of $\exists 1 \leqslant j \leqslant s\, \phi(i,j)$ in the same way. Finally 4. expresses that W'_s codes the truth value of $\forall 1 \leqslant i \leqslant s\, \exists 1 \leqslant j \leqslant s\, \phi(i,j)$ and 5. says that this value is 1.

$\exists$LA-induction proves that matrices W and W' always exist, so we have that $\exists$LA proves

$$\exists X\, , W, X', W' \Phi(X, W, X', W') \leftrightarrow \forall 1 \leqslant i \leqslant s\, \exists 1 \leqslant j \leqslant s\, \phi(i, j).$$

Extending this argument to formulas of any quantifier complexity, we have

Theorem 4.1 *If Ψ is any formula consisting of a sequence of bounded index quantifiers followed by a formula that can be expressed as a term, then Ψ is equivalent to an* $\exists$LA *formula, provably in* $\exists$LA.

We now have everything we need to show that $\exists$LA interprets V^1, which is a two-sorted version of Buss' theory S^1_2 [3]. See [5, 7] for a description of V^1. It is a variant of the second order theories V^i_j introduced in [3].

We give a translation from the language of V^1 into the language of LA. The two sorts of V^1 are numbers and strings, we translate these into index elements and binary row matrices respectively. The algebraic properties of the numbers in V^1 follow immediately from the properties of the index elements in LA (if we use the rich language for indices).

To show that the induction in V^1 holds, we will show how to translate the Σ^B_1 formulas of V^1 into $\exists$LA formulas. (The formulas Σ^B_0 are V^1 formulas with bounded index quantifiers, and no string quantifiers; the formulas Σ^B_1 are V^1 formulas of the form $\exists X_1 \leq t_1 \cdots \exists X_n \leq t_n \alpha$, where $\alpha \in \Sigma^B_0$.)

We translate quantifier free formulas ϕ in the language of V^1 into LA terms t_ϕ such that t_ϕ is 1 if ϕ is true and 0 if ϕ is false. If ϕ is an atomic formula in the language of indices, we translate it into $\mathrm{cond}(\phi, 1, 0)$. If ϕ is of the form "$i \in X$" we translate it as $e(X, 1, i)$. For boolean connectives, inductively we translate $\phi \wedge \psi$ as $t_\phi \cdot t_\psi$ and $\neg\phi$ as $1 - t_\phi$. Then by theorem 4.1 in $\exists$LA every Σ^B_0 formula is equivalent to an $\exists$LA formula. Hence every Σ^B_1 formula is equivalent to an $\exists$LA formula and so induction holds for such formulas.

Theorem 4.1 *If we are allowed to use the full language of arithmetic for indices, then V^1 is conservative over* ∃LA. □

V^1 is a variant of S_2^1, and can be thought of as formalizing "polynomial time reasoning". So provided that we work over the two-element field, it is not hard to give in V^1 a proof of correctness for the natural polynomial time algorithm for Gaussian elimination. This can be extended to other fields whose elements can straightforwardly be coded in binary. Since ∃LA can interpret V^1 by the previous arguments, it follows that ∃LA can prove the correctness of Gaussian elimination as well, over these fixed fields.

∃LA is field independent (so field elements do not have to be encoded), and obviously more suitable for algebraic reasoning than V^1. In the next section we show directly that if we use our rich language, then ∃LA formalizes and proves the Gaussian elimination algorithm correct.

We note that the interpretation in this section could have been obtained easily in ∀LAP, since matrix powering already gives us a way of multiplying the elements of a row matrix together (see [7]).

5 Gaussian elimination

Theorem 5.1 *If we use the full index language and include a term for field inverse, then every formula in the language of* LA *containing only bounded index quantifiers is equivalent, in* ∃LA*, to an* ∃LA *formula.*

Proof We show that any quantifier free LA formula can be expressed as a block of bounded index quantifiers followed by $t = 1$, for some term t of type ring whose value, for any arguments, is either 1 or 0.

