

# Factoring Polynomials over Finite Fields

More precisely: Factoring and testing irreducibility of sparse polynomials over small finite fields

Richard P. Brent  
MSI, ANU

joint work with  
Paul Zimmermann  
INRIA, Nancy

27 August 2009

# Outline

- Introduction
  - ▶ Polynomials over finite fields
  - ▶ Irreducible and primitive polynomials
  - ▶ Mersenne primes
- Part 1: Testing irreducibility
  - ▶ Irreducibility criteria
  - ▶ Modular composition
  - ▶ Three algorithms
  - ▶ Comparison of the algorithms
  - ▶ The “best” algorithm
  - ▶ Some computational results
- Part 2: Factoring polynomials
  - ▶ Distinct degree factorization
  - ▶ Avoiding GCDs, blocking
  - ▶ Another level of blocking
  - ▶ Average-case complexity
  - ▶ New primitive trinomials

# Polynomials over finite fields

We consider univariate polynomials  $P(x)$  over a finite field  $F$ . The algorithms apply, with minor changes, for any small positive characteristic, but since time is limited we assume that the characteristic is two, and  $F = \mathbb{Z}/2\mathbb{Z} = \text{GF}(2)$ .

$P(x)$  is *irreducible* if it has no nontrivial factors. If  $P(x)$  is irreducible of degree  $r$ , then [Gauss]

$$x^{2^r} = x \pmod{P(x)}.$$

Thus  $P(x)$  divides the polynomial  $\mathcal{P}_r(x) = x^{2^r} - x$ . In fact,  $\mathcal{P}_r(x)$  is the product of all irreducible polynomials of degree  $d$ , where  $d$  runs over the divisors of  $r$ .

## Counting irreducible polynomials

Let  $N(d)$  be the number of irreducible polynomials of degree  $d$ . Thus

$$\sum_{d|r} dN(d) = \deg(\mathcal{P}_r) = 2^r .$$

By Möbius inversion we see that

$$rN(r) = \sum_{d|r} \mu(d) 2^{r/d} .$$

Thus, the number of irreducible polynomials of degree  $r$  is

$$N(r) = \frac{2^r}{r} + O\left(\frac{2^{r/2}}{r}\right) .$$

Since there are  $2^r$  polynomials of degree  $r$ , the probability that a randomly selected polynomial is irreducible is  $\sim 1/r \rightarrow 0$  as  $r \rightarrow +\infty$ .

*Almost all* polynomials over (fixed) finite fields are reducible.

## Analogy

Polynomials of degree  $r$  are analogous to integers of  $r$  digits. By the prime number theorem, the number of  $r$ -digit primes in base  $b$  is about

$$\int_{b^{r-1}}^{b^r} \frac{dt}{\ln t} = \left( \frac{b^r - b^{r-1}}{r \ln b} \right) \left( 1 + O\left(\frac{1}{r}\right) \right).$$

The Riemann Hypothesis implies an error term  $O(rb^{r/2})$  as  $r \rightarrow +\infty$  for the integral on the left [von Koch].

The [conditional] error bound for integers is not quite as good as the [proved] error bound for finite fields. Littlewood's 1914 result is

$$\psi(x) - x = \Omega_{\pm}(x^{1/2} \log \log \log x)$$

so presumably

$$\pi(x) - \text{Li}(x) = \Omega_{\pm}(x^{1/2} \log \log \log x / \log x)$$

and we can not expect to get as good a bound for integers as for finite fields, even if RH is true.

# Representing finite fields

Irreducible polynomials over finite fields are useful in several applications. As one example, observe that, if  $P(x)$  is an irreducible polynomial of degree  $r$  over  $\text{GF}(2)$ , then

$$\text{GF}(2)[x]/P(x) \cong \text{GF}(2^r).$$

In other words, the ring of polynomials mod  $P(x)$  gives a representation of the finite field with  $2^r$  elements.

If, in addition,  $x$  is a generator of the multiplicative group, that is if every nonzero element of  $\text{GF}(2)[x]/P(x)$  can be represented as a power of  $x$ , then  $P(x)$  is said to be *primitive*.

# Primitive polynomials and shift registers

Primitive polynomials can be used to obtain linear feedback shift registers (LFSRs) with maximal period  $2^r - 1$ , where  $r$  is the degree of the polynomial. These have applications to stream ciphers and pseudo-random number generators.

Testing primitivity can be difficult, because we need to know the prime factorization of  $2^r - 1$ . Of course, this is trivial if  $2^r - 1$  is prime (a *Mersenne prime*).

The number of primitive polynomials of degree  $r$  over  $\text{GF}(2)$  is

$$\frac{\phi(2^r - 1)}{r} \leq N(r) \leq \frac{2^r - 2}{r},$$

with equality when  $2^r - 1$  is prime.

# Sparsity

In applications we often want  $P(x)$  to be *sparse*, that is to have only a small number of nonzero coefficients. Ideally  $P(x)$  should be a *trinomial*

$$x^r + x^s + 1, \quad r > s > 0.$$

However, for the application to random number generators, sparsity is a two-edged sword, because it implies unwanted correlations. For example, a random number generator based on the trinomial  $x^r + x^s + 1$  has  $x_n$  depending on  $(x_{n-r}, x_{n-s})$ . This was the motivation for my random number generator *xorgens* (the topic of another talk).



# Mersenne primes

A *Mersenne prime* is a prime of the form  $2^n - 1$ , for example 3, 7, 31, 127, 8191, ...

There are *conjectured* to be infinitely many Mersenne primes. The number for  $n \leq N$  is conjectured to be of order  $\log N$ .

