# Large-scale verification of Vandiver's conjecture

David Harvey

November 9, 2008

# Plan for the talk

- Number-theoretic background
  (Excellent reference: Washington's *Cyclotomic Fields*.)
- Some algorithms
- The software
- The hardware

# Number-theoretic background

# Notation

$p$ = an odd prime

$\zeta$ = primitive $p$-th root of unity

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

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

$A$, $A^+$ = 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 $p$.

Also known as the Kummer–Vandiver conjecture.

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 a 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 about 88,080,384.
- For $p < 163{,}577{,}856$, have done the hard part (computing the 'irregular indices'), haven't verified Vandiver yet.

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

I'll say more about this computation later.

# Naive heuristics

Suppose that $h_p^+$ 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 example there is good empirical evidence that $h_p \neq 1$ about 39.35% $(= 1 - e^{-1/2})$ of the time.

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

# Galois module structure of $A_p$

Decompose $A_p$ according to the orthogonal idempotents

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

where $\omega : (\mathbf{Z}/p\mathbf{Z})^\times \to \mathbf{Z}_p^\times$ is the Teichmuller character (lifts $a$ to a root of unity $\omega(a) \equiv a \pmod{p}$).

Obtain the decomposition $A_p = \bigoplus_{i=0}^{p-2} \varepsilon_i A_p$.

# Galois module structure of $A_p$

- $\varepsilon_0 A_p = \varepsilon_1 A_p = 0$.
- Ribet's theorem:

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

  where $B_k$ is the $k$-th Bernoulli number.
- Vandiver's conjecture is equivalent to

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

- The odd and even eigenspaces are related by a *reflection theorem*. If $i$ is even, then

$$\dim_p(\varepsilon_i A_p) \leq \dim_p(\varepsilon_{p-i} A_p) \leq 1 + \dim_p(\varepsilon_i A_p).$$

# Irregular primes

$p$ 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 $p$.

The *index of irregularity*, $i(p)$, is the number of irregular indices that $p$ has.

Ribet's theorem says that the non-trivial components $\varepsilon_i A_p$ (for odd $i$) correspond precisely to the irregular indices for $p$.

# 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 $\varepsilon_5 A_{37} \neq 0$.

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.

# Obligatory example Sage session

Let $J$ be an non-principal ideal of $\mathbf{Q}(\zeta_{37})$. Then the class of $J$ must lie in $\varepsilon_5 A_{37}$, and $J^{37} \sim (1)$. We should have

