Large-scale verification of Vandiver's conjecture

David Harvey
(joint work with Joe Buhler)

December 5, 2008

Plan for the talk

- Number-theoretic background and heuristics
  (Excellent reference: Washington's *Cyclotomic Fields*.)
- Some algorithms
- Software and hardware

Number-theoretic background and heuristics

Notation

p = an odd prime

ζ = primitive p-th root of unity

$K = \mathbf{Q}(\zeta)$

$K^+ = \mathbf{Q}(\zeta) \cap \mathbf{R} = \mathbf{Q}(\zeta + \zeta^{-1})$ = maximal real subfield of K

A, A^+ = class groups of K, K^+

A_p, A_p^+ = p-parts of A, A^+

h, h^+, h_p, h_p^+ = orders of A, A^+, A_p, A_p^+

$G = \mathrm{Gal}(K/\mathbf{Q}) \cong (\mathbf{Z}/p\mathbf{Z})^\times$

$\sigma_a = (\zeta \mapsto \zeta^a) \in G$ for $a \in (\mathbf{Z}/p\mathbf{Z})^\times$.

Vandiver's conjecture

Vandiver's conjecture asserts that $h_p^+ = 1$ for all p.
(Equivalently $p \nmid h^+$.)

Also known as the Kummer–Vandiver conjecture. Arose in connection with early work on Fermat's last theorem.

Kummer verified it by hand for $p < 200$.

Vandiver verified it with a desk calculator up to about 600.

Lehmer verified it up to about 5000 in the late 1940s (one of the first pure mathematics calculations performed on an electronic computer).

$\vdots$

Most recent is Buhler et al (2001), verified up to 12,000,000.

Vandiver's conjecture

Current project (joint work with Joe Buhler):

- Aim: check it for all $p < 39 \cdot 2^{22} = 163{,}577{,}856$.
- Done so far: verified completely up to $2^{27} = 134{,}217{,}728$.
- For $p < 163{,}577{,}856$, have done the hard part (computing the 'irregular indices'), Vandiver verification is in progress.

The cost to verify up to X with state-of-the-art algorithms is about $O(X^2 \log X)$, so this computation is about 200 times larger than the 2001 attempt.

I'll say more about the computation later.

Naïve heuristics

Suppose that h^+ is "uniformly distributed" modulo p. Then

$$\#\{\text{counterexamples} \leq X\} \approx \sum_{p \leq X} \frac{1}{p} \approx \log\log X.$$

Maybe this accounts for not seeing any counterexamples yet.

But "uniformly distributed" is a dangerous assumption. For the whole class group (not just the plus part) there is good empirical evidence that $p|h$ about 39.35% $(= 1 - e^{-1/2})$ of the time.

We can (heuristically) explain this behaviour by studying the $\mathbf{Z}_p[G]$-module structure of A_p.

Galois module structure of A_p

Decompose

$$A_p = \bigoplus_{i=0}^{p-2} e_i A_p$$

according to the orthogonal idempotents

$$e_i = \frac{1}{p-1} \sum_{a=1}^{p-1} \omega^i(a) \sigma_a^{-1} \in \mathbf{Z}_p[G], \qquad 0 \le i \le p-2,$$

where $\omega : (\mathbf{Z}/p\mathbf{Z})^\times \to \mathbf{Z}_p^\times$ is the Teichmüller character.

Note: $e_i A_p$ is the submodule where σ_a acts as $\omega^i(a)$ for all a.

Galois module structure of A_p

Fact: $e_0 A_p = e_1 A_p = 0$.

For remaining *odd* eigenspaces, have Ribet's theorem:

$$e_i A_p \neq 0 \iff p \mid B_{p-i}, \qquad i = 3, 5, \ldots, p-2,$$

where B_k is the k-th Bernoulli number.

p is called *irregular* if $p \mid B_k$ for some $k = 2, 4, \ldots, p-3$.

