Proving properties of matrices over $\mathbb{Z}_2$

Michael Soltys

March 14, 2012

Abstract

We prove assorted properties of matrices over $\mathbb{Z}_2$, and outline the complexity of the concepts required to prove these properties. The goal of this line of research is to establish the proof complexity of matrix algebra. It also presents a different approach to linear algebra: one that is formal, consisting in algebraic manipulations according to the axioms of a ring, rather than the traditional semantic approach via linear transformations.

Keywords: Proof complexity, matrix identities, Frege and extended Frege.

1 Introduction

We are interested in the proof complexity of matrix algebra over the field of two elements GF(2). In particular, we examine the identity $AB = I \rightarrow BA = I$ which has been proposed as a candidate for separating the Frege and extended Frege propositional proof systems. We investigate the properties of matrices that can be proven over the simplest of fields, with the hope that understanding these properties will yield a low-complexity proof of $AB = I \rightarrow BA = I$.

All matrices are considered to be over the field of two elements $\{0, 1\}$; in the literature this field is denoted as $\mathbb{Z}_2$ or as GF(2). Let $M_{n\times m}(\mathbb{F})$ be the set of $n \times m$ matrices over a field $\mathbb{F}$; let $M(n) := M_{n\times n}(\mathbb{Z}_2)$, i.e., $M(n)$ is the set of all square $n \times n$ matrices over $\mathbb{Z}_2$. If the size of a matrix A is not specified, we assume $A \in M(n)$. We use A_{ij} to denote entry (i, j) of the matrix A.

We let A^t denote the *transpose* of matrix A, i.e., the matrix whose (i, j) entry is entry (j, i) of A. Let $I_n, 0_n$ be the *identity* matrix and *zero* matrix, respectively, in $M(n)$.

Given a matrix $A \in M(n)$, we often find it useful to represent A it terms of its *principal minor*, denoted M_A, as follows:

$$\left[\begin{array}{cc} a & R_A \\ S_A & M_A \end{array}\right],$$

where a is the top-left entry of A, i.e., $a = A_{11}$, and

$$\begin{aligned} R_A &= \left[\begin{array}{cccc} A_{12} & A_{13} & \dots & A_{1n} \end{array}\right], \\ S_A &= \left[\begin{array}{cccc} A_{21} & A_{31} & \dots & A_{n1} \end{array}\right]^t. \end{aligned} \tag{1}$$

The results in this paper can be interpreted in various extensions of the theory **LA**, for example **LAP** or $\exists$**LA**, where **LAP**. The theory **LA** is capable of rudimentary ring reasoning about matrices; the theory **LAP** adds matrix powering, $P(A, i) = A^i$, and $\exists$**LA** permits induction on Σ_1^B-formulas, which are formulas with existential quantification over matrices of bounded size. Over $\mathbb{Z}_2$, **LA** corresponds to $\mathbf{AC}^0(2)$, **LAP** to $\mathbf{NC}^2$ (in fact, slightly weaker), and $\exists$**LA** corresponds to polynomial-time reasoning. See section 13, the appendix, and [SC04] for more details.

Alternatively, we can employ theories corresponding to the complexity class $\oplus\mathbf{L} = \mathbf{AC}^0(\det_2)$, as defined in [CF10] where the theory is called $\mathbf{V}\oplus\mathbf{L}$. The advantage of these theories over the **LA**-family of theories is that **LA** is field independent, and it is not clear that **LA** can put to use the fact that $\mathbb{Z}_2$ is a particularly simple field.

2 Matrix identities

A motivation for this paper is to understand the complexity of concepts required to prove:

$$AB = I \rightarrow BA = I, \qquad \text{('Inverse Identity')}$$

and related matrix identities (see [SC04]), and to understand the proof complexity of combinatorial matrix algebra in general. That is, we would like to know what is the complexity of the concepts required to reason about basic linear algebra, and about the combinatorial applications of matrix algebra—as presented, for example, in [BR91].

There are two main motivations for this line of research; first, in reverse mathematics we are interested in the weakest logical theories capable of formalizing linear algebra. But mainly, Cook proposed the 'Inverse Identity' as a candidate for separating the Frege and extended Frege propositional proof systems—this is one of the principal open questions in theoretical computer science.

Lemma 1 *If $AB = I$ and we can show that A has* some *left-inverse, then $BA = I$; furthermore, this can be shown in* **LA**. *In other words,*

$$\mathbf{LA} \vdash (AB = I \wedge \exists C(CA = I)) \rightarrow BA = I.$$

Proof: $AB = I$ implies $ABA = A$, so $A(BA - I) = 0$; so if A has *some* left inverse, call it C, then $CA(BA - I) = 0$, so $BA = I$. □

3 Powers and products

Over the field $\mathbb{Z}_2$, $c + c = 0$, and hence, for any $A \in M(n)$, we have that $A + A = 0_n$. Lemmas 2 and 3 come from [Cob58]. The lemmas in this section can be shown in **LAP** augmented by the index function $y = 2^x$, which we define as $\exp(0) = 1$ and $\exp(i+1) = 2 \cdot \exp(i)$. This permits repeated squaring of a matrix, polynomially many times of course.

Lemma 2 $(I + A)^{2^i} = I + A^{2^i}$.

Proof: $(I + A)^2 = (I + A)(I + A) = I + A + A + A^2 = I + A^2$. □

Lemma 3 $(I + A)^{2^i} = I \iff A^{2^i} = 0$.

Proof: $A^{2^i} = 0 \iff I + A^{2^i} = I \iff (I + A)^{2^i} = I$, where the last equality follows by lemma 2. □

Lemma 4 *If* $(I + A)^{2^i - 1} = I$ *and* $AB = I$*, then* $A^{2^i - 1} = I$.

