BIRS 2012 – Nils Bruin: 1. Skolem's method

June 11, 2012

1.1 : Basic observations on Thue equations

Suppose $f(x,y) \in \mathbb{Z}[x,y]$ is homogeneous, square-free of degree d and let $c \in \mathbb{Z}$. We want to find the solutions of

$$f(x,y) = c \text{ for } x,y \in \mathbb{Z}. \tag{1}$$

Thue already proved that if $d \geq 3$ then there are only finitely many solutions. Let us assume that the coefficient f_d of x^d in $f(x,y)$ is non-zero (we can always ensure this is the case via an $\mathrm{GL}_2(\mathbb{Z})$-transformation on x,y, which preserves integrality of solutions). To avoid some technical complications, we assume that $f_d = 1$.

We consider the algebra $L = \mathbb{Q}[z]/(f(z,1)$ and denote θ for a root of $f(z,1)$ in L. If $f(z,1)$ is irreducible then L is a number field. Otherwise, by virtue of $f(z,1)$ being square-free, L is a product of number fields, corresponding to the irreducible factors. Nothing but generality is lost by limiting to the case where L is a number field.

We write $\mathcal{O}_L$ for the ring of integers of L. We have that $\mathcal{O}_L^\times = \mathcal{O}_{L,\mathrm{tors}}^\times \times \langle \epsilon_1, \ldots, \epsilon_r \rangle$, where $\mathcal{O}_{L,\mathrm{tors}}^\times$ is the finite subgroup of torsion units and $\epsilon_1, \ldots, \epsilon_r$ is a system of fundamental units.

The main observation for most approaches to Thue equations is that

$$f(x,y) = N_{L/\mathbb{Q}}(x - \theta y).$$

Thus, we are looking for $x - \theta y \in \mathcal{O}_L$ of norm c. It is straightforward to determine a finite number of elements $\gamma \in \mathcal{O}_L$ such that for any solution x,y there is a γ such that

$$x - \theta y = \gamma \epsilon_1^{n_1} \cdots \epsilon_r^{n_r}$$

We can expand the right hand side with respect to the $\mathbb{Q}$-basis $\{1, \theta, \ldots, \theta^{d-1}\}$ for L. We write $\mathbf{n} = (n_1, \ldots, n_r)$ and obtain

$$x - \theta y = Q_{0,\gamma}(\mathbf{n}) + Q_{1,\gamma}(\mathbf{n})\theta + \cdots + Q_{d-1,\gamma}(\mathbf{n})\theta^{d-1}.$$

Therefore, we can express x,y entirely in terms of $\mathbf{n}$ and obtain $d-2$ equations in $n_1, \ldots, n_r$, so if $r \leq d-2$, which only fails when L is a totally real number field, then it is not unreasonable to expect that these equations only have a finite number of solutions. Of course, the nature of the function $Q_{i,\gamma}(\mathbf{n})$ is unclear at this moment.

1.2 : Skolem's p-adic approach

Let $p > 2$ be a rational prime not dividing the discriminant of $f(z,1)$ or c. That means that $\mathcal{O}_L \otimes \mathbb{Z}_p = \mathbb{Z}_p[\theta]$, that $\mathcal{O}_L/p\mathcal{O}_L$ is a product of finite fields and that the elements γ we considered before are units in $\mathcal{O}_L \otimes \mathbb{Z}_p$.

We consider the reduction map

$$\mathcal{O}_L^\times \to (\mathcal{O}_L/p\mathcal{O}_L)^\times$$

a denote its kernel by $\Lambda_p = \langle \eta_1, \ldots, \eta_r \rangle$. This kernel is torsion-free and of finite index in $\mathcal{O}_L^\times$. Thus, at the expense of having to consider more values γ, it is sufficient to consider equations

$$\frac{x - \theta y}{\gamma} = \eta_1^{n_1} \cdots \eta_r^{n_r}.$$

In order to prove that (1) has only finitely many solutions, it suffices to prove that if there is a solution x_0, y_0 for γ, then there are only finitely many other solutions for that γ, since if there are no such solutions, we definitely have a finite number of them. So without loss of generality we can assume that $\gamma = x_0 + \theta y_0$.