$$(\sigma_{20}(J))^2 J \sim (J^{20^5})^2 J \sim (1).$$

since $2 \times 20^5 \equiv -1 \bmod 37$. Let's check it:

```
sage: proof.number_field(False)
sage: K.<z> = CyclotomicField(37)
sage: G = K.class_group()       # about 2 minutes

sage: J = G.gen().ideal(); J
Fractional ideal (94351, z - 40856)

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

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

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

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

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

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 $E^+/C^+$!)

# Structure of $E^+$

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

Decompose $E_p^+$ as a $\mathbf{Z}_p[G]$-module; it turns out that

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

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

(This is consistent with Dirichlet's unit theorem, which says that $\operatorname{rank}_{\mathbf{Z}} E^+ = (p-3)/2$.)

# Structure of $E^+$

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

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

Let

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

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* $\varepsilon_i E_p^+ \cong \mathbf{Z}_p$; it might lie in $p\mathbf{Z}_p$.

Vandiver's conjecture is equivalent to the statement that each $S_i$ *does* generate $\varepsilon_i E_p^+$.

# 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 a Poisson distribution again...

... so Vandiver's conjecture should fail for a (fairly large) positive proportion of primes!

This conclusion seems unlikely given the numerical evidence.

# More heuristics

However, there is an obstruction.

Fact: if $S_i \in p\mathbf{Z}_p$, then $p \mid B_i$.

Taking this into account, the number of counterexamples $\leq X$ should be about

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

# 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 is 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 $p < 163{,}577{,}856$.

# More heuristics (trust me, I'm a mathematician)

One average, expect *one* counterexample before $10^{14}$.

TACC's archival storage facility (1 petabyte) can barely store a single polynomial for this computation.

Moore's law $\implies$ get to $10^{14}$ by about 2084 AD.

Expect *two* counterexamples before $10^{100}$.

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

Universe has insufficiently many particles to represent each polynomial.

Expect *three* counterexamples before $10^{750}$.

Moore's law $\implies$ get to $10^{750}$ in 10000 years.

# 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 $\varepsilon_k E_p^+$.

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

# 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 essentially solved this sequentially for $B_2, B_4, B_6, \ldots$, yielding running 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.

# 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 this to a single polynomial multiplication of length $\sim p/2$ over $\mathbf{Z}/p\mathbf{Z}$.

# Verifying Vandiver's conjecture

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

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

To test whether $S_k$ is a $p$-th power, we only need consider

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

which approximates $S_k$ modulo $(E_p^+)^p$.

# 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 $\mathbf{Z}/\ell\mathbf{Z}$.

If this test fails for one $\ell$, we could try a different $\ell$ — but so far this has never been necessary.

# Verifying Vandiver's conjecture

Besides Vandiver's conjecture, we also compute the lambda invariant from Iwasawa theory. Essentially we check that $A_p$ is as small as possible consistent with the value of $i(p)$ (i.e. that each nontrivial $\varepsilon_i A_p$ is no bigger than $\mathbf{Z}/p\mathbf{Z}$).

# The software

# The software

The most expensive part of the computation is finding the Bernoulli numbers modulo $p$.

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

To make best use of the 64-bit processor, we do everything modulo two primes simultaneously.

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
- Included in recent versions of Sage, but no direct interface yet
- 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 things still missing

# zn_poly multiplication performance

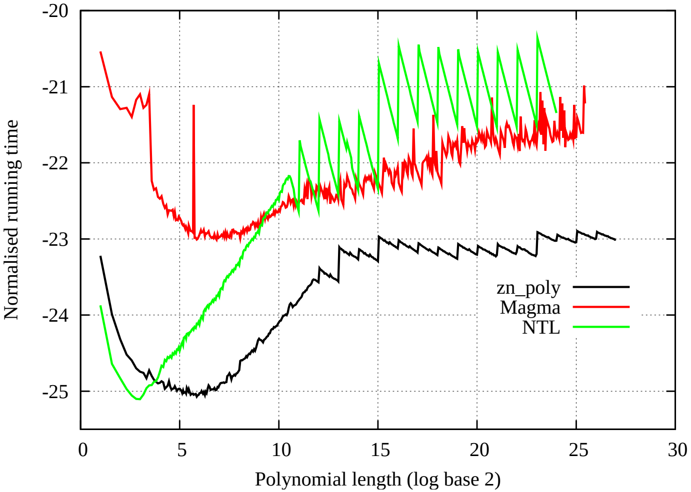


Figure: Multiplication of polynomials modulo a 48-bit modulus (Opteron)

# Multiplication algorithms

Multiplication algorithms:

- For small degree (say $\leq 4000$, depending on modulus size), uses ordinary or multipoint Kronecker substitution (H., 2008) reducing the problem to integer multiplication (via GMP).
- For large degree uses Schönhage–Nussbaumer convolution. Reduces length $n$ multiplication to $O(\sqrt{n})$ multiplications of length $O(\sqrt{n})$.
- The Schönhage–Nussbaumer convolution uses a cache-friendly adaptation (H., 2008) of the truncated FFT (van der Hoeven, 2005) for smooth performance.
- Future versions will also use naive classical multiplication for low degree (currently under development).

# Series inversion algorithms

For series inversion, uses a $1.5M(n)$ algorithm based on the middle product (Hanrot–Quercia–Zimmerman, 2004).

The middle product is implemented via the transposition principle (includes a transposed truncated FFT and IFFT...).

# Integer multiplication

Integer multiplication:

- New GMP assembly code, written especially for the Opteron
- About 25–30% faster than Gaudry's well-known patch
- Written by Torbjörn Granlund and H.
- Should be released in GMP 4.3, hopefully later this year (more likely next year)

# The hardware

# Small–to–medium machines

- My laptop (2 $\times$ 2.0GHz Core 2 Duo, 1GB RAM)
- sage.math (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)

# Slightly larger machines

TACC clusters:

- Lonestar: 1300 nodes.
  - Each node $= 4 \times 2.66$GHz Xeon (Woodcrest), 8GB RAM.
  - Total cores $= 5200$, total RAM $= 10$ TB.
  - We used $\approx 119000$ core-hours.
- Ranger: 3936 nodes.
  - Each node $= 16 \times 2.3$GHz Opteron (Barcelona), 32GB RAM.
  - Total cores $= 62976$, total RAM $= 123$ TB.
  - We used $\approx 69000$ core-hours.

About **21 core-years** altogether.

On both machines, have 2GB RAM per core. If $p \approx 163{,}577{,}856$, one polynomial of length $p/2$ requires 0.6 GB to store. Not much room to move! Managing memory was the biggest challenge of the computation.

# Machines drawn to scale

My laptop

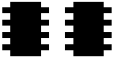

sage.math

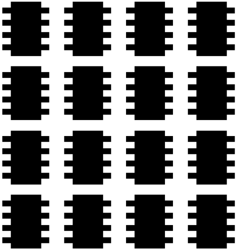

# Machines drawn to scale

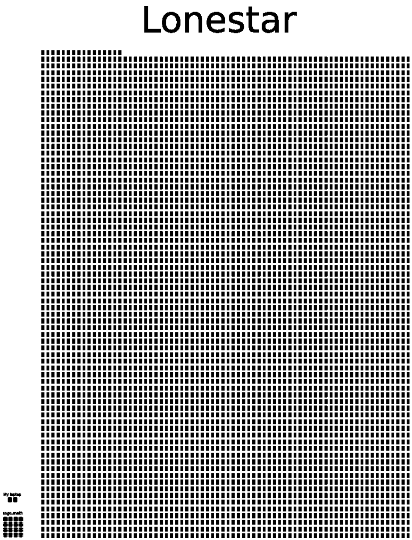

# Machines drawn to scale

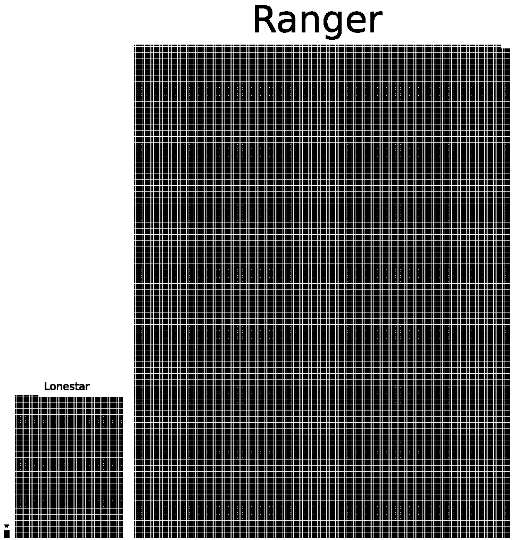

Thank you!