Proof: Suppose that $(I + A)^{2^i - 1} = I$; multiply both sides by $(I + A)$ to obtain $(I + A)^{2^i} = (I + A)$ and using lemma 2 we have $I + A^{2^i} = I + A$, and so $A^{2^i} = A \Rightarrow A^{2^i} B = AB \Rightarrow A^{2^i - 1} = I$. □

Lemma 5 $AB = BA \iff (I + A)(I + B) = (I + B)(I + A)$.

Lemma 6 *If* A, B *are inverses, i.e.,* $AB = BA = I$*, then* $(A+B)^{2^i} = A^{2^i} + B^{2^i}$.

Proof: First note that $(I+A)(I+B) = I+A+B+AB = I+A+B+I = A+B$; similarly $(I + B)(I + A) = I + B + A + BA = B + A$, and in particular $(I + A)(I + B) = (I + B)(I + A)$. Therefore:

$$\begin{aligned}
(A + B)^{2^i} &= (I + A)^{2^i} (I + B)^{2^i} \\
&\overset{(*)}{=} (I + A^{2^i})(I + B^{2^i}) \\
&= I + A^{2^i} + B^{2^i} + A^{2^i} B^{2^i} \\
&= I + A^{2^i} + B^{2^i} + I \\
&= A^{2^i} + B^{2^i},
\end{aligned}$$

where the $(*)$ equality follows from lemma 2. □

4 Idempotence, nilpotence and zero-divisors

A matrix A is *idempotent* if $A^2 = A$, it is *nilpotent* if $A^i = 0$ for some $i > 0$ and it is a *right-zero-divisor* (respectively, *left-zero-divisor*) if there exists a matrix $C \neq 0$ such that $AC = 0$ (respectively, $CA = 0$). If A is a right-zero-divisor then it is also a left-zero-divisor, but the C might differ; for example, if

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}, \quad D = \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix},$$

then $AC = 0$, but $CA \neq 0$ while $DA = 0$.

If $AB = I$ then **LA** proves that BA is idempotent, and also that neither A can be a left-zero-divisor nor B can be right-zero-divisors; for example, if $AB = I$ and $BC = 0$, then $ABC = 0$, so $C = 0$.

By lemma 2, if A is idempotent, so is $(I + A)$.

If $AB = I$, then **LAP** shows that neither A nor B can be nilpotent. Suppose $A^i = 0$. Then $A^i B^i = 0$ but $A^i B^i = I$.

5 Symmetric matrices

Let A by a *symmetric* matrix, i.e., $A = A^t$, that is A equals its *transpose*. We say that x, y are A-*orthogonal* if $x^t Ay = 0$; we sometimes write this as $\langle x, y\rangle_A = x^t Ay$.

Lemma 7 *If A is a symmetric matrix, then $V = \{x \in \mathbb{Z}_2^n : \langle x, x\rangle_A = 0\}$ is a vector space.*

Proof: If A is symmetric then:

$$\langle x, y\rangle_A = x^t Ay = (x^t Ay)^t = y^t A^t (x^t)^t = y^t Ax = \langle y, x\rangle_A.$$

Then, if $x, y \in V$ then:

$$\begin{aligned}\langle (x + y), (x + y)\rangle_A &= (x + y)^t A(x + y)\\ &= x^t Ax + y^t Ax + x^t Ay + y^t Ay\\ &= \langle x, x\rangle_A + \langle y, x\rangle_A + \langle x, y\rangle_A + \langle y, y\rangle_A,\end{aligned}$$

and $\langle x, x\rangle_A = \langle y, y\rangle_A = 0$ since we assumed that $x, y \in V$, and

$$\langle y, x\rangle_A + \langle x, y\rangle_A = 0,$$

as well since $\langle y, x\rangle_A = \langle x, y\rangle_A$. □

Lemma 8 **LA** *proves that if A is symmetric, and $AB = I$, then $BA = I$, and furthermore B is also symmetric.*

Proof: $AB = I$ then $I = I^t = (AB)^t = B^t A^t = B^t A$, as A is symmetric. By lemma 1 we have $BA = I$. On the other hand, from $B^t A = I$ we have $B^t AB = B$ and so $B^t = B$, and hence B is also symmetric. □

Let $d(A)$ denote the diagonal of a symmetric matrix A, i.e.,

$$d(A) = [A_{11} A_{22} \ldots A_{nn}].$$

In an unpublished note, Filmus ([Fil10]) make an interesting observation, which we present as lemma 9. We give a simplified proof, however, in the style of Gaussian Elimination.

Lemma 9 $\exists$**LA** *proves that for all symmetric A, $\exists v$ such that $Av = d(A)$.*

Proof: The proof is by induction on n, $A \in M(n)$. For $n = 1$ let $v = 1$. For $n > 1$ let

$$A = \left[\begin{array}{cc} a & X \\ X^t & M \end{array} \right],$$

and we consider the following cases: if $X = 0$ then by IH $\exists v_M$ such that $Mv_M = d(M)$, where M is the principal submatrix of A (and M is also symmetric if A is symmetric). So let $v = [1 \quad v_M]$, and then $Av = d(A)$. If $X \neq 0$ and $a = 1$ then let

$$C = \left[\begin{array}{cc} 1 & 0 \\ X^t & I_{n-1} \end{array} \right],$$

and observe that

$$\begin{aligned} A' &= CAC^t \\ &= \left[\begin{array}{cc} 1 & 0 \\ X^t & I_{n-1} \end{array} \right] \left[\begin{array}{cc} a & X \\ X^t & M \end{array} \right] \left[\begin{array}{cc} 1 & X \\ 0 & I_{n-1} \end{array} \right] \\ &= \left[\begin{array}{cc} 1 & 0 \\ X^t & I_{n-1} \end{array} \right] \left[\begin{array}{cc} a & aX + X \\ X^t & X^tX + M \end{array} \right] \\ &= \left[\begin{array}{cc} a & aX + X \\ aX^t + X^t & aX^tX + X^tX + X^tX + M \end{array} \right] \\ &= \left[\begin{array}{cc} a & aX + X \\ aX^t + X^t & aX^tX + M \end{array} \right] \\ &= \left[\begin{array}{cc} 1 & 0 \\ 0 & X^tX + M \end{array} \right] \end{aligned}$$