First we rewrite each matrix equality $X = Y$ as

$$c(X) =_i c(Y) \wedge r(X) =_i r(Y) \wedge \forall 1 \leqslant i \leqslant c(X)\, \forall 1 \leqslant j \leqslant r(X)\, e(X,i,j) =_r e(Y,i,j).$$

Then we move the index quantifiers to the front. Field inverse gives us a way of representing equality between ring terms as a term: $1-(t-u)(t-u)^{-1} = 1$ if

$t = u$, and 0 otherwise. We can express equality or ordering of index elements as a term using cond, as in the last section. This covers the atomic formulas; we can express logical combinations of atomic formulas as a term in the usual way.

Thus any formula with only bounded index quantifiers can be translated into a form to which we can apply theorem 4.1. $\square$

We want to show that the textbook Gaussian elimination algorithm can be proven correct in $\exists$LA. Given a matrix A with n rows, the Gaussian elimination algorithm states that there exists a sequence of elementary matrices $E_1, E_2, \ldots, E_k$ such that the matrix $E_k \cdots E_2 E_1 A$ is in reduced row-echelon form.

Fix the $n \times m$ matrix A that we want to reduce. We need to be able to talk about sequences of matrices, but we can do that by concatenating a sequence of l many $n \times n$ matrices into a single $n \times ln$ matrix X. We can define $X(i) =_{def} \lambda kj\langle n, n, e(X, k, n \cdot (i-1) + j)\rangle$ to be the i-th matrix encoded in X, so $X = [X(1), X(2), \ldots, X(l)]$.

Let Y be the matrix of the same size as X such that

$$Y(1) = X(1) \wedge (\forall i < l)Y(i+1) = Y(i)X(i)$$

so $Y(n)$ is equal to the product of the matrices encoded by X. It can be shown in $\exists$LA that such a Y exists. Thus we can define a function Π in $\exists$LA which computes the product of all the matrices encoded by a given matrix X.

A matrix C is in reduced row-echelon form if all the empty rows are at the bottom, the first non-zero entry of any row is a 1, this leading one is the only non-zero entry in its column, and if the position of this leading 1 is (i, j), then all the elements in positions (p, q), where $p > i$ and $q \leq j$ are zero.

We say that a matrix is in *i-partial reduced row-echelon form* if its first i columns are in reduced row-echelon form. We let the formula $PRRE(C, i)$ assert that the matrix C is in i-partial reduced row-echelon form. Note that this is expressible using bounded index quantifiers, and thus is expressible as an ∃LA formula. Similarly there is an ∃LA formula expressing that a matrix is elementary and an ∃LA formula $EL(X)$ expressing that every matrix encoded by X is elementary.

We prove by induction on i that for all $0 \le i \le m$,

$$\exists X \leqslant i(n+1),\ EL(X) \wedge PRRE((\Pi X)A, i).$$

The base case $i = 0$ is trivial. For the inductive step, suppose we have a witness X for i and want to find a witness X' for $i+1$.

Let j be the first empty row in the first i columns of $(\Pi X)A$. If the entries j to n of the $i+1$st column of $(\Pi X)A$ are all 0, then we do not need to do anything and can set $X' = X$. Otherwise let k be a row in this range with a non-zero entry in this column. We construct a sequence of $\leqslant n+1$ elementary matrices that move this entry to the jth row, normalize it to 1, and subtract suitable multiples of it from the other rows to make their entries 0 (as in the proof of theorem 3.1). We append these matrices to X, and use induction to show that each one has the desired effect as we add it.

At the end we have X with $PRRE((\Pi X)A, m)$, so $(\Pi X)A$ is in reduced row-echelon form and we are done.

6 Separations

We show that LAP is not conservative over LA and ∃LA is not conservative over LAP. However the ideas used in this section are not capable of showing that these theories do not all prove the same quantifier free sentences (in particular the commutativity of inverses).