Note that our conditions imply that such a solution would have to have the same image in $\mathcal{O}_L/p\mathcal{O}_L$, so such a solution would be of the form

$$(x_0 + px_1) + \theta(y_0 + py_1).$$

Thus, we are left with solving equations of the form

$$1 + p\frac{(x_1 - \theta y_1)}{(x_0 - \theta y_0)} = \eta_1^{n_1} \cdots \eta_r^{n_r}, \text{ with } x_0, y_0 \text{ given.}$$

Skolem's method hinges on the observation that even for $x_1, y_1, n_1, \ldots, n_r \in \mathbb{Z}_p$, such an equation has only finitely many solutions. Note that both sides are congruent to 1 modulo p, so they lie inside the radius of convergence of the p-adic power series

$$\operatorname{Log}(1 + z) = z - \frac{1}{2}z^2 + \frac{1}{3}z^3 + \cdots.$$

Taking logarithms of both sides yields

$$\operatorname{Log}\left(1 + p\frac{x_1 - \theta y_1}{x_0 - \theta y_0}\right) = n_1 \operatorname{Log}(\eta_1) + \cdots + n_r \operatorname{Log}(\eta_r).$$

which, when we expand with respect to the $\mathbb{Z}_p$-basis $\{1, \theta, \ldots, \theta^{d-1}\}$, gives us d equations, linear in $n_1, \ldots, n_r$ and power series in x_1, y_1. One can solve this system, but the fact that we are required to look at bivariate power series is slightly awkward. We define $\eta_0 = 1 + p$. Then η_0 is a one-unit in $\mathbb{Z}_p^\times$, i.e., a unit that is congruent to 1 modulo p. The multiplicative group of one-units $1 + p\mathbb{Z}_p$ is isomorphic to the additive group $\mathbb{Z}_p$, via $\operatorname{Log}(z)$, and η_0 is a $\mathbb{Z}_p$-generator of it. That means for any $\lambda \in \mathbb{Z}_p$ there is a $n_0 \in \mathbb{Z}_p$ such that

$$(1 + p\lambda) = \eta_0^{n_0}.$$

Thus, we can rewrite our original equation as

$$(1 + p\lambda)\left(1 + p\frac{x_1 - \theta y_1}{x_0 - \theta y_0}\right) = \eta_0^{n_0} \cdots \eta_r^{n_r}.$$

We see that the left hand side equals

$$1 + p\frac{x_1 + (x_0 + px_1)\lambda - \theta(y_1 + (y_0 + py_0)\lambda)}{x_0 - \theta y_0},$$

so assuming that $y_0 \not\equiv 0 \pmod{p}$, we can set

$$\lambda = -\frac{y_1}{y_0 + py_1} = \eta_0^{n_0}$$
$$t = x_1 + (x_0 + px_1)\lambda$$

Note that $x_1, y_1 \in \mathbb{Z}_p$ if and only if $\lambda, t \in \mathbb{Z}_p$ and substituting these values in we see that our equation becomes

$$\operatorname{Log}\left(1 + p\frac{t}{x_0 - \theta y_0}\right) = n_0 \operatorname{Log}(\eta_0) + \cdots + n_r \operatorname{Log}(\eta_r).$$

If $y_0 \equiv 0 \pmod{p}$ then we must have $x_0 \not\equiv 0 \pmod{p}$ and we can apply the same trick with the roles of the x_i and y_i swapped, to obtain

$$\operatorname{Log}\left(1 + p\frac{t\theta}{x_0 - \theta y_0}\right) = n_0 \operatorname{Log}(\eta_0) + \cdots + n_r \operatorname{Log}(\eta_r).$$