The GIMPS project is searching systematically for Mersenne primes. So far 47 Mersenne primes are known, the largest being

$$2^{43112609} - 1 .$$

If  $2^n - 1$  is prime we say that  $n$  is a *Mersenne exponent*. A Mersenne exponent is necessarily prime, but not conversely ( $2^{11} - 1 = 23 \times 89$ ).

## Part 1: Testing irreducibility

Since irreducible polynomials are “rare” but useful, we are interested in algorithms for testing irreducibility.

$P(x)$  of degree  $r > 1$  is irreducible iff

$$x^{2^r} = x \pmod{P(x)}$$

and, for all prime divisors  $d$  of  $r$ , we have

$$\text{GCD} \left( x^{2^{r/d}} - x, P(x) \right) = 1 .$$

The second condition is required to rule out the possibility that  $P(x)$  is a product of irreducible factors of some degree(s)  $k = r/d$ ,  $d|r$ . This condition does not significantly change anything, so let us assume that  $r$  is prime. (In our examples  $r$  is a Mersenne exponent, so necessarily prime.) Then  $P(x)$  is irreducible iff

$$x^{2^r} = x \pmod{P(x)}.$$

## Another assumption

All the algorithms involve computations mod  $P(x)$ , that is, in the ring  $\text{GF}(2)[x]/P(x)$ .

In the complexity analysis we assume that  $P(x)$  is *sparse*, that is, the number of nonzero coefficients is small. Thus, reduction of a polynomial mod  $P(x)$  can be done in linear time.

The algorithms to be discussed still work without this assumption, but the complexity analysis no longer applies because more time is spent in the reductions mod  $P(x)$ .

In applications  $P(x)$  is often a *trinomial*

$$P(x) = x^r + x^s + 1, \quad r > s > 0.$$

although sometime *pentanomials*  $x^p + x^q + x^r + x^s + 1$  are considered.

## Irreducible and primitive trinomials

We have given formulas for the number of irreducible or primitive *polynomials* of degree  $r$  over  $\text{GF}(2)$ , but there is no known formula for the number of irreducible or primitive *trinomials*.

Since the number of irreducible polynomials  $N(r) \approx 2^r/r$ , the probability that a randomly chosen polynomial of degree  $r$  will be irreducible is about  $1/r$ .

It is *plausible* to assume that the same applies to trinomials. There are  $r - 1$  trinomials of degree  $r$ , so we might expect  $O(1)$  of them to be irreducible. More precisely, we might expect a Poisson distribution with some constant mean  $\mu$ .

This plausible argument is **too simplistic**, as shown by Swan's theorem. On the next slide we state a simplified version of Swan's Theorem that is relevant to trinomials.

## Swan's theorem (Corollary 5)

### Theorem

Let  $r > s > 0$ , and assume  $r + s$  is odd. Then  $T_{r,s}(x) = x^r + x^s + 1$  has an even number of irreducible factors over  $\text{GF}(2)$  in the following cases:

- a)  $r$  even,  $r \neq 2s$ ,  $rs/2 = 0$  or  $1 \pmod{4}$ .
- b)  $r$  odd,  $s$  not a divisor of  $2r$ ,  $r = \pm 3 \pmod{8}$ .
- c)  $r$  odd,  $s$  divisor of  $2r$ ,  $r = \pm 1 \pmod{8}$ .

In all other cases  $x^r + x^s + 1$  has an odd number of irreducible factors.

**Other cases:** if both  $r$  and  $s$  are even, then  $T_{r,s}(x)$  is a square.  
If both  $r$  and  $s$  are odd, apply the Theorem to  $T_{r,r-s}(x)$ .

# Implications of Swan's Theorem

For  $r$  is an odd prime, case (b) of Swan's Theorem says that the trinomial has an *even* number of irreducible factors, and hence must be *reducible*, if  $r = \pm 3 \pmod{8}$ , provided we exclude the special cases  $s = 2$  and  $r - s = 2$ .

For prime  $r = \pm 1 \pmod{8}$ , the heuristic Poisson distribution does seem to apply, with mean  $\mu \approx 3$ . Similarly for primitive trinomials, with a correction factor  $\phi(2^r - 1)/(2^r - 2)$ .

## Historical note

Swan (1962) rediscovered results of Pellet (1878) and Stickelberger (1897), so the name of the theorem depends on your nationality.

## Testing irreducibility — first algorithm

Our first and simplest algorithm for testing irreducibility is just *repeated squaring*:

```
Q(x) ← x;  
for j ← 1 to r do  
  Q(x) ← Q(x)2 mod P(x);  
if Q(x) = x then  
  return irreducible  
else  
  return reducible.
```

The operation  $Q(x) \leftarrow Q(x)^2 \bmod P(x)$  can be performed in time  $O(r)$ . The constant factor is small.

Since the irreducibility test involves  $r$  squarings, the overall time is  $O(r^2)$ .

## Polynomial multiplication

Before describing other algorithms for irreducibility testing, we digress to discuss polynomial multiplication, matrix multiplication, and modular composition.

To multiply two polynomials  $A(x)$  and  $B(x)$  of degree (at most)  $r$ , the “classical” algorithm takes time  $O(r^2)$ . There are faster algorithms, e.g. Karatsuba, Toom-Cook, and FFT-based algorithms.

For polynomials over  $GF(2)$ , the asymptotically fastest known algorithm is due to Schönhage. (The Schönhage-Strassen algorithm does not work in characteristic 2.)

Schönhage’s algorithm runs in time

$$M(r) = O(r \log r \log \log r) .$$

In practice, for  $r \approx 32\,000\,000$ , a multiplication takes about 480 times as long as a squaring.