where we remind the reader that over $\mathbb{Z}_2$, $x + x = 0$, and so $X^tX + X^tX = 0$ and since $a = 1$, $aX + X = X + X = 0$. Using the IH we know that $\exists w$ such that $A'w = d(A')$, i.e., $CAC^tw = d(CAC^t)$, and since $CC = I_n$ (over $\mathbb{Z}_2$), we have

$$\begin{aligned} A(C^tw) &= Cd(CAC^t) \\ &= C \left[\begin{array}{c} 1 \\ d(X^tX + M) \end{array} \right] \\ &= C \left[\begin{array}{c} 1 \\ X^t + d(M) \end{array} \right] \end{aligned}$$

were $d(X^tX+M) = X^t+d(M)$ follows (over $\mathbb{Z}_2$) from the fact that the diagonal of X^tX equals $[\; x_1x_1 \quad x_2x_2 \quad \ldots \quad x_{n-1}x_{n-1} \;]$ which in turn is just X since $x_ix_i = x_i$,

$$\begin{aligned} &= \left[\begin{array}{c} 1 \\ X^t + X^t + d(M) \end{array} \right] \\ &= \left[\begin{array}{c} 1 \\ d(M) \end{array} \right]. \end{aligned}$$

Since $a = 1$, $A(C^t w) = d(A)$ and so letting $v = C^t w$ we are done showing that in the case $X \neq 0$ and $a = 1$, $\exists v$ such that $Av = d(A)$.

If, on the other hand, $X \neq 0$ and $a = 0$, there are two possibilities. First, $d(M) = 0$, in which case $v = 0$ and $Av = 0 = d(A)$, or there is some diagonal entry of M, say M_{ii}, which is not equal to zero; let P_i be the identity matrix with row 1 and row i permuted. Then, we can repeat the argument for the second case with $A' = (CP_i)A(CP_i)^t$, since $A' = C(P_i A P_i^t)C^t$ which has the effect of bringing $M_{ii} \neq 0$ (and hence $M_{ii} = 1$) to the position $(1, 1)$ of A. □

Symmetric matrices over finite fields have been considered in [Mac69], where, in section I, the author shows the following interesting results—originally due to A. A. Albert, and can be found in [Alb38]:

Theorem 1 *If $A \in M(n)$ is an invertible symmetric matrix of* $\mathrm{GF}(2^m)$ *then A can be factored in the form $A = M^t M$ if and only if $d(A) \neq 0$, i.e., some entry on the diagonal of A is not zero.*

The proof of theorem 1, as presented in [Mac69], can be formalized in $\exists$**LA**.

6 Trace

The proof of lemma 10 can be formalized in **LAP** with exp (as defined in section 3), while lemma 11 can be (obviously) formalized in **LA**.

Lemma 10 $\mathrm{tr}(A) = \mathrm{tr}(A^{2^i})$ *for all i.*

Proof: Note that

$$\begin{bmatrix} a & R_A \\ S_A & M_A \end{bmatrix}^2 = \begin{bmatrix} a^2 + R_A S_A & aR_A + M_A R_A \\ aS_A + M_A S_A & S_A R_A + M_A^2 \end{bmatrix},$$

so

$$\begin{aligned} \mathrm{tr}(A^2) &= a^2 + R_A S_A + \mathrm{tr}(S_A R_A + M_A^2) \\ &= a + R_A S_A + \mathrm{tr}(S_A R_A) + \mathrm{tr}(M_A^2) \\ &= a + R_A S_A + R_A S_A + \mathrm{tr}(M_A^2), \end{aligned}$$

and again $R_A S_A + R_A S_A = 0$ and $\mathrm{tr}(M_A^2) = \mathrm{tr}(M_A)$ by induction; an induction that can be carried out in **LA**. On the other hand, $\mathrm{tr}(A) = \mathrm{tr}(A^{2^i})$ can be carried out in **LAP**. □

Lemma 11 $\mathrm{tr}(A^t A) = \mathrm{tr}(AA^t) = \sum_{i,j} A_{ij}$.

Proof: $\mathrm{tr}(A^t A) = \sum_i (A^t A)_{ii} = \sum_{i,j} A^t_{ij} A_{ji} = \sum_{i,j} A_{ji} A_{ji} = \sum_{ij} A_{ij}$. □

In fact, from the proof of lemma 11 we see that $(A^t A)_{ii}$ is the sum of the elements in row i of A.

7 Annihilating polynomials

We say that $p(x) \neq 0$ is an *annihilating polynomial* of A if $p(A) = 0$. Of course, the annihilating polynomial par excellence of any matrix is its characteristic polynomial; this is the famous Cayley-Hamilton theorem that can be shown in $\exists$**LA** (see [SC04]).

Lemma 12 **LA** *proves that* $p(A)^2 = p(A^2)$.

Proof: $p(A)^2 = \left(\sum_i a_i A^i\right)^2 = \sum_{i,j} a_i a_j A^{i+j} = \sum_i a_i^2 A^{2i}$, and we lost terms where $i \neq j$ since $a_i a_j A^{i+j} + a_j a_i A^{j+i} = 0$, so $p(A)^2 = \sum_i a_i (A^i)^2 = p(A^2)$. □

It follows from lemma 12 that $p(A)^{2^i} = p(A^{2^i})$, and that this can be shown in **LAP** with exp. The following lemma is an immediate consequence of the previous one.