If we write

$$\mathrm{Log}(1+p\frac{t}{x_0-\theta y_0}) = L_0(t)+\theta L_1(t)+\cdots+\theta^{d-1}L_{d-1}(t) \text{ with } L_i(t)\in\mathbb{Z}_p[[t]]$$

and

$$\mathrm{Log}(\eta_j) = b_{0j}+b_{1j}\theta+\cdots+b_{d-1,j}\theta^{d-1} \text{ with } b_{ij}\in p\mathbb{Z}_p,$$

then we obtain a system of equations

$$\begin{pmatrix} b_{00} & \cdots & b_{0r}\\ b_{10} & \cdots & b_{1r}\\ \vdots & \ddots & \vdots\\ b_{d-1,0} & \cdots & b_{d-1,r}\end{pmatrix}\begin{pmatrix} n_0\\ \vdots\\ n_r\end{pmatrix} = \begin{pmatrix} L_0(t)\\ L_1(t)\\ \vdots\\ L_{d-1}(t)\end{pmatrix}$$

We see that if $r+1<d$, then we can compute a non-trivial $\mathbb{Q}_p$-linear relation between the $L_i(t)$ and hence probably a non-trivial power series equation for t. In fact, one can prove this equation *will* be non-trivial.

In nearly all cases, the following lemma suffices.

1.3 Lemma: Let $L(z)=\sum_{n=0}^{\infty} a_n z^n \in \mathbb{Z}_p[\![x]\!]$ be a power series with $\lim_{n\to\infty}\mathrm{ord}_p(a_n)=\infty$. If

$$L(z)\equiv a_0+a_1 z \pmod{p^m}$$

with $\mathrm{ord}_p(a_1)<m$, then $z=-a_0/a_1$ is the only possible root of $L(z)$ in $\mathbb{Z}_p$.

Proof. Straightforward Hensel lifting argument. □

1.4 Example: Consider $f(x,y)=x^3-2y^3=1$. Then $L=\mathbb{Q}(\theta)=\mathbb{Q}(\sqrt[3]{2})$, the unit rank is 1 and $\epsilon_1=\theta-1$. We consider $p=5$ and the solution $(x_0,y_0)=(-1,-1)$. Then $\gamma=\epsilon_1$ and we obtain the system

$$\begin{pmatrix} 55 & 0\\ 0 & 100\\ 0 & 10\end{pmatrix}\begin{pmatrix} n_0\\ n_1\end{pmatrix} \equiv \begin{pmatrix} 5t\\ 5t+75t^2\\ 5t+25t^2\end{pmatrix} \pmod{5^3},$$

leading to a power series equation in t approximated by

$$(5t+75t^2)-10(5t+25t^2)\equiv 80t+75t^2\equiv 0 \pmod{5^3}.$$

Modulo 5^2 we see that Lemma 1.3 applies. Thus we see that the only solution $x,y\in\mathbb{Z}$ to the equation $x^3-2y^3=1$ that has $(x,y)\equiv(-1,-1) \pmod 5$ is the solution $x_0,y_0=-1,-1$ itself.

1.5 : Dirichlet sieving

In the previous section we have seen a p-adic method that, given a solution $x_0,y_0\in\mathbb{Z}$ to a Thue equation $f(x,y)=c$, can in all likelyhood prove that there are no other such solutions that are congruent to it modulo p. We are left with formulating a method that can show that certain congruence classes do *not* contain a solution.

As we saw, we can determine a finite set Γ such that any solution x_0,y_0 is of the form

$$x_0-\theta y_0=\gamma\epsilon_1^{n_1}\cdots\epsilon_r^{n_r}.$$

We recall that we write $\Lambda_p\subset\mathbb{Z}^r$ for the kernel of the homomorphism

$$\begin{array}{ccc} \mathbb{Z}^r & \to & (\mathcal{O}_L/p\mathcal{O}_L)^\times\\ (n_1,\ldots,n_r) & \mapsto & \epsilon_1^{n_1}\cdots\epsilon_r^{n_r}\end{array}$$

By looking at the equation modulo p, we can determine a set $V_p\subset\mathbb{Z}^r/\lambda_p$ that contains the reduction of any solution. The set V_p will have about p^2 elements, so it likely contains congruence classes that do not contain actual solutions. However, notice that we can combine information from several primes. If $\Lambda_p + \Lambda_q \neq \mathbb{Z}^r$, then $V_p \cap V_q$ could actually consist of less cosets of $\Lambda_p \cap \Lambda_q$ than one would expect. On an industrial scale, one picks a set of suitable primes S and computes

$$\bigcap_{p \in S} V_p \subset \mathbb{Z}^r / (\bigcap_{p \in S} \Lambda_p).$$

The heuristic that for a suitably chosen set S, this intersection is likely very small, and hence likely only contains cosets that actually correspond to actual solutions, is based on the following observation.