These separations do not make use of the symbol for field inverse or of anything in the index language except for ordering, successor and predecessor. However the arguments still work in the richer language.

Let $\mathbb{F}$ be the field obtained from the two element field by adjoining all the nth roots of unity for all $n \in \mathbb{N}$. Let $\mathbb{G}$ be the set of all (standard) matrices over $\mathbb{F}$. The three sorted structure $(\mathbb{N}, \mathbb{F}, \mathbb{G})$ can be considered as a structure in the language of LAP, giving the symbols of LAP their standard interpretation.

Let $A = (I, F, G)$ be an elementary extension of $(\mathbb{N}, \mathbb{F}, \mathbb{G})$ containing nonstandard index elements. We will make some models of LA by taking the closure in A of I together with a subset of F under terms in the language.

We characterize carefully the set of terms we want to close under.

Definition 6.1 *Ring terms are functions of the form $I^m \times F^n \to F$ for $m, n \in \mathbb{N}$. We think of the ring tuple from F^n as a set of parameters; we will be interested in the behaviour of a term when we fix the ring parameters and let the index arguments range over all index elements.*

The set T of ring terms is defined inductively as follows, where we use $\bar{i}$ to stand for index elements and x for a ring element:

1. *The identity function $x \mapsto x$ on ring elements is in T;*
2. *If $t_1, t_2 \in T$, then $t_1 + t_2, t_1 \cdot t_2, -t_1 \in T$; $t_1^{-1} \in T$ if we are including $^{-1}$ in our language;*
3. *If $t_1(\bar{i}), t_2(\bar{i}) \in T$ and $\phi(\bar{i})$ is any formula whose atomic formulas are of type index, then the function $\mathrm{cond}(\phi(\bar{i}), t_1(\bar{i}), t_2(\bar{i}))$ is in T;*
4. *If $t(i_1, \ldots, i_k) \in T$ and $f(i_{k+1}, \ldots, i_{k+l})$ is any term of type index then $t(i_1, \ldots, i_{k-1}, f(i_{k+1}, \ldots, i_{k+l}))$ is in T;*
5. *If $t(i_1, \ldots, i_k) \in T$ then the function*
$$i_1, \ldots, i_k \mapsto \sum_{j=1}^{i_1} t(j, i_2, \ldots, i_k)$$
is in T. In the language of LA *we write this function formally as $\Sigma(\lambda ij\langle 1, i_1, t(j, i_2, \ldots, i_k)\rangle)$.*

We also close T under padding out a function with extra arguments and under permuting the arguments.

Lemma 6.2 *Given a subset $X \subseteq F$ we can build a model B of* LA *as follows. The index part of B is the same as the index part I of A. The ring part F_B of B contains all elements of F of the form $t(\bar{i}, \bar{x})$ for t a ring term, $\bar{i} \subseteq I$ and $\bar{x} \subseteq X$. The matrix part G_B of B contains all matrices in G of the form $\lambda j_1 j_2 \langle k_1, k_2, t(j_1, j_2, \bar{i}, \bar{x})\rangle$ for t a ring term, $\bar{i}, k_1, k_2 \subseteq I$ and $\bar{x} \subseteq X$.*

The functions and relations of B are the ones induced by A.

Proof Firstly, by the definition of ring terms the elements of B are closed under all terms in the language of LA. Hence B is a substructure of A, so all the universal axioms of LA are true in B because they were true in A.

To show that the matrix equality axioms hold in B, suppose there are matrices X and Y in B with $B \models X \neq Y$. Then X and Y must differ in A at some coordinates (i, j), and since the index parts of A and B are the same X and Y differ at (i, j) in B.

It remains to show that open induction holds in B. So let $\phi(i)$ be any quantifier-free formula with one free index variable i and parameters from B. Let $j \in I$.