Such an integer k is called an *irregular index* for p; by Ribet's theorem these correspond precisely to the non-trivial odd eigenspaces of A_p.

The *index of irregularity*, $i(p)$, is the number of irregular indices that p has (number of non-trivial odd eigenspaces).

Galois module structure of A_p

Vandiver's conjecture concerns the *even* eigenspaces; it claims that

$$e_i A_p = 0, \qquad i = 2, 4, \ldots, p-3.$$

Note: the odd and even eigenspaces are related by a *reflection theorem*. If i is even, then

$$\dim_p(e_i A_p) \leq \dim_p(e_{p-i} A_p) \leq 1 + \dim_p(e_i A_p).$$

For example, if Vandiver's conjecture is true, then the odd eigenspaces must have p-rank ≤ 1.

Irregular primes (examples)

The smallest irregular prime is $p = 37$. We have

$$37 \mid B_{32} = \frac{-7709321041217}{510},$$

so $k = 32$ is an irregular index for 37, and in fact $i(37) = 1$. Ribet's theorem implies that $e_5 A_{37} \neq 0$.

The largest known $i(p)$ is 7, which first occurs for $p = 3{,}238{,}481$. Ribet's theorem says that the p-rank of A_p is at least 7.

Example Sage session

Construct $\mathbf{Q}(\zeta_{37})$ and compute class group:

```
sage: proof.number_field(False)        # assume GRH
sage: K.<zeta> = CyclotomicField(37)
sage: G = K.class_group()              # about 3 minutes (via PARI)
sage: G.order()
37
```

Let J be a non-principal ideal:

```
sage: J = G.gen().ideal(); J
Fractional ideal (94351, zeta - 40856)
sage: J.is_principal()
False
```

Consider the image of J under σ_{20}:

```
sage: sigmaJ = K.ideal(94351, zeta^20 - 40856); sigmaJ
Fractional ideal (94351, zeta + 16284)
```

Example Sage session

By Ribet's theorem, J must lie in $e_5 A_{37}$, so σ_{20} should act on J as multiplication by $20^5 \equiv 18 \bmod 37$. We should have

$$(\sigma_{20}(J))^2 J \sim (J^{18})^2 J \sim (1).$$

Let's check it:

```
sage: L = sigmaJ * sigmaJ * J; L
Fractional ideal (zeta^35 + zeta^33 + zeta^32 + zeta^29 + zeta^28 +
                 2*zeta^27 + zeta^26 + zeta^25 + 2*zeta^24 + zeta^23 +
                 zeta^21 - zeta^19 - zeta^17 + zeta^15 - zeta^14 +
                 zeta^12 + zeta^11 + zeta^10 + zeta^9 + zeta^7 +
                 zeta^6 + zeta^4 + 2*zeta + 1)

sage: L.is_principal()
True
```

Heuristics for irregular primes

Assume that B_k is "uniformly distributed" modulo p (for k even), i.e. is divisible by p with probability $1/p$.

Then

$$P(i(p) = r) = \binom{\frac{1}{2}(p-3)}{r} \left(1 - \frac{1}{p}\right)^{\frac{1}{2}(p-3)-r} \left(\frac{1}{p}\right)^{r}$$
$$\to \frac{e^{-1/2}}{2^r r!} \text{ as } p \to \infty.$$

Poisson distribution with parameter $1/2$.

Heuristics for irregular primes

Empirical data strongly supports the Poisson hypothesis (but we can't even prove there are infinitely many regular primes!):

$i(p)$	$\#p$	fraction	Poisson prediction
0	5,559,267	0.6066532	0.6065307
1	2,779,293	0.3032894	0.3032653
2	694,218	0.0757563	0.0758163
3	115,060	0.0125559	0.0126361
4	14,425	0.0015741	0.0015795
5	1,451	0.0001583	0.0001580
6	112	0.0000122	0.0000132
7	5	0.0000005	0.0000009

Table: Irregularity statistics for $p < 163{,}577{,}856$

Cyclotomic units

The best way to verify Vandiver's conjecture for a single p is via the *cyclotomic units* of K.