## Matrix multiplication

Let  $\omega$  be the exponent of matrix multiplication, so we can multiply  $n \times n$  matrices in time  $O(n^{\omega+\varepsilon})$  for any  $\varepsilon > 0$ . The best result is Coppersmith and Winograd's  $\omega < 2.376$ , though in practice we would use the classical ( $\omega = 3$ ) or Strassen ( $\omega = \log_2 7 \approx 2.807$ ) algorithm.

Since we are working over  $\text{GF}(2)$ , our matrices have single-bit entries. This means that the classical algorithm can be implemented very efficiently using full-word operations (32 or 64 bits at a time). Nevertheless, Strassen's algorithm is faster if  $n$  is larger than about 1000.

Good in practice is the “Four Russians” algorithm [Arlazarov, Dinic, Kronod & Faradzev, 1970]. It computes  $n \times n$  Boolean matrix multiplication in time  $O(n^3 / \log n)$ .

We can use the Four Russians' algorithm up to some threshold, say  $n = 1024$ , and Strassen's recursion for larger  $n$ , combining the advantages of both.

## Modular composition

The *modular composition* problem is: given polynomials  $A(x)$ ,  $B(x)$ ,  $P(x)$ , compute

$$C(x) = A(B(x)) \bmod P(x).$$

If  $\max(\deg(A), \deg(B)) < r = \deg(P)$ , then we could compute  $A(B(x))$ , a polynomial of degree at most  $(r-1)^2$ , and reduce it modulo  $P(x)$ . However, this wastes both time and space.

Better is to compute

$$C(x) = \sum_{j \leq \deg(A)} a_j (B(x))^j \bmod P(x)$$

by Horner's rule, reducing mod  $P(x)$  as we go, in time  $O(rM(r))$  and space  $O(r)$ . Using Schönhage's algorithm for the polynomial multiplications, we can compute  $C(x)$  in time  $O(r^2 \log r \log \log r)$ .

## Faster modular composition

Using an algorithm of Brent & Kung (1978), based on an idea of Paterson and Stockmeyer, we can reduce the modular composition problem to a problem of matrix multiplication. If the degrees of the polynomials are at most  $r$ , and  $m = \lceil r^{1/2} \rceil$ , then we have to perform  $m$  multiplications of  $m \times m$  matrices. The matrices are over the same field as the polynomials (that is,  $\text{GF}(2)$  here).

The Brent-Kung modular composition algorithm takes time

$$O(r^{(\omega+1)/2}) + O(r^{1/2}M(r)),$$

where the first term is for the matrix multiplications and the second term is for computing the relevant matrices.

Assuming Strassen's matrix multiplication, the first term is  $O(r^{1.904})$  and the second term is  $O(r^{1.5} \log r \log \log r)$ . Thus, the second term is asymptotically negligible (but maybe not in practice).

## Using modular composition

Let  $A_k(x) = x^{2^k} \bmod P(x)$ . Then a modular composition algorithm can be used to compute  $A_k(A_m(x)) \bmod P(x)$ . Since

$$A_k(A_m(x)) = \left(x^{2^m}\right)^{2^k} \bmod P(x) = A_{m+k}(x),$$

we can compute  $x^{2^r} \bmod P(x)$  with about  $\log_2(r)$  modular compositions instead of  $r$  squarings.

For example, if  $r = 17$ , we have  
(all computations in  $\text{GF}(2)[x]/P(x)$ ):

$$\begin{aligned} A_1(x) &= x^2, && \text{(trivial)} \\ A_2(x) &= A_1(A_1(x)) = x^4, && (\equiv 1 \text{ squaring}) \\ A_4(x) &= A_2(A_2(x)) = x^{16}, && (\equiv 2 \text{ squarings}) \\ A_8(x) &= A_4(A_4(x)) = x^{256}, && (\equiv 4 \text{ squarings}) \\ A_{16}(x) &= A_8(A_8(x)) = x^{2^{16}}, && (\equiv 8 \text{ squarings}) \\ A_{17}(x) &= A_{16}(x)^2 = x^{2^{17}}, && (1 \text{ squaring}) \end{aligned}$$

using only 4 modular composition steps.

## Second algorithm

To summarise, we can compute  $A_r(x) = x^{2^r} \bmod P(x)$  by the following recursive algorithm that uses the binary representation of  $r$  (*not* that of  $2^r$ ):

```
if  $r = 0$  then
  return  $x$ 
else if  $r$  even then
  { $U(x) \leftarrow A_{r/2}(x)$ ;
  return  $U(U(x)) \bmod P(x)$ }
else
  return  $A_{r-1}(x)^2 \bmod P(x)$ .
```

The algorithm takes about  $\log_2(r)$  modular compositions. Hence, if Strassen's algorithm is used in the Brent-Kung modular composition algorithm, we can test irreducibility in time  $O(r^{1.904} \log r)$ .

## Third algorithm

Recently, Kedlaya and Umans (2008) proposed an asymptotically fast modular composition algorithm that runs in time  $O_\varepsilon(r^{1+\varepsilon})$  for any  $\varepsilon > 0$ .

The algorithm is complicated, involving iterated reductions to multipoint multivariate polynomial evaluation, multidimensional FFTs, and the Chinese remainder theorem.

For details, see the papers on Umans's web site

<http://www.cs.caltech.edu/~umans/research.htm>

Using the Kedlaya-Umans fast modular composition instead of the Brent-Kung reduction to matrix multiplication, we can test irreducibility in time  $O_\varepsilon(r^{1+\varepsilon})$ .