Consider the commutative diagram

$$\begin{array}{ccc}
\{x, y \in \mathbb{Z} : x - \theta y \in \langle \epsilon_1, \ldots, \epsilon_r \rangle\} & \longrightarrow & \mathbb{Z}^r \\
\downarrow & & \downarrow \\
\prod_{p \in S} \{x, y \in \mathbb{F}_p : x - \theta y \in \langle \epsilon_1, \ldots, \epsilon_r \rangle \pmod{p}\} & \longrightarrow & \prod_{p \in S} \mathbb{Z}^r / \Lambda_p
\end{array}$$

The key is that the group $\prod_{p \in S} \mathbb{Z}^r / \Lambda_p$ is very far from cyclic if its components have many factors in common in their group orders, whereas the image of $\mathbb{Z}^r$ is of course only a subgroup generated by r generators.

In practice this method works extremely well.

1.6 Example: We return to our equation $f(x, y) = x^3 - y^3 = 1$. In this case, the only value for γ we need is $\gamma = 1$. We pick $p = 5$. We find

n	$(\theta - 1)^n \pmod{5}$
0	1
1	$\theta + 4$
2	$\theta^2 + 3\theta + 1$
3	$2\theta^2 + 3\theta + 1$
4	$\theta^2 + 3\theta + 3$
5	$2\theta^2 + 4$
6	$3\theta^2 + 4\theta$
7	$\theta^2 + \theta + 1$
8	1

We see that only for $n \equiv 0, 1 \pmod{5}$ we have that $(\theta - 1)^n$ is of the form $x - \theta y \pmod{5}$. This corresponds to the actual solutions $(x, y) = (1, 0), (-1, 1)$. So in this case the information at one prime allows us to limit only to the residue classes that contain actual solutions.

Had we made the less fortunate choice on $p = 11$, we would have found

$$\begin{aligned}
(\theta - 1)^0 &\equiv 1 && \pmod{11} \\
(\theta - 1)^1 &\equiv \theta - 1 && \pmod{11} \\
(\theta - 1)^{14} &\equiv 4\theta + 3 && \pmod{11} \\
(\theta - 1)^{19} &\equiv 6\theta + 4 && \pmod{11} \\
(\theta - 1)^{40} &\equiv 1 && \pmod{11}
\end{aligned}$$

However, combined with

$$\begin{aligned}
(\theta-1)^0 &\equiv 1 && \pmod{17}\\
(\theta-1)^1 &\equiv \theta-1 && \pmod{17}\\
(\theta-1)^{44} &\equiv 4\theta+15 && \pmod{17}\\
(\theta-1)^{64} &\equiv 15\theta && \pmod{17}\\
(\theta-1)^{81} &\equiv 2\theta+8 && \pmod{17}\\
(\theta-1)^{96} &\equiv 1 && \pmod{17}
\end{aligned}$$

we see that $\gcd(40,96)=8$. From $p=11$ we find that $n\equiv 0,1,6,5 \pmod 8$ and for $p=17$ we find that $n\equiv 0,1,4,0,1 \pmod 8$. Combined, we see that only $n=0,1 \pmod{96}$ need to be considered.