Let E, E^+ be the unit groups of K, K^+.

Let $C^+ \subseteq E^+$ be the group of *real cyclotomic units*. It is generated by elements of the form

$$\zeta^{\frac{(1-a)}{2}} \frac{1-\zeta^a}{1-\zeta} = \frac{\sin(\pi a/p)}{\sin(\pi/p)}, \qquad 1 \leq a \leq p-1.$$

Cyclotomic units

Fact: C^+ is of finite index of E^+, and $h^+ = [E^+ : C^+]$.

Vandiver's conjecture is equivalent to the statement that the p-part of E^+/C^+ is trivial.

Note: A^+ is not in general isomorphic to E^+/C^+ as Galois modules. It is unknown whether they are always isomorphic as abelian groups (according to Washington). For the p-parts, it is known that equality of orders holds for each eigenspace (of course Vandiver claims they are all trivial!).

Structure of E^+

Dirichlet's unit theorem $\implies$ $\text{rank}_{\mathbf{Z}} E^+ = (p-3)/2$.

Let $E_p^+ = \mathbf{Z}_p \otimes E^+$.

As a $\mathbf{Z}_p[G]$-module, we have the decomposition

$$E_p^+ = \bigoplus_{\substack{i=2 \\ i \text{ even}}}^{p-3} e_i E_p^+,$$

where each $e_i E_p^+ \cong \mathbf{Z}_p$.

Structure of E^+

The cyclotomic units can be used to explicitly write down elements of each component $e_i E_p^+$.

Let $g \in (\mathbf{Z}/p\mathbf{Z})^\times$ be a primitive root, and let

$$S_i = \prod_{a=1}^{p-1} \left(\zeta^{(1-g)/2} \frac{1-\zeta^g}{1-\zeta} \right)^{\omega(a)^i \sigma_a^{-1}} \quad \in e_i E_p^+.$$

Then S_i is a p-adic limit of cyclotomic units, and is non-trivial (the latter depends on the fact that $L_p(1, \omega^i) \neq 0$).

However, S_i might not *generate* $e_i E_p^+ \cong \mathbf{Z}_p$; it might lie in $p\mathbf{Z}_p$.

Vandiver's conjecture $\iff$ each S_i *does* generate $e_i E_p^+$.

More heuristics

This suggests another heuristic: suppose that S_i lies in $p\mathbf{Z}_p$ with probability $1/p$ for each i.

There are $(p-3)/2$ indices to choose from. We obtain the same Poisson distribution as before, so Vandiver's conjecture should fail for about 39.35% of primes!

Obviously this heuristic is broken. There must be an obstruction...

More heuristics

Fact: if $S_i \in p\mathbf{Z}_p$, then $p \mid B_i$. (Related to reflection theorem.)

So there are only $i(p)$ (not $\frac{p-3}{2}$) chances for S_i to lie in $p\mathbf{Z}_p$.

Assuming this is the only obstruction, the number of Vandiver counterexamples $\leq X$ should be about

$$\begin{aligned}
&\sum_{p \leq X} \sum_{r=0}^{\infty} P(i(p) = r) \times P(\text{some } S_i \in e_i E_p^+) \\
&= \sum_{p \leq X} \sum_{r=0}^{\infty} \left(\frac{e^{-1/2}}{2^r r!} \right) \left(1 - \left(1 - \frac{1}{p} \right)^r \right) \\
&= \sum_{p \leq X} 1 - e^{\frac{-1}{2p}} \approx \sum_{p \leq X} \frac{1}{2p} \\
&\sim \frac{1}{2} \log \log X.
\end{aligned}$$

More heuristics

For example:

- About 1.396 counterexamples less than 12,000,000.
- About 1.467 counterexamples less than 163,577,856.