**Warning:** the “ $O_\varepsilon(\dots)$ ” notation indicates that the implicit constant depends on  $\varepsilon$ . In this case, it is a rather large and rapidly increasing (probably exponential) function of  $1/\varepsilon$ .

# Comparison of the algorithms

*So the last shall be first,  
and the first last*

*Matthew 20:16*

The theoretical time bounds predict that the third algorithm should be the fastest, and the first algorithm the slowest. However, this is only for *sufficiently large* degrees  $r$ .

In practice, for  $r$  up to at least  $4.3 \times 10^7$ , the situation is reversed! The first algorithm is the fastest, and the third algorithm is the slowest.

# Parallel implementation

A minor drawback of the first (squaring) algorithm is that it is hard to speed up on a parallel machine. The other algorithms are much easier to parallelise.

Fortunately, this is not so relevant when we are considering many trinomials, as we can let different processors of a parallel machine work on different trinomials in parallel.



## Example, $r = 32\,582\,657$

Following are actual or estimated times on a 2.2 Ghz AMD Opteron 275 for  $r = 32\,582\,657$  (a Mersenne exponent).

- 1 Squaring (actual): 64 hours
- 2 Brent-Kung (estimates):
  - ▶ classical: 265 hours (19% mm)
  - ▶ Strassen: 254 hours (15% mm)
  - ▶ Four Russians: 239 hours (10% mm)  
(plus Strassen for  $n > 1024$ )
- 3 Kedlaya-Umans (estimate):  $> 10^{10}$  years

The Brent-Kung algorithm would be the fastest if the matrix multiplication were dominant; unfortunately the  $O(r^{1/2}M(r))$  overhead term dominates.

Since the overhead scales roughly as  $r^{1.5}$ , we estimate that the Brent-Kung algorithm would be faster than the squaring algorithm for  $r > 7 \times 10^8$  (approximately).

## Note on Kedlaya-Umans

Éric Schost writes:

*The Kedlaya-Umans algorithm reduces modular composition to the multipoint evaluation of a multivariate polynomial, assuming the base field is large enough.*

*The input of the evaluation is over  $F_p$ ; the algorithm works over  $\mathbb{Z}$  and reduces mod  $p$  in the end. The evaluation over  $\mathbb{Z}$  is done by CRT modulo a bunch of smaller primes, and so on. At the end-point of the recursion, we do a naive evaluation on all of  $F_{p^m}$ , where  $p$  is the current prime and  $m$  the number of variables. So the cost here is  $\geq p^m$ .*

*[Now he considers choices of  $m$  in the case  $r = 32\,582\,657$ ; all give  $p^m \geq 1.36 \times 10^{27}$ .]*

Our estimate of  $> 10^{10}$  years is based on a time of 1 nsec per evaluation (very optimistic).

## The “best” algorithm

Comparing the second algorithm with the first, observe that the modular compositions do not all save equal numbers of squarings. In fact the *last* modular composition saves  $\lfloor r/2 \rfloor$  squarings, the *second-last* saves  $\lfloor r/4 \rfloor$  squarings, etc.

Each modular composition has the same cost. Thus, if we can use only one modular composition, *it should be the one that saves the most squarings*.

If we use  $\lfloor r/2 \rfloor$  squarings to compute  $x^{2^{\lfloor r/2 \rfloor}} \bmod P(x)$ , then use one modular composition (and one further squaring, if  $r$  is odd), we can compute  $x^{2^r} \bmod P(x)$  faster than with any of the algorithms considered so far, provided  $r$  exceeds a certain threshold.

In the example, the time would be reduced from 64 hours to 44 hours, a saving of 31%.

Doing two modular compositions would reduce the time to 40 hours, a saving of 37%.

## Computational results

In 2007-8 Paul Zimmermann and I conducted a search for irreducible trinomials  $x^r + x^s + 1$  whose degree  $r$  is a (known) Mersenne exponent. Since  $2^r - 1$  is prime, *irreducible* implies *primitive*. The previous record degree of a primitive trinomial was  $r = 6\,972\,593$ .

$r$	$s$
24 036 583	8 412 642, 8 785 528
25 964 951	880 890, 4 627 670, 4 830 131, 6 383 880
30 402 457	2 162 059
32 582 657	5 110 722, 5 552 421, 7 545 455

**Table:** Ten new primitive trinomials  $x^r + x^s + 1$  of degree a Mersenne exponent, for  $s \leq r/2$ .

We used the first algorithm to test irreducibility of the most difficult cases. Most of the time was spent discarding the vast majority of trinomials that have a small factor, using a new factoring algorithm with good average-case behaviour (the topic of the second half of this talk).

## Part 2: Factoring

The problem of factoring a univariate polynomial  $P(x)$  over a finite field  $F$  often arises in computational algebra. An important case is when  $F$  has small characteristic and  $P(x)$  has high degree but is *sparse* (has only a small number of nonzero terms).

Since time is limited, I will make the same assumptions as in Part 1:  $F = \text{GF}(2)$  and  $P(x)$  is sparse, typically a trinomial

$$P(x) = x^r + x^s + 1, \quad r > s > 0,$$

although the ideas apply more generally.

The aim is to give an algorithm with good *amortized complexity*, that is, one that works well *on average*. Since we are restricting attention to trinomials, we average over all trinomials of fixed degree  $r$ .

Equivalently, we can use probabilistic language, and assume a uniform distribution over all trinomials of fixed degree  $r$ .

