Enumeration of Number Fields

John Voight
University of Vermont

Department Colloquium
University of Washington
19 February 2008

Abstract

How quickly can one enumerate number fields of fixed degree with bounded absolute discriminant? We discuss some mathematically and computationally interesting aspects of this question. For totally real number fields, a particular case of interest, we exhibit an algorithm which improves upon known methods by the use of elementary calculus (Rolle's theorem and Lagrange multipliers).

Problem

Today, we will be concerned with the following problem.

Problem. Let $B, n \in \mathbb{Z}_{>0}$. Compute the set $NF(n, B)$ of all number fields F of degree $n = [F : \mathbb{Q}]$ with absolute discriminant $|d_F| \leq B$, up to isomorphism.

The set $NF(n, B)$ is finite, as we shall see shortly.

(Number fields are represented in the usual way by a minimal polynomial of a primitive element, which is represented by a sequence of rational numbers.)

We consider this an algorithmic problem: we are interested in not only its theoretical complexity but also desire a practical program which takes as input the integers n, B and outputs the list of fields $NF(n, B)$.

Motivation

Problem. For $B, n \in \mathbb{Z}_{>0}$, compute

$$NF(n, B) = \{F : [F : \mathbb{Q}] = n,\ |d_F| \leq B\}.$$

This very basic problem arises very naturally.

Number fields are the raw material of number theory, and tables are useful to the extent that one views the subject as an experimental science. And indeed, there are still many mysteries to unfold!

My particular interest comes from arithmetic geometry, and the enumeration of Shimura curves of genus at most two.

But you might be interested for your own reasons!

Example

Example. If $n = 2$, then we are listing quadratic fields $\mathbb{Q}(\sqrt{d})$ (equivalently, polynomials $x^2 - d$) of discriminant d with $|d| \leq B$.

Of course, $0, 1 \neq d \in \mathbb{Z}$ is a such a discriminant if and only if

$$d \equiv 0, 1 \pmod 4 \text{ and } \frac{d}{\gcd(4, d)} \text{ is squarefree.}$$

Therefore,

$$\begin{aligned} NF(n, B) = \{\mathbb{Q}(\sqrt{d}) : d \equiv 2, 3 \pmod 4,\ d \text{ squarefree},\ 4d \leq B\} \\ \cup \{\mathbb{Q}(\sqrt{d}) : d \equiv 1 \pmod 4,\ d \text{ squarefree},\ d \leq B\}. \end{aligned}$$

Hence

$$\#NF(n, B) \sim 2 \cdot \frac{6}{\pi^2} \left(\frac{2}{3} \cdot \frac{B}{4} + \frac{1}{3} \cdot B \right) \sim \frac{6}{\pi^2} B, \quad \text{as } B \to \infty.$$

Problem, revisited

One may be also interested in specifying the signature $\mathrm{sig}(F) = (r_1, r_2)$ of F, where r_1, r_2 as usual denote the number of real, complex places of F (with $r_1 + 2r_2 = n$).

Problem. For $B, n \in \mathbb{Z}_{>0}$, compute

$$NF(n, B, \sigma) = \{F : [F : \mathbb{Q}] = n,\ \mathrm{sig}(F) = \sigma,\ |d_F| \leq B\}.$$

Analogously, one may also wish to specify the $\mathbb{Q}_p$-isomorphism class of the completion F_p for a finite set of primes p, or perhaps some aspect of the ramification or splitting behavior of p in F.

One may be interested in computing all number fields with bounded root discriminant $\delta_F = |d_F|^{1/[F:\mathbb{Q}]} \leq b$ (for b small enough to apply the Odlyzko bounds). Lastly, one may replace $\mathbb{Q}$ by an alternative base field E.

Our interest will be in totally real fields, i.e. $\mathrm{sig}(F) = (n, 0)$. Live demo!

Some previous work

- There are databases of number fields computed by the KASH group and the Pari group. These tables contain fields of all signatures with $n \leq 7$ and varying discriminant bound B. (In our application, we need n up to 10 and larger values of B.)

- Klüners and Malle have created a database for number fields, containing representative polynomials for all transitive Galois groups up to degree 15. (This is in a different spirit.)

- For small degrees, we have fast algorithms: Belabas (following Davenport-Heilbronn) has given a quasi-linear time algorithm for cubic fields; Cohen-Diaz y Diaz-Olivier use Kummer theory for quartic fields.

- A smattering of results (up to degree $n = 9$) list fields of a given degree and signature with smallest discriminant. For other small degrees (quintics, impritive sextics, ...), the construction of tables is described.