Lemma 13 *If* $p(x)$ *is an annihilating polynomial for* A*, then* $p(x)$ *is also an annihilating polynomial for* A^{2^i} *for all* i*.*

Proof: Suppose $p(A) = 0$. Then $0 = p(A)^{2^i} = p(A^{2^i})$. □

Lemma 14 *If* $AB = I$ *and* $p(x)$ *is an annihilating polynomial for* A*, then we can prove, in* **LAP***, that* $BA = I$*. That is,*

$$\mathbf{LAP} \vdash (AB = I \wedge p \neq 0 \wedge p(A) = 0) \rightarrow BA = I.$$

Proof: Suppose that $p(A) = a_0 I + a_1 A + \cdots + a_k A^k = 0$. As $p \neq 0$, let a_i be the smallest non-zero coefficient of this polynomial; so, in effect,

$$p(A) = a_i A^i + a_{i+1} A^{i+1} + \cdots + a_k A^k.$$

Consider

$$\begin{aligned} 0 = p(A)B^i &= a_i I + a_{i+1} A + \cdots + a_k A^{k-i} \\ &= I + A(a_{i+1} I + \cdots + a_k A^{k-i-1}) \\ &= I + (a_{i+1} I + \cdots + a_k A^{k-i-1})A, \end{aligned}$$

i.e., $(a_{i+1} I + \cdots + a_k A^{k-i-1})$ is the (two-sided) inverse of A, and so $AB = I$ implies $ABA = A$ which implies $A(BA - I) = 0$, and now using the inverse of A we obtain $BA - I = 0$ and so $BA = I$. □

8 Pigeonhole principle and counting

Let PHP denote the Pigeonhole principle. In this section we present three proofs of the 'Inverse Identity' that can be formalized in **LAP** augmented with PHP over sets of exponential size. What is interesting about these "counting arguments" is that they dispense with linear algebra in proving the 'Inverse Identity'; they rely on the finiteness of the underlying field, and basic ring properties of matrix addition and multiplication. As such, they offer a limited proof-complexity insight into the 'Inverse Identity'.

Proof I of the 'Inverse Identity'

This is a simple proof of the 'Inverse Identity', which extends easily to any finite field. It uses PHP over sets of size 2^{n^2} for matrices in $M(n)$.

Consider the sequence $I, A, A^2, A^3, \ldots, A^{2^{n^2}}$. Since $|M(n)| = 2^{n^2}$ it follows by PHP that there exist $0 \leq i < j \leq 2^{n^2}$ such that $A^i = A^j$; but then, using $AB = I$, we have $A^i B^i = A^j B^i \Rightarrow I = A^{j-i}$, where $j - i > 0$, and so A^{j-i-1} is a (left and right) inverse of A, and using lemma 1 we are done.

Proof II of the 'Inverse Identity'

Let $\Phi : M(n) \longrightarrow M(m)$ be a mapping.

Lemma 15 *If $n > m$, then $\exists Y \in M(m)$ such that*

$$|\{X \in M(n) : \Phi(X) = Y\}| \geq 2^{(n-m)^2}.$$

Proof: Let $S_\Phi(Y) := \{X \in M(n) : \Phi(X) = Y\}$. Since

$$M(n) = \bigcup_{Y \in M(m)} S_\Phi(Y)$$

it follows that $|M(n)| \leq |M(m)| \cdot \max\{|S_\Phi(Y)| : Y \in M(m)\}$, and so we have that $2^{(n-m)^2} \leq 2^{n^2 - m^2} \leq \max\{|S_\Phi(Y)| : Y \in M(m)\}$. □

Suppose that $AB = I$, and let $\Phi_A : M(n+1) \longrightarrow M(n)$ be a mapping defined as follows:

$$\Phi_A(C) := c_0 A + c_1 A^2 + \cdots + c_{(n+1)^2-1} A^{(n+1)^2-1}, \tag{2}$$

where A is assumed to be $n \times n$, and entry c_j is $C_{1+\mathrm{div}(j,n),\mathrm{rem}(j,n)}$, i.e., given q, r such that $j = qn + r$ where $0 \leq r < n$, then c_j is entry $c_{q+1,r}$. By lemma 15 we have that there is a $Y \in M(n)$ such that there exist $C \neq C'$ mapping to it, i.e., $\Phi_A(C) = \Phi_A(C')$, and so $\Phi_A(C) + \Phi_A(C') = 0$. Therefore, we obtain an annihilating polynomial for A of degree $(n+1)^2 - 1$; by lemma 14 we are done.

Proof III of the 'Inverse Identity'

We define $\Phi_A(C)$ as in (2), and we present a variation of the argument in proof II. Since $|M(n+1)| = 2^{(n+1)^2}$, one of the following two must hold:

$$|\{C \in M(n+1) : (\Phi_A(C))_{11} = 0\}| \geq 2^{(n+1)^2}/2,$$
$$|\{C \in M(n+1) : (\Phi_A(C))_{11} = 1\}| \geq 2^{(n+1)^2}/2.$$

We pick the set for which it is true; we now repeat the argument with the C's in that set and $(\Phi_A(C))_{12}$. We end up with C, C' such that $\Phi_A(C) = \Phi_A(C')$, and once again obtain an annihilating polynomial for A.

9 Gaussian Elimination

We give a proof of the 'Inverse Identity' based on the Gaussian elimination algorithm; this proof can be formalized in $\exists\mathbf{LA}$, and hence it is a polynomial time proof. In fact, in [Sol02, TS05] it has been shown that polysize extended Frege can prove the correctness of the Gaussian elimination procedure (over $\mathbb{Z}_2$, and over bigger fields as well). $\exists\mathbf{LA}$ allows induction over formulas asserting the existence of matrices; such formulas can express the existence of row and column operations necessary to carry out the Gaussian elimination algorithm. Hence $\exists\mathbf{LA}$ proves the correctness of Gaussian elimination (every matrix can be put in upper triangular form), and this in turn can be employed to prove the 'Inverse Identity'.

Suppose that $AB = I$; we show that A has some left-inverse (lemma 1). Recall the definition of S_A in (1). If S_A is zero then repeat the argument inductively on $M_A M_B = I_{n-1}$.