# Simplifications

## Square-free factorisation

We assume that  $P(x)$  is *square-free*. This is trivial for trinomials; in general we can consider  $P(x)/\text{GCD}(P(x), P'(x))$ .

## Distinct-degree factorization

We only consider *distinct-degree factorization*. That is, if  $P(x)$  has several factors of the same degree  $d$ , the algorithm will produce the product of these factors. The Cantor-Zassenhaus algorithm can be used to split this product into distinct factors. This is usually cheap because in most cases the product has small degree or consists of just one factor.

# Small factors and certificates

## Factor of smallest degree

To simplify the complexity analysis and speed up the algorithm in the common application of searching for irreducible polynomials, I only consider the time required to find *one* nontrivial factor or output “irreducible”.

## Certificates of reducibility

A nontrivial factor (preferably of smallest degree) gives a “reducibility certificate” that can quickly be checked. For example, if we claim that there are exactly two primitive trinomials of degree 859433, this statement can easily be checked using the certificates for degree 859433.

Kumada et al [*Math. Comp.* 2000] failed to do this, and missed the primitive trinomial

$$x^{859433} + x^{170340} + 1$$

because of a bug in their program!

## Factorization in $\text{GF}(2)[x]$

From now on we write “+” instead of “-” (they are equivalent in  $\text{GF}(2)[x]$ ).

As we already mentioned,  $x^{2^d} + x$  is the product of all irreducible polynomials of degree dividing  $d$ . For example,

$$x^{2^3} + x = x(x + 1)(x^3 + x + 1)(x^3 + x^2 + 1).$$

Thus, a simple (slow) algorithm to find a factor of smallest degree of  $P(x)$  is to compute  $\text{GCD}(x^{2^d} + x, P(x))$  for  $d = 1, 2, \dots$ . The first time that the GCD is nontrivial, it contains a factor of minimal degree  $d$ . If the GCD has degree  $> d$ , it must be a product of factors of degree  $d$ .

If no factor has been found for  $d \leq r/2$ , where  $r = \deg(P(x))$ , then  $P(x)$  must be irreducible.

Note that  $x^{2^d}$  should not be computed explicitly; instead compute  $x^{2^d} \bmod P(x)$  by repeated squaring.



## Application to trinomials

Some simplifications are possible when  $P(x) = x^r + x^s + 1$  is a trinomial.

- We can skip the case  $d = 1$  because a trinomial can not have a factor of degree 1.
- Since  $x^r P(1/x) = x^r + x^{r-s} + 1$ , we only need consider  $s \leq r/2$ .
- In the cases of interest ( $r$  and  $s$  not both even),  $P(x)$  is squarefree.
- By applying Swan's theorem, we can usually show that the trinomial under consideration has an odd number of factors; in this case we only need check  $d \leq r/3$ .

# Complexity of squares and multiplications

In Part 1 we already considered the complexity of computing squares and products in  $\text{GF}(2)[x]/P(x)$ . Recall that, with our usual assumption that  $P(x)$  is sparse, squaring can be performed in time

$$S(r) = \Theta(r) \ll M(r)$$

and multiplication can be performed in time

$$M(r) = O(r \log r \log \log r) .$$

In the complexity estimates we assume that  $M(r)$  is a sufficiently smooth and well-behaved function.

## Complexity of GCD

For GCDs we use a sub-quadratic algorithm that runs in time  $G(r) = O(M(r) \log r)$ .

More precisely,

$$G(2r) = 2G(r) + O(M(r)) ,$$

so

$$M(r) = O(r \log r \log \log r) \Rightarrow G(r) = \Theta(M(r) \log r) .$$

In practice, for  $r \approx 2.4 \times 10^7$  and our implementation on a 2.2 Ghz Opteron,

$$S(r) \approx 0.005 \text{ seconds,}$$

$$M(r) \approx 2 \text{ seconds,}$$

$$G(r) \approx 80 \text{ seconds,}$$

$$M(r)/S(r) \approx 400 ,$$

$$G(r)/M(r) \approx 40 .$$

## Avoiding GCD computations

In the context of integer factorization, Pollard (1975) suggested a blocking strategy to avoid most GCD computations and thus reduce the amortized cost; von zur Gathen and Shoup (1992) applied the same idea to polynomial factorization.

The idea of blocking is to choose a parameter  $\ell > 0$  and, instead of computing

$$\text{GCD}(x^{2^d} + x, P(x)) \text{ for } d \in [d', d' + \ell),$$

compute

$$\text{GCD}(p_\ell(x^{2^{d'}}, x), P(x)),$$

where the *interval polynomial*  $p_\ell(X, x)$  is defined by

$$p_\ell(X, x) = \prod_{j=0}^{\ell-1} (X^{2^j} + x).$$

In this way we replace  $\ell$  GCDs by **one GCD and  $\ell - 1$  multiplications** mod  $P(x)$ .

# Backtracking

The drawback of blocking is that we may have to backtrack if  $P(x)$  has more than one factor with degrees in  $[d', d' + \ell)$ , so  $\ell$  should not be too large. The optimal strategy depends on the expected size distribution of factors and the ratio of times for GCDs and multiplications.

## New idea - multi-level blocking

We introduce a finer level of blocking to replace most **multiplications** by **squarings**, which speeds up the computation in  $\text{GF}(2)[x]/P(x)$  of the interval polynomials  $\rho_m(x^{2^d}, x)$ , where

$$\rho_m(X, x) = \prod_{j=0}^{m-1} (X^{2^j} + x) = \sum_{j=0}^m x^{m-j} s_{j,m}(X),$$

$$s_{j,m}(X) = \sum_{0 \leq k < 2^m, w(k)=j} X^k,$$

and  $w(k)$  denotes the Hamming weight of  $k$ .