Many of these results and tables are at least ten years old!

General strategy

The general strategy to compute $NF(n, B)$ is well-known, and runs as follows. Throughout, let F be a number field with $[F : \mathbb{Q}] = n$, and let $\mathbb{Z}_F$ denote its ring of integers.

We use the geometry of numbers. We define the Minkowski norm

$$T_2 : F \to \mathbb{R}$$

$$\alpha \mapsto T_2(\alpha) = \sum_{i=1}^{n} |\sigma_i(\alpha)|^2 = \sum_{i=1}^{n} |\alpha_i|^2$$

where σ_i runs over the infinite places of F (equivalently, $\alpha_i = \sigma_i(\alpha) \in \mathbb{C}$ runs over the conjugates of α).

The norm T_2 gives $\mathbb{Z}_F$ the structure of a lattice of rank n with covolume (or determinant) $\sqrt{|d_F|}$.

General strategy

$\mathbb{Z}_F$ is a lattice of rank n under the T_2-norm, with determinant $\sqrt{|d_F|}$. A shortest lattice vector in $\mathbb{Z}_F$ is 1, and the sublattice $\mathbb{Z} \subset \mathbb{Z}_F$ is pure so we may consider the quotient lattice $\mathbb{Z}_F/\mathbb{Z}$. An application of Minkowski's theorem then yields the following result (known as Hunter's theorem).

Proposition. *There exists $\alpha \in \mathbb{Z}_F \setminus \mathbb{Z}$ such that $0 \leq \mathrm{Tr}(\alpha) \leq n/2$ and*

$$T_2(\alpha) \leq \frac{\mathrm{Tr}(\alpha)^2}{n} + \gamma_{n-1} \left(\frac{|d_F|}{n} \right)^{1/(n-1)}$$

where γ_{n-1} denotes the $(n-1)$st Hermite constant.

The Hermite constant measures "the densest lattice" in $\mathbb{R}^n$, defined by

$$\gamma_n = \max_{\substack{L \subset \mathbb{R}^n \\ \det(L)=1}} \min_{0 \neq v \in L} ||v||.$$

We know the values $\gamma_n^n = 1, 4/3, 2, 4, 8, 64/3, 64, 256$ for $n = 1, \ldots, 8$, given by the lattices $A_1, A_2, A_3, D_4, D_5, E_6, E_7, E_8$, respectively. The value of $\gamma_{24} = 4$ is given by the Leech lattice (Cohn-Kumar) and the best known upper bounds for γ_n are given by Cohn-Elkies.

General strategy

From Hunter's theorem, if $F \in NF(n, B)$, then one finds $\alpha \in \mathbb{Z}_F \setminus \mathbb{Z}$ such that $T_2(\alpha) = \sum_i |\alpha_i|^2 \leq C$ for some $C \in \mathbb{R}_{>0}$ depending only on n, B. Hence we obtain bounds on the power sums

$$|S_k(\alpha)| = \left|\sum_{i=1}^n \alpha_i^k\right| \leq T_k(\alpha) = \sum_{i=1}^n |\alpha_i|^k \leq nC^{k/2},$$

hence bounds on the coefficients $a_i \in \mathbb{Z}$ of the characteristic polynomial of α

$$f(x) = x^n + a_{n-1}x^{n-1} + \cdots + a_0 = \prod_{i=1}^n (x - \alpha_i)$$

by Newton's relations:

$$S_k + \sum_{i=1}^{k-1} a_{n-i} S_{k-i} + k a_{n-k} = 0.$$

This yields a finite set $NS(n, B)$ of possible $f(x) \in \mathbb{Z}[x]$.

If F is imprimitive, we may have $\mathbb{Q} \subset \mathbb{Q}(\alpha) \subsetneq F$ for α as given by Hunter's theorem. Using a relative version of Hunter's theorem (Martinet), one can proceed in an analogous manner.

Size of $NS(n, B)$

We have $S_k(\alpha) = \sum_{i=1}^{n} \alpha_i^k$ and the Newton relations

$$S_k + \sum_{i=1}^{k-1} a_{n-k+i} S_i + k a_{n-k} = 0$$

so that $S_1, \ldots, S_k$ determine a_{n-k}. A little work yields $|S_k(\alpha)| \leq C^{k/2}$.

