Presburger arithmetic and its extensions

18–21 May 2026 *Lecturer:* Dmitry Chistikov

(please send corrections to d.chistikov@warwick.ac.uk)

This document is not complete lecture notes. I hope it can be useful in conjunction with other sources (which are referenced here and in the homework assignment).

On Gaussian elimination

The standard Gaussian elimination procedure can be used to solve systems of equations (equalities) $A\vec{x} = \vec{c}$, where A is an $m \times n$ integer matrix, $\vec{c}$ is an integer vector of dimension m, and $\vec{x}$ is a vector of n unknowns.

Indeed, the existence of a solution in $\mathbb{Q}^n$ can be determined in polynomial time. Why? Firstly, the procedure requires at most a cubic number of arithmetic operations. Secondly, numbers that arise in the procedure have small bit size, so the cost of each arithmetic operation is also polynomial. Proof of this second fact requires a bit of work.

A key observation is that, if no rows and columns are *permuted* by the procedure, leading principal minors of the matrix A remain unchanged throughout. The leading principal minor of order k, denoted μ_k, is the determinant of the $k \times k$ submatrix at the intersection of the first k rows and first k columns of A. Indeed, each iteration only subtracts multiples of one row from rows below it.

Further, denote for $i > k$ and $j > k$ by

$$a_{ij}^{(k)} = \begin{vmatrix} a_{11} & \dots & a_{1k} & a_{1j} \\ \vdots & \ddots & \vdots & \vdots \\ a_{k1} & \dots & a_{kk} & a_{kj} \\ a_{i1} & \dots & a_{ik} & a_{ij} \end{vmatrix}$$

the leading principal minor of order k *bordered* by the ith row and jth column. These minors (subdeterminants) are also unchanged by (the forward phase of) Gaussian elimination. After k iterations, one can see that $a_{ij}^{(k)} = a'_{ij} \cdot \mu_k$, where a'_{ij} is the new entry in position (i, j). Therefore, $a'_{ij} = a_{ij}^{(k)}/\mu_k$, and thus every entry is a ratio of two minors (subdeterminants).

It remains to note that the absolute value of an $r \times r$ subdeterminant does not exceed $r! \cdot \|A\|_{\max}^r$, where $\|A\|_{\max} = \max_{i,j} |a_{ij}|$. The bit size of any such integer is $O(r) \cdot (\log r + \log \|A\|_{\max})$, at most polynomial in the bit size of the input.

We have not included in our calculations the right-hand sides of the equations, but it is not difficult to see that they will not affect the conclusion.

Division-free Gaussian elimination is a variant of the standard procedure in which, in order to eliminate x from equations $ax + by + \ldots = 0$ and $cx + dy + \ldots = 0$, we "cross-multiply" them by c and a, respectively, and then subtract one from the other. Assume $a \neq 0$; otherwise a different element is chosen as the pivot. The first equation in the system stays unchanged and the second becomes

$$(ad - bc)\, y + \ldots = 0.$$

The same rule is used to eliminate x from any further equations; after this step, the first equation is the only one that contains x.

A *fraction-free* version, while additionally employing division (by so-called *Bareiss factors*), keeps all coefficients integral and their bit size polynomial.

References: [Bareiss 1968], [Chistikov 2024, Sec. 5.3].

As it turns out, we can further build a version of quantifier elimination procedure for *integer linear programming*, also based on (a nondeterministic version of) Gaussian elimination:

Theorem 4 (Chistikov, Mansutti, Starchak 2024)**.** *There is an algorithm running in nondeterministic polynomial time which, given a system of linear inequalities* $\varphi\colon A\vec{u} + B\vec{v} \le \vec{c}$*, outputs in each nondeterministic branch (denoted* β*) another system* $\psi_\beta\colon F_\beta \cdot \vec{v} \le \vec{g}_\beta$ *such that* $\exists \vec{u}\, \varphi \Leftrightarrow \bigvee_\beta \psi_\beta$*. Here* $\vec{u}$ *and* $\vec{v}$ *are vectors of integer variables,* A, B, F_β *are integer matrices (for each* β*), and* $\vec{c}$ *and* $\vec{g}_\beta$ *are integer vectors (for each* β*).*

The algorithm is given in the publication. It extends Gaussian elimination by "adding" two ingredients: slack variables (for handling inequalities rather than just equalities) and Bareiss factors (from the fraction-free algorithm). We sketch the key idea of the proof.