Note that  $s_{j,m}(X^2) = s_{j,m}(X)^2$  in  $\text{GF}(2)[x]/P(x)$ . Thus,  $\rho_m(x^{2^d}, x)$  can be computed with cost  $m^2 S(r)$  if we already know  $s_{j,m}(x^{2^{d-m}})$  for  $0 < j \leq m$ .

## Effect of multi-level blocking

Using multi-level blocking, we replace  $m$  multiplications and  $m$  squarings by one multiplication and  $m^2$  squarings. Choosing

$$m \approx \sqrt{M(r)/S(r)},$$

(about 20 if  $M(r)/S(r) \approx 400$ ), the speedup over single-level blocking is about  $m/2 \approx 10$ .

# Fast initialization

The polynomials

$$s_{j,m}(x) = \sum_{0 \leq k < 2^m, w(k)=j} x^k$$

satisfy a “Pascal triangle” recurrence relation

$$s_{j,m}(x) = s_{j,m-1}(x^2) + x \times s_{j-1,m-1}(x^2)$$

with boundary conditions

$$s_{j,m}(x) = 0 \text{ if } j > m ,$$

$$s_{0,m}(x) = 1 .$$

Thus, we can compute

$$\{s_{j,m}(x) \bmod P(x) \mid 0 \leq j \leq m\}$$

in time  $O(m^2r)$ , even though the definition of  $s_{j,m}(x)$  involves  $O(2^m)$  terms.



# Recapitulation

To summarize, we use two levels of blocking:

- The outer level replaces most GCDs by multiplications.
- The inner level replaces most multiplications by squarings.
- The blocking parameter  $m \approx \sqrt{M(r)/S(r)}$  is used for the inner level of blocking.
- A different parameter  $\ell = km$  is used for the outer level of blocking.

## Example




Figure:  $\ell = 15, m = 5$

In the example,  $S = 1/25, M = 1, G = 10$

No blocking: cost  $15G + 15S = 150.6$

1-level blocking:  $G + 14M + 15S = 24.6$

2-level blocking:  $G + 2M + 75S = 15.0$

More realistically, suppose  $\ell = 80, m = 20,$   
 $S = 1/400, M = 1, G = 40$

No blocking: cost  $80G + 80S = 3200.2$

1-level blocking:  $G + 79M + 80S = 119.2$

2-level blocking:  $G + 3M + 1600S = 47.0$

# Sieving

A *small* factor is one with degree  $d < \frac{1}{2} \log_2 r$ , so  $2^d < \sqrt{r}$ .

It would be inefficient to find small factors in the same way as large factors. Instead, let  $d' = 2^d - 1$ ,  $r' = r \bmod d'$ ,  $s' = s \bmod d'$ . Then

$$P(x) = x^r + x^s + 1 = x^{r'} + x^{s'} + 1 \bmod (x^{d'} - 1),$$

so we only need compute

$$\text{GCD}(x^{r'} + x^{s'} + 1, x^{d'} - 1).$$

The cost of finding small factors is negligible (both theoretically and in practice), so will be ignored.

In the definition, the fraction  $\frac{1}{2}$  is rather arbitrary; it can be replaced by  $1 - \varepsilon$  for any  $\varepsilon > 0$ .

## Distribution of degrees of factors

In order to predict the expected behaviour of our algorithm, we need to know the expected distribution of degrees of irreducible factors. Our complexity estimates here are based on the assumption that trinomials of degree  $r$  behave like the set of all polynomials of the same degree, up to a constant factor:

**Assumption 1.** Over all trinomials  $x^r + x^s + 1$  of degree  $r$  over  $\text{GF}(2)$ , the probability  $\pi_d$  that a trinomial has no nontrivial factor of degree  $\leq d$  is at most  $c/d$ , where  $c$  is a constant and  $1 < d \leq r$ .

This assumption is plausible and in agreement with experiments, though not proven. Under the assumption, we use an amortized model to obtain the total complexity over all trinomials of degree  $r$ .

From Assumption 1, the probability that a trinomial does not have a small factor is  $O(1/\log r)$ .

## Simpler approximation

Let  $p_d = \pi_{d-1} - \pi_d$  be the probability that the smallest nontrivial factor of a randomly chosen trinomial has degree  $d \geq 2$ . Although not strictly correct, the following is a good approximation.

**Assumption 2.**  $p_d$  is of order  $1/d^2$ , provided  $d$  is not too large.

I will use Assumption 2 because it simplifies the amortized complexity analysis, but the same results can be obtained from Assumption 1 using summation by parts.

Some empirical evidence for Assumptions 1–2 in the case  $r = 6\,972\,593$  is given on the next slide. Results for other large Mersenne exponents are similar.

# Statistics for $r = 6972593$

$d$	$d\pi_d$	$d^2\rho_d$
2	1.33	1.33
3	1.43	1.71
4	1.52	1.52
5	1.54	1.84
6	1.60	1.47
7	1.60	1.85
8	1.67	1.29
9	1.64	2.10
10	1.65	1.73
100	1.77	
1000	1.76	
10000	1.88	
100000	1.62	
226887	2.08	
$r - 1$	2.00	

# Analogies

The following have similar distributions in the limit as  $n \rightarrow \infty$ :

- 1 Degree of smallest irreducible factor of a random monic polynomial of degree  $n$  over a finite field (say  $\text{GF}(2)$ ).
- 2 Size of smallest cycle in a random permutation of  $n$  objects.
- 3 Size (in base- $b$  digits) of smallest prime factor in a random integer of  $n$  digits.

