Lecture 33 Matrix Rank

Recall that the *rank* of an $m \times n$ matrix A over a field k is the maximum number of linearly independent rows (or columns) of A. It is the dimension of the image of the linear map $k^n \to k^m$ defined by A; equivalently, it is n minus the dimension of the kernel (the set of vectors in k^n annihilated by the map).

Once we have an NC algorithm to calculate the rank of a matrix, the door is opened for a wide variety of other NC algorithms in linear algebra. For example, to compute a basis for the vector space spanned by the columns of some matrix, we can compute the ranks of all sets of columns $\{c_1, \ldots, c_i\}$, $1 \leq i \leq n$, and add c_i to the basis only if the rank of $\{c_1, \ldots, c_i\}$ is one greater than the rank of $\{c_1, \ldots, c_{i-1}\}$.

We will start with the algorithm of Ibarra, Moran, and Rosier [53], which computes the rank of a matrix over the complex numbers $\mathcal{C}$.

Recall the following lemma from the last lecture:

Lemma 33.1 *Let C be an $n \times n$ matrix over any field. If* $\operatorname{rank} C^2 = \operatorname{rank} C$, *then we can compute* $\operatorname{rank} C$ *in NC by computing the characteristic polynomial* $\det(xI - C)$ *and finding the highest power of x that divides it, say x^d. Then* $\operatorname{rank} C = n - d$.

Let A be a matrix over $\mathcal{C}$, not necessarily square. The *conjugate transpose* of A, denoted $\overline{A}^T$, is the transpose of A with every entry replaced by its complex conjugate. Recall that the conjugate $\overline{z}$ of a complex number z is obtained by reflecting in the real axis: if $z = a + ib$, where a and b are real,

then $\overline{z} = a - ib$. Note that the product $z\overline{z}$ is always a nonnegative real number:

$$(a + ib) \cdot (a - ib) \;=\; a^2 + b^2 \ .$$

Let $B = \overline{A}^T A$. We will prove that A and B have the same rank; moreover, rank $B^2 =$ rank B and B is square, so Lemma 33.1 applies.

The matrix B is of a particularly nice form: it is *Hermitian*, which means that $B = \overline{B}^T$. A Hermitian matrix is the complex analog of a symmetric matrix.

Lemma 33.2 *For any complex vector $y \in \mathcal{C}^n$, $\overline{y}^T y = 0$ iff $y = 0$.*

Proof. If $y = (a_1, a_2, \ldots, a_n)$ then

$$\begin{aligned} \overline{y}^T y &= (\overline{a}_1, \overline{a}_2, \ldots, \overline{a}_n) \cdot (a_1, a_2, \ldots, a_n) \\ &= \sum_{i=1}^{n} \overline{a}_i \cdot a_i \ . \end{aligned}$$

Now $\overline{a}_i \cdot a_i$ is always a nonnegative real number, and it can only be zero if $a_i = 0$. The sum of nonnegative reals can only be zero if each term, and hence each a_i, is zero. □

Recall that the *kernel* of a linear map is the set of vectors that are mapped to the origin. Thus if the linear map is represented by the matrix A, then

$$\ker A \;=\; \{x \in \mathcal{C}^n \mid Ax = 0\} \ .$$

The rank of A is the dimension of the image of A, which is the same as $n - \dim \ker A$. The following lemma shows that it is sufficient to find the dimension of the kernel of $\overline{A}^T A$.

Lemma 33.3 $\ker A = \ker \overline{A}^T A$.

Proof.
($\subseteq$) If $x \in \ker A$ then $Ax = 0$, which implies that $\overline{A}^T Ax = 0$.
($\supseteq$) Suppose $x \in \ker \overline{A}^T A$. Then

$$\begin{aligned} \overline{A}^T Ax = 0 \;&\rightarrow\; \overline{x}^T \overline{A}^T Ax = 0 \\ &\rightarrow\; \overline{(Ax)}^T Ax = 0 \\ &\rightarrow\; Ax = 0 \qquad \text{by Lemma 33.2.} \end{aligned}$$

□

Lemma 33.4 *If B is Hermitian, then* rank $B =$ rank B^2.

Proof. It suffices to show that the kernels of B and B^2 coincide. Surely $\ker B \subseteq \ker B^2$. Now suppose $x \in \ker B^2$. Then

$$\begin{aligned} B^2x = 0 &\rightarrow \overline{x}^T BBx = 0 \\ &\rightarrow \overline{x}^T \overline{B}^T Bx = 0 \quad \text{because } B \text{ is Hermitian, so } B = \overline{B}^T \\ &\rightarrow \overline{(Bx)}^T Bx = 0 \\ &\rightarrow Bx = 0 \qquad \text{by Lemma 33.2.} \end{aligned}$$

Therefore $x \in \ker B$. □

Putting all this together, here is the algorithm for computing the rank of A: compute the square matrix $B = \overline{A}^T A$, compute the characteristic polynomial using Csanky or Chistov, and find the highest power x^d of x that divides it. The rank of A is $n - d$. As we have seen, all these steps can be performed in NC.

33.1 Mulmuley's Algorithm

For a complex matrix A, we showed that A, $\overline{A}^T A$, and $(\overline{A}^T A)^2$ all have the same rank, thus we can apply Lemma 33.1 to the square Hermitian matrix $\overline{A}^T A$. In the special case of real matrices, this says that A, $A^T A$, and $(A^T A)^2$ all have the same rank, and we can apply Lemma 33.1 to the symmetric matrix $A^T A$.