After an iteration of the main loop of the algorithm that targets x, let $y_1, \dots, y_r$ be unassigned slack variables in constraints involving x. The main challenge is to show that, if there is an assignment that satisfies the system before the iteration, then this property is preserved after it. Before the iteration, there is an assignment ν of integer values to the variables of the system as well as to slack variables that satisfies the system. In one iteration, the algorithm chooses (guesses) a nonnegative value from a finite range for one slack variable. We need to show that the range chosen for the values of slack suffices, that is, is big enough to preserve satisfiability.

Consider the auxiliary rational vector

$$\vec{q}^{\,\nu} = (q_1, \dots, q_r) \coloneqq \left(\frac{\nu(y_1)}{|a_1|}, \dots, \frac{\nu(y_r)}{|a_r|} \right),$$

where a_s is the coefficient at x in the equation tackled on this iteration (here $y_1, \dots, y_r$ are slack variables). Assume with no loss of generality that ν minimises $\min_{1 \le e \le r} q_e$ and this minimum is q_1. It suffices to show $q_1 < \mathrm{mod}(\varphi')$, where φ' is our system of equations at the beginning of the current iteration and $\mathrm{mod}(\cdot)$ is the (positive) least common multiple of all divisors appearing in the system. Equivalently, we need to show $\nu(y_1) < |a_1| \cdot \mathrm{mod}(\varphi')$.

Assume for the sake of contradiction that $q_1 \ge \mathrm{mod}(\varphi')$. Denote $a = a_1$ and let m be the coefficient at the slack variable y_1. Starting with ν, we construct a new assignment ν' that agrees with ν on all actual variables except x as well as on all slack variables eliminated previously. Suppose the tackled equation is $ax + my_1 + \tau' = 0$. We restrict our attention to the case $a > 0$. Choose

$$\begin{aligned} \nu'(x) &= \nu(x) + m \cdot \mathrm{mod}(\varphi'), \\ \nu'(y_1) &= \nu(y_1) - a \cdot \mathrm{mod}(\varphi') < \nu(y_1). \end{aligned}$$

Adding a multiple of $\mathrm{mod}(\varphi')$ ensures that all divisibility constraints are kept satisfied. Suppose that other equations have the form $a_e \cdot x + m \cdot y_e + \tau'' = 0$. Choose

$$\nu'(y_e) = \nu(y_e) \pm a_e \cdot \mathrm{mod}(\varphi') = |a_e| \cdot (q_e \pm \mathrm{mod}(\varphi')) \ge 0.$$

Here and above signs are chosen so that all the equations are satisfied and the inequality $\nu'(y_1) \ge 0$ holds. The chained inequality $0 \le \nu'(y_1) < \nu(y_1)$ is a contradiction as long as the divisibility constraints are satisfied and all slack is nonnegative. Note that $q_e \ge q_1$ by our assumption about q_1.

On Presburger arithmetic with exponentiation

We outline the main ideas behind a quantifier elimination procedure for Presburger arithmetic extended with the function $x \mapsto 2^{|x|}$, where $|\cdot|$ is the absolute value function on integers. We denote this logic by $\mathcal{Pres}\mathcal{Exp}$. The absolute value is just a device to avoid partially defined functions.

The first observation is that simple enough atomic formulas are easy to deal with. We call them *power comparisons:* for example, $2^{|x|} < 1000$, or $2^{|x|} > 5 \cdot 2^{|y|}$, or $2^{|x|} = 4 \pmod 7$. These can be linearised painlessly (and easily, as long as we are interested in proving decidability only).

A variable x is unproblematic if it only occurs linearly or if it only occurs in power comparisons. If all variables are unproblematic, we can use quantifier elimination procedures for Presburger arithmetic to handle these formulas from $\mathcal{Pres}\mathcal{Exp}$.

Consider a formula $\exists \vec{x}\ \varphi(\vec{x}, \vec{z})$, where φ is quantifier-free. Take a variable x from the quantifier block $\exists \vec{x}$. We call this variable *problematic* if it appears outside power comparisons and has an exponential occurrence.