Chance of success for current project was maybe 7%.

Actually it's worse than it looks, since the first few (regular) primes account for the bulk of those estimates.

Taking into account the actual values of $i(p)$ for each p, we obtain an estimate of 0.748 counterexamples for $p < 163{,}577{,}856$.

Some unreasonable extrapolations

On average, expect *one* counterexample for $p < 10^{14}$.

Moore's law $\implies$ get to 10^{14} by about 2084 AD (I will be 104).

Need about 1 petabyte (1 million gigabytes) memory to handle a single prime in this range.

Expect *two* counterexamples for $p < 10^{100}$.

Moore's law $\implies$ get to 10^{100} in 1000 years.

The universe has insufficiently many particles to satisfy memory requirements of current algorithms.

Some algorithms

Some algorithms

Two steps to verify Vandiver's conjecture for given p:

1. Compute $B_0, B_2, \ldots, B_{p-3}$ modulo p, to locate the irregular indices for p.
2. For each irregular index k, check whether S_k is a p-th power in $e_k E_p^+$.

Step 1 is *much* more expensive than step 2.

(Along the way we also check other cyclotomic invariants, in particular that each nontrivial $e_i A_p$ is no bigger than $\mathbf{Z}/p\mathbf{Z}$, i.e. that A_p is the smallest it can be consistent with $i(p)$.)

Computing Bernoulli numbers modulo p

Two methods for computing $B_0, B_2, \ldots, B_{p-3}$ modulo p:

- The "power series method".
- The "Voronoi congruence method".

Both have complexity $O(p \log^2 p)$ (ignoring $\log \log p$ terms).

But different implied constants and memory usage.

The power series method

Simplest version: use the identity

$$\frac{x}{e^x - 1} = \sum_{k \geq 0} \frac{B_k}{k!} x^k.$$

Uses a single power series inversion over $\mathbf{Z}/p\mathbf{Z}$ of length $\sim p$.

Fast power series arithmetic yields running time $O(p \log^2 p)$.

(Pre-1990 algorithms used some recurrence like

$$B_n = -\frac{1}{n+1} \sum_{k=0}^{n-1} \binom{n+1}{k} B_k,$$

and computed sequentially $B_2, B_4, B_6, \ldots$ in time $O(p^2)$.)

The power series method

There are redundancies, e.g. $B_k = 0$ for $k = 3, 5, \ldots, p-2$. Can exploit this via identities like

$$\frac{x^2}{\cosh x - 1} = -2 + \sum_{k=0}^{\infty} \frac{(2n-1)B_{2n}}{(2n)!} x^{2n}.$$

Only need power series inversion of length $\sim p/2$.

More sophisticated 'multisectioning' versions exist. We used one that involves:

- One series inversion of length $\sim p/8$.
- Four series multiplications of length $\sim p/8$.

This strategy saves a lot of memory (and possibly time, but this is unclear...)

The Voronoi congruence method

Let $g \in \mathbf{Z}/p\mathbf{Z}$ be a primitive root, and let

$$h(x) = \left\{ \frac{x}{p} \right\} - g \left\{ \frac{g^{-1}x}{p} \right\} + \frac{g-1}{2}.$$

Use the following identity:

$$B_{2k} \equiv \frac{4k}{1 - g^{2k}} \sum_{j=0}^{(p-3)/2} g^{2jk} \frac{h(g^j)}{g^j} \pmod{p}.$$

This may be interpreted as a DFT (number-theoretic transform) of the function $j \mapsto h(g^j)/g^j$ over $\mathbf{Z}/p\mathbf{Z}$.

Use Bluestein's FFT algorithm to convert the DFT to a single polynomial multiplication of length $\sim p/2$ over $\mathbf{Z}/p\mathbf{Z}$.