Otherwise, if $S_A \neq 0$, let P be a permutation matrix defined as follows: if $a = 1$, $P = I_n$; if $a = 0$, then P swaps the first row of A with a row whose first entry is non-zero; P is just the identity matrix with the corresponding rows swapped. Also let

$$C = \left[\begin{array}{cc} 1 & 0 \\ S_{PA} & I_{n-1} \end{array} \right],$$

that is, C is the identity matrix where the first column, except for the top entry, is replaced by the corresponding entries in PA. Observe that $CC = PP = I_n$. Then

$$\begin{aligned} AB = I_n &\Rightarrow (CP)AB(PC) = (CP)I_n(PC) \\ &\Rightarrow (CPA)(BPC) = C(PP)C = CC = I_n, \end{aligned}$$

and

$$CPA = \left[\begin{array}{cc} 1 & 0 \\ S_{PA} & I_{n-1} \end{array} \right] \left[\begin{array}{cc} 1 & R_{PA} \\ S_{PA} & M_{PA} \end{array} \right] = \left[\begin{array}{cc} 1 & R_{PA} \\ 0 & S_{PA}R_{PA} + M_{PA} \end{array} \right].$$

We now repeat the argument inductively on $M_{CPA}M_{BPC} = I_{n-1}$. Notice that this proof is intrinsically sequential, as at each step we construct a new matrix, and we need the previous steps to do that.

10 Quasi-triangular matrices

In this section we explore the following question: what is the weakest condition on matrices A, B that allows a proof of the 'Inverse Identity' in $\mathbf{LA}$? We already know from lemma 8 that if A, B are symmetric matrices then $\mathbf{LA}$ proves the 'Inverse Identity' for such matrices. Here, inspired by the Gaussian elimination proof in the previous section, we define the notion of a "quasi-triangular pair" of matrices.

Given $A \in M(n)$, let $M_{A,i}$ be its i-th minor. That is $M_{A,0} = A$, $M_{A,1} = M_A$, and for $i \geq 2$, $M_{A,i} = M_{M_{A,i-1}}$, i.e., $M_{A,i}$ is A with the first i rows and columns removed. Let $a_i, R_{A,i}, S_{A,i}$ be defined as follows: $a_i = (M_{A,i-1})_{ii} = A_{ii}$, and $R_{A,i} = R_{M_{A,i-1}}$ and $S_{A,i} = S_{M_{A,i-1}}$. We say that a matrix A is *quasi-triangular* if for all i we have that $R_{A,i}$ is zero or $S_{A,i}$ is zero; it is clear how this is a generalization of the notion of a triangular matrix.

We define recursively what it means for a matrix A to be a *quasi-transpose* of a matrix B; if $A, B \in M(1)$, then A is a quasi-transpose of B if $A = B$. For $i > 1$, A is a quasi-transpose of B if $A_{11} = B_{11}$ and either $R_A = R_B \wedge S_A = S_B$, or $R_A = S_B^t \wedge S_A = R_B^t$, and M_A is a quasi-transpose of M_B.

Lemma 16 *The matrix A is quasi-triangular if A is the quasi-transpose of a triangular matrix.*

Finally, we say that matrices (A, B) are a *quasi-triangular pair* if for all i at least one of $\{R_{A,i}, S_{A,i}, R_{B,i}, S_{B,i}\}$ is zero. Observe that if (A, B) is a quasi-triangular pair then so is (M_A, M_B); indeed, the point of this definition is to find as weak a condition on A, B as possible to ensure that if $AB = I_n$, then $M_A M_B = I_{n-1}$, which allows induction on **LA** formulas and consequently an **LA** proof of the 'Inverse Identity'.

Lemma 17 **LA** *proves that if (A, B) is a quasi-triangular pair, and $AB = I$, then $BA = I$.*

Proof: Suppose that $A, B \in M(n)$, and they are a quasi-triangular pair. We prove that $\mathbf{LA} \vdash AB = I \to BA = I$, by induction on n, by cases on the definition of a "quasi-triangular pair."

$$AB = \begin{bmatrix} a & R_A \\ S_A & M_A \end{bmatrix} \begin{bmatrix} b & R_B \\ S_B & M_B \end{bmatrix} = \begin{bmatrix} ab + R_A S_B & aR_B + R_A M_B \\ bS_A + M_A S_B & S_A R_B + M_A M_B \end{bmatrix} \tag{3}$$

Case 1: $S_A = 0$ or $R_B = 0$, then we can see from equation (3) that

$$AB = \begin{bmatrix} ab + R_A S_B & aR_B + R_A M_B \\ bS_A + M_A S_B & M_A M_B \end{bmatrix},$$

since $S_A R_B = 0 \iff S_A = 0 \vee R_B = 0$, and since $AB = I_n$, it follows that $M_A M_B = I_{n-1}$, and given that (A, B) is a quasi-triangular pair, so is (M_A, M_B), and hence by induction $M_B M_A = I_{n-1}$.

Suppose now that $S_A = 0$. Then,

$$BA = \begin{bmatrix} ba & bR_A + R_B M_A \\ aS_B & S_B R_A + I_{n-1} \end{bmatrix},$$

and from (3) we see that $0 = bS_A + M_A S_B = M_A S_B$, and since $M_B M_A = I_{n-1}$ it follows that $S_B = 0$, so $ab = 1$ and so $ba = 1$. Finally,

$$bR_A + R_B M_A = R_A + R_B M_A = R_A + R_A M_B M_A = R_A + R_A = 0.$$

Therefore, $BA = I_n$. The case where $R_B = 0$ is symmetric.