Suppose that $B \models \phi(0)$ and that $B \models \phi(i) \supset \phi(i+1)$ for all $i \in I$. Then this is also true if we replace B with A because B is a substructure of A and ϕ is quantifier-free. Hence by open induction in A we know that $A \models \phi(j)$. Hence this must also be true in B. $\square$

Now let B be the structure given by this lemma if we take X to be all of F. Note that if $X = F$, then B is the same as A in the index and ring part, but G_B is a proper subset of G, i.e., A has more matrices.

Lemma 6.3 *Every matrix in B contains only a finite number of different ring elements.*

Proof We will show that for every ring term t there is a constant $s \in \mathbb{N}$ such that if we fix the ring arguments and let the index arguments range over all of I, the range of t contains at most s many different ring elements.

This is proved by induction on the complexity of t, and the only case that raises any difficulties is case 5 of the definition of a ring term. So suppose $t(\bar{i})$ is a ring term, that we have fixed the ring parameters and that only the elements $a_1, \ldots, a_s$ can appear in the range of t as we vary $\bar{i}$. Because we are in a field of characteristic 2, the sum of any combination of these elements can be reduced to the form $c_1a_1 + \ldots + c_sa_s$ where each c_l is either 0 or 1. Hence there are at most 2^s different values that such a sum can take. So if we let u be the ring term

$$i_1, \ldots, i_k \mapsto \sum_{j=1}^{i_1} t(j, i_2, \ldots, i_k)$$

then the range of u (once we have fixed the ring arguments) has size at most 2^s. □

Lemma 6.4 *There is an* $\exists$LA *formula* $\Phi(x, n)$ *expressing that the powers* $x, x^2, \ldots, x^n$ *are distinct.*

Proof Let $\mathrm{Pow}(x, n, W)$ be the formula

$$W = \lambda ij\langle 1, n, \mathrm{cond}(j = 1, x, xW_{j-1})$$

expressing that W encodes the first n powers of x. Let $\chi(X)$ be the formula:

$$\exists Y,\ \lambda ij\langle r(X), c(X), X_{i,j}Y_{i,j}\rangle = \lambda ij\langle r(X), c(X), 1\rangle.$$

expressing that some Y codes the inverse of every element of X, so in any model of the true theory of $(\mathbb{N}, \mathbb{F}, \mathbb{G})$ it will be true if and only if none of the elements of X is zero.

Now take $\Phi(x, n)$ to be the formula

$$\exists W,\ \mathrm{Pow}(x, n, W) \wedge \chi(\lambda ij\langle n, n, \mathrm{cond}(i = j, 1, W_i - W_j)\rangle).$$

The matrix defined by the λ term here has its (i, j)th entry zero (for $i \neq j$) if and only if the ith and jth powers of x are the same. □

Lemma 6.5 *The sentence "for all x and all n there is a row matrix W consisting of the first n powers of x" is false in B.*

Proof Let n be a nonstandard element of I. The sentence $\forall i\, \exists x\, \Phi(x, i)$ is true in $(\mathbb{N}, \mathbb{F}, \mathbb{G})$ because given i we can take x to be a root of unity so that $x, x^2, \ldots, x^i$ are all distinct: let $j = i$ if i is odd, and $j = i+1$ if i is even, and let x be the jth primitive root of unity. Then x is a root of the polynomial $x^j - 1$, and in the field of characteristic two the formal derivative of $x^j - 1$ is just x^{j-1} (since j is odd). Since $x^j - 1$ and x^{j-1} are relatively prime (by the Euclidean Algorithm), it follows that all the roots of $x^j - 1$ are distinct (see [6]). This is also true in A, so we can find an x in A such that the first n powers of x are distinct.