## Analogies — more details

More precisely, let  $P_d$  be the limiting probability that the smallest irreducible factor has degree  $> d$ , that the smallest cycle has length  $> d$ , or that the smallest prime factor has  $> d$  digits, in cases 1–3 respectively. Then

$$P_d \sim c/d \text{ as } d \rightarrow \infty$$

(the constant  $c$  is different in each case).

For example, in case 3, let  $x = b^d$ ; then

$$P_d = \prod_{\text{prime } p < x} \left(1 - \frac{1}{p}\right) \sim \frac{e^{-\gamma}}{\ln x} = \left(\frac{e^{-\gamma}}{\ln b}\right) \frac{1}{d}$$

by the theorem of Mertens.



## Outer level blocking strategy

The blocksize in the outer level of blocking is  $\ell = km$ . We take an increasing sequence

$$k = k_0j \text{ for } j = 1, 2, 3, \dots,$$

where  $k_0m$  is of order  $\log r$  (since small factors will have been found by sieving). This leads to a quadratic polynomial for the interval bounds.

There is nothing magic about a quadratic polynomial, but it is simple to implement and experiments show that it is reasonably close to optimal.

## Optimal blocking strategy?

Using the data that we have obtained on the distribution of degrees of smallest factors of trinomials, and assuming that this distribution is insensitive to the degree  $r$ , we could obtain a strategy that is close to optimal.

The choice  $k_0 j$  with suitable  $k_0$  is simple and not too far from optimal. The number of GCD and  $\text{sqr}/\text{mul}$  operations is usually within a factor of 1.5 of the minimum possible.

## Expected cost of sqr/mul

Recall that the inner level of blocking replaces  $m$  multiplications by  $m^2$  squarings and one multiplication, where  $m \approx \sqrt{M(r)/S(r)}$  makes the cost of squarings about equal to the cost of multiplications.

For a smallest factor of degree  $d$ , the expected number of squarings is  $m(d + O(\sqrt{d}))$ . Averaging over all trinomials of degree  $r$ , the expected number is

$$O\left(m \sum_{d \leq r/2} \frac{d + O(\sqrt{d})}{d^2}\right) = O(m \log r) .$$

Thus, the expected cost of sqr/mul operations per trinomial is

$$\begin{aligned} & O\left(S(r) \log r \sqrt{M(r)/S(r)}\right) \\ &= O\left(\log r \sqrt{M(r)S(r)}\right) \\ &= O\left(r(\log r)^{3/2}(\log \log r)^{1/2}\right) . \end{aligned}$$

## Expected cost of GCDs

Suppose that  $P(x)$  has smallest factor of degree  $d$ . The number of GCDs required to find the factor, using our (quadratic polynomial) blocking strategy, is  $O(\sqrt{d})$ . By Assumption 2, the expected number of GCDs for a trinomial with no small factor is

$$1 + O\left(\sum_{(\lg r)/2 < d \leq r/2} \frac{\sqrt{d}}{d^2}\right) = 1 + O\left(\frac{1}{\sqrt{\log r}}\right).$$

Thus, the expected cost of GCDs per trinomial is

$$O(G(r)/\log r) = O(M(r)) = O(r \log r \log \log r).$$

This is asymptotically  $\ll$  expected cost of  $\text{sqr}/\text{mul}$  operations

## Cost of GCDs in practice

In practice, for  $r \approx 4.3 \times 10^7$ , GCDs take about 65% of the time versus 35% for `sqr/mul`.

The asymptotic analysis is misleading. This is because

$$\sqrt{\frac{\log r}{\log \log r}}$$

is a very slowly growing function of  $r$ .

## Comparison with classical algorithms

For simplicity I will use the  $\tilde{O}$  notation which ignores log factors. For example, instead of

$$O(n \log n (\log \log n)^2)$$

we can write  $\tilde{O}(n)$ .

The “classical” algorithm takes an expected time  $\tilde{O}(r^2)$  per trinomial, or  $\tilde{O}(r^3)$  to cover all trinomials of degree  $r$ .

The new algorithm takes expected time  $\tilde{O}(r)$  per trinomial, or  $\tilde{O}(r^2)$  to cover all trinomials of degree  $r$ .

## Comparison in practice

In practice, the new algorithm is faster by a factor of about 160 for  $r = 6\,972\,593$ , and by a factor of about 1000 for  $r = 43\,112\,609$ .

Thus, comparing the computation for  $r = 43\,112\,609$  with that for  $r = 6\,972\,593$ : using the classical algorithm would take about 240 times longer (impractical), but using the new algorithm saves a **factor of 1000**.

Generally, our search for eight different Mersenne exponents  $r \in \{3\,021\,377, \dots, 43\,112\,609\}$  took less time for larger  $r$ , due to incremental improvements in the search program!

## Computational results

In Sept 2008 Paul Zimmermann and I started searching for primitive trinomials of degree 43 112 609 (the largest known Mersenne exponent).

Dan Bernstein and Tanja Lange joined in the search and contributed CPU cycles. With their help, the search was completed by April 2009.

We found four new primitive trinomials  $x^r + x^s + 1$ ,  $r = 43\,112\,609$ :

$$s = 3\,569\,337, 4\,463\,337, 17\,212\,521, 21\,078\,848$$

Testing irreducibility took about 119 hours per trinomial on a 2.2 Ghz AMD Opteron (bogong), using our first algorithm. The “best” algorithm would take about 69 hours (saving 42%).

Most of the time (about 22 processor-years) was spent eliminating reducible trinomials at an average rate of about 32 sec per trinomial ( $\times 431\,12609/2$  trinomials).