Two other small improvements. An application of the arithmetic-geometric mean gives $|a_0| = |N_{F/\mathbb{Q}}(\alpha)| \leq (C/n)^{n/2}$. Next, according to the Newton relations we see that the value of S_k is determined modulo k by $S_1, \ldots, S_{k-1}$, hence there are at most $\left(2C^{k/2} + 1\right)/k$ values of S_k. In sum, we have

$$\begin{aligned} \#NS(n, B) &= O\left(\tfrac{n}{2}\left(\tfrac{2C}{2}\right) \cdots \left(\tfrac{2C^{(n-1)/2}}{n-1}\right)\left(\tfrac{C}{n}\right)^{n/2}\right) \\ &= O\left(\frac{2^n}{n^{(n-2)/2}(n-1)!} C^{(n-1)(n+2)/4}\right) = O\left(\left(\frac{2e}{n^{7/4}}\right)^n B^{(n+2)/4}\right) \end{aligned}$$

since $C = \frac{\mathrm{Tr}(\alpha)^2}{n} + \gamma_{n-1}\left(\frac{|d_F|}{n}\right)^{1/(n-1)} = O\left(\left(\frac{B}{n}\right)^{1/(n-1)}\right)$.

Note the size of $NS(n, B)$ is sharp for $n = 2$: we saw that $\#NF(n, B) \sim \frac{6}{\pi^2} B$ for $B \to \infty$.

Conjectural size of $NF(n, B)$

In summary: To compute $NF(n, B)$, we enumerate polynomials over a set $NS(n, B)$ which is of size $O(B^{(n+2)/4})$. What is the expected size of the output? We have the following folklore conjecture, which has been attributed to Linnik.

Conjecture. *For any $n \in \mathbb{Z}_{>0}$ and signature $\sigma = (r_1, r_2)$, there exists $c_{n,\sigma} \in \mathbb{R}_{>0}$ such that*

$$\#NF(n, B, \sigma) \sim c_{n,\sigma} B.$$

(One can also specify p-adic data and one makes a similar conjecture, or consider extensions of a fixed ground field E.)

This conjecture has been proven for $n = 3$ (Davenport-Heilbronn) and for $n = 4, 5$ (Bhargava). If one believes this conjecture, then the naïve algorithm runs in exponential time (as a function of n). For large n, the best known result is that $\#NF(n, B) = O\left(B^{\exp(C\sqrt{\log n})}\right)$ for some absolute constant C. This result may never in practice outperform the naïve method...? So for higher (intermediate) degrees, we seek to chip away at the implied constant.

Totally real fields: Bounds from Hunter's theorem

We now restrict to the case where F is primitive and totally real. We have, by Hunter's theorem, that

$$-\left\lfloor \frac{n}{2} \right\rfloor \leq a_{n-1} = -\operatorname{Tr}(\alpha) \leq 0$$

and combining Newton's relations with the fact that

$$S_2(\alpha) = \sum_i \alpha_i^2 = \sum_i |\alpha_i|^2 = T_2(\alpha)$$

we obtain the lower bound

$$a_{n-2} = \frac{1}{2}a_{n-1}^2 - \frac{1}{2}T_2(\alpha) \geq \frac{1}{2}\left(1 - \frac{1}{n}\right)a_{n-1}^2 - \frac{\gamma_{n-1}}{2}\left(\frac{B}{n}\right)^{1/(n-1)}.$$

We next need a good upper bound on a_{n-2}.

Totally real fields: Upper bounds for a_{n-2}

To find an upper bound for a_{n-2}, we note that $T_2(\alpha) = \mathrm{Tr}(\alpha^2)$ and that α^2 is a totally positive algebraic integer.

Theorem (Smyth). *If θ is a totally positive algebraic integer, then*

$$\mathrm{Tr}(\theta) > 1.7719[\mathbb{Q}(\theta) : \mathbb{Q}]$$

unless θ is a root of one of the following polynomials:

$$x - 1,\ x^2 - 3x + 1,\ x^3 - 5x^2 + 6x - 1,$$
$$x^4 - 7x^3 + 13x^2 - 7x + 1,\ x^4 - 7x^3 + 14x^2 - 8x + 1.$$

(N.B. $\delta = (3 \pm \sqrt{5})/2$ are the roots of the second polynomial; $2\cos(2\pi/7)+2$ and its conjugates are the roots of third; and the latter are quadratic extensions of $\mathbb{Q}(\delta)$.)