The Voronoi congruence method

This method has a fairly low constant in the running time, but 'multisectioning' opportunities only available when $\frac{1}{2}(p-1)$ has small factors. Memory constraints rule this out for large enough p.

The Voronoi congruence is also useful for verification purposes; can evaluate $B_k \pmod{p}$ for one pair (p, k) in time $O(p)$. We store several B_k for each p from the main computation, and then check them with an independent implementation on different hardware later.

(Spinoff project: one can compute B_k as an exact rational number, using only modular information, faster than the usual 'zeta function algorithm', using this type of formula.)

Verifying Vandiver's conjecture

Suppose k is an irregular index for p (i.e. $p \mid B_k$). Recall that

$$S_k = \prod_{a=1}^{p-1} \left(\zeta^{(1-g)/2} \frac{1-\zeta^g}{1-\zeta} \right)^{\omega(a^{-1})^k \sigma_a}.$$

To test whether S_k is a p-th power, we can approximate modulo $(E_p^+)^p$, and consider only

$$S_k^* = \prod_{a=1}^{p-1} \left(\zeta^{a(1-g)/2} \frac{1-\zeta^{ag}}{1-\zeta^a} \right)^{a^{p-1-k}},$$

which is now just a cyclotomic unit in K^+.

Verifying Vandiver's conjecture

To test whether S_k^* is a p-th power, we choose some degree 1 prime ideal $\tilde{\ell}$ in K and check whether S_k^* is a p-th power in $\mathcal{O}_K/\tilde{\ell}$.

This corresponds to choosing a prime $\ell \equiv 1 \pmod{p}$, choosing a p-th root of unity $t \in \mathbf{Z}/\ell\mathbf{Z}$, and then checking whether

$$\prod_{a=1}^{p-1} \left(t^{a(1-g)/2} \frac{1 - t^{ag}}{1 - t^a} \right)^{a^{p-1-k}}$$

is a p-th power in $\mathbf{Z}/\ell\mathbf{Z}$. Very simple test, involving only rational arithmetic.

If it is not a p-th power, then Vandiver holds for this eigenspace.

If it *is* a p-th power, we could try a different ℓ — but so far this has never been necessary.

Software and hardware

The software

The most expensive part of the computation is finding the Bernoulli numbers modulo p (to obtain the irregular indices).

This boils down to fast polynomial arithmetic in $\mathbf{Z}/p\mathbf{Z}[x]$ — in particular polynomial multiplication and series inversion.

To make best use of the 64-bit processor, we do everything modulo two primes simultaneously ($27 + 27 < 64$).

Parallelisation was handled with a simple MPI program (two primes per task).

zn_poly

We used the zn_poly polynomial arithmetic library:

- A C library, released under GPL
- Available from `http://cims.nyu.edu/~harvey/zn_poly/`
- Under development for about a year
- Supports any modulus that fits into an unsigned long (performance is best for odd moduli)
- Good support for multiplication, series inversion, middle products in high degree case
- Automatically tuned thresholds for all algorithms
- Under heavy development, lots of obvious things still missing

zn_poly multiplication performance

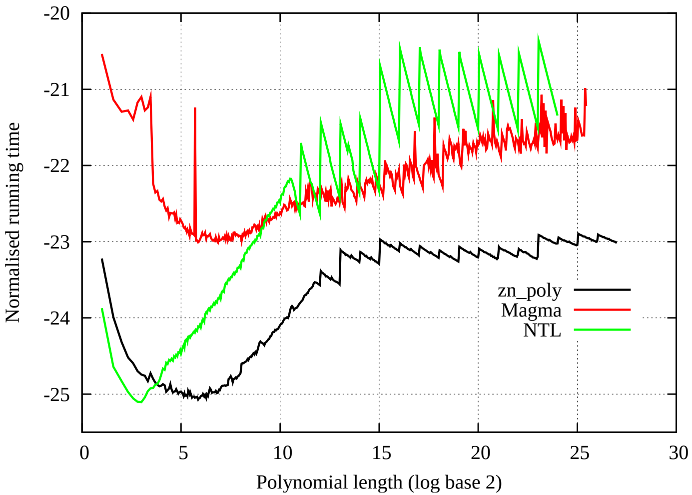

Figure: Multiplication of polynomials modulo a 48-bit modulus (Opteron)