Case 2: $S_B = 0$ or $R_A = 0$, then since $ab + R_A S_B = 1$ it follows that $ab = 1$, and so $a = b = 1$. Therefore,

$$I_n = AB = \begin{bmatrix} 1 & R_B + R_A M_B \\ S_A + M_A S_B & S_A R_B + M_A M_B \end{bmatrix},$$

thus $M_A S_B = S_A$ and $R_A M_B = R_B$, and hence

$$M_A \underbrace{S_B R_A}_{=0} M_B + M_A M_B = I_{n-1},$$

and so $M_A M_B = I_{n-1}$. By induction we conclude that $M_B M_A = I_{n-1}$, and we are done. □

11 Lanczos algorithm

The “Block Lanczos Algorithm for Finding Dependencies over a Finite Field” was invented by Peter L. Montgomery ([Mon95]). Recall that if $AB = I$, then from basic algebraic manipulations we obtain that $A(BA - I) = 0$; thus, if we could show that A has *any* left-inverse, we would be able to show the ‘Inverse Identity’. As symmetric matrices have a lot of nice properties, perhaps one could work with the symmetric matrix $\hat{A} = A^t A$ instead of working with A. Note that showing that $\hat{A}$ has a left-inverse would still prove, from $A(BA - I) = 0$, that $BA = I$. This is because $A(BA - I) = 0$ implies $A^t A(BA - I) = 0$, and so $\hat{A}(BA - I) = 0$. This section suggests a connection between the “Block Lanczos Algorithm” and our ‘Inverse Identity’.

Let A be a symmetric matrix and suppose that we have a set of m vectors $\{w_1, w_2, \ldots, w_m\}$, and that they satisfy the following three conditions:

$$\begin{aligned} & w_i^t A w_i \neq 0 \quad 1 \leq i \leq m \\ & w_i^t A w_j = 0 \quad i \neq j \\ & \exists X(AW = WX) \end{aligned} \tag{4}$$

where the second condition says that the w_i's are A-orthogonal, as defined in section 5, and where in the last condition $W = \begin{bmatrix} w_1 & w_2 & \ldots & w_m \end{bmatrix}$, i.e., W is a matrix whose rows are the w_i's. Note that saying $\exists X(AW = WX)$ is equivalent to stating that for all w_i, $Aw_i \in \operatorname{span}\{w_1, w_2, \ldots, w_m\}$. Suppose further that $\exists y(b = Wy)$ and define the vector v as follows:

$$v := \sum_{i=1}^{m} \frac{w_i^t b}{w_i^t A w_i} w_i. \tag{5}$$

Theorem 2 **LA** $\vdash [(4) \wedge \exists y(b = Wy) \wedge (5)] \rightarrow Av = b$.

PROOF: We prove the theorem in the case $\mathbb{F} = \mathbb{Z}_2$. Over the field $\mathbb{Z}_2$ we have the implication $w_i^t A w_i \neq 0 \Rightarrow w_i^t A w_i = 1$, and hence definition (5) can be restated as $v := \sum_{i=1}^m (w_i^t b) w_i$. Then:

$$\begin{aligned} W^t(Av - b) &= W^t(A(\sum_{i=1}^m (w_i^t b)w_i) - b) \overset{(*)}{=} (\sum_{i=1}^m (w_i^t b) W^t A w_i) - W^t b \\ &= (\sum_{i=1}^m (w_i^t b)(w_i^t A w_i) e_i) - W^t b \overset{(**)}{=} (\sum_{i=1}^m (w_i^t b) e_i) - W^t b = W^t b - W^t b = 0, \end{aligned}$$

where $(*)$ and $(**)$ follow from (4) since $w_j^t A w_i = 0$ for $j \neq i$ and $w_i^t A w_i = 1$, respectively.

Thus we obtain $W^t(Av - b) = 0$ and hence, for any $n \times m$ matrix X, $(WX)^t(Av - b) = 0$. By (4) $\exists X(AW = WX)$, and so $(AW)^t(Av - b) = 0$, and since A is symmetric it follows that

$$W^t A(Av - b) = 0. \tag{6}$$

By assumptions $\exists y(Av - b = Wy)$, i.e., $Av - b = \sum_{i=1}^m y_i w_i$, and so multiplying on the left by $w_i^t A$ we conclude, for every i, that $y_i = 0$. Hence $Av - b = 0$ and so $Av = b$. □

Note that given (4) and (5) and $\exists y(b = Wy)$, then $Av = b$ for any field $\mathbb{F}$, finite or infinite.

Suppose that $AB = I$, and consider the matrix $\hat{A} = A^t A$. Represent B as $\begin{bmatrix} b_1 & b_2 & \dots & b_n \end{bmatrix}$ where b_i is the i-th column of B. Then observe:

$$b_i^t \hat{A} b_j = b_i^t A^t A b_j = (Ab_i)^t (Ab_j) = e_i^t e_j = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$

and thus B satisfies the first two Lanczos conditions for $\hat{A} = A^t A$ given in (4). The third Lanczos condition is equivalent to $BA = I$.

12 Open questions

If $AB = I$, can we show that $A^t A$ (which is symmetric) is also invertible? Since $\mathbf{LA} \vdash AB = I \rightarrow A(BA - I) = 0$, it follows that $(A^t A)(BA - I) = 0$ is provable from $AB = I$ in **LA**. Once we have a left-inverse for $A^t A$ we would have $BA = I$. It would be interesting to see if Lanczos' algorithm can be of help here.

Matrices over $\mathbb{Z}_2$ are often employed as adjacency matrices; that is, given a graph $G = ([n], E)$, where $E \subseteq [n] \times [n]$, it can be represented by $A_G \in M(n)$ as follows: $(i, j) \in E \iff A_{ij} = 1$. But matrices over $\mathbb{Z}_2$ can be also seen as *incidence* matrices where given a collection of n elements, and a collection $X_1, X_2, \dots, X_m$ of subsets of $[n]$, $i \in X_k \iff A_{ij} = 1$, where $A \in M_{n \times m}(\mathbb{Z}_2)$. An interesting paper examining incidence matrices is [Rys60].