We conclude that (aside from these cases)

$$a_{n-2} < \tfrac{1}{2}a_{n-1}^2 - 0.88595n.$$

Totally real fields: Rolle's theorem

Now, given values for $a_{n-1}, a_{n-2}, \ldots, a_{n-k}$, we deduce bounds for a_{n-k-1}. Let $f_k(x) = \dfrac{f^{(n-k)}(x)}{(n-k)!} = g_k(x) + a_{n-k}$. Consider for illustration the case $k = 3$. Then

$$g_3(x) = \tfrac{n(n-1)(n-2)}{6}x^3 + \tfrac{(n-1)(n-2)}{2}a_{n-1}x^2 + (n-2)a_{n-2}x.$$

Let $\beta_1 < \beta_2$ denote the roots of $f_2(x)$.

Thus

$$f_3(\beta_1) = g_3(\beta_1) + a_{n-3} > 0$$

$$f_3(\beta_2) = g_3(\beta_2) + a_{n-3} < 0$$

hence $-g_3(\beta_1) < a_{n-3} < -g_3(\beta_2)$.

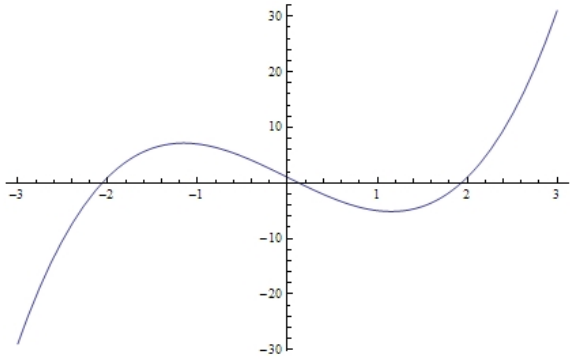

Totally real fields: Rolle's theorem

To obtain these bounds, we have applied Rolle's theorem, in the following form: If $g(x) \in \mathbb{R}[x]$ has distinct real roots, then the roots of $g'(x)$ are distinct and are interlaced with those of $g(x)$.

(The replacement when $\mathbb{R} = \mathbb{C}$ is the Gauss-Lucas theorem, which states that the roots of $g'(x)$ lie in the convex hull of the set of roots of $g(x)$. There is a weak p-adic replacement as well.)

In a similar way, if $\beta_1 < \cdots < \beta_{k-1}$ denote the roots of $f_{k-1}(x)$, then we find that

$$\max_{\substack{1 \le i \le k-1 \\ i \equiv k \,(2)}} -g_k(\beta_i) < a_{n-k} < \min_{\substack{1 \le i \le k-1 \\ i \not\equiv k \,(2)}} -g_k(\beta_i).$$

Computing these bounds requires a computation of real roots. By far, the fastest algorithm is to use Newton's method, since we have an interval where the root can lie. For our application, we need only a few significant digits (4 or 6 suffice), and then careful treatment of round-off error!

Totally real fields: Lagrange multipliers

We can obtain further bounds as follows. Let S denote the set of all totally real polynomials $f(x) \in \mathbb{R}[x]$ with specified coefficients $a_{n-1}, \ldots, a_{n-k-1}$ with $k \geq 3$. Note that since the roots of such f are bounded, the set $S \subset \mathbb{R}^n$ is closed and bounded, and hence there exists an f with the largest (resp. smallest) possible root, denoted β_k (resp. β_0).

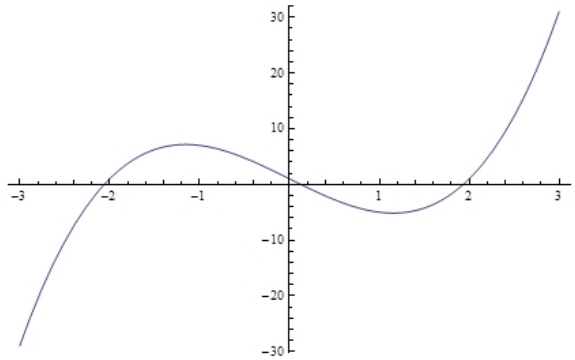

Then $f_k(\beta_k) = g_k(\beta_k) + a_{n-k} > 0$ with a similar inequality for β_0.

These combine with the above to yield neatly:

$$\max_{\substack{0 \leq i \leq k \\ i \equiv k\,(2)}} -g_k(\beta_i) < a_{n-k} < \min_{\substack{0 \leq i \leq k \\ i \not\equiv k\,(2)}} -g_k(\beta_i).$$

Totally real fields: Lagrange multipliers

One can compute the largest (resp. smallest) root β_k (resp. β_0) by an application of Lagrange multipliers.

We solve the following problem: Maximize α_n subject to

$$S_i(\alpha) = \sum_{i=1}^{n} \alpha_i = s_i$$

for $i = 1, \ldots, k-1$.