Hardware

Small-to-medium machines:

- My laptop (2 $\times$ 2.0GHz Core 2 Duo, 1GB RAM)
- sage.math @ UW (16 $\times$ 1.8GHz Opteron, 64GB RAM)
- alhambra @ Harvard (16 $\times$ 2.6GHz Opteron, 96GB RAM)
- Joe Buhler's cluster (20 $\times$ 3.4GHz Pentium 4, 1GB RAM each)

Hardware (to scale)

My laptop

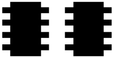

sage.math

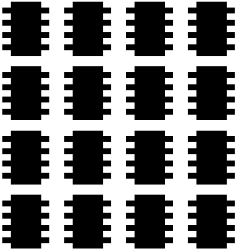

More hardware

Also used some *slightly larger* machines at TACC (Texas Advanced Computing Center, University of Texas, thanks to Fernando Rodriguez Villegas):

- Lonestar: 1300 nodes.
  - Each node $= 4 \times 2.66$GHz Xeon (Woodcrest), 8GB RAM.
  - Total cores $= 5200$, total RAM $= 10$ TB.
  - We used ≈ 119000 core-hours.

More hardware (to scale)

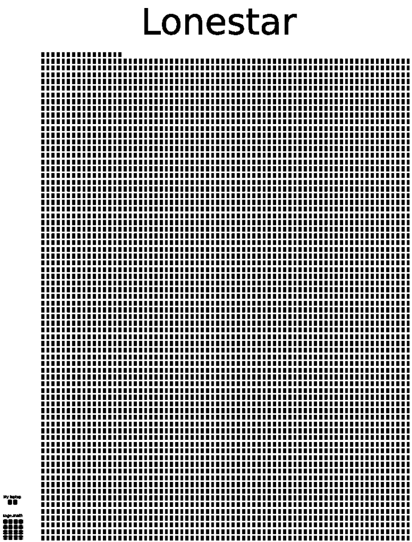

Hardware

- Ranger: 3936 nodes.
  - Each node = 16×2.3GHz Opteron (Barcelona), 32GB RAM.
  - Total cores = 62976, total RAM = 123 TB.
  - 4th most powerful computer worldwide in June 2008.
  - We used ≈ 69000 core-hours.

Hardware (to scale)

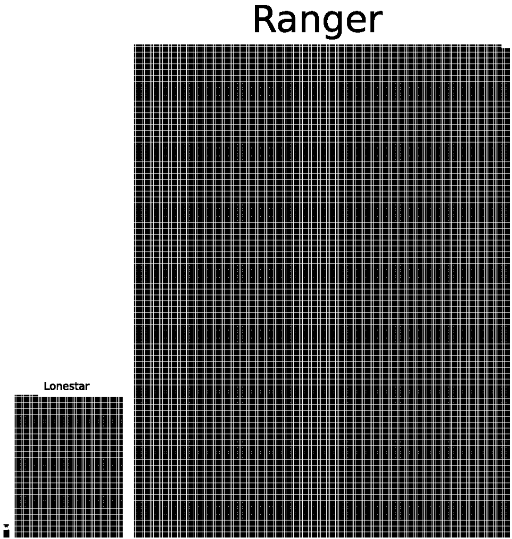

Hardware

About **21 core-years** altogether.

On both machines, have 2 GB RAM per core. If $p \approx 163{,}577{,}856$, one polynomial of length $p/2$ occupies 0.6 GB. Managing memory was the biggest challenge of the computation.

Thank you!