Unfortunately, this does not work for fields of finite characteristic. For example, over the field $\mathcal{Z}_5$, $(1,2) \cdot (1,2) = 0$; moreover, the matrix

$$A = \begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix} \tag{45}$$

is symmetric and of rank 1, but $A^T A = A^2 = 0$.

This pathological state of affairs was partially resolved by Borodin, von zur Gathen, and Hopcroft [15], who gave a probabilistic NC algorithm, and Chistov [18] who gave a nonuniform deterministic algorithm. Mulmuley [82] gave the first deterministic NC algorithm, which we describe here.

Mulmuley keeps inner products from vanishing by throwing in indeterminates. The idea being exploited here is that even over fields of finite characteristic, a polynomial expression in the indeterminate x vanishes if and only if all its coefficients vanish; in other words, the indeterminate x is *transcendental* over the field k (it is not the root of any polynomial with coefficients in k). If not too many indeterminates are used and the degrees of the polynomials involved are not too big, then the computations can still be done efficiently in parallel, except that now we work *symbolically*, using polynomial arithmetic. This may require another factor of n more processors, but can still be done in NC.

Officially, we will be working in the transcendental extension $k(x)$ of k. This is isomorphic to the field of rational functions over k described in Lecture 32. It is the smallest field that contains k and the single indeterminate x, and is unique up to isomorphism. The rational function p/q can be represented as the pair of polynomials (p, q), and the operations $+$ and $\cdot$ in $k(x)$ can be done using polynomial arithmetic on the numerators and denominators:

$$\begin{aligned} \frac{p_1}{q_1} + \frac{p_2}{q_2} &= \frac{p_1 q_2 + p_2 q_1}{q_1 q_2} \\ \frac{p_1}{q_1} \cdot \frac{p_2}{q_2} &= \frac{p_1 p_2}{q_1 q_2} . \end{aligned}$$

To test equality, we need to reduce these fractions to lowest terms by factoring out the gcd of the numerator and denominator, but this can be done in NC, as will be shown in the next lecture.

To illustrate the technique, consider the matrix A of (45) over $\mathcal{Z}_5$. Instead of working with A, we can work instead with the matrix XA, where x is an indeterminate and

$$X = \begin{bmatrix} 1 & 0 \\ 0 & x \end{bmatrix} .$$

Then

$$XA = \begin{bmatrix} 1 & 0 \\ 0 & x \end{bmatrix} \cdot \begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix} = \begin{bmatrix} 1 & 2 \\ 2x & 4x \end{bmatrix} .$$

The matrix X has entries in $k(x)$ and is nonsingular, therefore XA has the same rank as A. Moreover, the matrix

$$(XA)^2 = \begin{bmatrix} 1 & 2 \\ 2x & 4x \end{bmatrix}^2 = \begin{bmatrix} 1+4x & 2+3x \\ 2x+3x^2 & 4x+x^2 \end{bmatrix}$$

also has rank 1 (the second row is $2x$ times the first). Since $(XA)^2$, XA, and A all have the same rank, we can apply Lemma 33.1 to XA and the problem is solved.

This works in general. Let A be an $m \times n$ matrix, $m \geq n$, over an arbitrary field k. We note that going from k to $k(x)$ does not affect the rank of A, since the rank of A is r iff A has a nonzero $r \times r$ minor (determinant of an $r \times r$ submatrix), and this computation does not care whether we are over k or $k(x)$.

We can assume without loss of generality that A is square and symmetric; if not, we consider instead the square symmetric matrix

0	A
A^T	0

of size $(m+n) \times (m+n)$, whose rank is exactly twice that of A.

Now assume A is $n \times n$ and symmetric. Let X be the $n \times n$ diagonal matrix with $1, x, x^2, \ldots, x^{n-1}$ on the diagonal.

Lemma 33.5 (Mulmuley [82]) *The matrices A, XA and $(XA)^2$ all have the same rank.*

Proof. Certainly

$$\operatorname{rank} XAXA \;\leq\; \operatorname{rank} XA \;\leq\; \operatorname{rank} A\ .$$

Since X is nonsingular, $\operatorname{rank} XAXA = \operatorname{rank} AXA$, therefore in order to show that $\operatorname{rank} A \leq \operatorname{rank} XAXA$ it suffices to show that $\operatorname{rank} A \leq \operatorname{rank} AXA$. Assume for a contradiction that there is a vector $u \in k(x)^n$ such that $Au \neq 0$ but $AXAu = 0$. By multiplying through by the denominators of the elements of u, we can assume without loss of generality that $u = u(x) \in k[x]$. Let $v = Au$ and let $u(y)$ be $u(x)$ with x replaced by a new indeterminate y. Then $v(y) = Au(y)$ and

$$v(y)^T X v(x) \;=\; u(y)^T AXAu(x) \;=\; 0\ . \tag{46}$$

But if d is the maximum degree of any element of v and t is the maximum index for which v_t is of degree d, then the coefficient of $y^d x^{t-1} x^d$ is nonzero. This can be seen by writing out $v(y)^T X v(x)$ as a sum

$$v(y)^T X v(x) \;=\; \sum_{i=1}^{n} v(y)_i x^{i-1} v(x)_i$$

and noting that there is exactly one nonzero term in this expression with the monomial $y^d x^{t-1} x^d$, which cannot be canceled. This contradicts (46). □