By the method of Lagrange multipliers, we find easily that if $(\alpha_i)_{i=1}^n \in \mathbb{R}^n$ yields such a maximum (resp. minimum), then there are at most $k-2$ distinct values among $\alpha_1, \ldots, \alpha_n$.

For example, in the case $k = 3$, we must solve the equations

$$(n-1)\alpha_1 + \alpha_n = s_1 = -a_{n-1} \quad \text{and} \quad (n-1)\alpha_1^2 + \alpha_n^2 = s_2 = a_{n-1}^2 - 2a_{n-2}$$

which yields simply

$$\beta_0, \beta_3 = -\frac{a_{n-1}}{n} \mp \frac{n-1}{n}\sqrt{a_{n-1}^2 - 2\left(1 + \frac{1}{n-1}\right)a_{n-2}}.$$

Totally real fields: Lagrange multipliers

Maximize α_n subject to $S_i(\alpha) = s_i$ for $i = 1, \ldots, k-1$: there are at most $k-2$ distinct values among $\alpha_1, \ldots, \alpha_n$.

For $k = 4$, for each partition $r + s = n - 1$, one obtains a system of equations which via elimination theory yield a (quite lengthy) degree 6 equation for α_n. For $k \geq 5$, we can continue in this way but instead must solve the system numerically; in practice, we do not significantly improve on the bounds for β_0, β_k whenever $k \geq 6$.

The method of Lagrange multipliers originates with Pohst and applies to all signatures in a different form. The combined Rolle's theorem-Lagrange multiplier bounds are by far the best available in the case of totally real fields, in our experience.

Implementation: Some issues

Speed: Speed is of the absolute essence—we are doing basic operations zillions of times. Maple and Mathematica cannot compete for this reason.

Number field arithmetic: Bare C lacks number field arithmetic. Pari and Magma (based on KANT) do: For "large examples", Magma is by far the quickest (in our experience). However, in our situation, our polynomials are small and both systems have highly optimized algorithms for testing irreducibility and computing p-maximal orders.

Real root finding: Magma also has amazingly fast real root finding; it is, however, a black box, and there is no fast Newton's method! Pari only computes real roots and lets you to guess if a root is real or not.

Optimized code: So it is up to us to implement a lightning fast Newton's method! One can work with the Pari library in C to write optimized code. In Magma, one cannot do this very easily without traveling to Sydney!

Cost: Magma is not free, hence not on UVM machines. Pari is free.

Intuitive: Pari can often be counterintuitive to use and to program.

Implementation: Sage

Having seemingly eliminated every alternative, we turn to Sage.

Sage includes Pari, so it has number field arithmetic. It uses Python, which is a very friendly modern (object-oriented) programming language. It is free. It incorporates Cython (Behnel-Bradshaw-Stein), which easily allows one to write optimized C code for repeated tasks.

Despite being a relatively new system (so that some functionality is limited), since Sage is open source it is easy to contribute yourself! Even though one must think about issues like how best to coerce between a C `int`, a Python integer, and a Sage `Integer`, there is a very active and generous development community willing to help!

It has the further advantage that there is a package for distributed computing called DSage (Qiang). Since the problem of computing $NF(n, B)$ is proudly parallelizable, a distributed computation on 30 Windows machines in a lab at UVM using DSage is currently underway (at night and on weekends!).

Data: Comparisons

We performed further tests on `sage.math.washington.edu`. We computed $\#NF(6, 15^6, (6, 0)) = 458$ in Magma, Pari, and Sage. The Sage code, which has seen the most optimization, runs in time 5m5s. The Magma code was not optimized any further after we discovered the real roots issue, so this is not a fair comparison: we aborted it after a half hour. The Pari code (run from GP—running it from gp2c did not give any significant advantage), ran in time 16m9s, a similarly unfair comparison.

We can recreate the list of 154 totally real fields in $NF(7, 150 \cdot 10^6, (7, 0))$, which appears in the tables of the KASH group, in just 44m51s.

For nonic fields, for comparison purposes we computed the 2 primitive totally real fields in $NF(9, 12 \cdot 10^9, (9, 0))$ in about a day and a half. This compares well to the previous computation of these fields (Takeuchi), which was reported to take one and a half months (in 1999)! Even this computation took only an hour when distributed!

Thanks!

Thanks to: Jürgen Klüners, Noam Elkies, Claus Fieker, and David Dummit for useful discussions; William Stein, Robert Bradshaw, Craig Citro, Yi Qiang, and the rest of the Sage development team for computational support; and Larry Kost and Helen Read for their technical assistance in the UVM lab.

And thank you for your attention!