## Recent results

We thought we were finished with this project, but in April the GIMPS project found another large Mersenne prime,

$$2^{42643801}.$$

Note that  $42\,643\,801 < 43\,112\,609$ . The GIMPS project does not guarantee to find Mersenne primes in increasing order of size!

Fortunately the new cluster *orac* arrived just in time. In June Paul Zimmermann and I started a search for primitive trinomials of degree  $r = 42\,643\,801$ .

So far the search is 99% complete, and we have found five new primitive trinomials.

# Conclusion

The new double-blocking strategy works well and, combined with fast multiplication and GCD algorithms, has allowed us to find new primitive trinomials of record degree. This would have been impractical using the classical algorithms.

The same ideas work over finite fields  $\text{GF}(p)$  for small prime  $p > 2$ , and for factoring sparse polynomials  $P(x)$  that are not necessarily trinomials: all we need is that the time for  $p$ -th powers (mod  $P(x)$ ) is much less than the time for multiplication (mod  $P(x)$ ).

# Acknowledgement

Thanks to

- Éric Schost for his comments on the work of Kedlaya and Umans.
- Dan Bernstein and Tanja Lange for contributing computer time to the search for degree 43 112 609.
- Alan Steel for independently verifying most of our new primitive trinomials using Magma.
- Victor Shoup for his package NTL, which was used to debug our software and check our certificates.
- Judy-anne for her help with beamer!

## References

V. L. Arlazarov, E. A. Dinic, M. A. Kronod & I. A. Faradzev, On economical construction of the transitive closure of an oriented graph, *Soviet Math. Dokl.* **11** (1975), 1209–1210.

W. Bosma & J. Cannon, *Handbook of Magma Functions*, School of Mathematics and Statistics, University of Sydney, 1995.

<http://magma.maths.usyd.edu.au/>

R. P. Brent, P. Gaudry, E. Thomé & P. Zimmermann, Faster multiplication in  $GF(2)[x]$ , *Proc. ANTS VIII 2008, Lecture Notes in Computer Science* **5011**, 153–166.

<http://wwwmaths.anu.edu.au/~brent/pub/pub232.html>

R. P. Brent & H. T. Kung, Fast algorithms for manipulating formal power series, *J. ACM* **25** (1978), 581–595.

<http://wwwmaths.anu.edu.au/~brent/pub/pub045.html>

R. P. Brent, S. Larvala & P. Zimmermann, A fast algorithm for testing reducibility of trinomials mod 2 and some new primitive trinomials of degree 3021377, *Math. Comp.* **72** (2003), 1443–1452.

<http://wwwmaths.anu.edu.au/~brent/pub/pub199.html>

R. P. Brent & P. Zimmermann, A multi-level blocking distinct-degree factorization algorithm, *Finite Fields and Applications: Contemporary Mathematics* **461** (2008), 47–58.

<http://wwwmaths.anu.edu.au/~brent/pub/pub230.html>

R. P. Brent & P. Zimmermann, Ten new primitive binary trinomials, *Math. Comp.* **78** (2009), 1197–1199.

<http://wwwmaths.anu.edu.au/~brent/pub/pub233.html>

D. G. Cantor and H. Zassenhaus, A new algorithm for factoring polynomials over finite fields, *Math. Comp.* **36** (1981), 587–592.

D. Coppersmith & W. Winograd, Matrix multiplication via arithmetic progressions, *J. Symb. Comput.* **9** (1980), 251–280.

J. von zur Gathen & J. Gerhard, *Modern Computer Algebra*, Cambridge Univ. Press, 1999.

J. von zur Gathen and J. Gerhard, Polynomial factorization over  $F_2$ , *Math. Comp.* **71** (2002), 1677–1698.

J. von zur Gathen and V. Shoup, Computing Frobenius maps and factoring polynomials, *Computational Complexity* **2** (1992), 187–224.  
<http://www.shoup.net/papers/>

K. Kedlaya & C. Umans, Fast modular composition in any characteristic, *Proc. FOCS 2008*, 146–155.  
<http://www.cs.caltech.edu/~umans/research.htm>

T. Kumada, H. Leeb, Y. Kurita and M. Matsumoto, New primitive  $t$ -nomials ( $t = 3, 5$ ) over  $GF(2)$  whose degree is a Mersenne exponent, *Math. Comp.* **69** (2000), 811–814. Corrigenda: *ibid* **71** (2002), 1337–1338.

A.-E. Pellet, Sur la décomposition d'une fonction entière en facteurs irréductibles suivant un module premier  $p$ , *Comptes Rendus de l'Académie des Sciences Paris* **86** (1878), 1071–1072.

J. M. Pollard. A Monte Carlo method for factorization, *BIT* **15** (1975), 331–334,

A. Schönhage, Schnelle Multiplikation von Polynomen über Körpern der Charakteristik 2, *Acta Inf.* **7** (1977), 395–398.

É. Schost, *Fast irreducibility test*, personal communication, 4 June 2008.

V. Shoup, NTL: A library for doing number theory.

<http://www.shoup.net/ntl/>

L. Stickelberger, Über eine neue Eigenschaft der Diskriminanten algebraischer Zahlkörper, *Verhandlungen des ersten Internationalen Mathematiker-Kongresses*, Zürich, 1897, 182–193.

R. G. Swan, Factorization of polynomials over finite fields, *Pacific J. Math.* **12** (1962), 1099–1106.

G. Woltman *et al*, GIMPS, The Great Internet Mersenne Prime Search.  
<http://www.mersenne.org/>