13 Appendix

The logical theory **LA** is strong enough to prove the ring properties of matrices such as $A(BC) = (AB)C, A+B = B+A$, but weak enough so that the theorems of **LA** translate into propositional tautologies with short Frege proofs. **LA** has three sorts of object: *indices* (i.e., natural numbers), *ring elements*, and *matrices*, where the corresponding variables are denoted $i, j, k, ...$; $a, b, c, ...$; and $A, B, C, ...$, respectively. The semantic assumes that objects of type ring are from a fixed but arbitrary ring (for the purpose of this paper we are only interested in $\mathbb{Z}_2$, which is a field), and objects of type matrix have entries from that ring.

Terms and formulas are built from the following function and predicate symbols, which together comprise the language $\mathcal{L}_{\mathbf{LA}}$:

$$\begin{gathered} 0_{\text{index}}, 1_{\text{index}}, +_{\text{index}}, *_{\text{index}}, -_{\text{index}}, \mathtt{div}, \mathtt{rem}, \\ 0_{\text{ring}}, 1_{\text{ring}}, +_{\text{ring}}, *_{\text{ring}}, -_{\text{ring}}, {}^{-1}, \mathtt{r}, \mathtt{c}, \mathtt{e}, \Sigma, \\ \leq_{\text{index}}, =_{\text{index}}, =_{\text{ring}}, =_{\text{matrix}}, \text{cond}_{\text{index}}, \text{cond}_{\text{ring}} \end{gathered} \tag{7}$$

The intended meaning should be clear, except in the case of $-_{\text{index}}$, cut-off subtraction, defined as $i - j = 0$ if $i < j$. For a matrix A: $\mathtt{r}(A), \mathtt{c}(A)$ are the numbers of rows and columns in A, $\mathtt{e}(A, i, j)$ is the ring element A_{ij} (where $A_{ij} = 0$ if $i = 0$ or $j = 0$ or $i > \mathtt{r}(A)$ or $j > \mathtt{c}(A)$), $\Sigma(A)$ is the sum of the elements in A. Also $\text{cond}(\alpha, t_1, t_2)$ is interpreted **if** α **then** t_1 **else** t_2, where α is a formula all of whose atomic sub-formulas have the form $m \leq n$ or $m = n$, where m, n are terms of type index, and t_1, t_2 are terms either both of type index or both of type ring. The subscripts $_{\text{index}}$, $_{\text{ring}}$, and $_{\text{matrix}}$ are usually omitted, since they ought to be clear from the context.

We use n, m for terms of type index, t, u for terms of type ring, and T, U for terms of type matrix. Terms of all three types are constructed from variables and the symbols above in the usual way, except that terms of type matrix are either variables $A, B, C, ...$ or λ-terms $\lambda ij\langle m, n, t\rangle$. Here i and j are variables of type index bound by the λ operator, intended to range over the rows and columns of the matrix. Also m, n are terms of type index *not* containing i, j (representing the numbers of rows and columns of the matrix) and t is a term of type ring (representing the matrix element in position (i, j)).

Atomic formulas have the forms $m \leq n, m = n, t = u, T = U$, where the three occurrences of $=$ formally have subscripts $_{\text{index}}, _{\text{ring}}, _{\text{matrix}}$, respectively. General formulas are built from atomic formulas using the propositional connectives $\neg, \vee, \wedge$ and quantifiers $\forall, \exists$.

13.1 Axioms and rules of LA

For each axiom listed below, every legal substitution of terms for free variables is an axiom of **LA**. Note that in a λ term $\lambda ij\langle m, n, t\rangle$ the variables i, j are bound. Substitution instances must respect the usual rules which prevent free variables from being caught by the binding operator λij. The bound variables i, j may be renamed to any new distinct pair of variables.

13.1.1 Equality Axioms

These are the usual equality axioms, generalized to apply to the three-sorted theory **LA**. Here $=$ can be any of the three equality symbols, x, y, z are variables of any of the three sorts (as long as the formulas are syntactically correct). In A4, the symbol f can be any of the non-constant function symbols of **LA**. However A5 applies only to $\leq$, since this in the only predicate symbol of **LA** other than $=$.

A1 $x = x$
A2 $x = y \rightarrow y = x$
A3 $(x = y \wedge y = z) \rightarrow x = z$
A4 $x_1 = y_1, ..., x_n = y_n \rightarrow f x_1 ... x_n = f y_1 ... y_n$
A5 $i_1 = j_1, i_2 = j_2, i_1 \leq i_2 \rightarrow j_1 \leq j_2$

13.1.2 Axioms for indices

These are the axioms that govern the behavior of index elements. The index elements are used to access the entries of matrices, and so we need to define some basic number theoretic operations.

A6 $i + 1 \neq 0$
A7 $i * (j + 1) = (i * j) + i$
A8 $i + 1 = j + 1 \rightarrow i = j$
A9 $i \leq i + j$
A10 $i + 0 = i$
A11 $i \leq j \wedge j \leq i$
A12 $i + (j + 1) = (i + j) + 1$
A13 $[i \leq j \wedge j \leq i] \rightarrow i = j$
A14 $i * 0 = 0$
A15 $[i \leq j \wedge i + k = j] \rightarrow j - i = k$
A16 $\neg(i \leq j) \rightarrow j - i = 0$
A17 $[\alpha \rightarrow \mathrm{cond}(\alpha, i, j) = i] \wedge [\neg\alpha \rightarrow \mathrm{cond}(\alpha, i, j) = j]$

13.1.3 Axioms for a ring

These are the axioms that govern the behavior for ring elements; addition and multiplication, as well as additive inverses. We do not need multiplicative inverses.