1.7 : Geometric interpretation

It is instructive to interpret Skolem's method in a geometric setting. A Thue equation is built from a *homogeneous* form. That suggests a projective variety playing a role. However, the equation itself is not homogeneous. Let us consider an example. Take the affine curve

$$C'\colon x^2y-xy^2+3xy+1=0$$

with projective closure

$$C\colon X^2Y-XY^2+3XYZ+Z^3=0.$$

It is clear what *integral* points on C' are: Points for which $x,y\in\mathbb{Z}$. On the projective curve C we can represent any rational point using integers because we can clear denominators. Thus, on a projective variety, integral and rational points are the same thing. We can recognize the integral points on C' from integral points on C, though: These are points that can be represented by integers $(X_0:Y_0:Z_0)$ with Z_0 a unit. Our particular example is a genus 0 curve, as shown by the parametrization

$$\begin{array}{ccc}\mathbb{P}^1 & \to & C\\ (U:V) & \to & (U^3:V^3:UV(U-V)).\end{array}$$

Thus, we see that the integral points on C' correspond to solutions to

$$f(U,V)=UV(U-V)=\pm 1 \text{ with } U,V\in\mathbb{Z}.$$

Note that C' is $C\setminus\{Z=0\}$. Since $C\simeq\mathbb{P}^1$ via the parametrization above, we find

$$C'\simeq\mathbb{P}^1\setminus\{UV(U-V)=0\}.$$

This applies in general: Solving Thue equations amounts to finding the integral points on projective lines minus points.

1.8 Multiplicative Groups: Note that $\mathbb{P}^1\setminus\{0,\infty\}=\mathbb{G}_m$, the multiplicative group. The integer points on $\mathbb{G}_m$ are the units of the base ring, by definition. One way to express this is as

$$\begin{array}{ccc}\mathbb{P}^1\setminus\{UV=0\} & \to & \dfrac{\mathbb{G}_m\times\mathbb{G}_m}{\mathbb{G}_m}\\ (U:V) & \mapsto & (U:V)\end{array}$$

With more points removed, we can map into a higher dimensional algebraic group

$$\begin{array}{ccc}\mathbb{P}^1\setminus\{UV(U-V)=0\} & \to & \dfrac{\mathbb{G}_m\times\mathbb{G}_m\times\mathbb{G}_m}{\mathbb{G}_m}\\ (U:V) & \mapsto & (U:V:U-V)\end{array}$$

1.9 Twisted tori: We assume that x^d has a non-zero coefficient in $f(x,y)$. Let $L = \mathbb{Q}[x]/(f(x,1)$ and let θ be the class of x in L. Then $\mathbb{G}_m(L) = L^\times$. We can make an algebraic group T over $\mathbb{Q}$ such that $T(\mathbb{Q}) \simeq \mathbb{G}_m(L)$ in the following way. We use that $\{1, \theta, \dots, \theta^{d-1}\}$ is a $\mathbb{Q}$-basis for L. Given two elements $\alpha = a_0 + a_1\theta + \cdots + a_{d-1}\theta^{d-1}$ and $\beta = b_0 + b_1\theta + \cdots + b_{d-1}\theta^{d-1}$, we can write out

$$\alpha\beta = c_0 + c_1\theta + \cdots + c_{d-1}\theta^{d-1}$$

with $c_i \in \mathbb{Q}[a_0, \dots, a_{d-1}, b_0, \dots, b_{d-1}]$. This gives us an algebraic group law, defined over $\mathbb{Q}$. Similarly, we can write out the norm form $F(a_0, \dots, a_{d-1}) = N(\alpha)$. As a variety T is $\mathbb{A}^d \setminus \{F = 0\}$.

The scalar inclusion $\mathbb{Q} \subset L$ is expressed as

$$\begin{array}{ccc} \mathbb{G}_m & \rightarrow & T \\ a & \mapsto & (a, 0, \dots, 0) \end{array}$$

Over L, we would have the map $\mathbb{P}^1 \setminus \{f(x,y) = 0\} \rightarrow \mathbb{G}_m$ defined by $(x : y) \mapsto x - \theta y$. This induces the map

$$\begin{array}{ccc} \mathbb{P}^1 \setminus \{f(x,y) = 0\} & \rightarrow & \dfrac{T}{\mathbb{G}_m} \\ (x : y) & \mapsto & (x : -y : 0 : \dots : 0) \end{array}$$

With a little work we can check that essentially the integer points from one side have to map to integer points on the other. Dirichlet's Theorem implies that the integer points on a torus form a finitely generated group.