Now A contains a unique matrix W consisting of the sequence of these powers, but by construction this matrix is not in B since it contains infinitely many different ring elements. $\square$

Corollary 6.6 LAP *is not* $\forall\exists$LA *conservative over* LA. $\square$

For the other separation result, to build a model of LAP that is not a model of $\exists$LA we need to add a case to the inductive definition of a ring term to handle the addition of matrix powering to the language:

6. For $k \geqslant 2$, if $t(i_1, \ldots, i_k)$ is a ring term then so is

$$i_1, \ldots, i_k, i_{k+1}, i_{k+2} \mapsto e(P(i_{k+1}, \lambda j_1 j_2 \langle i_{k+2}, i_{k+2}, t(j_1, j_2, i_3, \ldots, i_k)\rangle), i_1, i_2).$$

Lemma 6.7 *If $X \subseteq F$ we can form a model C of* LAP *by taking the closure of X and I in A under all ring terms (including those defined using the new case 6), just as in lemma 6.2.* $\square$

Lemma 6.8 *Every ring element of F that is in the range of a ring term with parameters from our original field $\mathbb{F}$ is itself in $\mathbb{F}$.*

Proof Suppose our parameters are $b_1, \dots, b_m \in \mathbb{F}$. There is a finite subfield E of $\mathbb{F}$ containing $b_1, \dots, b_m$. If t is a ring term with these parameters, then it is true in $(\mathbb{N}, \mathbb{F}, \mathbb{G})$ that everything in the range of t is still in E, because in the standard model none of the operations described by ring terms can take us out of the field E. Because E is finite this property of t (with these parameters) is expressible by a first order formula, hence it is also true in A. $\square$

Theorem 6.9 LAP $\not\vdash$ $\exists$LA.

Proof Let $C = (I, F_C, G_C)$ be the substructure of A given by lemma 6.7, taking X to be all of the field $\mathbb{F}$. Then $C \models$ LAP, and by lemma 6.8 $F_C = \mathbb{F}$.

Let i be any element of I. Then if $i \in \mathbb{N}$ there is an element x of $\mathbb{F}$ whose powers $x, x^2, \dots, x^i$ are all distinct (namely an ith or $i+1$st primitive root of unity). If $i \notin \mathbb{N}$ then there is no such element x, because every element of $\mathbb{F}$ has finite order. Hence $C \models \exists x\, \Phi(x, 0)$ and $C \models \forall i, \exists x\, \Phi(x, i) \supset \exists x\, \Phi(x, i+1)$, but it is not the case that $C \models \forall i\, \exists x\, \Phi(x, i)$. So C is not a model of $\exists$LA. $\square$

Acknowledgements: The authors would like to thank Steve Cook for the very helpful conversations that led to this work.

References

[1] Stuart J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Information Processing Letters*, 18(3):147–150, 1984.

[2] Maria L. Bonet, Samuel R. Buss, and Toniann Pitassi. Are there hard examples for frege systems? *Feasible Mathematics*, II:30–56, 1994.

[3] Samuel R. Buss. *Bounded Arithmetic*. Bibliopolis, Naples, 1986.

[4] Stephen A. Cook. A taxonomy of problems with fast parallel algorithms. *Information and Computation*, 64(13):2–22, 1985.

[5] Stephen A. Cook. Proof complexity and bounded arithmetic. Notes for CSC2429S, "Proof Complexity and Bounded Arithmetic", given at the Fields Institute (available on line at `www.cs.toronto.edu/~sacook`), 1998.

[6] David S. Dummit and Richard M. Foote. *Abstract Algebra*. Prentice Hall, 1991.

[7] Michael Soltys. *The Complexity of Derivations of Matrix Identities*. PhD thesis, University of Toronto, 2001. Available from the ECCC server; see `www.eccc.uni-trier.de/eccc-local/ECCC-Theses/soltys.html`.

[8] Michael Soltys and Stephen A. Cook. The proof complexity of linear algebra. In *Seventeenth Annual IEEE Symposium on Logic in Computer Science (LICS 2002)*, 2002.