A18 $0 \neq 1 \wedge a + 0 = a$
A19 $a + (-a) = 0$
A20 $1 * a = a$
A21 $a + b = b + a$
A22 $a * b = b * a$
A23 $a + (b + c) = (a + b) + c$
A24 $a * (b * c) = (a * b) * c$

A25 $a * (b + c) = a * b + a * c$
A26 $[\alpha \to \mathrm{cond}(\alpha, a, b) = a] \wedge [\neg\alpha \to \mathrm{cond}(\alpha, a, b) = b]$

13.1.4 Axioms for matrices

Axiom A27 states that $\mathtt{e}(A, i, j)$ is zero when i, j are outside the size of A. Axiom A28 defines the behavior of constructed matrices. Axioms A29-A32 define the function Σ recursively by first defining it for row vectors, then column vectors ($A^t := \lambda ij\langle \mathtt{c}(A), \mathtt{r}(A), A_{ji}\rangle$), and then in general using the decomposition (8). Finally, axiom A33 takes care of empty matrices.

A27 $(i = 0 \vee \mathtt{r}(A) < i \vee j = 0 \vee \mathtt{c}(A) < j) \to \mathtt{e}(A, i, j) = 0$
A28 $\mathtt{r}(\lambda ij\langle m, n, t\rangle) = m \wedge \mathtt{c}(\lambda ij\langle m, n, t\rangle) = n \wedge [1 \le i \wedge i \le m \wedge 1 \le j \wedge j \le n]$
$\to \mathtt{e}(\lambda ij\langle m, n, t\rangle, i, j) = t$
A29 $\mathtt{r}(A) = 1, \mathtt{c}(A) = 1 \to \Sigma(A) = \mathtt{e}(A, 1, 1)$
A30 $\mathtt{r}(A) = 1 \wedge 1 < \mathtt{c}(A) \to \Sigma(A) = \Sigma(\lambda ij\langle 1, \mathtt{c}(A) - 1, A_{ij}\rangle) + A_{1\mathtt{c}(A)}$
A31 $\mathtt{c}(A) = 1 \to \Sigma(A) = \Sigma(A^t)$
A32 $1 < \mathtt{r}(A) \wedge 1 < \mathtt{c}(A) \to \Sigma(A) =\mathtt{e}(A, 1, 1) + \Sigma(\mathtt{R}(A)) + \Sigma(\mathtt{S}(A)) + \Sigma(\mathtt{M}(A))$
A33 $\mathtt{r}(A) = 0 \vee \mathtt{c}(A) = 0 \to \Sigma A = 0$

Where

$$\begin{aligned} &\mathtt{R}(A) := \lambda ij\langle 1, \mathtt{c}(A) - 1, \mathtt{e}(A, 1, i + 1)\rangle, \\ &\mathtt{S}(A) := \lambda ij\langle \mathtt{r}(A) - 1, 1, \mathtt{e}(A, i + 1, 1)\rangle, \\ &\mathtt{M}(A) := \lambda ij\langle \mathtt{r}(A) - 1, \mathtt{c}(A) - 1, \mathtt{e}(A, i + 1, j + 1)\rangle. \end{aligned} \tag{8}$$

13.1.5 Rules for LA

In addition to all the axioms just presented, **LA** has two rules: matrix equality and induction.

Matrix equality rule
From the three premises:

1. $\mathtt{e}(T, i, j) = \mathtt{e}(U, i, j)$
2. $\mathtt{r}(T) = \mathtt{r}(U)$
3. $\mathtt{c}(T) = \mathtt{c}(U)$

we conclude $T = U$.

The only restriction is that the variables i, j may not occur free in $T = U$; other than that, T and U can be arbitrary matrix terms. Our semantics implies that i and j are implicitly universally quantified in the top formula. The rule allows us to conclude $T = U$, provided that T and U have the same numbers of rows and columns, and corresponding entries are equal.

Induction rule

$$\frac{\alpha(i) \to \alpha(i+1)}{\alpha(0) \to \alpha(n)}$$

Here $\alpha(i)$ is any formula, n is any term of type index, and $\alpha(n)$ indicates n is substituted for free occurrences of i in $\alpha(i)$. (Similarly for $\alpha(0)$.)

This completes the description of **LA**. We finish this section by observing the substitution property in the lemma below. We say that a formula S' of **LA** is a *substitution instance* of a formula S of **LA** provided that S' results by substituting terms for free variables of S. Of course each term must have the same sort as the variable it replaces, and bound variables must be renamed as appropriate.

Lemma 18 *Every substitution instance of a theorem of* **LA** *is a theorem of* **LA**.

This follows by straightforward induction on **LA** proofs. The base case follows from the fact that every substitution instance of an **LA** axiom is an **LA** axiom.

References

[Alb38] A. Adrian Albert. Symmetric and alternating matrices in an arbitrary field, i. *Transactions of the American Mathematical Society*, 43(3):386–436, May 1938.

[BR91] Richard A. Brualdi and Herbert J. Ryser. *Combinatorial Matrix Theory.* Cambridge University Press, 1991.

[CF10] Stephen Cook and Lila Fontes. Formal theories for linear algebra. Presented at the Federated Logic Conference, 2010.

[Cob58] S. M. Cobb. On powers of matrices with elements in the field of integers modulo 2. *The mathematical gazette*, 42(342):267–271, December 1958.

[Fil10] Yuval Filmus. Range of symmetric matrices over GF(2). Unpublished note, University of Toronto, January 2010.

[Mac69] Jessie MacWilliams. Orthogonal matrices over finite fields. *The American Mathematical Monthly*, 76(2):152–164, February 1969.

[Mon95] Peter L. Montgomery. A block Lanczos algorithm for finding dependencies over GF(2). In *EUROCRYPT*, pages 106–120, 1995.

[Rys60] H. J. Ryser. Matrices of zeros and ones. *Bulletin of the American Mathematical Society*, 66(6):442–464, February 1960.

[SC04] Michael Soltys and Stephen A. Cook. The complexity of derivations of matrix identities. *Annals of Pure and Applied Logic*, 130(1–3):277–323, December 2004.

[Sol02] Michael Soltys. Extended Frege and Gaussian Elimination. *Bulletin of the Section of Logic*, 31(4):1–17, 2002.

[TS05] Neil Thapen and Michael Soltys. Weak theories of linear algebra. *Archive for Mathematical Logic*, 44(2):195–208, 2005.