Suppose $\vec{x} = (x, y)$; then the initial idea is to rewrite the formula $\exists \vec{x}\ \varphi(\vec{x}, \vec{z})$ with a big disjunction of formulas of the form $\exists x\ \exists y\ \varphi_i$ such that in each φ_i either x is unproblematic or y is unproblematic. We do not quite achieve this, but come close.

Theorem 5 (based on Semenov 1980)**.** *Take a formula $\exists x\ \varphi(\vec{x}, \vec{z})$ in $\mathcal{Pres}\mathcal{Exp}$, where all variables in the vector $\vec{x}$ are problematic. There exist finitely many formulas $\psi_i(\vec{x}, \vec{z}, \vec{w})$ in $\mathcal{Pres}\mathcal{Exp}$ such that:*

- $\exists \vec{x}\ \varphi(\vec{x}, \vec{z}) \Leftrightarrow \forall \vec{w}\ \exists \vec{x}\ \bigvee_i \psi_i(\vec{x}, \vec{z}, \vec{w})$,
- *variables $\vec{w}$ are fresh, and*
- *in every ψ_i some $2^{|x|}$ (for x within the vector $\vec{x}$) only appears in power comparisons.*

The proof of the theorem is effective, and based on this proof one can show that $\mathcal{Pres}\mathcal{Exp}$ is decidable. Intuitively, even though new quantified variables are introduced, they are only placed into the "next" (i.e., second innermost) quantifier block. In fact, if φ has no free variables ($\vec{z}$), then no new variables ($\exists \vec{w}$) are introduced either. This ensures that a suitable quantifier elimination–like procedure terminates.

We have sketched the main idea behind the proof of Theorem 5 in our Thursday lecture. We considered a single inequality involving $2^{|x|}$. By splitting into many cases, we can assume that $2^{|x|}$ dominates powers of other existentially quantified variables that occur in the same inequality. The inequality compares a linear combination of several powers, $2^{|u|}$ for variables u within the block $\exists \vec{x}$, to the value of a term that only depends on free variables, $\tau(\vec{z})$. (In the more general case, variables u may also occur linearly.)

We can compare the value of $2^{|x|}$ to the value of τ approximately. In $\mathcal{Pres}\mathcal{Exp}$, comparing *bit lengths* of these two values is an appropriate choice. For an integer N, denote by $\lambda(N)$ the highest power of 2 not exceeding $|N|$: $\lambda(0) = 0$ and, for all N, $\lambda(N) \leq |N| < 2\lambda(N)$. We split into several cases according to whether $2^{|x|} > \lambda(\tau)$, $2^{|x|} = \lambda(\tau)$, or $2^{|x|} < \lambda(\tau)$. We may need up to two additional cases to consider the

scenarios $2^{|x|}/\lambda(\tau) \in \{\frac{1}{2}, 2\}$ separately. If either of $2^{|x|}$ and $\lambda(\tau)$ is bigger than the other by a factor of 4 or more, the truth value of the inequality can determined unambiguously. And in the case of exact equality we may substitute $\lambda(\tau)$ for $2^{|x|}$, eliminating an occurrence of a problematic variable.

The new formula produced in this way uses the signature of $\mathscr{Pres}\mathscr{E}\boldsymbol{xp}$ extended with the function λ. But we can use the definition of λ in $\mathscr{Pres}\mathscr{E}\boldsymbol{xp}$ to our advantage: all occurrences of λ can be eliminated by introducing new variables. Indeed, we can replace $\lambda(\tau)$ with $2^{|w|}$ for a fresh w and assert that, for all values of w, *if* the chained inequality $2^{|w|} \leq |\tau| < 2 \cdot 2^{|w|}$ is satisfied, then the formula is true. But this is, in effect, just adding two more cases $2^{|w|} > |\tau|$ and $|\tau| \geq 2 \cdot 2^{|w|}$ to the big case split (big disjunction) and inserting a universal quantifier $\forall w$ outside $\exists \vec{x}$. The case $\lambda(\tau) = 0$ should also be considered separately.

References: [Benedikt, Chistikov, Mansutti 2023, Algo. 4 and Sec. 4].